The S-Hamiltonian Cycle Problem

Antoine Amarilli
Univ. Lille, Inria, CNRS, Centrale Lille, UMR 9189 CRIStAL, F-59000 Lille, France

Arthur Lombardo
DI ENS, ENS, PSL University, CNRS, Inria, Paris, France
Univ. Lille, CNRS, Inria, Centrale Lille, UMR 9189 CRIStAL, F-59000 Lille, France

Mikaël Monet
Univ. Lille, Inria, CNRS, Centrale Lille, UMR 9189 CRIStAL, F-59000 Lille, France

Abstract

Determining if an input undirected graph is Hamiltonian, i.e., if it has a cycle that visits every vertex exactly once, is one of the most famous NP-complete problems. We consider the following generalization of Hamiltonian cycles: for a fixed set S of natural numbers, we want to visit each vertex of a graph G exactly once and ensure that any two consecutive vertices can be joined in k hops for some choice of $k \in S$. Formally, an *S-Hamiltonian cycle* is a permutation $(v_0, \ldots, v_{n-1})$ of the vertices of G such that, for $0 \leq i \leq n-1$, there exists a walk between v_i and $v_{i+1 \bmod n}$ whose length is in S. (We do not impose any constraints on how many times vertices can be visited as intermediate vertices of walks.) Of course Hamiltonian cycles in the standard sense correspond to $S = \{1\}$. We study the *S-Hamiltonian cycle problem* of deciding whether an input graph G has an S-Hamiltonian cycle. Our goal is to determine the complexity of this problem depending on the fixed set S. It is already known that the problem remains NP-complete for $S = \{1, 2\}$, whereas it is trivial for $S = \{1, 2, 3\}$ because any connected graph contains a $\{1, 2, 3\}$-Hamiltonian cycle.

Our work classifies the complexity of this problem for most kinds of sets S, with the key new results being the following: we have NP-completeness for $S = \{2\}$ and for $S = \{2, 4\}$, but tractability for $S = \{1, 2, 4\}$, for $S = \{2, 4, 6\}$, for any superset of these two tractable cases, and for S the infinite set of all odd integers. The remaining open cases are the non-singleton finite sets of odd integers, in particular $S = \{1, 3\}$. Beyond cycles, we also discuss the complexity of finding S-Hamiltonian paths, and show that our problems are all tractable on graphs of bounded cliquewidth.

2012 ACM Subject Classification Mathematics of computing → Graph theory

Keywords and phrases Graph, Cycle, Hamiltonian

Acknowledgements This work was funded in part by the French government under management of Agence Nationale de la Recherche (ANR) as part of the France 2030 program, reference ANR-23-IACL-0008 (PR[AI]RIE-PSAI).

1 Introduction

The Hamiltonian cycle problem is a well-known problem in graph theory and complexity theory. It asks, given an undirected graph, whether it contains a *Hamiltonian cycle*, i.e., a cycle that visits every vertex exactly once. This problem is well-known to be NP-complete, and the same is true if we want to find a *Hamiltonian path*, i.e., a path that visits every vertex exactly once.

Faced with the hardness of this problem, one natural relaxation is to allow greater distance bounds between pairs of consecutive vertices instead of requiring them to be adjacent. The simplest case is to allow hops of length 1 or 2 (instead of just 1). Equivalently, we consider the *square* G^2 of the input graph G, in which we connect any two vertices that can be joined by a path of length at most 2, and we ask whether G^2 is Hamiltonian. It is still NP-complete to determine whether an input graph admits such a relaxed cycle [25], even though many graph families are known whose square is always Hamiltonian: this is the case of two-connected graphs [12, 19] and other families [14, 1, 8, 9, 10]. (We note that there are also other different (and seemingly unrelated) ways to define generalizations of Hamiltonicity, e.g., taking the iterated line graph as in [7, 23].)

The natural next question is to allow hops of 1, 2, or 3. However, it is then known that the cube G^3 of every connected graph G is Hamiltonian [17, 24], and this result is a commonly used trick for efficient enumeration algorithms (the so-called "even-odd trick" [17, 24] (see also [26]). The case of hops of length at most 3 immediately implies that relaxed Hamiltonian cycles always exist in connected graphs for any other set of the form $S = \{1, \ldots, k\}$ with $k > 3$.

However, this does not address the question of other fixed sets of allowed hop lengths. For instance, we can consider $S = \{2\}$: this asks whether we can find a permutation $(v_0, \ldots, v_{n-1})$ of the vertices of the input graph G in which any two consecutive vertices (including the endpoints v_{n-1} and v_0) can be joined by a walk of length *exactly* two. We call such a cycle a *{2}-Hamiltonian cycle*. This question relates to the Hamiltonicity of the *common neighborhood graph* of G [15], i.e., the graph where we connect pairs of vertices that are joined by such walks. Hossein et al. [15, Theorem 4.1] in particular give a sufficient condition for graphs to have a $\{2\}$-Hamiltonian cycle, but they know that the condition is not necessary – so to our knowledge there is no characterization of the graphs admitting a $\{2\}$-Hamiltonian cycle and no complexity bounds on the problem of recognizing them. Alternatively, one can consider, say, $S = \{1, 2, 4\}$: do all graphs admit a $\{1, 2, 4\}$-*Hamiltonian cycle*, i.e., a cycle where any two consecutive vertices are joined by a walk with some length in S?

Contributions. In this paper we study which graphs admit S-Hamiltonian cycles, and the complexity of recognizing them, depending on the fixed set S. We show that, depending on the choice of S, the problem behaves very differently: there are sets S for which every (connected) graph admits an S-Hamiltonian cycle, other cases where the decision problem is NP-hard, and last some cases where the decision problem is non-trivial but efficiently solvable. We entirely classify these behaviors, except for one remaining family of open cases.

The main results come from the sets $S = \{2\}$, $S = \{2, 4\}$, $S = \{2, 4, 6\}$ and $S = \{1, 2, 4\}$. For $S = \{2\}$, we prove that the $\{2\}$-Hamiltonian cycle problem is NP-complete. This proceeds by a reduction from the classical $\{1\}$-Hamiltonian path problem: we reduce it to the $\{2\}$-Hamiltonian path problem with specified endpoints by building the incidence graph of the input graph and attaching a triangle, and we then reduce that problem to the $\{2\}$-Hamiltonian cycle problem by attaching a specific gadget. For $S = \{2, 4\}$, we show that the {2,4}-Hamiltonian cycle problem is NP-hard under Cook reductions. Specifically, we reduce from the {1,2}-Hamiltonian cycle problem (which is known to be NP-hard [25]) and build graphs by attaching a specific gadget on every pair of sufficiently close vertices. As for $S = \{2, 4, 6\}$, we show that a connected graph admits an {2,4,6}-Hamiltonian cycle if and only if it is non-bipartite: this gives a linear-time recognition algorithm, and we show that a witnessing cycle can also be built in linear time when one exists. Finally, for $S = \{1, 2, 4\}$, we prove that every connected graph admits a {1,2,4}-Hamiltonian cycle by first showing the same result on trees. The proof is by induction on the number of vertices, and it also yields a linear-time algorithm to build a witnessing cycle.

The remaining family of unclassified cases are the non-singleton finite sets of odd integers, i.e., the sets $S = \{1, \ldots, 2k + 1\}$ for $k > 1$, for which the complexity of the S-Hamiltonian cycle problem remains open.

We also discuss three variants of the S-Hamiltonian cycle problem, the first being the variants when the set S is infinite, which are all tractable. The second variant is the study of the complexity of S-Hamiltonian path, i.e., we study the complexity of deciding whether an input graph contains an S-Hamiltonian path (possibly with specified endpoints); we show that in most cases this has the same complexity as the S-Hamiltonian cycle problem. Finally, we also study the tractability of the S-Hamiltonian cycle problem when restricting the graphs that are allowed as inputs: we show that, for any fixed finite set S, the problem is tractable when the input graphs are required to have bounded cliquewidth

Paper outline. In Section 2, we introduce the definitions and notations that are used in this paper, and give more formal details about the S-Hamiltonian cycle problem and its variants. In Section 3, we present our main result and a decision tree that delineates the tractable, intractable, and open cases. We then give the detailed proofs for each case. First, Section 4 is dedicated to the proofs of the NP-completeness of the $\{2\}$-Hamiltonian cycle and $\{2,4\}$-Hamiltonian cycle problems. Then, Section 5 deals with the $\{1,2,4\}$-Hamiltonian cycle problem, and Section 6 with the $\{2,4,6\}$-Hamiltonian cycle problem. We then discuss problem variants in Section 7. More specifically, in Section 7.1, we study the complexity of the S-Hamiltonian cycle problem when S is an infinite set. We discuss the complexity for the S-Hamiltonian path variants of our problem in Sections 7.2 and 7.3. Finally, in Section 7.4, we study the case of graphs of bounded cliquewidth. We conclude in Section 8. For lack of space, most detailed proofs are deferred to the appendix.

2 Preliminaries and Problem Statement

Standard graph notions. All graphs in this paper are undirected, do not feature self-loops, are non-empty (i.e., have at least one vertex), and are simple graphs (i.e., without parallel edges). Furthermore, we always assume that graphs are connected. Formally then, a *graph* $G = (V, E)$ consists of a finite set of vertices V (also denoted $V(G)$) and a set of edges E (also denoted $E(G)$) which is a set of pairs of vertices, i.e., the edge $\{u, v\}$ connects the distinct vertices u and v, and we also say the edge is *incident* to u and to v, and that u and v are *adjacent*. A *spanning subgraph* of G is a graph (V, E') with $E' \subseteq E$ that is connected. A *tree* is an acyclic graph, and it is a *rooted tree* if a vertex has been designated as the root. A *spanning tree* T of G is a spanning subgraph of G which is a tree.

A *walk* in a graph G is a sequence of vertices $(v_0, v_1, \ldots, v_k)$ such that for each $i \in \{0, \ldots, k-1\}$ the vertices v_i and v_{i+1} are adjacent. The *length* of the walk is defined as the number of edges traversed, which is k. A *simple path* in a graph G is a walk where all vertices are distinct, and a *simple cycle* is a simple path where the first and the last vertices are adjacent. The length of a cycle is then defined as one more than the length of its defining walk. We require graphs to be *connected*, i.e., for any choice of two vertices there exists a walk that connects them.

The *incidence graph* of $G = (V, E)$ is the graph $G' = (V', E')$ whose vertices are the vertices and edges of G and where each edge of G is connected in G' to its incident vertices; formally $V' := V \uplus E$ and $E' := \{\{u, e\} \mid e \in E, u \in e\}$. The *line graph* of G, denoted $L(G) = (V', E')$, is the graph on the edges of G where two edges of G are connected in G' if they are incident in G to a common vertex; formally $V' := E$ and $E' := \{\{e, e'\} \mid e, e' \in E, e \cap e' \neq \emptyset\}$. We call G *bipartite* if its vertex set can be partitioned into two non-empty sets X and Y, called the *parts*, such that each edge of G connects a vertex in X to a vertex in Y. We recall that a graph is bipartite if and only if it does not admit an odd-length simple cycle, and that checking whether an input graph is bipartite can be done in linear time.

Problem definition. We now formally define the main notions studied in this paper. In the following, we assume S is a non-empty finite subset of $\mathbb{N}^+$, the positive integers (we will study the case of infinite S in Section 7.1).

▶ **Definition 2.1** (**S-path**). *Let $G = (V, E)$ be a graph, and let $S \subseteq \mathbb{N}^+$. A sequence of vertices $P = (v_1, v_2, \ldots, v_n)$ is called an* S-path *of G if, for each $i \in \{1, \ldots, n-1\}$, the vertices v_i and v_{i+1} are connected by a walk of length $\ell \in S$, i.e., there exists some length $\ell \in S$ and some walk in G of length ℓ that starts at v_i and ends at v_{i+1}. Furthermore, P is called a* simple S-path *if all its vertices are distinct.*

▶ **Definition 2.2** (**S-Hamiltonian path** and **S-Hamiltonian cycle**). *Let $G = (V, E)$ be a graph, and let $S \subseteq \mathbb{N}^+$. An* S-Hamiltonian path *of G is a simple S-path that contains all the vertices of G. An* S-Hamiltonian cycle *of G is an S-Hamiltonian path whose first and last vertices are connected by a walk of a length $\ell \in S$ in G.*

For example, the classical Hamiltonian cycle problem corresponds to the case $S = \{1\}$, as two vertices are connected by a walk of length 1 precisely when they are adjacent.

Alternative definitions. Note that our definitions above require the existence of a *walk* of a length in S connecting any two consecutive vertices. The rest of the paper only considers this choice of definition, but we briefly discuss here some alternative possible choices, in order to dispel potential confusion.

First, we do *not* require that the witnessing walk is a *shortest path*. In other words, we do not require that the *distance* between two consecutive vertices u and v belongs to S, only that *some possible walk* has length in S: there may be shorter walks connecting u and v with a length not in S. Requiring shortest paths as connecting walks would be a more stringent requirement in general, even though there are some sets S for which both definitions coincide (e.g., sets of the form $S = \{1, \ldots, k\}$ for some $k > 0$). We note that, already for the set $S = \{2\}$, requiring shortest paths between consecutive vertices would be different, and it would relate to the Hamiltonicity of the *exact-distance square* of the input graph [2]. We do not study this question in the present work.

Second, we do *not* require that the witnessing walks between consecutive vertices are *simple paths*. Requiring simple paths would again be a more stringent requirement in general –though again it would make no difference for sets of the form $S = \{1, \ldots, k\}$, and it would also make no difference for $S = \{2\}$. Again, in this work we do not study this alternative definition requiring simple paths as connecting walks.

Third, we reiterate that our definitions do not pose any restriction on the number of times that a vertex may occur in the connecting walks. In particular, note that vertices that have already been visited earlier in the permutation can still be traversed as intermediate vertices of later connecting walks.

Last, we point out one important consequence of our choice to allow connecting walks that may feature repeated vertices: when two vertices are connected by a walk of length $\ell > 0$, then they are also connected by a walk of length $\ell + 2$, simply by going back-and-forth on the last edge. Thus, in our problem, we can always assume without loss of generality that the set S is *closed under subtraction of* 2, i.e., whenever $\ell \in S$ is an allowed distance and $\ell - 2 > 0$ then the distance $\ell - 2$ is also allowed. In other words, whenever S contains a number ℓ, then we assume that it also contains all smaller numbers of the same parity – for clarity we will always write down these numbers explicitly.

Computational complexity. The main problem we study is the S-Hamiltonian cycle problem for fixed sets S, which asks, given a graph G, whether G admits an S-Hamiltonian cycle. We also study in Section 7 the complexity of two variants the *S-Hamiltonian path problem*, for which we need to decide whether G has an S-Hamiltonian path, and the *S-Hamiltonian path problem with specified endpoints*, which asks whether G admits an S-Hamiltonian path that starts and ends at some specified endpoints given as input. For all these three problems, the complexity is always measured as usual as a function of the input graph. We note that these problems are always in NP for any choice of fixed S: the certificate is the permutation of vertices $s = (v_1, \ldots, v_n)$, and it is easy to check in polynomial time that the certificate is correct. We will show cases when the problem is NP-hard, and other cases when the problem can be efficiently solved: either because it is trivial (i.e., all graphs have an S-Hamiltonian cycle), or because there is an efficient algorithm to decide it. In the efficient cases, we will also study the complexity of the problem of efficiently finding a witness (e.g., an S-Hamiltonian cycle) when one exists, and give tractable algorithms for this task.

For some problems, we will only be able to show NP-hardness under *Cook reductions* (i.e., polynomial-time Turing reductions), rather than the more commonly used *Karp reductions* (i.e., polynomial-time many-one reductions). Recall that a many-one reduction is a reduction that maps each instance of a problem to exactly one instance of another problem whereas a Turing reduction is allowed to use the oracle of the target problem multiple times on different instances. Thus, saying that a problem is NP-complete under Cook reductions is a weaker statement than saying that it is NP-complete under Karp reductions, but it still implies that there is no polynomial-time algorithm for the problem unless P = NP.

3 Main Result

Having formally defined our various problems, we now present in this section the main result of this paper, which classifies the complexity of the S-Hamiltonian cycle problem, up to some remaining open cases. Recall that sets S are considered to be finite, unless stated otherwise – we revisit this choice in Section 7. Results on the other variants (S-Hamiltonian path, with or without specified endpoints), and on the generalization of our problems when S can be infinite, will also be presented in Section 7.

▶ **Theorem 3.1.** *For every finite and non-empty set $S \subseteq \mathbb{N}^+$ the S-Hamiltonian cycle-problem is either NP-complete under Cook reductions, in P, trivial (i.e., true on all graphs), or open, as depicted in Figure 1. Further, when the problem is in P or trivial, we can compute a witnessing S-Hamiltonian cycle in linear time in the input graph.*

We explain in the remaining of this section the roadmap to prove Theorem 3.1. We start by recalling the relevant known results from the literature.

▶ **Proposition 3.2** ([18]). *The {1}-Hamiltonian cycle problem is NP-complete.*

▶ **Proposition 3.3** ([25]). *The {1,2}-Hamiltonian cycle problem is NP-complete.*

The following will be particularly useful to us:

▶ **Proposition 3.4** ([17, 24]). *The {1,2,3}-Hamiltonian cycle problem is trivial. More precisely, for every graph G and $u, v \in V(G)$ with $u \neq v$, there exists a {1,2,3}-Hamiltonian path from u to v in G.*

Now, we present the new results that complete the classification of Theorem 3.1.

▶ **Theorem 3.5.** *The {2}-Hamiltonian cycle problem is NP-complete.*

▶ **Theorem 3.6.** *The {2,4}-Hamiltonian cycle problem is NP-complete under Cook reductions.*

▶ **Theorem 3.7.** *Every connected graph admits a {1,2,4}-Hamiltonian cycle. Furthermore, such a cycle can be found in linear time.*

▶ **Theorem 3.8.** *The graphs that have a {2,4,6}-Hamiltonian cycle are exactly the non-bipartite graphs, which implies that the {2,4,6}-Hamiltonian cycle problem can be solved in linear time. Furthermore, on graphs having a {2,4,6}-Hamiltonian cycle, we can construct one in linear time.*

From these results, we can decide the complexity of any set S thanks to the decision process illustrated by the decision tree in Figure 1, where we assume without loss of generality that S is closed under subtraction of 2. In this decision tree, we denote by P that the problem corresponding to the sets S considered is polynomial, by *NP-c* that it is NP-complete (possibly under Cook reductions), by *Trivial* that every graph admits such a cycle, and by *open* that we do not know yet. The new results are in the bold nodes. The detailed proof of the correctness of that decision process, which is also a proof of Theorem 3.1, is in Appendix A. That said, the only non-obvious reasoning in this proof is the following. First, for sets containing 6, we know by Theorem 3.8 that the lengths 2, 4, and 6 are enough to find a S-Hamiltonian cycle in any non-bipartite graph. On the other hand, having more even lengths will never allow to get a S-Hamiltonian cycle in a bipartite graph for obvious reasons, thus sets S with only even numbers that are supersets of {2,4,6} can be grouped together. Second, all sets S that are supersets of {1,2,3} can be grouped together because we know by Proposition 3.4 that the lengths $1, 2$ and 3 are enough to find an S-Hamiltonian cycle in any graph. Third, all sets S that are supersets of {1,2,4} can be grouped together for the same reason thanks to Theorem 3.7.

The next three sections are dedicated to proving the four main theorems of this section.

4 NP-Complete Cases

In this section, we start the proof of Theorem 3.1 by giving the NP-hardness results for the cases $S = \{2\}$ (in Section 4.1) and $S = \{2, 4\}$ (in Section 4.2).

4.1 The Case $S = \{2\}$

To show that the {2}-Hamiltonian cycle problem is NP-hard, we proceed in two steps. We first reduce the {1}-Hamiltonian path problem with specified endpoints which is known to be NP-complete [18], to the {2}-Hamiltonian path problem with specified endpoints. Then, we reduce the latter problem to the {2}-Hamiltonian cycle problem. Let us show the first step:

▶ **Theorem 4.1.** *The {2}-Hamiltonian path problem with specified endpoints is NP-complete.*

Proof sketch. Let G be a graph and $\alpha \neq \beta$ be the specified endpoints in $V(G)$ for the Hamiltonian path problem. Construct a new graph H by first taking the incidence graph of G (which features an *edge-vertex* for each edge of G) and then connecting β to a dangling triangle. Pick an arbitrary edge-vertex β' adjacent to α in H. We can then conclude by establishing the following: G has a {1}-Hamiltonian path connecting α and β if and only if H has a {2}-Hamiltonian path connecting α and β'.

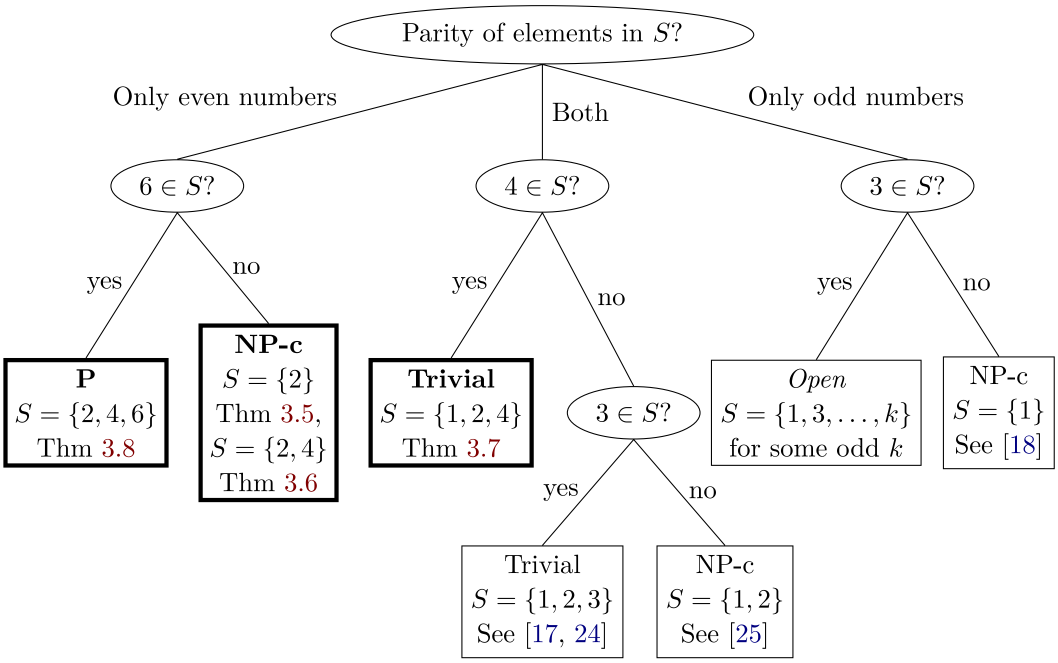

Figure 1 Decision tree for the complexity of the S-Hamiltonian cycle problem for S a finite and non-empty subset of $\mathbb{N}^+$.

In particular, for the forward direction, to obtain a $\{2\}$-Hamiltonian path in H we first follow the original $\{1\}$-Hamiltonian path of G, which is now a simple $\{2\}$-path from α to β in H, then goes through the triangle to reach an edge-vertex. From there, we use a strengthening of [3, Theorem 6.5.4] that states that the line graph of a Hamiltonian graph is Hamiltonian, which allows us to visit all edge-vertices. ◀

We now show the hardness of the $\{2\}$-Hamiltonian cycle problem by reducing from the $\{2\}$-Hamiltonian path with specified endpoints:

▶ **Theorem 3.5.** *The $\{2\}$-Hamiltonian cycle problem is NP-complete.*

Proof sketch. Let G be an input graph and $\alpha, \beta \in V(G)$ be two different vertices. The reduction is to attach the gadget D from Figure 2 to a copy of G by adding the edges $\{\alpha, d_0\}$ and $\{\beta, d_1\}$. We then prove that G has a $\{2\}$-Hamiltonian path from α to β if and only if the new graph H has a $\{2\}$-Hamiltonian cycle. This is done by observing that the vertices d_3, d_5, d_7 of the gadget D have very constrained neighborhoods in any $\{2\}$-Hamiltonian cycle of H, which forces a specific ordering of the vertices of D in such a cycle. ◀

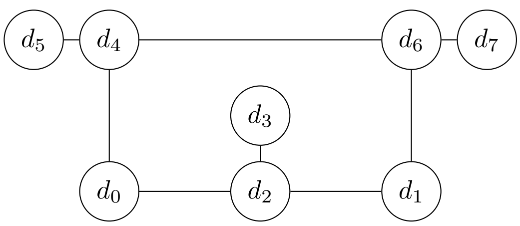

Figure 2 Gadget D for the proof of Theorem 3.5, found with the help of a computer program.

4.2 The Case $S = \{2, 4\}$

We now turn to the case of the {2,4}-Hamiltonian cycle problem. We show NP-hardness via a Cook reduction from the {1,2}-Hamiltonian cycle problem; formally:

▶ **Theorem 3.6.** *The {2,4}-Hamiltonian cycle problem is NP-complete under Cook reductions.*

Our hardness proof uses Cook reductions for their ability to reduce an instance to multiple instances with multiple oracle calls. While we suspect that {2,4}-Hamiltonian cycle should also be NP-hard under the more standard notion of Karp reductions, we have not been able to show this.

Theorem 3.6 will directly follow from the following result.

▶ **Proposition 4.2.** *Given an input graph G and two vertices $x \neq y$ at distance at most 2 in G, we can construct in polynomial time a graph $H_{x,y}$ such that:*

- *if G has a {1,2}-Hamiltonian cycle where x and y are consecutive, then $H_{x,y}$ has a {2,4}-Hamiltonian cycle,*
- *if $H_{x,y}$ has a {2,4}-Hamiltonian cycle, then G has a {1,2}-Hamiltonian cycle.*

Note that, when $H_{x,y}$ has a {2,4}-Hamiltonian cycle, then the {1,2}-Hamiltonian cycle in G is not necessarily one where x and y are consecutive. Nevertheless, the result of Proposition 4.2 suffices to design a Cook reduction that shows the NP-hardness of the {2,4}-Hamiltonian cycle problem:

Proof of Theorem 3.6. We do a Cook reduction from the {1,2}-Hamiltonian cycle problem, which as we already mentioned in Proposition 3.2 is NP-hard. Given an input G to the {1,2}-Hamiltonian cycle problem, assuming without loss of generality that it has at least two vertices, we consider every pair x, y of vertices at distance at most 2 in G and construct the corresponding graph $H_{x,y}$: this amounts to polynomially many graphs which are all built in polynomial time. Now, if G has a {1,2}-Hamiltonian cycle C then for any choice of contiguous vertices x, y of C, we know that $H_{x,y}$ has a {2,4}-Hamiltonian cycle. Conversely, if some $H_{x,y}$ has a {2,4}-Hamiltonian cycle then G has a {1,2}-Hamiltonian cycle. This implies that G has a {1,2}-Hamiltonian cycle if and only if some $H_{x,y}$ has a {2,4}-Hamiltonian cycle, so we can indeed solve {1,2}-Hamiltonian cycle in polynomial time with access to an oracle for {2,4}-Hamiltonian cycle. This establishes the Cook reduction and concludes the proof. ◀

Thus, all that is left to show Theorem 3.6 is to prove Proposition 4.2. To do so, we will need the following intermediate lemma which is a stronger version of [20, Theorem 2]. This lemma originally proves that the square of a line graph is always Hamiltonian, i.e., that the line graph of any (connected) graph always has a {1,2}-Hamiltonian cycle, and we need a slightly stronger version that allows us to choose two vertices of the line graph to be consecutive in the {1,2}-Hamiltonian cycle.

▶ **Lemma 4.3** (Strengthening of Theorem 2 of [20])**.** *Let G be a graph with at least 3 edges. Let $L(G)$ be the line graph of G. Let α and β be two vertices of $L(G)$ connected by a walk of length 2 in $L(G)$. Then there exists a {1,2}-Hamiltonian cycle in $L(G)$ where α and β are consecutive.*

With this lemma in hand, we are now ready to prove Proposition 4.2.

Proof sketch of Proposition 4.2. Given an input graph G on which we want to determine the existence of a {1,2}-Hamiltonian cycle and two of its vertices x and y, we build $H_{x,y}$ by taking the incidence graph of G (formed of node-vertices and edge-vertices) and attaching

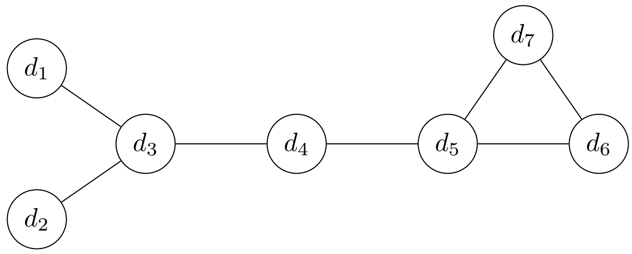

Figure 3 Gadget D for the proof of Proposition 4.2

the gadget from Figure 3 with edges from d_1 and d_2 respectively to the vertices x' and z' corresponding to x and to a vertex z which is a neighbor of x towards y (specifically, we take $z = y$ if x and y are at distance 1, and we take z to be a common neighbor of x and y if they are at distance 2).

For the forward direction, we show that a {1,2}-Hamiltonian cycle C of G where x and y are adjacent can be lifted to a {2,4}-Hamiltonian cycle C' of $H_{x,y}$. At a high level, C shows us how to visit the node-vertices of $H_{x,y}$, and we use D to toggle to the edge-vertices and back. Further, we use Lemma 4.3 to build the portion of the cycle that visits the edge-vertices of $H_{x,y}$ while starting and ending at appropriate vertices.

For the backward direction, given a {2,4}-Hamiltonian cycle C' of $H_{x,y}$, we use the fact that $H_{x,y}$ is almost bipartite except for the triangle $\{d_5, d_6, d_7\}$ in D. This allows us to show that C' can only switch one time from node-vertices to edge-vertices and one time back, so that C' must contain a contiguous {2,4}-Hamiltonian path N that visits all the node-vertices, which is preceded and succeeded by visits to D. Further considerations on the structure of $H_{x,y}$ ensure that N must in fact be a {1,2}-Hamiltonian cycle of G, i.e., its starting and ending vertices correspond to vertices at a distance of at most 2 in G. ◀

5 Trivial Case: $S = \{1, 2, 4\}$

In this section we show that every graph admits a {1,2,4}-Hamiltonian cycle which we can construct in linear time.

▶ **Theorem 3.7.** *Every connected graph admits a {1,2,4}-Hamiltonian cycle. Furthermore, such a cycle can be found in linear time.*

Up to taking a spanning tree, this is equivalent to showing that every tree admits a {1,2,4}-Hamiltonian cycle, which we do in the following lemma (with an additional property that is needed for its own proof).

▶ **Lemma 5.1.** *Let T be a tree rooted at r such that $|V(T)| \geq 2$. Then there exists a {1,2,4}-Hamiltonian cycle $(r, v_1, \ldots, v_k)$ of T such that v_1 is a child of r. Furthermore, such a cycle can be computed in linear time.*

Proof sketch. The idea of the proof is to do an induction on the number of vertices of T. We distinguish the children of r in T between those children that are leaves and those children that are roots of non-trivial subtrees. If there are leaf children, then we can simply visit them in order with a 1-hop from r and 2-hops between each of them, easily ensuring the requirement that the second vertex of the cycle is a child of r. For each non-leaf child v_i of r, we use the induction hypothesis to get a cycle C_i which visits the subtree of T rooted at v_i while starting at v_i and visiting a child of v_i immediately afterwards. Up to conjugating these cycles and reversing every other cycle, we can stitch these cycles together to visit the subtrees rooted at non-leaf children of r: we start at the first such vertex v_1 (reached via a 1-hop from r, or a 2-hop from the last leaf child of r), follow C_1 (after conjugation and reversal), and finish on a child w_1 of v_1, then we do 4-hop to a child w_2 of v_2 and do C_2 in reverse order to finish on v_2. At the end of this process, we finish either on the last non-leaf child of r or a child thereof, and we can close the cycle and go back to r with a 1-hop or a 2-hop.

This inductive proof can be transformed into a linear-time algorithm that constructs the desired {1,2,4}-Hamiltonian cycle of T. ◀

6 Polynomial Case: $S = \{2, 4, 6\}$

In this section, we show that the {2,4,6}-Hamiltonian cycle problem can be solved in linear time because the graphs with a {2,4,6}-Hamiltonian cycle are precisely the non-bipartite graphs:

▶ **Theorem 3.8.** *The graphs that have a {2,4,6}-Hamiltonian cycle are exactly the non-bipartite graphs, which implies that the {2,4,6}-Hamiltonian cycle problem can be solved in linear time. Furthermore, on graphs having a {2,4,6}-Hamiltonian cycle, we can construct one in linear time.*

To prove this theorem, we will first show that every non-bipartite graph admits a {2,4,6}-Hamiltonian cycle.

▶ **Lemma 6.1.** *Given a non-bipartite graph G, we can construct a {2,4,6}-Hamiltonian cycle of G in linear time.*

Proof sketch. Let G be a non-bipartite graph. Then G must contain a simple cycle of odd length. Let $C = (v_0, v_1, \ldots, v_{2k})$ be a shortest odd cycle in G. We construct a spanning forest rooted at each v_i by performing a simultaneous breadth-first search starting from all v_i. This way, we obtain a collection of trees $T_0, T_1, \ldots, T_{2k}$, each rooted at their respective v_i, and covering all vertices of G. We can use Proposition 3.4 about the triviality of the {1,2,3}-Hamiltonian path problem to create, in each T_i, a {1,2,3}-path of the tree obtained from taking only vertices at odd depth (resp., at even depth) in T_i, which translates to a {2,4,6}-path of those vertices in T_i. We traverse the cycle C twice. During the first traversal, we visit the vertices v_i with even index i, and during the second traversal, we visit the vertices v_i with odd index i. When visiting a vertex v_i, we also enumerate all vertices of T_i at even depth, and all vertices of T_{i+1} at odd depth. (The indices are taken modulo $2k + 1$.) This yields a {2,4,6}-Hamiltonian cycle of G. ◀

We can now prove Theorem 3.8, which concludes the proof of Theorem 3.1:

Proof of Theorem 3.8. Given a graph G, we first determine in linear time whether G is bipartite. If G is bipartite, then every walk of length in $S = \{2, 4, 6\}$ will keep us in the same part of the bipartition, so we cannot hope to ever visit the other side. Otherwise, Lemma 6.1 gives us a {2,4,6}-Hamiltonian cycle in linear time. ◀

7 Variants

In this section we consider variants of the S-Hamiltonian cycle problem studied in Theorem 3.1. We start with the case of infinite sets S, for which we can easily show that the S-Hamiltonian cycle problem is either trivial or solvable in linear time. We then move from cycles to paths, and discuss the complexity of the S-Hamiltonian path problem and of the S-Hamiltonian path problem with specified endpoints: we talk in more detail about the case $S = \{1, 2, 4\}$ for paths with specified endpoints, which requires some new reasoning. Last, we discuss the question of whether we can make S-Hamiltonian cycle tractable restricting the class of input graphs.

7.1 Infinite Sets

In this section, we discuss the complexity of the S-Hamiltonian cycle problem for infinite sets S. The crucial observation is that most infinite sets S closed under subtraction of 2 are trivial, because they contain either $\{1, 2, 3\}$ or $\{1, 2, 4\}$. (We give the formal details about this classification in Appendix E.1.) The two non-trivial sets are the set of all even integers, and the set of all odd integers. The case of the set of all even integers follows easily from Theorem 3.8. As for the set of all odd integers, we give its characterization below:

▶ **Theorem 7.1.** *Let S denote the set of all odd integers, and let G be a graph. Then G admits an S-Hamiltonian cycle if and only one of the two following cases holds: G is non-bipartite, or G is bipartite with parts of equal cardinality.*

Proof sketch. Let G be a graph. If G is non-bipartite, then we can connect any two vertices by a walk of odd length. Indeed, we can go through an odd cycle as many times as needed to adjust the parity of the length of that walk. Hence, any ordering of the vertices of G is an S-Hamiltonian cycle.

If G is bipartite, the only odd-length walks are the ones connecting vertices from one part to the other. Hence, in order to have an S-Hamiltonian cycle, the ordering of the vertices must alternate between the two parts, which implies that the two parts must have equal size. Conversely, if the two parts have equal size, then any ordering that alternates between the two parts is an S-Hamiltonian cycle because any pair of vertices from different parts are connected by a walk of odd length. ◀

7.2 Hamiltonian Path Variants

In this section, we discuss the complexity of the two variants of the S-Hamiltonian cycle problem introduced in Section 2: the S-Hamiltonian path problem, and the S-Hamiltonian path problem with specified endpoints. The complexity of these problems on the various possible finite sets S does not always follow from our results on the S-Hamiltonian cycle problem, so we explain the situation in this subsection and in the next subsection.

First, for any set S, it is clearly always possible to reduce from the S-Hamiltonian path problem or from the S-Hamiltonian cycle problem to the S-Hamiltonian path problem with specified endpoints, under Cook reductions. The former is achieved by calling the oracle for S-Hamiltonian path with specified endpoints on all pairs of vertices, and the latter is achieved by calling that oracle on all pairs of vertices connected by a walk having a length in S. However, some of the proofs presented in this paper can be modified to show a better complexity for the S-Hamiltonian path problems.

As most of these problems behave the same way as their S-Hamiltonian cycle counterpart, we will highlight some cases here: we present the detailed explanations of every set for both variants (and their proofs when applicable) in Appendix E.2. The only result that has been proved to be different is for the {1,2,4}-Hamiltonian path with specified endpoints. Whereas {1,2,4}-Hamiltonian cycle was trivial (Theorem 3.7), for {1,2,4}-Hamiltonian path with specified endpoints the problem can be solved in linear time by excluding a small explicit family of negative instances. We present this result in the next section.

The only other sets S where the complexities of the path and cycle problems are not known to coincide is $S = \{2, 4\}$ and $S = \{2, 4, 6\}$, where we leave some problems open: see Appendix E.2 for more details.

7.3 {1,2,4}-Hamiltonian Path Problem with Specified Endpoints

In this section, we prove that the {1,2,4}-Hamiltonian path problem with specified endpoints can be solved in linear time. For those proofs, we will need to introduce the notions of *star* and *bistar*. A *star centered at a vertex* x is a graph with at least two vertices such that every vertex different from x is adjacent only to x. A *bistar centered at vertices* a *and* b is a graph such that a and b are adjacent and every vertex different from a and b is adjacent to only either a or b (but not both), and both a and b have degree at least 2. These two definitions are illustrated in Figure 4. We now introduce a lemma that will help us prove the desired theorem.

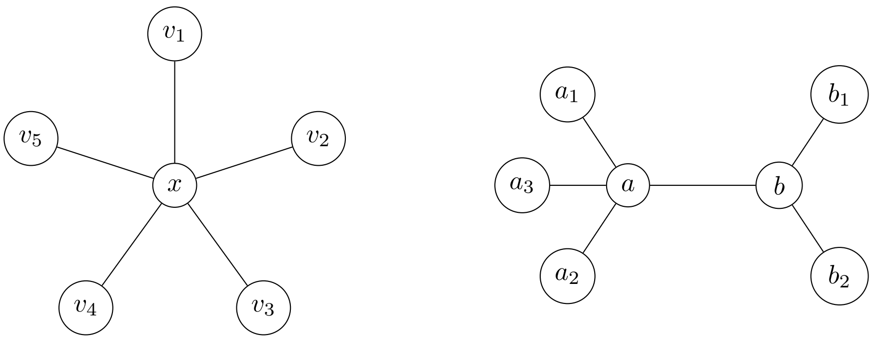

Figure 4 A star centered at x (left) and a bistar centered at a and b (right).

▶ **Lemma 7.2.** *Let T be a tree with $|V(T)| > 2$ and $a, b \in V(T)$ two vertices of T such that a and b are not adjacent. Then there exists a {1,2,4}-Hamiltonian path between a and b.*

Proof sketch. The proof is done by induction on the number of vertices of T. The idea is to divide the tree into multiple vertex disjoint subtrees depending on the shape of T, and to use the induction hypothesis as well as Lemma 5.1 to construct simple $\{1, 2, 4\}$-paths of these subtrees, which can then be concatenated to form the desired {1,2,4}-Hamiltonian path between a and b. ◀

We can now show our characterization for the {1,2,4}-Hamiltonian path problem with specified endpoints:

▶ **Theorem 7.3.** *Let G be a graph and a and b be two different vertices of G. Then G admits a {1,2,4}-Hamiltonian path from a to b if and only if G is not a bistar centered at a and b.*

Proof sketch. Let G be a graph and a and b be two different vertices of G. If they are non-adjacent, we conclude by Lemma 7.2. If a and b are adjacent and G is not a bistar, then up to exchanging a and b we find a neighbor v_1 of a which is not b and which has itself a neighbor v_2 which is also not b. We take a spanning tree of G including the edges $\{a, v_1\}$ and $\{v_1, v_2\}$, and we use Lemma 7.2 and finally obtain a {1,2,4}-Hamiltonian path of G.

Now, if G is a bistar, a path from a can start from neighbors of a (resp. neighbors of b) but will not be able to reach neighbors of b (resp. a) without taking b (which must be at the end). Thus a bistar centered at a and b cannot contain a {1,2,4}-Hamiltonian path from a to b, which concludes the proof. ◀

7.4 Restricted Graph Classes

In this section, we last turn to one last question: can the S-Hamiltonian cycle problem be made tractable if we assume that the input graphs are in restricted graph classes? This is motivated by the fact that the Hamiltonian cycle problem, while NP-hard in general, is known to be tractable on many restricted classes of graphs, in particular bounded-treewidth and bounded-cliquewidth graphs [4, 11]. Hence, one can ask whether the same is true of the S-Hamiltonian cycle problem for sets S other than $S = \{1\}$.

There are related results in this direction for the case $S = \{1, 2\}$ and when the input graph is a tree. Namely, it is known that a tree T has a {1,2}-Hamiltonian cycle if and only if it is a so-called caterpillar graph [13]. Further, the trees having a {1,2}-Hamiltonian path have been characterized in [22] and can be recognized in linear time.

We now claim that, for arbitrary finite sets S, tractability extends beyond trees to bounded-treewidth and even bounded-cliquewidth graphs. Namely:

▶ **Theorem 7.4.** *For any fixed finite set $S \subseteq \mathbb{N}^+$ and bound $k \in \mathbb{N}$, the S-Hamiltonian cycle problem on input graphs of cliquewidth at most k is in polynomial time.*

Proof sketch. The proof uses the fact that bounded cliquewidth is preserved by MSO interpretations [6, 5]. So, from the input graph G, we construct the graph G_S that has the same vertices as G and connects two vertices if and only if they are connected by a walk of a length in S: this is definable by an MSO interpretation, so G_S also has bounded cliquewidth. It is now easy to see that G_S admits a Hamiltonian cycle if and only if G admits an S-Hamiltonian cycle. From there, thanks to the result from [4], there exists an XP algorithm to solve the Hamiltonian cycle problem on graphs with bounded cliquewidth, which allows us to conclude. ◀

However, bounded cliquewidth is not the end of the story for the Hamiltonian cycle problem, as it is also known to be tractable on some classes of unbounded cliquewidth, e.g., proper interval graphs [16, Theorem 5]. Hence, a natural question for future work would be to explore whether the S-Hamiltonian cycle problem is also tractable on other classes of unbounded cliquewidth.

8 Conclusion and Future Work

We have studied the S-Hamiltonian cycle problem and we determined its complexity for most finite sets S. We classified the complexity of every result into a decision tree that allows to restrict our attention to fewer cases. Among those cases, some were known results, and others have been proven in this paper, namely the NP-completeness when $S = \{2\}$ and $S = \{2, 4\}$, the triviality when $S = \{1, 2, 4\}$, and the polynomial-time solvability when $S = \{2, 4, 6\}$. Furthermore, we also addressed the case of infinite sets S, which are either trivial or polynomial-time solvable; and we also studied the variants of paths, with or without specified endpoints, for which we obtained similar complexity classifications for most cases.

We now mention potential future work. The most immediate open question is to classify the complexity of the S-Hamiltonian cycle problem for $S = \{1, 3\}$ (and of other finite sets of odd integers): we suspect that this problem is NP-hard, but we have not been able to show it. However, it is important to notice that in case of bipartite graphs (and for the same reasons as in Theorem 7.1 on the infinite set of all odd numbers), a necessary condition to admit an S-Hamiltonian cycle for a finite set S of odd integers is to have both parts of the same size. That is nonetheless not a sufficient condition, as there may be an "imbalance" of vertices of a certain part of the bipartition in some local part of the graph, that hinders the alternation of parts in our cycle. Yet, we do not know how to conclude from here, and the complexity of S-Hamiltonian cycle remains open, even when the input graph is restricted to be bipartite, or to be non-bipartite. Concerning the path variants of our problems, obviously the same as above holds true for S a finite set of only odd numbers, but there are also the cases of $\{2,4\}$-Hamiltonian path and $\{2,4,6\}$-Hamiltonian path with specified endpoints that we have left open.

Another direction would be to study the alternate definition that does not allow back-and-forth travels in intermediate walks of S-paths. This definition is harder to grasp, and there are a lot more sets to consider because we can no longer assume without loss of generality that S is closed by subtraction of 2. One could also consider the more stringent definition that requires the back-and-forth travels to form simple paths, or those where the lengths in S are required to correspond to distances between the two vertices (i.e., prohibiting the existence of shorter paths).

Last, another natural question would be to look into directed graphs. However, it seems that for any given finite set S, there exist arbitrarily large graphs that do not contain an S-Hamiltonian cycle. This suggests that the general complexity answers for directed graphs might be totally different from those on undirected graphs.

References

1 Mohamed El Kadi Abderrezzak, Evelyne Flandrin, and Zdeněk Ryjáček. Induced $S(K_{1,3})$ and hamiltonian cycles in the square of a graph. *Discrete Mathematics*, 207(1-3), 1999. doi:10.1016/S0012-365X(99)00115-6.

2 Yandong Bai, Pedro P. Cortés, Reza Naserasr, and Daniel A. Quiroz. Characterizing and recognizing exact-distance squares of graphs. *Discrete Mathematics*, 347, 2024. doi:10.1016/j.disc.2023.113493.

3 R. Balakrishnan and K. Ranganathan. *A Textbook of Graph Theory*. Springer New York, 2012. doi:10.1007/978-1-4614-4529-6.

4 Benjamin Bergougnoux, Mamadou Moustapha Kanté, and O-joung Kwon. An optimal XP algorithm for Hamiltonian cycle on graphs of bounded clique-width. *Algorithmica*, 82(6), 2020. doi:10.1007/s00453-019-00663-9.

5 Achim Blumensath and Bruno Courcelle. Recognizability, hypergraph operations, and logical types. *Information and Computation*, 204(6), 2006. doi:10.1016/j.ic.2005.11.006.

6 Édouard Bonnet, Jan Dreier, Jakub Gajarskỳ, Stephan Kreutzer, Nikolas Mählmann, Pierre Simon, and Szymon Toruńczyk. Model checking on interpretations of classes of bounded local cliquewidth. In *Proceedings of the 37th Annual ACM/IEEE Symposium on Logic in Computer Science*, 2022. doi:10.1145/3531130.3533367.

7 Gary Chartrand. On Hamiltonian line-graphs. *Transactions of the American Mathematical Society*, 134(3), 1968. doi:10.2307/1994875.

8 Jan Ekstein. Hamiltonian cycles in the square of a graph. *The Electronic Journal of Combinatorics*, 18(1), 2011. URL: https://www.combinatorics.org/ojs/index.php/eljc/article/view/v18i1p203, doi:10.37236/690.

9 Jan Ekstein and Herbert Fleischner. A best possible result for the square of a 2-block to be Hamiltonian. *Discrete Mathematics*, 344(1), 2021. doi:10.1016/j.disc.2020.112158.

10 Jan Ekstein and Herbert Fleischner. The most general structure of graphs with Hamiltonian or Hamiltonian connected square. *Discrete Mathematics*, 347(1), 2024. doi:10.1016/j.disc.2023.113702.

11 Wolfgang Espelage, Frank Gurski, and Egon Wanke. How to solve NP-hard graph problems on clique-width bounded graphs in polynomial time. In *Proceedings of the International Workshop on Graph-Theoretic Concepts in Computer Science*. Springer, 2001. doi:10.1007/3-540-45477-2_12.

12 Herbert Fleischner. The square of every two-connected graph is Hamiltonian. *Journal of Combinatorial Theory, Series B*, 16, 1974. doi:10.1016/0095-8956(74)90091-4.

13 Frank Harary and Allen Schwenk. Trees with Hamiltonian square. *Mathematika*, 18, 1971. doi:10.1112/s0025579300008494.

14 George Hendry and Walter Vogler. The square of a connected $S(K_{1,3})$-free graph is vertex pancyclic. *Journal of Graph Theory*, 9(4), 1985. doi:10.1002/jgt.3190090415.

15 Samaneh Hossein-Zadeh, Ali Iranmanesh, Asma Hamzeh, and Mohammad Ali Hosseinzadeh. On the common neighborhood graphs. *Electronic Notes in Discrete Mathematics*, 45:51–56, 2014. doi:10.1016/j.endm.2013.11.011.

16 Louis Ibarra. A simple algorithm to find Hamiltonian cycles in proper interval graphs. *Information Processing Letters*, 109, 2009. doi:10.1016/j.ipl.2009.07.010.

17 Jerome J. Karaganis. On the cube of a graph. *Canadian Mathematical Bulletin*, 11(2), 1968. doi:10.4153/CMB-1968-037-0.

18 R. M. Karp. On the computational complexity of combinatorial problems. *Networks*, 5, 1975. doi:10.1002/net.1975.5.1.45.

19 Janina Müttel and Dieter Rautenbach. A short proof of the versatile version of Fleischners theorem. *Discrete Mathematics*, 313(19), 2013. doi:10.1016/j.disc.2012.07.032.

20 Ladislav Nebeský. On the line graph of the square and the square of the line graph of a connected graph. *Časopis pro pěstování matematiky*, 098, 1973. doi:10.21136/cpm.1973.117794.

21 Sang-il Oum and Paul Seymour. Approximating clique-width and branch-width. *Journal of Combinatorial Theory, Series B*, 96(4), 2006. doi:10.1016/j.jctb.2005.10.006.

22 Jakub Radoszewski and Wojciech Rytter. Hamiltonian paths in the square of a tree. In *Proceedings of the International Symposium on Algorithms and Computation*. Springer, 2011.

23 Eminjan Sabir and Elkin Vumar. Spanning connectivity of the power of a graph and Hamilton-connected index of a graph. *Graphs and Combinatorics*, 30(6), 2014. doi:10.1007/s00373-013-1362-4.

24 Milan Sekanina. On an ordering of the set of vertices of a connected graph. *Publ. Fac. Sci. Univ. Brno*, 412, 1960.

25 P. Underground. On graphs with Hamiltonian squares. *Discrete Mathematics*, 21, 1978. doi:10.1016/0012-365x(78)90164-4.

26 Takeaki Uno. Two general methods to reduce delay and change of enumeration algorithms. *National Institute of Informatics (in Japan) Technical Report E*, 4, 2003. URL: https://repository.nii.ac.jp/record/385/files/03-004E.pdf.

A Proofs for Section 3 (Main Result)

▶ **Theorem 3.1.** *For every finite and non-empty set $S \subseteq \mathbb{N}^+$ the S-Hamiltonian cycle-problem is either NP-complete under Cook reductions, in P, trivial (i.e., true on all graphs), or open, as depicted in Figure 1. Further, when the problem is in P or trivial, we can compute a witnessing S-Hamiltonian cycle in linear time in the input graph.*

Proof. Given a finite and non-empty set S, we assume without loss of generality that it is closed by subtraction of 2. We then distinguish three cases: either S contains only even numbers, or it contains even and odd numbers, or it contains only odd numbers.

Case 1: S contains only even numbers. Then there are two cases: either S contains 6 or it does not. First, if $6 \in S$, then the characterization is that an S-Hamiltonian cycle exists in a given graph G if and only if G is non-bipartite. Specifically, if G is bipartite, then the allowed lengths in S will not allow us to move from one part of the bipartition to the other, so there is no S-Hamiltonian cycle. Conversely, if G is not bipartite, then we show an algorithm in Theorem 3.8 that efficiently creates a {2,4,6}-Hamiltonian cycle, which is also an S-Hamiltonian cycle. Second, if $6 \notin S$, then we have either $S = \{2\}$ or $S = \{2, 4\}$ and both cases are proved to be NP-complete in Theorem 3.5 and Theorem 3.6 respectively.

Case 2: S contains even and odd numbers. In this case, S must contain 1 and 2. If those are the only elements of S, then the {1,2}-Hamiltonian cycle problem is NP-complete, as stated in Proposition 3.3. Otherwise, S must additionally contain either 3 or 4, in which case every graph contains an S-Hamiltonian cycle. Indeed, if S also contains 3, we saw in Proposition 3.4 that every graph admits a {1,2,3}-Hamiltonian cycle, which must also be an S-Hamiltonian cycle. Otherwise, if S also contains 4, then we prove in Theorem 3.7 that every graph admits a {1,2,4}-Hamiltonian cycle which is an S-Hamiltonian cycle as well.

Case 3: S contains only odd numbers. Then if $S = \{1\}$ we know that the problem is NP-complete as stated in Proposition 3.2. Otherwise, $S = \{1, \dots, 2k + 1\}$ with $k > 1$ and this is the case that our work leaves open. ◀

B Proofs for Section 4 (NP-Complete Cases)

B.1 Proofs for $S = \{2\}$

▶ **Theorem 4.1.** *The {2}-Hamiltonian path problem with specified endpoints is NP-complete.*

Proof. We prove the claim by reduction from the {1}-Hamiltonian path problem with specified endpoints, which is NP-hard [18].

Let G be the input graph, and let $\alpha \neq \beta$ be the specified endpoints. We can assume without loss of generality that G has at least two vertices, as otherwise the answer is trivial. Without loss of generality, up to modifying G, we can assume that α and β are not adjacent. Indeed, if they are, we can add a new vertex to G connected only to β and make it the new β, which does not change the existence of a {1}-Hamiltonian path in G.

We will build in polynomial time a new graph H and choose two vertices $x \neq y$ in $V(H)$ so as to ensure the following equivalence (*): G has a {1}-Hamiltonian path from α to β if and only if H has a {2}-Hamiltonian path from x to y. The construction of H is as follows: Take the incidence graph G' of G and connect β to a triangle (Figure 5). Formally:

- Build the incidence graph G' by doing the following:
  - For each vertex $v \in V(G)$, add a corresponding vertex $v' \in G'$. We call these *node-vertices.*

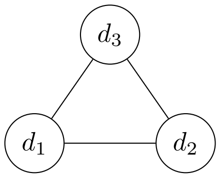

Figure 5 Triangle gadget D used in the proof of Theorem 4.1

- For each edge $\{u, v\} \in E(G)$, add a new vertex w_{uv} to G' and connect it to both u' and v'. We call these *edge-vertices*.
- Let H be G' together with a copy of the triangle gadget D from Figure 5 and an edge $\{\beta', d_1\}$ connecting d_1 to the vertex β' of H that corresponds to β in G.

Last, define $x \coloneqq \alpha'$, and define y to be any edge-vertex adjacent to x in H. (Such an edge-vertex necessarily exists because the input graph G is connected so α must have some incident edge in G.)

The construction of H given above is clearly in polynomial time, so all that remains is to establish correctness, i.e., proving the equivalence (*) stated above. Before diving into the proof, we will first see which pairs of vertices are connected by a walk of length 2 in H. One can easily check that they are precisely the following:

- Node-vertices whose originals are adjacent in G: $u', v' \in V(H)$ where $uv \in E(G)$, with the walk $u' - w_{uv} - v'$ of length 2.
- Edge-vertices whose original edges share a common vertex: w_{uv}, w_{vt} with $u, v, t \in V(G)$ and $uv, vt \in E(G)$, via the intermediate vertex $w_{uv} - v' - w_{vt}$.
- Special cases thanks to D:
  - all vertices of $\{d_1, d_2, d_3\}$ are connected to each other;
  - β' is connected to both d_2 and d_3;
  - d_1 is connected (through β') to every edge-vertex corresponding to an edge of G incident to β.

We will now show equivalence (*). We first prove the forward direction. Suppose that G contains a $\{1\}$-Hamiltonian path $P = (v_1, v_2, \ldots, v_k)$ such that $\alpha = v_1$ and $\beta = v_k$. We need to prove that H contains a $\{2\}$-Hamiltonian path from x to y. To achieve this, we will construct a $\{2\}$-Hamiltonian path P' in H by starting from the node-vertices corresponding to the vertices in P and then adding the vertices of D and the edge-vertices. We start by initializing $P' = (v'_1, \ldots, v'_k)$. We next add the vertices of D, which will allow us to reach an edge-vertex: $P' = (v'_1, \ldots, v'_k, d_2, d_3, d_1)$.

To complete the definition of P', we will now sort the edge-vertices into sets depending on their latest adjacent node-vertex according to P. We will then add all those sets in reverse order of their latest adjacent node-vertex in P. This is possible because the line graph of a Hamiltonian graph is also Hamiltonian [3], and because the graph that connects two edge-vertices if they share a common node-vertex is exactly the line graph of G. This process is a stronger version of the one proved in [3, Theorem 6.5.4] because it further allows us to choose our endpoints in the line graph.

Formally, we first create the sets P_i for $i \in \{3, \ldots, k\}$. Then, because P is Hamiltonian, it must contain all node-vertices. This means that every edge-vertex is of the form $w_{v_i v_j}$ where $v_i, v_j \in P$. Thus, for each edge-vertex $w_{v_i v_j}$ except for y and $w_{v_1 v_2}$, we add it to the set $P_{\max(i,j)}$. Now, all the P_i are disjoint and their union contains all edge-vertices of H except y and $w_{v_1 v_2}$. Furthermore, notice that for all $i \in \{3, \ldots, k\}$, $w_{v_{i-1} v_i} \in P_i$.

We can now add those sets of edge-vertices to P' in this order: For each i from k down to 3, add the vertices of P_i to P' by adding $w_{v_{i-1}v_i}$ last. Then, if $y \neq w_{v_1v_2}$, add $w_{v_1v_2}$ and then add y to P'. Else, add y to P'.

The reader can then easily check that P' is a $\{2\}$-Hamiltonian path in H from x to y.

We now show the backward direction of (*). Suppose that H contains a $\{2\}$-Hamiltonian path Q starting at x and ending at y. As we have seen before, the neighbors of d_2 and d_3 in Q can only be d_2 (for d_3), d_3 (for d_2) and d_1 and β' (for both). Thus, d_2 and d_3 must be consecutive in Q, otherwise the first one will be preceded and followed by d_1 and β' in some order, and we will not be able to have neighbors for the second one. We also know that d_2 and d_3 are preceded by one of d_1 and β', and followed by the other one. Let Q_D be the contiguous subsequence of D containing the four vertices $\{d_1, d_2, d_3, \beta'\}$. From our study of walks of length 2 in H, we know that to go from a node-vertex to an edge-vertex, one must go through Q_D. Now, because x is a node-vertex, there can only be node-vertices until Q_D, and because we cannot reach d_1 directly from a node-vertex, we know that Q_D in fact must start with β', so it must end with d_1. Similarly, because y is an edge-vertex, there can only be edge-vertices between d_1 and y. Thus, letting Q' be the prefix of Q until β' (i.e., until the first vertex of Q_D), we know that Q' must visit only node-vertices, and it must visit all of them because afterwards Q can only visit edge-vertices. Thus, Q' starts with $x = \alpha'$, it ends with $y = \beta'$, it contains precisely all node-vertices, and adjacent pairs in Q' are connected by a walk of length 2 in H hence (from our study of length-2 walks in H) the corresponding vertices in G are adjacent. Thus, taking the vertices of G in the order of the corresponding vertices in Q', we have obtained the desired $\{1\}$-Hamiltonian path from α to β in G. This completes the reduction and concludes the proof. ◀

We now provide the full proof of the following result.

▶ **Theorem 3.5.** *The $\{2\}$-Hamiltonian cycle problem is NP-complete.*

Proof. We reduce from the $\{2\}$-Hamiltonian path problem with specified endpoints, which we just proved to be NP-hard in Theorem 4.1. Let D be the specific gadget given in Figure 2.

Given an input graph G and endpoint vertices $\alpha \neq \beta$ of G, Let H be the graph obtained by copying G and connecting it to D by adding the two edges $\{\alpha, d_0\}$ and $\{\beta, d_1\}$. The construction of H from G can obviously be done in polynomial time, and it suffices to show that the reduction is correct by proving the following claim (*): G has a $\{2\}$-Hamiltonian path from α to β if and only if H has a $\{2\}$-Hamiltonian cycle.

We first prove the forward direction of (*). Suppose G has a $\{2\}$-Hamiltonian path P starting at α and ending at β. Then one can check that $P' := P(d_6, d_5, d_0, d_3, d_1, d_7, d_4, d_2)$ is a $\{2\}$-Hamiltonian cycle of H.

We now prove the backward direction of (*). Suppose H has a $\{2\}$-Hamiltonian cycle C. We will examine the vertices of C, and talk of *neighbors* to mean pairs of vertices that occur consecutively in C (including possibly as the first and last vertices)– note that these vertices are then connected by a walk of length 2 but they are not necessarily neighbors in H. We observe that the only vertices connected to d_3 by a walk of length 2 are d_0 and d_1. Thus, these must be its two neighbors in C. For similar reasons, d_5 must have d_0 and d_6 as its neighbors in C and d_7 must have d_1 and d_4 as its neighbors in C. Because every vertex must have exactly 2 neighbors in C, we deduce that d_0 must be next to both d_3 and d_5 while d_1 must be next to both d_3 and d_7. This means that, up to reversing and/or conjugating the cycle, we know that the sequence of vertices $S = (d_6, d_5, d_0, d_3, d_1, d_7, d_4)$ must occur consecutively in C. Furthermore, d_2 must be right before or right after this subsequence. If it were not, the previous and next vertices of the sequence would be α and β and we would either go on d_2 and be stuck on it, or go to different vertices and have no way to ever come back to d_2. Assume that d_2 is right before S, the other case is analogous. Then the only vertex connected to d_4 by a walk of length 2 that is not in the sequence is α, and we deduce that α is the vertex right after S. The same holds for d_2 and β, thus β is the vertex before d_2. By removing every vertex of D we are then left with a {2}-Hamiltonian path of H that starts at α and ends at β, which is also a {2}-Hamiltonian path of G that starts at α and ends at β. This is what we wanted to obtain, so the backward direction of (*) is established, which concludes the proof. ◀

B.2 Proofs for $S = \{2, 4\}$

▶ **Proposition 4.2.** *Given an input graph G and two vertices $x \neq y$ at distance at most 2 in G, we can construct in polynomial time a graph $H_{x,y}$ such that:*

- *if G has a {1,2}-Hamiltonian cycle where x and y are consecutive, then $H_{x,y}$ has a {2,4}-Hamiltonian cycle,*
- *if $H_{x,y}$ has a {2,4}-Hamiltonian cycle, then G has a {1,2}-Hamiltonian cycle.*

Proof. Let G be a graph, and $x \neq y$ two vertices of G at distance at most 2. Let z be a neighbor of x towards y (defined formally below). Let us now build the graph $H_{x,y}$ by taking the incidence graph of G and connect each of x and z respectively to the vertices d_1 and d_2 of the gadget D shown in Figure 3.

Formally, we construct $H_{x,y}$ as follows:

- Build the incidence graph $I(G)$ by doing the following:
  - For each vertex $v \in V(G)$, add a corresponding vertex $v' \in I(G)$. We call these *node-vertices*.
  - For each edge $\{u, v\} \in E(G)$, add a new vertex w_{uv} to $I(G)$ and connect it to both u' and v'. We call these *edge-vertices*.
- Let $H_{x,y}$ be $I(G)$ together with a copy of the gadget D from Figure 3. Define z this way: if x and y are adjacent then $z = y$, else let z be any common neighbor of x and y. Now make $H_{x,y}$ connected by adding two new edges $\{x', d_1\}$ and $\{z', d_2\}$.

We now prove that this construction satisfies the two required properties, starting with the first item of the proposition statement. Let C be a {1,2}-Hamiltonian cycle of G where x and y are consecutive, and let us show that $H_{x,y}$ has a {2,4}-Hamiltonian cycle. Up to conjugating and reversing C, we can suppose without loss of generality that C starts at x and that y immediately follows, so let us write $C = (x, y, v_1, v_2, \ldots, v_k)$. Let us build the {2,4}-Hamiltonian cycle of $H_{x,y}$ by initializing C' with x' and adding enough vertices of D to be able to reach an edge-vertex, namely: $C' = (x', d_5, d_4)$. Here, the intermediate walks are $x' - d_1 - d_3 - d_4 - d_5$ and $d_5 - d_7 - d_6 - d_5 - d_4$, both of which have length 4.

Now, consider the graph G' that is a copy of G where we add two vertices d'_3 and d'_5 and the three edges $\{x, d'_3\}$, $\{z, d'_3\}$, and $\{d'_3, d'_5\}$. Observe that the edges of G' precisely correspond to the edge-vertices of $H_{x,y}$ plus the three vertices d_1, d_2, and d_4. By construction G' has at least 3 edges, and further the edges $\{d'_3, d'_5\}$ and $\{z, d'_3\}$ are connected by a walk of length 2 in the line graph $L(G')$ of G', as witnessed by the walk $\{d'_3, d'_5\} - \{x, d'_3\} - \{z, d'_3\}$. By Lemma 4.3, there is a {1,2}-Hamiltonian cycle C'' in $L(G')$ where $\{d'_3, d'_5\}$ and $\{z, d'_3\}$ are consecutive. We can transform C'' into a simple $\{2, 4\}$-path P'' of $H_{x,y}$ that starts at d_4, visits all the edge-vertices of $H_{x,y}$ and d_1, and ends on d_2, by replacing each vertex of $L(G')$ by the corresponding edge-vertex in $H_{x,y}$ and replacing $\{d'_3, d'_5\}$, $\{x, d'_3\}$, and $\{z, d'_3\}$, by d_4, d_1, and d_2, respectively. We append P'' to C' (without the first vertex d_4 as it is already in C'). Then from there, we go to d_6 through $d_2 - d_3 - d_4 - d_5 - d_6$, then we go to d_7 through $d_6 - d_5 - d_7$, and finally we go to d_3 through $d_7 - d_6 - d_5 - d_4 - d_3$.

We can now complete the construction of C' by concatenating our current C' with $C_{end} = (y', v'_1, v'_2, \ldots, v'_k)$, where we go to y' through $d_3 - d_2 - y'$ if we took $z = y$ (a walk of length 2), or otherwise through $d_3 - d_2 - z' - w_{zy} - y'$ (a walk of length 4). Then, all consecutive vertices in C_{end} are connected by a walk of length 2 or 4 because they were connected by a walk of length 1 or 2 in G by definition of C being a {1,2}-Hamiltonian cycle of G. Last, there is a walk of length 2 or 4 from v'_k to the first vertex x' of C' for the same reason. Thus, C' is indeed a {2,4}-Hamiltonian cycle of $H_{x,y}$.

We now show the second item of the proposition statement: let C' be a {2,4}-Hamiltonian cycle of $H_{x,y}$, and let us prove that G has a {1,2}-Hamiltonian cycle. Up to conjugating C', we can assume without loss of generality that C' starts at a node-vertex v. Because C' must at some point visit an edge-vertex and come back to v, it means that C' must switch between node-vertices and edge-vertices at least twice (note that if C' starts with v and ends with an edge-vertex then this is also counted as a switch, namely, a switch between the last vertex and the first vertex v as we close the cycle). Now, if we remove the vertices of D from $H_{x,y}$, we are left with a graph that is bipartite between node-vertices and edge-vertices. Thus, as long as no vertex of D is visited (in C' or as an intermediate vertex of a walk in C'), and because we are only allowed walks of even length, we know that C' must visit either only node-vertices or only edge-vertices between visits to vertices of D (as vertices of C' or as intermediate vertices). More precisely, because the only way to switch between node-vertices and edge-vertices is to traverse the triangle $\{d_5, d_6, d_7\}$ (without which $H_{x,y}$ would be bipartite), and because the maximal walk length 4 is less than the distance between the node-vertices and the vertices of this triangle, we even know that C' can only change between node-vertices and edge-vertices by a visit to at least one vertex of D which occurs in C' (i.e., not just as an intermediate vertex). We will now show that we can isolate in C' a contiguous sequence N which visits all node-vertices.

To justify this claim, it is easy to see by construction that D only has two vertices which may occur in C' consecutively to a node-vertex, namely, the vertices d_3 and d_5, which are the only vertices of D connected by a walk of length 2 or 4 to a node-vertex. Thus, consider the first time that C' visits a vertex of D. As C' was visiting node-vertices until then, the first vertex of D that occurs in C' is d_3 or d_5. We argue that the other vertex of d_3, d_5 cannot occur before C' exits D. Indeed, assuming by contradiction that it does, there are three possibilities. The first case is that the first vertex of C' after vertices of D is a node-vertex, and then we can never re-enter D and so we can never visit the edge-vertices, a contradiction. The second case is that the first vertex of C' after vertices of D is an edge-vertex, but then we need to re-enter D to switch back to node-vertices and we will have no way to exit D to a node-vertex because both d_3 and d_5 are taken. The third case is that C' ends after this visit to D, but this is impossible because the edge-vertices were not visited. Hence, in all three cases we know that C' exits D without visiting the other vertex from d_3, d_5. In particular, C' exits D to an edge-vertex. We then know that the other vertex from d_3, d_5 must be used to exit D when C' will switch back to the node-vertices. This reasoning implies that C' cannot switch between node-vertices and edge-vertices more than two times. In conclusion, up to conjugating C', we can write it as $C' = ND_1ED_2$ where N consists precisely of all the node-vertices, where D_1 and D_2 are only composed of vertices from D, and where E contains of all the edge-vertices (note that we have not ruled out that vertices from D may appear in E, but this does not matter).

We now focus our attention on N and show that it yields a {1,2}-Hamiltonian cycle in G.

To see why, observe that consecutive node-vertices in C must correspond to walks of length 2 or 4 in $H_{x,y}$. These correspond to walks of length 1 or 2 in G, or to walks that go via D but the only such walk between node-vertices in $H_{x,y}$ is the walk of length 4 connecting x' and z' for which the corresponding vertices x and z were adjacent in G.

Thus, if we take the original vertices of G corresponding to the node-vertices in the order they appear in N, we get a {1,2}-Hamiltonian path in G, from the original vertex a of the first vertex right after D_2 in C' (i.e., the first vertex of N), denoted a', to the original vertex b of the last vertex right before D_1 in C', denoted b'. To complete the proof, we need to show that this path is a cycle, i.e., that a and b are at connected by a walk of length 1 or 2 in G. Let us analyze which possible node-vertices a' and b' can be. There are two types of node-vertices that are connected to D by walks of length 2 or 4:

- x' and z' can reach either d_3 or d_5.
- Any node-vertex different from x' or z' whose original in G is adjacent to x or z can reach D only on d_3.

Recall now from our previous reasoning that D_1 must start with one of d_3, d_5 and D_2 must end with the other vertex among d_3, d_5. Thus, one of a' and b' occurs consecutively to d_5, so one of a' and b' must be x' or z'. The other vertex must then be either the other one of x' and z' or a node-vertex whose original in G is adjacent to x or z. Recall now that z was picked to be adjacent to x in G, so in all the possibilities considered we know that a and b in G must be at a distance of at most 2.

Thus G contains a {1,2}-Hamiltonian cycle, which concludes the proof. ◀

We now give the full proof of the following lemma, which adapts the proof of [20, Theorem 2].

▶ **Lemma 4.3** (Strengthening of Theorem 2 of [20]). *Let G be a graph with at least 3 edges. Let $L(G)$ be the line graph of G. Let α and β be two vertices of $L(G)$ connected by a walk of length 2 in $L(G)$. Then there exists a {1,2}-Hamiltonian cycle in $L(G)$ where α and β are consecutive.*

Proof. Let G be a graph with at least 3 edges and $L(G)$ be its line graph. Let α and β be two vertices of $L(G)$ connected by a walk of length 2 in $L(G)$. We first show that it is sufficient to prove the claim for trees: Consider a spanning tree T_1 of G. Color the edges of T_1 in blue in G. Subdivide each uncolored edge of G by adding a new vertex in the middle of the edge, and color one of the two new edges in red and the other in blue. Let T_2 be the graph composed of all the blue edges. Obviously, T_2 is a tree with at least 3 edges. It is easy to see that $L(T_2)$ is isomorphic to a spanning subgraph of $L(G)$. This means that we just need to prove that for any tree T with at least three edges, and for any vertex a and b of $L(T)$ such that a and b are at walk distance 2 in $L(T)$, there exists a {1,2}-Hamiltonian cycle in $L(T)$ where a and b are consecutive. Indeed, this would imply that $L(G)$ also has such a {1,2}-Hamiltonian cycle simply by taking the {1,2}-Hamiltonian cycle of $L(T_2)$ and following the same sequence of vertices. Thus, we can now focus on proving the claim for trees.

Let T be a tree with at least 3 edges. Let $L(T)$ be the line graph of T. Let α and β be two vertices of $L(T)$ such that there exists a walk of length 2 between them in $L(T)$. We will show that there exists a {1,2}-Hamiltonian cycle in $L(T)$ where α and β are consecutive. We prove the claim by induction on the number of edges of T. The base case is when T has 3 edges, and then there are two possibilities. First, if T is a path with 4 vertices, then $L(T)$ is a path with 3 vertices and the claim is immediate. Second, if T is a star with 3 branches, then $L(T)$ is the complete graph on 3 vertices, and the claim is also immediate. For the induction case, fix $n > 3$, suppose the claim holds for any tree with strictly fewer than n

edges, and let T be a tree with n edges. We will show that the claim is true on T. Let $L(T)$ be the line graph of T and let α and β be vertices of $L(T)$ such that there exists a walk of length 2 between them in $L(T)$.

There are two cases: either T is a path, or T is not a path. First, if T is a path, then $L(T)$ is the path with $n-1$ vertices $(v_1, \ldots, v_{n-1})$, and we can take the {1,2}-Hamiltonian cycle $C = (v_1, v_3, \ldots, v_{n-1}, v_{n-2}, \ldots, v_2)$ if $n-1$ is odd, and the same with v_{n-1} and v_{n-2} swapped if n is even. This choice of C ensures in fact that every pair of vertices at distance 2 in $L(T)$ are consecutive in C, hence this is true in particular of α and β.

Second, if T is not a path, then T has at least 3 leaves. We connect these leaves to a vertex of degree > 2 by iteratively taking the unique neighbor of the current vertex until we reach such a vertex. (The vertices of degree > 2 that we reach in this fashion may or may not be distinct.) The three paths constructed in this way do not share any edges, so let us pick one that contains neither the edge corresponding to α nor the edge corresponding to β. Formally, letting v_k be the chosen leaf and v_0 the vertex of degree at least 3 that we reach, we obtain by this argument a simple path $v_0, \ldots, v_k$ in T where v_0 has degree > 2, all vertices v_i with $0 < i < k$ have degree 2, v_k has degree 1, and the edges of that path are all different from α and β.

Let us consider the tree T_0 obtained from T by deleting the vertices $v_1, \ldots, v_k$ and all edges incident to them, which we then root at v_0. By $u_1, \ldots, u_i$ we denote the children of v_0 in T_0; as v_0 had degree > 2 in T, we know that $i \geq 2$. Now, because T_0 has at most $n-1$ edges, and $L(T_0)$ contains α and β, we can use the induction hypothesis on T_0 to obtain a {1,2}-Hamiltonian cycle C in $L(T_0)$ where α and β are consecutive. We will now modify C to be a {1,2}-Hamiltonian cycle of $L(T)$ where α and β are still consecutive, by inserting back the edges of the path $v_0, \ldots, v_k$ which we had removed.

Let x and y be two vertices of $L(T_0)$ such that x and y are consecutive in C (also possibly as first and last vertices), such that y is incident to v_0, i.e. $y = \{v_0, u_y\}$ for some child u_y of v_0 in T_0, and such that the edge of T_0 corresponding to x is incident to one of the u_i in T_0. Such a pair (x, y) must exist because we will at some point need to go from the subtree of T_0 rooted at u_i to the subtree rooted at u_j for some $i \neq j$. (Note that here we use the fact that v_0 has degree ≥ 2 in T_0.) There are three ways to achieve this with two consecutive vertices of $L(T_0)$ at distance ≤ 2: either by going from $\{v_0, u_i\}$ to $\{v_0, u_j\}$, or by going from $\{u_i, w_i\}$ to $\{v_0, u_j\}$ (where w_i is a neighbor of u_i which is different from v_0) or by going from $\{v_0, u_i\}$ to $\{u_j, w_j\}$ (where w_j is a neighbor of u_j which is different from v_0). In any case, we can pick x and y accordingly.

Now, we want to insert the missing vertices between x and y. However, if $\{x, y\} = \{\alpha, \beta\}$, this will not work, as we want to keep them consecutive. To this end, we will show how to modify x and y to ensure that at most one of them is in $\{\alpha, \beta\}$, such that they still respect the properties listed above. If $\{x, y\} = \{\alpha, \beta\}$, the other vertex consecutive to y in C is either corresponding to an edge incident to a u_l (which is forced if u_y is a leaf in T_0), or it is deeper in subtree rooted at u_y. In the first case, this other neighbor of y can be called x in the remainder of the proof. In the other case, the only way to exit the subtree rooted at u_y is to go from an edge incident to u_y that is not y towards an edge incident to v_0, and they can become the new x and y respectively.

We now have x and y such that at most one of them is in $\{\alpha, \beta\}$, such that y is incident to v_0, and such that x is incident to a u_i. By P we denote the $\{1,2\}$-path in $L(T)$ such that if $k = 1$, then $P = (x, \{v_0, v_1\}, y)$, and if $k \geq 2$, then $P = (x, \{v_0, v_1\}, \{v_2, v_3\}, \ldots, \{v_{g-3}, v_{g-2}\}, \{v_{g-1}, v_g\}, \{v_h, v_{h-1}\}, \ldots, \{v_2, v_1\}, y)$, where g is the greatest odd integer not exceeding k and h is the greatest even integer not exceeding k. If in C we replace $\{x, y\}$ by

P, we obtain a {1,2}-Hamiltonian cycle in $L(T)$ and because we insert nothing between α and β, they are still consecutive in this resulting cycle. Hence, we have shown that there is a {1,2}-Hamiltonian cycle in $L(T)$ where α and β are consecutive, which concludes the proof of the induction hypothesis. Thus, we have shown the desired claim by induction, which concludes the proof. ◀

C Proofs for Section 5 (Trivial Case: $S = \{1, 2, 4\}$)

▶ **Lemma 5.1.** *Let T be a tree rooted at r such that $|V(T)| \geq 2$. Then there exists a {1,2,4}-Hamiltonian cycle $(r, v_1, \ldots, v_k)$ of T such that v_1 is a child of r. Furthermore, such a cycle can be computed in linear time.*

Proof. We proceed by induction on $|V(T)|$.

Base case. The base case is when $|V(T)| = 2$, in which case T is the tree with one edge and the claim is trivial.

Induction step. Let $n \in \mathbb{N}, n \geq 3$, and assume the statement holds for all trees with at most $n-1$ vertices. Let T be a tree rooted at r with n vertices.

Partition the children of r into leaf children $u_1, \ldots, u_\ell$ and non-leaf children $v_1, \ldots, v_m$, where $\ell, m \geq 0$ and $\ell + m \geq 1$ because $n \geq 3$. (In fact we even know that $\ell + m \geq 2$, but we do not need it.) For each i, denote by T_i the subtree of T rooted at v_i, i.e., the connected component of $T \setminus \{r\}$ that contains v_i. Note that, as v_i is not a leaf, we must have $|V(T_i)| \geq 2$ for all i.

Because every T_i has less than n vertices, we know by induction hypothesis that each T_i admits a {1,2,4}-Hamiltonian cycle C_i starting at v_i and starting with a 1-hop to a child of v_i in T_i, denoted w_i.

We will now construct the {1,2,4}-Hamiltonian cycle of T in two stages. First, we will visit all leaf children of r (if some exist), and second we will visit all vertices in each subtree T_i one after the other (if some exist).

Stage 1: The leaf children. If $\ell > 0$, let us traverse the leaves in order: start from r and move to u_1 (walk of length 1). Then for $j = 1, \ldots, \ell - 1$ move from u_j to u_{j+1} via $u_j - r - u_{j+1}$ (walk of length 2). If $m = 0$, we close the cycle with $u_\ell - r$ (walk of length 1) and we obtain a {1,2,4}-Hamiltonian cycle of T which starts at a child of r, so we are finished. Otherwise, if $m > 0$, continue to stage 2.

Stage 2: All the T_i. First, go to v_1. If $\ell > 0$, this is done via $u_\ell - r - v_1$ (walk of length 2); if $\ell = 0$, this is done via $r - v_1$ (walk of length 1). Then traverse the {1,2,4}-Hamiltonian cycle C_1 of T_1 after conjugation and reversal, going from v_1 to w_1 while visiting all vertices of T_1. Next, traverse the other non-leaf children trees $T_2, \ldots, T_m$ as follows:

1. For $i = 2, \ldots, m$, we alternate those two operations:
   - **a.** If i is even, then we go to w_i from w_{i-1} through $w_{i-1} - v_{i-1} - r - v_i - w_i$, which is a walk of length 4. We traverse the cycle C_i of T_i in order (after conjugation) to go from w_i to v_i while visiting all vertices of T_i.
   - **b.** If i is odd, then we go to v_i from v_{i-1} through $v_{i-1} - r - v_i$, which is a walk of length 2. We traverse the cycle C_i of T_i (after conjugation) in reverse order, starting from v_i and ending at w_i, while visiting all vertices of T_i.

   This alternation continues until all subtrees have been visited.
2. Finally, after processing T_m, we return to r from either w_m or v_m depending on the parity of m. If m is odd, this is a walk of length 2 through $w_m - v_m - r$; if m is even, this is a walk of length 1.

The constructed sequence therefore is a {1,2,4}-Hamiltonian cycle of T starting at r where the second vertex is a child of r, which establishes the inductive claim.

This construction can be implemented by a recursive algorithm which considers each vertex exactly once, thus runs in linear time, which concludes the proof. ◀

We also provide the proof of the Theorem 3.7, using the previous lemma.

Proof of Theorem 3.7. Let G be a connected graph. If $|V(G)| = 1$, then G is a single vertex, say v, and (v) is a {1,2,4}-Hamiltonian cycle of G. Otherwise, choose a spanning tree T of G. Using Lemma 5.1, we know that T admits a {1,2,4}-Hamiltonian cycle $(r, v_1, \dots, v_k)$ where r is the root of T and v_1 is a child of r in T. Because T is a subgraph of G, $(r, v_1, \dots, v_k)$ is also a {1,2,4}-Hamiltonian cycle of G. ◀

D Proofs for Section 6 (Polynomial Case: $S = \{2, 4, 6\}$)

▶ **Lemma 6.1.** *Given a non-bipartite graph G, we can construct a {2,4,6}-Hamiltonian cycle of G in linear time.*

Proof. Let $G = (V, E)$ be a non-bipartite graph. Then G must contain a cycle of odd length. Let $C = (v_0, v_1, \dots, v_{2k})$ be a shortest odd cycle in G, where $k \geq 1$.

Step 1: Constructing a spanning forest rooted at the cycle. Consider the graph $G' := G \setminus C$ obtained from G by removing all edges between vertices of C. We construct a spanning forest $\{T_0, T_1, \dots, T_{2k}\}$ rooted at the cycle vertices as follows:

1. Initialize each tree $T_i := \{v_i\}$.
2. Perform a simultaneous breadth-first search starting from all v_i.
3. Whenever a vertex $u \in V \setminus C$ is reached for the first time from some root v_i, assign u to T_i.
4. Continue until every vertex in V is assigned to exactly one T_i.

By construction, each T_i is connected, the trees are disjoint, and their union covers all vertices of G, forming a spanning forest rooted at the cycle vertices.

Step 2: Constructing Hamiltonian paths within each tree. For $0 \leq i \leq 2k$ and u a vertex of T_i, let $d_{T_i}(u)$ be the depth of u in the tree T_i. Let $V_i^{\text{even}} := \{u \in T_i \mid d_{T_i}(u) \bmod 2 = 0\}$ be the set of vertices at even depth in T_i, and let $V_i^{\text{odd}} := \{u \in T_i \mid d_{T_i}(u) \bmod 2 = 1\}$ be the set of vertices at odd depth in T_i. We claim that we can construct a path in every T_i that starts at any vertex of V_i^{even} (resp. V_i^{odd}), ends on any other vertex of V_i^{even} (resp. V_i^{odd}) and visits all vertices of V_i^{even} (resp. V_i^{odd}) exactly once using only walks of length 2, 4, or 6. To achieve this, we first have to build the tree T_i^{even} (resp. T_i^{odd}) by taking the vertices of V_i^{even} (resp. V_i^{odd}) and connecting them if they are at distance 2 in T_i. Notice that T_i^{even} and T_i^{odd} are indeed trees, where the parent of each vertex is its grandparent in T_i. Recall now that we saw in Proposition 3.4 that any tree admits a {1,2,3}-Hamiltonian path between any 2 given vertices. We can use this result to create a {1,2,3}-Hamiltonian path in T_i^{even} (resp. T_i^{odd}) between any two of its vertices. By taking the same vertices in the same order as the path we just constructed, we will end up with a simple $\{2, 4, 6\}$-path in T_i between any 2 given vertices because every distance is just multiplied by 2.

By using this property, we define for each i a simple $\{2, 4, 6\}$-path taking all the vertices of V_i^{even} except for v_i, denoted P_i^{even}, this way: If $V_i^{\text{even}} = \{v_i\}$, P_i^{even} is empty. Else, if V_i^{even} contains a vertex at depth 4 in T_i, P_i^{even} starts at any vertex in V_i^{even} at depth 4 in T_i and ends at any vertex at depth 2 in T_i. Finally, the last case is when V_i^{even} contains only

v_i and vertices at depth 2 in T_i, in which case we define P_i^{even} as a path between any two vertices at depth 2 in T_i. Similarly, we can define for each i a simple $\{2,4,6\}$-path P_i^{odd} that takes all the vertices of V_i^{odd} in this way: If V_i^{odd} is empty, P_i^{odd} is empty. Else, if V_i^{odd} contains a vertex at depth 3 in T_i, P_i^{odd} starts at any vertex in V_i^{odd} at depth 1 in T_i and ends at any vertex at depth 3 in T_i. Finally, the only other case is when V_i^{odd} only contains vertices at depth 1 in T_i, and then P_i^{odd} is a path between any two vertices at depth 1 in T_i.

Step 3: Creating the cycle. We now construct the {2,4,6}-Hamiltonian cycle by interleaving visits to the cycle C and the Hamiltonian paths that we constructed for V_i^{even} and V_i^{odd} within each tree. Informally, we will go around the cycle twice, the first time visiting every vertex v_i with even index i (skipping the odd indexes). During the second traversal we will skip the vertices visited in the first traversal and visit the vertices with odd index. When stopping on a v_i, we enumerate P_i^{even}, and when skipping a v_i, we enumerate P_i^{odd}. This allows us to visit every vertex exactly once, and end up on the starting vertex. The formal construction of the cycle is as follows:

1. Start at v_0.
2. For $i = 0$ to $k-1$:
   a. Visit the vertices from P_{2i}^{even}, starting at a vertex at depth 2 or 4 in T_{2i} (or none if P_{2i}^{even} is empty) and ending on a vertex at depth 2 in T_{2i} (or at v_{2i} itself if P_{2i}^{even} is empty).
   b. If P_{2i+1}^{odd} is not empty, go to its first vertex, which is at depth 1 in T_{2i+1}, through a walk of length 4 (or 2 if P_{2i}^{even} was empty). Then, traverse P_{2i+1}^{odd}, ending on a vertex at depth 1 or 3 in T_{2i+1}.
   c. Move to v_{2i+2}. If P_{2i+1}^{odd} was empty, this is a walk of length 2 through $v_{2i} - v_{2i+1} - v_{2i+2}$. Otherwise, this is a walk of length 2 or 4 from the last vertex of P_{2i+1}^{odd}.
3. We can now go through P_{2k}^{even}, ending on a vertex at depth 2 in T_{2k}, or stay at v_{2k} if P_{2k}^{even} is empty. Once this is done, we have visited all vertices of the V_i^{even} for every even i and all vertices of the V_i^{odd} for every odd i.
4. We then go to the first vertex of P_0^{odd} which is at depth 1 in T_0, through a walk of length 4 (or 2 if P_{2k}^{even} was empty) and traverse P_0^{odd} ending at a vertex at depth 1 or 3 in T_0.
5. We now just have to repeat the previous loop while swapping the roles of even and odd. Formally, for $i = 0$ to $k-1$:
   a. Move to v_{2i+1}. This is a walk of length 2 or 4 from either the last vertex of P_{2i}^{odd} or from v_{2i} if P_{2i}^{odd} was empty.
   b. If P_{2i+1}^{even} is not empty, visit the vertices of P_{2i+1}^{even} in order, ending on a vertex at depth 2 in T_{2i+1}. Else stay at v_{2i+1}.
   c. Go to the first vertex of P_{2i+2}^{odd} which is at depth 1 in T_{2i+2} through a walk of length 4 from the last vertex of P_{2i+1}^{even} (or of length 2 from v_{2i+1} if P_{2i+1}^{even} was empty) and traverse P_{2i+2}^{odd}, ending at a vertex at depth 1 or 3 in T_{2i+2}.
6. Finally, return to v_0 from the last vertex of P_{2k}^{odd} or from v_{2k}, which is a walk of length 2 or 4. This completes the cycle and concludes the proof. ◀

E Proofs for Section 7 (Variants)

E.1 Proofs for Infinite Sets

Let S be an infinite set of positive integers that is closed under subtraction of 2. Then, if $S = \mathbb{N}^+$, then S contains $\{1,2,3\}$ and the problem is trivial by Proposition 3.4. Otherwise, either S contains all the even numbers, or S contains all the odd numbers (or both). In the first case, then S may also contain some odd numbers, but in particular it must contain 1, so that S contains $\{1,2,4\}$, which means that every graph admits an S-Hamiltonian cycle by Theorem 3.7. In the second case, then S may also contain some even numbers, but in particular it must contain 2, so that S contains $\{1,2,3\}$, which means that every graph admits an S-Hamiltonian cycle by Proposition 3.4. Hence, the only non-trivial cases are when S is the set of all even numbers, and when S is the set of all odd numbers.

First, for S the set of all even numbers, the characterization essentially follows from Section 6. Indeed, as S contains $\{2,4,6\}$, we know by Lemma 6.1 that every non-bipartite graph has an S-Hamiltonian cycle and that we can compute it in linear time. Now, it is easy to see that bipartite graphs cannot have an S-Hamiltonian cycle for similar reasons to the proof of Theorem 3.8: walks of even length cannot make us switch from one part of the bipartition to the other. Hence, we have shown:

▶ **Theorem E.1.** *Let S denote the set of all even integers, and let G be a graph. Then G admits an S-Hamiltonian cycle if and only if G is non-bipartite.*

Then, for S the set of all odd numbers, the classification does not follow from our earlier results. We will now show the following:

▶ **Theorem 7.1.** *Let S denote the set of all odd integers, and let G be a graph. Then G admits an S-Hamiltonian cycle if and only one of the two following cases holds: G is non-bipartite, or G is bipartite with parts of equal cardinality.*

Proof. Let G be a graph. We will first show that G admits an S-Hamiltonian cycle if G is non-bipartite.

Suppose G is non-bipartite. Then let $C = (v_1, \ldots, v_n)$ be any ordering of $V(G)$. We claim C is an S-Hamiltonian cycle of G. Indeed, since G is non-bipartite, it contains an odd cycle O. For any two consecutive vertices v_i and v_{i+1} in C, we will show there exists an odd-length walk between them in G. Let P_1 be any walk from v_i to some vertex $u \in V(O)$, and let P_2 be any walk from u to v_{i+1}. If $|P_1| + |P_2|$ is odd, then P_1P_2 (the concatenation of P_1 with P_2) is a walk of odd length between v_i and v_{i+1}. If $|P_1| + |P_2|$ is even, then traverse O once between P_1 and P_2, and P_1OP_2 is a walk of length $|P_1| + |O| + |P_2|$, which is odd. Since any two consecutive vertices in C are connected by an odd-length walk in G, C is an S-Hamiltonian cycle of G.

Now it remains to show that if G is bipartite, then G has an S-Hamiltonian cycle if and only if its bipartition has equally-sized parts.

Suppose G is bipartite with parts A and B. First, we show that if G admits an S-Hamiltonian cycle, then $|A| = |B|$. Let $C = (v_1, \ldots, v_n)$ be an S-Hamiltonian cycle of G. In a bipartite graph, a walk between two vertices has even length (resp., odd length) if and only if it joins two vertices that are in the same parts (resp., in different parts). Since each consecutive pair (v_i, v_{i+1}) in C must be connected by an odd-length walk, they must belong to different parts. This means C must alternate between A and B, so we must have $|A| = |B| = n/2$.

Conversely, suppose G is bipartite with equally-sized parts A and B, where $|A| = |B| = n/2$. Consider any ordering $C = (a_1, b_1, a_2, b_2, \ldots, a_{n/2}, b_{n/2})$ that alternates between vertices of A and B. Since consecutive vertices in C belong to different parts, and because any two vertices in different parts are connected by a walk which has odd length, we know that there is indeed an odd-length walk between each consecutive pair. Therefore, C is an S-Hamiltonian cycle of G. ◀

E.2 Proofs for Hamiltonian Path Variants

In this section, we detail the status of the path versions of our main problem (with and without specified endpoints), for all sets S, starting with the known results:

- The {1}-Hamiltonian path problem, with or without specified endpoints, is well-known to be NP-hard through an easy reduction from the {1}-Hamiltonian cycle problem.
- Concerning the {1,2}-Hamiltonian path problems, it is not directly a known result but it is easy to see that the reduction proposed for the {1,2}-Hamiltonian cycle problem in [25] also allows us to reduce the {1}-Hamiltonian path problem (resp., the {1}-Hamiltonian path problem with specified endpoints) to the {1,2}-Hamiltonian path problem (resp., the {1,2}-Hamiltonian path problem with specified endpoints).
- As for the {1,2,3}-Hamiltonian path problems, as we reviewed in Proposition 3.4 the proof from [17, 24] states that every graph has a {1,2,3}-Hamiltonian path starting and ending at any pair of vertices, which means that both problems are trivial.

Now, we talk about what can be derived from our new results:

- Because we needed it in Section 4.1, we already proved that the {2}-Hamiltonian path problem with specified endpoints is NP-hard. As for the variant without specified endpoints, we will prove it to also be NP-hard in Theorem E.2.
- As every graph contains a {1,2,4}-Hamiltonian cycle, it must also be the case that it contains a {1,2,4}-Hamiltonian path, simply obtained by breaking the cycle. However, the situation for specified endpoints {1,2,4}-Hamiltonian path is quite different. Indeed, we show in Theorem 7.3 that this decision problem is non-trivial but that it is actually solvable in linear time.
- The characterization of the {2,4,6}-Hamiltonian path problem is exactly the same as the one for the {2,4,6}-Hamiltonian cycle problem, which is that either the given graph is bipartite and thus it cannot contain such a path, or it is not and then it must contain a {2,4,6}-Hamiltonian cycle which is also a {2,4,6}-Hamiltonian path of that graph. Concerning the specified endpoints version, we believe it is possible to modify the algorithm of the proof of Theorem 3.8 to be able to start and end anywhere (so that the problem would be solvable in linear time, and the path with specified endpoints would exist precisely on non-bipartite graphs, for any pair of endpoints in that case), but we leave this problem open.
- The last problems we are left with are the {2,4}-Hamiltonian path problems (with and without specified endpoints). For the {2,4}-Hamiltonian path problem with specified endpoints, the problem admits a Cook reduction from the {2,4}-Hamiltonian cycle problem, so it is NP-hard under Cook reductions. For the {2,4}-Hamiltonian path problem without specified endpoints, we leave its complexity open, but we conjecture that it is NP-hard like {2,4}-Hamiltonian cycle.
- Finally, the {1,3}-Hamiltonian path problems (with or without prescribed endpoints), and the problems with finite sets $S = \{1, \ldots, 2k + 1\}$ of odd integers for some $k > 0$, are left open by the present work, like the corresponding S-Hamiltonian cycle problem.

E.3 Proofs for the {2}-Hamiltonian Path Problem

We prove in this appendix that the {2}-Hamiltonian path problem is NP-complete, by a reduction from the {2}-Hamiltonian path problem with specified endpoints.

▶ **Theorem E.2.** *The {2}-Hamiltonian path problem is NP-complete.*

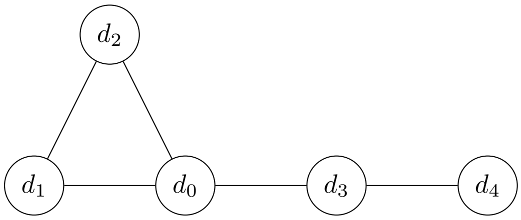

Figure 6 Dangling gadget D for the proof of Theorem E.2

Proof. We reduce from the $\{2\}$-Hamiltonian path problem with specified endpoints which we proved to be NP-complete in Theorem 4.1.

Let G be a graph and $\alpha, \beta \in V(G)$ two different vertices of G. Let us construct in polynomial time a graph H and show that G has a $\{2\}$-Hamiltonian path between α and β if and only if H contains a $\{2\}$-Hamiltonian path.

We construct H by connecting G to two copies of the gadget D represented in Figure 6 by adding an edge from α to the vertex d_0 in one copy, and an edge from β to the vertex d_0 in the other copy.

Formally, H is constructed by copying G and adding D and D' to G (with D' a copy of D such that $V(D') = \{d'_0, d'_1, d'_2, d'_3, d'_4\}$), which we will connect to G with the edges (α, d_0) and (β, d'_0).

We first show the forward direction of the reduction. Let us assume that G has a $\{2\}$-Hamiltonian path P starting at α and ending at β. Then one can easily verify that the path $P' = (d_4, d_0, d_2, d_3, d_1)P(d'_1, d'_3, d'_2, d'_0, d'_4)$ is a $\{2\}$-Hamiltonian path of H.

We now show the more interesting backward direction. Suppose H has a $\{2\}$-Hamiltonian path Q. Because d_4 and d'_4 only have one vertex at distance 2 (resp., d_0 and d'_0), they must be the endpoints of Q. Without loss of generality, up to reversing Q, suppose d_4 is the start and d'_4 is the end. From d_4 we must go to d_0 and now we have the possibility of going out of D. If we do, then the vertices d_1, d_2, d_3 of D are not yet visited, and the only vertex of G to which they are connected by a walk of length 2 is α. This means we can only enter D again by using α which prevents us from going out afterwards and implies that we cannot reach d'_4 which must be at the end of Q. For this reason, we need to visit all vertices of D before going to α. We will now eventually leave D and it will necessarily be via α and we will not re-enter D again afterwards. By a symmetric reasoning, we can only take vertices of D' after β and all vertices of D' must be after β.

In summary, we have shown that Q is composed of three $\{2\}$-paths:

- One from d_4 to α by taking only vertices of D;
- one from α to β by taking only vertices of G;
- and one from β to d'_4 by taking only vertices of D'.

Then by removing the vertices before α and those after β we get a $\{2\}$-Hamiltonian path from α to β in G.

This completes the proof that the $\{2\}$-Hamiltonian path problem is NP-complete. ◀

E.4 Proofs for the {1,2,4}-Hamiltonian Path Problem with Specified Endpoints

▶ **Lemma 7.2.** *Let T be a tree with $|V(T)| > 2$ and $a, b \in V(T)$ two vertices of T such that a and b are not adjacent. Then there exists a {1,2,4}-Hamiltonian path between a and b.*

Proof. We will prove this result by induction on the number of vertices of T.

Base case. There is only one tree of size 3, which is the path of size 3. The only two vertices that are not adjacent are the vertices at both ends of the path, and we can achieve a {1,2,4}-Hamiltonian path between them by taking two steps of length 1.

Induction step. Assume that the statement holds for all trees with strictly fewer than $n \geq 4$ vertices. Let T be a tree with $|V(T)| = n$, and let $a, b \in V(T)$ be two non-adjacent vertices.

Because T is a tree, there is a unique simple {1}-path between a and b, which we write $P_{ab} = (a, v_1, \ldots, v_k, b)$. Because a is not adjacent to b, we have $k \geq 1$. Let T' be the forest obtained from T by removing every edge of P_{ab}, and let T_a be the connected component of T' containing a.

If $k \geq 2$, then we can define two paths P_a and P_b that will form a {1,2,4}-Hamiltonian path of T when concatenated. If a is a leaf, then we define P_a as (a). Else, T_a has at least two vertices and we use Lemma 5.1 to get a {1,2,4}-Hamiltonian cycle of T_a $(a, a_1, \ldots, a_m)$ where a_1 is adjacent to a. We then define P_a as $(a, a_m, \ldots, a_1)$. By removing every vertex of T_a from T, we get a tree $T_{v_1 b}$ in which v_1 and b are non-adjacent and $|V(T_{v_1 b})| < n$ because at least we removed a. This means we can use the induction hypothesis on $T_{v_1 b}$ to get a {1,2,4}-Hamiltonian path from v_1 to b, which we call P_b. Concatenating P_a and P_b yields a path from a to b which goes through every vertex of T. Every step in P_a or P_b is of length 1, 2 or 4 due to their construction. Furthermore, the transition between them is either from a to v_1 (of length 1) if a is a leaf, or else from a_1 to v_1 through a (of length 2). Thus, we have constructed a {1,2,4}-Hamiltonian path of T from a to b.

If $k = 1$, let c be the common neighbor of a and b. Let T_c be the connected component of T' containing c and T_b be the one containing b. There are two cases: either $T_c = \{c\}$, or T_c contains at least two vertices. In the first case, we can define P_a in the same way as above. Namely, if a is a leaf, $P_a = (a)$. Alternatively, we use Lemma 5.1 to get a {1,2,4}-Hamiltonian cycle of T_a $(a, a_1, \ldots, a_m)$ where a_1 is adjacent to a and we define $P_a = (a, a_m, \ldots, a_1)$. We then define P_b similarly. Namely, if b is a leaf, $P_b = (b)$. Else we use the same lemma to get a {1,2,4}-Hamiltonian cycle of T_b $(b, b_1, \ldots, b_l)$ and we can define P_b as $(b_1, \ldots, b_l, b)$. Then $P = P_a(c)P_b$ is a {1,2,4}-Hamiltonian path of T from a to b. In the first case, we have $|V(T_c)| \geq 2$ and we can use Lemma 5.1 to get a {1,2,4}-Hamiltonian cycle of T_c $(c, c_1, \ldots, c_j)$ where c_1 is adjacent to c. If both a and b are leaves, then $P = (a, c, c_j, \ldots, c_1, b)$ is a {1,2,4}-Hamiltonian path of T from a to b. Otherwise, if both T_a and T_b are stars centered at a and b respectively, such that $V(T_a) = \{a, a_1, \ldots, a_m\}$ and $V(T_b) = \{b, b_1, \ldots, b_l\}$, then $P = (a, a_1, \ldots, a_m, b_1, \ldots, b_l, c, c_j, \ldots, c_1, b)$ is a {1,2,4}-Hamiltonian path of T from a to b. Finally, if one of T_a or T_b is not a star centered at a or b, then one of them must have a vertex at distance 2 from a or b. Without loss of generality, suppose T_a has a vertex a_1 at distance 2 from a. Then we can use the induction hypothesis on T_a to get a {1,2,4}-Hamiltonian path P_a from a to a_1. We now define P_b exactly as above and $P = P_a(c_1, \ldots, c_j, c)P_b$ is a {1,2,4}-Hamiltonian path from a to b. ◀

▶ **Theorem 7.3.** *Let G be a graph and a and b be two different vertices of G. Then G admits a {1,2,4}-Hamiltonian path from a to b if and only if G is not a bistar centered at a and b.*

Proof. Let G be a graph and a and b be two different vertices of G. If a and b are not adjacent in G, then let T be a spanning tree of G. In T, a and b are also not adjacent. Thus, we can use Lemma 7.2 to get a {1,2,4}-Hamiltonian path of T from a to b, which is also a {1,2,4}-Hamiltonian path of G from a to b.

Suppose a and b are adjacent and that G is not a bistar centered at a and b, then there exists two vertices $v_1, v_2 \in G$ different from a or b such that v_2 is adjacent to v_1 and v_1 is adjacent to a or b. Without loss of generality, suppose v_1 is adjacent to a. Then let T be a spanning tree of G such that $\{a, v_1\}$ and $\{v_1, v_2\}$ are edges of T. Let T' be the tree obtained from T by removing the edge $\{a, b\}$, let T_a be the connected component of T' containing a and T_b the one containing b. Then we can use Lemma 7.2 to get a {1,2,4}-Hamiltonian path P_a from a to v_2 in T_a. We can also use Lemma 5.1 to get a {1,2,4}-Hamiltonian cycle $(b, b_1, \ldots, b_l)$ where b and b_1 are adjacent. Then $P = P_a(b_1, \ldots, b_l, b)$ is a {1,2,4}-Hamiltonian path of T and of G from a to b.

Finally, we must show that a bistar cannot contain a {1,2,4}-Hamiltonian path connecting its two centers. Suppose G is a bistar centered at a and b and let $(a_1, \ldots, a_m)$ with $m \geq 1$ be the neighbors of a that are different from b, and let $(b_1, \ldots, b_l)$ with $l \geq 1$ be the neighbors of b that are different from a. From a we can reach any vertex. If our path starts by going from a to some a_i for some i, then the only vertices it can reach from now on are either other a_j for $j \neq i$ or b. Going to another a_j does not change that fact which means there will be no way to reach a b_i without taking b. Because b must be taken last, we have excluded the case of a path that starts with some a_i. However, the only other solution is to start by taking a b_i, and this does not work for the exact same reasoning: we cannot visit the a_j without visiting a. Thus there is no {1,2,4}-Hamiltonian path from a to b in G. ◀

E.5 Proofs for Bounded Cliquewidth

▶ **Theorem 7.4.** *For any fixed finite set $S \subseteq \mathbb{N}^+$ and bound $k \in \mathbb{N}$, the S-Hamiltonian cycle problem on input graphs of cliquewidth at most k is in polynomial time.*

Proof. Let $S = \{d_1, \ldots, d_m\} \subseteq \mathbb{N}^+$ be the fixed finite set of authorized lengths. Let k be the constant cliquewidth bound, and let $\mathcal{C}$ be a class of finite graphs of cliquewidth bounded by k.

We define for each $d_i \in S$ the following FO formula: $\phi_{d_i}(x, y) = \exists z_2 \cdots \exists z_{d_i}(E(x, z_2) \wedge E(z_2, z_3) \wedge \cdots \wedge E(z_{d_i}, y))$ with $E(x, y)$ the edge relation of the graph G. We also define the following formula: $\phi(x, y) = \bigvee_{d_i \in S} \phi_{d_i}(x, y)$. For each $d_i \in S$, the formula $\phi_{d_i}(x, y)$ defines the constraint that x and y are connected in G by a walk of length d_i, and the formula $\phi(x, y)$ then defines the constraint that x and y are connected in G by a walk of length ℓ for some $\ell \in S$.

This allows us to define the following MSO interpretation: For $G \in \mathcal{C}$, define a new graph H with the same vertex set $V(H) = V(G)$, and with edge set $E(H)$ defined by $E(H) = \big\{\{u, v\} | (u, v) \in V(G)^2 \mid G \models \phi(u, v)\big\}$.

We now show that G admits an S-Hamiltonian cycle if and only if H admits a Hamiltonian cycle. First, if G admits an S-Hamiltonian cycle, then there exists a cycle C in G that visits every vertex exactly once and such that for every two consecutive vertices u and v in C, there exists a walk of length in S between u and v. By definition of H, this means that there is an edge between u and v in H. Hence, the cycle C is also a Hamiltonian cycle in H. Conversely, if H admits a Hamiltonian cycle, then there exists a cycle C in H that visits every vertex exactly once and such that for every two consecutive vertices u and v in C, there is an edge between u and v in H. By definition of H, this means that there exists a walk of length in S between u and v in G. Hence, the cycle C is also an S-Hamiltonian cycle in G.

Because $\mathcal{C}$ has bounded cliquewidth, and because MSO interpretations preserve bounded cliquewidth [6, 5], letting $\mathcal{C}'$ be the class of the graphs H that can be obtained from $G \in \mathcal{C}$ by the above MSO interpretation, then $\mathcal{C}'$ also has bounded cliquewidth: let k' be a bound on the cliquewidth of the graphs in $\mathcal{C}'$. Finally, because the Hamiltonian cycle problem is in XP for graphs of bounded cliquewidth [4], there exists an algorithm that solves the S-Hamiltonian cycle problem on $\mathcal{C}$ in XP time, defined in the following way. Given a graph $G \in \mathcal{C}$, we construct the graph H as defined above, which takes polynomial time since S is fixed, and H must have cliquewidth bounded by k'. Using the result from [21, Theorem 1.1], we can get a decomposition of H of width at most $f(k')$ for some function f. Then, using the result from [4, Theorem 1], we can solve the Hamiltonian cycle problem on H in time $O(n^{g(f(k'))})$ for some function g.

This gives us an algorithm to solve the S-Hamiltonian cycle problem on $\mathcal{C}$ in time $O(n^{g(f(k'))})$, which is polynomial since k' is a constant depending only on S and k. ◀