# On the Complexity of Bidirected Interleaved Dyck-Reachability 

YUANBO LI, Georgia Institute of Technology, USA<br>QIRUN ZHANG, Georgia Institute of Technology, USA<br>THOMAS REPS, University of Wisconsin-Madison, USA

Many program analyses need to reason about pairs of matching actions, such as call/return, lock/unlock, or set-field/get-field. The family of Dyck languages $\left\{D_{k}\right\}$, where $D_{k}$ has $k$ kinds of parenthesis pairs, can be used to model matching actions as balanced parentheses. Consequently, many program-analysis problems can be formulated as Dyck-reachability problems on edge-labeled digraphs. Interleaved Dyck-reachability (InterDyck-reachability), denoted by $D_{k} \odot D_{k}$-reachability, is a natural extension of Dyck-reachability that allows one to formulate program-analysis problems that involve multiple kinds of matching-action pairs. Unfortunately, the general InterDyck-reachability problem is undecidable.

In this paper, we study variants of InterDyck-reachability on bidirected graphs, where for each edge $\langle p, q\rangle$ labeled by an open parenthesis " $(a$ ", there is an edge $\langle q, p\rangle$ labeled by the corresponding close parenthesis " $) a$ ", and vice versa. Language-reachability on a bidirected graph has proven to be useful both (i) in its own right, as a way to formalize many program-analysis problems, such as pointer analysis, and (ii) as a relaxation method that uses a fast algorithm to over-approximate language-reachability on a directed graph. However, unlike its directed counterpart, the complexity of bidirected InterDyck-reachability still remains open.

We establish the first decidable variant (i.e., $D_{1} \odot D_{1}$-reachability) of bidirected InterDyck-reachability. In $D_{1} \odot D_{1}$-reachability, each of the two Dyck languages is restricted to have only a single kind of parenthesis pair. In particular, we show that the bidirected $D_{1} \odot D_{1}$-reachability problem is in PTIME. We also show that when one extends each Dyck language to involve $k$ different kinds of parentheses (i.e., $D_{k} \odot D_{k}$-reachability with $k \geq 2$ ), the problem is NP-hard (and therefore much harder).

We have implemented the polynomial-time algorithm for bidirected $D_{1} \odot D_{1}$-reachability. $D_{1} \odot D_{1}$ reachability provides a new over-approximation method for bidirected $D_{k} \odot D_{k}$-reachability in the sense that $D_{k} \odot D_{k}$-reachability can first be relaxed to bidirected $D_{1} \odot D_{1}$-reachability, and then the resulting bidirected $D_{1} \odot D_{1}$-reachability problem is solved precisely. We compare this $D_{1} \odot D_{1}$-reachability-based approach against another known over-approximating $D_{k} \odot D_{k}$-reachability algorithm. Surprisingly, we found that the over-approximation approach based on bidirected $D_{1} \odot D_{1}$-reachability computes more precise solutions, even though the $D_{1} \odot D_{1}$ formalism is inherently less expressive than the $D_{k} \odot D_{k}$ formalism.

CCS Concepts: • Mathematics of computing $\rightarrow$ Graph algorithms; • Theory of computation $\rightarrow$ Program analysis.

Additional Key Words and Phrases: Formal Language Graph Reachability, Interleaved-Dyck Language, Static Analysis, Complexity

Authors' addresses: Yuanbo Li, School of Computer Science, Georgia Institute of Technology, 266 Ferst Dr NW, Atlanta, GA, 30332, USA, yuanboli@gatech.edu; Qirun Zhang, School of Computer Science, Georgia Institute of Technology, 266 Ferst Dr NW, Atlanta, GA, 30332, USA, qrzhang@gatech.edu; Thomas Reps, Computer Science Department, University of Wisconsin-Madison, 1210 West Dayton St., Madison, WI, 53706, USA, reps@cs.wisc.edu.


This work is licensed under a Creative Commons Attribution 4.0 International License.
© 2021 Copyright held by the owner/author(s).
2475-1421/2021/1-ART59
https://doi.org/10.1145/3434340

## ACM Reference Format:

Yuanbo Li, Qirun Zhang, and Thomas Reps. 2021. On the Complexity of Bidirected Interleaved DyckReachability. Proc. ACM Program. Lang. 5, POPL, Article 59 (January 2021), 28 pages. https://doi.org/10. 1145/3434340

## 1 INTRODUCTION

Many static-analysis problems can be formulated as a reachability problem with respect to a formal language $L$ on an edge-labeled graph $G$. Two nodes in $G$ are $L$-reachable iff there exists a path between them, and the string spelled out by the path is a word in $L$. An $L$-reachability algorithm computes all $L$-reachable pairs in $G$. A Dyck language $D_{k}$, which consists of all strings of well-balanced parentheses over $k$ kinds of parentheses, is perhaps the most widely used formal language in $L$-reachability formulations [Reps 1998]. Many practical analyses use Dyck-reachability to express (ordered) matching actions, such as call/return (i.e., context-sensitivity) [Reps 2000; Yan et al. 2011], lock/unlock [Kahlon 2009; Ramalingam 2000], file-open/file-close [Späth et al. 2019], or set-field/get-field (i.e., field-sensitivity) [Arzt et al. 2014; Yan et al. 2011]. Kodumal and Aiken [2004] observed that "almost all of the applications of [L-reachability, where $L$ is a context-free language] in program analysis are based on Dyck languages."

Interleaved Dyck-reachability ("InterDyck-reachability" for short) is a more expressive formalism than Dyck-reachability. An InterDyck-reachability problem is an $L$-reachability problem in which $L$ involves two interleaved Dyck languages. Let $D_{p}$ and $D_{b}$ be the Dyck languages of balanced parentheses and balanced brackets, respectively. The InterDyck language $D_{p} \odot D_{b}$ is the language in which $D_{p}$ and $D_{b}$ are interleaved. For example, the string " $\left[{ }_{1}\left(1[2)_{1}(2]_{2}\right]_{1}\right)_{2}$ " in $D_{p} \odot D_{b}$ represents the interleaving of the string " $\left({ }_{1}\right)_{1}(2)_{2}$ " in $D_{p}$ and the string " $\left[{ }_{1}\left[{ }_{2}\right]_{2}\right]_{1}$ " in $D_{b}$. The InterDyck-reachability formalism can be applied to many client applications, such as alias analysis [Cheng and Hwu 2000; Sridharan and Bodík 2006], taint analysis [Arzt et al. 2014; Huang et al. 2015], and typestate analysis [Späth et al. 2019]. However, precise InterDyck-reachability is undecidable in general [Reps 2000]. Thus, practical analysis usually over-approximates the exact solution.

In this paper, we study variants of InterDyck-reachability restricted to bidirected graphs, where for each edge $\langle p, q\rangle$ labeled by a open parenthesis " $(a$ ", there is an edge $\langle q, p\rangle$ labeled by the corresponding close parenthesis ") $)_{a}$ ", and vice versa. Although $L$-reachability problems on bidirected graphs are a special case of general $L$-reachability problems, they play an important role in program analysis. For instance, graph bidirectedness allows pointer analysis to be formulated as a CFLreachability problem [Reps 1998, §4.4]. Many context- and field-sensitive alias analyses are based on bidirected InterDyck-reachability [Xu et al. 2009; Yan et al. 2011]. Moreover, bidirected Dyckreachability problems can be solved by a linear-time algorithm [Chatterjee et al. 2018], whereas the best algorithm for the directed counterpart runs in $O\left(n^{3} / \log n\right)$ time [Chaudhuri 2008]. Recent work by [Li et al. 2020] shows that graph simplification based on bidirected Dyck-reachability can be used to speed up any InterDyck-reachability algorithm. Bidirected InterDyck-reachability can also be used as a relaxation of directed InterDyck-reachability (i.e., by introducing all necessary labeled reversed edges into the underlying graph): the solution of the bidirected variant over-approximates the solution of the directed variant.

Unfortunately, despite tremendous progress on algorithms for solving bidirected problems, the question of whether bidirectedness makes any InterDyck-reachability variant decidable has remained open. This paper fills this gap by studying the computational complexity of two bidirected

InterDyck-reachability variants (i.e., $D_{1} \odot D_{1}$-reachability and $D_{k} \odot D_{k}$-reachability ${ }^{1}$ ). In particular, the paper makes two main contributions.

- We identify the first decidable variant of InterDyck-reachability, namely, bidirected $D_{1} \odot D_{1}$ reachability. We show that bidirected $D_{1} \odot D_{1}$-reachability is in PTIME.
- For the bidirected $D_{k} \odot D_{k}$-reachability problem with $k \geq 2$, we establish that it is NP-hard, and thus likely to be a much harder problem than bidirected $D_{1} \odot D_{1}$-reachability.

We implemented the bidirected $D_{1} \odot D_{1}$-reachability algorithm, and applied it to a real-world alias analysis [Xu et al. 2009; Yan et al. 2011]. Our empirical evaluation uncovered an interesting phenomenon. We compared the precision of (i) relaxing a context- and field-sensitive alias analysis to a context- and field-insensitive variant-by relaxing bidirected $D_{k} \odot D_{k}$-reachability to bidirected $D_{1} \odot D_{1}$-reachability-and applying the polynomial-time $D_{1} \odot D_{1}$-reachability algorithm against (ii) a known over-approximation algorithm for bidirected $D_{k} \odot D_{k}$-reachability [Zhang and Su 2017]. Our empirical results show that the exact solution from (i) is more precise than the overapproximating solution obtained from (ii), even though the $D_{1} \odot D_{1}$ formalism is inherently less expressive than the $D_{k} \odot D_{k}$ formalism.

Organization. Section 2 defines several concepts used in the rest of the paper. Section 3 presents the key idea behind our proof of the PTIME complexity of bidirected $D_{1} \odot D_{1}$-reachability. Section 4 establishes that $D_{1} \odot D_{1}$-reachability is in PTIME. Section 5 establishes the NP-hardness result for bidirected $D_{k} \odot D_{k}$-reachability. Section 6 presents experimental results. Section 7 discusses related work. Section 8 concludes.

## 2 PRELIMINARIES

### 2.1 Formal-Language-Reachability Problems

A large number of program-analysis problems can be formulated as formal-language-reachability problems. A formal-language-reachability problem in an edge-labeled graph extends ordinary graph reachability: for a node to be considered reachable (e.g., from a distinguished start node), it must be connected by a path along which the edge labels spell out a word in a given formal language $L$. Language $L$ is used to specify potential "flows of information" in the graph. More precisely, the restriction to $L$-paths provides a mechanism for excluding paths in the graph along which information could not flow.

Definition 2.1 (Formal-Language-Reachability Problem). Let $L$ be a formal language over an alphabet $\Sigma$, and $G$ be a graph in which each edge is labeled with a symbol from $\Sigma$. Each path corresponds a word formed by concatenating the symbols of the edges of the path (in the same order as the edges). A path is an L-path iff its corresponding word is in the language L. For each L-path from node $u$ to $v$, we say that $v$ is $L$-reachable from $u$ and the two nodes are an $L$-reachable pair. The all-pairs $L$-reachability problem is to compute all L-reachable pairs in $G$.

Example 2.2. Consider the context-free language $D_{1}$ defined in Figure 1a. Figure 1b gives an edge-labeled graph. In this graph, the path from $s$ to $t$ that traverses the cycle exactly once forms the word " $((()))$ " in language $D_{1}$. Consequently, $t$ is $D_{1}$-reachable from $s$.

[^0]\[

$$
\begin{aligned}
& S::=S S \\
& S::=(S) \\
& S::=\epsilon
\end{aligned}
$$
\]

(a) Dyck language $D_{1}$.

(b) Dyck-reachability.

Fig. 1. Dyck language $\left(D_{1}\right)$ and Dyck-reachability ( $D_{1}$-reachability).

### 2.2 Dyck Languages and Interleaved Dyck Languages

This section introduces the specific family of languages that we focus on in this paper. An interleaved Dyck language (InterDyck language) is constructed from multiple Dyck languages. A Dyck language generates a set of strings of balanced parentheses:

Definition 2.3 (Dyck Language). A Dyck language $D_{k}$ with $k$ different kinds of parenthesis pairs can be defined by the following grammar:

$$
S::=S S\left|\left({ }_{1} S\right)_{1}\right|\left({ }_{2} S\right)_{2}|\ldots|\left({ }_{k} S\right)_{k} \mid \epsilon .
$$

The language $D_{1}$ defined in Figure 1a is a special Dyck language with only one kind of parenthesis pair. In general, program-analysis applications can use the different kinds of parenthesis pairs in $D_{k}$ to track calls/returns through different call-sites of functions and reads/writes to different fields [Reps 2000; Zhang and Su 2017].

Many practical program analyses need to track multiple matching properties simultaneously [Arzt et al. 2014; Huang et al. 2015; Ramalingam 2000; Späth et al. 2019; Sridharan and Bodík 2006]. For example, a context- and field-sensitive alias analysis [Xu et al. 2009; Yan et al. 2011] needs to consider only the paths in which each method call is matched with the corresponding method return (i.e., context-sensitivity); at the same time, it should consider the paths in which data assigned to a field of an object can only be passed via an access on the same field (i.e., field-sensitivity). One Dyck language can be used to capture context-sensitivity, and a second Dyck language can be used to capture field-sensitivity. Consequently, when an analysis needs to keep track of two properties simultaneously, we formulate the problem as an InterDyck-reachability problem: an InterDyck language accepts all words that consist of the interleaving of two Dyck words. We define the interleaving operator and the InterDyck languages as follows:

Definition 2.4 (Interleaving Operator). Suppose that two languages $L_{1}$ and $L_{2}$ are defined over alphabets $\Sigma_{1}$ and $\Sigma_{2}$, respectively. The interleaving operator $\odot: L_{1} \times L_{2} \rightarrow\left(\Sigma_{1} \cup \Sigma_{2}\right)^{*}$ on words is defined as follows:

$$
\begin{aligned}
a \odot \epsilon & =\{a\} \text { for } a \in L_{1} \\
\epsilon \odot b & =\{b\} \text { for } b \in L_{2} \\
c_{1} a \odot c_{2} b & =\left\{c_{1} w \mid w \in\left(a \odot c_{2} b\right)\right\} \cup\left\{c_{2} w \mid w \in\left(c_{1} a \odot b\right)\right\} \\
& \quad \text { for } c_{1} \in \Sigma_{1}, c_{2} \in \Sigma_{2}, a \in L_{1}, b \in L_{2} .
\end{aligned}
$$

The interleaving operator can also be overloaded as a binary operator on languages:

$$
L_{1} \odot L_{2}=\bigcup_{a \in L_{1}, b \in L_{2}} a \odot b
$$

Definition 2.5 (Interleaved Dyck language). Given two Dyck languages $D_{m}$ and $D_{n}$ with disjoint alphabets, the interleaved Dyck language ("InterDyck language") constructed from $D_{m}$ and $D_{n}$ is defined as the language $D_{m} \odot D_{n}$.

```
class T { T f; T g; }
T getF(T a){ return a.f; }
void setF(T b, T val){ b.f = val; }
...
T x,y;
setF(y,x);
z = getF(y);
```

(a) Code snippet.

(b) Graph representation.

Fig. 2. Alias analysis via InterDyck-language reachability.
Example 2.6. Consider the two Dyck languages $D_{p}$ and $D_{b}$ defined by the following grammars:

$$
\begin{aligned}
& D_{p}: S::=S S|(S)| \epsilon \\
& D_{b}: S:=S S|[S]| \epsilon .
\end{aligned}
$$

The interleaved Dyck language $D_{p} \odot D_{b}$ contains such words as "([)]", "([)(])", etc.

### 2.3 Bidirected InterDyck-Reachability

We proceed to define the InterDyck-reachability problem on bidirected graphs.
Definition 2.7 (Bidirected InterDyck-Reachability). Consider an L-reachability problem on a digraph $G=(V, E)$, where $L$ is the InterDyck language $D_{p} \odot D_{b}$. Let $O$ and $C$ denote the sets of all open- and close-parenthesis symbols from $D_{p} \odot D_{b}$, respectively. For each edge $e=\langle u, v\rangle \in E$, if the edge label is an open parenthesis " $(i$ " $\in O$, there must exists a reverse edge $\bar{e}=\langle v, u\rangle$ with edge label $"{ }^{\prime}{ }^{\prime} \in C$, and vice versa. We say that $G$ is a bidirected graph (for the InterDyck language $D_{p} \odot D_{b}$ ) and the reachability problem is a bidirected InterDyck-reachability problem.

### 2.4 Applying InterDyck-Reachability in Program Analysis

We illustrate how InterDyck-reachability is used in program analysis by sketching how it can be used in a context- and field-sensitive alias analysis for a Java-style language [Yan et al. 2011] shown in Figure 2. We use one Dyck language $D_{p}^{\prime}$ to capture context-sensitivity, and another Dyck language $D_{b}^{\prime}$ to capture field-sensitivity. In $D_{p}^{\prime}$, an open parenthesis " $(i$ " represents a method call at line $i$; a close parenthesis ") ${ }^{\prime}$ " models a method return to the call-site at line $i$. In $D_{b}^{\prime}$, an open bracket " $[f \text { " represent a write to the field } f \text {; a close bracket " }]_{f}$ " represents a read to the field $f$. Consider the program in Figure 2a. The alias analysis checks whether two variables can point to the same object. In the program, the $f$ field of variable $y$ is set to the value of variable $x$, and then variable $z$ is set to the value of the $f$ field of $y$. Thus, variables $x$ and $z$ can potentially (and in this case, must) point to the same object.

Figure 2 b gives the corresponding graph. The graph is augmented with inverse edges. Specifically, for each edge $\langle p, q\rangle$ labeled with " $\left[f\right.$ ", there always exists an inverse edge $\langle q, p\rangle$ labeled with "] ${ }_{f}$ ", and vice versa. Similarly, each "( $i$ "-labeled edges is accompanied with the corresponding ") ${ }_{i}$ "-labeled edges in the inverse direction. The bidirectedness is a prerequisite for CFL-reachability-based pointer/alias analysis [Reps 1998; Sridharan et al. 2005], otherwise, two nodes may not be reachable from each other even by standard graph reachability. To check whether variables $x$ and $z$ can be aliases, it suffices to decide whether there exists an InterDyck-path from node $x$ to $z$. In graph $G$, there is a valid InterDyck-path " $\left.{ }_{6}[f)_{6}(7]_{f}\right)_{7}$ ", indicating that the variable $x$ and $z$ can be aliases in the original program. Note that the alias analysis based on InterDyck-reachability is a simplification of the alias reachability presented by Sridharan and Bodík [2006]. The field edges approximate the field loads and stores in the flow graph [Sridharan and Bodík 2006; Sridharan et al. 2005].

The approximation may lead to spurious aliasing, as discussed by Xu et al. [2009, §4]. However, experimental results show that, in practice, the overall performance is better than the algorithms proposed by Sridharan et al. [2005] and Sridharan and Bodík [2006].

## 3 OVERVIEW OF THE BIDIRECTED $D_{1} \odot D_{1}$-REACHABILITY PROOF

This section presents the key ideas behind the proof that bidirected $D_{1} \odot D_{1}$-reachability is polynomial-time solvable. As in most graph-reachability problems, in $D_{1} \odot D_{1}$-reachability it is infeasible to collect up all $D_{1} \odot D_{1}$-paths because the number of such paths can be infinite, and the lengths of such paths can be unbounded as well. Instead, the goal is to identify all $D_{1} \odot D_{1}$-reachable node pairs (Definition 2.1). The argument for identifying all $D_{1} \odot D_{1}$-reachable pairs is achievable in PTIME is structured into two parts:

- We first show that every bidirected $D_{1} \odot D_{1}$-reachable pair $(u, v)$ is also connected by a $D_{1} \odot D_{1}$-path of a special form, which we call a shallow $D_{1} \odot D_{1}$-path (Definition 3.2).
- We then give a polynomial-time tabulation algorithm to find all $(u, v)$ pairs that are connected by at least one shallow $D_{1} \odot D_{1}$-path.


### 3.1 Bidirectedness and Shallow Paths

Bidirectedness plays a pivotal role in the proof. In particular, bidirectedness enables a path to use the reverse edges, which allows a path to go back and forth along a sequence of edges in the graph. As we illustrate below, for a given path $P$ that contains some cycles, it is possible to construct a new path $P^{\prime}$ in which the cycles are arranged differently. Cycle rearrangement is the key insight that enables us to prove the existence of shallow paths, which we now define.

A shallow path is characterized by the number of unmatched open parentheses in each prefix of the path. To specify the property that a shallow path must satisfy, we first define a quantity to characterize the unmatched open parentheses in a given path prefix.

Definition 3.1 (Tuple of Unmatched Open Parentheses). Let $G=(V, E)$ be a bidirected graph for a bidirected $D_{1} \odot D_{1}$-reachability problem, and let $P=v_{0} v_{1} \ldots v_{l}$ be a $D_{1} \odot D_{1}$-path in $G$. We define the tuple of unmatched open parentheses at the $i^{\text {th }}$ node of $P$, denoted by unmatch $(P, i)$, to be the tuple ( $a, b$ ), where

- $a$ is the number of unmatched open parentheses in the first Dyck language in the length-i prefix $v_{0} v_{1} \ldots v_{i}$ of the path $P$,
- $b$ is the number of unmatched open parentheses in the second Dyck language in the length-i prefix $v_{0} v_{1} \ldots v_{i}$ of the path $P$.
Definition 3.2 (Shallow Path). Let $G=(V, E)$ be a bidirected graph for a $D_{1} \odot D_{1}$-reachability problem. A $D_{1} \odot D_{1}$-path $P$ in $G$ is shallow if, for each $i$ between 1 and $|P|$, the tuple $(a, b)=$ unmatch $(P, i)$ satisfies $a \leq 6 n$ or $b \leq 6 n$, where $n=|V|$ is the number of nodes in $G$.

Figure 3 illustrates the state space of unmatched open parentheses in $D_{1} \odot D_{1}$-paths, along with the trajectory in that space of a specific shallow $D_{1} \odot D_{1}$-path. In particular, because the number of unmatched open parentheses for the two Dyck languages must never simultaneously exceed $6 n$ in a shallow path, the trajectory can never enter the $(6 n, \infty) \times(6 n, \infty)$ region of Figure 3.

Given a $D_{1} \odot D_{1}$-path between two nodes $u$ and $v$ in bidirected graph $G$, there is always a shallow $D_{1} \odot D_{1}$-path that connects $u$ and $v$. This property is captured by the following theorem:

Theorem 3.3. In a bidirected $D_{1} \odot D_{1}$-reachability problem, if there exists a $D_{1} \odot D_{1}$-path from node $u$ to node $v$, then there exists a shallow $D_{1} \odot D_{1}$-path from $u$ to $v$.

We defer the proof of Theorem 3.3 to Section 4.1, but note here that to establish the existence of the shallow path, the key observations used in the proof are as follows:


Fig. 3. State space for unmatched open parentheses in $D_{1} \odot D_{1}$-paths, and the trajectory of a shallow path.




Fig. 4. Cycle rearrangement in a bidirected graph. In this graph, we assume the bound for the number of the unmatched parentheses is 30 . Cycle rearrangement moves one cycle $C_{2}$ from a later position to the current position to decrease the current number of unmatched parentheses in the path.

- The bidirectedness of the graph allows a path to go back and forth along a sequence of edges.
- This flexibility always allows us to transform a non-shallow path into a shallow one, by rearranging cycles in the original $D_{1} \odot D_{1}$-path.
The following example provides some intuition about these points, particularly about how cycle rearrangement is used to transform a non-shallow $D_{1} \odot D_{1}$-path into a shallow one.

Example 3.4 (Cycle Rearrangement). Consider the graph shown in Figure 4. There exists a $D_{1} \odot D_{1-}$ path from node $a$ to node $c$ that first goes through cycle $C_{1}$ seven times and then cycle $C_{2}$ five times. As we follow the path and reach node $b$ for the first time, the tuple of unmatched open parentheses/brackets is $(35,34)$. From this $D_{1} \odot D_{1}$-path, we wish to construct a shallow(er) path, which we will do by exploiting bidirectedness.

To keep the example simple, let us assume that the bound is 30 (rather than $6 n$ ). In general, our construction makes use of bidirectedness of the graph to go back and forth along a sequence of edges that connect two cycles in the path.

In our example, we need to rearrange the positions of the cycles so that the number of unmatched open parentheses never goes higher than 30 . When the original path reaches node $b$ for the first time, the number of unmatched parentheses/brackets for $D_{1}$ and $D_{1}$ both exceed 30 . We can construct an alternative path in which the number of unmatched symbols for at least one of the two languages is no larger than 30 ; this path is labeled as the "rearranged path" in Figure 4. The rearranged path is based on the original path; the difference is that one occurrence of cycle $C_{2}$ is shifted earlier in the path. The rearranged path begins by going around cycle $C_{1}$ six times (instead of seven); then uses the edge $\langle a, b\rangle$ to reach node $b$; goes around cycle $C_{2}$ once; then uses the (reversed) edge $\langle b, a\rangle$ to return to node $a$; goes around cycle $C_{1}$ a final time; and finishes up by returning to $b$, going around $C_{2}$ four times, and going to node $c$. Because the first occurrence of cycle $C_{2}$ cancels some open parentheses and open brackets, the number of unmatched symbols for at least one of the two
languages-in this case, both languages-is no larger than 30 at each point in the rearranged path. In general, it turns out that we can always transform a non-shallow $D_{1} \odot D_{1}$-path into a shallow one via such cycle rearrangements (Section 4.1).

### 3.2 From Shallow Paths to a Tabulation Algorithm

Theorem 3.3 reduces the problem of bidirected $D_{1} \odot D_{1}$-reachability to the problem of finding all pairs of nodes that are connected by a shallow $D_{1} \odot D_{1}$-path. In this section, we sketch why shallow $D_{1} \odot D_{1}$-reachability can be solved in PTIME. At a high level, the reason is that shallowness-i.e., at each position in the path, the number of unmatched parentheses for at least one of the languages is bounded-enables us to use a summary-based approach to store information about certain segments of shallow paths. These summary edges can be tabulated in polynomial time, as can information about all node pairs that are connected by shallow $D_{1} \odot D_{1}$-paths.

In Section 4.2, we give a tabulation algorithm for finding summary edges and all node pairs that are connected by shallow $D_{1} \odot D_{1}$-paths (Algorithm 2). Figure 3 can be used to illustrate the concept behind the tabulation strategy used in the algorithm.

- For a shallow $D_{1} \odot D_{1}$ path, because it never steps into the $(6 n, \infty) \times(6 n, \infty)$ region, once it steps out of the $[0,6 n] \times[0,6 n]$ region from the $a=6 n$ line, it always stays in the $(6 n, \infty) \times[0,6 n]$ region until it comes back to the line $a=6 n$. It indicates that even though the path segments in the $(6 n, \infty) \times[0,6 n]$ region can have an unbounded number of open parentheses in one Dyck language, eventually, these parentheses are all matched with the close parentheses. An analogous observation holds for the $[0,6 n] \times(6 n, \infty)$ region.
- Such segments can be represented by summary edges, of which there are only a polynomial number, and we can compute all the summary edges for path segments in the $(6 n, \infty) \times[0,6 n]$ and $[0,6 n] \times(6 n, \infty)$ regions in polynomial time.
- We can explore all paths within the $[0,6 n] \times[0,6 n]$ region-including all summary edges on the $[0,6 n] \times\{6 n\}$ and $\{6 n\} \times[0,6 n]$ borders-in polynomial time.
- From the information tabulated in this fashion, we can enumerate all shallow $D_{1} \odot D_{1}$ reachable pairs in polynomial time.

Theorem 3.5. For a given bidirected $D_{1} \odot D_{1}$-reachability problem, if there is a shallow $D_{1} \odot D_{1}$-path between nodes $u$ and $v$, Algorithm 2 will identify $u, v$ are shallow $D_{1} \odot D_{1}$-reachability.

Together with Theorem 3.3, we have the following corollary:
Corollary 3.6. Algorithm 2 solves the bidirected $D_{1} \odot D_{1}$-reachability problem precisely.

## 4 PROOF OF THE PTIME SOLVABILITY OF BIDIRECTED $D_{1} \odot D_{1}$-REACHABILITY

This section presents the details of the PTIME proof for bidirected $D_{1} \odot D_{1}$-reachability. In Section 4.1, we prove Theorem 3.3, demonstrating how to construct a shallow path, given a $D_{1} \odot D_{1}$ reachable path. In Section 4.2, we prove Theorem 3.5, establishing that our tabulation algorithm identifies all shallow $D_{1} \odot D_{1}$-reachable pairs, and analyze the time complexity of the algorithm. Our result shows that the bidirected $D_{1} \odot D_{1}$-reachability problem is in PTIME.

### 4.1 Shallow-Path Property

This section proves that between any $D_{1} \odot D_{1}$-reachable nodes $u, v$, there always exists a shallow path that connects $u$ and $v$. We prove this property by showing that, given any $D_{1} \odot D_{1}$-path $P$ in a bidirected graph, we can construct a shallow $D_{1} \odot D_{1}$-path $P^{\prime}$ based on $P$. In the shallow-path construction, we exploit the bidirectedness of the graph. The bidirected edges enable the constructed path $P^{\prime}$ to go back and forth over segments of $P$, which provides the flexibility to rearrange in $P^{\prime}$


Fig. 5. Regions and region transitions during shallow-path construction.
the position of cycles from the original path $P$. We divide the tuples of unmatched open parentheses for a given path position into four regions as follows (see Figure 5):

Definition 4.1 (Unmatched Number Regions). We classify a $D_{1} \odot D_{1}$-path $P$ at each node position as being in one of four states. In particular, for the path $P$ at the $i^{t h}$ node position, we define four regions according to the values in the tuple ( $a, b$ ) of unmatched open parentheses at the $i^{\text {th }}$ node, as depicted in Figure 5:

- Forbidden Region: When the tuple $(a, b)$ satisfies $a>6 n$ and $b>6 n$, path $P$ is in the forbidden region at the $i^{\text {th }}$ node. A shallow path never steps into the forbidden region.
- Warning Region: When the tuple $(a, b)$ satisfies $a \geq 4 n$ and $b \geq 4 n$, but is not in the forbidden region, path $P$ is in the warning region at the $i^{\text {th }}$ node.
- Safe Region A: When the tuple $(a, b)$ satisfies $a \geq 2 n$ and $b \geq 2 n$, but is not in the warning region or the forbidden region, path $P$ is in safe region $A$ at the $i^{\text {th }}$ node.
- Safe Region B: If at the $i^{\text {th }}$ node the path $P$ is not in any of the regions defined above, it is in safe region $B$ at the $i^{\text {th }}$ node.

We devise a method to construct a shallow path $P^{\prime}$ based on the original path $P$. The starting point for the construction is the observation that if the constructed path $P^{\prime}$ never steps into the forbidden region, $P^{\prime}$ is a shallow $D_{1} \odot D_{1}$-path. The construction therefore "steers" $P^{\prime}$ away from the forbidden region. The key insight behind the method is that when we find path $P^{\prime}$ stepping into the warning region, we can move a cycle from a later position to the current position so that the constructed path $P^{\prime}$ is forced out of the warning region after the rearrangement. Because the warning region acts as a buffer around the forbidden region, $P^{\prime}$ never has a chance to step into the forbidden region. The proof is constructive, so we describe the proof in the form of algorithms.

Shallow Path Construction Algorithm. Algorithm 1 describes our approach to shallow-path construction. In the algorithm, we use a tuple ( $s_{1}, s_{2}$ ) with $s_{1}, s_{2} \in\{+,-, 0, ?\}$ to characterize the type of a cycle. When the cycle has more open-parentheses edges than close-parenthesis edges in the first Dyck language, $s_{1}$ is ' + '. If it has more close parentheses than open parentheses, $s_{1}$ is ' - '. For cycles with equal number of open and close parentheses, $s_{1}$ is ' 0 '. When the relation between the numbers of open and close parentheses of the first Dyck language is unimportant, $s_{1}$ is '?'. The value of $s_{2}$ characterizes the cycle's type with respect to the second Dyck language in an analogous fashion. In the algorithm, we treat the nodes $v_{i}, v_{j}, i \neq j$, in a path $P=v_{0} v_{1} \ldots v_{|P|}$ as different nodes due to their different indexes, even though they may correspond to the same node in the actual graph. This treatment helps the algorithm distinguish cycles according to the

```
Algorithm 1: Shallow-Path Construction
    Input : Edge-labeled bidirected graph \(G=(V, E)\),
                An \(D_{1} \odot D_{1}\)-path \(P=v_{0} v_{1} \cdots v_{|P|}\)
    Output:A shallow \(D_{1} \odot D_{1}\)-path \(P^{\prime}\)
    \(i \leftarrow 0\)
    \(L_{1}, L_{2} \leftarrow \varnothing\)
    \(P^{\prime} \leftarrow\) empty path
    while \(i \neq|P|\) do
        prev \(\leftarrow\) get_region \(\left(P^{\prime}\right)\)
        \(P^{\prime} \leftarrow P^{\prime} \cdot v_{i}\)
        \(i \leftarrow i+1\)
        cur \(\leftarrow\) get_region \(\left(P^{\prime}\right)\)
        \((a, b) \leftarrow\) unmatch \(\left(P^{\prime},\left|P^{\prime}\right|\right)\)
        if prev is safe region \(A\), cur is warning region then
            if \(a=4 n\) then
                cycle rearrangement for cycle type \((-\), ?)
            if \(b=4 n\) then
                    cycle rearrangement for cycle type (?, -)
        if prev is safe region \(A\) and cur is safe region \(B\) then
            if there exists \(a(-,-)\) cycle \(C\) in \(L_{1} \cup L_{2}\) then
                reverse rearrangement for \(C\)
        if \(v_{i}\) is in some cycle \(C\) in the lists \(L_{1}, L_{2}\) then
            \(i=i+|C|\)
            \(L_{1} \leftarrow L_{1} \backslash\{C\}, L_{2} \leftarrow L_{2} \backslash\{C\}\)
    return \(P^{\prime}\)
```

sequence of (indexed) nodes that they consist of. For example, two cycles $C_{1}=v_{i} v_{i+1} v_{i+2} v_{i+3}$ and $C_{2}=v_{j} v_{j+1} v_{j+2} v_{j+3}$ are considered as two different cycles, even though they may actually represent the same cyclic sequence of edges in the graph.

In Algorithm 1, the variable $i$ tracks the index in the original path $P$ that corresponds to the position in the new path $P^{\prime}$. We also introduce two rearrangement lists $L_{1}$ and $L_{2}$. In particular, $L_{1}$ records the cycles that are moved earlier in $P^{\prime}$ to lower the number of unmatched parentheses for the first Dyck language; $L_{2}$ records the moved cycles for the second Dyck language. The cycles in the lists $L_{1}$ and $L_{2}$ have been already used in the new path $P^{\prime}$, but not yet in the current prefix of the original path $P$. Algorithm 1 constructs the shallow path as follows:

- Lines 5-9. In most cases, the newly constructed path $P^{\prime}$ simply follows the original path $P$. In lines 6-7, the algorithm appends node $v_{i}$ to $P^{\prime}, i$ is incremented to indicate that node $v_{i}$ has already been used in $P^{\prime}$. In lines 5 and $8-9$, the algorithm collects the information of unmatched tuple $(a, b)$ of $P^{\prime}$ and the region of $P^{\prime}$ to help decide whether $P^{\prime}$ need to deviate from path $P$ in the construction. The construction deviates from path $P$ base on three cases.
- Lines $\mathbf{1 0 - 1 4}$. The algorithm detects that the constructed path $P^{\prime}$ steps into the warning region from safe region $A$. To prevent the path going further and possibly stepping into the forbidden region, the algorithm "moves" a cycle that occurs later in path $P$ to the end of path $P^{\prime}{ }^{2}$ As shown in lines $11-12$, if the path steps into the warning region with $a=4 n-$ i.e., the

[^1]number of unmatched open parentheses for the first Dyck language is $4 n$-the algorithm uses a (,- ?) cycle $C$, which will bring down the number of unmatched parentheses for the first Dyck language, thereby forcing $P^{\prime}$ to leave the warning region. Then the rearrangement also adds cycle $C$ to $L_{1}$, to record that this cycle has already been used in $P^{\prime}$, but has not yet been reached in the original path $P$. Lines 13-14 are the analogous steps for paths stepping into the warning region with $4 n$ unmatched open parentheses of the second Dyck language. We call this cycle-moving behavior "cycle rearrangement." The details of cycle rearrangement are discussed in Procedure 1 later in the section, where we also demonstrate that (i) it is always feasible to perform cycle rearrangement, and (ii) during the rearrangement, path $P^{\prime}$ cannot step into the forbidden region.

- Lines 15-17. The algorithm detects that the path $P^{\prime}$ steps into the safe region $B$ from safe region $A$. If there is a $(-,-)$ cycle $C$ recorded in either $L_{1}$ or $L_{2}$, the algorithm moves the reverse cycle of $C$, denoted by $\bar{C}$, to the end of $P^{\prime}$, thereby cancelling the rearrangement effects for cycle $C$. It also removes cycle $C$ from its rearrangement lists $L_{1}$ and $L_{2}$. We call this reverse-cycle movement "reverse rearrangement." We present the details of reverse rearrangement in Procedure 2 later in the section. Reverse rearrangement guarantees the validity of the constructed path $P^{\prime}$, i.e., during the construction of $P^{\prime}$, there are always more open parentheses than close parentheses in both $D_{1}$ languages.
- Lines 18-20. The algorithm finds that the original path $P$ encounters a cycle that has already been used in the constructed path $P^{\prime}$. It skips over the cycle in $P$, and removes the cycle from the rearrangement lists $L_{1}$ and $L_{2}$.
Figure 5 illustrates all possible transitions between regions in shallow-path construction. Due to the invariants of the construction, a shallow path $P^{\prime}$ can never step into the forbidden region. Consequently, $P^{\prime}$ is a shallow $D_{1} \odot D_{1}$-path. In the remainder of the section, we describe the details of cycle rearrangement and reverse rearrangement.

Cycle Rearrangement. When the constructed path $P^{\prime}$ steps into the warning region, Algorithm 1 performs a cycle rearrangement. The goal of cycle rearrangement is to lower the number of unmatched open parentheses in the constructed path $P^{\prime}$ by moving a (,- ?) or (?, - ) cycle that occurs in a latter position of $P$ to the current position of $P^{\prime}$, thereby decreasing the number of unmatched open parentheses for the appropriate Dyck language, and forcing path $P^{\prime}$ to step out of the warning region after the rearrangement.

Procedure 1 performs the cycle rearrangement. The rearrangement is for a cycle of some specified type, held in the input parameter type. The cycle-rearrangement algorithm attaches the closest simple cycle $C$ of the required type to the end of path $P^{\prime}$. Suppose at the beginning of Procedure 1 , a path $P^{\prime}$ ends with a node $v_{i}$. Procedure 1 finds the path in $P$ from $v_{i}$ to cycle $C$. This connecting path is accumulated in variable bridge, and the path fragment appended to $P^{\prime}$ is the sequence bridge $\cdot C \cdot \overline{\text { bridge. Note that the effects of bridge and its reverse } \overline{\text { bridge }} \text { cancel each other, in the }}$ sense that the net change in the tuple of unmatched open parentheses at the end of $P^{\prime}$ versus at the end of $P^{\prime} \cdot$ bridge $\cdot C \cdot \overline{\text { bridge }}$ is due to the unmatched parentheses in $C$. In addition, Procedure 1 also records the information that cycle $C$ has been used in $P^{\prime}$ by adding cycle $C$ to the appropriate rearrangement list, $L_{1}$ or $L_{2}$, according to the sought-for cycle type: if type is (,- ?), $C$ is added to $L_{1}$; if type is (?, - ), $C$ is added to $L_{2}$.

Specifically, in lines 3-17, Procedure 1 considers simple cycles that are successively further and further away from $v_{i}$, while gathering up a cycle-free connecting path bridge between node $v_{i}$ and the current cycle under consideration. The test in lines 5-6 causes all cycles that are already in the rearrangement lists $L_{1}$ and $L_{2}$ to be skipped over because they have already been used in path $P^{\prime}$. To maintain bridge as a cycle-free path, whenever the algorithm identifies a cycle in lines 7-16,

```
Procedure 1: Cycle Rearrangement
    Input : Edge-labeled bidirected graph \(G=(V, E)\),
                The original \(D_{1} \odot D_{1}\) path \(P=v_{0} v_{1} \cdots v_{|P|}\)
                Currently constructed shallow path \(P^{\prime}\)
                Rearrangement lists \(L_{1}, L_{2}\)
                Current corresponding node \(v_{i}\)
                The type of cycle to rearrange: type \(\in\{(?,-),(-, ?)\}\)
    Output:Updated shallow path \(P^{\prime}\)
                Modified rearrangement lists \(L_{1}, L_{2}\)
    bridge \(\leftarrow v_{i}\)
    pos \(\leftarrow i\)
    while true do
        pos \(\leftarrow\) pos +1
        if \(v_{\text {pos }}\) is in any of the cycles of \(L_{1}\) or \(L_{2}\) then
            continue
        if \(v_{\text {pos }}\) as a node already exists in bridge then
            There exists a cycle \(C\) in bridge
            bridge \(\leftarrow\) bridge \(-C\)
            if \(C\) belongs to the desired cycle type type then
            \(P^{\prime} \leftarrow P^{\prime}+\) bridge \(+C+\overline{\text { bridge }}\)
            if type is \((-\), ?) then
                    \(L_{1} \leftarrow L_{1} \cup\{C\}\)
            if type is \((?,-)\) then
                \(L_{2} \leftarrow L_{2} \cup\{C\}\)
            break
        bridge \(\leftarrow\) bridge \(+v_{\text {pos }}\)
    return \(P^{\prime}, L_{1}, L_{2}\)
```

the cycle is deleted from connecting path bridge. If the cycle satisfies the cycle-type requirement, bridge $\cdot C \cdot \overline{\text { bridge }}$ is appended to $P^{\prime}$ (line 11), and $C$ is recorded in the appropriate rearrangement list (lines 12-15). Otherwise, the algorithm looks for the next cycle.

Note that it is always feasible and valid to carry out the cycle-rearrangement operation when it is invoked by Algorithm 1. Consider the tuple ( $a, b$ ) of unmatched open parentheses for node $v_{i}$ at the end of path $P^{\prime}$. We know that $a>4 n$ and $b>4 n$ because $v_{i}$ is in the warning region. Then the cycle-rearrangement algorithm can always find a cycle with more close parentheses than open parentheses, otherwise the original path cannot be a valid $D_{1} \odot D_{1}$-path. Then we show the validity of cycle rearrangement in terms of the shallow-path construction, i.e., during cycle rearrangement, path $P^{\prime}$ cannot step into the forbidden region. When Algorithm 1 invokes cycle rearrangement, path $P^{\prime}$ has just stepped into the warning region. Without loss of generality, let us assume that it was the first Dyck language that just crossed the $4 n$ threshold of unmatched open parentheses; i.e., $a=4 n+1$. We observe that

- because bridge is acyclic, $\mid$ bridge $\mid \leq n-1$, and
- for the simple cycle $C,|C| \leq n$.

Consequently, the total number of unmatched open parentheses at any point in the fragment


```
Procedure 2: Reverse Rearrangement
    Input :Edge-labeled bidirected graph \(G=(V, E)\),
            The original \(D_{1} \odot D_{1}\) path \(P=v_{0} v_{1} \cdots v_{|P|}\)
            Currently constructed shallow path \(P^{\prime}\)
            Rearrangement list \(L_{1}, L_{2}\)
            Current corresponding node \(v_{i}\)
            The simple cycle \(C\) to be reversed
    Output:Updated shallow path \(P^{\prime}\)
            Modified Rearrangement List \(L_{1}, L_{2}\)
    bridge \(\leftarrow v_{i}\)
    pos \(\leftarrow i\)
    while true do
        pos \(\leftarrow\) pos +1
        if \(v_{\text {pos }}\) is in \(C\) then
            bridge \(\leftarrow\) bridge \(+v_{\text {pos }}\)
            \(P^{\prime} \leftarrow\) bridge \(+\bar{C}+\overline{\text { bridge }}\)
            if \(C \in L_{1}\) then
                    Remove \(C\) from \(L_{1}\)
                else
                    Remove \(C\) from \(L_{2}\)
            break
        bridge \(\leftarrow\) bridge \(+v_{\text {pos }}\)
        if \(v_{\text {pos }}\) as a node already exists in bridge then
            There exists a cycle \(C\) in bridge
            bridge \(\leftarrow\) bridge \(-C\)
    return \(P^{\prime}, L_{1}, L_{2}\)
```

unmatched parentheses for the first Dyck language cannot exceed $6 n$ at any point, meaning that the path $P^{\prime} \cdot$ bridge $\cdot C \cdot \overline{\text { bridge }}$ never enters the forbidden region.

Reverse Rearrangement. In Algorithm 1, some simple cycles of $P$ of type (,- ?) or (?, - ) are moved to earlier relative positions in $P^{\prime} . P^{\prime}$ could use up all of the unmatched open parentheses, in which case one or both components of the tuple of unmatched open parentheses at some position of $P^{\prime}$ could become negative; this indicates that $P^{\prime}$ becomes an invalid $D_{1} \odot D_{1}$-path. The shallow path construction uses reverse rearrangement to ensure the validity of the constructed path $P^{\prime}$. In particular, reverse rearrangement attaches to $P^{\prime}$ the reverse cycle $\bar{C}$ of a moved (,-- ) cycle $C$ that is already part of $P^{\prime}$. The net effect of reverse rearrangement is to raise the number of open parentheses at the end of $P^{\prime}: \bar{C}$ is a $(+,+)$ cycle.
Procedure 2 presents the reverse-rearrangement method. The loop body (lines 4-16) carries out the search for the cycle $C$. Note that the search starts from position $i$ (line 2 ), which represents the corresponding position in the original path $P^{\prime}$ of the current position in the constructed path $P^{\prime}$. If the algorithm finds $C$, it appends bridge $\cdot \bar{C} \cdot \overline{\text { bridge }}$ to $P^{\prime}$. The net effect is equivalent to appending the reverse cycle $\bar{C}$ at the end of $P^{\prime}$. In lines 14-16, if the algorithm encounters an irrelevant cycle, it is removed from bridge. This guarantees bridge is always acyclic.

Note that the reverse-rearrangement operation is always feasible when Algorithm 1 invokes it. The connecting path bridge is acyclic, so we have $\mid$ bridge $\mid \leq n-1$. $C$ is a simple cycle, so $|C| \leq n$. Consequently, the total number of close parentheses in bridge $\cdot \bar{C} \cdot \overline{\text { bridge }}$ is at most $2 n-1$. When


Fig. 6. The reason why $L_{1}=\varnothing$ by the time $P^{\prime}$ encounters a (,- ?) cycle. The dashed arrow represents the path $P^{\prime}$ from the moment $m_{1}$ when it finishes the last cycle rearrangement, to the moment $m_{2}$ when the value of $a$ becomes no larger than $2 n$.

Algorithm 1 invokes reverse rearrangement, path $P^{\prime}$ has just stepped into safe region B. Thus, for the tuple of unmatched open parentheses $(a, b)$, it has $a \geq 2 n$ and $b \geq 2 n$. Without loss of generality, let us assume that it was the first Dyck language that just touched the $2 n$ threshold of unmatched open parentheses; i.e., $a=2 n$ and $b>2 n$. When bridge $\cdot \bar{C} \cdot \overline{\text { bridge }}$ is attached to $P^{\prime}$, the total number of unmatched open parentheses for the first Dyck language cannot reach 0 at any point, meaning that the path $P^{\prime} \cdot$ bridge $\cdot \bar{C} \cdot \overline{\text { bridge }}$ never uses up all the unmatched parentheses.

Correctness of Shallow Path Construction. To present the correctness of the shallow path construction, we first verify the constructed path is indeed shallow. Then we show that the constructed path is a valid InterDyck path. We verify the constructed path $P^{\prime}$ is shallow by showing that it never steps into the forbidden region. Algorithm 1 invokes cycle rearrangement to steer path $P^{\prime}$ from the warning region back to safe region A in lines 10-14. Similarly, for each invocation of reverse rearrangement in lines 15-17 of Algorithm 1, it steers $P^{\prime}$ from safe region B back to safe region A . The reason is that the reversal of a $(-,-)$ cycle is a $(+,+)$ cycle, and appending a $(+,+)$ cycle directs the trajectory of $P^{\prime}$ in the north-east direction in Figure 5. As we described above, during the construction, $P^{\prime}$ never steps into the forbidden region, thus it is a shallow path.

Then we show that during the construction, the path $P^{\prime}$ will never use up unmatched open parentheses/brackets, thus it is a valid InterDyck path. We start with a Lemma.

Lemma 4.2. Let $(a, b)=$ unmatch $\left(P^{\prime}, i\right)$. If $L_{1} \neq \varnothing$, then $a>2 n$; if $L_{2} \neq \varnothing$, then $b>2 n$.

Proof. By the symmetry of $a$ and $b$, it suffices to prove just one of the two cases-if $L_{1} \neq \varnothing$, then $a>2 n$. We prove the lemma by contradiction. Suppose that the path $P^{\prime}$ constructed so far satisfies $L_{1} \neq \varnothing$ and $a \leq 2 n$. Let us compare two moments $m_{1}$ and $m_{2}$ in the construction of $P^{\prime}:($ (i) $m_{1}$, earlier in the construction of $P^{\prime}$, just after the most-recent cycle rearrangement, and (ii) $m_{2}$, the (hypothetical) current moment when we have $L_{1} \neq \varnothing$ and $a \leq 2 n$. At $m_{1}$, we must have $a>3 n$. The reason is because a cycle rearrangement happens when path $P^{\prime}$ steps into the warning region (and hence $a>4 n$ ), and a cycle rearrangement always attaches a simple cycle with length less than $n$. Thus, as the construction proceeded from $m_{1}$ to $m_{2}$, the number $a$ of unmatched open symbols for the first Dyck language changed from $a>3 n$ to $a \leq 2 n$. In other words, $a$ dropped in value by more than $n$ from moment $m_{1}$ to moment $m_{2}$.

We first demonstrate that such a drop can happen only if $P^{\prime}$ passes through some ( - , ?) cycle $C_{1}$. To establish this claim, suppose that $P^{\prime}$ does not go through a (,- ?) cycle; then the path appended to $P^{\prime}$ as the construction proceeded from $m_{1}$ to $m_{2}$ can be decomposed into several ( + ,?) cycles and a set of path segments $S_{\text {chains }}$, where there are no repetitions of nodes among any of the path segments in $S_{\text {chains }}$. Each ( + ,?) cycle can only increase $a$, so only the path segments in $S_{\text {chains }}$ can decrease the value of $a$. Because $S_{\text {chains }}$ has no repeated nodes, the total number of nodes is less
than $n$, which contradicts the fact that the decrease of $a$ needs to be larger than $n$. Consequently, as the construction proceeded from $m_{1}$ to $m_{2}, P^{\prime}$ must have been extended with some (,- ?) cycle $C_{1}$.

All cycles recorded in $L_{1}$ have the type ( - ,?). Because of the greedy search carried out by Procedure 1 , cycle rearrangement is always performed with respect to the (,- ?) cycle that is closest to the position in path $P$ at the time Procedure 1 is invoked. Consequently, at moment $m_{1}$ every cycle recorded in $L_{1}$ comes earlier in path $P$ than cycle $C_{1}$. Figure 6 depicts this situation, where cycle $C_{2}$ is an arbitrary member of $L_{1}$. The dashed arrow indicates how as the construction of $P^{\prime}$ proceeds from $m_{1}$ to $m_{2}$, each cycle in $L_{1}$ (such as $C_{2}$ ) is skipped over and removed from $L_{1}$-see lines 18-20 of Algorithm 1. This aspect of Algorithm 1 will drain $L_{1}$ of all its members before the algorithm gets to append cycle $C_{1}$ to $P^{\prime}$. As a consequence, $L_{1}=\varnothing$ at moment $m_{2}$, which contradicts the assumption that $L_{1} \neq \varnothing$. We have shown that it is impossible to have $L_{1}=\varnothing \wedge a \leq 2 n$, and thus if $L_{1} \neq \varnothing$, then $a>2 n$ as was to be shown.

The following theorem establishes that the constructed path $P^{\prime}$ is a valid $D_{1} \odot D_{1}$-path.
Theorem 4.3. In the newly constructed path $P^{\prime}$, at each node position $j$, let $(a, b)=\operatorname{unmatch}\left(P^{\prime}, j\right)$. Then $a \geq 0 \wedge b \geq 0$. Consequently, $P^{\prime}$ is a valid $D_{1} \odot D_{1}$-path.

Proof. We prove this theorem by contradiction, showing that it is impossible to have $a<0$ or $b<0$. Without loss of generality, suppose that at some moment during the construction of $P^{\prime}$, $a<0$ at the end of $P^{\prime}$-i.e., $P^{\prime}$ moved from safe region B to a situation in which $a<0$. In this case, by Lemma 4.2, we must have $L_{1}=\varnothing$-otherwise, $a$ must be greater than $2 n$.

Now consider the other rearrangement list, $L_{2}$. First, suppose that $L_{2}$ is also $\varnothing$-i.e., both $L_{1}=\varnothing$ and $L_{2}=\varnothing$. In this situation, either (i) no cycle rearrangements have been performed so far in creating $P^{\prime}$, (ii) for each rearranged cycle $C$ that appears in $P^{\prime}$, a reverse rearrangement with respect to $C$ has also been performed in $P^{\prime}$, or (iii) for each rearranged cycle $C$ that appears in $P^{\prime}$, the while-loop in lines 4-20 of Algorithm 1 has also encountered $C$ and removed it from $L_{1}$ or $L_{2}$ in line 20. Note that in Algorithm 1, the value of variable $i$ establishes the correspondence between the current position in $P^{\prime}$ and a prefix $v_{0} v_{1} \ldots v_{i}$ of the original path $P$. When $L_{1}=\varnothing \wedge L_{2}=\varnothing$, if $P^{\prime}$ contains a cycle $C$ from a rearrangement, it either contains the cycle $\bar{C}$ from a subsequent reverse rearrangement, or the prefix $v_{0} v_{1} \ldots v_{i}$ of path $P$ also contains the cycle $C$. The prefix $v_{0} v_{1} \ldots v_{i}$ of $P$ and $P^{\prime}$ are "synchronized" in the weak sense that unmatch $(P, i)=$ unmatch $\left(P^{\prime},\left|P^{\prime}\right|\right)$. In particular, the number of unmatched open parentheses at the ends of $v_{0} v_{1} \ldots v_{i}$ and $P^{\prime}$ must be equal. Consequently, it is impossible to have $a<0$, otherwise at location $i$ in $P$, the number of unmatched open parenthesis would also be negative. It means that the original path $P$ is not a valid $D_{1} \odot D_{1}$-path, which contradicts the underlying assumption of the shallow-path construction.

Now suppose that $L_{2} \neq \varnothing$. Except during cycle rearrangement, path $P^{\prime}$ is either in safe region B or safe region A throughout the shallow-path construction. Cycles can only be inserted in the rearrangement list $L_{2}$ by cycle rearrangements. After a cycle rearrangement, the path $P^{\prime}$ is in the safe region A. Note that if $P^{\prime}$ ever enters safe region A, it can only re-enter safe region B after all $(-,-)$ cycles have been drained from $L_{2}$ and $L_{1}$. The reason is because whenever a reverse rearrangement is performed in lines $15-17$ of Algorithm $1, P^{\prime}$ is steered back to safe region A. Thus, $P^{\prime}$ can only re-enter safe region B if there are no more (,-- ) cycles in $L_{1} \cup L_{2}$ Consequently, whenever $P^{\prime}$ is in safe region B , if there are cycles recorded in $L_{2}$ they must have the type $(\{0,+\},-)$. Now consider path $P^{\prime}:$ it contains each of the $(\{0,+\},-)$ cycles that are recorded in $L_{2}$. We now compare the number of unmatched open parentheses at the end of prefix $v_{0} v_{1} \ldots v_{i}$ of $P$ and at the end of $P^{\prime}$ (as constructed so far). Because all the cycles in $L_{2}$ have type ( $\{0,+\},-$ ), path $P^{\prime}$ can have some additional unmatched open parentheses compared to path $P$ at position $i$. If at the end of $P^{\prime}$ we have $a<0$, then the number of unmatched open parentheses at position $i$ of $P$ must also be
negative. That would mean that path $P$ is not a valid $D_{1} \odot D_{1}$-path, which again contradicts the underlying assumption of the shallow-path construction.

Theorem 4.3 guarantees that the constructed path is always a valid $D_{1} \odot D_{1}$-path. The shallowpath construction also ensures that $P^{\prime}$ cannot step into the forbidden region. Consequently, the path $P^{\prime}$ is a shallow $D_{1} \odot D_{1}$-path.

### 4.2 Bidirected $D_{1} \odot D_{1}$-Reachability Algorithm

This section presents our algorithm to compute all pairs connected by shallow $D_{1} \odot D_{1}$-paths for bidirected $D_{1} \odot D_{1}$-reachability problem, which-by the results of Section 4.1-provides the solution to the bidirected $D_{1} \odot D_{1}$-reachability problem. As shown in Figure 3, for a shallow path, we can enumerate all the intermediate states of the path in the $[0,6 n] \times[0,6 n]$ region. For the path segments that pass out of the square region $[0,6 n] \times[0,6 n]$, we track them using the "CO summary", which is defined as follows:

Definition 4.4 (Close-Open (CO) summary). A CO summary counts a bounded number of unmatched close parentheses and unmatched open parentheses, while allowing any number of matched brackets. The CO summary of unbalanced parentheses, denoted by $\mathrm{CO}_{)}$, is defined by the grammar:

```
\(C_{0} O_{0}:=\epsilon\)
\(C_{i} O_{j}::=\) ' \(\left.{ }^{\prime} C_{i} O_{j}\right]\) ' for \(0 \leq i, j \leq 6 n\)
\(C_{i} O_{j}::={ }^{\prime}{ }^{\prime} C_{i-1} O_{j}\) for \(1 \leq i \leq 6 n\)
\(C_{i} O_{j}::=C_{i} O_{j-1}\) '(' for \(1 \leq j \leq 6 n\)
\(C_{i} O_{j}::=C_{a} O_{b} C_{c} O_{d} \quad\) where \(i=a+\max (c-b, 0), j=d+\max (b-c, 0), 0 \leq i, j, a, b, c, d \leq 6 n\).
```

The language of CO summary involves the alphabets of two Dyck languages. Above, we used parentheses to represent the symbols of the first Dyck language and brackets for the symbols of the second Dyck language. $n$ denotes the number of nodes in the graph $G$.

The CO summary of the unbalanced brackets is also defined in a similar manner, denoted by $\mathrm{CO}_{][ }$, in which the roles of parentheses and brackets are interchanged. Both $\mathrm{CO}_{\text {) }}$ ( and $\mathrm{CO}_{\text {] }}$ are context-free languages, and hence the all-pairs $\mathrm{CO}_{\text {) }}$ ( and all-pairs $\mathrm{CO}_{[\mathrm{I}}$-reachability problems can be computed via standard CFL-reachability algorithms in polynomial time [Reps 1998]. We refer to such information as "CO summary information."

We then argue that the CO summary information can summarize all path segments that pass out of, and then return to, the $[0,6 n] \times[0,6 n]$ region in Figure 3. The path segment $u \leadsto v$ in Figure 3 is an example of the kind of path segment that will be summarized by CO summary information. In a $\mathrm{CO}_{)}$(-path, there can be no more than $6 n$ unmatched close and open parentheses, whereas the brackets are always matched. Similarly, in a $\mathrm{CO}_{[\mathrm{l}}$-path, there can be no more than $6 n$ unmatched close and open brackets, whereas the parentheses are always matched. Consider again the path segment $u \leadsto v$, it has no more than $6 n$ unmatched close and open brackets, but the parentheses are always matched. In other words, $u \leadsto v$ is a $\mathrm{CO}_{[ }[-$path. A dual kind of path segment that passes out of the central $[0,6 n] \times[0,6 n]$ region, traverses the $[0,6 n] \times(6 n, \infty)$ region, and then returns to the $[0,6 n] \times[0,6 n]$ region has no more than $6 n$ unmatched close and open parentheses, but the brackets are always matched. Consequently, it is a $\mathrm{CO}_{)}$(-path. Then, the finite sets of pairs obtained by computing all-pairs $\mathrm{CO}_{\text {) }}$ ( and all-pairs $\mathrm{CO}_{][ }$-reachability can summarize all possible path segments of shallow $D_{1} \odot D_{1}$-paths in $G$ that pass out of, and return to, the central $[0,6 n] \times[0,6 n]$ region.

Assume that we have the access to the CO summary information $S_{1}=\mathrm{CO}_{][ }$and $S_{2}=\mathrm{CO}_{\text {) }}$, which can be computed in polynomial-time, Algorithm 2 identifies all $D_{1} \odot D_{1}$ shallow-path-reachable pairs. Algorithm 2 constructs a new undirected graph $G^{\prime}$. Each node in $G^{\prime}$ has a label $(v, a, b)$. The

```
Algorithm 2: Precise \(D_{1} \odot D_{1}\) Algorithm
    Input :Edge-labeled bidirected graph \(G=(V, E)\),
                CO summary information \(S_{1}=\mathrm{CO}_{\text {] }}\)
                CO summary information \(S_{2}=\mathrm{CO}\) )
    Output:A set of \(D_{1} \odot D_{1}\) reachable pairs Sol
    Sol \(\leftarrow \varnothing\).
    \(G^{\prime} \leftarrow\) a new graph with \(n * 6 n * 6 n\) nodes and no edges.
    In graph \(G^{\prime}\), give each node a label \((i, j, k)\) where \(1 \leq i \leq n, 1 \leq j, k \leq 6 n\)
    foreach \(u \rightarrow v \in E\) do
        for \(j=1\) to \(6 n-1\) do
            for \(k=1\) to \(6 n\) do
                add an undirected edge between node \((u, j, k)\) and \((v, j+1, k)\) in \(G^{\prime}\)
    foreach \(u \xrightarrow{[ } v \in E\) do
        for \(j=1\) to \(6 n\) do
            for \(k=1\) to \(6 n-1\) do
                add an undirected edge between node \((u, j, k)\) and \((v, j, k+1)\) in \(G^{\prime}\)
    foreach summary \(u \xrightarrow{C_{a} O_{b}} v \in S_{1}\) do
        for \(j=a\) to \(\min (6 n-b+a, 6 n)\) do
            for \(k=1\) to \(6 n\) do
                add an undirected edge between node \((u, j, k)\) and \((v, j-a+b, k)\) in \(G^{\prime}\).
    foreach summaryu \(\xrightarrow{C_{a} O_{b}} v \in S_{2}\) do
        for \(j=1\) to \(6 n\) do
            for \(k=a\) to \(\min (6 n-b+a, 6 n)\) do
                add an undirected edge between node \((u, j, k)\) and \((v, j, k-a+b)\) in \(G^{\prime}\)
    Run an undirected-graph-reachability algorithm on \(G^{\prime}\)
    foreach node pair \(((u, 0,0),(v, 0,0)) \in G^{\prime}\) do
        if \((u, 0,0)\) and \((v, 0,0)\) are connected \(\in G^{\prime}\) then
            \(\mathrm{Sol} \leftarrow \mathrm{Sol} \cup\{(\mathrm{u}, \mathrm{v})\}\)
    return Sol
```

first entry, $v$, is a node in the original graph $G$, whereas $a$ and $b$ are counts of unmatched open parentheses and unmatched open brackets, respectively. We only create nodes in $G^{\prime}$ for which $0 \leq a, b \leq 6 n$, and thus $G^{\prime}$ has $36 n^{3}+12 n^{2}+n$ nodes in total.

- Lines $\mathbf{4 - 1 1}$. For every open-parenthesis-edge between $v_{1}, v_{2}$ in the original graph, we create an undirected edge between $\left(v_{1}, i, j\right)$ and $\left(v_{2}, i+1, j\right)$ as described in lines 4-7. Similarly, in lines $8-11$, for an open-bracket-edge between $v_{1}, v_{2}$, we create an undirected edge between $\left(v_{1}, i, j\right)$ and $\left(v_{2}, i, j+1\right)$. Thus, all path segments of shallow $D_{1} \odot D_{1}$-paths in $G$ that stay within the central $[0,6 n] \times[0,6 n]$ region will be connected in $G^{\prime}$. For example, if between nodes $u, v$ in $G$ there is a path with two unmatched open parentheses and three unmatched open brackets, then there will be a path in $G^{\prime}$ between the nodes $(u, 0,0)$ and $(v, 2,3)$.
- Lines $\mathbf{1 2 - 1 9}$. To capture the path segments that pass out of and return to the $[0,6 n] \times[0,6 n]$ square, we use the $C O$ summary information to guide the edge construction. In lines $12-15$, if
nodes $u, v$ are connected by a $C_{p} O_{q}$ summary in $S_{1}=\mathrm{CO}_{\mathrm{]}}$, then we connect node ( $u, i, j$ ) and $(v, i, j-p+q)$. Similarly in lines $16-19$, the algorithm performs the similar construction for $C O$ summary information in $\left.S_{2}=\mathrm{CO}\right)($. By this means, all path segments of shallow $D_{1} \odot D_{1}$-paths in $G$ that pass out of and return to the $[0,6 n] \times[0,6 n]$ square are represented in the new undirected graph $G^{\prime}$.
- Lines 20-23. Bidirected $D_{1} \odot D_{1}$-reachability is now performed by a standard algorithm for undirected graph reachability on $G^{\prime}$. As shown in lines 21-23, if nodes $(u, 0,0)$ and $(v, 0,0)$ are connected, nodes $u, v$ are $D_{1} \odot D_{1}$-reachable in the bidirected $D_{1} \odot D_{1}$-reachability problem.
Complexity Analysis of the Bidirected $D_{1} \odot D_{1}$-reachability Algorithm. Here we analyze the time complexity of our algorithm for bidirected $D_{1} \odot D_{1}$-reachability. The algorithm has two steps: (i) compute the CO summary information, and (ii) perform Algorithm 2. To compute the CO summary information in step (i), we can use any standard CFL-reachability algorithm. One standard algorithm has $O\left(n^{3}\right)$ time complexity when the grammar size is considered to be a constant [Reps 1998]. However, in our case the size of the grammar is not a constant: the CO summary grammar has $n^{2}$ non-terminals and $O\left(n^{4}\right)$ productions. A variation of the time-complexity argument for the standard algorithm can be used to show that the running time for the CFL-reachability step is $O\left(n^{7}\right)$. For step (ii), the for loops in lines 4-7 and 8-11 of Algorithm 2 have time complexity $O\left(n^{3}\right)$. Lines 12-15 and 16-19 insert edges into graph $G^{\prime}$. $G^{\prime}$ has $O\left(n^{3}\right)$ nodes, and in the worst case the number of edges is $O\left(n^{6}\right)$; thus, in the worst case, the number of insertions is $O\left(n^{6}\right)$. Finally, running a standard algorithm for undirected-graph reachability on a graph with $m$ edges takes $O(m)$ time, and thus the time complexity of line 20 is $O\left(n^{6}\right)$. To sum up, the total running time for the precise $D_{1} \odot D_{1}$-reachability algorithm is $O\left(n^{7}\right)$; therefore, the $D_{1} \odot D_{1}$-reachability problem is in PTIME.


## 5 BIDIRECTED $D_{K} \odot D_{K}$-REACHABILITY COMPLEXITY

The InterDyck-reachability problem on directed graphs has been known to be undecidable [Reps 2000]. Note that the graph construction used in the undecidability proof of Reps [2000] cannot be applied to the bidirected case. This section shows that even with the bidirected relaxation, the problem is still at least NP-hard. We establish the hardness by reducing an NP-complete problem, paths avoiding forbidden pairs, to the bidirected $D_{k} \odot D_{k}$-reachability problem.

### 5.1 Overview

We first present an overview of the bidirected $D_{k} \odot D_{k}$-reachability proof. In our reduction from the paths avoiding forbidden pairs (PAFP) problem to bidirected $D_{k} \odot D_{k}$-reachability, we construct a new graph such that $D_{k} \odot D_{k}$-reachability between nodes of the new graph provides the solution for the corresponding PAFP instance. The graph for $D_{k} \odot D_{k}$-reachability consists of three major parts: the transformed graph, the preprocessing graph, and the checker graph. To facilitate the discussion of the reduction, we first reduce the PAFP problem to another variant of bidirected InterDyck-reachability: the $N$-fold $D_{1}$-reachability problem (Section 5.2). The reduction to $N$-fold $D_{1}$-reachability serves as a bridge to the reduction to bidirected $D_{k} \odot D_{k}$-reachability. In particular, it contains the same three parts as the bidirected $D_{k} \odot D_{k}$-reachability reduction. However, one of the parts, the checker graph, is much simpler and more intuitive, thus paving the way for the illustration of the reduction to bidirected $D_{k} \odot D_{k}$-reachability (Section 5.3). We start by introducing the PAFP problem, which is defined as follows:

Definition 5.1 (Paths Avoiding Forbidden Pairs). Given a graph $G=(V, E)$ and two nodes $s$, $t \in V$ and a set of node pairs $F=\left\{f_{1}, \ldots, f_{|F|}\right\} \subseteq(V \times V)$, the paths avoiding forbidden pairs problem is to decide whether there exists a path froms to $t$ such that the path contains at most one node in each node pair in $F$. Such paths are called F-paths.


Fig. 7. A path avoiding forbidden pairs instance.


Fig. 8. Overview for the $N$-fold $D_{1}$ reachability reduction.
The complexity of PAFP problems has been extensively studied in the literature [Gabow et al. 1976; Kovác 2014]. It is known that the complexity for many variants of the problem are NP-complete, and here we focus on the variant on directed acyclic graphs for our reduction.

Theorem 5.2 ( [Gabow et al. 1976]). The paths with forbidden pairs problem on directed acyclic graphs is NP-complete.

Example 5.3 (Paths Avoiding Forbidden Pairs). Figure 7 provides an instance of the PAFP problem. Consider the directed acyclic graph in Figure 7. It is associated with a forbidden-pair set $F=$ $\{(s, b),(s, c),(b, d)\}$. The instance is to decide whether there exists a path from node $s$ to node $t$ such that the path only involves at most one of the nodes in each forbidden pair in the set $F$. This instance has a positive answer because the path $s \rightarrow a \rightarrow d \rightarrow t$ is an $F$-path, i.e., it involves at most one node in each forbidden pair.

### 5.2 Complexity of Bidirected $N$-Fold $D_{1}$-Reachability

This section establishes the NP-hardness of the $N$-fold $D_{1}$-reachability problem on bidirected graphs. It serves as an warm-up construction for $D_{k} \odot D_{k}$-reachability on bidirected graphs. We define the bidirected $N$-fold $D_{1}$-reachability problem as follows.

Definition 5.4 (Bidirected $N$-fold $D_{1}$-reachability). An instance of the $N$-fold $D_{1}$ reachability problem is an instance of L-reachability where $L$ is an interleaved Dyck language $L::=\underbrace{}_{1} \odot D_{1} \odot \cdots \odot D_{1}$. An instance of the bidirected $N$-fold $D_{1}$-reachability problem is a problem
$N$ times
instance on a graph that is bidirected for the alphabet symbols in the $N$ Dyck languages.
We then present the polynomial time reduction from PAFP problem to $N$-fold $D_{1}$-reachability problem. In this case, the upper bound for $N$ is $O\left(|V|^{2}\right)$, where $V$ is the set of nodes of the graph $G=(V, E)$ in the PAFP problem. In the reduction, we construct a new graph with three major parts: the transformed graph $G_{t}$, the preprocessing graph $G_{p}$, and the checker graph $G_{c}$, illustrated in Figure 8. In the graph constructed for $N$-fold $D_{1}$-reachability, the preprocessing and transformed graphs are connected by an edge labeled by the $\epsilon$ symbol from the last node in $G_{p}$ to the $s_{\text {in }}$ node in $G_{t}$. The transformed graph and the checker graph are connected by another " $\epsilon$ "-edge from the $t_{\text {out }}$ node in $G_{t}$ to the start node in $G_{c}$. We design the graphs with the following properties in mind.

- Transformed Graph. The transformed graph $G_{t}$ is based on the original graph in the PAFP instance. There exists an $F$-path between nodes $s$ and $t$ in the PAFP problem iff there exists
a path in the transformed graph from $s_{\text {in }}$ to $t_{\text {out }}$ that spells out a word consisting of only non-repeated open parentheses.
- Preprocessing Graph. The preprocessing graph $G_{p}$ serves to append an arbitrary prefix of up to $|F|+|V|$ open parentheses. $F$ is the set of the forbidden pairs and $V$ is the node set of the PAFP graph. The constructed graph has $|F|+|V|$ different kinds of open parentheses. The preprocessing graph and the transformed graph can produce a path that uses every kind of open parenthesis exactly once iff the PAFP instance has a solution. This path could start from an arbitrary node in the preprocessing graph and end at node $t_{\text {out }}$ in $G_{t}$.
- Checker Graph. The checker graph $G_{c}$ will consume exactly one instance of every kind of open parenthesis, thereby checking whether $G_{p}$ and $G_{t}$ can produce such a path.
This construction reduces a PAFP instance to a corresponding instance of the $N$-fold $D_{1}$-reachability problem. If any node in the preprocessing graph is $N$-fold $D_{1}$-reachable to the end node of the checker graph, then in the original PAFP problem, there is an $F$-path from $s$ to $t$. If the PAFP problem has a solution, then at least one node in the preprocessing graph is $N$-fold $D_{1}$-reachable to the end node of the checker graph. The rest of the section describes the three major parts of the constructed graph for $N$-fold $D_{1}$-reachability.

Transformed-Graph Construction. To construct the transformed graph $G_{t}$, for every node $v$ in the original graph, we introduce two nodes $v_{\text {in }}$ and $v_{\text {out }}$ in the transformed graph. In the PAFP instance, the forbidden pair set $F=\left\{f_{1}, f_{2}, \ldots, f_{|F|}\right\}$ has a subset $F_{v}=\left\{f_{i_{1}}, \ldots, f_{i\left|F_{v}\right|}\right\}$ that contains all the forbidden pairs that involve node $v$. Then we connect $v_{\text {in }}$ to $v_{\text {out }}$ with a path that spells out the word that consists of all the open parentheses that have an index corresponding to a forbidden pair in $F_{v}$, followed by a final open parenthesis that represents node $v$ itself, e.g., " $\left(i_{1}\left(i_{2} \ldots i_{\left.\right|_{f v} \mid}\left(\right.\right.\right.$ node ${ }_{v}$ ". If there is no forbidden pair involving node $v$, i.e., $F_{v}=\varnothing$, then we connect $v_{\text {in }}$ with $v_{\text {out }}$ with a single edge labeled "( node $_{e}$ ". The edges in the transformed graph are bidirected. Figure 9 provides an example of the construction, for a node $a$ in the original graph. Suppose that in the forbidden-pair set $F$, the subset of all forbidden pairs involving node $a$ is $F_{a}=\left\{f_{i}, f_{j}, f_{k}\right\}$. We then add a path from $a_{\text {in }}$ to $a_{\text {out }}$ that spells out the word " $\left(i_{j}\left({ }_{k}\right.\right.$ node $_{a}$ ". For every edge $e=(u, v)$ in the original PAFP graph, we insert a bidirected $\epsilon$ edge from $u_{\text {out }}$ to $v_{\text {in }}$ in the transformed graph $G_{t}$. The transformed graph $G_{t}$ ensures the property in the following lemma.

Lemma 5.5. In the transformed graph, there exists a path from s to $t$ with only non-repeated open parentheses if and only if there exists an F-path in the original graph from sto $t$.

Proof. The "if" direction. Suppose that in the original paths avoiding forbidden pair problem, there exists an $F$-path from node $s$ to $t$. Then we can always find a simple $F$-path in the graph. Because the $F$-path involves at most one node in each node pair, for a simple $F$-path, due to the construction its corresponding path in the transformed graph contains at most one instance of each kind of open-parenthesis symbol.

The "only if" direction. Suppose that we have a path whose word contains at most one instance of each open-parenthesis symbol. Then there is always a simple path satisfying this "at most one instance" property as well. For the simple path, the construction forbids the possibility to exploit the bidirected edges in the transformed graph. Between every pair of nodes $v_{\text {in }}$ and $v_{\text {out }}$, to use the reversed edges from $v_{\text {out }}$ to $v_{\text {in }}$, the $)_{\text {node }}$ parenthesis edge must be matched with the corresponding open parenthesis, which only exists in the edge from $v_{\text {in }}$ to $v_{\text {out }}$. If there exists such an open parenthesis in the path, the path is not a simple path, because both of the nodes $v_{\text {in }}$ and $v_{\text {out }}$ are used twice. Thus, the simple path with the word containing each open parenthesis at most once corresponds to an $F$-path in the PAFP instance.


Fig. 9. Node construction in transformed graph.


Fig. 10. Transformed graph construction example for the PAFP instance in Example 5.3.


Fig. 11. Preprocessing graph.
Example 5.6 (Transformed-Graph Construction). Here we present a concrete example of the transformed-graph construction, based on the PAFP from Figure 7. In Figure 10, given the original graph and the forbidden-pair set $F$, we can construct the corresponding transformed graph following the construction described above.

This example illustrates the property of the transformed graph. For an $F$-path $s \rightarrow a \rightarrow d \rightarrow t$ in the original graph, there is a corresponding path $s_{\mathrm{in}} \rightarrow s_{\mathrm{out}} \rightarrow a_{\mathrm{in}} \rightarrow a_{\mathrm{out}} \rightarrow d_{\mathrm{in}} \rightarrow d_{\mathrm{out}} \rightarrow t_{\mathrm{in}} \rightarrow$ $t_{\text {out }}$ spells out the word " $\left.{ }_{1} \mathcal{2}_{2}\right)_{s}\left(a{ }_{3}\right)_{d}(t)$ ", which contains each open parenthesis at most once.

Preprocessing-Graph Construction. The transformed graph only produces paths with at most one edge for each kind of open parenthesis. However, the checker graph verifies whether there exists a path that contains every kind of open parenthesis exactly once. Preprocessing graph $G_{p}$ appends the other missing open parentheses that are not included in the transformed graph. For convenience, we denote the node set $V=\left\{v_{1}, \ldots, v_{|V|}\right\}$ and their corresponding parentheses as $\left(v_{1}, \ldots\left(_{v_{V \mid} \mid}\right.\right.$ In the preprocessing graph $G_{p}$, we have $|F|+|V|$ nodes $p_{1}, \ldots, p_{|F|+|V|}$. Each pair of nodes $p_{i}$ and $p_{i+1}$ are connected by $|F|+|V|$ parallel edges, with the labels " $\left(1\right.$ ", " $\left(2\right.$ ", $\ldots$, " $\left(|F|\right.$ ", " $\left(v_{1}\right.$ ", ... " $\left(v_{|V|}\right.$ ", as shown in Figure 11. The edges constructed are bidirected. In the preprocessing graph, a path from node $p_{i}$ to $p_{n}$ can append $n-i$ arbitrary open parenthesis to a path that starts at $p_{n}$. Due to the construction, we obtain the following lemma:

Lemma 5.7. In the constructed graph, $s_{i n}$ is $N$-fold $D_{1}$-reachable to $t_{\text {out }}$ if and only if there exists a path from some node $p_{i}$ in the preprocessing graph to the node start in the checker graph with exactly one open parenthesis for each kind.

Checker-Graph Construction. In the checker graph, there is one path from node start to node end that consumes every kind of open parenthesis exactly once. To construct such a checker graph for $N$-fold $D_{1}$-reachability, we only need to create a path that spells out " $\left.\left.\left.\left.)_{1}\right)_{2} \ldots\right)_{|F|}\right)_{v_{1}} \ldots\right)_{v_{|V|}}$ " from node start to node end (see Figure 12). Note that each constructed edge should be bidirected. In the checker graph, each close parenthesis consumes one corresponding open parenthesis. The checker


Fig. 12. Checker graph for $N$-fold $D_{1}$-reachability.


Fig. 13. Overview of checker graph $G_{\boldsymbol{c}}$ for bidirected $D_{k} \odot D_{k}$-reachability.


Fig. 14. Checker component $C_{i}$ for bidrected $D_{k} \odot D_{k}$-reachability. The ciin node contains $|F|+|V|-1$ parallel loops. The purpose is to transform the open parenthesis " $\left(k\right.$ " to " $[k \text { " in the other language. The ") })_{i}$ "-edge between the $c i_{i n}$ node and $c i_{\text {out }}$ node consumes exactly one " ( $i$ " parenthesis. The parallel loops at the $c i_{o u t}$ node can transform the "[ $k$ "-edges back to parenthesis edges. The total effect of the $C_{i}$ component is to consume an " ( $i$ " parenthesis regardless its position in the path.
graph contains exactly one occurrence of each kind of close parenthesis, thus matching with a path that has exactly one occurrence of each kind of open parenthesis.

When we put the three component graphs together, if there exists an $F$-path, then there exists an path from some node $p_{i}$ in the preprocessing graph to the node start of the checker graph with exactly one occurrence of each kind of open parenthesis. The continuation of the path to the end node of the checker graph causes the path to be accepted as an $N$-fold $D_{1}$-path. Conversely, if we find any $N$-fold $D_{1}$-path between any node in the preprocessing graph and the end node of the checker graph, there must exist a path in the transformed graph that contains every open parenthesis at most once. We can interpret that path as a solution in the paths avoiding forbidden pairs problem. We have reduced the paths avoiding forbidden pairs problem to the bidirected $N$-fold $D_{1}$-reachability problem, and thus the bidirected $N$-fold $D_{1}$-reachability is NP-hard.

Theorem 5.8. The bidirected $N$-fold $D_{1}$-reachability problem is NP-hard.

### 5.3 Complexity of Bidirected $D_{k} \odot D_{k}$-Reachability

The reduction from a PAFP instance to an instance of bidirected $D_{k} \odot D_{k}$-reachability again involves constructing a graph with three major parts. The preprocessing graph and the transformed graph used in the reduction remain the same as for $N$-fold $D_{1}$-reachability. The difference in the construction for the bidirected $D_{k} \odot D_{k}$-reachability problem is a more complicated checker graph.

We first provide an overview of the structure of the checker graph $G_{c}$ in Figure 13. There are $|F|+|V|$ components between the start and end nodes, each constructed to consume one kind of open parenthesis exactly once. In the reduction for the bidirected $D_{k} \odot D_{k}$-reachability problem, all the open parentheses that appear in the preprocessing graph and the transformed graph are symbols in the first Dyck language. The use of interleaving with symbols from the second Dyck language only comes into play in the checker graph. In contrast, in the construction for $N$-fold $D_{1}$-reachability, the open parentheses in the preprocessing graph and the transformed graph came
from separate $D_{1}$ languages. As we will see, for the bidirected $D_{k} \odot D_{k}$-reachability problem, in the checker graph, we also need to handle some issues related to the ordering of the open parentheses along the path through the preprocessing graph and the transformer graph.

We denote the $i^{\text {th }}$ graph component of the checker graph by $C_{i}$. In graph $C_{i}$, the goal is to consume one ( ${ }_{i}$ parenthesis from the path that came through the preprocessing graph and the transformer graph. Overall, the goal of the checker graph is to consume every open-parenthesis symbol (of the first Dyck language) exactly once. For each checker component $C_{i}$, as illustrated in Figure 14, the checker needs to consume one occurrence of the symbol " ${ }_{i}$ ". We illustrate the functionality of such a component with a concrete example. Suppose that there is a path " $1_{1}\left(_{2}()_{3}\right.$ " before the checker $C_{2}$. The purpose of the loop in the $c i_{i n}$ node (in this case $i=2$ ) in Figure 14 is to transform the open parentheses after " $(2$ " to correspondingly numbered bracket symbols. After the loop, the path can be " ${ }_{1}\left({ }_{2}()_{3}\right)_{3}\left[3\right.$ ". It then "consumes" the " $\left(2 \text { " symbol, and the path becomes " } 1_{1}\left({ }_{2}(3)_{3}\right]_{3}\right)_{2}$ ". Finally, the loop on the second node in Figure 14 will transform the bracket edges back to the corresponding parenthesis edges: the path " ${ }_{1}\left({ }_{2}(3) \text { becomes " }\left(1_{1}\left({ }_{2}\right)_{3}\right)_{3}\left[{ }_{3}\right)_{2}\right]_{3}(3$ " after the checker component $C_{2}$. Notice that because the path labeled " $\left.\left(1_{1}\left({ }_{2}()_{3}\right)_{3}[)_{3}\right)_{2}\right]_{3}(3$ " can always be matched by the same close-parenthesis string as the path labeled " $\left(1(3)\right.$ ", we can regard checker component $C_{2}$ as having "consumed" exactly one open-parenthesis symbol from the middle of the string, namely " ( 2 ". The checker graph for the construction checks whether a path contains each open-parenthesis symbol exactly once. By an argument similar to the one given in Section 5.2, we can show that this construction reduces the PAFP problem to bidirected $D_{k} \odot D_{k}$-reachability.

Theorem 5.9. The bidirected $D_{k} \odot D_{k}$-reachability problem is NP-hard.
In addition, we can use a $D_{2}$ language to simulate the behavior of a $D_{k}$ language in polynomial time. For example, a naïve approach is to use " $\left(_{1}(2 \ldots)(2\right.$ " to represent " $(i$." Thus, we have:

$$
\underbrace{}_{i \text { times }}
$$

Corollary 5.10. The bidirected $D_{2} \odot D_{2}$-reachability problem is NP-hard.

## 6 EVALUATION

We have implemented the precise bidirected $D_{1} \odot D_{1}$-reachability algorithm in Algorithm 2 and applied it to a context- and field-sensitive alias analysis for Java. The bidirected $D_{1} \odot D_{1}$-reachability algorithm provides a new over-approximation approach for bidirected $D_{k} \odot D_{k}$-reachability. Specifically, it first relaxes bidirected $D_{k} \odot D_{k}$-reachability to bidirected $D_{1} \odot D_{1}$-reachability and then solves the $D_{1} \odot D_{1}$-reachability problem precisely. To explore the precision of the bidirected $D_{1} \odot D_{1}$-reachability algorithm, our evaluation addresses the following research questions:

- As an over-approximation method for $D_{k} \odot D_{k}$-reachability, what is the precision of bidirected $D_{1} \odot D_{1}$-reachability compared with other approximation algorithms?
- How much over-approximation is there when bidirected $D_{1} \odot D_{1}$-reachability is used to over-approximate bidirected $D_{k} \odot D_{k}$-reachability?


### 6.1 Experimental Setup

Context-sensitive field-sensitive alias analysis. Our experiments adopt an existing context-sensitive field-sensitive alias analysis for Java [Xu et al. 2009; Yan et al. 2011]. We apply it to the standard Dacapo benchmark [Blackburn et al. 2006]. Section 2.4 describes the alias analysis with an example. Recall that the analysis is formulated as a bidirected InterDyck-reachability problem.

Over-approximating $D_{k} \odot D_{k}$-reachability algorithm. We compare the precision of our $D_{1} \odot D_{1-}$ reachability algorithm with a $D_{k} \odot D_{k}$-reachability algorithm based on linear conjunctive language reachability (LCL-reachability) [Zhang and Su 2017]. The $D_{k} \odot D_{k}$ languages are a subclass of the

Table 1. Precision and performance results for bidirected $D_{1} \odot D_{1}$-reachability, bidirected LCL-reachability, and bidirected $D^{\prime}$-reachability algorithms. Note that $D^{\prime}$-reachability under-approximates $D_{k} \odot D_{k^{-}}$ reachability. The precision column reports the numbers of $D_{k} \odot D_{k}$-reachable pairs.

|  | Precision |  |  | Time (s) |  |  |
| :--- | :---: | :--- | :--- | :--- | :--- | :--- |
|  | $D_{1} \odot D_{1}$ | LCL | $D^{\prime}$ | $D_{1} \odot D_{1}$ | LCL | $D^{\prime}$ |
| antlr | 1526931 | 1548761 | 1516159 | 22584.9 | 0.59 | 0.74 |
| bloat | 1593965 | 1614797 | 1585031 | 19249.7 | 0.62 | 0.83 |
| chart | - | 4866011 | 4711077 | - | 1.17 | 1.65 |
| eclipse | 1528643 | 1561457 | 1512137 | 32858.5 | 0.87 | 0.77 |
| fop | 3503778 | 3667004 | 3427676 | 44784.1 | 2.69 | 1.49 |
| hsqldb | 1515720 | 1536380 | 1507386 | 29122.4 | 0.56 | 0.69 |
| jython | 1568236 | 1592280 | 1557182 | 18782.9 | 0.73 | 0.89 |
| luindex | 1518813 | 1537773 | 1510299 | 16472.0 | 0.54 | 0.71 |
| lusearch | 1520970 | 1541082 | 1511882 | 27845.5 | 0.56 | 0.75 |
| pmd | 1574309 | 1593967 | 1566059 | 30241.3 | 0.56 | 0.76 |
| xalan | 1515312 | 1534172 | 1507196 | 21930.4 | 0.54 | 0.72 |

LCL languages. The LCL-reachability algorithm is inherently an over-approximation algorithm. We evaluate the LCL-reachability algorithm on the bidirected graphs.

An under-approximating $D_{k} \odot D_{k}$-reachability algorithm. To understand the false positives in both the $D_{1} \odot D_{1}$-reachability and the LCL-reachability algorithms, we compare the solution set with that of an under-approximating $D^{\prime}$-reachability algorithm. Let the InterDyck language be $D_{p} \odot D_{b}$. We can treat it as a Dyck language $D^{\prime}$ with alphabet $\Sigma_{D^{\prime}}=\Sigma_{D_{p}} \cup \Sigma_{D_{b}}$. The new Dyck language $D^{\prime}$ is a subset of the $D_{p} \odot D_{b}$ language: only well-balanced words are accepted; interleaved words are rejected. Therefore, we can run the $D_{k}^{\prime}$-reachability algorithm to obtain a lower bound on the number of pairs for both $D_{k} \odot D_{k}$ over-approximations.

Implementation We use the implementation described in the work of Yan et al. [2011] to derive the alias-analysis graphs from the DaCapo suite [Blackburn et al. 2006]. All graphs are pre-processed by a graph-simplification algorithm for InterDyck-reachability [Li et al. 2020]. We have implemented both the LCL-reachability algorithm and the precise bidirected $D_{1} \odot D_{1}$-reachability algorithm in $\mathrm{C}++$. All executables are compiled by $\mathrm{g}++-5.4$ with -O 2 optimization. The experiments were conducted on a machine with a Xeon E5-2650 CPU with 96 GB memory, running Ubuntu 16.04.

### 6.2 Precision Comparison with the LCL-Reachability Algorithm

To evaluate the precision, we compared the numbers of reported reachable pairs between the precise $D_{1} \odot D_{1}$-reachability algorithm and the LCL-reachability algorithm. Both algorithms overapproximate $D_{k} \odot D_{k}$-reachability. Table 1 presents the precision and performance results for both algorithms. For the $D_{1} \odot D_{1}$-reachability algorithm, we are able to obtain precise solutions to the relaxed problems, except for the chart benchmark, which consumes too much memory. In Table 1, we observe that the numbers of reachable pairs reported by the bidirected $D_{1} \odot D_{1}$-reachability algorithm is consistently smaller than those reported by LCL-reachability. On average, bidirected $D_{1} \odot D_{1}$-reachability obtains a solution that is about $2.0 \%$ smaller than the one obtained via the LCL-reachability algorithm. Table 1 indicates an interesting result: at least for context-sensitive field-sensitive alias analysis, the benefit of having a precise $D_{1} \odot D_{1}$-reachability algorithm can outweigh its disadvantage of being based on a less expressive formalism.

### 6.3 Precision Comparison with the $D^{\prime}$-Reachability Algorithm

To understand the degree of over-approximation introduced by using $D_{1} \odot D_{1}$-reachability to solve a relaxation of $D_{k} \odot D_{k}$-reachability, we also compared the precision of the $D_{1} \odot D_{1}$-reachability
algorithm with the under-approximating $D^{\prime}$-reachability algorithm. Table 1 shows that the gap between the numbers of reachable pairs found in the $D_{1} \odot D_{1}$-reachability algorithm and the numbers found by the under-approximating $D^{\prime}$-reachability algorithm is relatively small. On average, the precise $D_{1} \odot D_{1}$-reachability algorithm produced solutions that were $0.89 \%$ larger than the ones obtained via $D^{\prime}$-reachability. This result indicates that in the simplified $D_{k} \odot D_{k}$ graphs, the over-approximation introduced by ignoring distinctions between different parenthesis symbols in each Dyck language is insignificant.

### 6.4 Discussion

From Table 1, we can also observe that the bidirected $D_{1} \odot D_{1}$-reachability algorithm is very slow compared to the LCL-reachability algorithm. Recall that the time complexity of $D_{1} \odot D_{1}$-reachability algorithm is $O\left(n^{7}\right)$, which poses significant challenges for any practical use. Our work focuses on establishing the PTIME result for $D_{1} \odot D_{1}$-reachability. In particular, the $D_{1} \odot D_{1}$-reachability algorithm used in the experiments is directly adopted from the PTIME proof. On the other hand, our experimental study suggests that the $D_{1} \odot D_{1}$-reachability algorithm offers a precise overapproximation for $D_{k} \odot D_{k}$-reachability, after some pre-processing [Li et al. 2020]. Designing an efficient bidirected $D_{1} \odot D_{1}$-reachability algorithm remains an interesting future direction.

It is perhaps surprising to see that the bidirected $D_{1} \odot D_{1}$-reachability algorithm can achieve better precision than the LCL-reachability algorithm for $D_{k} \odot D_{k}$-reachability. Note that LCLreachability is perhaps the most precise $D_{k} \odot D_{k}$-reachability algorithm based on empirical evaluations [Zhang and Su 2017]. In our experiments, the input graphs of all algorithms are pre-processed by a graph-simplification algorithm [Li et al. 2020]. The graph simplification itself can also contribute to the precision improvements. Specifically, graph simplification utilizes the parenthesis-index information in $D_{k} \odot D_{k}$ to remove irrelevant edges. For the original bidirected graphs where $n$ is significantly larger, our $O\left(n^{7}\right)$-time $D_{1} \odot D_{1}$-reachability algorithm is too expensive to evaluate.

## 7 RELATED WORK

The computational complexity of InterDyck-reachability problem has been extensively studied. It has been shown that the InterDyck-reachability problem on directed graphs is undecidable [Ramalingam 2000; Reps 2000]. The undecidability proof by Reps [2000] reduces the Post Correspondence Problem (PCP) to InterDyck-reachability. Given two lists of words $L_{1}=\left\{x_{1}, \ldots, x_{k}\right\} . L_{2}=$ $\left\{y_{1}, \ldots, y_{k}\right\}$, the PCP problem is to decide whether there exists a list of indices $\left\{i_{1}, \ldots, i_{l}\right\}$ such that the string $x_{i_{1}} x_{i_{2}} \ldots x_{i_{l}}$ is the same as $y_{i_{1}} y_{i_{2}} \ldots y_{i_{l}}$. The proof use a Dyck language $D_{2}$ to encode the strings, with ' $1_{1}$ ' parenthesis representing ' 0 ' and the ' $(2$ ' parenthesis representing ' 1 '. The other Dyck language keeps track of the indices of the chosen words. A valid index list exists if and only if there exists an InterDyck-path in the graph representing $x_{i_{1}} x_{i_{2}} \ldots x_{i_{l}} \overline{y_{i} y_{i_{l-1}}} \ldots \overline{y_{i_{1}}}$.However, in the bidirected variants, the same proof technique does not work. In a directed graph, graphs can always force a path to first finish all $x_{i}$ words and then go through $y_{i}$ words. The bidirectedness of the graph introduces additional paths. A valid InterDyck-reachable path may represent the following string: $x_{i_{1}} x_{i_{2}} \overline{y_{i_{2}}} x_{i_{3}} \overline{y_{i_{3}} y_{i_{1}}}$. For these InterDyck-paths, we cannot reconstruct solutions of the PCP instance. Thus, the reduction technique does not apply to the bidirected variant, and the complexity for bidirected $D_{k} \odot D_{k}$-reachability still remains open.
InterDyck-reachability is the natural formalism when an analysis needs to keep track of two balanced-parenthesis properties. Various practical approaches have been proposed to overapproximate InterDyck-reachability solutions. The traditional approach includes using a contextfree language (CFL) to approximate the InterDyck language, and then solving the CFL-reachability problem [Sridharan and Bodík 2006; Sridharan et al. 2005; Yan et al. 2011]. Recently, the conjunctive-language-reachability approach was proposed to approximate InterDyck-reachability [Zhang and

Su 2017]. It uses a conjunctive language to describe an InterDyck language precisely, and provides an approximation algorithm for the reachability problem for linear conjunctive languages. The synchronized push-down systems (SPDSs) of Späth et al. [2019] also provides an approach to solve InterDyck-reachability. An SPDS uses a push-down automaton with two stacks. Each stack can handle the parenthesis matching for one Dyck language. To over-approximate $D_{k} \odot D_{k}$-reachability, SPDSs computes the $D_{k} \odot D_{k}$-reachability twice, each time ignoring one of the two stacks in the system. SPDSs collect the intersection of the two solution sets as the final result. SPDSs can also recognize incomplete InterDyck-paths, i.e., prefixes of InterDyck-paths.

In contrast to over-approximation algorithms, Madhusudan and Parlato [2011] propose a bounded tree-width approach, which can provide an under-approximate solution for InterDyck-reachability problems. InterDyck-reachability can be considered to be a variant of the state-reachability problem for multi-stack pushdown automata. In the context of concurrent program analysis, a natural semantic restriction for under-approximation is to consider only the states reachable within a bounded number of context switches [Qadeer and Rehof 2005]. The emptiness problem for multistack pushdown automata with bounded context-switches can be further reduced to the emptiness problem for graph automata for a special class of graphs (the so-called multiply nested words) with bounded tree-width. The bounded tree-width approach decides the emptiness problem for such graph automata. It provides a technique to explore state reachability for multi-stack pushdown automata with a bounded number of context switches, thus providing an under-approximation for the InterDyck-reachability problem.

## 8 CONCLUSION

This paper identifies an important class of InterDyck-reachability problems. Even though there are many practical applications of bidirected InterDyck-reachability, the computational complexity had been unaddressed. Our results in this paper start to fill in the picture. To the best of our knowledge, we have identified the first decidable variant of any InterDyck-reachability problem-namely, bidirected $D_{1} \odot D_{1}$-reachability-and showed that it is solvable in PTIME. We also showed that bidirected $D_{k} \odot D_{k}$-reachability is NP-hard.

We experimented with an implementation of the algorithm for precise bidirected $D_{1} \odot D_{1}$ reachability. We applied it to a context- and field-sensitive alias analysis. The experiment compared the precision of (i) relaxing the problem to bidirected $D_{1} \odot D_{1}$-reachability and solving the relaxed problem precisely, with (ii) a different over-approximation algorithm for bidirected $D_{k} \odot D_{k^{-}}$ reachability. Surprisingly, the $D_{1} \odot D_{1}$-reachability approach produced more precise results than the $D_{k} \odot D_{k}$-reachability approach. This experiment may show the benefit of obtaining a precise solution, even though the bidirected $D_{1} \odot D_{1}$-reachability formalism is inherently less expressive than the bidirected $D_{k} \odot D_{k}$-reachability formalism.

## ACKNOWLEDGEMENTS

We would like to thank the anonymous POPL reviewers for valuable feedback on earlier drafts of this paper, which helped improve its presentation. This work was supported in part by the United States National Science Foundation (NSF) under Grants No. 1917924 and No. 1420866; by a gift from Rajiv and Ritu Batra; by ONR under grants N00014-17-1-2889 and N00014-19-1-2318. Any opinions, findings, conclusions, or recommendations expressed in this publication are those of the authors, and do not necessarily reflect the views of the sponsoring agencies.

## REFERENCES

Steven Arzt, Siegfried Rasthofer, Christian Fritz, Eric Bodden, Alexandre Bartel, Jacques Klein, Yves Le Traon, Damien Octeau, and Patrick D. McDaniel. 2014. FlowDroid: precise context, flow, field, object-sensitive and lifecycle-aware taint analysis for Android apps. In ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '14, Edinburgh, United Kingdom - June 09-11, 2014. 259-269.
S. M. Blackburn, R. Garner, C. Hoffman, A. M. Khan, K. S. McKinley, R. Bentzur, A. Diwan, D. Feinberg, D. Frampton, S. Z. Guyer, M. Hirzel, A. Hosking, M. Jump, H. Lee, J. E. B. Moss, A. Phansalkar, D. Stefanović, T. VanDrunen, D. von Dincklage, and B. Wiedermann. 2006. The DaCapo Benchmarks: Java Benchmarking Development and Analysis. In OOPSLA '06: Proceedings of the 21st annual ACM SIGPLAN conference on Object-Oriented Programing, Systems, Languages, and Applications. 169-190.
Krishnendu Chatterjee, Bhavya Choudhary, and Andreas Pavlogiannis. 2018. Optimal Dyck reachability for data-dependence and alias analysis. Proc. ACM Program. Lang. 2, POPL (2018), 30:1-30:30.
Swarat Chaudhuri. 2008. Subcubic algorithms for recursive state machines. In Proceedings of the 35th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2008, San Francisco, California, USA, January 7-12, 2008, George C. Necula and Philip Wadler (Eds.). ACM, 159-169.
Ben-Chung Cheng and Wen-mei W. Hwu. 2000. Modular interprocedural pointer analysis using access paths: design, implementation, and evaluation. In Proceedings of the 2000 ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI 2000). ACM, 57-69.
Harold N. Gabow, Shachindra N. Maheswari, and Leon J. Osterweil. 1976. On Two Problems in the Generation of Program Test Paths. IEEE Trans. Software Eng. 2, 3 (1976), 227-231.
Wei Huang, Yao Dong, Ana Milanova, and Julian Dolby. 2015. Scalable and precise taint analysis for Android. In Proceedings of the 2015 International Symposium on Software Testing and Analysis (ISSTA 2015). ACM, 106-117.
Vineet Kahlon. 2009. Boundedness vs. Unboundedness of Lock Chains: Characterizing Decidability of Pairwise CFLReachability for Threads Communicating via Locks. In Proceedings of the 24th Annual IEEE Symposium on Logic in Computer Science, LICS 2009. IEEE Computer Society, 27-36.
John Kodumal and Alexander Aiken. 2004. The set constraint/CFL reachability connection in practice. In Proceedings of the ACM SIGPLAN 2004 Conference on Programming Language Design and Implementation 2004, Washington, DC, USA, fune 9-11, 2004. ACM, 207-218.
Jakub Kovác. 2014. On the Complexity of Rearrangement Problems under the Breakpoint Distance. 7. Comput. Biol. 21, 1 (2014), 1-15.

Yuanbo Li, Qirun Zhang, and Thomas W. Reps. 2020. Fast graph simplification for interleaved Dyck-reachability. In Proceedings of the 41st ACM SIGPLAN International Conference on Programming Language Design and Implementation, PLDI 2020, London, UK, June 15-20, 2020, Alastair F. Donaldson and Emina Torlak (Eds.). ACM, 780-793.
P. Madhusudan and Gennaro Parlato. 2011. The tree width of auxiliary storage. In Proceedings of the 38th ACM SIGPLANSIGACT Symposium on Principles of Programming Languages, POPL 2011, Austin, TX, USA, January 26-28, 2011. ACM, 283-294.
Shaz Qadeer and Jakob Rehof. 2005. Context-Bounded Model Checking of Concurrent Software. In Tools and Algorithms for the Construction and Analysis of Systems, 11th International Conference, TACAS 2005, Held as Part of the foint European Conferences on Theory and Practice of Software, ETAPS 2005, Edinburgh, UK, April 4-8, 2005, Proceedings (Lecture Notes in Computer Science, Vol. 3440). Springer, 93-107.
G. Ramalingam. 2000. Context-sensitive synchronization-sensitive analysis is undecidable. ACM Trans. Program. Lang. Syst. 22, 2 (2000), 416-430.
Thomas W. Reps. 1998. Program analysis via graph reachability. Inf. Softw. Technol. 40, 11-12 (1998), 701-726. https: //doi.org/10.1016/S0950-5849(98)00093-7
Thomas W. Reps. 2000. Undecidability of context-sensitive data-independence analysis. ACM Trans. Program. Lang. Syst. 22, 1 (2000), 162-186.
Johannes Späth, Karim Ali, and Eric Bodden. 2019. Context-, flow-, and field-sensitive data-flow analysis using synchronized Pushdown systems. Proc. ACM Program. Lang. 3, POPL (2019), 48:1-48:29.
Manu Sridharan and Rastislav Bodík. 2006. Refinement-based context-sensitive points-to analysis for Java. In Proceedings of the ACM SIGPLAN 2006 Conference on Programming Language Design and Implementation, Ottawa, Ontario, Canada, Fune 11-14, 2006. ACM, 387-400.
Manu Sridharan, Denis Gopan, Lexin Shan, and Rastislav Bodík. 2005. Demand-driven points-to analysis for Java. In Proceedings of the 20th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2005, October 16-20, 2005, San Diego, CA, USA. ACM, 59-76.
Guoqing (Harry) Xu, Atanas Rountev, and Manu Sridharan. 2009. Scaling CFL-Reachability-Based Points-To Analysis Using Context-Sensitive Must-Not-Alias Analysis. In ECOOP 2009- Object-Oriented Programming, 23rd European Conference,

Genoa, Italy, July 6-10, 2009. Proceedings, Sophia Drossopoulou (Ed.). 98-122.
Dacong Yan, Guoqing (Harry) Xu, and Atanas Rountev. 2011. Demand-driven context-sensitive alias analysis for Java. In Proceedings of the 20th International Symposium on Software Testing and Analysis, ISSTA 2011, Toronto, ON, Canada, July 17-21, 2011, Matthew B. Dwyer and Frank Tip (Eds.). ACM, 155-165.
Qirun Zhang and Zhendong Su. 2017. Context-sensitive data-dependence analysis via linear conjunctive language reachability. In Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, January 18-20, 2017, Giuseppe Castagna and Andrew D. Gordon (Eds.). ACM, 344-358.


[^0]:    ${ }^{1}$ Except in a few cases, we are not explicit about the actual parenthesis symbols used, referring to $D_{1} \odot D_{1}$ when each Dyck language has exactly one kind of parenthesis pair, and to $D_{k} \odot D_{k}$ when each language has more than one kind of parenthesis pair. It would be more precise to say $D_{k_{1}} \odot D_{k_{2}}$, but we prefer to avoid the double subscripting. In examples, when we wish to distinguish between the two languages, we use $D_{p} \odot D_{b}$ (" p " for "parentheses" and "b" for "brackets").

[^1]:    ${ }^{2}$ The algorithm "moves" the cycle in the sense previously explained in Example 3.4.

