

# Kidney Exchange: Faster Parameterized Algorithms and Tighter Lower Bounds

Palash Dey

Indian Institute of Technology  
Kharagpur, India  
palash.dey@cse.iitkgp.ac.in

Abhishek Sahu

National Institute of Science Education and Research  
Bhubaneswar, India  
abhiksheksahu@niser.ac.in

Aritra Banik

National Institute of Science Education and Research  
Bhubaneswar, India  
aritra@niser.ac.in

Sujoy Bhore

Indian Institute of Technology  
Bombay, India  
sujoy@cse.iitb.ac.in

## ABSTRACT

The kidney exchange mechanism allows many patient-donor pairs who are otherwise incompatible with each other to come together and exchange kidneys along a cycle. However, due to infrastructure and legal constraints, kidney exchange can only be performed in small cycles in practice. In reality, there are also some altruistic donors who do not have any paired patients. This allows us to also perform kidney exchange along paths that start from some altruistic donor. Unfortunately, the computational task is NP-complete. To overcome this computational barrier, an important line of research focuses on designing faster algorithms, both exact and using the framework of parameterized complexity.

The standard parameter for the kidney exchange problem is the number  $t$  of patients that receive a healthy kidney. The current fastest known deterministic FPT algorithm for this problem, parameterized by  $t$ , is  $O^*(14^t)$ . In this work, we improve this by presenting a deterministic FPT algorithm that runs in time  $O^*((4e)^t) \approx O^*(10.88^t)$ . This problem is also known to be W[1]-hard parameterized by the treewidth of the underlying undirected graph. A natural question here is whether the kidney exchange problem admits an FPT algorithm parameterized by the pathwidth of the underlying undirected graph. We answer this negatively in this paper by proving that this problem is W[1]-hard parameterized by the pathwidth of the underlying undirected graph. We also present some parameterized intractability results improving the current understanding of the problem under the framework of parameterized complexity.

## KEYWORDS

Kidney exchange; parameterized complexity; FPT algorithm; W-hardness; para-NP-hardness

### ACM Reference Format:

Palash Dey, Aritra Banik, Abhishek Sahu, and Sujoy Bhore. 2026. Kidney Exchange: Faster Parameterized Algorithms and Tighter Lower Bounds. In

All authors made equal contributions in this work.



This work is licensed under a Creative Commons Attribution International 4.0 License.

*Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). <https://doi.org/10.65109/VKIC3106>

*Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 9 pages. <https://doi.org/10.65109/VKIC3106>

## 1 INTRODUCTION

Kidney failure is a global health crisis, affecting millions worldwide. Patients suffering from acute renal failure typically have two main treatment options: dialysis or kidney transplantation. However, dialysis is associated with a significantly lower quality of life, and the average lifespan of recipients on dialysis is only around 10 years. As a result, kidney transplantation is the preferred choice for many recipients. Unfortunately, the gap between the demand for and supply of kidneys is vast. In the U.S., approximately 40,000 people are added to the kidney transplant waiting list every year, but only about 20,000 find a compatible donor [35]. This discrepancy has led to longer waiting times, with the current median wait for a kidney transplant from a deceased donor averaging 4.05 years [34]. Tragically, thousands of patients are removed from the waiting list each year because they either die or become too ill to undergo the surgery.

In this challenging context, Kidney Paired Donation (KPD), also known as the Kidney Exchange Program, was introduced in 2000 to provide an innovative solution to the problem of incompatible donor-recipient pairs [3, 30]. In KPD, recipients who have an incompatible donor can exchange their donor with another pair, forming a “barter market” [1, 3, 6] that increases the likelihood of successful kidney transplants. This system allows participants to find compatible donors more easily, often resulting in higher-quality matches that may lead to longer-lasting kidneys [32]. Since its introduction, KPD has expanded globally, with countries like the U.S., U.K., and India adopting the program, resulting in thousands of successful transplants and a significantly higher chance of matching better-quality kidneys [10].

The central challenge in any KPD system is the Kidney Exchange Problem, how to allocate kidneys to maximize the number of transplants. This is typically modeled as a directed graph where each vertex represents a patient-donor pair. A directed edge from vertex  $u$  to  $v$  indicates that the donor from  $u$  is compatible with the recipient at  $v$ .

Kidney exchanges usually happen through cycles, where each recipient receives a kidney from the donor of the previous vertex. Since donors can withdraw once their paired recipient is treated, all

surgeries in a cycle must be performed simultaneously to prevent defaults. Due to logistical constraints (each transplant involves two surgeries), only small cycles are allowed [1, 5, 7, 16, 27, 28].

In addition to cycles, some programs allow chains initiated by altruistic donors (non-directed donors, or NDDs) who do not have a paired recipient. In a chain, a donor gives to a compatible recipient, whose paired donor donates to another recipient, and so on. Unlike cycles, chains can proceed non-simultaneously, since a broken chain does not leave any recipient worse off [15, 31]. In the worst case, affected pairs can re-enter the exchange [31].

With altruistic donors, the problem becomes finding a collection of disjoint cycles and chains that maximizes coverage of patient-donor pairs while complying with logistical constraints. This forms the core optimization problem in kidney exchanges.

Since the computational problem is NP-complete, there is a long line of work focusing on designing heuristics that work reasonably well in practice. While existing clearing algorithms perform well in practice, understanding their theoretical limits provides insight into whether faster or more scalable mechanisms are even possible. Xiao and Wang [36] were the first to propose an exact algorithm for the kidney exchange problem that is guaranteed to output an optimal solution. Maiti and Dey [26] extended this line of work by initiating the study of parameterized complexity of the kidney exchange problem. This work improves the current state of the art in this line of research.

## 1.1 Our Contribution

Parameter	KIDNEY EXCHANGE
$t$	FPT ( $O^*(10.88^t)$ ) [Theorem 1]
$t + \ell_p + \ell_c$	No poly kernel [Theorem 2]
pathwidth	W[1]-hard ( $\ell_c = 0$ ) [Theorem 3]
pathwidth	W[1]-hard ( $\ell_p = 0$ ) [Theorem 4]
DFVS+ $\ell_p+\ell_c$	Para-NP Hard [Corollary 1]

**Table 1: Summary of results.**  $t$  : the number of patients receiving a kidney;  $\ell_p(\ell_c)$  : maximum allowed length of a kidney exchange path (cycle); DFVS: size of any minimum directed feedback vertex set.

In this work, we obtain the fastest known deterministic FPT algorithm for the KIDNEY EXCHANGE problem with respect to the parameter  $t$ , the number of patients helped. The previous best algorithms of Johnson et al. [19] ran in randomized time  $O^*(4^t)$  and deterministic time  $O^*(14^t)^1$ . Our approach improves the deterministic running time to  $O^*((4e)^t) \approx O^*(10.88^t)$ . The key ingredients are a careful combination of *color coding*-derandomized via  $t$ -perfect hash families, and a *subset-based dynamic program* that systematically builds vertex-disjoint cycles and altruistic paths. By adapting the fastest color-coding subroutines for detecting cycles and (altruistic paths) of prescribed colors, we achieve the stated bound.

<sup>1</sup> $O^*(\cdot)$  notation hides polynomial factors in  $n$ .

**THEOREM 1.** KIDNEY EXCHANGE parameterized by  $t$ , the number of patients who can be helped, admits a deterministic FPT algorithm running in time  $O^*((4e)^t) \approx O^*(10.88^t)$ .

We next examine the prospects for polynomial kernels for the same parameter  $t$ . It is worth noting that exploring the existence of a polynomial kernel for a practically significant problem, such as kidney exchange, is highly motivating, as kernelization algorithms can substantially reduce the size of problem instances commonly encountered in practice. While it is relatively straightforward to design FPT algorithms for our problem without aiming for a specific running time, obtaining a polynomial kernel turns out to be unlikely under the standard complexity theoretic assumption  $\text{NP} \subseteq \text{coNP/poly}$ . In particular, by a straightforward reduction from the DIRECTED  $k$ -PATH problem of Bodlaender, Jansen, and Kratsch [11]—which is known not to admit a polynomial kernel or compression unless  $\text{NP} \subseteq \text{coNP/poly}$ —we show that such a compression remains impossible even for the larger parameter of  $t + \ell_p + \ell_c + |\mathcal{B}|$ . This provides a nearly complete picture with respect to the natural solution-size parameter.

**THEOREM 2.** Under the assumption  $\text{NP} \not\subseteq \text{coNP/poly}$ , KIDNEY EXCHANGE admits no polynomial kernel (nor polynomial compression) parameterized by  $t + \ell_p + \ell_c + |\mathcal{B}|$  where  $\mathcal{B}$  is the set of altruistic vertices.

We further strengthen the hardness landscape established by Johnson et al. [19], who proved that the problem remains W[1]-hard when parameterized by the treewidth of the underlying graph. Our results show that the problem remains W[1]-hard even when parameterized by the strictly larger parameter of pathwidth, and that even the addition of  $\ell_c$  or  $\ell_p$  alone does not make the problem (fixed-parameter)tractable. Specifically, we provide two separate hardness results for KIDNEY EXCHANGE problem: one on bounded-pathwidth DAGs (where  $\ell_c = 0$ ) and another on bounded-pathwidth graphs with no altruistic vertices (where  $\ell_p = 0$ ). These results follow from a parameter-preserving reduction from the BIN PACKING problem with unary-encoded weights [20].

It is worth noting that when parameterized by all three parameters together, the problem becomes FPT, as a corollary of the result by Maiti et al. [26], who proved fixed-parameter tractability when parameterized by *treewidth* and  $\max \ell_p, \ell_c$ . Together, these results not only provide a more precise boundary of tractability but also give a complete picture of the parameterized complexity of KIDNEY EXCHANGE with respect to *pathwidth*,  $\ell_p$ , and  $\ell_c$ .

**THEOREM 3.** KIDNEY EXCHANGE is W[1]-hard when parameterized by the undirected pathwidth of the input graph, even when the graph is acyclic.

**THEOREM 4.** KIDNEY EXCHANGE is W[1]-hard when parameterized by the undirected pathwidth of the input graph, even when no altruistic vertex is present.

While Theorem 3 already implies W-hardness when parameterized by the size of a Directed Feedback Vertex Set (DFVS), we strengthen this result by establishing para-NP-hardness even for the larger parameter DFVS +  $\ell_p + \ell_c$ , by showing that the problem remains NP-hard on DAGs even when  $\ell_p$  is bounded. Note that since the associated graph is a DAG, one may assume  $\ell_c = 0$  without loss of generality. Our proof uses a carefully designed two-step

polynomial reduction from the 3-PARTITION problem, where each set of a given weight corresponds to an equivalent altruistic path of proportional length in the resulting instance of the kidney exchange problem.

**THEOREM 5.** *The KIDNEY EXCHANGE problem on DAGs is NP-hard even when restricted to instances with a fixed constant path-length bound  $\ell_p$ .*

**Corollary 1.** *The KIDNEY EXCHANGE problem is para-NP-hard with respect to  $DFVS+\ell_p+\ell_c$ .*

## 1.2 Related Work

Rapaport first proposed the idea of kidney exchange [30]. By 2005, over 60,000 patients were listed on the UNOS transplant registry [32], and the New England Renal Transplant Oversight Committee approved a regional clearinghouse for exchanges in 2004 [5]. Roth et al. [4] introduced strategy-proof, constrained-efficient mechanisms under 0–1 preferences. Some studies limit exchanges to cycles [13, 22, 33], while others allow both cycles and chains [17, 28, 36]. All variants can be framed as graph packing problems, with connections to barter markets and set packing explored by Jia et al. [21].

The problem is NP-hard, even with short cycles or chains [1, 8, 23], and is para-NP-hard when parameterized by cycle and path lengths. Approximation algorithms have been proposed via set packing [21, 23].

Practical methods include integer programming [28], used by the UK’s NHS for national donor matching. Dickerson et al. [15] emphasized the need to bound chain lengths for feasibility. Glorie et al. [17] developed a branch-and-price algorithm supporting multi-objective optimization. Li et al. [24] introduced a polynomial-time algorithm for Lorenz-dominant fractional matchings-fair, incentive-compatible, and implementable via lotteries. Approximation methods for short cycles were given by Biro et al. [9], while Riascos-Alvarez et al. [22] proposed a decomposition method for large instances, allowing prioritization between chains and cycles.

Dickerson et al. [15] also introduced the vertex type parameter to capture structural similarity in real graphs. Xiao and Wang [36] built on this with an exact algorithm of complexity  $\mathcal{O}(2^n n^3)$  and an FPT algorithm parameterized by vertex types, assuming unbounded chain and cycle lengths. Maiti and Dey [26] initiated the study of parameterized complexity of the kidney exchange problem and presented a  $\mathcal{O}^*(2^{O(t)})$  time algorithm where  $t$  is the maximum number of patients that receive a kidney; they also showed some other algorithmic and hardness results. Hebert-Johnson et al. [18] improved this running time and presented a randomized algorithm for this problem running in time  $\mathcal{O}^*(4^t)$  and some hardness results.

Lin et al. [25] studied the cycle-only case and proposed a randomized parameterized algorithm based on the number of recipients and maximum cycle length. We refer to [?] for a recent survey on kidney exchange.

## 2 PRELIMINARIES

For an integer  $k$ , we denote the sets  $\{0, 1, \dots, k\}$  and  $\{1, 2, \dots, k\}$  by  $[k]_0$  and  $[k]$  respectively.

A kidney exchange problem is formally represented by a directed graph  $\mathcal{G} = (\mathcal{V}, \mathcal{A})$  which is known as the *compatibility graph*. A subset  $\mathcal{B} \subseteq \mathcal{V}$  of vertices denotes *altruistic donors* (also called *non-directed donors*); the other set  $\mathcal{V} \setminus \mathcal{B}$  of vertices denotes a patient-donor pair who wish to participate in the kidney exchange program. We have a directed edge  $(u, v) \in \mathcal{A}$  if the donor of the vertex  $u \in \mathcal{V}$  has a kidney compatible with the patient of the vertex  $v \in \mathcal{V} \setminus \mathcal{B}$ . Kidney exchange happens either (i) along a *trading-cycle*  $u_1, u_2, \dots, u_k$  where the patient of the vertex  $u_i \in \mathcal{V} \setminus \mathcal{B}$  receives a kidney from the donor of the vertex  $u_{i-1} \in \mathcal{V} \setminus \mathcal{B}$  for every  $2 \leq i \leq k$  and the patient of the vertex  $u_1$  receives the kidney from the donor of the vertex  $u_k$ , or (ii) along a *trading-chain*  $u_1, u_2, \dots, u_k$  where  $u_1 \in \mathcal{B}, u_i \in \mathcal{V} \setminus \mathcal{B}$  for  $2 \leq i \leq k$  and the patient of the vertex  $u_j$  receives a kidney from the donor of the vertex  $u_{j-1}$  for  $2 \leq j \leq k$ . Due to operational reasons, all the kidney transplants along a trading-cycle or a trading-chain should be performed simultaneously. This puts an upper bound on the length  $\ell$  of feasible trading-cycles and trading-chains. We define the length of a path or cycle as the number of edges in it. The kidney exchange clearing problem is to find a collection of feasible trading-cycles and trading-chains that maximizes the number of patients who receive a kidney. Formally, it is defined as follows.

**Definition 2.1 (KIDNEY EXCHANGE).** Given a directed graph  $\mathcal{G} = (\mathcal{V}, \mathcal{A})$  with no self-loops, an altruistic vertex set  $\mathcal{B} \subset \mathcal{V}$ , two integers  $\ell_p$  and  $\ell_c$  denoting the maximum length of respectively paths and cycles allowed, and a target  $t$ , compute if there exists a collection  $\mathcal{C}$  of disjoint cycles of length at most  $\ell_c$  and paths with starting from altruistic vertices only each of length at most  $\ell_p$  which cover at least  $t$  non-altruistic vertices. We denote an arbitrary instance of KIDNEY EXCHANGE by  $(\mathcal{G}, \mathcal{B}, \ell_p, \ell_c, t)$ .

### 2.1 Graph Theoretic Terminologies

In a graph  $\mathcal{G}$ ,  $\mathcal{V}[\mathcal{G}]$  denotes the set of vertices in  $\mathcal{G}$  and  $\mathcal{E}[\mathcal{G}]$  denotes the set of edges in  $\mathcal{G}$ .  $\mathcal{G}[\mathcal{V}']$  denotes the induced subgraph on  $\mathcal{V}'$  where  $\mathcal{V}' \subseteq \mathcal{V}[\mathcal{G}]$ . Two vertices  $u$  and  $v$  in a directed graph  $\mathcal{G}$  are called vertices of the same type if they have the same set of in-neighbors and the same set of out-neighbors. If there are no self-loops in  $\mathcal{G}$ , vertices of the same type form an independent set. Treewidth measures how treelike an undirected graph is. We refer to [14] for an elaborate description of treewidth, tree decomposition, and nice tree decomposition. Since our graph is directed, whenever we mention treewidth of our graph, we refer to the treewidth of the underlying undirected graph; two vertices  $u$  and  $v$  are neighbors of the underlying undirected graph if and only if either there is an edge from  $u$  to  $v$  or from  $v$  to  $u$ . Also refer to Table 1 for the important notations.

### 2.2 Parameterized Complexity

A tuple  $(x, k)$ , where  $k$  is the parameter, is an instance of a parameterized problem. *Fixed parameter tractability* (FPT) refers to solvability in time  $f(k) \cdot p(|x|)$  for a given instance  $(x, k)$ , where  $p$  is a polynomial in the input size  $|x|$  and  $f$  is an arbitrary computable function of  $k$ . We use the notation  $\mathcal{O}^*(f(k))$  to denote  $\mathcal{O}(f(k) \text{poly}(|x|))$ .

We say a parameterized problem is para-NP-hard if it is NP-hard even for some constant values of the parameter.

$\mathcal{B}$	set of altruistic vertices	$t$	target number of patients to receive kidneys
$\ell_p$	length of the longest path allowed	$\ell_c$	length of the longest cycle allowed
$\tau$	treewidth of underlying undirected graph	$\theta$	number of vertex types
$\Delta$	maximum degree of underlying undirected graph	$\ell$	$\max\{\ell_p, \ell_c\}$

Table 2: Notation table.

*Definition 2.2 (Kernelization).* [14] A kernelization algorithm for a parameterized problem  $\Pi \subseteq \Gamma^* \times \mathbb{N}$  is an algorithm that, given  $(x, k) \in \Gamma^* \times \mathbb{N}$ , outputs, in time polynomial in  $|x| + k$ , a pair  $(x', k') \in \Gamma^* \times \mathbb{N}$  such that (a)  $(x, k) \in \Pi$  if and only if  $(x', k') \in \Pi$  and (b)  $|x'|, k' \leq g(k)$ , where  $g$  is some computable function. The output instance  $x'$  is called the kernel, and the function  $g$  is referred to as the size of the kernel. If  $g(k) = k^{O(1)}$ , then we say that  $\Pi$  admits a polynomial kernel.

It is well documented that the presence of a polynomial kernel for certain parameterized problems implies that the polynomial hierarchy collapses to the third level (or, more accurately,  $\text{CoNP} \subseteq \text{NP/Poly}$ ). As a result, polynomial-sized kernels are unlikely to be present in these problems. We use a polynomial-parameter-preserving reduction to demonstrate kernel lower bounds.

### 3 A FASTER DETERMINISTIC FPT PARAMETERIZED BY $t$

In this section, we design the fastest known deterministic FPT algorithm for the problem parameterized by  $t$ , consisting of three main steps. In the first step, we apply a color-coding procedure that ensures every vertex involved in the solution receives a distinct color which plays a crucial role in the second step, where we design a dynamic programming (DP) algorithm that systematically checks whether a solution exists in which exactly one vertex from each color subset is selected. During the gradual construction of entries corresponding to larger color sets, we use as a subroutine, a slight modified version of the standard DP algorithm [14] that runs in  $2^{|S|}$  time to determine whether a cycle or altruistic path exists that uses all colors in  $S$  exactly once. The formal algorithm is presented below.

#### 3.1 Color Coding

We employ the color-coding method of Alon, Yuster, and Zwick [2]. In this technique, each vertex of  $G$  is assigned one of  $t$  colors uniformly at random. With probability at least  $e^{-t}$ , the  $t$  vertices forming the desired solution all receive distinct colors. This randomization coloring can be eliminated (made deterministic) by using families of  $t$ -perfect hash functions of size  $e^t \cdot t^{O(\log t)}$  [29]. Henceforth, we assume access to such an  $(n, t)$ -perfect hash family. By iterating over all colorings in the family, we may assume without loss of generality that in some coloring, the vertices of the solution are colored distinctly. For every coloring in the family we employ the following DP algorithm that checks if a colorful solution is possible for the KIDNEY EXCHANGE instance:

#### 3.2 Dynamic Programming Formulation

Let  $[t] = \{1, 2, \dots, t\}$  denote the set of colors. For every subset  $T' \subseteq [t]$ , define

$$DP(T') = \begin{cases} 1 & \text{there exists a collection of vertex-disjoint cycles and altruistic paths that uses exactly one vertex from each color in } T'. \\ 0 & \text{otherwise} \end{cases}$$

The recurrence is as follows:

$$DP(T') = 1 \iff \exists T'' \subset T' \text{ such that } DP(T'') = 1 \text{ and either}$$

- there is a directed cycle with exactly one vertex from each color in  $T' \setminus T''$ , or
- there is an altruistic path with exactly one vertex from each color in  $T' \setminus T''$ .

For finding such cycles/altruistic paths in time  $2^{|T' \setminus T''|}$ , we use the subsequent step.

#### 3.3 Detecting Cycles and Altruistic Paths of Prescribed Colors

The task of checking whether there exists a cycle or altruistic path using exactly a set of colors  $S \subseteq [t]$  can be solved via a slight modification of the standard color-coding dynamic programming algorithm for detecting colorful paths [14]. The standard DP algorithm runs in  $2^{|S|}$  time and determines whether a directed path exists that uses all colors exactly once. For an altruistic path, the DP formulation is restricted so that the path must start at an altruistic vertex, while the rest of the algorithm remains unchanged. For a colorful cycle, one can guess an arc and then, in the DP algorithm, search for a path that uses all colors exactly once and starts and ends at the endpoints of the guessed arc. Thus these algorithms run in time  $O^*(2^{|S|})$  and can be used to determine required colorful altruistic paths/cycles.

#### 3.4 Run Time Analysis

Given a colored graph (with  $t$  colors), the DP algorithm computes all DP entries for all subsets  $T' \subseteq [t]$ . For each  $T'$ , we consider all partitions  $T'' \subseteq T'$  where we look for a colorful solution that must use all colors in  $T''$  and check for an additional colorful cycle/altruistic path using all colors in  $T' \setminus T''$  (which takes time  $2^{|T' \setminus T''|}$ ). Hence, the total running time is therefore bounded by

$$\sum_{T' \subseteq [t]} \sum_{T'' \subseteq T'} 2^{|T' \setminus T''|} \cdot n^{O(1)}$$

We now simplify the double sum. For each  $i = |T'|$ , there are  $\binom{t}{i}$  choices of  $T'$ , and for each such  $T'$ :

$$\sum_{T'' \subseteq T'} 2^{|T' \setminus T''|} = \sum_{j=0}^i \binom{i}{j} \cdot 2^{i-j} = 3^i.$$

Thus the total runtime to compute the DP table is

$$\sum_{i=0}^t \binom{t}{i} \cdot 3^i = (1+3)^t = 4^t.$$

Hence provided a colored input graph, the algorithm checks if the graph has a solution of size  $t$  in time  $O^*(4^t)$  time. Using families of  $t$ -perfect hash functions of size  $e^t \cdot t^{O(\log t)}$  [29] (these many instances of colorful graphs) for coloring, we can bound the overall running time of our algorithm by

$$O^*(e^t \cdot t^{O(\log t)} \cdot 4^t).$$

**THEOREM 1.** *KIDNEY EXCHANGE parameterized by  $t$ , the number of patients who can be helped, admits a deterministic FPT algorithm running in time  $O^*((4e)^t) \approx O^*(10.88^t)$ .*

#### 4 NO POLY KERNEL PARAMETERIZED BY

$$t + \ell_p + \ell_c (+|\mathcal{B}|)$$

In this section, we provide a polynomial-time, parameter-preserving reduction from DIRECTED  $k$ -PATH, which asks whether a directed path of length  $k$  exists and is known not to admit a polynomial compression (see Proposition 4.1 below) under standard complexity theoretic assumptions, to an equivalent instance of KIDNEY EXCHANGE, thereby establishing the desired hardness result.

**PROPOSITION 4.1.** [12] *DIRECTED  $k$ -PATH (Long Directed Path parameterized by  $k$ ) admits neither a polynomial kernel nor a polynomial compression unless  $\text{NP} \subseteq \text{coNP/poly}$ .*

**Reduction:** Given an input instance  $(G = (V, E), k)$  of DIRECTED  $k$ -PATH, we create an equivalent instance  $(\mathcal{G} = (V \cup \{a\}, E \cup \{(a, v) \mid v \in V\}), \mathcal{B} = \{a\}, \ell_p = k+1, \ell_c = 0, t = k+1)$  of KIDNEY EXCHANGE, i.e., we add an additional single altruistic vertex  $a$  with arcs from  $a$  to every vertex  $v \in V$ . We set  $\ell_c = 0, \ell_p = k+1$ , and require  $t = k+1$  transplants.

**LEMMA 4.2.**  *$(G = (V, E), k)$  of DIRECTED  $k$ -PATH is a yes instance if and only if  $(\mathcal{G} = (V \cup \{a\}, E \cup \{(a, v) \mid v \in V\}), \mathcal{B} = \{a\}, \ell_p = k+1, \ell_c = 0, t = k+1)$  is a yes instance of KIDNEY EXCHANGE.*

**PROOF.** In the forward direction, suppose  $(G, k)$  is a yes instance and contains a directed simple path  $P = v_1, v_2, \dots, v_k$  on  $k$  vertices. Then in  $\mathcal{G}$ ,  $P' = a, v_1, v_2, \dots, v_k$  is an altruistic path of length  $k+1$  making  $(\mathcal{G}, \mathcal{B}, \ell_p, \ell_c, t)$  a yes instance of KIDNEY EXCHANGE.

Conversely, suppose the constructed KIDNEY EXCHANGE instance is a yes instance. Since  $\ell_c = 0$ , no cycles are allowed, and because there is exactly one altruistic vertex  $a$ , the solution must contain a single altruistic path starting at  $a$  of length  $k+1$ , necessarily of the form  $P' = a, v_1, v_2, \dots, v_k$ , where the  $v_i$  are distinct vertices in  $V$ . Removing the initial vertex  $a$  yields the directed simple path  $P = v_1, v_2, \dots, v_k$  in  $G$  making  $(G, k)$  a yes instance.  $\square$

Since  $k \in O^*(t + \ell_p + \ell_c + |\mathcal{B}|)$ , we immediately get the following as a corollary.

**THEOREM 2.** *Under the assumption  $\text{NP} \not\subseteq \text{coNP/poly}$ , KIDNEY EXCHANGE admits no polynomial kernel (nor polynomial compression) parameterized by  $t + \ell_p + \ell_c + |\mathcal{B}|$  where  $\mathcal{B}$  is the set of altruistic vertices.*

#### 5 W[1]-HARD NESS PARAMETERIZED BY PATHWIDTH

In this section, we present our claimed hardness results, showing that KIDNEY EXCHANGE remains W[1]-hard when parameterized by either pathwidth $+\ell_c$  or pathwidth $+\ell_p$ . Specifically, we provide these two hardness results for KIDNEY EXCHANGE separately: one for bounded-pathwidth DAGs (where  $\ell_c = 0$ ) and another for bounded-pathwidth graphs with no altruistic vertices (where  $\ell_p = 0$ ). These results follow from the known W[1]-hard of the BIN PACKING problem with unary-encoded weights [20]. Note that, Maiti et al. [26] showed that KIDNEY EXCHANGE is FPT when parameterized by the combined parameter of *treewidth*,  $\ell_p$ , and  $\ell_c$ , and consequently also FPT when parameterized by the combined parameter of *pathwidth*,  $\ell_p$ , and  $\ell_c$ . Together with our result, this provides a complete characterization of the parameterized complexity of KIDNEY EXCHANGE with respect to *pathwidth*,  $\ell_p$ , and  $\ell_c$ . We begin with the formal definition of the BIN PACKING problem, followed by our claimed parameter-preserving reductions to KIDNEY EXCHANGE.

**Definition 5.1** (BIN PACKING). Given a set  $\mathcal{U}$  of  $n$  items with positive integral sizes  $w_1, \dots, w_n$  and  $k$  bins each of size  $\frac{W}{k}$  where  $W = \sum_{i=1}^n w_i$ , compute if there exists a partition of the items into  $k$  parts such that the total size of every part is exactly  $\frac{W}{k}$ . We denote an arbitrary instance of BIN PACKING by  $(\mathcal{U}, k)$ .

We know that BIN PACKING is W[1]-hard parameterized by  $k$  even when all the numbers in the input are encoded in unary [20]. The high-level idea of our reduction is as follows. Each item from the bin packing becomes a small chain of donors and recipients, and each bin becomes a longer chain that starts from an altruistic donor. We connect these parts in such a way that deciding which item chains attach to which altruistic chain is exactly the same as deciding which items go into which bins. If the items can be evenly packed into bins, then the corresponding kidney exchange instance also has a perfect matching of transplants, and vice versa. Because our construction maintains the acyclicity of the graph and does not significantly increase its structural complexity, this demonstrates that the kidney exchange problem remains challenging even for graphs with low pathwidth. We now present the formal proof.

**THEOREM 3.** *KIDNEY EXCHANGE is W[1]-hard when parameterized by the undirected pathwidth of the input graph, even when the graph is acyclic.*

**PROOF.** We reduce from BIN PACKING parameterized by  $k$ . Let  $(\mathcal{U} = \{w_i : i \in [n]\}, k)$  be an arbitrary instance of BIN PACKING. We assume without loss of generality that  $w_i \geq 3k^2n^2$  for every  $i \in [n]$ ; if it is not the case, then we multiply  $w_i$  with  $3k^2n^2$  for every  $i \in [n]$ . We now construct an instance  $(\mathcal{G}, \mathcal{B}, \ell_p, \ell_c, t)$  of KIDNEY EXCHANGE. For every  $w_i \in \mathcal{U}$ , we have a directed path  $p_i$  starting from  $u_{i,s}$  to  $u_{i,\lambda}$  consisting of  $w_i$  vertices. We also have a directed path  $q_j = (a_{j,1}, \dots, a_{j,n})$  for every  $j \in [k]$ . Other than the edges in the above paths, we also have the following edges. For every  $i \in [n]$ ,

we have an edge from  $a_{j,i}$  to  $u_{i,s}$ . We also have an edge from  $u_{i,\lambda}$  to  $a_{j,i+1}$  for every  $i \in [n-1]$ . This finishes the description of  $\mathcal{G}$ . We note that  $\mathcal{G}$  is an acyclic graph with a following topological order

$$a_{1,1}, \dots, a_{k,1}, \vec{\mathcal{V}}(p_1), a_{1,2}, \dots, a_{k,2}, \vec{\mathcal{V}}(p_2), \dots, \vec{\mathcal{V}}(p_n)$$

where  $\vec{\mathcal{V}}(p_i)$  is the vertices in  $p_i$  ordered as the directed path  $p_i$  from  $u_{i,s}$  to  $u_{i,\lambda}$  for  $i \in [n]$ . The set  $\mathcal{B}$  of altruistic vertices is  $\{a_{j,1} : j \in [k]\}$ . We set  $\ell_p = \frac{W}{k} + n$  where  $W = \sum_{i=1}^n w_i$  and  $\ell_c$  arbitrarily – value of  $\ell_c$  does not matter since the graph of the KIDNEY EXCHANGE instance is acyclic. We define  $t = |\mathcal{V}(\mathcal{G}) \setminus \mathcal{B}|$ , that is, the target is to cover every non-altruistic vertices in  $\mathcal{G}$ . We now claim that the two instances are equivalent.

In one direction, let us assume that the BIN PACKING instance is a YES instance. Let  $(X_1, \dots, X_k)$  be a partition of  $\mathcal{U}$  such that the total size of items in  $X_j$  is  $\frac{W}{k}$  for every  $j \in [k]$ . For every  $j \in [k]$ , we define a path  $\mathcal{P}_j$  as

$$\mathcal{P}_j = a_{j,1}, Q_1, a_{j,2}, Q_2, \dots, a_{j,n}, Q_n$$

where  $Q_i$  is  $\vec{\mathcal{V}}(p_i)$  if  $w_i \in X_j$  and the empty set otherwise. We first observe that, since  $(X_1, \dots, X_k)$  is a partition of  $\mathcal{U}$ ,  $(\mathcal{V}(\mathcal{P}_1), \dots, \mathcal{V}(\mathcal{P}_k))$  is a partition of  $\mathcal{V}(\mathcal{G})$ . Also, since the total size of items in  $X_j$  is  $\frac{W}{k}$ , each  $\mathcal{P}_j$  contains  $\frac{W}{k} + n$  vertices which is  $\ell_p$ . Hence, the KIDNEY EXCHANGE instance is a YES instance.

In the other direction, let us assume that the KIDNEY EXCHANGE instance is a YES instance. Since  $\ell_p = \frac{W}{k} + n, t = |\mathcal{V}(\mathcal{G}) \setminus \mathcal{B}|$ , and  $|\mathcal{B}| = k$ , there exist  $k$  disjoint paths  $\mathcal{P}_1, \dots, \mathcal{P}_k$  in  $\mathcal{G}$  each of length  $\ell_p = \frac{W}{k} + n$  covering all the vertices of  $\mathcal{G}$ . For every  $i \in [n]$ , we observe that each path  $\mathcal{P}_j, j \in [k]$  must either contain all the vertices of the path  $p_i$  or none of its vertices, since  $\mathcal{P}_j$  can neither end inside  $p_i$  as then the next vertex cannot be covered by any other path nor move from inside  $p_i$  to outside  $p_i$  as there is no such edge in  $\mathcal{G}$ . We now define  $X_j = \{w_i : i \in [n], \mathcal{P}_j \text{ contains the path } p_i\}$  for  $j \in [k]$ . We claim that  $\sum_{x \in X_j} x = \frac{W}{k}$ . Indeed, otherwise  $\mathcal{P}_j$  covers either at most  $\frac{W}{k} + n^2 - 3k^2n^2$  vertices or at least  $\frac{W}{k} + n^2 + 3k^2n^2$  vertices since  $w_i \geq 3k^2n^2$  for every  $i \in [n]$ . We now refute both possibilities, thereby proving the claim. The path  $\mathcal{P}_j$  cannot cover at least  $\frac{W}{k} + n^2 + 3k^2n^2$  vertices since  $\ell_p = \frac{W}{k} + n$ . Also, the path  $\mathcal{P}_j$  cannot cover at most  $\frac{W}{k} + n^2 - 3k^2n^2$  vertices, since then the remaining  $k-1$  paths each of length at most  $\ell_p$  cannot cover all the non-altruistic vertices in  $\mathcal{G}$ . Hence, the KIDNEY EXCHANGE instance is a YES instance too.

Finally, we observe that the pathwidth of  $\mathcal{G}$  is at most  $k-1$  since the following sequence of bags each containing at most  $k$  vertices is a valid path decomposition of  $\mathcal{G}$ .

$$\{a_{1,1}, \dots, a_{k,1}\}, \vec{\mathcal{V}}(p_1), \{a_{1,2}, \dots, a_{k,2}\}, \vec{\mathcal{V}}(p_2), \dots, \vec{\mathcal{V}}(p_n)$$

Now the result follows since BIN PACKING is  $W[1]$ -hard parameterized by  $k$ .  $\square$

We can modify the reduced instance of Theorem 3 to show that KIDNEY EXCHANGE is  $W[1]$ -hard with respect to pathwidth even when there is no altruistic vertex. We add an edge from the last vertex of the path  $p_n$  to each altruistic vertex. Everything else remains the same. We present the detailed proof below.

**THEOREM 4.** *KIDNEY EXCHANGE is  $W[1]$ -hard when parameterized by the undirected pathwidth of the input graph, even when no altruistic vertex is present.*

**PROOF.** From an arbitrary instance of BIN PACKING, we construct a KIDNEY EXCHANGE instance exactly as in the proof of Theorem 3 except we also add an edge from  $u_{n,\lambda}$  (which is the last vertex of the path  $p_n$ ) to  $a_{j,1}$  for every  $j \in [k]$ , there is no altruistic vertex,  $\ell_c = \frac{W}{k} + n$ ,  $\ell_p$  is anything (since there is no altruistic vertex, the value of  $\ell_p$  is irrelevant), and  $t = \mathcal{V}(\mathcal{G})$ . We now prove equivalence the KIDNEY EXCHANGE instance with the BIN PACKING instance.

In one direction, let us assume that the BIN PACKING instance is a YES instance. Let  $(X_1, \dots, X_k)$  be a partition of  $\mathcal{U}$  such that the total size of items in  $X_j$  is  $\frac{W}{k}$  for every  $j \in [k]$ . For every  $j \in [k]$ , we define a cycle  $C_j$  as

$$C_j = a_{j,1}, Q_1, a_{j,2}, Q_2, \dots, a_{j,n}, Q_n, a_{j,1}$$

where  $Q_i$  is  $\vec{\mathcal{V}}(p_i)$  if  $w_i \in X_j$  and the empty set otherwise. We first observe that, since  $(X_1, \dots, X_k)$  is a partition of  $\mathcal{U}$ ,  $(\mathcal{V}(C_1), \dots, \mathcal{V}(C_k))$  is a partition of  $\mathcal{V}(\mathcal{G})$ . Also, since the total size of items in  $X_j$  is  $\frac{W}{k}$ , each  $C_j$  contains  $\frac{W}{k} + n$  vertices which is  $\ell_c$ . Hence, the KIDNEY EXCHANGE instance is a YES instance.

In the other direction, let us assume that the KIDNEY EXCHANGE instance is a YES instance. Since  $\ell_c = \frac{W}{k} + n, t = |\mathcal{V}(\mathcal{G}) \setminus \mathcal{B}|$ , and  $|\mathcal{B}| = k$ , there exist  $k$  disjoint cycles  $C_1, \dots, C_k$  in  $\mathcal{G}$  each of length  $\ell_c = \frac{W}{k} + n$ . For every  $i \in [n]$ , we observe that each cycle  $C_j, j \in [k]$  must either contain all the vertices of the path  $p_i$  or none of its vertices, since  $C_j$  can neither end inside  $p_i$  as then the next vertex cannot be covered by any other path nor move from inside  $p_i$  to outside  $p_i$  as there is no such edge in  $\mathcal{G}$ . We now define  $X_j = \{w_i : i \in [n], C_j \text{ contains the path } p_i\}$  for  $j \in [k]$ . We claim that  $\sum_{x \in X_j} x = \frac{W}{k}$ . Indeed, otherwise  $\mathcal{P}_j$  covers either at most  $\frac{W}{k} + n^2 - 3k^2n^2$  vertices or at least  $\frac{W}{k} + n^2 + 3k^2n^2$  vertices since  $w_i \geq 3k^2n^2$  for every  $i \in [n]$ . We now refute both the possibilities, thereby proving the claim. The cycle  $C_j$  cannot cover at least  $\frac{W}{k} + n^2 + 3k^2n^2$  vertices since  $\ell_c = \frac{W}{k} + n$ . Also, the cycle  $C_j$  cannot cover at most  $\frac{W}{k} + n^2 - 3k^2n^2$  vertices, since then the remaining  $k-1$  cycles each of length at most  $\ell_c$  cannot cover all the non-altruistic vertices in  $\mathcal{G}$ . Hence, the KIDNEY EXCHANGE instance is a YES instance too.

Finally, we observe that the pathwidth of  $\mathcal{G}$  is at most  $2k-1$  since the following sequence of bags each containing at most  $k$  vertices is a valid path decomposition of  $\mathcal{G}$  where  $\mathcal{A} = \{a_{1,1}, \dots, a_{k,1}\}$  and  $\mathcal{A} \cup \vec{\mathcal{V}}(p_i) := \mathcal{A} \cup \{u_{i,s}\}, \dots, \mathcal{A} \cup \{u_{i,\lambda}\}$  for  $i \in [n]$ .

$$\{a_{1,1}, \dots, a_{k,1}\}, \mathcal{A} \cup \vec{\mathcal{V}}(p_1), \mathcal{A} \cup \{a_{1,2}, \dots, a_{k,2}\}, \mathcal{A} \cup \vec{\mathcal{V}}(p_2), \dots, \mathcal{A} \cup \vec{\mathcal{V}}(p_n)$$

Now the result follows since BIN PACKING is  $W[1]$ -hard parameterized by  $k$ .  $\square$

## 6 NP-HARDNESS ON DAGS

To establish NP-hardness in our setting, we first reduce a classical 3-PARTITION instance into an equivalent instance of FIXED-SIZE-3-PARTITION, where each feasible group is required to contain exactly

three elements that collectively achieve a specified *target sum*. This reduction is crucial for our subsequent reduction to KIDNEY EXCHANGE, in which altruistic paths of fixed lengths in the constructed solution correspond precisely to the groups attaining the *target sum*.

### 6.1 Reduction from 3-PARTITION to FIXED-SIZE-3-PARTITION

In the standard 3-PARTITION problem, we are given a multiset  $A = \{a_1, \dots, a_{3m}\}$  of positive integers and a target value  $B$ . The task is to decide whether  $A$  can be split into exactly  $m$  disjoint triples whose elements each sum to  $B$ . This problem remains NP-hard even under strong assumptions such as  $B/4 < a_i < B/2$  and for constant values of  $B$ . For our reduction, we do not rely on any such promise; instead, we enforce triple cardinality groups via an additive shift.

*Definition 6.1.* (FIXED-SIZE-3-PARTITION) In this problem, we are given a multiset  $A = \{a_1, \dots, a_{3m}\}$  of positive integers and a target value  $B$ , with the additional property that any subset of elements summing to  $B$  contains exactly three elements. The question is whether there exists a partition of  $A$  into  $m$  disjoint groups, each having a total sum equal to  $B$ .

Below, we provide a polynomial-time reduction from 3-PARTITION to FIXED-SIZE-3-PARTITION, thereby establishing the NP-hardness of the latter problem.


The key idea is to enforce that every feasible group contains exactly three elements by applying a uniform additive shift to all item values. By choosing a constant  $C$  much larger than any of the original values, any subset of fewer than three elements cannot reach the new target sum (which is larger than  $3C$  and smaller than  $4C$ ), while any subset of more than three elements necessarily exceeds it. Thus, only groups of size three can potentially satisfy the target constraint, and among such groups, feasibility still depends on the specific magnitudes of the  $a_i$ . Once three elements are selected, subtracting the shift restores the original sum, ensuring that each valid group corresponds precisely to a triple in the original instance. The formal construction and subsequent arguments are presented below.

**Construction.** Given an instance  $(A, B)$  of 3-PARTITION, where  $A = \{a_1, \dots, a_{3m}\}$ , we define a constant  $C := 10B$  and construct an equivalent instance  $(A' = \{a'_1, a'_2, \dots, a'_{3m}\}, B' = 3C + B)$  of FIXED-SIZE-3-PARTITION where  $a'_i := a_i + C$  for  $i \in [3m]$ .

**LEMMA 6.2.**  $(A, B)$  is a yes-instance of 3-PARTITION if and only if  $(A', B')$  is a yes-instance of FIXED-SIZE-3-PARTITION.

**PROOF.** Before presenting the proof, we note that since all  $a_i$  are positive integers, the input instance of 3-PARTITION cannot contain any element  $a_i > B$ . With this observation, consider any group of size at most 2 in the constructed instance of FIXED-SIZE-3-PARTITION, can have a maximum sum of  $2B + 2C < 3C < B'$ . On the other hand, any group containing four elements has a minimum possible sum of  $4C > B'$ , which further implies that  $k_r \leq 3$ . Hence, the size of any group achieving the target sum  $B'$  must be exactly 3.

**(If).** Suppose  $(A, B)$  is a yes-instance of 3-PARTITION and  $A$  can be partitioned into  $m$  disjoint triples  $\{\{a_{r_1}, a_{r_2}, a_{r_3}\}\}_{r \in [m]}$  with  $\sum_{j=1}^3 a_{r_j} = B$ . Then for each such triple the corresponding triples



**Figure 1: Reduction from FIXED-SIZE-3-PARTITION problem**

of  $\{\{a'_{r_1}, a'_{r_2}, a'_{r_3}\}\}_{r \in [m]}$  forms a valid partition of  $A'$  into  $m$  groups since  $\sum_{j=1}^3 a'_{r_j} = \sum_{j=1}^3 (a_{r_j} + C) = B + 3C = B'$ , making  $(A', B')$  a yes-instance of FIXED-SIZE-3-PARTITION.

**(Only if).** Conversely, suppose  $(A', B')$  is a yes-instance of FIXED-SIZE-3-PARTITION and  $A'$  admits a partition into  $m$  groups each of size 3 as argued in the beginning part of our proof,  $\{\{a'_{r_1}, a'_{r_2}, a'_{r_3}\}\}_{r \in [m]}$  each with sum  $B' = 3C + B$ . But then  $\{\{a_{r_1}, a_{r_2}, a_{r_3}\}\}_{r \in [m]}$  forms a valid partition of  $A$  into  $m$  groups since  $\sum_{j=1}^3 a_{r_j} = \sum_{j=1}^3 (a'_{r_j} - C) = B' - 3C = B$ , making  $(A, B)$  is a yes-instance of 3-PARTITION.  $\square$

Since the magnitudes of all elements including the target sum  $B'$  increase by at most a factor linear in  $B$ , the reduction is polynomial in the input size. And 3-PARTITION being NP-hard, immediately implies the following.

**THEOREM 6.3.** FIXED-SIZE-3-PARTITION is NP-hard even for constant values of  $B'$ .

### 6.2 Reduction from FIXED-SIZE-3-PARTITION to KIDNEY EXCHANGE on DAGs

We next prove that the KIDNEY EXCHANGE problem on directed acyclic graphs is NP-hard, even for a fixed(constant) value of  $\ell_p$  by showing a polynomial reduction from FIXED-SIZE-3-PARTITION problem.

**Construction.** Let  $(A' = \{a'_1, \dots, a'_{3m}\}, T')$  be an instance of FIXED-SIZE-3-PARTITION. We construct an equivalent instance  $(\mathcal{G}, \mathcal{B}, \ell_p, \ell_c, t)$  of KIDNEY EXCHANGE as follows.

- For each  $a'_i \in A'$ , create a directed (element-)path  $P_i$  of  $a'_i$  many vertices (see Figure 1).
- Add  $m$  altruistic vertices  $\mathcal{B} = \{s_1, \dots, s_m\}$ . From each  $s_j$ ,  $j \in [m]$ , add an outgoing edge/arc to the first vertex of every  $P_i$ ,  $i \in [3m]$ .
- Index the paths  $P_1, \dots, P_{3m}$  and, for every  $i < j$ , add an outgoing edge/arc (concatenation arcs) from the last vertex of  $P_i$  to the first vertex of  $P_j$ . This preserves acyclicity and provides an opportunity for paths to concatenate.
- Set  $\ell_p := B' + 1$  and  $t = |V(\mathcal{G})| = \sum_{i \in [3m]} a'_i + m$ . We consider vertex-disjoint altruistic paths of length at most  $\ell_p$  and are interested in solutions that cover all vertices of  $\mathcal{G}$ .

**LEMMA 6.4.**  $(A', B')$  is a yes-instance of FIXED-SIZE-3-PARTITION if and only if  $(\mathcal{G}, \mathcal{B}, \ell_p, \ell_c, t)$  is a yes instance of KIDNEY EXCHANGE.

**PROOF. (If).** Suppose  $(A', B')$  is a yes-instance of FIXED-SIZE-3-PARTITION where  $A'$  can be partitioned into  $m$  disjoint triples  $\{a'_{r_1}, a'_{r_2}, a'_{r_3}\}_{r \in [m]}$  each summing to  $T'$ . For each triple  $\{a'_{r_1}, a'_{r_2}, a'_{r_3}\}$  (where we assume w.l.o.g. that  $r_1 < r_2 < r_3$ ), we construct an altruistic path in solution of length exactly  $\ell_p$  that starts at  $s_r$  followed by the paths  $P_{r_1}$ ,  $P_{r_2}$ , and  $P_{r_3}$ . These  $m$  altruistic paths together cover all vertices while each being of length exactly  $\ell_p$ .

**(Only if).** Conversely, suppose  $(\mathcal{G}, \mathcal{B}, \ell_p, \ell_c, t)$  is a yes-instance of KIDNEY EXCHANGE, admitting a collection of  $m$  vertex-disjoint altruistic paths that together cover all vertices of  $\mathcal{G}$ . Note that each such altruistic path must be of length exactly  $\ell_p$  in order to cover all vertices of the graph. Moreover, each altruistic path begins at an altruistic vertex  $s_r$  and immediately enters some element-path  $P_i$  at its first vertex, since the only outgoing edges from altruistic vertices lead to first/initial vertices of the paths constructed from elements. By construction, each element-path  $P_i$  is a directed path with only incoming edges to its first vertex, and only outgoing edges from its last vertex. Consequently, an altruistic path can only traverse entire element-paths; it cannot enter or exit from an internal vertex of an element-path. Additionally, no altruistic path can end at any vertex other than the last vertex of an element-path, since neither the internal vertices nor the last vertex can be covered by any other altruistic path. In particular, this ensures that the element-paths are assigned integrally/completely to the altruistic paths. Finally, the concatenation arcs enforce that any altruistic path traverses element-paths in increasing index order. The integral/complete assignment of element paths to altruistic paths along with the fact that every altruistic path has length exactly  $\ell_p = 1 + B'$  enforces that every altruistic path contains precisely 3 element paths. Therefore the altruistic paths induce a partition of the index set  $\{1, \dots, 3m\}$  into exactly  $m$  groups (one group per altruistic path), where a group corresponding to an altruistic path consists of the indices of the element-paths it traverses in increasing order. This implies that the sum total of all non-altruistic vertices in any such group of size 3 is exactly equal to  $B'$ . Hence this partition of  $3m$  indices also gives a partition of  $A'$  into  $m$  groups each summing up to  $B'$ , making  $(A', B')$  a yes instance of FIXED-SIZE-3-PARTITION.  $\square$

**Complexity.** The reduction introduces  $O(\sum_i a'_i)$  vertices and  $O((\sum_i a'_i)^2)$  edges, both polynomial in the input size. And,  $\ell_p$  being  $B' + 1$  implies the following.

**THEOREM 5.** *The KIDNEY EXCHANGE problem on DAGs is NP-hard even when restricted to instances with a fixed constant path-length bound  $\ell_p$ .*

**Corollary 1.** *The KIDNEY EXCHANGE problem is para-NP-hard with respect to  $DFVS + \ell_p + \ell_c$ .*

## 7 CONCLUSION

In this work, we advanced the theoretical understanding of the Kidney Exchange problem by establishing both faster parameterized algorithms and tighter lower bounds. Specifically, we presented the fastest known deterministic fixed-parameter tractable (FPT) algorithm parameterized by the number of patients helped, improving the previous best running time from  $O^*(14^t)$  to  $O^*((4e)^t) \approx O^*(10.88^t)$ . Our approach combines color-coding, derandomized

through  $t$ -perfect hash families, with a subset-based dynamic programming formulation that systematically constructs feasible cycles and altruistic chains. This yields a clean, purely combinatorial algorithm that is efficient, transparent, and asymptotically optimal among deterministic approaches of this kind.

Beyond algorithmic improvements, we showed that the problem admits no polynomial kernel when parameterized by  $t + \ell_p + \ell_c + |\mathcal{B}|$ , unless  $\text{NP} \not\subseteq \text{coNP/poly}$ , thus ruling out the possibility of efficient instance compression under standard complexity assumptions. We further strengthened the known hardness results by proving that Kidney Exchange is  $W[1]$ -hard parameterized by pathwidth, even under strong structural restrictions such as acyclicity or the absence of altruistic donors. This settles a natural open question regarding the tractability of the problem on graphs of bounded pathwidth. Complementing these results, we also demonstrated para-NP-hardness with respect to the combined parameters  $DFVS, \ell_p, \ell_c$ , thereby completing a near-exhaustive map of the problem's parameterized complexity landscape.

Taken together, our findings delineate a sharp boundary between tractable and intractable regimes for Kidney Exchange under a variety of structural and quantitative parameters. From a broader perspective, these results strengthen the theoretical foundations of algorithmic kidney exchange, clarifying what forms of efficiency are achievable through parameterization and what barriers are provably inherent.

There are several promising avenues for future research. First, an immediate direction is to further improve the running time of FPT algorithms, either by tightening constants in the exponent or by exploring alternative algorithmic paradigms such as representative sets or inclusion–exclusion–based methods. Another direction is to investigate lower bounds under the Exponential Time Hypothesis (ETH) or Strong ETH, which could clarify whether our current algorithm is asymptotically optimal. It would also be valuable to reduce the polynomial dependence on the input size  $n$ , as this term often dominates in practical settings.

From a modeling perspective, our work focuses on the classical cycle-and-chain formulation. Extending these techniques to richer and more realistic models—including weighted or utility-based exchanges, dynamic or online formulations, and fairness- or priority-aware variants—remains an important open challenge. Such extensions would bridge the gap between parameterized theory and the design of scalable, value-aligned kidney exchange algorithms used in national registries. Finally, exploring kernelization-inspired preprocessing heuristics or hybrid parameterizations that combine structural and solution-size parameters may provide new insights into balancing theoretical rigor with real-world efficiency.

## ACKNOWLEDGEMENT

Aritra Banik acknowledges support from the Anusandhan National Research Foundation (ANRF) (erstwhile Science, Education, and Research Board (SERB)), Government of India, via the project MTR/2022/000253. Palash Dey thanks the Anusandhan National Research Foundation (ANRF) (erstwhile Science, Education, and Research Board (SERB)), Government of India, for supporting this work through Core Research Grant under file no. CRG/2022/003294.

## REFERENCES

- [1] David J. Abraham, Avrim Blum, and Tuomas Sandholm. 2007. Clearing algorithms for barter exchange markets: enabling nationwide kidney exchanges. In *Proc. 8th ACM Conference on Electronic Commerce (EC-2007)*, Jeffrey K. MacKie-Mason, David C. Parkes, and Paul Resnick (Eds.). ACM, 295–304. <https://doi.org/10.1145/1250910.1250954>
- [2] Noga Alon, Raphael Yuster, and Uri Zwick. 1995. Color-Coding. *J. ACM* 42, 4 (1995), 844–856. <https://doi.org/10.1145/210332.210337>
- [3] Alvin Roth, Tayfun Sönmez, Utku Ünver. 2004. Kidney exchange. *Quarterly Journal of Economics* 119, 2 (2004), 457–488.
- [4] Alvin Roth, Tayfun Sönmez, Utku Ünver. 2005. A kidney exchange clearinghouse in new england. *Am. Econ. Rev* 95, 2 (2005), 376–380.
- [5] Alvin Roth, Tayfun Sönmez, Utku Ünver. 2005. Pairwise Kidney exchange. *J. Econ. Theory* 125, 2 (2005), 151–188.
- [6] Alvin Roth, Tayfun Sönmez, Utku Ünver. 2007. Efficient kidney exchange: Co-occurrence of wants in markets with compatibility-based preferences. *Am. Econ. Rev* (2007), 828–851.
- [7] Itai Ashlagi and Alvin E Roth. 2021. Kidney exchange: An operations perspective. *Management Science* 67, 9 (2021), 5455–5478.
- [8] Rémy Belmonte, Teshu Hanaka, Masaaki Kanzaki, Masashi Kiyomi, Yasuaki Kobayashi, Yusuke Kobayashi, Michael Lampis, Hirotaka Ono, and Yota Otachi. 2022. Parameterized Complexity of (A, 1)-Path Packing. *Algorithmica* 84, 4 (2022), 871–895. <https://doi.org/10.1007/S00453-021-00875-Y>
- [9] Péter Biro, David F Manlove, and Romeo Rizzi. 2009. Maximum weight cycle packing in directed graphs, with application to kidney exchange programs. *Discrete Math Algorithms Appl* 1, 04 (2009), 499–517.
- [10] Péter Biro, Joris Van de Klundert, David Manlove, William Pettersson, Tommy Andersson, Lisa Burnapp, Pavel Chromy, Pablo Delgado, Piotr Dworczak, Bernadette Haase, et al. 2021. Modelling and optimisation in European kidney exchange programmes. *European Journal of Operational Research* 291, 2 (2021), 447–456.
- [11] Hans L. Bodlaender, Bart M. P. Jansen, and Stefan Kratsch. 2013. Kernel bounds for path and cycle problems. *Theor. Comput. Sci.* 511 (2013), 117–136. <https://doi.org/10.1016/J.TCS.2012.09.006>
- [12] Hans L. Bodlaender, Bart M. P. Jansen, and Stefan Kratsch. 2014. Kernelization Lower Bounds by Cross-Composition. *Theory of Computing Systems* 54, 3 (2014), 441–456.
- [13] Miguel Constantino, Xenia Klimentova, Ana Viana, and Abdur Rais. 2013. New insights on integer-programming models for the kidney exchange problem. *Eur. J. Oper. Res.* 231, 1 (2013), 57–68.
- [14] Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. 2015. *Parameterized Algorithms*. Springer. <https://doi.org/10.1007/978-3-319-21275-3>
- [15] John P Dickerson, David F Manlove, Benjamin Plaut, Tuomas Sandholm, and James Trimble. 2016. Position-indexed formulations for kidney exchange. In *Proc. 2016 ACM Conference on Economics and Computation*. 25–42.
- [16] Rachel Freedman, Jana Schaich Borg, Walter Sinnott-Armstrong, John P Dickerson, and Vincent Conitzer. 2020. Adapting a kidney exchange algorithm to align with human values. *Artificial Intelligence* 283 (2020), 103261.
- [17] Kristiaan M Glorie, J Joris van de Klundert, and Albert PM Wagelmans. 2014. Kidney exchange with long chains: An efficient pricing algorithm for clearing barter exchanges with branch-and-price. *Manuf. Serv. Oper. Manag.* 16, 4 (2014), 498–512.
- [18] Úrsula Hébert-Johnson, Daniel Lokshtanov, Chinmay Sonar, and Vaishali Suri-narayanan. 2024. Parameterized Complexity of Kidney Exchange Revisited. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI 2024, Jeju, South Korea, August 3-9, 2024*. [ijcai.org, 76–84](https://www.ijcai.org/76-84). <https://www.ijcai.org/proceedings/2024/9>
- [19] Úrsula Hébert-Johnson, Daniel Lokshtanov, Chinmay Sonar, and Vaishali Suri-narayanan. 2024. Parameterized Complexity of Kidney Exchange Revisited. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI-24*, Kate Larson (Ed.). International Joint Conferences on Artificial Intelligence Organization, 76–84. <https://doi.org/10.24963/ijcai.2024/9> Main Track.
- [20] Klaus Jansen, Stefan Kratsch, Dániel Marx, and Ildikó Schlotter. 2013. Bin packing with fixed number of bins revisited. *J. Comput. Syst. Sci.* 79, 1 (2013), 39–49. <https://doi.org/10.1016/J.JCSS.2012.04.004>
- [21] Zhipeng Jia, Pingzhong Tang, Ruosong Wang, and Hanrui Zhang. 2017. Efficient near-optimal algorithms for barter exchange. In *Proc. 16th Conference on Autonomous Agents and MultiAgent Systems*. 362–370.
- [22] Xenia Klimentova, Filipe Alvelos, and Ana Viana. 2014. A new branch-and-price approach for the kidney exchange problem. In *Proc. International Conference on Computational Science and Its Applications*. Springer, 237–252.
- [23] Michael Krivelevich, Zeev Nutov, Mohammad R Salavatipour, Jacques Verstraete, and Raphael Yuster. 2007. Approximation algorithms and hardness results for cycle packing problems. *ACM T Algorithms* 3, 4 (2007), 48–es.
- [24] Jian Li, Yicheng Liu, Lingxiao Huang, and Pingzhong Tang. 2014. Egalitarian pairwise kidney exchange: fast algorithms via linear programming and parametric flow. In *AAMAS*. 445–452.
- [25] Mugang Lin, Jianxin Wang, Qilong Feng, and Bin Fu. 2019. Randomized parameterized algorithms for the kidney exchange problem. *Algorithms* 12, 2 (2019), 50.
- [26] Arnab Maiti and Palash Dey. 2022. Parameterized Algorithms for Kidney Exchange. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, Luc De Raedt (Ed.). [ijcai.org, 405–411](https://doi.org/10.24963/ijcai.2022/58). <https://doi.org/10.24963/ijcai.2022/58>
- [27] Vicky Mak-Hau. 2017. On the kidney exchange problem: cardinality constrained cycle and chain problems on directed graphs: a survey of integer programming approaches. *J Comb Optim* 33, 1 (2017), 35–59.
- [28] David F Manlove and Gregg O’malley. 2015. Paired and altruistic kidney donation in the UK: Algorithms and experimentation. *J. Exp. Algorithmics* 19 (2015), 1–21.
- [29] Moni Naor, Leonard J. Schulman, and Aravind Srinivasan. 1995. Splitters and Near-Optimal Derandomization. In *Proceedings of the Twenty-Seventh Annual ACM Symposium on Theory of Computing (STOC)*. ACM, 182–191. <https://doi.org/10.1145/225058.225119>
- [30] Felix T Rapaport. 1986. The case for a living emotionally related international kidney donor exchange registry. In *Transplantation proceedings*, Vol. 18. 5.
- [31] David Gamarnik Ross Anderson, Itai Ashlagi and Alvin E Roth. 2015. Finding long chains in kidney exchange using the traveling salesman problem. *Proc. National Academy of Sciences* 112, 3 (2015), 663–668.
- [32] Dorry L Segev, Sommer E Gentry, Daniel S Warren, Brigitte Reeb, and Robert A Montgomery. 2005. Kidney paired donation and optimizing the use of live donor organs. *Jama* 293, 15 (2005), 1883–1890.
- [33] Tayfun Sönmez and M Utku Ünver. 2014. Altruistically unbalanced kidney exchange. *J. Econ. Theory* 152 (2014), 105–129.
- [34] Darren Stewart, Tatenda Mupfudze, and David Klassen. 2023. Does anybody really know what (the kidney median waiting) time is? *American Journal of Transplantation* 23, 2 (2023), 223–231.
- [35] Dylan Walsh. 2021. A beautiful application: Using economics to make kidney exchanges more efficient and fair. *Stanford Graduate School of Business* (2021).
- [36] Mingyu Xiao and Xuanbei Wang. 2018. Exact Algorithms and Complexity of Kidney Exchange. In *IJCAI*. 555–561.