Information-Theoretic Topology-Hiding Broadcast: Wheels, Stars, Friendship, and Beyond

D'or Banoun
Reichman University, Herzliya, Israel

Elette Boyle
Reichman University, Israel
NTT Research, Sunnyvale, CA, USA

Ran Cohen
Reichman University, Herzliya, Israel

Abstract

Topology-hiding broadcast (THB) enables parties communicating over an incomplete network to broadcast messages while hiding the network topology from within a given class of graphs. Although broadcast is a privacy-free task, it is known that THB for certain graph classes necessitates computational assumptions, even against "honest but curious" adversaries, and even given a single corrupted party. Recent works have tried to understand when THB can be obtained with *information-theoretic* (IT) security (without cryptography or setup assumptions) as a function of properties of the corresponding graph class.

We revisit this question through a case study of the class of *wheel* graphs and their subgraphs. The n^{th} wheel graph is established by connecting n nodes who form a cycle with another "center" node, thus providing a natural extension that captures and enriches previously studied graph classes in the setting of IT-THB.

We present a series of new findings in this line. We fully characterize feasibility of IT-THB for any class of subgraphs of the wheel, each possessing an embedded *star* (i.e., a well-defined center connected to all other nodes). Our characterization provides evidence that IT-THB feasibility may correlate with a more fine-grained degree structure – as opposed to pure connectivity – of the corresponding graphs. We provide positive results achieving *perfect* IT-THB for new graph classes, including ones where the number of nodes is unknown. Further, we provide the first feasibility of IT-THB on non-degenerate graph-classes with $t > 1$ corruptions, for the class of *friendship* graphs (Erdös, Rényi, Sós'66).

2012 ACM Subject Classification Security and privacy → Information-theoretic techniques; Theory of computation → Cryptographic protocols

Keywords and phrases broadcast, topology-hiding protocols, information-theoretic security

Digital Object Identifier 10.4230/LIPIcs.ITC.2024.1

Funding *D'or Banoun*: Supported in part by AFOSR Award FA9550-21-1-0046.
Elette Boyle: Supported in part by AFOSR Award FA9550-21-1-0046 and ERC Project HSS (852952).
Ran Cohen: Supported in part by NSF grant no. 2055568, by ISF grant 1834/23, and by the Algorand Centres of Excellence programme managed by Algorand Foundation. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of Algorand Foundation.

1 Introduction

Topology-hiding protocols over an incomplete communication network guarantee that colluding parties do not learn additional information about the topology of the network graph (from within a given class of graphs), beyond their own neighbor-set [12]. Such protocols may be of interest in settings where the communication structure itself is sensitive information,

5th Conference on Information-Theoretic Cryptography (ITC 2024).
Editor: Divesh Aggarwal; Article No. 1; pp. 1:1–1:13
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

such as in social networks, or peer-to-peer networks based on geographical position. Perhaps the most fundamental goal is that of achieving topology-hiding *broadcast* (THB), where a designated sender wishes to convey an input to all participating parties.

Although broadcast is a privacy-free task, THB turned out to be a challenging goal on its own. It was recently shown that THB for certain graph classes necessitates computational assumptions, even in the "honest but curious" *semi-honest* setting (when corrupted parties follow the protocol honestly but try to learn more information from their joint view), and even given a *single* corrupted party [6, 5]. This lies in stark contrast to the topology-revealing case, in which broadcast is trivially achievable in the semi-honest setting.

Obtaining topology hiding based on computational assumptions has been the subject of a fruitful collection of works, leading to various THB, and in turn, general topology-hiding secure multiparty computation (THC) protocols [12, 8, 2, 1, 9, 6, 10, 11, 3]. It is known by now how to construct THB protocols for the class of all graphs (of polynomial size) that are secure against *any* subset of semi-honest corruptions under standard number-theoretic cryptographic hardness assumptions such as DDH, QR, and LWE,[1] or from unstructured assumptions such as constant-round constant-rate oblivious transfer [3].

Motivated by an analogous question within secure multi-party computation, the work of [5] asked whether existence of an honest majority can enable *information-theoretically* secure THB protocols in certain settings, without relying on cryptographic assumptions and withstanding computationally unbounded adversaries. We refer to this as IT-THB. The work of [5] ruled out 1-secure IT-THB on a path with four nodes (which is 1-connected) but devised a perfect 1-secure information-theoretic THC on cycles of known length (which are 2-connected); see Figure 1. Given these initial evidence, they conjectured that feasibility of IT-THB may depend on the *connectivity*[2] of the graphs within the class: namely, that $(t+1)$-connectivity is sufficient and/or necessary for t-secure IT-THB.

The special case of $t = 1$ was further investigated by [4], who proved that the conjecture holds in this case for the stronger notion of THC. They showed that information-theoretic THC with security against a single semi-honest corruption is possible if and only if the connectivity of every graph in the class is at least 2. However, they additionally showed that the conjecture does *not* hold for THB, by constructing a perfectly secure THB against a single corruption for the butterfly graph class (Figure 1), where each graph is only 1-connected.

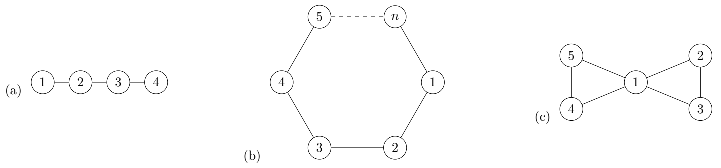

Figure 1 (a) Class $\mathcal{G}_{\mathsf{4\text{-}path}}$, of all isomorphisms of 4 nodes on a path; 1-secure THB over $\mathcal{G}_{\mathsf{4\text{-}path}}$ implies key agreement. (b) Class $\mathcal{G}_{\mathsf{cycle}}(n)$ of all isomorphisms of n nodes on a cycle; admits 1-secure IT-THB. (c) Class $\mathcal{G}_{\mathsf{butterfly}}$ of all isomorphisms of 5 nodes on a butterfly graph (two triangles with a common node); contains 1-connected graphs yet admits 1-secure IT-THB.

The results of [5, 4] open a rich domain of questions. As [4] showed, high connectivity is not the "right" criterion for feasibility of THB (in contrast to THC), and alternative graph-properties may serve as candidate conjectures. Therefore, our first question is:

Given a graph-class, which graph properties characterize feasibility of 1*-secure IT-THB?*

[1] DDH stands for the decisional Diffie-Hellman assumption, QR for the quadratic residuosity assumption, and LWE for the learning with errors assumption.

[2] We consider node-connectivity; that is, a graph is k-connected if and only if every pair of nodes is connected by k *vertex-disjoint* paths.

Zooming into [4], their general positive result, of 1-IT-THB over 2-connected graphs, has a nonzero (yet exponentially small) error probability. This means that THB with *perfect* security is only known for cycles [5], for the butterfly graph [4], and for graphs with at most four nodes [4]. Is it possible that for graphs with $n > 5$ nodes the source of perfect 1-THB is the highly symmetric structure of cycles, and other graph classes inherently require a positive error?

Are there additional graph-classes that support perfectly secure THB?

Finally, all feasibility results for IT-THB are secure against a single corruption. Indeed, 2-secure THB on a 4-node rectangle, possibly with a missing edge, requires oblivious transfer [6], and a 2-secure THB on a cycle with 7 nodes (or more) requires key agreement [5]. The statistically secure THB protocols for 2-connected graphs from [4] completely break if there are two corruptions, and in the butterfly class two corruptions trivialize the problem, as there is no information to hide. One may wonder if IT-THB simply cannot withstand multiple corruptions that provide several points of view about the graph topology, except for degenerate cases where the topology is already revealed by the corrupted parties' neighbor-sets. This leads to our third question:

Are there graph-classes that support IT-THB with more than a single corruption?

1.1 Our Contributions

In this work, we conduct an investigation of these questions through a case study of the class of *wheel graphs* and their subgraphs. The n^{th} wheel graph W_n is established by connecting a single node (the "center") to n nodes who form a cycle, as depicted below. The wheel graph-class $\mathcal{G}_{\mathsf{wheel}}(n)$ consists of all isomorphisms of the wheel graph, i.e., all assignments of the labels $\{1, \ldots, n+1\}$ to the nodes of the wheel graph W_n.

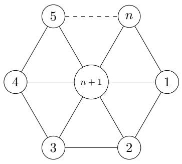

Wheel graphs and their subgraphs form a natural extension that captures and enriches previously studied graph classes in the setting of IT-THB: for example, paths, cycles, triangles, and butterfly graphs. Interestingly, although $\mathcal{G}_{\mathsf{wheel}}(n)$ has increased connectivity over the n-cycles, the corresponding state-of-the-art THB protocols for $\mathcal{G}_{\mathsf{wheel}}(n)$ are slightly worse. Note that the cycle protocol cannot simply be run directly, as parties on the perimeter of the graph do not know – in fact, must not know – which neighbor is the center node.

Several challenges arise when hiding the topology of $\mathcal{G}_{\mathsf{wheel}}(n)$. First, consider a node v on the perimeter; such a node has three neighbors, one of which is the center. To hide the identity of the center node, either the protocol does not utilize the power of the center, or each of the non-center neighbors must emulate the behavior of the center toward v, and further, v must emulate the center toward all its neighbors. Second, consider the center node; this node is connected to all other parties but must not learn how the parties on the perimeter are connected among themselves. Further, an adversary that corrupts two parties on the perimeter without a common neighbor must not learn their relative distance on the perimeter. Note that an adversary that corrupts $n-2$ nodes in a wheel graph knows the entire topology from the corrupted nodes' neighbor-sets; however, for $t \leq n-3$ corruptions not all is revealed (i.e., when there are 4 honest parties).[3]

[3] When considering arbitrary admissible graphs with $n+1$ nodes (as defined below), there is more information to hide; therefore, an adversary that corrupts n nodes knows the entire topology but for

Characterization of wheels and subgraphs with an embedded star. Our first result shows that perfectly secure THB is possible against a single semi-honest corruption on the class of wheel graphs $\mathcal{G}_{\mathsf{wheel}}(n)$, as well as on certain classes of its subgraphs. Concretely, given any family of subgraphs of the wheel with $n+1$ nodes, with an embedded star in each graph (i.e., where the center is fully connected and has degree n), we show that IT-THB with one corruption is possible if either the *minimal degree* of non-center nodes in the family is greater than 1, or if it is 1 but so is the *maximal degree*. Surprisingly, we show that this characterization is tight for any such subclasses that are closed under isomorphism (i.e., for each graph topology in the class, all relabelings of this graph are also contained in the class); that is, if the maximal degree is greater than 1 but the minimal degree is 1, then THB on this class implies key agreement.

This would suggest that feasibility of IT-THB may correlate with a more fine-grained degree structure, as opposed to connectivity, of graphs.

More concretely, we begin by defining *admissible subgraphs* as subgraphs of the wheel graph W_n in which the degree of the center is n and the degree of every other node is either 2 or 3. The *butterfly graph* is an example for an admissible subgraph for $n = 4$, as well as the $(2n+1)$-node *friendship graph* F_n,[4] see Figure 2.

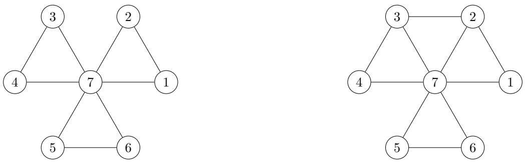

Figure 2 Examples of admissible subgraphs of $\mathcal{G}_{\mathsf{wheel}}(6)$. On the left is a friendship graph in which every non-center node has degree 2, and on the right is a subgraph where every non-center node has degree 2 or 3.

When considering graphs with an embedded star, i.e., with a fully connected center, *non-admissible graphs* are those who contain a non-center node of degree 1. The extreme example is the *star graph* in which the center node is connected to n nodes, and no other edges exist, see Figure 3.

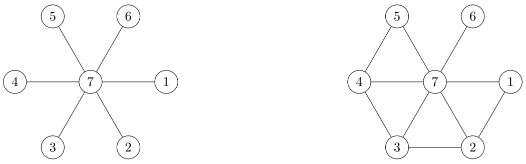

Figure 3 Example of non-admissible subgraphs of $\mathcal{G}_{\mathsf{wheel}}(6)$. On the left is the star graph with 7 nodes. On the right is a subgraph with a single node of degree 1.

Our characterization nearly shows that IT-THB is possible for a given graph-class with a fully connected center if and only if it consists only of admissible subgraphs. The single exception is the graph class $\mathcal{G}_{\mathsf{star}}(n)$ that only contain star graphs, which are not admissible; this class is degenerate (trivially providing topology hiding) since any node can identify the center and derive the whole topology.

$t \leq n-2$ not all is revealed (i.e., when there are 2 honest parties) .

[4] The friendship graph F_n, introduced in [7], is a planar, undirected graph with $2n+1$ nodes and $3n$ edges. F_n can be constructed by joining n triangles with a common node.

▶ **Theorem 1** (IT-THB for admissible graphs with fixed size, informal). *Let $n \in \mathbb{N}$ with $n \geq 4$, and let $\mathcal{G} \subseteq \mathcal{G}_{\mathsf{wheel}}(n)$ be a graph-class in which every graph has $n+1$ nodes and the center has degree n.*

Then, if either $\mathcal{G} = \mathcal{G}_{\mathsf{star}}(n)$ or if $\mathcal{G}$ consists of admissible graphs, there exists perfectly secure IT-THB against a single semi-honest corruption over $\mathcal{G}$. Otherwise, THB over $\mathcal{G}$ secure against a single semi-honest corruption exists if and only if key agreement exists.

Theorem 1 demonstrates another interesting phenomena: a nontrivial example of a graph-class in which $\mathcal{G}$ is the union of two sub-classes $\mathcal{G}_1$ and $\mathcal{G}_2$, such that each sub-class admits an IT-THB, yet the there is no IT-THB for $\mathcal{G}$. Specifically, while $\mathcal{G}_{\mathsf{wheel}}(n)$ and $\mathcal{G}_{\mathsf{star}}(n)$ each individually admits 1-IT-THB, any 1-THB protocol on $\mathcal{G}_{\mathsf{wheel}}(n) \cup \mathcal{G}_{\mathsf{star}}(n)$ requires key agreement.

Generalizing to variable-size subgraphs. We proceed to analyze subgraphs of $\mathcal{G}_{\mathsf{wheel}}(n)$ that are generated by removing some of the nodes. Note that when removing the center node, the resulting subgraph is either a cycle with n nodes $\mathcal{G}_{\mathsf{cycle}}(n)$, which supports 1-secure perfect THB, or a path with up to n nodes that necessitates key agreement. Therefore, we focus on keeping the center and removing nodes from the perimeter. An interesting observation is that when removing k neighboring nodes from the perimeter, the result is an admissible subgraph of the wheel with $n+1-k$ nodes with one edge removed from the perimeter. Similarly, removing arbitrary k nodes yields a subgraph of the wheel with $n+1-k$ nodes with m edges removed from the perimeter, where m is the number of sets of neighboring nodes that are removed.

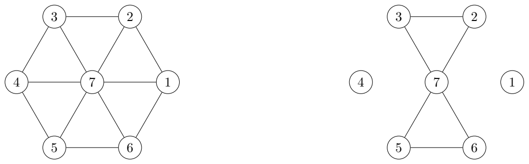

Figure 4 On the left is a wheel graph. On the right is the resulting graph when removing nodes ① and ④ together with their corresponding edges. The result is an admissible graph F_2.

A more interesting question is thus to characterize families of such subgraphs whose number of nodes is not a priori known. We remark that topology hiding on graphs of unknown size can be surprisingly complex: For example, THB with an additional sender-anonymity guarantee for the simple class of 2-paths and 3-paths implies *infinitely often oblivious transfer* [4, Thm 5.4].

We utilize a useful property of the protocol used for proving Theorem 1 (discussed further in Section 2) that effectively hides the number of nodes from non-center parties. We show that the protocol can be applied also to the current setting to obtain perfect IT-THB.

▶ **Theorem 2** (IT-THB for admissible graphs with varying size, informal). *Let $n \in \mathbb{N}$ and let $\mathcal{G}$ be a graph-class such that every $(V,E) \in \mathcal{G}$ is a subgraph of the wheel graph, and it holds that $4 \leq |V| \leq n+1$ and there is a center node with degree $|V|-1$. Then,*

- *if the maximal degree of non-center nodes is 1, i.e., $\mathcal{G}$ consists only of stars (possibly of different size), or*
- *if the minimal degree of non-center nodes is 2 or 3, i.e., $\mathcal{G}$ consists only of admissible graphs, or*
- *if $\mathcal{G}$ consists both of stars and admissible graphs but they are of different sizes,*

there exists perfectly secure IT-THB *against a single semi-honest corruptions over* $\mathcal{G}$*. Otherwise,* THB *over* $\mathcal{G}$ *secure against a single semi-honest corruption exists if and only if key agreement exists.*

We note that Theorem 2 subsumes Theorem 1; therefore, in the technical sections we directly prove Theorem 2.

Tolerating many corruptions: the case of friendship graphs. The feasibility results thus far were limited to a single corruption. The reason lies in the structure of the protocol, which enables two colluding parties with two common neighbors to learn which of them is the center; see Section 2 for an illustration. Therefore, it still remains open whether IT-THB tolerating $t > 1$ corruption is possible, aside from degenerate cases in which the topology is fully determined from neighbor-sets of any t nodes.

We proceed to analyze an interesting class of subgraphs of a wheel graph with varying size, which consists of *friendship graphs*. Recall that for $n \geq 1$, the friendship graph F_n is a $(2n+1)$-nodes graph constructed by joining n triangles with a common node. They were named after the friendship theorem [7], which states that if in a finite set of people every pair has one common friend, then there exists one person who is friend with everyone. We consider a class consisting of friendship graphs of different sizes. Note that the connectivity of each of those graphs is 1, and by their structure every two nodes can only have one common neighbor, so the attack discussed above no longer applies. We prove that indeed perfect IT-THB tolerating *any* number of corruptions can be achieved on this class. For an integer k, consider the graph class $\mathcal{G}_{\mathsf{friendship}}(k)$ containing all isomorphisms of the friendship graph F_k.

▶ **Theorem 3** (t-IT-THB over friendship graphs, informal)**.** *Let* $n \in \mathbb{N}$ *with* $n \geq 2$*, let* $t < 2n+1$*, and consider a graph-class* $\mathcal{G} \subseteq \bigcup_{k=2}^{n} \mathcal{G}_{\mathsf{friendship}}(k)$*. There exists a perfectly secure* THB *protocol against* t *semi-honest corruptions over* $\mathcal{G}$*.*

We remark that Theorem 3 presents the first feasibility of information-theoretic THB on non-degenerate graph-classes with $t > 1$ corruptions.

Organization of the paper. Due to severe space restrictions, we defer most of the technical content, including the construction of our protocols, the formal statements, and the security proofs, to the full version of the paper. We proceed to provide an overview of the techniques in Section 2.

2 Technical Overview

We move on to describing some of our techniques. We begin by explaining in Section 2.1 the high-level ideas of the protocols used for our positive result. Next, in Section 2.2, we describe our usage of the *phantom-jump* technique from [4] for our negative result.

2.1 Feasibility Results: The "Oblivious Centralized Coordination" Technique

Our protocols are inspired by the THB protocol for the butterfly graph from [4]. We extend it in several aspects to support more involved graph classes that contain an embedded star, i.e., a well-defined center connected to all other nodes. In the overview below, we begin by describing the simpler case of friendship graphs, and then proceed to the wheel graph, and to arbitrary admissible graphs.

Starting point: the butterfly graph. Recall that the butterfly graph (Figure 1) is in fact the friendship graph F_2: a 5-node graph consisting of two triangles connected by a common center node. The high-level idea is to use the center node for coordinating the protocol. The protocol runs multiple instances of *reliable message transmission* (RMT), one for every potential receiver. In each RMT instance, the sender P_S sends its message to all its neighbors in the first step. Note that each party knows whether it is a neighbor of P_S, so it knows whether it should receive a message or not in the first round. At that point it is guaranteed that the center node holds the message and so can deliver it to the receiver (in case the receiver is not the center).

This, of course, will reveal to the receiver who is the center node. Therefore, the center must do so in an *oblivious* way, without exposing itself. In the butterfly graph, if the receiver P_R is not the center it has one more neighbor other than the center. The approach taken in [4] is to secret share the message m with the additional neighbor, and have each neighbor deliver one share. However, the center does not know who that neighbor is. Therefore, the center node prepares 2-out-of-2 shares of the message m for each potential neighbor, i.e., each non-receiver party.

To help the center hide its identity, each other party assists by acting as the center and preparing 2-out-of-2 shares of zero (so called, *blinding terms for addition*) for each of its non-receiver neighbors (a non-center party has either one or two non-receiver neighbors). Next, the receiver receives four values from each of its neighbors (recall that in the butterfly graph there are four nodes other than the receiver, see Figure 1), such that the center sends the sum of the share m for each party with the share of zero it received from that party, and the second neighbor sends the sum of the share received from its non-receiver neighbor with the share of zero sent to this neighbor, along with three random values (one for each other party). The receiver can then select the correct pair which corresponds to its true neighbors. Thus, P_R can reconstruct m without knowing which of its neighbors is the center.

This approach is secure as long as the receiver is not the center. However, if P_R is the center, it may learn the neighbor-set of other nodes (e.g., by inspecting which pairs of values sum up to 0). This is solved by adding *suitable offset* values, which are multiplied by *blinding terms for multiplication*, and only come into play if P_R is the center. Specifically, if P_R is *not* the center, then P_R will send the same offset to both its neighbors (this will ensure that the offset will be canceled out). If P_R is the center, then P_R will send a *different* offset to each neighbor (this requires working over a larger field, e.g., $\mathbb{F}_4$, to support a different value per party); in this case, the pairs of values seen by P_R will induce a linear system of two equations with two variables, and the different offsets will guarantee that the system has full rank and always has a solution. This, in turn, will prove that the center cannot identify which pairs of parties are connected.

The friendship graph. As discussed above, we view the butterfly graph as two triangles connected in a joint node; that is, as the friendship graph F_2. In the full version, we prove that the 1-THB protocol for the butterfly graph-class $\mathcal{G}_{\mathsf{butterfly}}$ (consisting of all isomorphisms of 5 nodes to F_2) extends in natural way to 1-THB for the class of friendship graph $\mathcal{G}_{\mathsf{friendship}}(n)$, for $n \geq 2$, consisting of all isomorphisms of $2n+1$ nodes to F_n.[5] Namely, the receiver now receives a vector of $2n$ values from each of its neighbors, and those values are uniformly distributed conditioned on the corresponding values of its neighbors that sum up to the message. Further, recall that in case the receiver is the center, it must provide a different offset to each of its neighbors; hence, the underlying field $\mathbb{F}_q$ must grow and satisfy $q \geq 2n$.

[5] Note that for $n = 1$ a friendship graph is just a triangle, and there is no well-defined center.

Friendship of variable size. A second observation is that for non-center parties, the protocol behaves in a "local" manner, in the sense that the neighbors of a non-center node are neighbors on their own. When P_R is not the center, this enables the receiver's neighbors to jointly construct the shares of the message in a coordinated (yet oblivious) way. Only the center's actions truly depend on the actual number of parties, while non-center parties only need to know an upper bound on the number of parties.

In the full version, we prove that this locality property makes the protocol suitable for a variable number of nodes (i.e., a variable number of triangles). Non-center nodes proceed as if the graph has n triangles (where n is an upper bound), and the center node emulates missing nodes in its head. Formally, we consider the graph $F_{k,n}$ as an *augmented friendship graph* of $2n+1$ nodes, where $2k+1$ nodes form a connected component which is the friendship graph F_k, and all other $2n+1-(2k+1)=2(n-k)$ nodes are singletons (isolated parties). Each isolated party simply outputs 0 in this protocol (unless it is the sender, in which case it outputs its input), and the *agreement* and *validity* properties are only required for the connected component of the sender.

Friendship with many corruptions. Another interesting observation, is that locality enables tolerating an arbitrary number of $t < 2n+1$ corruptions, *without* any adjustments to the protocol. We prove this in the full version. Intuitively, to see why, we distinguish between an honest center and a corrupt center.

In case the center is honest, then once there is more than a single corruption, the adversary can immediately identify who the center is. This is not considered a violation of privacy, since this can be deduced just by observing the common neighbor of the corrupted parties, and *without* observing any protocol messages. When focusing on each triangle now, if both non-center nodes are corrupted there is nothing to hide within the triangle, whereas if none of the non-center nodes is corrupted the adversary learns nothing new from the protocol. The case where there is a single corrupted non-center in the triangle reduces to the single corruption case from before.

In case the center is corrupted, and there is another non-center corrupted party, then all the information in its triangle is already known, regardless of whether the second non-center party is honest or not. Further, consider the set of honest parties that have an honest neighbor, then the center together with all other corrupt parties do not learn the connectivity of this set.

We note that despite the technical simplicity of this result, it bares a more significant conceptual contribution, as it provides the first feasibility of $\mathsf{IT\text{-}THB}$ with more than one corruption beyond trivial graph classes.

Beyond friendship: the wheel graph. We proceed to extend the *oblivious centralized coordination* technique to more involved graph classes that admit an embedded star. As before, we begin by considering a *single* corruption. One can view the $(2n+1)$-nodes friendship graph F_n as a subgraph of the wheel W_{2n} in which every non-center node has degree 2. The wheel graph presents the other extreme in some sense, as every non-center node has degree 3.

A first attempt to extend the protocol to this new regime, is to use 3-out-of-3 secret sharing instead of 2-out-of-2. Stated differently, before, in F_n, if P_R is not the center it receives a *vector* of $2n$ values from each of its two neighbors such that the matching pair of values sum up to the message and all other values are independently and uniformly distributed. When considering the $(n+1)$-nodes wheel graph W_n, if P_R is not the center then it has 3 neighbors, and it receives a *matrix* of $n \times n$ values from each of its neighbors such that the corresponding entries in these matrices[6] sum up to the message and all other values are independently and uniformly distributed.

However, as opposed to the friendship regime, once a non-center node has degree 3 the protocol loses its locality property, as now not all neighbors of the receiver are neighbors on their own, and so the matrices are not "synchronized" like the vectors in the previous case. Indeed, if done without care, this approach leads to an attack. The reason is that the preparation of entry (v, w) for the matrix of party P_u is done as follows: if v and w are not neighbors of u sample a random value; if only one is a neighbor use the value that this party sent before (to ensure it will cancel out); and if both parties are neighbors of u then take the sum of their values. Therefore, the receiver can identify repeating entries in a matrix to deduce pairs of neighboring parties, as illustrated in Figure 5.

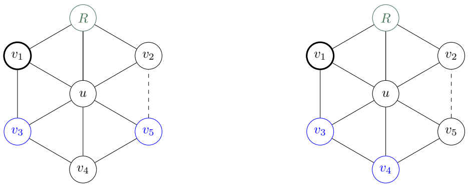

Figure 5 Illustration of an attack on a naïve protocol for $\mathcal{G}_{\mathsf{wheel}}(n)$. The receiver P_R has three neighbors: v_1, v_2, and the center u. Say that v_1 has another neighbor v_3, which has a third neighbor v_4, which has a third neighbor v_5. Then, P_R receives a matrix from v_1; however, since v_3 sends a single value to v_1, the entry (v_3, v_4) will be the same as the entry (v_3, v_5). This means that both v_4 and v_5 are *not* neighbors of v_1.

Our solution to this issue is to have each pair of neighbors (none of which is P_R) generate a vector of n *correlated values*, as opposed to a single value. This is done by having each party sample a vector of random values and send it to each of its non-receiver neighbors. In fact, those correlated values make the *blinding terms* and the *suitable-offset terms* redundant, so these values are no longer used in this protocol. In the full version, we prove that the resulting protocol is secure for the class of wheel graphs.

Admissible graphs. Having established 1-THB for the case when non-center nodes have degree 2 (friendship graphs) and the case where they have degree 3 (wheel graphs), we proceed to combine the ideas together and support any admissible graph. Intuitively, since the protocols share a similar structure, one can hope to execute both options concurrently. That is, the parties run two independent executions: one for the case where P_R has two neighbors, and one for the case where P_R has three neighbors. This, however, is vulnerable to an attack, since when a receiver has three neighbors it can find correlations in the messages it receives for the degree-2 execution and identify who the center is, as illustrated in Figure 6.

[6] That is, for neighbors u, v, w take entry (u, v) from the matrix of w, entry (v, w) from the matrix of u, and entry (w, u) from the matrix of v. In the protocol, we ensure the matrices are symmetric, i.e., $\mathbf{M}[u, v] = \mathbf{M}[v, u]$.

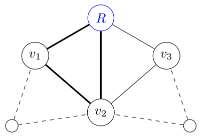

Figure 6 Illustration of an attack on a non-careful protocol for admissible graphs. Consider a non-center receiver P_R with neighbors v_1, v_2, and v_3; assume that v_2 is the center. Further consider running the friendship protocol over this graph. The left diagram, illustrates the view P_R obtains for the triangle with v_1 and v_2: here P_R will obtain the message m. The middle diagram, illustrates the view P_R obtains for the triangle with v_2 and v_3: again, P_R will obtain the message m. The right diagram, illustrates the view P_R obtains for the triangle with v_1 and v_3: here, there is no direct edge between v_1 and v_3; hence, P_R will *not* obtain the message m. Therefore, P_R can identify that v_2 is the center.

The main idea in overcoming this attack, is that although we need to run two executions in order to hide the degree of the receiver (when it is not the center), we only need one execution to deliver the message to the receiver, and the second does not need to convey any information. Further, the receiver already knows its degree, so it knows which execution is the "right" one, and can sabotage the "redundant" one. Specifically:

- In case the receiver's degree is 3, in the degree-2 execution it will send a *different* offset for each neighbor (and the degree-3 execution will be executed correctly).
- In case the receiver's degree is 2, in the degree-2 execution it will correctly send the *same* offset to its neighbors (and the degree-3 execution will not leak any information because the receiver does not have three neighbors).

In the full version, we prove that the resulting protocol is secure against one corruption for graph-classes consisting of admissible graphs.

Many corruptions. The protocol described above establishes feasibility of 1-IT-THB for any graph-class consisting of admissible graphs (even of variable size). This feasibility is tight for a single corruption, as stated in Theorem 2. It is tempting though to extend the resiliency of the protocol, similarly to the class of friendship graphs that support any number of corruptions. It turns out that the non-local nature of non-friendship, admissible graphs enables an attack on the protocol when the adversary controls two nodes.

We illustrate the attack in Figure 7. Consider a pair of corrupted parties ② and ④ , and assume that none of them is the center. Further, assume that each has degree 3, and that they have two common neighbors, denoted ③ and ⑥ . Clearly, by the structure of the graph, ② and ④ together can deduce that either ③ is the center, or ⑥ is the center.

However, when running in this setting the 1-secure protocol described above, the colluding parties may learn correlations that will expose which of their common neighbors is the center. Specifically, recall that when ③ is the receiver, it sends to its neighbors ② and ④ the suitable-offset values. In case ③ is the center, the offset value for ② is the same as the one for ④ , whereas in case ③ is *not* the center these are different values.

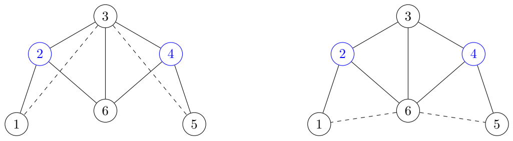

Figure 7 Attack on non-friendship admissible graphs with two corruptions.

We emphasize that in friendship graphs every non-center node has degree two; hence, the scenario from Figure 7 cannot occur. We leave it as an open question to find a protocol that is resilient to $t > 1$ corruptions for non-friendship admissible graphs.

2.2 Impossibility Results: The "Phantom Jump" Technique

The phantom-jump technique, introduced in [4], was used to show that key agreement is necessary for 1-secure THB over the class $\mathcal{G}_{\mathsf{triangle}}$ consisting of a triangle, with possibly one of its edges missing (see Figure 8). In this class, if a party has two neighbors it does not know whether its neighbors are directly connected or not, but a party with one neighbor knows the entire topology.

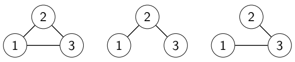

Figure 8 The class $\mathcal{G}_{\mathsf{triangle}}$ from [4], consisting of a triangle, with possibly one of its edges missing.

In the full version we prove the lower bound of Theorem 1 (namely that 1-THB on the union of an admissible graph-class of size $n + 1$ with $\mathcal{G}_{\mathsf{star}}(n)$ necessitates key agreement) by a direct reduction to the impossibility in [4]. Below we explain in a more explicit manner how the phantom-jump technique from [4] is used in this argument. We illustrate this for $\mathcal{G} = \mathcal{G}_{\mathsf{wheel}}(4) \cup \mathcal{G}_{\mathsf{star}}(4)$ where both graphs consist of 5 nodes.

The high-level idea, going back to [5], is to construct a key-agreement protocol from a 1-secure THB protocol π for $\mathcal{G}$. Recall the desired key-agreement protocol is run between two parties, Alice and Bob, and concludes with the parties outputting a bit $b \in \{0, 1\}$, such that a channel eavesdropper listening to communications cannot predict the value of b with non-negligible advantage. To construct a key-agreement protocol from π, Alice begins by choosing two long random strings m_1 and m_2 and sending them to Bob in the clear. Next, Alice and Bob continue in phases as follows:

- In each phase Alice and Bob locally toss coins A and B, respectively.
- They proceed to run two executions of π in which Alice always emulates ① and Bob always emulates ②. In addition, if $A = 0$ then Alice emulates ③ , ④ , and ⑤ as neighbors of ①, who acts as the center of the star, and ③ broadcasting m_1 in the first run; otherwise she emulates ③ , ④ , and ⑤ as neighbors of ①, who acts as the center of the star, and ③ broadcasting m_2 in the second run. Similarly, if $B = 1$ then Bob emulates ③ , ④ , and ⑤ as neighbors of ②, who acts as the center of the star, and ③ broadcasting m_1 in the first run; otherwise he emulates ③ , ④ , and ⑤ as neighbors of ②, who acts as the center of the star, and ③ broadcasting m_2 in the second run. See Figure 9 for an illustration.
- If parties ① and ② output m_1 in the first run and m_2 in the second, Alice and Bob output their bits A and B, respectively; otherwise, they execute another phase.

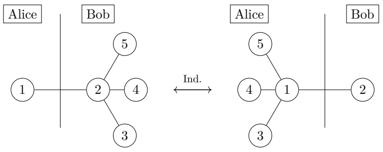

Figure 9 Using wheels and stars to construct a key-agreement protocol.

Clearly, if $A = B$ in some iteration then Alice and Bob will output the same coin, and by the assumed security of π, the eavesdropper Eve will not be able to learn who emulated ③ , ④ , and ⑤ in the first run and who in the second. If $A \neq B$, then in at least one of the runs nobody emulates the broadcaster ③ , so with overwhelming probability Alice and Bob will detect this case and execute another iteration.

In more detail, when $A = B$ the view of Eve consists of the communication between ① and ② , as depicted in Figure 9. By THB security, when ② acts as the center it cannot distinguish between the star and the wheel; in particular, the distribution of the messages on the channel between ① and ② is indistinguishable in both cases. Again, by THB security, when ① is not the center of the wheel it cannot know which of its neighbors is the center, so it cannot distinguish between the center being ② or ③ ; in particular, the distribution of the messages on the channel between ① and ② is indistinguishable in both cases. Similarly, when ② is not the center of the wheel, it cannot distinguish between the center being ① or ③ ; in particular, the distribution of the messages on the channel between ① and ② is indistinguishable in both cases. Finally, when ① acts as the center it cannot distinguish between the star and the wheel; in particular, the distribution of the messages on the channel between ① and ② is indistinguishable in both cases. By a simple hybrid argument it follows that the messages between ① and ② are indistinguishable when communicating in a star topology when ① is the center and when ② is the center, and it follows that the distinguishing advantage of Eve is negligible. See Figure 10 for an illustration of the hybrid argument.

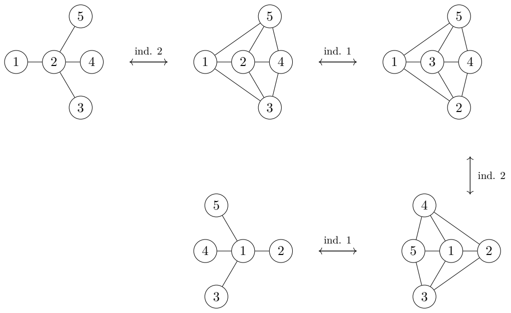

Figure 10 Hybrid steps in the phantom jump over wheels and stars.

References

1 Adi Akavia, Rio LaVigne, and Tal Moran. Topology-hiding computation on all graphs. In *37th Annual International Cryptology Conference (CRYPTO), part I*, pages 447–467, 2017.

2 Adi Akavia and Tal Moran. Topology-hiding computation beyond logarithmic diameter. In *36th Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT), part III*, pages 609–637, 2017.

3 Marshall Ball, Alexander Bienstock, Lisa Kohl, and Pierre Meyer. Towards topology-hiding computation from oblivious transfer. In *Proceedings of the 21st Theory of Cryptography Conference (TCC), part I*, pages 349–379, 2023.

4 Marshall Ball, Elette Boyle, Ran Cohen, Lisa Kohl, Tal Malkin, Pierre Meyer, and Tal Moran. Topology-hiding communication from minimal assumptions. In *Proceedings of the 18th Theory of Cryptography Conference (TCC), part II*, pages 473–501, 2020.

5 Marshall Ball, Elette Boyle, Ran Cohen, Tal Malkin, and Tal Moran. Is information-theoretic topology-hiding computation possible? In *Proceedings of the 17th Theory of Cryptography Conference (TCC), part I*, pages 502–530, 2019.

6 Marshall Ball, Elette Boyle, Tal Malkin, and Tal Moran. Exploring the boundaries of topology-hiding computation. In *37th Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT), part III*, pages 294–325, 2018.

7 Paul Erdös, Alfréd Rényi, and Vera T. Sós. On a problem of graph theory. *Studia Sci. Math. Hungar.*, 1:215–235, 1966.

8 Martin Hirt, Ueli Maurer, Daniel Tschudi, and Vassilis Zikas. Network-hiding communication and applications to multi-party protocols. In *36th Annual International Cryptology Conference (CRYPTO), part II*, pages 335–365, 2016.

9 Rio LaVigne, Chen-Da Liu Zhang, Ueli Maurer, Tal Moran, Marta Mularczyk, and Daniel Tschudi. Topology-hiding computation beyond semi-honest adversaries. In *Proceedings of the 16th Theory of Cryptography Conference (TCC), part II*, pages 3–35, 2018.

10 Rio LaVigne, Chen-Da Liu Zhang, Ueli Maurer, Tal Moran, Marta Mularczyk, and Daniel Tschudi. Topology-hiding computation for networks with unknown delays. In *Proceedings of the 23rd International Conference on the Theory and Practice of Public-Key Cryptography (PKC), part II*, pages 215–245, 2020.

11 Shuaishuai Li. Towards practical topology-hiding computation. In Shweta Agrawal and Dongdai Lin, editors, *28th International Conference on the Theory and Application of Cryptology and Information Security (ASIACRYPT), part I*, pages 588–617, 2022.

12 Tal Moran, Ilan Orlov, and Silas Richelson. Topology-hiding computation. In *Proceedings of the 12th Theory of Cryptography Conference (TCC), part I*, pages 159–181, 2015.