

Further Collapses in TFNP

Mika Göös
EPFL

Alexandros Hollender
University of Oxford

Siddhartha Jain
EPFL

Gilbert Maystre
EPFL

William Pires
McGill University

Robert Robere
McGill University

Ran Tao
McGill University

February 15, 2022

Abstract. We show $\text{EOPL} = \text{PLS} \cap \text{PPAD}$. Here the class EOPL consists of all total search problems that reduce to the END-OF-POTENTIAL-LINE problem, which was introduced in the works by Hubáček and Yogevev (SICOMP 2020) and Fearnley et al. (JCSS 2020). In particular, our result yields a new simpler proof of the breakthrough collapse $\text{CLS} = \text{PLS} \cap \text{PPAD}$ by Fearnley et al. (STOC 2021). We also prove a companion result $\text{SOPL} = \text{PLS} \cap \text{PPADS}$, where SOPL is the class associated with the SINK-OF-POTENTIAL-LINE problem.

1 Introduction

Our main results are two collapses of total NP search problem (TFNP) classes.

Theorem 1. $\text{EOPL} = \text{PLS} \cap \text{PPAD}$.

Theorem 2. $\text{SOPL} = \text{PLS} \cap \text{PPADS}$.

Let us explain what these collapses mean and how they fit into the diverse complexity zoo of search problem classes, as summarised in [Figure 1](#). The classes PLS, PPAD, PPADS are classical. They were all introduced in the original pioneering works [[MP91](#), [JPY88](#), [Pap94](#)] that founded the theory of TFNP. To define these classes, it is most convenient to describe a canonical complete problem for each class. (See [Section 2](#) for more formal definitions).

PLS: SINK-OF-DAG (SoD). We are given *implicit access* to a directed graph $G = (V, E)$ that is acyclic, has out-degree at most 1, and has exponentially many nodes, $|V| = 2^n$. The graph is described by a $\text{poly}(n)$ -sized circuit: for any node $v \in V$, we can compute its unique *successor* (out-neighbour) u , if any, and also an integer *potential*, which is guaranteed to increase along the direction of the edge (v, u) . The goal is to find a *sink* node (out-degree 0).

PPAD: END-OF-LINE (EOL). We are given access to a directed graph $G = (V, E)$ that has in/out-degree at most 1, and has $|V| = 2^n$ nodes. The graph is described by a $\text{poly}(n)$ -sized circuit: for any $v \in V$, we can compute its *successor* u and *predecessor* u' , if any. We are guaranteed that if v 's successor is u , then u 's predecessor is v , and vice versa. In addition, we are given the name of a *distinguished source* node v^* (in-degree 0, out-degree 1). The goal is to find any source or sink other than v^* .

PPADS: SINK-OF-LINE (SOL). Same as EOL except the goal is to find a sink.

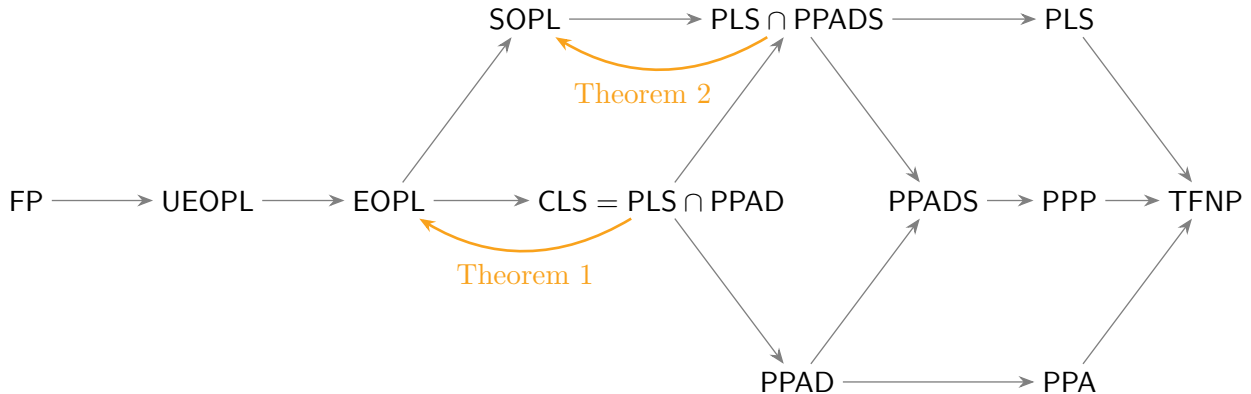


Figure 1: Class diagram for TFNP with new inclusions highlighted. An arrow $A \rightarrow B$ denotes $A \subseteq B$.

Modern classes. Research in the past decade has studied several relatively weak classes of search problems that lie below PLS and PPAD. The intersection class $PLS \cap PPAD$ is, of course, one immediate such example. This class, however, feels quite artificial at first glance. It does not seem to admit any “natural” complete problem. Motivated by this, Daskalakis and Papadimitriou [DP11] introduced the *continuous local search* class $CLS \subseteq PLS \cap PPAD$, which, by its very definition, admits natural complete problems related to the local optimisation of continuous functions over the real numbers (computed by arithmetic circuits). The class CLS is exceptional in that it captures the complexity of *real continuous* optimisation problems, while most classical search problem classes are designed to capture *combinatorial principles*, often phrased in terms of directed graphs.

In order to understand CLS from a more combinatorial perspective, Hubáček and Yogevev [HY20] and Fearnley, Gordon, Mehta, and Savani [FGMS20] introduced the class $EOPL \subseteq CLS$, whose complete problem is the namesake END-OF-POTENTIAL-LINE (EOPL) problem, defined below. (The paper [HY20] initially defined a more restricted “metered” version of this problem, but we use the formulation from [FGMS20], which they prove is equivalent to the one from [HY20].) It is also natural to define a *sink-only* version of EOPL as suggested by [GKRS18].

EOPL: END-OF-POTENTIAL-LINE (EOPL). We are given access to a directed graph $G = (V, E)$ that is acyclic, has in/out-degree at most 1, and has $|V| = 2^n$ nodes; that is, G is a disjoint union of directed paths. The graph is described by a $\text{poly}(n)$ -sized circuit: for any node we can compute its successor and predecessor, if any, and also an integer potential, which is guaranteed to increase along the directed edges. In addition, we are given the name of a distinguished source v^* . The goal is to find any source or sink other than v^* .

SOPL: SINK-OF-POTENTIAL-LINE (SOPL). Same as EOPL except the goal is to find a sink.

It is comforting to know that the definition of EOPL is robust: Ishizuka [Ish21] showed that a version of EOPL that guarantees $\text{poly}(n)$ many distinguished sources is still equivalent (via polynomial-time reductions) to the above standard version with a single source.

Fearnley et al. [FGMS20] also defined a more restricted subclass $UEOPL \subseteq EOPL$ where the complete problem is UNIQUE-EOPL, a version of EOPL with a *unique* directed path. They showed that this class contains many important search problems with unique witnesses, such as unique sink orientations, linear complementary problems, ARRIVAL [DGK⁺17, GHH⁺18]. Other problems known to lie in UEOPL are a restricted version of the Ham-Sandwich problem [CCM20] and a pizza cutting problem [Sch21]. Fearnley et al. [FGMS20] conjecture that $UEOPL \neq EOPL$.

A surprising collapse. In a breakthrough, Fearnley, Goldberg, Hollender, and Savani [FGHS21] showed that, despite appearances to the contrary, $\text{CLS} = \text{PLS} \cap \text{PPAD}$. This goes against the conjecture of Daskalakis and Papadimitriou [DP11] that the classes are distinct, a belief which underlied much of their original motivation for introducing CLS. The nontrivial direction of the collapse is a reduction from a canonical complete problem $\text{SoD} \wedge \text{EoL} \in \text{PLS} \cap \text{PPAD}$ (defined below) to a problem $\text{KKT} \in \text{CLS}$, which involves computing a Karush–Kuhn–Tucker point of a smooth function. We may summarise the main technical result of Fearnley et al. [FGHS21] as

$$\text{SoD} \wedge \text{EoL} \leq \text{KKT} \qquad \text{which implies } \text{PLS} \cap \text{PPAD} \subseteq \text{CLS}. \quad (1)$$

Here we use \leq to denote a polynomial-time reduction between search problems. The operator \wedge produces the *meet* of two search problems: the input to problem $A \wedge B$ is a pair (x, y) where x is an instance of A and y is an instance of B and the goal is to output either a solution to x or to y . Then $\text{SoD} \wedge \text{EoL}$ is the canonical (albeit “unnatural”) complete problem for $\text{PLS} \cap \text{PPAD}$ [DP11].

Our new collapses. Our main results, [Theorems 1](#) and [2](#), follow from two new reductions, the first one of which strengthens the reduction (1) from [FGHS21]:

$$\text{SoD} \wedge \text{EoL} \leq \text{EoPL} \qquad \text{which implies } \text{PLS} \cap \text{PPAD} \subseteq \text{EOPL}, \quad (2)$$

$$\text{SoD} \wedge \text{SoL} \leq \text{SoPL} \qquad \text{which implies } \text{PLS} \cap \text{PPADS} \subseteq \text{SOPL}. \quad (3)$$

These reductions are between purely combinatorially defined search problems. In the case of (2), this bypasses the continuous middle-man of CLS and makes our reduction relatively simple to describe. In particular, we get a new simpler proof of the breakthrough collapse of [FGHS21] by combining (2) with the inclusion $\text{EOPL} \subseteq \text{CLS}$ proved by [HY20]. Furthermore, the new collapse implies that problems related to Tarski’s fixpoint theorem [EPRY20] and to a colourful version of Carathéodory’s theorem [MMSS17] lie in EOPL.

A further surprise? Given that the collapse $\text{CLS} = \text{PLS} \cap \text{PPAD}$ was considered extremely surprising by most people, how surprised should we be by the further collapse

$$\text{EOPL} = \text{CLS} = \text{PLS} \cap \text{PPAD} ?$$

Fearnley et al. [FGMS20] wrote regarding EOPL vs. CLS that “we actually think it could go either way.” In the wake of their breakthrough, the paper [FGHS21] explicitly conjectured $\text{EOPL} \neq \text{CLS}$.

For the authors of the present paper, the new collapse did come as an utter shock. When we began work on this project, our intuitions convinced us that, again, $\text{EOPL} \neq \text{CLS}$, a conjecture which had just found its way to the second author’s PhD thesis [Hol21, Chapter 12]. In our convictions, we set out to prove this separation in the *black-box model* where, instead of circuits, the directed graphs are described by black-box oracles. We tried in vain for nine months. The upshot is that [Theorems 1](#) and [2](#) now crush this possibility, as they hold even in the black-box model.

2 A Unified View: The GRID Problem

In this section we formally define all the problems of interest. We take the unusual approach of defining a single problem (which we call the GRID problem) with various parameters which can be tweaked to obtain all of the problems we study in this paper. This mainly serves two purposes. First of all, it is particularly convenient for presenting our reductions, since it allows us to combine instances from different problems more easily. The second reason is that we believe that this unified view of seemingly very different problems is of independent interest.

The GRID problem. For $n \in \mathbb{N}$, let $[n] := \{1, 2, \dots, n\}$. We define a general problem on a grid $[N] \times [M]$, where N and M should be thought of as being (potentially) exponentially large. The problem involves A paths starting from column 1 ($[N] \times \{1\}$) and moving from column i ($[N] \times \{i\}$) to column $i + 1$ ($[N] \times \{i + 1\}$). On the last column ($[N] \times \{M\}$) there are at most B valid ends of paths. If paths are not allowed to merge, then by the Pigeonhole Principle $A > B$ ensures the existence of a solution, i.e., a path that does not end at a valid position on the last column. If paths are allowed to merge, then a solution is guaranteed to exist as long as $B = 0$. To make things more precise, the paths start from nodes 1 to A in the first column (i.e., $[A] \times \{1\}$), and the valid termination points are nodes 1 to B in the last column (i.e., $[B] \times \{M\}$).

In more detail, we are given a boolean circuit $S: [N] \times [M] \rightarrow [N] \cup \{\text{null}\}$, the *successor circuit*, which allows us to efficiently compute the outgoing edge at a node. If $S(x, y) = \text{null}$, then (x, y) does not have an outgoing edge. Otherwise, there is an outgoing edge from (x, y) to $(S(x, y), y + 1)$. The problem also has two parameters which are used to tweak the definition: r (*reversible*) and b (*bijective*). Intuitively, when $r = 1$, we are also given access to a *predecessor circuit* $P: [N] \times [M] \rightarrow [N] \cup \{\text{null}\}$, which, analogously to S , allows us to efficiently compute the incoming edge at a node. In particular, when $r = 1$, two paths cannot merge. If, on top of that, $b = 1$, then we do not allow any new paths apart from the original A paths, and we also require that all B valid ends of paths are actually reached by a path. The combination $r = 0, b = 1$ is not allowed. The formal definition is as follows.

Definition 1. In the GRID problem, given N, M, A, B with $N \geq A > B \geq 0$ and $M \geq 2$, boolean circuits $S, P: [N] \times [M] \rightarrow [N] \cup \{\text{null}\}$, and bits $r, b \in \{0, 1\}$, output any of the following:

1. $x \in [A]$ such that $S(x, 1) = \text{null}$, *(missing pigeon/source)*
2. $x \in [N]$ such that $S(x, M - 1) > B$, *(invalid hole/sink)*
3. $x \in [N]$ and $y \in [M - 2]$ such that
 $S(x, y) \neq \text{null}$ and $S(S(x, y), y + 1) = \text{null}$, *(pigeon interception/sink)*
4. If $r = 1$ and $b = 1$:
 - (a) $(x, y) \in ([N] \times [M - 1]) \setminus ([A] \times \{1\})$ such that
 $S(x, y) \neq \text{null}$ and $P(x, y) = \text{null}$, or *(pigeon genesis/source)*
 - (b) $x \in [B]$ such that $P(x, M) = \text{null}$. *(empty hole/sink)*

We also enforce the following two conditions syntactically:

- If $r = 0$, then $b = 0$ and $B = 0$.
- If $r = 1$, then the successor and predecessor circuits are consistent, which can be enforced as follows. The circuit S is replaced by the circuit \overline{S} , which on input (x, y) computes $x' := S(x, y)$ and outputs x' , unless $(x', y + 1) \notin [N] \times [M]$ or $P(x', y + 1) \neq x$, in which case it outputs null. Similarly, the circuit P is replaced by the circuit \overline{P} , which on input (x, y) computes $x' := P(x, y)$ and outputs x' , unless $(x', y - 1) \notin [N] \times [M]$ or $S(x', y - 1) \neq x$, in which case it outputs null.




Figure 2: Examples of GRID problems. Square nodes are valid starts of paths (top-most A nodes in the first column) and diamonds are valid ends of paths (top-most B nodes in the last column). Solutions are drawn in red. However, for visual clarity we highlight the actual sinks rather than the *sink predecessors* as in Definition 1. Nodes with a null successor are drawn without an outgoing pointer. (2a) has parameters $(r = 0, b = 0, A = 1, B = 0)$ and defines an SoD instance. Only the successor circuit is drawn, as the predecessor circuit is not used by SoD. In particular, directed paths can *merge*, such as for node $(2, 3)$. (2b) has parameters $(r = 1, b = 0, A = N, B = N - 1)$ and defines an INJ-PHP instance. The diamond with a green circle would be a solution of BIJ-PHP (with $b = 1$) but is not a solution of INJ-PHP. (2c) has parameters $(r = 1, b = 0, A = 1, B = 0)$ and defines an SOPL instance. Sources with green circles would be solutions of an EOPL instance (with $b = 1$).

Canonical complete problems as special cases of GRID. By imposing various restrictions on the GRID problem, we obtain the following canonical complete problems; see Figure 2. (Here INJ-PHP/BIJ-PHP stand for Injective/Bijective Pigeonhole Principle.)

- SoD: $r = 0, b = 0, A = 1, B = 0$. (PLS-complete)
- SOPL: $r = 1, b = 0, A = 1, B = 0$. (SOPL-complete)
- EOPL: $r = 1, b = 1, A = 1, B = 0$. (EOPL-complete)
- INJ-PHP: $r = 1, b = 0, M = 2, N = A = B + 1$. (PPADS-complete)
- BIJ-PHP: $r = 1, b = 1, M = 2, N = A = B + 1$. (PPAD-complete)

Remark 1. Here we have slightly abused notation by calling these problems SoD, SOPL and EOPL even though their original definitions (in [JPY88, GKRS18, FGMS20], respectively) do not use a grid structure, and instead come with an additional circuit computing the potential of any node. It is not too hard to see that these grid-versions of the problems are indeed polynomial-time equivalent to the original versions. The main idea is that the grid implicitly provides a potential value for every node (x, y) , namely its column number y . Thus, given such a problem on a grid, it is easy to define a potential circuit by simply assigning the potential value y to any node (x, y) of the grid.

The other direction is slightly more involved. Consider an instance of one of the original problems with vertex set $V = [N]$ and potential values lying in $P = [M]$. Without loss of generality, we can assume that along any edge the potential increases by exactly one. Indeed, this was proved explicitly by [FGMS20] when they reduced EOPL to EOML, and the same idea applies to SoD and SOPL as well. The reduction to the grid-version of the problem is then obtained by identifying a vertex $x \in V$ that has potential $p \in P$ with the node (x, p) on the $[N] \times [M]$ grid.

The following is essentially folklore (see, e.g., [BM04]), so we only provide a brief proof sketch.

Lemma 1. *INJ-PHP and BIJ-PHP are respectively PPADS- and PPAD-complete.*

Proof Sketch. To see that BIJ-PHP lies in PPAD, we can reduce to EOL (see, e.g., [DGP09] for a formal definition) with vertex set $V = [N] \times [2]$ as follows: add a directed edge from node $(x, 2)$ to node $(x, 1)$ for all $x \leq A - 1$. By taking (x, A) as the distinguished source node, this yields an EOL instance with the same solutions as the original BIJ-PHP instance. On the other hand, given an instance of EOL with vertex set $V = [N]$ and distinguished source node N (without loss of generality), we construct an instance of BIJ-PHP on $[N] \times [2]$ as follows: for any isolated vertex $x \in [N]$, create an edge from $(x, 1)$ to $(x, 2)$; for any edge from x to y , create an edge from $(x, 1)$ to $(y, 2)$. This simple reduction proves the PPAD-hardness of BIJ-PHP. The exact same constructions can be used to prove that INJ-PHP is PPADS-complete, by reducing to and from the SOL problem (formally defined by Beame et al. [BCE⁺98], who call it SINK). \square

In Section 6, we briefly explain how an extended version of the GRID problem can be used to also capture PPP, the class defined by Papadimitriou [Pap94] to capture a version of the Pigeonhole Principle where edges can only be computed efficiently in the forward direction.

3 Path-Pigeonhole Problems

In this section we use the GRID problem to define some interesting extensions of the two pigeonhole problems. Namely, we consider the case where, instead of just two columns, there are many columns. In a certain sense, this corresponds to allowing the pigeons to travel for a long time before reaching a hole. In particular, we can no longer efficiently tell in which hole a given pigeon will land. This allows us to show that the problems remain hard even when there are significantly more pigeons than holes. This fact, stated in Lemma 2 below, will be crucial to obtain our main result later.

Let $f: \mathbb{N} \rightarrow \mathbb{N}$ be a polynomial-time computable function with $f(t) > t$. In this section, we consider the following restrictions of GRID:

- PATH-INJ-PHP _{f} : $r = 1, b = 0, A = f(B)$.
- PATH-BIJ-PHP _{f} : $r = 1, b = 1, A = f(B)$.

The following lemma is an important ingredient for the proof of our main result.

Lemma 2. *Let $f(t) > t$ be polynomial-time computable. There exists a reduction INJ-PHP \leq PATH-INJ-PHP _{f} that maps an instance with parameters $(A, B) = (T + 1, T)$ to an instance with parameters $(A, B) = (f(T), T)$ and $(N, M) = (f(T), f(T) - T + 1)$.*

Proof. The idea behind this reduction is very simple. Intuitively, we have the ability to merge $T + 1$ paths into T paths, such that finding a “mistake”, i.e., a path that stops, requires solving the INJ-PHP instance. In particular, if we start with some N paths, where $N \geq T + 1$, then we can merge those paths into $N - 1$ paths by merging the first $T + 1$ paths into T paths, and leaving the remaining $N - (T + 1)$ paths unchanged. Applying this idea repeatedly, we can merge $f(T)$ paths into just T paths in $f(T) - T$ steps. This results in an instance of PATH-INJ-PHP _{f} with $f(T) - T + 1$ columns, where every solution yields a solution to the INJ-PHP instance; see Figure 3. More formally, let (S, P) denote an instance of INJ-PHP with parameters $A = T + 1$ and $B = T$. Without loss of generality, we can assume that no pigeon goes to the invalid hole, i.e., $S(x, 1) \neq T + 1$




Figure 3: The reduction $\text{INJ-PHP} \leq \text{PATH-INJ-PHP}_f$ in Lemma 2. We build a PATH-INJ-PHP_f instance by chaining several identical INJ-PHP instances side-by-side. Note that a solution to the PATH-INJ-PHP_f instance can be directly mapped to one for the original INJ-PHP instance.

for all $x \in [T + 1]$. We construct an instance $(\widehat{S}, \widehat{P})$ of PATH-INJ-PHP_f on the $[N] \times [M]$ grid, where $N = f(T)$, $M = f(T) - T + 1$, $A = f(T)$ and $B = T$. The successor circuit \widehat{S} is defined as follows:

$$\widehat{S}(x, y) := \begin{cases} S(x, 1) & \text{if } x \in [T + 1] \text{ and } y \in [M - 1], \\ x - 1 & \text{if } T + 2 \leq x \leq f(T) - y + 1 \text{ and } y \in [M - 1], \\ \text{null} & \text{otherwise,} \end{cases}$$

and the predecessor circuit \widehat{P} is then defined accordingly to be consistent with \widehat{S} . Both circuits can be constructed in polynomial time, given S and P , and given that f can be computed in polynomial time. It is straightforward to check that any solution of the constructed instance yields a solution to the original INJ-PHP instance. \square

The same proof idea also yields that $\text{BIJ-PHP} \leq \text{PATH-BIJ-PHP}_f$.

4 SOPL = PLS \cap PPADS

In this section, we prove Theorem 2, namely $\text{SOPL} = \text{PLS} \cap \text{PPADS}$. To prove this we provide a reduction from a $\text{PLS} \cap \text{PPADS}$ -complete problem to the canonical SOPL -complete problem.

Lemma 3. $\text{SoD} \wedge \text{INJ-PHP} \leq \text{SoPL}$.

Proof Sketch. There are two obstacles to a direct reduction from SoD to SoPL : (i) we can only compute edges in the forward direction (i.e., we only have access to a successor circuit), and (ii) multiple edges can point to the same node.

To resolve the first issue, we modify the original $[N] \times [M]$ grid by taking N copies of each node. This ensures that there is a separate copy of each node v for each potential predecessor on the previous column. As a result, edges from different predecessor nodes will point to different copies of v . This means that predecessor nodes can now also be computed efficiently. However, the second issue remains: since we have made N copies of each node, there are also N copies of each predecessor node, and thus N edges coming from that predecessor node that somehow have to merge into the corresponding copy of v .




Figure 4: The reduction $\text{SoD} \wedge \text{PATH-INJ-PHP} \leq \text{SoPL}$ in the proof of Lemma 3. Given instances of SoD and PATH-INJ-PHP, we build an SoPL instance whose solutions can be traced back to solutions of $\text{SoD} \wedge \text{PATH-INJ-PHP}$. To overcome the issue of merging paths in SoD, the nodes of the SoD instance are replaced with a copy of a PATH-INJ-PHP gadget (in blue). Those gadgets are ultimately built out of the initial INJ-PHP instance (not shown) using Lemma 2.

To overcome the second obstacle, we make use of the following high-level idea: use the INJ-PHP instance to “hide” the fact that multiple edges can point to a single node. Unfortunately, we cannot use the INJ-PHP instance to hide the fact that N paths merge into a single node. But, if we take N^2 copies of each original node, instead of just N , then we have N^2 paths and N target nodes. By Lemma 2, the INJ-PHP instance can be turned into a PATH-INJ-PHP_f instance that hides the fact that N^2 paths merge into N paths. Thus, we replace each original node v of the SoD instance by a gadget that has N^2 nodes in the left-most column and N nodes in the right-most column, and such that finding the sink of a path inside the gadget requires solving the INJ-PHP instance. Importantly, we only construct the paths inside this gadget when the original node v has a successor in the SoD instance. This ensures that when v is an isolated node, the corresponding gadget does not contain any edges. Figure 4 illustrates the construction.

Note that although we might have added many new sources to the graph (which are irrelevant for SoPL), it remains the case that from any sink of the new graph, we can extract either a solution to SoD or to INJ-PHP.

In the final construction, edges can indeed be computed in both directions efficiently. Namely, given any node, we can determine in polynomial time if it has an incoming and/or outgoing edge, as well as the identity of the potential predecessor and successor nodes. Here, we crucially use the fact that edges can be computed efficiently in both directions in the INJ-PHP instance. \square

Proof. Let S be an instance of SoD on the grid $[N] \times [M]$. We are also given an instance of INJ-PHP and without loss of generality we can assume that it has parameters $(A, B) = (N + 1, N)$. By Lemma 2, we can reduce this INJ-PHP instance to a $\text{PATH-INJ-PHP}_{t^2}$ instance with parameters

$(A, B) = (N^2, N)$. Without loss of generality, we can assume that this $\text{PATH-INJ-PHP}_{t^2}$ instance (S', P') is defined on the grid $[N^2] \times [M]$ (by padding the SOD or $\text{PATH-INJ-PHP}_{t^2}$ instance with additional columns, if needed). This is not important for the reduction, but will be convenient.

We will take N^2 copies of each node in the original SOD instance, and make M copies of each column. As a result, our SOPL instance will be defined on the $[N^3] \times [M^2]$ grid. It will be convenient to use some special notation to refer to points in this grid. For $\alpha \in [N^2]$ and $x \in [N]$, we use the notation (α, x) to denote the row $\alpha + (x - 1) \cdot N^2 \in [N^3]$. This corresponds to indexing the α th copy of row x of the original instance. We also introduce some additional notation to index these $[N^2]$ copies: for $i, j \in [N]$, we let $[i, j] := i + (j - 1) \cdot N \in [N^2]$. Thus, $([i, j], x)$ denotes the $[i, j]$ th copy of row x . The “[i, j]” notation essentially subdivides $[N^2]$ into N blocks containing N values each, which will be useful for routing incoming edges to the correct copy of a node. We introduce similar notation to denote the M copies of a column $y \in [M]$. Putting everything together, the notation $(\alpha, x; k, y) \in [N^2] \times [N] \times [M] \times [M]$ denotes the node $(\alpha + (x - 1) \cdot N^2, k + (y - 1) \cdot M) \in [N^3] \times [M^2]$. In particular, the notation $([i, j], x; k, y)$ is well-defined.

The circuits \widehat{S}, \widehat{P} of the SOPL instance on $[N^3] \times [M^2]$ are defined as follows:

$$\widehat{S}([i, j], x; k, y) := \begin{cases} ([i, x], S(x, y)) & \text{if } k = M \text{ and } j = 1, \\ (S'([i, j], k), x) & \text{if } k < M \text{ and } S(x, y) \neq \text{null}, \\ \text{null} & \text{otherwise} \end{cases}$$

$$\widehat{P}([i, j], x; k, y) := \begin{cases} ([i, 1], j) & \text{if } k = 1 \text{ and } y > 1 \text{ and } S(j, y - 1) = x, \\ (P'([i, j], k), x) & \text{if } k > 1 \text{ and } S(x, y) \neq \text{null}, \\ \text{null} & \text{otherwise} \end{cases}$$

where $(\alpha, z) \in [N^2] \times [N]$ is interpreted as an element in $[N^3]$ as above, and where we use the convention $(*, \text{null}) = (\text{null}, *) = \text{null}$. Using the fact that S' and P' are consistent, it can be checked that \widehat{S} and \widehat{P} are also consistent.

In order to argue about the correctness of the reduction, consider any sink $([i, j], x; k, y)$ of the SOPL instance. If $2 \leq k \leq M - 1$, then it must be that $([i, j], k)$ is a sink of the $\text{PATH-INJ-PHP}_{t^2}$ instance (S', P') . If $k = M$ and $j = 1$, then $([i, j], x; k, y)$ cannot be a sink, since $\widehat{P}([i, j], x; k, y) \neq \text{null}$ implies that $S(x, y) \neq \text{null}$, and thus $\widehat{S}([i, j], x; k, y) \neq \text{null}$. If $k = M$ and $j > 1$, then $([i, j], k)$ is an invalid sink on the last column of the $\text{PATH-INJ-PHP}_{t^2}$ instance, and so in particular a solution. If $k = 1$ and $S(x, y) \neq \text{null}$, then $([i, j], k)$ is a missing source on the first column of the $\text{PATH-INJ-PHP}_{t^2}$ instance, and so again a solution. Finally, if $k = 1$ and $S(x, y) = \text{null}$, then it must be that $S(j, y - 1) = x$ and thus (x, y) is a sink of the original SOD instance, and this is witnessed by the node $(j, y - 1)$. \square

5 EOPL = PLS \cap PPAD

In this section, we prove [Theorem 1](#), namely $\text{EOPL} = \text{PLS} \cap \text{PPAD}$. The equality $\text{SOPL} = \text{PLS} \cap \text{PPADS}$ ([Theorem 2](#)) proved in the previous section, together with the fact that $\text{PPAD} \subseteq \text{PPADS}$, immediately imply that

$$\text{SOPL} \cap \text{PPAD} = \text{PLS} \cap \text{PPAD}.$$

As a result, in order to prove [Theorem 1](#), it suffices to give a reduction from an $\text{SOPL} \cap \text{PPAD}$ -complete problem to an EOPL-complete problem:

Lemma 4. $\text{SOPL} \wedge \text{BIJ-PHP} \leq \text{EOPL}$.




Figure 5: A naive attempt at a reduction $\text{SOPL} \leq \text{EOPL}$. Although most solutions arise from the sinks of the original SOPL instance, a spurious solution is introduced at \bar{v}_0 , which does not correspond to any sink of the original instance.

Proof Sketch. A very natural attempt at a reduction from SOPL to EOPL is to try to remove all undistinguished sources, i.e., all sources except the trivial one. Then, clearly, any EOPL-solution would have to be a sink, and thus also a solution to SOPL.

There is a simple trick that *almost* achieves this. First, make a *reversed* copy of the SOPL instance, i.e., reverse the direction of all edges, and the ordering of the potential. Note that sources of the original instance have now become sinks in the reversed copy, and vice versa. Then, for each source node v of the original graph, add an edge pointing from its copy \bar{v} (which is a sink) to v .

The only problem with this reduction is that we have eliminated *all* sources of the original graph, including the distinguished one. In particular, the distinguished source v_0 of the original instance is no longer a source, since there is an edge from its copy \bar{v}_0 to v_0 . As a result, the reduction fails, because the instance of EOPL we have constructed does not have a distinguished source. Furthermore, we cannot hope to turn one of the new sources into a distinguished source, since any such source yields a solution to the original instance (where it is a sink).

In order to address this issue, we add a new node u and select it as our new distinguished source. Clearly, u is a solution of the instance, since it is a distinguished source that is not actually a source, but just an isolated node. Now, imagine that we remove the edge (\bar{v}_0, v_0) and instead introduce an edge (u, v_0) ; see Figure 5. Then, u is no longer a solution, but \bar{v}_0 becomes a sink, and thus a solution, instead. In other words, the reduction can pick whether it wants u or \bar{v}_0 to be a solution by changing this edge. Of course, in both cases, the resulting instance is very easy to solve, but this minor observation already provides the idea for the next step.

Take k copies of the instance we have constructed (before adding u). There are now k copies $v_0^{(1)}, \dots, v_0^{(k)}$ of the original distinguished source, and k copies of the reverse copy $\bar{v}_0^{(1)}, \dots, \bar{v}_0^{(k)}$. Remove the edges $(\bar{v}_0^{(i)}, v_0^{(i)})$ for $i = 1, \dots, k$. If we now introduce the new distinguished source u , we have $k + 1$ nodes that “need” an outgoing edge in order to not be solutions (namely, $u, \bar{v}_0^{(1)}, \dots, \bar{v}_0^{(k)}$) and k nodes that “need” an incoming edge (namely, $v_0^{(1)}, \dots, v_0^{(k)}$). Clearly, no matter how we introduce edges here, one of $u, \bar{v}_0^{(1)}, \dots, \bar{v}_0^{(k)}$ will not have an outgoing edge and will be a solution. However, we can use a BIJ-PHP instance to make it hard to find such a solution. Let K denote the parameter of the BIJ-PHP instance, i.e., $K + 1$ points are mapped to K points. Then, we let $k := K$ and add edges between $u, \bar{v}_0^{(1)}, \dots, \bar{v}_0^{(k)}$ and $v_0^{(1)}, \dots, v_0^{(k)}$ according to the BIJ-PHP instance. An example of the construction is depicted in Figure 6.

Now, it is easy to check that any undistinguished source or any sink of the resulting graph must yield a solution to the BIJ-PHP instance or a solution of the SOPL instance. In particular, if u is not a source, then this yields a solution to BIJ-PHP. \square

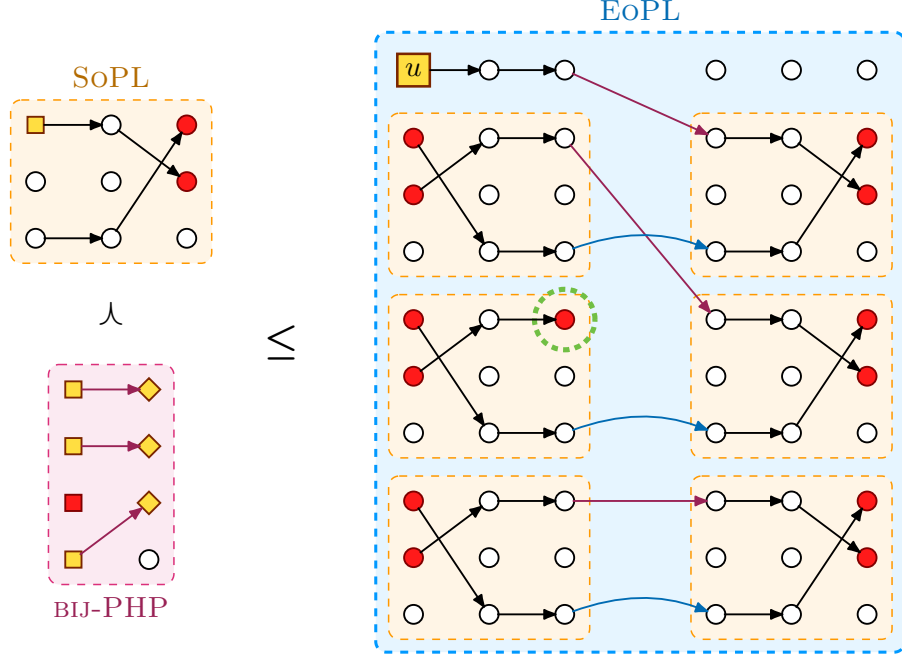


Figure 6: The reduction $\text{SOPL} \wedge \text{BIJ-PHP} \leq \text{EoPL}$ in Lemma 4. The BIJ-PHP instance (in pink) connects the newly introduced source u together with the distinguished sources and sinks of the copied SOPL instances. Non-distinguished sources and sinks are connected with blue edges. The node circled in green corresponds to a solution of the BIJ-PHP instance.

Proof. Let (S, P) be an instance of SOPL on the grid $[N] \times [M]$. Without loss of generality, we can assume that all sources occur on the first column, i.e., for any source $(x, y) \in [N] \times [M]$ it holds that $y = 1$. Indeed, by appropriately increasing N , for each source (x, y) we can add a path that starts on the first column and ends at (x, y) , thus effectively “transferring” the source to the first column. Let (S', P') be an instance of BIJ-PHP on the grid $[K + 1] \times [2]$ that maps $K + 1$ pigeons to K holes.

We take K copies of the SOPL instance and K copies of the reversed SOPL instance, all together in a single grid. This grid will be of the form $[KN] \times [2M]$. For clarity, we will use the notation $(i, x; y) \in [K] \times [N] \times [2M]$ to denote the element $(x + (i - 1) \cdot N, y) \in [KN] \times [2M]$. The i th copy of the instance will be embedded in $\{i\} \times [N] \times ([2M] \setminus [M])$, while the i th reversed copy will be in $\{i\} \times [N] \times [M]$. Formally, we define new successor and predecessor circuits \widehat{S}, \widehat{P} on $[KN] \times [2M]$ as follows:

$$\widehat{S}(i, x; y) := \begin{cases} (i, P(x, M - y + 1)) & \text{if } y \leq M, \\ (i, S(x, y - M)) & \text{if } y \geq M + 1 \end{cases}$$

$$\widehat{P}(i, x; y) := \begin{cases} (i, S(x, M - y + 1)) & \text{if } y \leq M, \\ (i, P(x, y - M)) & \text{if } y \geq M + 1 \end{cases}$$

where $(i, z) \in [K] \times [N]$ is interpreted as an element in $[KN]$, and where we use the convention $(i, \text{null}) = \text{null}$.

Since S and P are consistent, \widehat{S} and \widehat{P} are also consistent. Note that there are currently no edges between column M and column $M + 1$. In the second step of the reduction we add edges between these two columns as follows. For every $i \in [K]$ and $x \in [N] \setminus \{1\}$, if $(i, x; M + 1)$ is a

source, then we add an edge from $(i, x; M)$ to $(i, x; M + 1)$. Note that in that case $(i, x; M + 1)$ was a sink. The case where $x = 1$ is handled separately, because it corresponds to nodes that are copies of the distinguished source of the original SOPL instance. For any $i \in [K]$ and for $x = 1$, if $S'(i, 1) = j \neq \text{null}$, we add an edge from $(i, 1; M)$ to $(j, 1; M + 1)$. Note that here we also use P' (which is assumed to be consistent with S') to implement this edge in $(\widehat{S}, \widehat{P})$.

Finally, we introduce a new special node u on column M which will act as our new distinguished source. If $S'(K + 1, 1) = j \neq \text{null}$, then we add an edge from u to $(j, 1; M + 1)$. By extending the grid to be $[KN + 1] \times [2M]$, by renaming nodes and by “transferring” the source u to the first column as before, we can ensure that the distinguished source is $(1, 1)$.

It is easy to check that the new circuits \widehat{S}, \widehat{P} can be constructed in polynomial time. For the correctness of the reduction, note that any source or sink that occurs on columns $[2M] \setminus \{M, M + 1\}$ must correspond to a sink of the original SOPL instance. On the other hand, any source or sink that occurs on column M or $M + 1$ must correspond to a solution of the BIJ-PHP instance (namely, a pigeon without a hole, or a hole without a pigeon). This completes the reduction. \square

6 Discussion

As mentioned in the introduction, it remains open whether $\text{UEOPL} \stackrel{?}{=} \text{EOPL}$. Separating the two classes in the black-box model would be an important first step towards pinning down the complexity of the various natural problems contained in UEOPL , since it would provide strong evidence that these problems are unlikely to be complete for $\text{PLS} \cap \text{PPAD}$.

The techniques developed in this paper do not seem to yield any other major class collapse. Indeed, our reductions are all black-box, and the main classes are known to be distinct in that model [BCE⁺98, Mor01, BM04]. A notable exception is the question of whether PLS is a subset of PPP , or even of PPADS . This remains open even in the black-box model.

In the remainder of this section we briefly present some observations about the path pigeonhole problems, as well as a further consequence of our reduction techniques: a version of SOD where paths are not allowed to merge turns out to be $\text{PLS} \cap \text{PPP}$ -complete.

Path-Pigeonhole problems. Lemma 2 in particular establishes that PATH-INJ-PHP_f is PPADS -hard. Membership in PPADS can be shown by reducing to INJ-PHP using a construction similar to the reduction from EOL to BIJ-PHP in the proof of Lemma 1.

The statement of Lemma 2 also holds for $\text{BIJ-PHP} \leq \text{PATH-BIJ-PHP}_f$, and the proof is essentially the same. This shows that PATH-BIJ-PHP_f is PPAD -hard. However, it is unclear whether PATH-BIJ-PHP_f lies in PPAD . Indeed, using the same idea as for $\text{PATH-INJ-PHP}_f \leq \text{INJ-PHP}$ yields an instance with $A \gg B$, and we cannot increase B artificially here (whereas this is possible in INJ-PHP). Another way to state this is to say that we can reduce PATH-BIJ-PHP_f to an instance of EOL that has many distinguished source nodes, instead of just one. It is known that EOL with a polynomial number of distinguished sources remains PPAD -complete [GH21], but in general we will obtain an exponential number of such sources here.

Extending the GRID problem to capture PPP. The canonical PPP -complete problem is PIGEON-CIRCUIT [Pap94]: given a circuit mapping N pigeons to $N - 1$ holes, find a *collision*, i.e., two pigeons that are mapped to the same hole. Importantly, unlike in INJ-PHP or BIJ-PHP , we are not given a circuit to compute the mapping in the other direction, i.e., from holes to pigeons. In order to capture this problem, we extend the definition of GRID by introducing an additional parameter bit $c \in \{0, 1\}$, which stands for *collision*. We also introduce a new solution type:

5. If $r = 0$ and $c = 1$: $x_1, x_2 \in [N]$ and $y \in [M - 1]$ such that $x_1 \neq x_2$ and $S(x_1, y) = S(x_2, y) \neq \text{null}$, *(pigeon collision/merging)*

Furthermore, the syntactic condition “If $r = 0$, then $b = 0$ and $B = 0$ ” is replaced by the condition:

- If $r = 0$, then $b = 0$. If $r = 0$ and $c = 0$, then $B = 0$.

The PPP-complete problem PIGEON-CIRCUIT is then obtained by setting $r = 0, c = 1, M = 2, N = A = B + 1$. In fact, GRID remains in PPP even if we just set $r = 0, c = 1$ and leave the other parameters unfixed. This can be shown by using a construction similar to the reduction from EOL to BIJ-PHP in the proof of [Lemma 1](#).

SoD without merging. What is the complexity of SoD if paths are not allowed to merge? In other words, what is the complexity of the GRID problem with parameters $r = 0, c = 1, A = 1, B = 0$? Clearly, this restricted version still lies in PLS, and by the previous paragraph it also lies in PPP. Using the ideas developed in this paper, it can be shown that the problem is in fact $\text{PLS} \cap \text{PPP}$ -complete. To see this, note that using the simple construction in the proof of [Lemma 2](#) we can reduce PIGEON-CIRCUIT to a path-version of the problem where $f(T)$ pigeons are mapped to T holes. Then, the construction in the proof of [Lemma 3](#) can be used to reduce $\text{SoD} \wedge \text{PIGEON-CIRCUIT}$ to SoD without merging.

Acknowledgements

We thank Aviad Rubinfeld for his many questions during e-mail correspondence.

References

- [BCE⁺98] Paul Beame, Stephen Cook, Jeff Edmonds, Russell Impagliazzo, and Toniann Pitassi. The relative complexity of NP search problems. *Journal of Computer and System Sciences*, 57(1):3–19, 1998. doi:10.1006/jcss.1998.1575.
- [BM04] Joshua Buresh-Oppenheim and Tsuyoshi Morioka. Relativized NP search problems and propositional proof systems. In *Proceedings of the 19th IEEE Conference on Computational Complexity (CCC)*, pages 54–67, 2004. doi:10.1109/CCC.2004.1313795.
- [CCM20] Man-Kwun Chiu, Aruni Choudhary, and Wolfgang Mulzer. Computational Complexity of the α -Ham-Sandwich Problem. In *47th International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 31:1–31:18, 2020. doi:10.4230/LIPIcs.ICALP.2020.31.
- [DGK⁺17] Jérôme Dohrau, Bernd Gärtner, Manuel Kohler, Jiří Matoušek, and Emo Welzl. ARRIVAL: A zero-player graph game in $\text{NP} \cap \text{coNP}$. In *A Journey Through Discrete Mathematics*, pages 367–374. Springer, 2017. doi:10.1007/978-3-319-44479-6_14.
- [DGP09] Constantinos Daskalakis, Paul Goldberg, and Christos Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009. doi:10.1137/070699652.
- [DP11] Constantinos Daskalakis and Christos Papadimitriou. Continuous local search. In *Proceedings of the 22nd Symposium on Discrete Algorithms (SODA)*. Society for Industrial and Applied Mathematics, January 2011. doi:10.1137/1.9781611973082.62.
- [EPRY20] Kousha Etessami, Christos Papadimitriou, Aviad Rubinfeld, and Mihalis Yannakakis. Tarski’s theorem, supermodular games, and the complexity of equilibria. In *Proceedings of the 11th Innovations in Theoretical Computer Science Conference (ITCS)*, volume 151, pages 18:1–18:19, 2020. doi:10.4230/LIPIcs.ITCS.2020.18.

- [FGHS21] John Fearnley, Paul Goldberg, Alexandros Hollender, and Rahul Savani. The complexity of gradient descent: $\text{CLS} = \text{PPAD} \cap \text{PLS}$. In *Proceedings of the 53rd Symposium on Theory of Computing (STOC)*, 2021. doi:10.1145/3406325.3451052.
- [FGMS20] John Fearnley, Spencer Gordon, Ruta Mehta, and Rahul Savani. Unique end of potential line. *Journal of Computer and System Sciences*, 114:1–35, 2020. doi:10.1016/j.jcss.2020.05.007.
- [GH21] Paul Goldberg and Alexandros Hollender. The Hairy Ball problem is PPAD-complete. *Journal of Computer and System Sciences*, 122:34–62, 2021. doi:10.1016/j.jcss.2021.05.004.
- [GHH⁺18] Bernd Gärtner, Thomas Dueholm Hansen, Pavel Hubáček, Karel Král, Hagar Mosaad, and Veronika Slívová. ARRIVAL: Next stop in CLS. In *45th International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 107, pages 60:1–60:13. Schloss Dagstuhl, 2018. doi:10.4230/LIPICs.ICALP.2018.60.
- [GKRS18] Mika Göös, Pritish Kamath, Robert Robere, and Dmitry Sokolov. Adventures in monotone complexity and TFNP. In *Proceedings of the 10th Innovations in Theoretical Computer Science Conference (ITCS)*, volume 124, pages 38:1–38:19, 2018. doi:10.4230/LIPICs.ITCS.2019.38.
- [Hol21] Alexandros Hollender. *Structural results for total search complexity classes with applications to game theory and optimisation*. PhD thesis, University of Oxford, 2021. URL: <https://ora.ox.ac.uk/objects/uuid:67e2d80b-76bf-4b49-9b7d-8bbd91633dd7>.
- [HY20] Pavel Hubáček and Eylon Yogev. Hardness of continuous local search: Query complexity and cryptographic lower bounds. *SIAM Journal on Computing*, 49(6):1128–1172, 2020. doi:10.1137/17m1118014.
- [Ish21] Takashi Ishizuka. The complexity of the parity argument with potential. *Journal of Computer and System Sciences*, 120:14–41, sep 2021. doi:10.1016/j.jcss.2021.03.004.
- [JPY88] David Johnson, Christos Papadimitriou, and Mihalis Yannakakis. How easy is local search? *Journal of Computer and System Sciences*, 37(1):79–100, August 1988. doi:10.1016/0022-0000(88)90046-3.
- [MMSS17] Frédéric Meunier, Wolfgang Mulzer, Pauline Sarrabezolles, and Yannik Stein. The rainbow at the end of the line—a PPAD formulation of the colorful Carathéodory theorem with applications. In *Proceedings of the 28th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1342–1351, 2017. doi:10.1137/1.9781611974782.87.
- [Mor01] Tsuyoshi Morioka. Classification of search problems and their definability in bounded arithmetic. Master’s thesis, University of Toronto, 2001. URL: <https://www.collectionscanada.ca/obj/s4/f2/dsk3/ftp04/MQ58775.pdf>.
- [MP91] Nimrod Megiddo and Christos Papadimitriou. On total functions, existence theorems and computational complexity. *Theoretical Computer Science*, 81(2):317–324, April 1991. doi:10.1016/0304-3975(91)90200-L.
- [Pap94] Christos Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *Journal of Computer and System Sciences*, 48(3):498–532, June 1994. doi:10.1016/S0022-0000(05)80063-7.
- [Sch21] Patrick Schneider. The Complexity of Sharing a Pizza. In *32nd International Symposium on Algorithms and Computation (ISAAC)*, pages 13:1–13:15, 2021. doi:10.4230/LIPICs.ISAAC.2021.13.