

ARRIVAL: Next Stop in CLS*

Bernd Gärtner¹, Thomas Dueholm Hansen², Pavel Hubáček³, Karel Král³, Hagar Mosaad⁴,
and Veronika Slívová³

¹Department of Computer Science, ETH Zürich, Switzerland
gaertner@inf.ethz.ch

²Department of Computer Science, University of Copenhagen, Denmark
dueholm@di.ku.dk

³Computer Science Institute, Charles University, Prague, Czech Republic
{hubacek, kralka, slivova}@iuuk.mff.cuni.cz

⁴Department of Computer Science and Engineering, German University in Cairo, Egypt
hagar.omar@student.guc.edu.eg

February 22, 2018

Abstract

We study the computational complexity of **ARRIVAL**, a zero-player game on n -vertex switch graphs introduced by Dohrau, Gärtner, Kohler, Matoušek, and Welzl. They showed that the problem of deciding termination of this game is contained in $\text{NP} \cap \text{coNP}$. Karthik C. S. recently introduced a search variant of **ARRIVAL** and showed that it is in the complexity class **PLS**. In this work, we significantly improve the known upper bounds for both the decision and the search variants of **ARRIVAL**.

First, we resolve a question suggested by Dohrau et al. and show that the decision variant of **ARRIVAL** is in $\text{UP} \cap \text{coUP}$. Second, we prove that the search variant of **ARRIVAL** is contained in **CLS**. Third, we give a randomized $O(1.4143^n)$ -time algorithm to solve both variants.

Our main technical contributions are (a) an efficiently verifiable characterization of the unique witness for termination of the **ARRIVAL** game, and (b) an efficient way of sampling from the state space of the game. We show that the problem of finding the unique witness is contained in **CLS**, whereas it was previously conjectured to be **FPSPACE**-complete. The efficient sampling procedure yields the first algorithm for the problem that has expected runtime $O(c^n)$ with $c < 2$.

1 Introduction

Variants of switch graphs have applications and are studied for example in combinatorics and in automata theory (cf. [10] and the references therein). Dohrau et al. [5] introduced **ARRIVAL**, a natural computational problem on switch graphs, which they informally described as follows:

Suppose that a train is running along a railway network, starting from a designated origin, with the goal of reaching a designated destination. The network, however, is of a special nature: every time the train traverses a switch, the switch will change its position immediately afterwards. Hence, the next time the train traverses the same switch, the other direction will be taken, so that directions alternate with each traversal of the switch.

Given a network with origin and destination, what is the complexity of deciding whether the train, starting at the origin, will eventually reach the destination? [5]

*The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement n. 616787 and was supported by the project 17-09142S of GA ČR. P. Hubáček and V. Slívová were partially supported by PRIMUS grant PRIMUS/17/SCI/9. T.D. Hansen was partially supported by BARC which is funded by a VILLUM Investigator grant (16582) to Mikkel Thorup.

The above rather straightforward question remains unresolved. Dohrau et al. [5] showed that deciding ARRIVAL is unlikely to be NP-complete (by demonstrating that it is in $\text{NP} \cap \text{coNP}$), but it is currently not known to be efficiently solvable.

To determine whether the train eventually reaches its destination, it is natural to consider a *run profile*, i.e., the complete transcript describing how many times the train traversed each edge. Dohrau et al. [5] presented a natural integer programming interpretation of run profiles called *switching flows*, which have the advantage of being trivial to verify. The downside of switching flows is that they do not guarantee to faithfully represent the number of times each edge has been traversed; a switching flow might contain superfluous circulations compared to a valid run profile. Nevertheless, Dohrau et al. [5] proved that the existence of a switching flow implies that the train reaches its destination, and thus a switching flow constitutes an NP witness for ARRIVAL.

The coNP membership was shown by an insightful observation about the structure of switch graphs. Specifically, the train reaches its destination if and only if it never enters a node from which there is no directed path to d . The railway network can thus be altered so that all such vertices point to an additional “dead-end” vertex \bar{d} . The coNP witness is then simply a switching flow from the origin to the dead-end \bar{d} .

Given that the decision variant of ARRIVAL is in $\text{NP} \cap \text{coNP}$, it is natural to study the search complexity of ARRIVAL in the context of total search problems with the guaranteed existence of a solution, i.e., within the complexity class TFNP (which contains the search analogue of $\text{NP} \cap \text{coNP}$). Total search problems are classified into subclasses of TFNP using the methodology proposed by Papadimitriou [12] that clusters computational problems according to the type of argument assuring the existence of a solution. Karthik C. S. [13] noticed that the search for a switching flow is a prime candidate to fit into the hierarchy of TFNP problems. He introduced S-ARRIVAL, a search version of ARRIVAL that seeks a switching flow to either the destination d or the dead-end vertex \bar{d} , and showed that it is contained in the complexity class PLS [9] of total problems amenable to local search.

Fearnley et al. [6] recently studied multiple variants of reachability games on switch graphs and as one of their results gave a lower bound on the complexity of deciding ARRIVAL. Specifically, they showed that ARRIVAL is NL-hard.

1.1 Our Results

One of the open problems suggested by Dohrau et al. [5] was whether deciding the termination of ARRIVAL is contained in $\text{UP} \cap \text{coUP}$. Recall that given a railway network with an origin o and a destination d , the transcript of the route of the train captured in the run profile from o to d (if it exists) is unique. We show that it is possible to efficiently decide whether a switching flow corresponds to a run profile, which provides a positive answer to the above question and places ARRIVAL inside $\text{UP} \cap \text{coUP}$. We similarly also improve the upper bound on the search complexity of ARRIVAL: We show that S-ARRIVAL is contained in the complexity class CLS. Daskalakis and Papadimitriou [4] introduced CLS to classify problems that can be reduced to local search over continuous domains. CLS contains multiple important search problems such as solving simple stochastic games, finding equilibria in congestion games, and solving linear complementarity problems on P-matrices. For all of these problems, as well as for S-ARRIVAL, we currently do not have a polynomial time algorithm, and they are not known to be complete for some subclass of TFNP.

We establish the containment in CLS through a reduction to END-OF-METERED-LINE (EOML), a total search problem that was recently introduced by Hubáček and Yegorov [8] who also showed that it is in CLS. In EOML we are given a source in a directed graph with vertices of in-degree and out-degree at most one, and the task is to find a sink or a source different from the given trivial source. The access to the graph is given locally via information about the successor and predecessor of each vertex together with its distance from the trivial source (for the formal definition see Definition 19).

Our result makes it unlikely for S-ARRIVAL to be PLS-hard, which was one of the possibilities suggested by the containment in PLS shown by Karthik C. S. [13]. This is due to known black-box separations among subclasses of TFNP [11, 2], which suggest that CLS is a proper subclass of PLS. Note that our reduction from S-ARRIVAL to END-OF-METERED-LINE results in instances with a significantly restricted structure: the END-OF-METERED-LINE graph consists only of a single path and many isolated vertices. We believe that this structure may in future work be used to show that S-ARRIVAL is contained in FP.

Our reduction from S-ARRIVAL to END-OF-METERED-LINE also implies that we can use an algorithm by Aldous [1] to solve S-ARRIVAL. The algorithm is randomized and runs in $O(2^{n/2} \text{poly}(n))$ expected time on switch graphs with n vertices. This is the first algorithm with expected runtime $O(c^n)$ for $c < 2$. (A trivial $O(2^n \text{poly}(n))$ time algorithm can be obtained by following the path of the train through the network.) Aldous’ algorithm, in fact, solves any problem in PLS. It samples a large number of candidate solutions and then performs a local

search from the best sampled solution. The advantage of our reduction is that the resulting search space for END-OF-METERED-LINE is small enough to make Aldous’ algorithm useful, unlike in the previous reduction by Karthik C. S. [13] that showed containment in PLS.

Fearnley et al. [7] recently gave a reduction from P -matrix linear complementarity problems (PLCP) to END-OF-METERED-LINE. As in our case for ARRIVAL, this implies that Aldous’ algorithm can be used to solve PLCP. In fact this gives the fastest known randomized algorithm for PLCP, running in expected time $O(2^{n/2} \text{poly}(n))$ for input matrices of dimension $n \times n$. Fearnley et al. do not make this observation themselves, but it is straightforward to check that their reduction also gives an efficient representation of the search space. Although Aldous’ algorithm is very simple, it thus non-trivially improves the best runtime of algorithms for multiple problems. We believe that this way of applying Aldous’ algorithm is a powerful technique that will produce additional results in the future.

1.2 Technical remarks

Recall that a switching flow is a run profile with additional superfluous circulations compared to the valid run profile. Our main technical observation is a characterization of switching flows that correspond to the valid run profile. Given a switch graph G and a switching flow f , we consider the subgraph G^* induced over the railway network by the “last-used” edges; for every vertex v , we include in G^* only the outgoing edge that was, according to the switching flow, used by the train last time it left from v . Note that such last-used edges can be efficiently identified simply by considering the parity of the total number of visits at every vertex. When f is a valid run profile, then it is straightforward to see that the subgraph G^* is acyclic. We show that this property is in fact a characterization, i.e., any switching flow for which the induced graph G^* is acyclic must be a run profile. Given that this property is easy to check, we can use it to efficiently verify run profiles as UP witnesses. (The coUP witness is then a run profile to the dead-end at \vec{d} .)

For our reduction from S-ARRIVAL to END-OF-METERED-LINE we extend the above observation to partial switching flows that are not required to end at the destination. The vertices of the END-OF-METERED-LINE graph created by our reduction correspond to partial switching flows in the S-ARRIVAL instance. The directed edges connect partial run profiles to their natural successors and predecessors, i.e., the partial run extended or shortened by a single step of the train. Any switching flow that does not correspond to some partial run profile is an isolated vertex in the END-OF-METERED-LINE graph. Finally, the trivial source is the empty switching flow, and the distance from it can be computed for any partial run simply as the number of steps taken by the train so far. Given that there is only a single path in the resulting END-OF-METERED-LINE graph and that its sink is exactly the complete run, we get that the unique solution to the END-OF-METERED-LINE instance gives us a solution for the original instance of S-ARRIVAL.

To make the reduction efficiently computable, we need to address the verification of partial run profiles. As it turns out, partial run profiles can be efficiently verified using the graph G^* , in a similar way to complete run profiles discussed above. The main difference is that the graph of last-used edges for a partial run profile can contain a cycle, as the train might visit the same vertex multiple times on its route to the destination. However, we show that there is at most one cycle in G^* , which always contains the current end-vertex of the partial run. The complete characterization of partial run profiles (which covers also full run profiles) is given in Lemma 9, and the formal reduction is described in Section 4.2.1.

Finally, we show that every partial run profile is uniquely determined by its last-used edges and its end-vertex. This limits the size of the search space for the EOML instances that are produced by our reduction, which allows us to efficiently use Aldous’ algorithm [1] to solve ARRIVAL and S-ARRIVAL.

2 Preliminaries

In the rest of the paper we use the following standard notation. For $k \in \mathbb{N}$, we denote by $[k]$ the set $\{1, \dots, k\}$. For a graph $G = (V, E)$, we reserve $n = |V|$ for the number of vertices. The basic object that we study are switch graphs, as defined by Dohrau et al. [5].

Definition 1 (switch graph). A *switch graph* is a tuple $G = (V, E, s_0, s_1)$ where $s_0, s_1: V \rightarrow V$ and $E = \{(v, s_0(v)), (v, s_1(v)) \mid \forall v \in V\}$.¹

¹Whenever $s_0(v) = s_1(v)$ for some vertex $v \in V$ we depict them as multiple edges in figures.

Algorithm 1: RUN

Input : a switch graph $G = (V, E, s_0, s_1)$ and two vertices $o, d \in V$
Output for each edge $e \in E$, the number of times the train traversed e
:
1 $v \leftarrow o$ // position of the train
2 $\forall u \in V$ set $s_curr[u] \leftarrow s_0(u)$ and $s_next[u] \leftarrow s_1(u)$
3 $\forall e \in E$ set $r[e] \leftarrow 0$ // initialize the run profile
4 $step \leftarrow 0$
5 **while** $v \neq d$ **do**
6 $(v, w) \leftarrow s_curr[v]$ // compute the next vertex
7 $r[s_curr[v]]++$ // update the run profile
8 $swap(s_curr[v], s_next[v])$
9 $v \leftarrow w$ // move the train
10 $step \leftarrow step + 1$
11 **return** r

In order to avoid cumbersome notation, we slightly overload the use of s_0, s_1 and treat both as functions from vertices to edges; that is by $s_b(v)$ we denote the edge $(v, s_b(v))$ for $b \in \{0, 1\}$. We use this convention throughout the paper unless stated otherwise.

The ARRIVAL problem was formally defined by Dohrau et al. [5] as follows.

Definition 2 (ARRIVAL [5]). Given a switch graph $G = (V, E, s_0, s_1)$ and two vertices $o, d \in V$, the ARRIVAL problem is to decide whether the algorithm RUN (Algorithm 1) terminates, i.e., whether the train reaches the destination d starting from the origin o .

To simplify theorem statements and our proofs, we assume without loss of generality that both $s_0(d)$ and $s_1(d)$ end in d .

A natural witness for termination of the RUN procedure considered in previous work (e.g., [5]) is a switching flow. We extend the definition of a switching flow to allow for partial switching flows that do not necessarily end in the desired destination d .

Definition 3 ((partial) switching flow, end-vertex). Let $G = (V, E, s_0, s_1)$ be a switch graph. For $o, d \in V$, we say that $f \in \mathbb{N}^{2n}$ is a *switching flow from o to d* if the following two conditions hold.

Kirchhoff's Law (flow conservation):

$$\forall v \in V: \sum_{e=(u,v) \in E} f_e - \sum_{e=(v,w) \in E} f_e = [v = d] - [v = o],$$

where $[\cdot]$ is the indicator variable of the event in brackets.

Parity Condition:

$$\forall v \in V: f_{s_1(v)} \leq f_{s_0(v)} \leq f_{s_1(v)} + 1.$$

Kirchhoff's law means that o emits one unit of flow, d absorbs one unit of flow, and at all other vertices, in-flow equals out-flow. If $d = o$, we have a circulation.

Given an instance $(G = (V, E, s_0, s_1), o, d)$ of ARRIVAL, we say that f is a *switching flow* if it is a switching flow from o to d . A vector $f \in \mathbb{N}^{2n}$ is called a *partial switching flow* iff f is a switching flow from o to v for some vertex $v \in V$. We say that v is the *end-vertex* of the partial switching flow. We denote the end-vertex of f by v_f .

Definition 4 ((partial) run profile). A *run profile* is the switching flow r returned by the algorithm RUN (Algorithm 1) upon termination. A *partial run profile* is a partial switching flow corresponding to some intermediate value of r in the algorithm RUN (Algorithm 1).

Observation 5 (Dohrau et al. [5, Observation 1]). Each (partial) run profile is a (partial) switching flow.

Observation 6. An end-vertex v_f of a switching flow f is computable in polynomial time.

Proof. It is sufficient to determine which vertex has a net in-flow of one. □




Figure 1: An example of a switch graph G with a switching flow f (on the left) and the corresponding graph G_f^* (on the right). The s_0 edges are denoted by full arrows and the s_1 edges are denoted by dashed arrows. For each edge in G , the switching flow f is specified by the adjacent integer value.




Figure 2: An example of a switch graph G and cycles in the graphs G_f^* corresponding to partial run profiles after 3, 4, and 5 steps of the train (respectively from left to right). We use hatching to highlight the current end-vertex.

3 The Complexity of Run Profile Verification

Dohrau et al. [5] proved that it is possible to efficiently verify whether a given vector is a switching flow. In this section we show that we can also efficiently verify whether a switching flow is a run profile. Combining this with the results by Dohrau et al. [5], we prove that the decision problem of ARRIVAL is in $UP \cap coUP$ (see Section 4.1) and that the search problem of ARRIVAL lies in the complexity class CLS (see Section 4.2). As outlined in Section 1.2, our approach for verification of run profiles is based on finding a cycle in a natural subgraph of the railway network G defined below. Specifically, we consider the subgraph of G that contains only the last visited outgoing edge of each vertex, i.e., every vertex has out-degree at most one (see Figure 1 for an illustration).

Definition 7 (G_f^*). Let $(G = (V, E, s_0, s_1), o, d)$ be an instance of ARRIVAL, and let $f \in \mathbb{N}^{2n}$ be a partial switching flow. We define a graph $G_f^* = (V, E^*)$ as follows

$$E^* = \{s_0(v) : \forall v \in V \text{ s.t. } f(s_0(v)) \neq f(s_1(v))\} \cup \{s_1(v) : \forall v \in V \text{ s.t. } f(s_0(v)) = f(s_1(v)) > 0\}.$$

Observation 8. Given a partial switching flow f , the graph G_f^* can be computed in polynomial time.

Lemma 9. A partial switching flow f is a partial run profile iff $f_{s_0(d)} = f_{s_1(d)} = 0$ and one of the following two conditions holds:

1. There exists no cycle in G_f^* .
2. There exists exactly one cycle in G_f^* and this cycle contains the end-vertex of f .

The main idea of the proof is based on the following fact: a switching flow f which is *not* a run profile must contain a circulation (as shown by Dohrau et al. [5]). Let f be a switching flow that we get from a run profile r by adding some flows on cycles, then the last added circulation (the last added cycle) must form a cycle in the corresponding graph G_f^* . On the other hand, a cycle containing the end-vertex is formed in G_f^* whenever the train arrives to a previously visited vertex. An illustration of the graph G^* at consecutive steps of the algorithm RUN, with the corresponding evolution of the end-vertices and cycles, is given in Figure 2.

We prove the “ \Leftarrow ” implication of Lemma 9 by contradiction. Given a switching flow f we consider the longest run profile r that is everywhere at most f . We denote the difference $f - r$ by δ . We utilize the following

observation at two points in our proof. First, to prove that both f and r have the same end-vertex, that is $v_r = v_f$. Second, to prove that any cycle we find in the intersection of G_f^* and the non-zero edges of δ avoids v_f .

Observation 10. Let f be a switching flow and let r be the longest partial run profile such that all coordinates are at most f and denote $\delta = f - r$. Then $\delta_{s_0(v_r)} = \delta_{s_1(v_r)} = 0$.

Proof of Observation 10. Suppose that either $\delta_{s_0(v_r)} \neq 0$ or $\delta_{s_1(v_r)} \neq 0$, and that r is such that the next edge for the train to continue on is $s_0(v_r)$. From the maximality of r we get that $\delta_{s_0(v_r)}$ is equal to zero. Then we get

$$\begin{aligned} f_{s_1(v_r)} &> r_{s_1(v_r)} && \text{(since } \delta_{s_1(v_r)} \text{ is non-zero)} \\ &= r_{s_0(v_r)} && \text{(from the parity condition for } r) \\ &= f_{s_0(v_r)} && \text{(from the maximality of } r) \end{aligned}$$

This leads to a contradiction with the parity condition of Definition 3 as $f_{s_1(v_r)} > f_{s_0(v_r)}$.

The other case is similar: suppose that the train should continue with the edge $s_1(v_r)$. From the maximality of r we get that $\delta_{s_1(v_r)}$ is equal to zero. And thus we get

$$\begin{aligned} f_{s_0(v_r)} &> r_{s_0(v_r)} && \text{(since } \delta_{s_0(v_r)} \text{ is non-zero)} \\ &= r_{s_1(v_r)} + 1 && \text{(from the parity condition for } r) \\ &= f_{s_1(v_r)} + 1 && \text{(from the maximality of } r) \end{aligned}$$

This gives a contradiction with the parity condition of Definition 3 as $f_{s_0(v_r)} > f_{s_1(v_r)} + 1$. \square

Proof of Lemma 9. We start by proving that any partial run profile r has at most one cycle and that this cycle always contains the end-vertex v_r . We proceed by induction on the length of the run profile, i.e., the number of steps in the algorithm RUN. The base case is the initial run profile 0^{2n} for which $G_{0^{2n}}^*$ contains no cycle. For the induction step, let us assume that a vector r^i is a partial run profile and the vector r^{i+1} is the partial run profile after one step from r^i .

If $G_{r^i}^*$ contains a cycle then removing the edge from the end-vertex v_{r^i} to its successor in $G_{r^i}^*$ produces a cycle-free graph. Adding an edge between v_{r^i} and $v_{r^{i+1}}$ creates a graph with at most one cycle. Moreover, if there is a cycle then it has to contain $v_{r^{i+1}}$.

To prove the reverse implication, assume we have a partial switching flow f that satisfies the conditions from the statement of this lemma, i.e., G_f^* has at most one cycle and if the cycle is present it contains the end-vertex v_f . Let us denote by r the longest partial run profile such that all its coordinates are at most f , that is $r_e \leq f_e$ for each $e \in E$. Consider the difference $\delta = f - r$ of the switching flow and its longest run profile. We complete the proof by the following case analysis:

1. If the end-vertex v_f is the same as the end-vertex v_r ($v_f = v_r$), then the difference $\delta = f - r$ is a flow, and moreover it is a circulation. That is the Kirchhoff's law is satisfied in all vertices of G :

$$\forall v \in G: \sum_{e=(u,v) \in E} \delta_e - \sum_{e=(v,w) \in E} \delta_e = 0.$$

If δ is identically zero we are done as $f = r$, and thus f is a partial run profile.

Otherwise, we show that the circulation δ will result in a cycle in G_f^* . Namely we prove that G_f^* contains an outgoing edge from v with a non-zero value in δ iff v has a non-zero outgoing edge in δ . Using this and the fact that δ satisfies the Kirchhoff's law everywhere, we find a cycle in G_f^* .

Let u be any vertex of G such that only one of the edges $s_0(u)$ and $s_1(u)$ is non-zero in δ . We claim that then the non-zero edge is contained in the graph G_f^* . There are only two possible cases:

- (a) Either $r_{s_0(u)} = f_{s_0(u)}$, and thus $r_{s_0(u)} = r_{s_1(u)} - 1$ and $f_{s_0(u)} = f_{s_1(u)}$ (see Table 1a), and G_f^* contains the edge $s_1(u)$,
- (b) or $r_{s_1(u)} = f_{s_1(u)}$, and thus $r_{s_0(u)} = r_{s_1(u)}$ and $f_{s_0(u)} = f_{s_1(u)} + 1$ (see Table 1b), and G_f^* contains the edge $s_0(u)$.

	$s_0(u)$	$s_1(u)$
δ	0	1
r	a	$a - 1$
f	a	a

(a) Case (1.a).

	$s_0(u)$	$s_1(u)$
δ	1	0
r	a	a
f	$a + 1$	a

(b) Case (1.b).

Table 1: Case analysis from the proof of Lemma 9 when only one of the edges $s_0(u)$ and $s_1(u)$ is non-zero in δ .

On the other hand, if both $s_0(u)$ and $s_1(u)$ are non-zero in δ then G_f^* contains one of them. Thus we may choose any non-zero edge in δ which is in G_f^* and proceed by another adjacent non-zero directed edge (in δ) in G_f^* until we construct a directed cycle in G_f^* .

By the Observation 10, we know that all outgoing edges from $v_r = v_f$ are zero in δ , and thus the end-vertex v_f is not contained in the cycle we have found.

2. If the end-vertex v_f is not the same as the end-vertex v_r ($v_f \neq v_r$) then we get a contradiction with f being a partial switching flow. It follows from Observation 10 that

$$\sum_{e=(u,v_r) \in E} f_e - \sum_{e=(v_r,w) \in E} f_e \geq \sum_{e=(u,v_r) \in E} r_e - \sum_{e=(v_r,w) \in E} r_e = \begin{cases} 0 & v_r = o, \\ 1 & \text{otherwise,} \end{cases}$$

which is in contradiction with f being a partial switching flow for which the end-vertex v_f differs from v_r .

This concludes the proof of Lemma 9. \square

Lemma 11. It is possible to verify in polynomial time whether a vector is a run profile.

Proof. We can check that a vector f is a switching flow in polynomial time due to Dohrau et al. [5]. The construction of the graph G_f^* is polynomial by Observation 8. Lemma 9 gives us a polynomial time procedure to check if f is also a run profile as it is sufficient to check if G_f^* contains more than one cycle or whether it has a cycle not containing the end-vertex. This check can be done by a simple modification of the standard depth-first search on G_f^* . \square

4 The Computational Complexity of ARRIVAL

In this section we use our efficient structural characterization of run profiles from Lemma 9 to improve the known results about the computational complexity of ARRIVAL. Specifically, we show that the decision version of ARRIVAL is in $UP \cap coUP$ and the search version is in CLS.

4.1 The Decision Complexity of ARRIVAL

Our upper bound on the decision complexity of ARRIVAL follows directly from the work of Dohrau et al. [5] by application of Lemma 11.

Theorem 12. ARRIVAL is in $UP \cap coUP$.

Proof. The unique UP certificate for a YES-instance of ARRIVAL is the run profile r returned by the algorithm RUN. Clearly, for each YES-instance there exists only one such vector r and r does not exist for NO-instances. By Lemma 11, we can determine whether a candidate switching flow r is a run profile in polynomial time.

The coUP membership follows directly from the reduction of NO-instances of ARRIVAL to YES-instances of ARRIVAL as suggested by Dohrau et al. [5]. The reduction adds to the original graph G a new vertex \bar{d} , and for each vertex $v \in V$ such that there is no directed path from v to the destination d , the edges $s_0(v)$ and $s_1(v)$ are replaced with edges (v, \bar{d}) . This alteration of the original switch graph can be performed in polynomial time. Dohrau et al. [5] proved that the train eventually arrives either at d or \bar{d} . The unique coUP witness for ARRIVAL is then a run profile from o to the dead-end \bar{d} . \square

4.2 The Search Complexity of ARRIVAL

The search complexity of ARRIVAL was first studied by Karthik C. S. [13], who introduced a total search variant of ARRIVAL as follows.

Definition 13 (S-ARRIVAL [13]). Given a switch graph $G = (V, E, s_0, s_1)$ and a pair of vertices $o, d \in V$, define a graph G' as follows:

1. Add a new vertex \bar{d} .
2. For each vertex v such that there is no directed path from v to d , replace edges $s_0(v)$ and $s_1(v)$ with edges (v, \bar{d}) .
3. Edges $s_0(d), s_1(d), s_0(\bar{d})$, and $s_1(\bar{d})$ are self-loops.

The problem S-ARRIVAL is to find a switching flow in G' either from o to d or from o to \bar{d} .

The above Definition 13 is motivated by the proof of membership in $\text{NP} \cap \text{coNP}$ by Dohrau et al. [5]. Namely, in order to ensure that a solution for S-ARRIVAL always exists, it was necessary to add to the switch graph G the dead-end vertex \bar{d} .

Note that our method for efficient verification of run profiles from Lemma 11 allows us to define a more natural version of S-ARRIVAL directly on the graph G without any modifications. Instead of relying on the dead-end vertices, we can use the fact that a partial run profile with an edge that was visited for $2^n + 1$ times is an efficiently verifiable witness for NO-instances of ARRIVAL.

Definition 14 (S-ARRIVAL - simplified). Given a switch graph $G = (V, E, s_0, s_1)$ and a pair of vertices $o, d \in V$, the S-ARRIVAL problem asks us to find one of the following:

1. a run profile $\mathbf{r} \in \mathbb{N}^{[2^n]}$ from o to d , or
2. a run profile $\mathbf{r} \in \mathbb{N}^{[2^n]}$ from o to any $v \in V$ such that
 - $\mathbf{r}_{(u,v)} = 2^n + 1$, where u is the last vertex visited by the train before it reached the end-vertex v of \mathbf{r} , and
 - $\mathbf{r}_{e'} \leq 2^n$ for all $e' \neq (u, v)$.

The correspondence of the above version of S-ARRIVAL to the original one follows formally from the following lemma.

Lemma 15 (Karthik C. S. [13, Lemma 1]). For any $G = (V, E, s_0, s_1)$ and a pair of vertices $o, d \in V$. Let \mathbf{r} be a run profile (thus $v_{\mathbf{r}} = d$), then $\mathbf{r}_e \leq 2^n$ for each edge $e \in E$.

To argue membership of our version of S-ARRIVAL in TFNP, we need to show that both types of solutions in Definition 14 can be verified efficiently. Solutions of the first type are simply run profiles, and we have already shown that they can be verified in polynomial time in Lemma 11. In order to be able to verify solutions of the second type, it remains to argue that for any partial run profile, the immediate predecessor of its end-vertex can be determined in polynomial time.

Lemma 16. Let \mathbf{r} be a partial run profile after $R \geq 1$ steps and u be the vertex visited by the train at step $R - 1$. Then

1. either u is the unique predecessor of $v_{\mathbf{r}}$ in $G_{\mathbf{r}}^*$, or
2. there is a single cycle in $G_{\mathbf{r}}^*$ containing $v_{\mathbf{r}}$ and u is the predecessor of $v_{\mathbf{r}}$ on this cycle.

Proof. First, note that if u is the end-vertex one step before $v_{\mathbf{r}}$ becomes the end-vertex then $G_{\mathbf{r}}^*$ must contain the edge $(u, v_{\mathbf{r}})$, as it is the last edge used by the train to leave u . Thus, in the first case the immediate predecessor of $v_{\mathbf{r}}$ in the partial run \mathbf{r} is unambiguously given by the only predecessor of $v_{\mathbf{r}}$ in $G_{\mathbf{r}}^*$.

For the second case we show that $G_{\mathbf{r}}^*$ contains a directed cycle C (containing the end-vertex $v_{\mathbf{r}}$) and u is unambiguously given by the predecessor of $v_{\mathbf{r}}$ in $G_{\mathbf{r}}^*$ that lies on C . We find the cycle C by constructing the longest possible directed path $c_0 = v_{\mathbf{r}}, c_1, \dots, c_k$ in $G_{\mathbf{r}}^*$ without repeating vertices. Note that it cannot happen that c_k has no outgoing edge in $G_{\mathbf{r}}^*$. Otherwise, \mathbf{r} would have two different end-vertices $v_{\mathbf{r}}$ and c_k (as having no outgoing edge

in G_r^* means that the train has never left this vertex). By Lemma 9, the directed edge from c_k has to end in the end-vertex v_r , or else there would be a cycle in G_r^* that avoids v_r .

The algorithm RUN takes R steps to generate the run profile \mathbf{r} , i.e., $\sum_{e \in E} r_e = R$. Let $t_r : V \rightarrow \{0, 1, \dots, R-1\}$ be the function returning the last step after which a vertex was left by the train in the partial run profile \mathbf{r} . Observe that, except for the edge through which the train arrived to v_r ,² it holds for all edges $(x, y) \in G_r^*$ that $t_r(x) < t_r(x) + 1 \leq t_r(y)$. However, the above inequality cannot hold for all edges on the cycle C , and thus C has to contain the last used edge and the train had to be in c_k at step $R-1$. \square

Observation 17. S-ARRIVAL from Definition 13 reduces to simplified S-ARRIVAL from Definition 14.

Proof. Given a solution of the second type of the simplified S-ARRIVAL, i.e., the long run profile, we can get a run profile \mathbf{r} to \vec{d} in polynomial time. For each vertex u we can determine whether there is an oriented path from it to the destination d , and if there is no such path we set $r_{s_0(v)} = r_{s_1(v)} = 0$. We compute the end vertex v_r and set $s_0(v_r) = 1$. All other components of \mathbf{r} are set according to the original solution of the simplified S-ARRIVAL. \square

4.2.1 S-ARRIVAL is in CLS

Karthik C. S. [13] showed that S-ARRIVAL is contained in the class PLS. We improve this result and prove that S-ARRIVAL is in fact contained in CLS. As a by-product, we also obtain a randomized algorithm for S-ARRIVAL with runtime $\mathcal{O}(1.4143^n)$ which is the first algorithm for this problem with expected runtime $\mathcal{O}(c^n)$ for $c < 2$.

The class of total search problems that are amenable to “continuous” local search was defined by Daskalakis and Papadimitriou [4] using the following canonical problem.

Definition 18 (CLS [4]). CLS is the class of total search problems reducible to the following problem called CLOPT.

Given two arithmetic circuits $f : [0, 1]^3 \rightarrow [0, 1]^3$ and $p : [0, 1]^3 \rightarrow [0, 1]$, and two real constants $\varepsilon, \lambda > 0$, find either a point $x \in [0, 1]^3$ such that $p(f(x)) \leq p(x) + \varepsilon$ or a pair of points $x, x' \in [0, 1]^3$ certifying that either p or f is not λ -Lipschitz.

Instead of working with CLOPT, we use as a gateway for our reduction a problem called END-OF-METERED-LINE (EOML) which was recently defined and shown to lie in CLS by Hubáček and Yogev [8].

Definition 19 (END-OF-METERED-LINE). Given circuits $S, P : \{0, 1\}^m \rightarrow \{0, 1\}^m$, and $V : \{0, 1\}^m \rightarrow [2^m] \cup \{0\}$ such that $P(0^m) = 0^m \neq S(0^m)$ and $V(0^m) = 1$, find a string $x \in \{0, 1\}^m$ satisfying one of the following:

1. either $P(S(x)) \neq x$ or $S(P(x)) \neq x \neq 0^m$,
2. $x \neq 0^m$ and $V(x) = 1$,
3. either $V(x) > 0$ and $V(S(x)) - V(x) \neq 1$ or $V(x) > 1$ and $V(x) - V(P(x)) \neq 1$.

The circuits S, P from Definition 19 implicitly represent a directed graph with vertices labelled by binary strings of length m , where each vertex has both out-degree and in-degree at most one. The circuit P represents the predecessor and the circuit S represents the successor of a given vertex as follows: there is an edge from a vertex u to a vertex v iff $S(u) = v$ and $P(v) = u$. Finally, the circuit V can be thought of as an odometer that returns the distance from the trivial source at 0^m or value 0 for vertices lying off the path starting at the trivial source. The task in END-OF-METERED-LINE is to find a sink or a source different from the trivial one at 0^m (the solutions of the second and of the third type in Definition 19 ensure that V behaves as explained above).

We are now ready to present our reduction from S-ARRIVAL to END-OF-METERED-LINE.

Theorem 20. S-ARRIVAL can be reduced to END-OF-METERED-LINE, and thus it is contained in CLS.

Proof. Let (G, o, d) be an instance of S-ARRIVAL. We construct an instance of EOML that contains a vertex for each candidate partial switching flow over the switch graph G , i.e., for each vector with $2n$ coordinates and values from $[2^n + 1] \cup \{0\}$. The EOML instance will comprise of a directed path starting at the initial (empty) partial run profile 0^{2n} . Each vertex on the path has an outgoing edge to its consecutive partial run profile. Any vertex that does not correspond to a partial run profile becomes a self-loop. Finally, the valuation circuit V returns either the number of steps in the corresponding partial run profile or the zero value if the vertex does not correspond to a partial run profile.

Algorithm 2: SUCCESSOR

Input : a vector $\mathbf{x} \in ([2^n + 1] \cup \{0\})^{2n}$
Output : a vector $\mathbf{y} \in ([2^n + 1] \cup \{0\})^{2n}$
:
1 **if** \mathbf{x} is a partial run profile **then** // efficiently testable by Lemma 11
2 compute the end-vertex v_x
3 **if** $v_x = d$ or $x_e = 2^n + 1$ for some $e \in E$ **then**
4 **return** \mathbf{x} // the train terminates or runs for too long
5 **else**
6 $b \leftarrow \mathbf{x}(s_1(v_x)) - \mathbf{x}(s_0(v_x))$ // parity of the current visit at v_x
7 $e \leftarrow s_b(v_x)$ // the next edge to traverse
8 $y_h \leftarrow \begin{cases} x_h + 1 & \text{if } h = e \\ x_h & \text{otherwise} \end{cases}$ // run profile update
9 **return** \mathbf{y}
10 **else**
11 **return** \mathbf{x} // self-loop

Algorithm 3: PREDECESSOR

Input : a vector $\mathbf{x} \in ([2^n + 1] \cup \{0\})^{2n}$
Output : a vector $\mathbf{y} \in ([2^n + 1] \cup \{0\})^{2n}$
:
1 **if** $\mathbf{x} = \mathbf{0}$ **then**
2 **return** $\mathbf{0}$ // the trivial source
3 **if** \mathbf{x} is a partial run profile **then** // efficiently testable by Lemma 11
4 compute the end-vertex v_x
5 // correctness by Lemma 16
6 **if** v_x has a single predecessor in G_x^* **then**
7 $e \leftarrow$ the only incoming edge of v_x in G_x^*
8 **else**
9 $e \leftarrow$ the only incoming edge of v_x which lies on a directed cycle
10 **if** $x_e \leq 2^n + 1$ and $x_{e'} < 2^n + 1$ for all $e' \neq e$ **then**
11 $y_h \leftarrow \begin{cases} x_h - 1 & \text{if } h = e \\ x_h & \text{otherwise} \end{cases}$
12 **return** \mathbf{y}
12 **return** \mathbf{x} // self-loop

Algorithm 4: VALUATION

Input : a vector $\mathbf{x} \in ([2^n + 1] \cup \{0\})^{2n}$
Output : a value $v \in \mathbb{N}$
:
1 $v \leftarrow 0$ // default value for self-loops
2 **if** \mathbf{x} is a partial run profile **then** // efficiently testable by Lemma 11
3 $v \leftarrow 1 + \sum_{i=1}^{2n} x_i$
4 **return** v

Formal description of the circuits S , P , and V defining the above EOML graph is given by algorithms SUCCESSOR (Algorithm 2), PREDECESSOR (Algorithm 3), and VALUATION (Algorithm 4).

A polynomial bound on the size of the circuits S , P , and V follows directly from Observation 8 (computing G^*), Lemma 11 (testing whether a given vector is a partial run profile), Observation 6 (computing the end-vertex), and Lemma 16 (computing the previous position of the train).

²The inequality does not hold for v_r , since $t(v_r)$ has not been updated to the time R yet.

Lemma 9 and Lemma 16 imply that the EOML graph indeed consists of a single directed path and isolated vertices with self-loops. By the construction of V (it outputs the number of steps of the train), there are no solutions of the second or the third type (cf. Definition 19). Thus, the EOML instance has a unique solution which has to correspond to a run profile in the original S-ARRIVAL instance or to a partial run profile certifying that the train ran for too long (see the second type of solution in Definition 14). \square

5 An $O(1.4143^n)$ Algorithm for S-ARRIVAL

Consider any problem that can be put into the complexity class PLS, i.e., can be reduced to the canonical PLS-complete problem LOCALOPT (see also [13, Definition 1]):

Definition 21 (LOCALOPT). Given circuits $S: \{0, 1\}^m \rightarrow \{0, 1\}^m$, and $V: \{0, 1\}^m \rightarrow [2^m] \cup \{0\}$, find a string $x \in \{0, 1\}^m$ such that $V(x) \geq V(S(x))$.

Aldous [1] introduced the following simple algorithm that can be used to solve LOCALOPT: pick $2^{m/2}$ binary strings uniformly and independently at random from $\{0, 1\}^m$, and let x_{\max} be the selected string that maximizes the value $V(x)$. Starting from $x = x_{\max}$, repeatedly move to the successor $S(x)$, until $V(x) \geq V(S(x))$. He showed that the expected number of circuit evaluations performed by the algorithm before finding a local optimum is at most $O(m2^{m/2})$. Note that in the case of EOML it is possible that all sampled solutions are isolated vertices in the EOML graph. In this case the 0^m string has the best known value, and the search is started from there.

In the case of S-ARRIVAL, the PLS membership proof of Karthik C. S. constructs the circuits for successor and valuation with $O(n^2)$ input bits [13, Theorem 2]; the EOML instance constructed in Theorem 20—proving CLS and in particular PLS membership—yields circuits of $O(n^2)$ input bits as well. This number is too high to yield a randomized algorithm of non-trivial runtime.

This number of $O(n^2)$ input bits comes from the obvious encoding of a partial run profile: each of the $2n$ edges has a nonnegative integer flow value of at most $2^n + 1$. But in fact, a *terminating* run has at most $n2^n$ partial run profiles, as no *vertex-state pair* (v, s_{curr}) can repeat in Algorithm 1. In other words, a partial run profile f is determined by its end-vertex v_f as well as the positions of all switches at the time of the corresponding visit of v_f . This means that a partial run profile can be encoded with $n + \log_2 n$ bits, and if we had a PLS or CLS membership proof of S-ARRIVAL with circuits of this many inputs only, we could solve S-ARRIVAL in time $O(\text{poly}(n) \times 2^{n/2})$.

Next, we show that such membership proofs indeed exist. For this, we show that the above encoding (of a partial run profile by an end-vertex and the positions of all switches) can be efficiently decoded: given an end-vertex and the positions of all switches, we can efficiently compute a unique candidate for a corresponding partial run profile. The resulting encoding and decoding circuits can be composed with the ones in Theorem 20 to obtain PLS and CLS membership proofs with circuits of $n + \log_2 n$ input bits, and hence yield a randomized algorithm of runtime $O(\text{poly}(n) \times 2^{n/2}) = O(1.4143^n)$, as explained above.

5.1 Decoding Partial Run Profiles

We work with instances of S-ARRIVAL as in Definition 13, i.e., there is always a run profile either to d or to \bar{d} . This is without loss of generality [5, 13].

Definition 22. Let $G = (V, E, s_0, s_1)$ be a switch graph, and let $o, d, \bar{d} \in V$ be as in Definition 13. The *parity* of a run profile $f \in \mathbb{N}^{2n}$ is the vector $\mathbf{p}_f \in \{0, 1\}^{n-2}$ defined by $\mathbf{p}_v = \mathbf{f}_{s_0(v)} - \mathbf{f}_{s_1(v)} \in \{0, 1\}$, $v \in V \setminus \{d, \bar{d}\}$.

Note that we do not care about (the parity of) the switches at d and \bar{d} , since the algorithm stops as soon as d or \bar{d} is reached.

Here is the main result of this section. For a given target vertex $t \in V$ and given parity \mathbf{p} , there is exactly one candidate for a partial run profile from o to t with parity \mathbf{p} . Moreover, this candidate can be computed by solving a system of linear equations.

Lemma 23. Let (G, o, d, \bar{d}) be an instance of S-ARRIVAL, let $t \in V$ and $\mathbf{p} \in \{0, 1\}^{n-2}$. Then there exists exactly one vector $f \in \mathbb{R}^{2(n-2)}$ such that the following conditions hold.

Kirchhoff's Law (flow conservation):

$$\forall v \in V \setminus \{d, \bar{d}\}: \sum_{e=(u,v) \in E} \mathbf{f}_e - \sum_{e=(v,w) \in E} \mathbf{f}_e = [v = t] - [v = o] \quad (1)$$

where $[\cdot]$ is the indicator variable of the event in brackets.

Parity Condition: $p_f = p$, i.e.,

$$\forall v \in V \setminus \{d, \bar{d}\}: f_{s_0(v)} - f_{s_1(v)} = p_v. \quad (2)$$

Before we prove Lemma 23, let us draw a crucial conclusion: The unique partial run profile $f \in \mathbb{N}^{2n}$ with end-vertex t and parity p (if such f exists – note that we are only guaranteed a real-valued f) necessarily satisfies (1) and (2). Hence, we may use Lemma 23 to get the entries f_e for all edges except the ones leaving d and \bar{d} . Only if all the entries are nonnegative and integral and satisfy (1) at d and \bar{d} (under $f_{s_0(d)} = f_{s_1(d)} = f_{s_0(\bar{d})} = f_{s_1(\bar{d})} = 0$) do we have a candidate for a partial run profile. Hence, there is a unique candidate, and given t and a p , this candidate can be efficiently found.

Proof of Lemma 23. Set $m = 2(n - 2)$, $V' = V \setminus \{d, \bar{d}\}$ and let $A \in \mathbb{Z}^{m \times m}$ be the coefficient matrix of the linear system (1), (2) in the variables f_e . We show that A is invertible.

Let $q \in \mathbb{R}^m$ be the vector such that $q_{(v, s_i(v))} = -1$ if $s_i(v) = d$ and $q_{(v, s_i(v))} = 0$ otherwise. We show that q can be expressed as a linear combination of the rows of A in a unique way, from which invertibility of A and the statement of the lemma follow.

Let us use coefficients λ_v for each $v \in V'$ for the rows corresponding to the flow conservation constraints (1), and coefficients μ_v for each $v \in V'$ for the rows corresponding to the parity constraints (2). The column of A corresponding to variable $f_{(v, s_i(v))}$, has a -1 entry from the flow conservation constraint at v , and a 1 entry (if $i = 0$) or a -1 entry (if $i = 1$) from the parity constraint at v . If $s_i(v) \neq d, \bar{d}$, there is another 1 entry from the flow conservation constraint at $s_i(v)$. All other entries are zero. The equations that express q as a linear combination of rows of A are therefore the following.

$$\begin{aligned} \forall v \in V': \quad -\lambda_v + \mu_v + \lambda_{s_0(v)} \cdot [s_0(v) \neq d, \bar{d}] &= q_{(v, s_0(v))}, \\ \forall v \in V': \quad -\lambda_v - \mu_v + \lambda_{s_1(v)} \cdot [s_1(v) \neq d, \bar{d}] &= q_{(v, s_1(v))}. \end{aligned}$$

Or equivalently:

$$\lambda_v - \mu_v = \begin{cases} \lambda_{s_0(v)}, & \text{if } s_0(v) \neq d, \bar{d} \\ 1, & \text{if } s_0(v) = d \\ 0, & \text{if } s_0(v) = \bar{d} \end{cases} \quad v \in V', \quad (3)$$

$$\lambda_v + \mu_v = \begin{cases} \lambda_{s_1(v)}, & \text{if } s_1(v) \neq d, \bar{d} \\ 1, & \text{if } s_1(v) = d \\ 0, & \text{if } s_1(v) = \bar{d} \end{cases} \quad v \in V'. \quad (4)$$

We now show that there are unique coefficients λ_v, μ_v satisfying these equations. Let us define $\lambda_d = 1$ and $\lambda_{\bar{d}} = 0$. Adding corresponding equations of (3) and (4) then yields

$$\lambda_d = 1, \quad \lambda_{\bar{d}} = 0, \quad \lambda_v = \frac{1}{2} (\lambda_{s_0(v)} + \lambda_{s_1(v)}), \quad \forall v \in V'.$$

These are exactly the equations for the vertex values in a *stopping simple stochastic game* on the graph G with only average or degree-1 vertices and sinks d and \bar{d} (stopping means that d or \bar{d} are reachable from everywhere which is exactly what we require in a switch graph). Condon proved that these values are unique [3]. This also determines the μ_v 's uniquely. \square

6 Conclusion and Open Problems

We showed that candidate run profiles in ARRIVAL can be efficiently verified due to their structure. This allowed us to improve the known upper bounds for the search complexity of ARRIVAL and S-ARRIVAL. Here we mention some natural questions arising from our work.

- Are there any non-trivial graph properties that make ARRIVAL or S-ARRIVAL efficiently solvable? Given that we currently do not know of any polynomial time algorithm for ARRIVAL on general switch graphs, we could study the complexity of ARRIVAL on some interesting restricted classes of switch graphs.
- Are there other natural problems in $\text{UP} \cap \text{coUP}$ such that their corresponding search variant is reducible to EOML? Does END-OF-METERED-LINE capture the computational complexity of any TFNP problem with unique solution? Fearnley et al. [7] recently gave a reduction from the PLCP to EOML. Given that ARRIVAL and PLCP can be both reduced to EOML, yet another intriguing question is whether there exists any reduction between the two.

- As mentioned in Section 1.1, the reduction from PLCP to EOML by Fearnley et al. [7] implies that PLCP can be solved faster with Aldous' algorithm [1] than with any other known algorithm. It would be interesting to see whether Aldous' algorithm can similarly give improved runtimes for other problems than ARRIVAL and PLCP.

Acknowledgements. We wish to thank Karthik C. S. for helpful discussions and suggestions.

References

- [1] David Aldous. Minimization algorithms and random walk on the d-cube. *The Annals of Probability*, 11(2):403–413, 1983.
- [2] Josh Buresh-Oppenheim and Tsuyoshi Morioka. Relativized NP search problems and propositional proof systems. In *19th Annual IEEE Conference on Computational Complexity (CCC 2004), 21-24 June 2004, Amherst, MA, USA*, pages 54–67, 2004.
- [3] Anne Condon. The complexity of stochastic games. *Information and Computation*, 96(2):203 – 224, 1992.
- [4] Constantinos Daskalakis and Christos H. Papadimitriou. Continuous local search. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA, January 23-25, 2011*, pages 790–804, 2011.
- [5] Jérôme Dohrau, Bernd Gärtner, Manuel Kohler, Jiří Matoušek, and Emo Welzl. ARRIVAL: A zero-player graph game in $NP \cap coNP$. In Martin Loebl, Jaroslav Nešetřil, and Robin Thomas, editors, *A Journey Through Discrete Mathematics: A Tribute to Jiří Matoušek*, pages 367–374. Springer International Publishing, 2017.
- [6] John Fearnley, Martin Gairing, Matthias Mnich, and Rahul Savani. Reachability switching games. *CoRR*, abs/1709.08991, 2017. URL: <http://arxiv.org/abs/1709.08991>.
- [7] John Fearnley, Spencer Gordon, Ruta Mehta, and Rahul Savani. CLS: new problems and completeness. *CoRR*, abs/1702.06017, 2017. URL: <http://arxiv.org/abs/1702.06017>.
- [8] Pavel Hubáček and Eylon Yogev. Hardness of continuous local search: Query complexity and cryptographic lower bounds. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, pages 1352–1371, 2017.
- [9] David S. Johnson, Christos H. Papadimitriou, and Mihalis Yannakakis. How easy is local search? *J. Comput. Syst. Sci.*, 37(1):79–100, 1988.
- [10] Bastian Katz, Ignaz Rutter, and Gerhard J. Woeginger. An algorithmic study of switch graphs. *Acta Inf.*, 49(5):295–312, 2012.
- [11] Tsuyoshi Morioka. Classification of search problems and their definability in bounded arithmetic. *Electronic Colloquium on Computational Complexity (ECCC)*, (082), 2001.
- [12] Christos H. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *J. Comput. Syst. Sci.*, 48(3):498–532, 1994.
- [13] Karthik C. S. Did the train reach its destination: The complexity of finding a witness. *Inf. Process. Lett.*, 121:17–21, 2017.