École des Ponts
ParisTech

Thèse de doctorat

École doctorale MSTIC
Spécialité de doctorat : informatique

Thèse présentée et soutenue à l'École des Ponts ParisTech, le 20/12/2023, par :

Hélène LANGLOIS

Kernels and quasi-kernels in directed graphs

Composition du Jury

Irena RUSU — Examinatrice
Université de Nantes
Arnaud DE MESMAY — Examinateur
Université Gustave Eiffel
András SEBŐ — Examinateur
CNRS, Université Grenoble Alpes
Cristina BAZGAN — Rapportrice
Université Paris Dauphine
Nicolas TROTIGNON — Rapporteur
CNRS, École normale supérieure de Lyon

Encadrement de la thèse

Frédéric MEUNIER — Directeur de thèse
École des Ponts ParisTech
Stéphane VIALETTE — Directeur de thèse
CNRS, Université Gustave Eiffel

Remerciements

Mes premiers mots vont bien sûr à mes deux encadrants de thèse. J'ai eu la chance d'avoir été encadrée par ces deux chercheurs, de ceux dont l'enthousiasme constant nous fait aimer la recherche. J'ai toujours pris beaucoup de plaisir à vous voir pour parler de sciences et de tellement d'autres sujets avec sérieux et humour. Je souhaite à tous les doctorants de faire l'expérience d'un encadrement d'une telle qualité, autant scientifiquement qu'humainement. Merci Frédéric pour ton soutien, ton optimisme sincère et ton énergie débordante. Merci Stéphane pour ta bonne humeur et tes histoires.

Je remercie Christina Bazgan et Nicolas Trotignon pour avoir accepté de rapporter ma thèse et pour tous leurs précieux retours. Je remercie András Sebő d'avoir examiner sa soutenance et pour tous ses précieux retours qui ont permis d'améliorer ce manuscrit. Je remercie également Irena Rusu et Arnaud de Mesmay pour avoir accepté d'examiner ma soutenance.

Je tiens à remercier Romeo Rizzi, un collaborateur essentiel à ce travail qui m'a accueilli dans son université à Vérone. Nos discussions riches et passionnées ont été des jalons importants de ma thèse.

Merci également à l'ensemble du personnel de l'Ecole des Ponts et de l'Ecole Doctorale MSTIC pour leur accompagnement dans toutes les démarches. Merci notamment à Tony Lelièvre, Stéphanie Bonnel, Isabelle Simunic, Corinne Palescandolo et Mariam Sibide.

Merci à mes amis d'Orléans pour avoir répondu présents pour partager toutes mes joies, mes peines et mes angoisses qui ont rythmé ces années de doctorat. Pour citer Rémi, "j'ai trop hâte des ces 30 prochaines années avec vous". Merci Chloé pour toute cette tendresse et cette bienveillance. Merci Hélène pour ces conversations toujours intenses. Merci Juliette d'avoir participé à tant de bulletins météo du jour et de continuer à respecter le pacte d'écoute. Merci Lisa de continuer à me toucher par ta sensibilité. Merci Pauline de me faire danser. Merci Rémi de me faire vivre tes aventures par procuration. Merci Thomas pour ta sérénité et ton affection. Merci Timothé pour tous ces rendez-vous de septembre qui autorisent toujours à tout remettre en question. Merci Violette d'avoir envoyé et surtout reçu des milliers de vocaux pour être à côté de moi à 600 kilomètres de distance.

Parmi les personnalités qu'il m'a été donné de rencontrer à Orléans, je tiens à remercier également Kevin d'avoir cru avant moi que je ferais une thèse et Léna que j'ai si souvent appelée, merci pour les rires et les confidences.

Merci aux amis que j'ai rencontré à l'université et qui continuent de faire partie de ma vie. Merci pour tous ces moments privilégiés, avec Adrien sur ta terrasse à Montrouge, Antonin sur notre balcon à Palaiseau, avec Corentin dans les bars à Paris, Etienne au téléphone, Luis dans son appartement et Julien dans les parcs.

Merci à ceux que j'ai rencontrés à Champs-sur-Marne, avec qui j'ai pu rire, voyager et qui ont eux aussi été les victimes de mes débats internes. Merci Albéric pour ton apaisement et ta tendresse. Merci Alfred pour ta vision du monde passionnée. Merci Anton pour ta bonne humeur débordante. Merci Elias de toujours pleurer et rire à tout ce qui est drôle et surtout à ce qui ne l'est pas, de m'aider à embrasser mes nombreuses contradictions. Merci Emanuele d'être si attentif et attentionné. Merci Michel d'avoir vécu autant avec moi. Merci Roberta de réussir à être si intelligente et si bête à la fois. Merci Zoé pour tes conversations qui me font avancer et voyager, et pour ton immense soutien. Merci Dorian et Maxime pour l'organisation de ces séminaires du soir, qui donnaient une autre vision des réunions. Merci à toutes les autres personnes passées par le CERMICS que j'ai pû côtoyer.

Enfin, je remercie ma famille pour ces repas toujours animés. Merci à mes parents d' accompagner mes choix sans les orienter et merci pour votre curiosité qui m'a sans aucun doute donné le goût de la recherche. Merci Louise de continuer à chercher à me comprendre malgré nos différences. Merci Laure pour tous ces moments de grande folie et de grand calme, tous ressourçants.

Abstract

This thesis studies the concept of kernels in directed graphs, defined as subsets of vertices which are "independent" (meaning they don't contain adjacent vertices) and "absorbing" (every vertex outside the subset has at least one outneighbor in the subset). Kernels form a fundamental concept of digraph theory. Initially introduced by von Neumann and Morgenstern for game theory, they have also found applications in economics and logic. The existence of kernels in a digraph is not systematic (consider the directed cycle of length three) and is even NP-complete to decide. The first part of this thesis aims to explore the challenges and theoretical aspects of kernels.

The notion of quasi-kernels, which was introduced by Chvátal and Lovász in 1974, derives from a slight modification in the definition of kernels. A quasi-kernel is a subset of vertices that is independent and such that there is a directed path of length at most two from every vertex to that subset. Any digraph has a quasi-kernel that can be found in polynomial time. The research about quasi-kernels focuses on a conjecture called the "small quasi-kernel conjecture." It suggests that a sink-free digraph (where every vertex has a positive outdegree) has a quasi-kernel of size at most half of the vertex set. However, this conjecture is only confirmed for specific classes of digraphs. The second half of the thesis focuses on quasi-kernels, particularly on their algorithmic aspect, which has not been extensively explored yet.

Résumé étendu en français

Dans un graphe orienté, un sous-ensemble de sommets est *indépendant* s'il ne contient pas de sommets adjacents et *absorbant* si tout sommet est dans cet ensemble ou a un voisin sortant dans cet ensemble. Un *noyau* est un sous-ensemble de sommets qui est indépendant et absorbant. Formellement, un noyau d'un graphe orienté D est un sous-ensemble S de $V(D)$ qui est indépendant et tel que $V(D) \setminus S = N^{-}(S)$. Les figures 5 et 6 fournissent des exemples de noyaux (en rouge), l'un dans un cycle orienté de longueur quatre et l'autre dans un graphe quelconque. La figure 7 représente un cycle orienté de longueur trois, lequel n'a pas de noyau.

La notion de noyau est fondamentale en théorie des graphes orientés. Elle a été introduite par von Neumann et Morgenstern pour l'étude des stratégies gagnantes en théorie des jeux [74], et possède désormais des applications dans d'autres domaines, comme en économie ou en logique. Les noyaux ont été activement étudiés et sont encore au coeur de nombreuses questions scientifiques.

Nous avons vu que certains graphes orientés n'avaient pas de noyaux et c'est en fait NP-complet de décider si un graphe orienté a ou non un noyau [27]. Pour cette raison, la recherche autour des noyaux se concentre surtout sur des conditions suffisantes d'existence.

La première partie de cette thèse a pour but d'explorer ces différentes questions.

Chvátal et Lovász [26] ont montré en 1974 qu'une petite modification dans la définition de noyau assurait son existence systématique : ils ont prouvé que tout graphe orienté a un sous-ensemble indépendant Q de sommets tel que tout sommet n'appartenant pas à Q a un chemin orienté de longueur au plus deux vers un élément de Q. Un tel sous-ensemble est un *quasi-noyau*. Leur preuve fournit d'ailleurs un algorithme polynomial pour en trouver un dans tout graphe orienté (d'autres preuves simples existent [17]). Un noyau est évidemment un cas particulier de quasi-noyau. L'exemple de la figure 6 est donc, en particulier, un exemple de quasi-noyau d'un graphe orienté quelconque, et la figure 8 est un exemple d'un autre quasi-noyau dans le même graphe orienté. Contrairement aux noyaux, les quasi-noyaux ne possèdent pas encore d'applications, et la motivation de leur étude est purement théorique. Cependant, le domaine est actif (plusieurs papiers d'auteurs différents ont été publiés au cours de ce doctorat) avec une conjecture principale qui mène presque toute l'activité autour du sujet, la *conjecture des petits quasi-noyaux*, énoncée en 1976 par Erdős et Székely [41].

Conjecture ("Conjecture des petits quasi-noyaux")**.** *Tout graphe orienté sans puits D a un quasi-noyau de taille au plus $\frac{1}{2}|V(D)|$.*

Cette conjecture n'a été prouvée que pour certaines classes de graphes orientés. La seconde partie de cette thèse concerne l'étude des quasi-noyaux, et plus particulièrement de leurs aspects algorithmiques, sujet qui n'avait pas été traité précédemment.

Noyaux

Comme cela a déjà été évoqué, une partie importante de la recherche concernant les noyaux a été consacrée à chercher des conditions suffisantes d'existence. Le chapitre 1 est

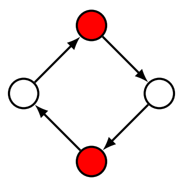

FIGURE 1 : Exemple d'un noyau dans un cycle orienté de longueur quatre.

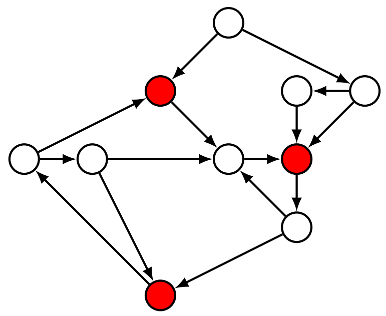

FIGURE 2 : Exemple d'un noyau dans un graphe orienté quelconque.

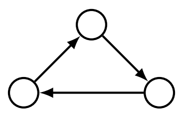

FIGURE 3 : Un cycle orienté de longueur trois n'a pas de noyau.

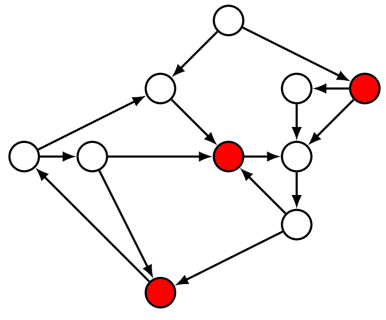

FIGURE 4 : Exemple d'un quasi-noyau dans un graphe orienté quelconque.

une brève introduction à la notion de noyaux. Il passe en revue les résultats classiques d'existence, mais aussi les principales applications et quelques questions ouvertes. En particulier, on y voit que les noyaux ont un lien particulier avec deux domaines différents, la théorie des jeux et l'étude des graphes parfaits.

Motivés par ce lien avec les graphes parfaits, Berge et Duchet ont introduit en 1983 [10] la notion de graphes orientés noyaux-parfaits. Un graphe orienté est *noyau-parfait* si tout sous-graphe orienté induit a un noyau (le graphe lui-même y compris). Une célèbre conjecture de Berge et Duchet, maintenant connue sous le nom du théorème de Boros et Gurvich [18], établit que sous certaines conditions sur l'orientation, la noyau-perfection coïncide avec la perfection.

Théorème (Boros et Gurvich [18])**.** *Toute orientation clique-acyclique d'un graphe parfait a un noyau.*

Le chapitre 2 commence par introduire le sujet de noyau-perfection et y contribue de deux manières différentes. D'abord, il traite d'une remarque de Boros et Gurvich qui laisse croire que deux sous-classes de graphes orientés — les "noyaux-solvables" et les "M-noyaux-solvables" sont distinctes. Cependant, comme prouvé dans la section 2.1, cela ne peut être vrai avec une telle définition de la M-noyau-solvabilité. En effet ils affirment que toute $M-$orientation clique-acyclique simple d'un anti-trou $\overline{C}_7$ a un noyau, ce qui ne peut être vrai avec la définition classique de $M-$orientation qui inclut les orientations clique-acyclique, comme le prouve l'obervation suivante.

Observation. *L'anti-trou $\overline{C}_7$ a une orientation clique-acyclique simple sans noyau.*

On la répare partiellement en montrant qu'il existe des graphes simple-noyau-solvables qui ne sont pas noyau-solvables.

Théorème. *Soit D une orientation simple clique-acylique d'un anti-trou impair $\overline{C}_{2k+1}$. Si $k \geqslant 4$, alors D a un noyau.*

La classe des graphes orientés M-noyaux-solvables est une sous-classe des graphes orientés simple-noyaux-solvables, mais la question de l'égalité ou non de ces deux classes reste ouverte. Ensuite, le théorème fournit des versions algorithmiques des théorèmes de Blidia et al. [13] concernant des opérations de graphes préservant l'existence de noyaux (théorème 2.3.1 et proposition 2.3.3). Le théorème suivant est sûrement aussi important que celui de Boros et Gurvich.

Théorème (Galeana-Sánchez et Neumann-Lara [44])**.** *Si tous les cycles orientés impairs ont deux cordes avec des têtes consécutives, il y a un noyau.*

Une version courte de la preuve de ce théorème est présentée au chapitre 3, avec une généralisation immédiate, les aspects algorithmiques de ce théorème, et la possibilité de vérifier les conditions et de calculer un noyau efficacement sont également discutés.

Il existe un autre théorème assurant l'existence d'un noyau. Il s'agit d'une généralisation du célèbre théorème de Gale-Shapley sur les mariages stables. On peut le formuler en termes de graphe orienté avec des arcs bleus et rouges comme suit.

Théorème (Sands, Sauer et Woodrow [71])**.** *Soit D un graphe orienté dont les arcs sont colorés en bleu et en rouge, de telle manière que la restriction à chaque couleur forme un graphe transitif ; alors D admet un noyau.*

Dans le chapitre 4, nous proposons des extensions et des variantes de ce théorème. Deux généralisations du théorème Sands–Sauer–Woodrow sont présentées.

Théorème. *Soit D un graphe orienté dont les arcs sont colorés en rouge et bleu, tel que les deux conditions suivantes sont satisfaites :*

- *Si (u, v) et (v, w) sont des arcs bleus, alors (u, w) est un arc bleu ou (w, u) et (w, v) sont des arcs rouges.*
- *Si (u, v) et (v, w) sont des arcs rouges, alors (u, w) est un arc rouge ou (v, u) et (w, u) sont des arcs bleus.*

Alors D a un noyau et il possible de le trouver en temps polynomial.

Théorème. *Soit D un graphe orienté dont les arcs sont colorés en bleu et rouge tel que les conditions suivantes sont satisfaites :*

- *Il n'y a pas de cycle dirigé monochromatique.*
- *Si $((v_1, v_2), (v_2, v_3), (v_3, v_4))$ est un chemin dirigé (ouvert ou fermé) tel que (v_1, v_2) est un arc rouge et (v_3, v_4) est un arc bleu alors ses sommets induisent au moins un autre arc ne terminant pas en v_2.*

Alors D a un noyau.

Pour chacune d'entre elle, nous essayons de pousser aussi loin que possible la technique de démonstration du théorème original. Pour l'une, nous obtenons même un algorithme en temps polynomial pour calculer un noyau. Pour cet algorithme, certains résultats élémentaires sur les posets (plus précisément, sur les posets d'antichaînes) jouent un rôle crucial ; voir la section 4.2.1.

En Annexe A se trouve un tableau résumant les complexités des principaux problèmes de complexité liés aux noyaux et un tableau répertoriant les principales conditions connues assurant l'existence d'un noyau.

Quasi-noyaux

Cette partie est le résultat d'un travail en collaboration avec Romeo Rozzi. Le chapitre 5 est une courte introduction aux quasi-noyaux et commence par deux preuves de l'existence d'un quasi-noyau dans tout graphe orienté. Comme évoqué précédemment, la conjecture des petits quasi-noyaux est la principale motivation de l'étude des quasi-noyaux. Elle reste encore ouverte à ce jour, et les résultats connus sont listés en section 5.2. En fait, même si la conjecture est formulée avec un ratio de $1/2$, l'existence d'un quasi-noyau avec tout autre ratio plus grand n'est pas connue. Dans le chapitre 6, nous contribuons à l'étude de cette question en prouvant le théorème suivant.

Théorème. *Tout split graphe orienté sans puits D a un quasi-noyau de taille au plus $\frac{3}{4}|V(D)|$.*

Dans ce même chapitre, nous prouvons également la conjecture avec un ratio asymptotiquement nul pour une sous-classe des splits.

Théorème. *Soit D une orientation d'un split graphe complet. Si D a un puits, alors il y a un unique quasi-noyau de taille minimum, composé des puits du graphe orienté. Si D n'a aucun puits, alors il a un quasi-noyau de taille au plus deux.*

Un des objectifs de cette thèse est d'initier l'étude algorithmique des quasi-noyaux. Le principal résultat établi est le suivant.

Théorème. *Décider si un graphe orienté a deux quasi-noyaux disjoints est un problème* NP*-complet.*

Ce résultat est motivé par une conjecture de Gutin et al. [47] – maintenant réfutée – affirmant que tout graphe orienté sans puits admet deux quasi-noyaux disjoints. L'existence de tels quasi-noyaux implique en particulier que l'un d'entre eux a une taille d'au plus la moitié des sommets. Un autre de nos résultats montre que calculer un quasi-noyau de taille minimale est difficile, même dans des graphes orientés très simples. Par exemple, on peut citer les deux résultats suivants.

Théorème. *Décider si une orientation d'un graphe split a un quasi-noyau de taille au plus k est un problème* W[2]*-complet.*

Théorème. *Décider si une orientation acyclique d'un graphe biparti a un quasi-noyau de taille au plus k est un problème* W[2]*-complet.*

Ces résultats, ainsi que leurs preuves, sont présentés dans le chapitre 7.

Selon le théorème de Courcelle [29], le calcul d'un quasi-noyau de taille minimale peut être effectué en temps polynomial pour les orientations de graphes ayant une largeur arborescente bornée, ainsi que pour le problème de décider s'il existe k quasi-noyaux deux à deux disjoints dans un graphe orienté. La complexité des algorithmes obtenus par une application directe du théorème de Courcelle est polynomiale, mais comporte un facteur constant élevé dépendant de la largeur arborescente. Nous démontrons que nous pouvons en réalité obtenir une complexité polynomiale avec un facteur constant raisonnable, avec les théorèmes suivants, sachant qu'une bonne décomposition arborescente peut être trouvée efficacement.

Théorème. *Si une bonne décomposition arborescente $(T, \mathcal{X})$ d'un graphe orienté de largeur w est donnée, trouver un quasi-noyau peut être fait en temps $\mathcal{O}(25^w|V(T)|)$.*

Théorème. *Si une bonne décomposition arborescente $(T, \mathcal{X})$ d'un graphe orienté de largeur w est donnée, trouver deux quasi-noyaux disjoints s'ils existent peut être fait en temps $\mathcal{O}(25^{2w}|V(T)|)$.*

Ce travail a été réalisé en collaboration avec Julien Baste et Antoine Castillon, et est présenté à la fin du chapitre 7.

En Annexe B se trouve un tableau résumant les complexités des principaux problèmes de complexité liés aux quasi-noyaux et un tableau répertoriant les principales conditions connues assurant l'existence d'un petit quasi-noyau.

Contents

List of Figures

Introduction

This thesis focuses on digraphs (directed graphs). We assume basic knowledge in this topic. Some terminology and notation are given at the end of this introduction. For those not defined in the latter, we refer the reader to the book by Bang-Jensen and Gutin [8]. As done in this book, the vertex set of an undirected graph G is denoted by $V(G)$ and its edge set by $E(G)$, and the vertex set of a digraph D is denoted by $V(D)$ and its arc set by $A(D)$.

In a digraph, a subset of vertices is *independent* if it does not contain adjacent vertices, and it is *absorbing* if every vertex is either in the subset or has an outneighbor in it. A *kernel* is a subset of vertices that is both independent and absorbing. Formally, a kernel of a digraph D is a subset S of $V(D)$ that is independent and such that $V(D) \setminus S = N^-(S)$. Figure 5 and Figure 6 illustrate examples of kernels represented in red, one in a directed cycle of length four and the other in an arbitrary graph. Figure 7 represents the example of the directed cycle of length three, which admits no kernel.

Kernels form a fundamental topic of the theory of digraphs. Introduced by von Neumann and Morgenstern in the context of board games analysis [74], they have found applications in other areas like economy and logic. They have been the subject of many research works and there are still several challenges about them. We have seen that there are graphs with no kernel, and actually it is even NP-complete to decide whether a digraph admits a kernel [27], that is why theoretical research about kernels mostly focuses on sufficient conditions for their existence.

The first half of the thesis aims at exploring a bit further these challenges.

Chvátal and Lovász [26] have shown in 1974 that a slight modification in the definition of kernel ensures its existence: they proved that any digraph has an independent subset Q of vertices such that every vertex not in Q has a directed path of length at most two to an element of Q. Such a subset is a *quasi-kernel*. Their proof provides a polynomial-time algorithm to find one in any digraph (alternate simple proofs exist [17]). Clearly, a kernel is a special case of a quasi-kernel. The example of Figure 6 is then a particular example of a quasi-kernel in an arbitrary digraph, and Figure 8 is an example of another quasi-kernel in the same digraph. Contrary to kernels, there are no known applications of quasi-kernels, and their motivation is essentially theoretical. Yet, it is a quite active topic (with several papers published by various authors during this PhD), with a main conjecture that drives almost all the activity around them, namely the *small quasi-kernel conjecture* formulated in 1976 by Erdős and Székely [41]. This conjecture states that a sink-free digraph D (i.e., every vertex of D has positive outdegree) has a quasi-kernel of size at most $\frac{1}{2}|V(D)|$. So far, this conjecture is only confirmed for narrow classes of digraphs.

The second half of the thesis is about quasi-kernels, in particular about their algorithmic aspect, something that has not been investigated before this thesis.

We provide now more details about the organization of the thesis and our main contributions.

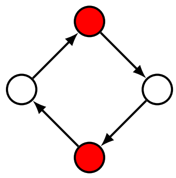

Figure 5: Example of a kernel in a directed cycle of length four.

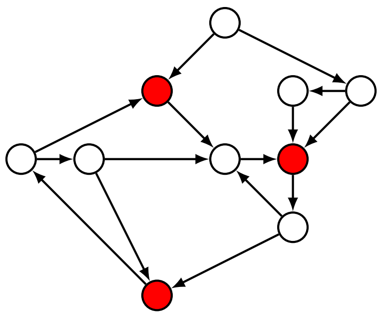

Figure 6: Example of a kernel in an arbitrary digraph.

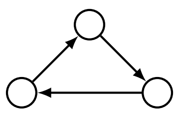

Figure 7: A directed cycle of length three has no kernel.

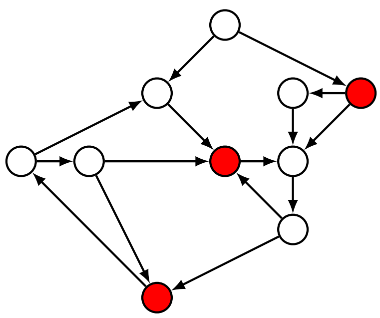

Figure 8: Example of a quasi-kernel in an arbitrary digraph.

Kernels

As mentioned above, an important stream of research about kernels has dealt with sufficient conditions for their existence. Chapter 1 aims at being a gentle introduction to the topic of kernels, which reviews the classical existence results, but also their main applications and some important open questions about them. We will in particular see there that kernels have strong connections with two somehow independent areas, namely game theory and perfect graphs.

Motivated by this relation with perfect graphs, Berge and Duchet have introduced in 1983 [10] the notion of kernel-perfect digraphs. A digraph is *kernel-perfect* if each induced subgraph admits a kernel. A celebrated conjecture of Berge and Duchet, now a theorem by Boros and Gurvich [18], states that some simple necessary conditions for kernel-perfectness are also sufficient if a graph is perfect. Chapter 2 provides a brief discussion on some aspects of kernel-perfection, and also contributes to that topic in two different ways. First, it deals with a statement by Boros and Gurvich [19] suggesting that two subclasses of kernel-perfect graphs—"kernel-solvable" and "M-kernel-solvable" (defined in Section 2.1)—are distinct. Yet, as we show in Section 2.1, it cannot be true with this definitions. We partially fix it by showing that there are "simple-kernel-solvable" graphs that are not kernel-solvable (Proposition 2.1.2). M-kernel-solvable graphs form a subclass of simple-kernel-solvable graphs, but it remains open whether it is a strict subclass. Second, it provides some algorithmic versions of theorems by Blidia et al. [13] on graph operations preserving the existence of kernels (Theorems 2.3.1 and Proposition 2.3.3).

There is another theorem ensuring the existence of a kernel. It is due to Galeana-Sánchez and Neumann-Lara [44] and ensures the following: *If every odd directed cycle has two chords with consecutive heads, then there is a kernel.* We propose in Chapter 3 a short version of their proof, together with immediate generalizations, and discuss a few algorithmic aspects of their theorem, namely, the possibility of checking the condition and of computing a kernel efficiently.

Another theorem ensuring the existence of a kernel is the Sands–Sauer–Woodrow theorem [71], which is actually a generalization of the celebrated Gale–Shapley theorem (about stable marriages). It can be stated in terms of a digraph with blue and red arcs: *Let D be a directed graph whose arcs are colored in blue and red and such that the restriction to each color forms a transitive digraph; then D admits a kernel.* In Chapter 4, we discuss and propose extensions and variations of this theorem. We obtain two generalizations of the Sands–Sauer–Woodrow theorem (Theorems 4.1.1 and 4.1.2). For each of them, we try to push as far as possible the proof technique of the original theorem. For one of the generalizations, we even get a polynomial-time algorithm for computing a kernel. For this result, some elementary (but new?) results on posets (more precisely, about posets of antichains) play a crucial role; see Section 4.2.1.

In Appendix A, one can find a table summarizing the complexities of the main complexity-related problems associated with kernels, and a table listing the main known conditions ensuring the existence of a kernel.

Quasi-kernels

This part results from a collaboration with Romeo Rizzi. Chapter 5 is a brief introduction to quasi-kernels and starts with two distinct proofs of the existence of a quasi-kernel in any digraph. As emphasized above, the small quasi-kernel conjecture is the major motivation for the study of quasi-kernels. It is still wide open and the known results are listed in Section 5.2. Actually, even if the conjecture is stated with a ratio 1/2, the existence of a quasi-kernel within any other larger fixed ratio is not known. In Chapter 6, we contribute to that question by showing that in every sink-free orientation of a split graph, there is a quasi-kernel of size 3/4 of the total number of vertices (Theorem 6.1.6). In this chapter, we also prove that if the split graph is complete, then there is a quasi-kernel of size at most two (Theorem 6.2.2). These two results, together with a few complementary results, are the content of the following paper, written in collaboration with my supervisors and Romeo Rizzi:

> Hélène Langlois, Frédéric Meunier, Romeo Rizzi, and Stéphane Vialette, *Quasi-kernels in split graphs*, to be submitted.

One of the starting objectives of this thesis was to provide the first contributions to the algorithmic aspects of quasi-kernels. This objective has been fulfilled. In the following paper, with the same collaborators,

> Hélène Langlois, Frédéric Meunier, Romeo Rizzi, and Stéphane Vialette, *Algorithmic aspects of small quasi-kernels*, International Workshop on Graph-Theoretic Concepts in Computer Science, Springer, 2022, pp. 370–382.

we prove two main results. The first one establishes that deciding whether a digraph admits two disjoint quasi-kernels is NP-complete. This result is motivated by a conjecture of Gutin et al. [47]—now disproved—stating that every sink-free digraph admits two disjoint quasi-kernels. Note that the existence of such quasi-kernels implies in particular that one of them is of size at most half of the vertices. The second implies that computing a minimum-size quasi-kernel is hard, even in very simple digraphs. For instance, it is W[2]-hard to compute a quasi-kernel of minimum size in orientations of split graphs, and also in acyclic orientations of bipartite graphs. These results with their proof are given in Chapter 7.

By Courcelle's theorem [29], computing a quasi-kernel of minimum size can be done in polynomial-time for orientations of graphs with bounded treewidth, and similarly for the problem of deciding whether there are k pairwise disjoint quasi-kernels in a digraph. The complexity of the algorithms obtained by a direct application of Courcelle's theorem is polynomial but has a huge constant factor depending on the treewidth. We prove that we can actually get a polynomial complexity with a reasonable constant factor. This work has been done in collaboration with Julien Baste and Antoine Castillon, and is given at then end of Chapter 7.

In Appendix B, one can find a table summarizing the complexities of the main complexity-related problems associated with quasi-kernels, and a table listing the main known conditions ensuring the existence of a small quasi-kernel.

Definition and notation

In this thesis all digraphs are finite and do not contain *parallel* arcs, namely arcs with the same head and tail. This way, an arc is identified with an ordered pair of vertices.

We consider a digraph D. A vertex v is an *inneighbor* (resp. *outneighbor*) of a vertex u if the arc (u, v) (resp. (v, u)) exists in the digraph. The set of all inneighbors (resp. outneighbors) of v is denoted by $N^-(v)$ (resp. $N^+(v)$) and is the *inneighborhood* (resp. *outneighborhood*) of v. The *closed* inneighborhood (resp. outneighborhood) of v, denoted by $N^-[v]$ (resp. $N^+[v]$), is defined as $N^-(v) \cup \{v\}$ (resp. $N^+(v) \cup \{v\}$). The *second inneighborhood* of v, denoted by $N^{--}(v)$, is defined as $N^-(N^-(v)) \setminus \{v\}$. The second outneighborhood could be defined analogously but will not be used in this thesis.

The inneighbors (resp. outneighbors) of a set S is $N^-(S) \coloneqq \{N^-(v)\colon v \in V(D)\} \setminus S$ (resp. $N^+(S) \coloneqq \{N^+(v)\colon v \in V(D)\} \setminus S$) and the closed inneighbordhood of S, denoted by $N^-[S]$ (resp. $N^+[S]$), is defined as $N^-(S) \cup S$ (resp. $N^+(S) \cup S$). The *second inneighbordhood* of S, denoted by $N^{--}(S)$, is defined as $N^-(N^-(S)) \setminus N^-[S]$. A vertex v is a *sink* (resp. *source*) if it has no outneighbhor (resp. inneighbor). A *walk* in D is a sequence $v_1v_2\dots v_k$ of vertices of D such that $(v_i, v_{i+1}) \in A(D)$. If the vertices $v_1, v_2, \dots, v_{k-1}$ are distinct, and $v_1 = v_k$, then the walk is a *cycle*. The *length* of a walk is the number of its arcs. A digraph D is *strongly connected* if for every pair $a, b \in V(D)$, there is a directed path from a to b. A set of vertices $C \subseteq V(D)$ is a *strongly connected component* of D if $D[C]$ is strongly connected and maximal for inclusion with this property.

An *hence-and-forth* pair of arcs is of the form $((u, v), (v, u))$. $(v, u) \in A(D)$. A digraph D is an *orientation* of a graph G if $V(D) = V(G)$ and for every $a, b \in E(G)$, (a, b), (b, a) or both are in $A(D)$. An orientation D is *simple* if there is no hence-and-forth pair of arcs.

We define the (directed) distance $d(v, w)$ from a vertex v to a vertex w as the minimum length of a directed path from v to w.

Part I

Kernels

Chapter 1

Kernels in a nutshell

We have already mentioned that it is NP-complete to decide whether a digraph admits a kernel [27]. However, many sufficient conditions have been proposed in the literature. This chapter, which aims at being a gentle introduction to the topic, presents the most famous conditions (Section 1.1), provides various applications (Section 1.2), and discusses the main challenges of this topic (Section 1.3).

1.1 Main conditions ensuring existence of a kernel

This section, aiming at presenting the main sufficient conditions from the literature ensuring the existence of a kernel, is subdivided into three subsections, each collecting sufficient conditions with a similar flavor:

- Conditions based on acyclicity of the digraph or on odd cycles.
- Conditions relying on the perfection of the underlying undirected graph.
- Conditions expressed in terms of a partition of the arc set into "red" and "blue" arcs.

It is worth emphasizing that all these conditions are about classes closed under taking induced subgraph. We are not aware of any sufficient condition from the literature that does not possess this property.

1.1.1 Acyclicity and odd cycles

1.1.1.1 Existence results

Among the most famous results on kernels, several ones ensure the existence of a kernel under a condition on cycles, and especially odd cycles. In this section, we aim at presenting the most important. We also proposed for most of them a short proof based on the notion of semi-kernel. (The idea of systematizing the use of semi-kernels for this kind of results goes back to the work of Galeana-Sánchez and Neumann-Lara [45].) A *semi-kernel* is a subset S of vertices that is independent and such that $N^+(S) \subseteq N^-(S)$. Note that a kernel is a semi-kernel, and that the empty set is also a semi-kernel. Its relevance is formalized by the following lemma.

Lemma 1.1.1. *If every non-empty induced subdigraph of a digraph D has a non-empty semi-kernel, then D has a kernel.*

Proof. We prove the result by induction on the number of vertices. Let D be a digraph whose non-empty induced subdigraphs all have a non-empty semi-kernel. If D has exactly one vertex, it has a kernel. So, suppose D has at least two vertices. By the assumption, D itself has a non-empty semi-kernel, which we denote by S. If S is a kernel, we are done. Otherwise, the digraph $D - N^-[S]$ is a non-empty induced subdigraph of D, which implies that all its induced subdigraphs have also a non-empty semi-kernel. By induction, $D - N^-[S]$ has a kernel. This kernel forms with S a kernel of D. □

The first theorem ever stated on kernels is the following.

Theorem 1.1.2 (Von Neumann and Morgenstern [74]). *Every acyclic digraph has a kernel.*

Proof. Consider an acyclic digraph. Every non-empty induced subdigraph has a sink, which is a semi-kernel of the subdigraph. Lemma 1.1.1 leads to the conclusion. □

This theorem has been generalized in many different ways. The next theorem is an example of such a generalization.

Theorem 1.1.3 (Duchet [36]). *Every digraph such that each directed cycle has at least one hence-and-forth pair of arcs has a kernel.*

Proof. Let D be a digraph such that each directed cycle has at least one of hence-and-forth pair of arcs. Let D' be a non-empty induced subdigraph of D. Remove from D' all its hence-and-forth pairs of arcs. This leads to a (non-empty) acyclic digraph. Pick any sink of this acyclic digraph. It is a semi-kernel of D'. Lemma 1.1.1 leads to the conclusion. □

From now on, all results of this subsection involves conditions on odd directed cycles. The next theorem is due to Richardson and generalizes Theorem 1.1.2. Originally, it was not stated in terms of kernels, but it is now considered as one of the most fundamental results about kernels.

Theorem 1.1.4 (Richardson [67]). *Every digraph with no odd directed cycle has a kernel.*

The proof will require the following lemma.

Lemma 1.1.5. *Every strongly connected digraph with no odd directed cycle is bipartite.*

Proof. Let D be a strongly connected digraph with no odd directed cycle. Let us prove the underlying graph of D has no odd cycle. By contradiction, consider $C = v_1, \ldots, v_m$ an odd cycle in the underlying graph and for each i, let W_i be a minimum walk from v_i to v_{i+1}. For each i, either $(v_i, v_{i+1}) \in A(D)$ and W_i has length one or $(v_{i+1}, v_i) \in A(D)$ and the length of W_i is odd (otherwise W_i together with (v_{i+1}, v_i) would be an odd directed cycle in D). Finally $W_1 W_2 \ldots W_m$ is an odd directed closed walk because each W_i is odd and m is odd as well. Consider the smallest directed closed walk of odd length. By minimality, this walk is a cycle. A contradiction. □

Proof of Theorem 1.1.4. Assume that D has no odd directed cycle. Let D' be a non-empty induced subdigraph of D. Let $B \subseteq V(D')$ be a strongly connected component of D' such that $N^+(B) = \varnothing$. According to Lemma 1.1.5, $D'[B]$ is bipartite. Pick any side of this bipartite graph. It is a semi-kernel of D'. Lemma 1.1.1 leads to the conclusion. □

There is an "odd directed cycle" version of Theorem 1.1.3.

Theorem 1.1.6 (Duchet [36]). *Every digraph such that each odd directed cycle has at least two hence-and-forth pairs of arcs has a kernel.*

Proof. Let D be a digraph such that each odd directed cycle has at least two hence-and-forth pairs of arcs. Let D' be a non-empty induced subdigraph of D. Remove from D' all its hence-and-forth pairs of arcs. Let $B \subseteq V(D')$ be a strongly connected component of D' such that $N^+(B) = \varnothing$. According to Lemma 1.1.5, $D'[B]$ is bipartite. Pick any side S of this bipartite graph. No pair of hence-and-forth pairs of arcs in D' has both endpoints in S since otherwise there would be an odd directed cycle with only one hence-and-forth pair of arcs. Therefore S is a semi-kernel of D'. Lemma 1.1.1 leads to the conclusion. □

The following theorem is probably the most powerful result in this stream of research. A quite short proof will be given in Chapter 3. Notably, it uses again semi-kernels in a way similar to what has been done for the results of the present section.

Theorem 1.1.7 (Galeana-Sánchez and Neumann-Lara [44])**.** *Let D be a digraph. Suppose that in D each odd directed cycle has at least two chords with consecutive heads. Then D has a kernel.*

1.1.1.2 Algorithmic considerations

We do not know whether it is possible to check in polynomial time the condition of Theorem 1.1.7. This will be further discussed in Chapter 3. Notice however that the conditions of all other theorems of this section can be checked in polynomial time. For Theorems 1.1.2 and 1.1.3, it is obvious. For Theorem 1.1.4, it can be done by checking that all strongly connected components are bipartite (see Lemma 1.1.5). For Theorem 1.1.6, remove all hence-and-forth pairs of arcs, and then check as for Theorem 1.1.4 if the remaining digraph together with any hence-and-forth pair of arcs contains an odd directed cycle.

Regarding the computation of a kernel itself, apart from Theorem 1.1.7, all proofs provide directly polynomial algorithms to compute a kernel.

1.1.2 Perfect graphs

For an undirected graph G, we denote by $\omega(G)$ its clique number (the maximum size of a clique) and by $\chi(G)$ its chromatic number (the minimum number of colors in a proper coloring). A graph G is *perfect* if $\omega(G') = \chi(G')$ for each induced subgraph G' of G (the graph G included). They form one of the most important graph classes, which is also a central notion in combinatorial optimization and information theory; see [72] for a survey. The strong perfect graph theorem, conjectured originally by Berge in 1961 [9] and proved by Chudnovsky et al. in 2006 [25], claims that a graph is perfect if and only if it does not contain an induced odd hole (chordless odd cycle of length at least 5) or an induced odd anti-hole (complement of an odd hole). The conjecture by Berge has lead to an intensive research activity around perfect graphs.

A directed graph is *clique-acyclic* if every directed cycle in a clique has at least one hence-and-forth pair of arcs. This condition is equivalent to require that every clique admits a kernel (i.e., a vertex absorbing all other vertices). Berge and Duchet [10] conjectured that every clique-acyclic orientation of a perfect graph admits a kernel. In Chapter 2, this conjecture will be further discussed, in particular for its relation with the notion of "kernel-perfection." This conjecture has been proved by Boros and Gurvich in 1996 [18]. A family of clique-acyclic orientations that plays an important role in the theory of kernels is formed by the M-clique-acyclic orientations. It has been introduced by Meyniel according to Duchet [37]. An orientation is an *M-clique-acyclic orientation* if every directed cycle of length three has at least two hence-and-forth pairs of arcs; see Figure 1.1. It is an easy fact that M-clique-acyclic orientations are indeed clique-acyclic. In this section, we state

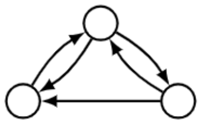

Figure 1.1: Every directed triangle looks like this in an M-clique-acyclic orientation.

formally the theorem of Boros and Gurvich and partial results about subclasses of perfect graphs and M-clique-acyclic orientations obtained before the resolution of the conjecture. We also briefly address some algorithmic aspects, which will anyway be more thoroughly discussed in Chapter 2.

1.1.2.1 Existence results

The next theorem is probably the most celebrated one in the area of kernels. Its original formulation was in terms of stable matchings in a bipartite graph. It is detailed in Section 1.2.3. The formulation below is obtained by translating the statement on the line graph of this bipartite graph. The *line graph* H of a graph G is a graph the vertices of which correspond to the edges of G, any two vertices of H being adjacent if and only if the corresponding edges of G are adjacent. It can easily be shown that line graphs of bipartite graphs are perfect [72].

Theorem 1.1.8 (Gale and Shapley [42] 1962)**.** *Every simple clique-acyclic orientation of the line graph of a bipartite graph admits a kernel.*

Maffray later found the following existence result generalizing the previous one.

Theorem 1.1.9 (Maffray [64] 1992)**.** *Every M-clique-acyclic orientation of a perfect line graph admits a kernel.*

Later on, different results have been found about the existence of kernels in particular orientations on different families of perfect graphs.

We cite here some main examples of existence results for orientations of perfect graphs.

A graph G is a *comparability* graph if there exists an orientation D of G such that if $(a, b) \in A(D)$ and $(b, c) \in A(D)$, then $(a, c) \in A(D)$. A graph G is *perfectly orderable* if there exists an acyclic orientation D of G with no induced P_4 $abcd$ in G such that $(a, b) \in A(D)$ and $(d, c) \in A(D)$.

Comparability graphs form a subclass of perfectly orderable graphs, which are perfect graphs [65].

Theorem 1.1.10 (Champetier [22] 1989)**.** *Every M-clique-acyclic orientation of a comparability graph admits a kernel.*

Theorem 1.1.11 (Blidia and Engel [14] 1992)**.** *Every M-clique-acyclic orientation of a perfectly orderable graph admits a kernel.*

A graph is *Meyniel* if every cycle of odd length at least 5 has at least two chords. Meyniel graphs form one of the most important subclasses of perfect graphs. A graph is *parity* if any two induced paths joining the same pair of vertices have the same parity. A parity graph is a Meyniel graph. Indeed, any odd cycle can be partitioned into two non-trivial paths of different parities. In a parity graph, there exists then a chord. Partitioning the cycle into two paths between the endpoints of this chord, this implies the existence of another chord.

Blidia proved the existence of a kernel in any M-clique-acyclic orientation of a parity graph [12]. This result has then been generalized as follows.

Theorem 1.1.12 (Blidia, Duchet, and Maffray [15] 1994)**.** *Every M-clique-acyclic orientation of a Meyniel graph admits a kernel.*

A graph is *i-triangulated* if every cycle of odd length at least five has at least two non crossing chords. An i-triangulated graphs is clearly a Meyniel graph. The following result is not included in the previous one, because of the more general orientation.

Theorem 1.1.13 (Maffray [63] 1986)**.** *Every clique-acyclic orientation of an i-triangulated graph admits a kernel.*

All the previous theorems are included in the following result established in 1996.

Theorem 1.1.14 (Boros and Gurvich [18] 1996)**.** *Every clique-acyclic orientation of a perfect graph admits a kernel.*

1.1.2.2 Algorithmic considerations

Deciding whether a graph is perfect can be done in polynomial time [23]. Deciding whether a graph belongs to any of the subclasses of perfect graphs mentioned in the theorems of Section 1.1.2.1, except for the perfectly orderable for which it is an NP-complete problem [50], can also be done in polynomial time:

- Line graph: linear [70].
- Comparability graph: $\mathcal{O}(|E(G)|a(G))$ where $a(G)$ is the arboricity of G [61].
- Parity graph: linear [21].
- Meyniel graph: $\mathcal{O}(|E(G)|^2)$ [20].
- i-triangulated graph: $\mathcal{O}(|V(G)||E(G)|)$ [69].

Deciding whether a simple orientation is clique-acyclic is clearly polynomial. Same thing for M-clique-acyclic orientations. This shows that the conditions of Theorems 1.1.8, 1.1.9, 1.1.10, and 1.1.12 can all be checked in polynomial time.

We are left with the two theorems stated for all clique-acyclic orientations. The condition of Theorem 1.1.14 is unlikely to be checked in polynomial time because checking that an orientation of a perfect graph is clique-acyclic is coNP-complete [7]. The complexity of checking the condition of Theorem 1.1.13 is not known.

Regarding the computation of a kernel, it is known to be polynomial for some special cases, which will be discussed in Chapter 2. We emphasize that no polynomial result is known for computing a kernel in a digraph satisfying the condition of the theorem of Boros and Gurvich (Theorem 1.1.14).

1.1.3 "Red-blue" conditions

1.1.3.1 Existence result

The following theorem is a famous generalization of the Gale–Shapley theorem (Theorem 1.1.8), each color corresponding to one side of the bipartite graph. The original statement was even more general since it was formulated for infinite digraphs.

Theorem 1.1.15 (Sands, Sauer, Woodrow [71])**.** *Let D be a directed graph whose arcs are colored with two colors. Then there is an independent set S of vertices of D such that, for every vertex x not in S, there is a monochromatic path from x to a vertex of S.*

This theorem has received many generalizations. Moreover, the proof technique was also fruitfully re-applied, as done for instance by Champetier to establish his theorem about comparability graphs (Theorem 1.1.10) or by Blidia and Engel for their theorem about perfectly orderable graphs (Theorem 1.1.11).

1.1.3.2 Algorithmic considerations

The condition of Theorem 1.1.15 can obviously be checked in polynomial time. A kernel can be computed in polynomial time by adapting the technique used by Abbas and Saoula [1] for their algorithmic version of Champetier's theorem.

1.2 Main applications

1.2.1 Board games analysis

The notion of a kernel was originally introduced by von Neumann and Morgenstern [74] as an abstract generalization of a concept of solution for strategic games. Kernels represent winning strategies in a Nim game, a strategic game in which two players remove objects turn by turn. On each turn, a player must remove at least one object. Depending on the version being played, the goal of the game is either to avoid or to take the last object.

As an example, consider a simple Nim game, the "matches game." There are n matches, each player removes one, two, or three matches turn by turn and the player removing the last one looses. The number of states in the game is finite. There are n different states n, $n-1$,..., or 1 representing the number of remaining matches.

Let find a winning strategy for the example of the matches game with $n = 12$.

The goal is to leave one matche on the table, to force the adversary to take the last one. To arrive at this state, the precedent winning state is to leave exactly five matches on the table, so the other player leaves between four and two matches and it is always possible to leave only one match. A step further, the precedent winning state is to leave nine matches, so the other player would leave between eight and six matches and it is possible to leave exactly five matches on the table. This analysis leads to the conclusion that if a player starts, he can be sure to win, and the strategy consists of starting by taking three matches, leaving then nine matches on the table, then leaving five and finally only one match.

A way to model a Nim game as a graph consists in considering the digraph where each vertex represents a state and there is an arc from a vertex to another if a valid move in the game leads one state to another. Starting from an initial state, represented by a source, the game evolves to a final state, represented by a sink.

In the example of the matches game, each vertex corresponds to the number of matches remaining on the table and each arc represents the removal of matches from the game, as represented in Figure 1.2.

A kernel in such a graph determines a winning strategy. If the player can achieve a state such that the corresponding vertex is in the kernel, he can win. By stability, the other player will move to a state not in the kernel, and by absorbance, it is then possible to make a move back to the kernel.

Any kernel in a graph modeling a Nim game corresponds to a winning strategy to the game.

1.2.2 Economics and Game Theory

In cooperative game theory, a *hedonic game* is a game that models the formation of coalitions of players when players have preferences over which group they belong to. Formally, it is a pair $(N, (\preccurlyeq_i)_{i \in N})$ of a finite set N of players, and a complete order relation over $\{S \subseteq N : i \in S\}$ for each player $i \in N$ of possible coalitions the player i belongs to. Hedonic games are well studied in economics, modeling human activities such as political organizations, and the focus lies on identifying sufficient conditions for the existence of

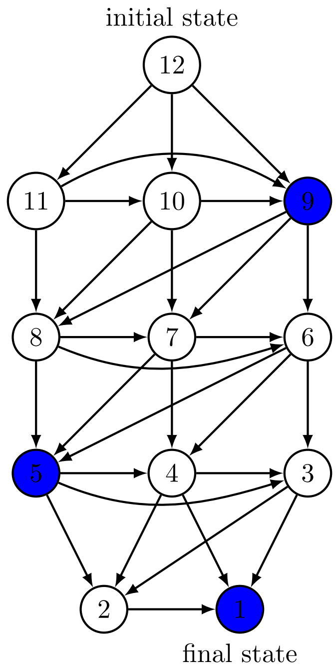

Figure 1.2: Graph corresponding to the matches game with 12 matches, where the kernel is represented in blue.

stable outcomes, i.e., such that there is no subset such that each member would improve its situation by making this subset a new coalition. It is possible to consider the influence of communications between players by defining the model of a hedonic game with a restricted communication structure. A *hedonic game with a restricted communication structure* is a triple $(N, (\preccurlyeq_i)_{i \in N}, \mathcal{F})$ where $(N, (\preccurlyeq_i)_{i \in N})$ is a hedonic game, and $\mathcal{F}$ is a feasible coalition system on N, i.e., a family of subsets of N such that $\{i\} \in \mathcal{F}$ for all $i \in N$ and $\varnothing \notin \mathcal{F}$. A partition π of players into subsets of $\mathcal{F}$ is *core stable* if no other subset of $\mathcal{F}$ can strictly improve the position of all the members of the new subset in its poset. Existence results about core stable outcomes for hedonic games can be formulated as existence results about kernels, as done by Igarashi [51].

Given a hedonic game Γ, we can define a digraph D^Γ with $V(D^\Gamma) = \mathcal{F}$ and $A(D^\Gamma) = \{(S,T) \colon \exists i \in S \cap T, S \not\succcurlyeq_i T\}$. The core stable feasible partitions of a hedonic game Γ are then the kernels of this digraph D^Γ, and the existence and complexity results about finding a core stable feasible partition arise from those about kernels.

1.2.3 Stable matchings

We have already mentioned the stable matching problem in Section 1.1.2.1 It can be stated as follows: given n men and n women, where each person has ranked all members of the opposite sex in order of preference, marry the men and women together such that there are no two people of opposite sex who would both marry each other rather than being married with their current partners. When there are no such pairs of people, the set of marriages is deemed stable. Note that this can be seen as a special case of a hedonic game.

Algorithms designed to address the stable marriage problem find practical use in many real-world scenarios, with one of the most notable examples being their application in matching graduating medical students with their hospital. The original algorithm has been designed in 1962 by David Gale and Lloyd Shapley [42], who established that it is always possible to find a solution to the stable marriage problem and to ensure the stability of all marriages, but its real-world relevance was not recognized until much later.

The study of the very practical allocation problem of assigning hospital to newly examined doctors has been made by Alvin Roth in the 1980's [68]. In 2012, the Nobel Memorial Prize in Economic Sciences was awarded to Lloyd Shapley and Alvin Roth "for the theory of stable allocations and the practice of market design" [28].

The algorithm, known as the Gale–Shapley algorithm, proceeds through a sequence of rounds as follows:

In the first round, each unengaged man proposes to its favorite woman. Then, each woman responds with a "maybe" to her preferred suitor and a "no" to all others. She becomes provisionally engaged to her favorite candidate at that point, and reciprocally, he becomes provisionally engaged to her.

In subsequent rounds, each unengaged man proposes to the most-preferred woman he has not yet approached, independently from her current engagement status. The woman responds with a "maybe" if she is not engaged yet or if she prefers this new candidate over her current provisional partner. In the latter case, she rejects her current partner, making him unengaged. This process continues until all individuals are engaged.

This algorithm is not only efficient but also guarantees a stable marriage. Its complexity is $\mathcal{O}(n^2)$, with n the number of men or women involved in the matching process.

Maffray [64] was the first to recognize that a stable matching is actually the kernel of a digraph. His reformulation of the Gale–Shapley theorem in term of kernels is given in Theorem 1.1.8. The digraph is constructed as follows. Consider the complete bipartite graph $B = (W, M; W \times M)$, where W represents the set of women and M the set of men. Call L the line graph of B. Each woman w corresponds to a clique C_w of L consisting of all vertices (m, w) with $m \in M$; each man m corresponds to a clique C_m of L consisting of all vertices (m, w) with $w \in W$. Now we orient L as follows: for each woman we orient the edges within C_w according to the preference of the woman, i.e., if she prefers man m_1 to man m_2, we orient the edge from vertex (w, m_2) to vertex (w, m_1) and so on.

A kernel in this orientation of L corresponds to a perfect stable marriage. Indeed, a kernel K in L corresponds to a maximal matching (and then maximum) in B and if there were an unstable couple in K then the corresponding vertex would have no successor in K, which contradicts the definition of a kernel.

1.2.4 Logic

Kernels have applications in two subfields of logic, namely symbolic logic and default logic. In both cases, kernels correspond to relevant objects of the subfield.

1.2.4.1 Symbolic logic

In symbolic logic, a theory is a set of sentences in a formal language. A model refers to an interpretation or a way of assigning meaning to the symbols and statements of a theory. As presented in a paper of Michał Walicki [75], from a theory it is possible to build a digraph such that there is a bijection between models of the theory and kernels of the digraph. In particular, a theory has no model, i.e., is *paradoxal* if and only if the corresponding digraph has no kernel. This has permitted to establish new results about paradoxes. Two examples of theories and their corresponding digraphs are provided in Example 1.2.1. Note that because a kernel is an independence set, a vertex a such that $(a, a) \in A(D)$ cannot be in any kernel.

Example 1.2.1 (Walicki [75]). *Let* Θ_1 *be the following theory:*

a. This and the next statement are false.	$a \Leftrightarrow \neg a \wedge \neg b$
b. The next statement is false.	$b \Leftrightarrow \neg c$
c. The previous statement is false.	$c \Leftrightarrow \neg b.$

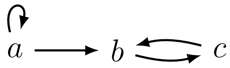

Figure 1.3: The digraph corresponding to the theory Θ_1

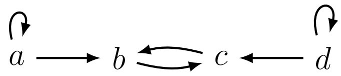

Figure 1.4: The digraph corresponding to the theory Θ_2

Making b true and a and c false is the only model so that Θ_1 involves no paradox. Its corresponding digraph is the one depicted in Figure 1.3. The singleton composed of vertex b is the only kernel of the digraph. Adding the statement

d. This and the previous statement are false. $d \Leftrightarrow \neg d \wedge \neg c$

to Θ_1 gives another theory Θ_2. In the latter, paradox is unavoidable. Its corresponding digraph, depicted in Figure 1.4, has no kernel.

1.2.4.2 Default logic

Default logic is a logic system used, for example, in artificial intelligence. It extends classical logic by allowing for the representation of default rules, which express statements that are assumed to be true by default but can be overridden by additional information. In default logic, there is a distinction between normal rules and default rules, with the latter serving as a kind of "default assumption." A default theory consists of a set of propositions, which can be believed beyond any doubt and a set of default rules, which provide a mechanism to jump to conclusion when there is no conflicting information available. As an example, the default rule formalizing that everything which is not a penguin [1] can be assumed to fly is written ¬penguin ∧ fly/fly. The central idea in default logic revolves around the concept of an "extension" of a default theory, which essentially represents what one can reasonably believe while staying in line with the default theory's principles.

In 1994, Yannis Dimopoulos and Vangelis Magirou [33] noticed the connection between the notions of extension and of kernel. Indeed, from a default theory, they build a digraph such that there is a kernel in it if and only if the default theory has an extension. An example is provided in Example 1.2.2. This result has been used to prove the NP-completeness of finding extensions and to build algorithm providing extensions in special cases.

Example 1.2.2 (Dimopoulos and Magirou [33])**.** *Consider the default theory Δ with the set of default rules $D = \{d_1 = a_1/a_1, d_2 = a_2 \wedge \neg a_1/a_2, d_3 = a_3 \wedge \neg a_1 \wedge \neg a_4/a_3, d_4 = a_4 \wedge \neg a_2 \wedge \neg a_6/a_4, d_5 = a_5 \wedge \neg a_3/a_5, d_6 = a_6 \wedge \neg a_5/a_6\}$ The digraph corresponding to this default theory is represented in Figure 1.5. The kernels of this digraph ($\{d_1, d_4, d_5\}$ and $\{d_2, d_3, d_6\}$) correspond to the extensions of Δ (conclusions of d_1, d_4, and d_5, i.e., $\{a_1, a_4, a_5\}$ and conclusions of d_2, d_3, and d_6, i.e., $\{a_2, a_3, a_6\}$).*

[1] La traduction française de "penguin" est "manchot." En effet, comme rappelé dans le livre "Le manchot qui en avait marre d'être pris pour un pingouin" de Nicolas Digard et Christine Roussey, les manchots, contrairement aux pingouins, ne volent pas.

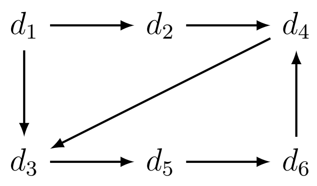

Figure 1.5: The digraph corresponding to the theory default theory Δ

1.2.5 Dinitz

A nice application of kernels in graph theory is the following theorem, which was originally conjectured by Jeff Dinitz in the seventies [40]. It was not solved until the nineties. The original problem has been formulated in terms of a latin square problem but can be reformulated into a graph theory problem.

Given a graph G and a set $C(v)$ of "allowed" colors for each vertex v (called a *list*), a *proper list coloring* is a choice function that maps every vertex v to a color in the list $C(v)$ such that the resulting coloring is proper.

The *list chromatic number* of a graph G, denoted by $\chi_\ell(G)$, is the minimum k such that if every color set $C(v)$ has size k, then a proper list coloring exists.

Theorem 1.2.3 (Dinitz theorem [46]). *The list chromatic number of the complete bipartite graph $K_{n,n}$ is n.*

The following easy statement relates colorings of a graph and kernels in some orientations. Combined with the Gale–Shapley theorem (Theorem 1.1.8), it is a key step in the proof of Theorem 1.2.3.

Lemma 1.2.4. *Let D be a directed graph, and suppose for all $v \in V(D)$ there is a color set $C(v)$ such that $|C(v)| \geqslant |N^+(v)| + 1$. Then, if every induced subdigraph of D has a kernel, G has a proper list coloring with respect to $(C(v))_v$.*

The proof of Theorem 1.2.3 can be turned into a polynomial-time algorithm.

1.3 Main challenges

Clearly, the next two questions are the main current challenges of the theory of kernels. In the two cases, the existence of kernel is ensured by a theorem, but a way to compute efficiently such a kernel is not known. Note that the well-posedness of the questions can be debated since for the first one, it is not known whether the condition can be checked in polynomial time, and for the second, it is a coNP-complete problem (Section 1.1.2.2 above). These aspects of the questions are discussed a bit further in Chapters 2 and 3.

Open question 1.3.1. *What is the complexity of computing a kernel in a digraph satisfying the condition of the theorem Galeana-Sánchez–Neumann-Lara (Theorem 1.1.7)?*

Open question 1.3.2. *What is the complexity of computing a kernel in a clique-acyclic orientation of a perfect graph?*

There are many special cases that are still open such as:

- What is the complexity of computing a kernel in a digraph such that every odd directed cycle has at least two crossing consecutive chords (defined in Chapter 3)? (Special case of open question 1.3.1.)

- What is the complexity of computing a kernel in a clique-acylic orientation of a perfectly orderable graph? (Special case of open question 1.3.2.)

There are other open questions regarding the complexity of computing kernels not requiring clique-acylicity, and thus no result of Section 1.1.2.2 applies on such a digraph. Some of them look accessible, like the following one.

Open question 1.3.3. *What is the complexity of deciding whether an orientation of an interval graph admits a kernel?*

Andres and Hochstättler have asked the same question for the more general class of perfect graphs, but, as far as we know, even for the very special case of interval graphs, the answer is not known. However, when the orientation is simple, deciding the existence of a kernel and computing one if it exists can be done in polynomial time for interval graphs, and even for chordal graphs [52].

Chapter 2

Some thoughts about kernel-solvability

A notion related to that of kernel-perfection is that of kernel-solvability. It involves clique-acyclic orientations, which have been defined in Section 1.1.2. For sake of convenience, we recall the definition: a digraph is *clique-acyclic* if every clique admits a kernel (i.e., a vertex absorbing all other vertices of the clique). A graph is *kernel-solvable* if every clique-acyclic orientation admits a kernel. The reader can refer to Figures 2 for the correspondence between these notions.

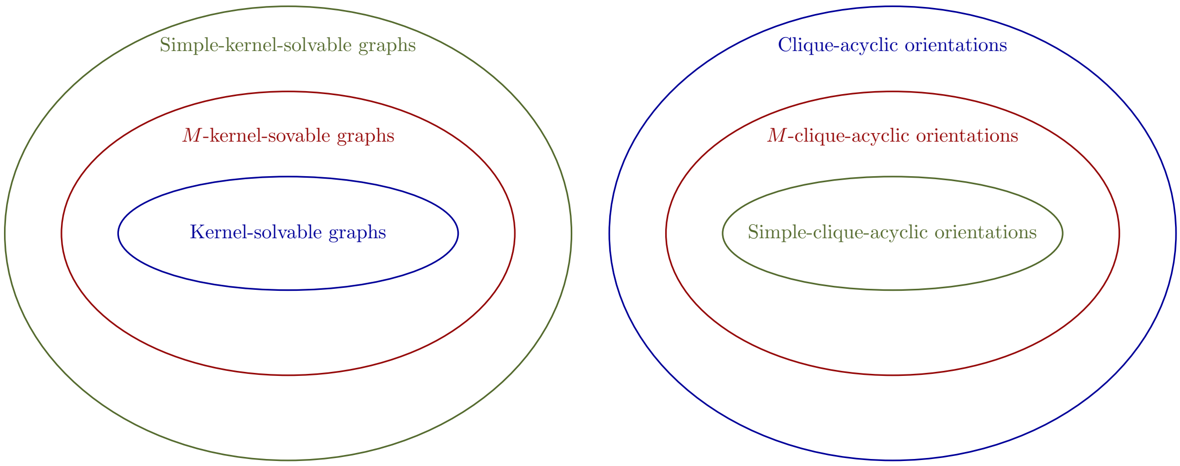

Figure 2.1: Inclusions between different graphs and orientations.

The conjecture by Berge and Duchet, mentioned in the Introduction, states that a graph is perfect if and only if it is kernel-solvable. It is an easy exercise to show that kernel-solvability is a property closed under taking induced subgraphs. Hence, a graph is kernel-solvable if and only if every clique-acyclic orientation is kernel-perfect.

The conjecture of Berge and Duchet is now proved, and this shows that a deep relation exists between perfection and kernel-perfection.

The 'if' direction is a consequence of the strong perfect graph theorem. Indeed, it is not too hard to check that there exist clique-acyclic orientations of odd holes and of odd anti-holes with no kernel. This will be further discussed in Section 2.1. It is worth noting that no other proof of that direction of the conjecture is known.

The 'only if' direction is the following theorem.

Theorem 1.1.14 (Boros and Gurvich [18])**.** *Every clique-acyclic orientation of a perfect graph admits a kernel.*

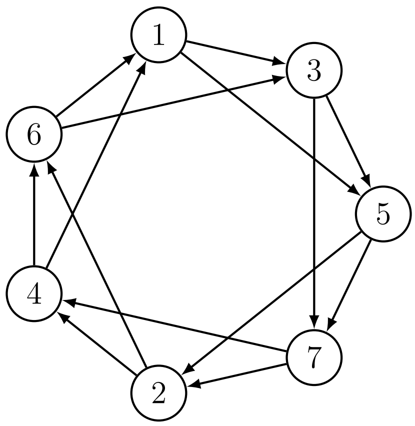

Figure 2.2: A simple clique-acyclic orientation of $\overline{C}_7$ with no kernel.

Boros and Gurvich proved their theorem with advanced notions of game theory. A simpler proof based on Scarf's lemma was later proposed by Aharoni and Holzman [2] and then simplified by Király and Pap [57] using Sperner's lemma. However, none of these proofs provides algorithms for computing a kernel. This raises the question of the complexity of kernel computation in a clique-acyclic orientation of a perfect graph. This question is further discussed in Section 2.2.

Actually, since every class appearing in Section 1.1.2.1 is closed under taking induced subdigraphs, the mentioned theorems imply the kernel-perfection of the digraphs of these classes.

Several graphs operations are known to preserve perfection and the proof of the strong perfect graph theorem relies on such operations. Similarly, Blidia et al. [13] showed that some of these operations preserve kernel-perfection. We will see in Section 2.3 that some also preserve the polynomiality of computing a kernel.

2.1 Kernels and odd anti-holes

We start with the following observation.

Observation 2.1.1. *The odd anti-hole $\overline{C}_7$ admits a simple clique-acyclic orientation with no kernel.*

Proof. Let D be the digraph with vertex set $V(D) \coloneqq \mathbb{Z}_7$, and with the following arcs set:

$$A(D) \coloneqq \{(i, i+2), (i, i+4) \colon i \in \mathbb{Z}_7\}.$$

This is a simple clique-acyclic orientation of $\overline{C}_7$. See Figure 2.2. This digraph has no kernel since every maximal independent set is a pair of the form $\{i, i+1\}$, which does not absorb $i+2$. □

We believe that Observation 2.1.1, although very simple, may be helpfull. Actually Boros and Gurvich [19, Observation 3.13] wrote that every M-clique-acyclic orientation of $\overline{C}_7$ (and thus in particular every simple clique-acyclic orientation) admits a kernel. This makes them claim that kernel-solvability and kernel-M-solvability are distinct. András Sebo actually noticed that even though they putted the common definition of M-clique-acyclic orientation in their paper, they were actually using a more constrainted one, requiering for the orientation that every non-directed cycle of length three has at least hence-and-forth pairs of arcs. In this manuscript, we will not use this other definition. Kernel-solvability has been defined in the beginning of that chapter. A graph is *M-kernel-solvable* if every M-clique-acylic orientation admits a kernel. This notion has been introduced by Duchet [37]. To this notions, we add the following one: a graph is *simple kernel-solvable* if every simple clique-acylic orientation admits a kernel. The connection between the previous notions is represented in Figure 2.

With the help of the strong perfect graph theorem, it is not too difficult to see that kernel-solvable graphs are perfect [19]. The key element here is to exhibit for every odd anti-hole a clique-acyclic orientation with no kernel. The next theorem shows that there is no hope to achieve the same result for simple clique-acyclic orientations, and that kernel-solvability and simple kernel-solvability are distinct notions, as implicitly claimed by Boros and Gurvich. We emphasize that the exact location of kernel-M-solvability with respect to simple kernel-solvability and kernel-solvability remains to investigate.

Theorem 2.1.2. *Let D be a simple clique-acyclic orientation of an odd anti-hole $\overline{C}_{2k+1}$. If $k \geqslant 4$, then D admits a kernel.*

Before giving the proof, we need preliminary lemmas.

In the next statement, any path has distinct endpoints.

Lemma 2.1.3. *Let C be a Hamiltonian cycle of a clique K_n with $n > 4$. Consider a subgraph L of K_n formed by pairwise vertex-disjoint paths and sharing no edge with C. Then it is possible to add edges to L to get a single path P sharing no edge with C and of length at least $n-2$. In case it is not possible to get a single path of length $n-1$, it is possible to ensure that the vertex missed by P is adjacent on C to both endpoints of P.*

Proof. We prove it by (decreasing) induction on the number of edges in L. The maximal number of edges L can have is $n-1$. So suppose that L has $n-1$ edges. In this case, it covers all vertices by one single path of length $n-1$. Suppose now that L has less than $n-1$ edges.

If it is formed by two paths or more, then it is possible to find one endpoint from one path and another endpoint from another path that are not adjacent on C. In that case, it is then possible to add an edge (not in C) to L between these endpoints so that the new subgraph is still formed by pairwise vertex-disjoint paths and shares no edge with C, and induction applies.

If it is formed by a single path, but misses at least two vertices of K_n, then one endpoint at least is not adjacent to a missed vertex on C. In that case again, it is then possible to add an edge (not in C) to L between these vertices so that it is still formed by a single path and shares no edge with C, and induction applies.

If it is formed by a single path, but misses exactly one vertex of K_n, either this missed vertex is adjacent to both endpoints of the path, and we are done, or we can add an edge (not in C) between this missed vertex and an endpoint of the path, so that we get a single path covering all the vertices. □

Lemma 2.1.4. *Let Γ be a directed cycle of length at least 7 in a simple clique-acyclic orientation of the odd anti-hole $\overline{C}_{2k+1}$ with $k \geqslant 4$. Then Γ has at least two chords with consecutive heads.*

Proof. Let $\ell \geqslant 3$ be such that the length is $2\ell+1$. Denote by H the (undirected) subgraph of $\overline{C}_{2k+1}$ induced by the vertices of Γ.

Suppose first that $\ell = k$. Then the chords of H form a $(2\ell-4)$-regular graph with $V(H)$ as vertex set (because the non-edges form a Hamiltonian cycle). Consider the arcs A' in D corresponding to the edges of this regular spanning subgraph of H. There are two consecutive vertices on Γ that are heads of arcs in A' (otherwise, we could locate $(\ell-2)(2\ell+1)$ heads among at most $(2\ell-4)\ell$ possible locations).

Suppose then that $\ell < k$. Then we apply Lemma 2.1.3 with $n = 2\ell + 1$, with C being the underlying undirected cycle of Γ, and with L being the non-edges of $\overline{C}_{2k+1}$ with both endpoints on C. In that case, there are three possibilities.

First possibility, it is possible to complete L with edges from $E(H) \setminus E(C)$ so as to get a path of length $n-1$, with endpoints that are non-adjacent on C. Then again the chords of H contain a $(2\ell-4)$-regular graph with $V(H)$ as vertex set (because L can then even be completed so as to get a cycle of length n, by adding the edge of H between the endpoints of L).

Second possibility, it is possible to complete L with edges from $E(H) \setminus E(C)$ so as to get a path of length $n-1$, with endpoints that are adjacent on C. The chords of H contain then a graph with $2\ell - 1$ vertices of degree $2\ell - 4$ and 2 vertices of degree $2\ell-3$ adjacent on C. Consider the arcs A' in D corresponding to the edges of this regular spanning subgraph of H. There are two consecutive vertices on Γ that are heads of arcs in A' (otherwise, we could locate $(\ell-2)(2\ell-1) + 2\ell - 3 = 2\ell^2 - 3\ell - 1$ heads among at most $(\ell-1)(2\ell-4) + 2\ell - 3 = 2\ell^2 - 4\ell + 1$ possible locations).

Third possibility, it is possible to complete L with edges from $E(H) \setminus E(C)$ so as to get a path of length $n-2$, with both endpoints adjacent on C to the vertex missed by L in H. The chords of H contain then a graph with $2\ell - 2$ vertices of degree $2\ell - 4$ and 2 vertices of degree $2\ell-3$, both adjacent on C to one vertex of degree $2\ell-2$. Consider the arcs A' in D corresponding to the edges of this regular spanning subgraph of H. There are two consecutive vertices on Γ that are heads of arcs in A' (otherwise, we could locate $(\ell-2)(2\ell-2)+2\ell-3+\ell-1 = 2\ell^2-3\ell$ heads among either at most $(\ell-1)(2\ell-4)+2\ell-2 = 2\ell^2-4\ell+2$ possible locations, or at most $(\ell-2)(2\ell-4)+2(2\ell-3) = 2\ell^2-4\ell+2$ possible locations). □

Proof of Theorem 2.1.2. Denote by $v_1, \ldots, v_{2k+1}$ the vertices of the odd anti-hole, so that $v_i v_{i+1}$ is a non-edge for all i (and with $2k+2 \coloneqq 1$).

We will apply the proof technique of the Galeana-Sánchez–Neumann-Lara theorem 1.1.7, but we choose in addition the vertex u with an extra property. To do so, we first establish that there is a v_i such that the edges $v_i v_{i-2}$ and $v_i v_{i+2}$ are both oriented towards v_i.

Suppose for a contradiction that for every i, the arc (v_i, v_{i-2}) or the arc (v_i, v_{i+2}) exists. Without loss of generality, we can assume that the arc (v_1, v_3) exists. Then (v_3, v_5), (v_5, v_7), ..., (v_{2k+1}, v_2), (v_2, v_4), etc. exist as well, i.e., (v_i, v_{i+2}) exists for all i. Since the orientation is clique-acyclic, we get that (v_i, v_{i+4}) exists for all i. Repeating this argument, we get that an edge $v_i v_j$ with $i < j$ and $2 \leqslant j - i \leqslant 2k - 1$ is oriented (v_i, v_j) precisely when $j - i$ is even. Since $k \geqslant 4$, the arc (v_1, v_7) exists in D. The arcs (v_7, v_4) and (v_4, v_1) also exist; this contradicts the simple clique-acyclic orientation of D.

Let u be the vertex v_i such that the edges $v_i v_{i-2}$ and $v_i v_{i+2}$ are both oriented towards v_i. We know that $D - N^-[u]$ admits a kernel K, because $\overline{C}_{2k+1}$ is minimally imperfect and the Boros–Gurvich theorem applies. We follow now exactly the proof of the Galeana-Sánchez–Neumann-Lara theorem, up to the step where it is shown that u does not belong to S.

So, suppose for a contradiction that u belongs to S. We keep the definition of P and v, and we still have that P together with (u, v) is a directed cycle Γ of odd length. By Lemma 2.1.4, if Γ is of length at least 7, then Γ has two chords with consecutive heads, which is not possible as detailed in the proof of the Galeana-Sánchez–Neumann-Lara theorem. The cycle Γ cannot be of length 3 by the clique-acyclicity of D.

We are left with the case when Γ is of length 5. Denote by H the (undirected) subgraph of $\overline{C}_{2k+1}$ induced by the vertices of Γ. Note that H is formed by the underlying undirected cycle of Γ, plus chords. If it has two chords with consecutive heads, we are done, as above. Actually, an easy case-checking shows that if it has two chords, the clique-acyclic condition does imply that there are consecutive heads. Since there is at least one chord (otherwise,

$k = 2$), there is exactly one chord. The complement of H is thus a path of length 4. The vertices of H are therefore of the form $v_j, v_{j+1}, \ldots, v_{j+4}$. Without loss of generality, the arcs of Γ are

$$(v_{j+2}, v_j), \quad (v_j, v_{j+3}), \quad (v_{j+3}, v_{j+1}), \quad (v_{j+1}, v_{j+4}), \quad (v_{j+4}, v_{j+2}),$$

and the chord is (v_{j+4}, v_j) by the clique-acyclicity of D. Since by definition $u = v_i$ has no outneighbor among $\{v_{i-2}, v_{i+2}\}$, the vertex u is necessarily v_j or v_{j+1}. Assume first that $u = v_j$. By construction of Γ, we have v_{j+3} and v_{j+4} in K. But $v_{j+4} \in N^-(u)$ contradicts the definition of K. Then, assume that $u = v_{j+1}$. By construction of Γ, we have v_{j+4} and v_j in K. But this contradicts the fact that K is an independent set.

This shows that u does not belong to S. Finishing exactly as in the proof of the Galeana-Sánchez–Neumann-Lara theorem, we get that S is a semi-kernel of D. Since $\overline{C}_{2k+1}$ is minimally imperfect, every proper subdigraph of D also admits a semi-kernel, because it admits a kernel by the Boros–Gurvich theorem. Lemma 1.1.1 shows then that D admits a kernel. □

2.2 Computing kernels in clique-acyclic orientations of perfect graphs

The theorem of Boros and Gurvich raises the question of the computation of a kernel in a clique-acyclic orientation of a perfect graph. This question has been identified as one the main questions about kernels at the end of Chapter 1 (Open question 1.3.2). As mentioned there, a formal definition of such a computational task is problematic because, even if deciding perfection is polynomial [24], deciding whether an orientation of a perfect graph is clique-acyclic is coNP-complete [7]. Therefore, a clique-acyclic orientation of a perfect graph cannot be a polynomial-time recognizable instance of a computational problem (assuming P $\neq$ NP). Even finding a polynomial certificate of the correctness of such an instance is probably not possible because it is unlikely that a coNP-complete problem is in NP. However, deciding whether an orientation is an M-clique-acyclic orientation is obviously in P, which makes the following a well-posed open question.

Open question 2.2.1. *What is the complexity of computing a kernel in an M-clique-acyclic orientation of a perfect graph?*

Since the theorem of Boros and Gurvich ensures that a kernel always exists for such digraphs, this computational problem belongs to TFNP. As it has been noted in the introduction of the chapter, none of the known proofs provides a concrete algorithm. It is not even known whether the problem belongs to one of the studied subclasses of TFNP. The proofs based on Scarf's lemma and Sperner's lemma might indicate at first glance a membership to the PPAD class. This is actually correct, but only if the cliques have bounded size; see [56], keeping in mind that a strong fractional kernel in a perfect graph can be turned into a kernel in polynomial time. (In the general case, the reduction to End-Of-Line implied by those proofs does not allow polynomial moves.)

For some subclasses of perfect graphs, the problem has been proved to be polynomial. Here are the known results. When the most general case of clique-acyclic orientations is covered, the result is stated at that level of generality.

- Clique-acyclic orientations of chordal graphs [52]. (In this case, such orientations can be recognized in polynomial time.)
- Simple clique-acylic-orientations of perfect claw-free graphs [52]. Line graphs of bipartite graphs are a special case (which coincides stable marriages).

- Clique-acylic orientations of bipartite graphs (easy result).
- Clique-acylic orientations of DE graphs [32, 52]. In this case, such orientations can be recognized in polynomial time.
- M-clique-acyclic orientations of comparability graphs [1].

2.3 Graph operations preserving the existence of kernels

At least four graph operations are known to preserve kernel-perfection: disjoint union (trivial), addition of a pending vertex (easy), clique-sum, and join. For the clique-sum, this is a result by Jacob [53], and for the join, this is a result by Blidia et al. [13]. For this latter operation, there is an extra (but natural) condition regarding the orientation of the new arcs. The first three operations preserve the polynomiality of kernel computation. By this, we mean that given a class $\mathcal{D}_0$ of kernel-perfect digraphs for which kernel computation is polynomial, its closure $\mathcal{D}$ by the operation is still a class of kernel-perfect digraphs for which kernel computation is polynomial. Again, for the disjoint union and the addition of a pending vertex, this is easy. For the clique-sum, this is a result by Pass-Lanneau et al. [52]. We establish a similar polynomiality result for the join operation. This work benefits from preliminary results of a research project done by Benjamin Pyryt and Marie Temple-Boyer at École des Ponts ParisTech in 2020.

The *join* of two vertex-disjoint (undirected) graphs G_1, G_2 is the (undirected) graph, denoted by $G_1 * G_2$, such that

$$V(G_1 * G_2) \coloneqq V(G_1) \cup V(G_2)$$
$$E(G_1 * G_2) \coloneqq E(G_1) \cup E(G_2) \cup \{v_1v_2 \colon v_1 \in V(G_1), v_2 \in V(G_2)\}.$$

A *join* of two vertex-disjoint digraphs D_1 and D_2 is an orientation D of the join of the underlying undirected graphs of D_1 and D_2 such that $D[V(D_1)] = D_1$ and $D[V(D_2)] = D_2$.

Theorem 2.3.1. *Let $\mathcal{D}_0$ be a family of kernel-perfect digraphs for which the computation of a kernel can be done in polynomial time. Let $\mathcal{D}$ be the family of digraphs obtained from $\mathcal{D}_0$ by repeated join and disjoint-union operations, with the following condition on the join operation: every directed cycle of length three intersecting both operands of the join has at least two hence-and-forth pairs of arcs.*

Then $\mathcal{D}$ is a family of kernel-perfect digraphs for which the computation of a kernel can be done in polynomial time.

Proof. Blidia et al. [13] have actually shown the following. Let D_1 and D_2 be two kernel-perfect digraphs, and D be a join of D_1 and D_2 satisfying the condition of the statement on the length three directed cycle. Let K_1 be a kernel of D_1. Define $Y \coloneqq \{v \in V(D_2) \colon V(D_1) \subseteq N^-(v)\}$, and let K_Y be a kernel of $D_2[Y]$ and K_Z be a kernel of $D_2\big[V(D_2) \setminus N^-[K_Y]\big]$. (We keep the notation from Blidia et al.) Then K_1 or $K_Y \cup K_Z$ is a kernel of D.

Suppose that the complexity of computing a kernel in a digraph from $\mathcal{D}_0$ is $\mathcal{O}(n^\alpha)$ with $\alpha \geqslant 3$, where n is the number of vertices. We prove by induction on n that the complexity of computing a kernel in a digraph from $\mathcal{D}$ is $\mathcal{O}(n^\alpha)$. Since finding two graphs such that its join form $\mathcal{D}$ (just consider the complement of the components of the complement of $\mathcal{D}$) and computing Y can be done in quadratic time, denote C the constant such that those operations are doable in less than Cn^2 operations. We denote by $f(n)$ the complexity of finding a kernel in a graph of $\mathcal{D}$ of size n and by n_1 (resp. n_2 and y) the size of D_1 (resp. D_2 and Y). For every digraph of $\mathcal{D}_0$, by definition, $f(n) \leqslant Cn^\alpha$. By induction, consider a digraph of size n in $\mathcal{D}$, we have then the following:

$$\begin{aligned} f(n) &\leqslant Cn^2 + f(n_1) + f(y) + f(n_2 - y) \\ &\leqslant Cn^2 + Cn_1^\alpha + Cy^\alpha + C(n_2 - y)^\alpha \\ &\leqslant Cn^2 + Cn_1^\alpha + Cn_2^\alpha \\ &\leqslant C(2n_1^2 + 2n_2^2) + Cn_1^\alpha + Cn_2^\alpha \\ &\leqslant C(\alpha n_1^{\alpha-1} + \alpha n_2^{\alpha-1}) + Cn_1^\alpha + Cn_2^\alpha \\ &\leqslant C(n_1 + n_2)^\alpha. \end{aligned}$$ □

A direct application of the previous theorem is the following corollary about tree-cographs. A graph is a *tree-cograph* if it can be constructed from trees by disjoint union and complement operations.

Corollary 2.3.2. *A kernel in an M-clique-acyclic-orientation of a tree-cograph can be computed in polynomial time.*

Proof. The class of tree-cographs is closed under taking induced subgraphs and the class can be obtained from trees and complement of trees by repeated join and disjoint union. Since a tree is acyclic and the stable sets of a complement of a tree are of size at most two, it is clearly polynomial to compute a kernel in any orientation of a tree or a complement of a tree. Using Theorem 2.3.1, we get the result. □

In their paper, Blidia et al. [13] also consider the *duplication* operation, which is an important operation of graphs. *Duplicating* a vertex u in a (undirected) graph G consists in adding a new vertex v with $N(v) \coloneqq N(u)$. When an edge uv is moreover added, then we talk about *adjacent* duplication. Otherwise, it is a *non-adjacent* duplication.

It is not clear how a directed version of duplication would preserve kernel-perfection. Blidia et al. [10] have shown that non-adjacent duplication preserves kernel-solvability and that adjacent duplication of a kernel-solvable graph leads to a graph with a kernel up to an extra condition on the orientation. Now that Theorem 1.1.14 has been proved, these results are just a consequence of the identification between kernel-solvable graphs and perfect graphs. However, their proof technique is relevant to establish some preservation of polynomial-time computation of kernels by duplication. We use it in the proof of the following result. A graph is *distance-hereditary* if it can be constructed from a single vertex by pending (operation consisting in adding a vertex adjacent to exactly one other vertex) and duplicating (operation consiting in adding a vertex having the same neighbourhood as an other vertex) operations. Distance-hereditary graphs are perfect graphs [62].

Proposition 2.3.3. *A kernel in a simple clique-acyclic-orientation of a distance-hereditary graph can be computed in polynomial time.*

Proof. Consider a clique-acyclic orientation of a distance-hereditary graph D. The following polynomial recursive procedure provides a kernel in D. Check if D is the result of a pending or a adjacent or non-adjacent duplicating operation.

Suppose the existence of a pending vertex $v \in V(D)$. If v is a sink, then consider K a kernel of $D - N^-[v]$, and $K \cup \{v\}$ is a sink of D. Else, consider K a kernel of $D - \{v\}$. If $K \cap N^+(v) \neq \varnothing$, K is a kernel of D, otherwise $K \cup \{v\}$ is a kernel of D.

Suppose now the existence of a duplication $v \in V(D)$ of a vertex $w \in V(D)$. Adapting the proofs of Blidia et al. [13] in the special case of simple orientations, we partition the set $N^-(v) \cup N^+(v) (= N^-(w) \cup N^+(w))$ into $X_1 = N^+(v) \cap N^+(w)$ and $X_2 = N^-(v) \cup N^-(w)$. Let us distinguish the cases of adjacent and non-adjacent duplications.

If the duplication is non-adjacent, consider the digraph D' obtained from $D - \{v\}$ by reorienting every arc between w and X_2 so that every vertex in X_2 in an outneighbor of w. We claim that D' is clique-acylic. Indeed, by contradiction, consider a new directed triangle. It is necessarily a directed triangle with $x \in X_1$, $y \in X_2$, and w resulting from the reorientation of the arc between y and w. This implies that $(y, v) \in A(D)$ and the vertices y, v and x form a directed triangle in D, which contradicts the definition of D. Consider then a kernel K of D'. If $w \in K$, one can check that $K \cup \{v\}$ forms a kernel of D, and if $w \notin K$, necessarily $K \cap X_1 \neq \varnothing$ and K is a kernel of D.

Now, without loss of generality, suppose that $(v, w) \in A(D)$. Then, by assumption on the orientation of D, every vertex y in X_2 is such that $(y, w) \in A(D)$. Consider now $D' = D - \{v\}$ and a kernel K of D'. One can check that K is a kernel of D. □

Chapter 3

The "two-chord" condition for odd directed cycles

The main contribution of this chapter is a streamlined version of the proof of the Galeana-Sánchez–Neumann-Lara theorem. This theorem was already introduced in Section 1.1 of Chapter 1. We state here again for sake of readability.

Theorem 1.1.7 (Galeana-Sánchez and Neumann-Lara [44])**.** *Let D be a directed graph. Suppose that in D each directed cycle of odd length has at least two chords with consecutive heads. Then D has a kernel.*

A special case of this condition is the "two crossing consecutive chords-condition." Two chords in a directed cycle are *crossing consecutive* if they are of the form (u_1, u_3), (u_2, u_4) with the vertices u_1, u_2, u_3, u_4 being consecutive in this order on the cycle. Before Galeana-Sánchez and Neumann-Lara have proved their theorem, Duchet and Meyniel [38] have established an intermediary generalization of Richardson's theorem (Theorem 1.1.4), which states that a directed graph in which every odd cycle has "two crossing consecutive chords" admits a kernel.

On the contrary, we could also look for generalizations of Theorem 1.1.7. Our simplification of the original proof will make clear that there are indeed such generalizations. This is done in Section 3.2, where we gather these generalizations in Theorem 3.2.1. Conjecturing that the existence of two chords in every odd directed cycle is enough to ensure the existence of a kernel is very tempting, and Meyniel formulated it in 1976. This conjecture has been then disproved by Galeana-Sánchez [43]. She actually established the following result, with an explicit construction.

Proposition 3.0.1 (Galeana-Sánchez [43])**.** *For every $k \geqslant 2$, there exists a digraph with no kernel and such that each odd directed cycle has at least k chords.*

The construction for $k = 2$ is given in Figure 3.1.

3.1 Proof

Proof of Theorem 1.1.7. The proof works by induction on the number of vertices.

The theorem is obviously correct for the graph reduced to a single vertex. Assume now that D has at least two vertices, and let u be any vertex of D. By induction, $D \setminus N^-[u]$ has a kernel K, and set $K' = K \cup \{u\}$. If K does not contain any vertex of $N^+(u)$, then K' is a kernel of D.

We can thus assume that K contains at least one vertex of $N^+(u)$, and we denote by I the set $K \cap N^+(u)$. Notice that I is an independent set. Let S be the set of vertices

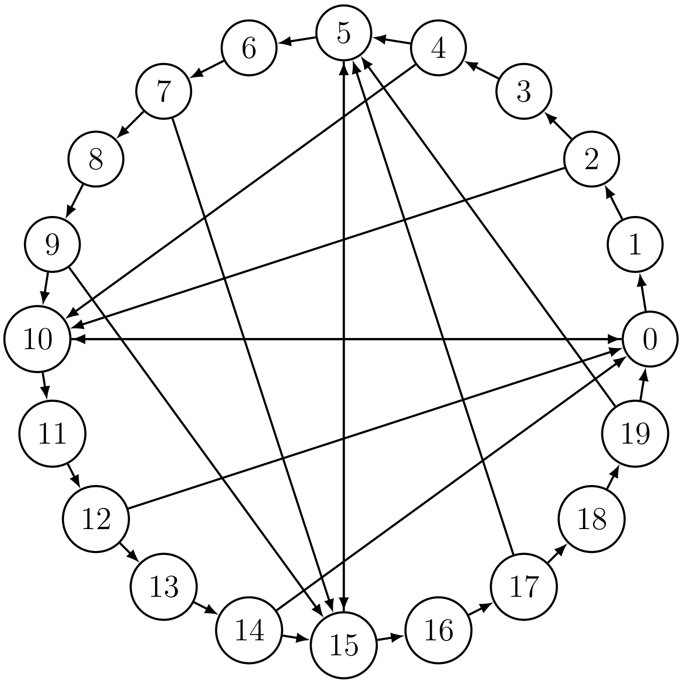

Figure 3.1: The counterexample of Meyniel's conjecture.

v of K' such that there exists a directed path from a vertex in I to v that satisfies the following two conditions:

(i) the path alternates between vertices in K' and vertices not in K'.

(ii) for all vertices w and w' such that w comes before w' on the path and such that w' is not the end vertex of the path, if $w \in N^+(w')$, then $w \notin K'$.

Note that S is non-empty and contains in particular I. We claim that S is actually a semi-kernel.

To show this, we first prove that u is not an element of S. Suppose for a contradiction that u belongs to S. Then there is a directed path P from I to u satisfying the two conditions (i) and (ii). Choose such a path of minimum length and denote by v its origin. It is of even length because of condition (i). Together with the arc (u, v), the path P forms a directed cycle of odd length. This cycle has two chords with consecutive heads by assumption. In particular, it has a chord whose head is in K'. The tail is not in K: if the head is u, it is by definition of K; if the head is not u, it is by independence of K. The tail is not u either: it would contradicts the minimality of P. Moreover, the tail cannot come before the head on P since this would contradict the minimality of P. Therefore, the tail comes after the head, which is in K. This contradicts condition (ii).

We show now that S is a semi-kernel. Since S does not contain u, it is a subset of K and is thus an independent set. Consider now a vertex s of S with an outneighbor w. Note that w is not in K and is distinct from u. By definition of s, there is a directed path Q from I to s satisfying the two conditions (i) and (ii). If w is in $N^-(u)$, then there is an arc from w to a vertex of $V(Q) \cap K$ since otherwise the path obtained by appending (s, w) and (w, u) to Q would satisfy the two conditions (i) and (ii) and this would imply that u belongs to S (indeed, by independence of K, no arc forbidden by (ii) can begin at s). If w is not in $N^-(u)$, then there is an arc from w to a vertex t in K, because this latter set is a kernel of $D \setminus N^-[u]$. Either there is such a t on Q and then $t \in S$, or considering any such t and appending (s, w) and (w, t) to Q satisfies the two conditions (i) and (ii), so $t \in S$. In any case, there is an arc from w to an element of S, and S is therefore a semi-kernel.

We have proved that D has a semi-kernel. By induction, every proper induced subdigraph of D admits a kernel, and thus a semi-kernel. Lemma 1.1.1 leads to the conclusion. □

3.2 Another result

The same proof actually shows that the following generalization of Theorem 1.1.7 is also true.

Theorem 3.2.1. *Let D be a digraph. Suppose that in D each odd directed cycle has*

- *two chords with consecutive heads, or*
- *two non-crossing chords of odd length "in the same direction" (see Figure 3.2), or*
- *two crossing chords, one being short and the other of odd length (see Figure 3.3).*

Then D has a kernel.

Proof. The only part in the proof of Theorem 1.1.7 where the condition about odd directed cycles is used is when we show that u does not belong to S. This is the part that we are going to adapt.

We still suppose for a contradiction that u belongs to S, and we keep the same definition for P and v. In the proof of Theorem 1.1.7, we have checked that the odd directed cycle formed by P and (u, v) cannot have two chords with consecutive heads. So, we are left with the possibilities offered by the two other items. Both involve an odd chord.

Consider such an odd chord (w, w') of the odd directed cycle formed by P and (u, v). The vertex w is distinct from u since otherwise this would contradict the minimality of P. The vertex w' is also distinct from u for the same reason. Moreover, the vertices w, w', and u cannot be in this order on the directed cycle, again for the same reason of minimality of P.

Thus, the vertices w and w' are such that w, u, and w' are distinct and in this order on the cycle. This shows that we cannot be in the possibility offered by the second item: we cannot have two non-crossing chords of odd length in the same direction. Moreover, the configuration of w and w' with respect to u makes that w and w' are both outside K'. But then any short crossing chord as in Figure 3.3 would connect two elements of K', which is impossible. □

It would be interesting to see whether we could further generalize Theorem 3.2.1 by adding to the list of possibilities for an odd directed cycle that of Theorem 1.1.6 (two hence-and-forth pairs of arcs).

3.3 Algorithmic considerations

3.3.1 Checking the condition

As mentioned in Section 1.1.1.2, we do not know whether the condition of the Galeana-Sánchez–Neumann-Lara theorem can be checked in polynomial time. However, we can prove that the less general condition of the special case by Duchet and Meyniel can be checked in polynomial time.

Proposition 3.3.1. *Deciding whether every odd directed cycle has two crossing consecutive chords can be done in polynomial time.*

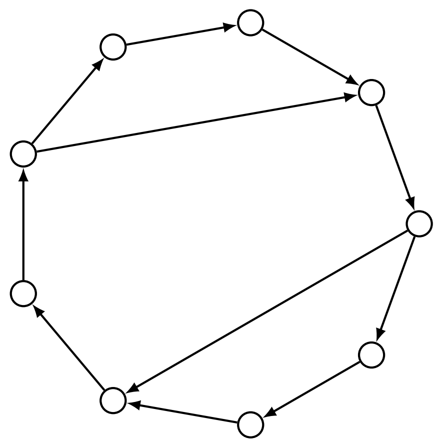

Figure 3.2: An odd directed cycle having two chords in the same direction of odd length.

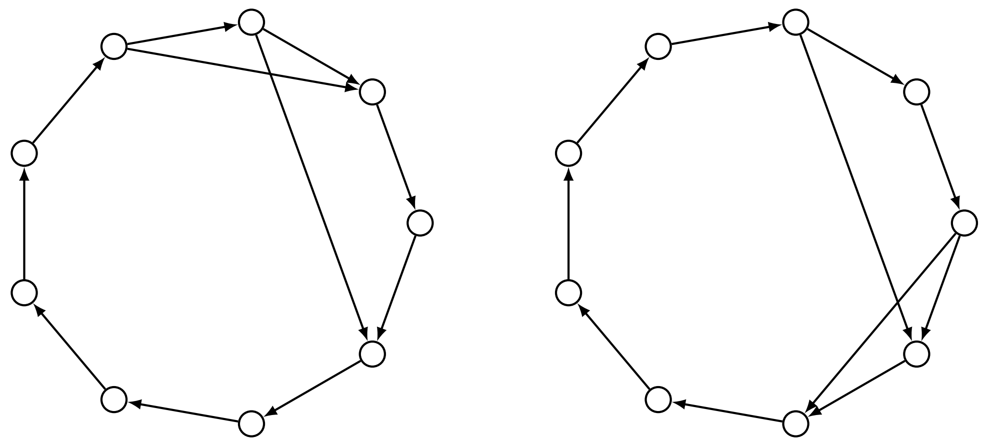

Figure 3.3: Odd directed cycles having two crossing chords, one short and the other of odd length.

Proof. Consider a digraph D and build D' as follows: Each quadruplet of vertices (a, b, c, d) in D such that there is a path of length three $a \to b \to c \to d$ without crossing chords becomes a vertex in D'; Put an arc between two vertices x, y in D' if and only if the last three vertices corresponding to x are the first three vertices corresponding to y.

Then, each odd directed cycle of D has at least two crossing consecutive chords if and only if D' has no odd directed cycle. Using Lemma 1.1.5 (which is actually an equivalence), it is sufficient to test if each strongly connected component is bipartite. This can be done by greedily trying to color with two colors. □

3.3.2 Computing a kernel

The question of whether a kernel can be computed in polynomial time when the directed graph satisfies the condition of Theorem 1.1.7 (or the special case of the crossing consecutive chords) is still open. In this thesis, this has been identified as one of the main open questions about kernels (see Section 1.3). It is not clear how to derive an efficient algorithm from the proof of that theorem (neither the original proof, nor the shorter one we propose). Yet, one of the crucial steps of the proof—namely, building a semi-kernel from a kernel in some subgraph—can be performed in polynomial time. An odd directed cycle not having two chords with consecutive heads will be called a *bad cycle*. The following algorithm takes a digraph D, a vertex u and a kernel in $D - u$ and returns a subset S of $V(D)$. As stated in Proposition 3.3.3, if D has no bad cycle it returns a semi-kernel in polynomial time.

Algorithm 1 Algorithm for computing a semi-kernel when we have almost a kernel

---

Require: a digraph D, a vertex u, a kernel K of $D - u$;
Ensure: S is a non-empty semi-kernel of D or there is a bad cycle;

1: $K' \leftarrow K \cup \{u\}$;
2: **if** $N^+(u) \cap K \neq \varnothing$ **then**
3:     $S \leftarrow K$;
4:     **return** S;
5: **if** $N^-(u) \cap K = \varnothing$ **then**
6:     $S \leftarrow K'$;
7:     **return** S;
8: $U \leftarrow \{u\}$;
9: **while** $N^+(U \cap K') \setminus N^-(U \cap K') \neq \varnothing$ **do**
10:     Pick $v \in N^+(U \cap K') \setminus N^-(U \cap K')$;
11:     $U \leftarrow U \cup \{v\} \cup (N^+(v) \cap K')$;
12: $S \leftarrow U \cap K'$;
13: **if** $N^-(u) \cap U \cap K \neq \varnothing$ **then**
14:     **return** "There is a bad cycle";
15: **else**
16:     **return** S;

Denote by v_i the vertex picked at line 10 at iteration i of Algorithm 1.

Lemma 3.3.2. *If $N^-(u) \cap K \neq \varnothing$ and $N^+(u) \cap K = \varnothing$, then, at the end of Algorithm 1, there exist for every vertex w in $U \cap K$ a sequence $i_1 < i_2 < \cdots < i_\ell$ and vertices $w_1, w_2, \ldots, w_\ell = w$ in $U \cap K$ such that $u, v_{i_1}, w_1, v_{i_2}, w_2, \ldots, w_{\ell-1}, v_{i_\ell}, w_\ell$ is a directed path of D.*

Proof. Assume $N^-(u) \cap K \neq \varnothing$ and $N^+(u) \cap K = \varnothing$. If the "while" loop is not executed at all, it means that $S = U \cap K' = \{u\}$ at the end of the algorithm, which implies that

$U \cap K = \varnothing$. We can thus assume that the "while" loop is executed at least once. We prove by induction that the property is true at every iteration. Since the set U is non-decreasing along the execution of the algorithm, the property will also be satisfied at the end of the algorithm.

Consider iteration 1 of the "while" loop, and let w be a vertex of K added to U. It cannot be the vertex v_1, which does not belong to K. It is thus a vertex of $N^+(v_1)$, and there is a directed path u, v_1, w in D, as required.

Now, consider an iteration i, and suppose that the property is satisfied up to iteration $i-1$. Let w be a vertex of K added to U. It cannot be v_i, which does not belong to K. It is thus a vertex of $N^+(v_i)$. The vertex v_i is by construction in the outneighborhood of a vertex w' of $U \cap K'$, which has been added to U' at an iteration $i' < i$. If $w' = u$, then $u, v_1, w_1 = w$ is the desired directed path. Suppose thus $w' \neq u$. The vertex w' belongs to $U \cap K$. By induction, there exist a sequence $i_1 < i_2 < \cdots < i_{\ell'}$ and vertices $w_1, w_2, \ldots, w_{\ell'} = w'$ in $U \cap K$ such that $u, v_{i_1}, w_1, v_{i_2}, w_2, \ldots, w_{\ell'-1}, v_{i_{\ell'}}, w_{\ell'}$ is a directed path of D. The sequence $u, v_{i_1}, w_1, v_{i_2}, w_2, \ldots, w_{\ell'-1}, v_{i_{\ell'}}, w_{\ell'}, v_i, w$ is the desired directed path of D. □

Proposition 3.3.3. *Consider a digraph D satisfying the condition of Theorem 1.1.7, a vertex $u \in V(D)$, and a kernel K of $D - u$. Algorithm 1 computes a semi-kernel of D in polynomial time.*

Proof. Note the algorithm terminates after at most $|V(D)|$ iterations because we cannot have twice the same vertex v at line 10. It remains to prove that the subset S returned is always a semi-kernel of D. If $N^+(u) \cap K \neq \varnothing$, then K is a kernel of D (and thus a semi-kernel) If $N^+(u) \cap K = \varnothing$ and $N^-(u) \cap K = \varnothing$, then $K' = K \cup \{u\}$ is a kernel of D (and thus a semi-kernel). So we assume from now on that $N^+(u) \cap K \neq \varnothing$ and $N^-(u) \cap K \neq \varnothing$. Denote by v_i the vertex v added by the algorithm at iteration i at line 10. Note that v_i is never a vertex of K since this latter set is an independent set and $N^+(u) \cap K = \varnothing$. We consider now two cases in turn.

- *$N^-(u) \cap U \cap K \neq \varnothing$ at the end of the algorithm.* Pick w in $N^-(u) \cap U \cap K$. According to Lemma 3.3.2, there exist a sequence $i_1 < i_2 < \cdots < i_\ell$ and vertices $w_1, w_2, \ldots, w_\ell$ in K such that $u, v_{i_1}, w_1, v_{i_2}, w_2, \ldots, w_{\ell-1}, v_{i_\ell}, w_\ell, u$ is a directed cycle (with $w_\ell = w$). Choose ℓ as small as possible. The length of the cycle is odd. Consider any chord of the cycle. It does not connect two vertices w_k since K is an independent set of D. It is not of the form $(v_{i_k}, w_{k'})$ with $k < k'$ by minimality of ℓ. For $k > k'$, the vertex $w_{k'}$ is added to U at iteration $i_{k'}$ or before and since the vertex v_{i_k} is not in $N^-(U \cap K)$ at iteration i_k, there is no chord of the form $(v_{i_k}, w_{k'})$ with $k > k'$ either. Therefore, the two heads of the crossing chords are u and v_{i_1}. By minimality of the of ℓ, the tail of the chord of u can not be an element of K and by definition, it can not be one of the $v_i's$. This proves that the cycle is bad.

- *$N^-(u) \cap U \cap K = \varnothing$ at the end of the algorithm.* In this case, S is set to $U \cap K'$, and no vertex of $U \cap K$ has u as outneighbor. Thus S is an independent set. Moreover, since the "while" loop has been left, we have $N^+(S) \subseteq N^-(S)$. □

Note that this proof provides a slightly different technique for finding the semi-kernel than the one used in the proof of Theorem 1.1.7. It is not clear how to make this latter polynomial.

Chapter 4

Blue/red arcs

One of the most celebrated generalizations of the Gale–Shapley theorem is the following theorem.

Theorem 1.1.15 (Sands, Sauer, and Woodrow [71])**.** *Let D be a directed graph whose arcs are colored with two colors. Then there is an independent set S of vertices of D such that, for every vertex x not in S, there is a monochromatic path from x to a vertex of S.*

An equivalent formulation of this theorem is the following: *Let D be a directed graph whose arcs are colored with two colors such that the restriction to each color forms a transitive digraph; then D admits a kernel.*

In this chapter we discuss and propose extensions and variations of this theorem.

For two vertices a, b, we denote by $a \xrightarrow{b} b$ (resp. $a \xrightarrow{r} b$) the existence of a blue (resp. red) arc from a to b.

4.1 Main results

Our first result consists in replacing the condition of Theorem 1.1.15 by the one depicted in Figure 4.1. The original condition is obtained by restricting each implication to its first alternative. It is also a generalization of Theorem 1.1.10 by Champetier ensuring that an M-clique-acyclic orientation of a comparability graph always admits a kernel, and of a theorem by Abbas and Saoula [1] that states the polynomiality of computing such a kernel under the same condition. Indeed, considering another orientation certifying that the graph is a comparability graph, we can color each arc in blue if the two orientations match and in red otherwise (as done by Champetier [22]), and it is easy to check that the condition of our result is satisfied.

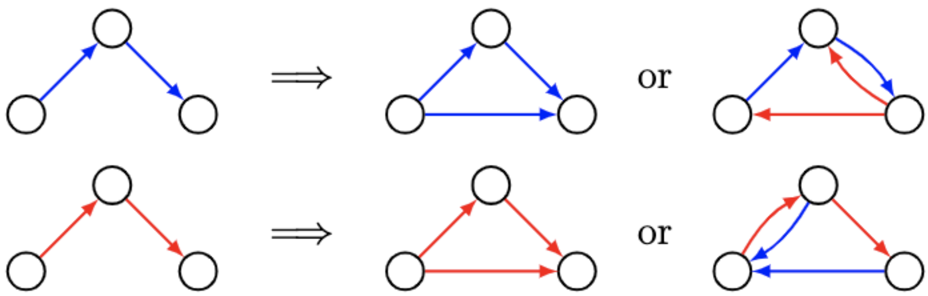

Figure 4.1: The condition in Theorem 4.1.1

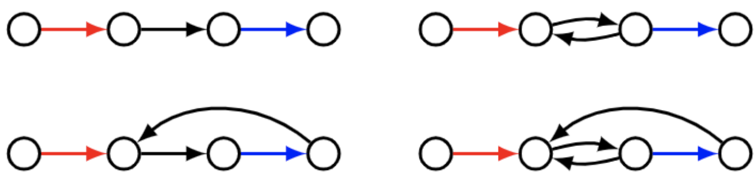

Figure 4.2: The forbidden induced structures of the second condition in Theorem 4.1.2

Theorem 4.1.1. *Let D be a digraph whose arcs are colored in blue and red, such that the following conditions are both satisfied (see Figure 4.1):*

(i) *If (u, v) and (v, w) are blue arcs, then (u, w) is a blue arc, or (w, u) and (w, v) are red arcs.*

(ii) *If (u, v) and (v, w) are red arcs, then (u, w) is a red arc, or (v, u) and (w, u) are blue arcs.*

Then D has a kernel, and it is possible to compute such a kernel in polynomial time.

Our second result is in the same vein as the previous one.

Theorem 4.1.2. *Let D be a digraph whose arcs are colored in blue and red, such that the following conditions are both satisfied (see Figure 4.2 for the forbidden structures):*

(i) *There is no monochromatic directed cycle.*

(ii) *If $((v_1, v_2), (v_2, v_3), (v_3, v_4))$ is a (closed or open) directed path such that (v_1, v_2) is red and (v_3, v_4) is blue, then its vertices induce at least another arc not ending at v_2.*

Then D has a kernel.

There is no implication between the two statements. Indeed, a digraph of size three such that every vertex has a blue arc to every other vertex satisfies conditions of Theorem 4.1.1 but has a monochromatic directed cycle. Also, a directed blue path of length two respects the conditions of Theorem 4.1.2 but not those of Theorem 4.1.1. Note that the proof of Theorem 4.1.2, even though it is very similar to the proof of Theorem 4.1.1, does not provide any efficient algorithm. The problem of finding a kernel in polynomial time in digraphs respecting conditions of Theorem 4.1.2 is still open.

4.2 Proof of Theorem 4.1.1

4.2.1 The poset of antichains

The following way of extending a partial order on antichains will be useful in the proof of Theorem 4.1.1. Let $(\mathcal{P}, \preccurlyeq)$ be a poset. Let $\mathcal{A}$ be the collection of all antichains of this poset. We extend $\preccurlyeq$ on $\mathcal{A}$ by setting $\alpha \preccurlyeq \alpha'$ for two antichains $\alpha, \alpha' \in \mathcal{A}$ whenever for each element x in α, there exists an element x' in α' such that $x \preccurlyeq x'$. We believe this construction and the following lemma are well-known but we have not been able to find any reference in the literature.

Lemma 4.2.1. *With this extended definition, $\preccurlyeq$ is a partial order on $\mathcal{A}$.*

Proof. Reflexivity and transitivity are immediate. We establish antisymmetry. Suppose for a contradiction that there exist two antichains $\alpha \neq \alpha'$ such that $\alpha \preccurlyeq \alpha'$ and $\alpha' \preccurlyeq \alpha$. Since $\alpha \neq \alpha'$, we can assume without loss of generality that there exists $x \in \alpha$ such that $x \notin \alpha'$. Since $\alpha \preccurlyeq \alpha'$, there exists $x' \in \alpha'$ such that $x \prec x'$. Since $\alpha' \preccurlyeq \alpha$, there exists $x'' \in \alpha$ such that $x' \preccurlyeq x''$, and thus such that $x \prec x''$; a contradiction. □

The next lemma will be useful to establish the polynomiality of the computation of the kernel. Actually, only the upper bound on the size of a chain is needed for the proof but we establish the existence of a chain matching the upper bound for sake of completeness.

Lemma 4.2.2. *If $\mathcal{P}$ is finite, then the maximal size of a chain in $(\mathcal{A}, \preccurlyeq)$ is $|\mathcal{P}| + 1$.*

Proof. Let $\alpha_1 \prec \alpha_2 \prec \cdots \prec \alpha_\ell$ be a chain of $(\mathcal{A}, \preccurlyeq)$. For every $i \in \{2, \ldots, \ell\}$, there exists $x_i \in \mathcal{P}$ such that $x_i \in \alpha_i \setminus \alpha_{i-1}$. We claim that all x_i's are distinct. Suppose for a contradiction that there exist $j < k$ such that $x_j = x_k$. Let $y \in \alpha_{k-1}$ be such that $x_j \preccurlyeq y$. Note that since x_k does not belong to α_{k-1}, we have actually $x_j \neq y$. Moreover, there exists $z \in \alpha_k$ such that $y \preccurlyeq z$. We have thus $x_k \prec z$, with both elements belonging to α_k; a contradiction. This shows that every chain in $(\mathcal{A}, \preccurlyeq)$ is of size at most $|\mathcal{P}| + 1$.

We prove now by induction on $|\mathcal{P}|$ that there exists a chain in $(\mathcal{A}, \preccurlyeq)$ of size $|\mathcal{P}| + 1$. This is obviously true if $\mathcal{P} = \varnothing$. Assume now that $\mathcal{P} \neq \varnothing$. Let x be a maximal element of $\mathcal{P}$. By induction, $\mathcal{A}$ possesses a chain of size $|\mathcal{P}|$ such that none of its elements contains x. Let α be the maximal antichain of $(\mathcal{P}, \preccurlyeq)$ in this chain of $(\mathcal{A}, \preccurlyeq)$. Add x to α, and remove from it all elements y such that $y \preccurlyeq x$. This leads to a new antichain α' such that $\alpha \prec \alpha'$, showing that there exists a chain in $(\mathcal{A}, \preccurlyeq)$ of size $|\mathcal{P}| + 1$. □

4.2.2 Two lemmas

Consider a digraph D as in Theorem 4.1.1. Such a digraph satisfies some properties, which we state as lemmas since they will be useful in the proof of Theorem 4.1.1.

Lemma 4.2.3. *There is a vertex v such that the implication $v \xrightarrow{r} w \implies w \to v$ holds for all vertices w.*

Proof. Consider the set R of red arcs (v, w) such that the arc (w, v) does not exist (i.e., the set of red arcs (v, w) that witness the fact the v does not satisfy the implication). We claim that the restriction of D to arcs in R is acyclic. Suppose for a contradiction that there is a cycle, and consider such a cycle of minimal length. It cannot be of length 2 because this would contradict the definition of R. It is thus of length at least 3. By (ii) and by minimality of the length, there is an arc of the cycle whose reverse arc exists in blue in the digraph. Such an arc contradicts the definition of R, this proves the acyclicity. The restriction to arcs in R has thus a sink. Such a sink necessarily satisfies the implication. □

Lemma 4.2.4. *If there is a blue dipath from a vertex u to a vertex v, then $u \xrightarrow{b} v$ or $v \xrightarrow{r} u$.*

Proof. Consider a minimum-length blue dipath P from a vertex u to a vertex v. If it is of length 1, then $u \xrightarrow{b} v$. So suppose that the length of P is at least 2. By (i) and by minimality of the length of P, there exists a red dipath from v to u that uses vertices of P in the opposite order as that induced by P. Note that if P is of odd length, the existence of such a path requires to use a red arc that is the reverse of an arc of P. Let Q be such a red dipath of minimal length. By (ii) and by minimality of the length of both P and Q, the dipath Q is of length 1 and we have $v \xrightarrow{r} u$. □

4.2.3 The proof

Proof of Theorem 4.1.1. Let D^b be the digraph obtained from D by keeping only the blue arcs and denote by $\mathcal{K}^b$ the collection of its strongly connected components. Let $K \preccurlyeq K'$ holds for $K, K' \in \mathcal{K}^b$ if there is a blue dipath from K to K'. This makes $(\mathcal{K}^b, \preccurlyeq)$ a poset. We extend the definition of $\preccurlyeq$ on the antichains of this poset as in Section 4.2.1. By Lemma 4.2.1, $\preccurlyeq$ is a partial order on these antichains. By Lemma 4.2.4, the strongly connected components of D^b intersected by an independent set I of D form an antichain of $(\mathcal{K}^b, \preccurlyeq)$, which we denote by α_I.

Let $\mathcal{I}$ be the set of independent sets I of D such that $I \xrightarrow{r} w \implies w \to I$ holds for all vertices w. By Lemma 4.2.3, $\mathcal{I}$ is non-empty (and an element from $\mathcal{I}$ can be determined by simply scanning the vertices of D). We describe now a procedure to modify such a $I \in \mathcal{I}$ when it is not a kernel, in order to get a new element I' in $\mathcal{I}$ such that $\alpha_I \prec \alpha_{I'}$. By finiteness, this will show the existence of a kernel. With Lemma 4.2.2, this will even imply the polynomiality of the method.

Suppose that $I \in \mathcal{I}$ is not a kernel and let $U := V(D) \setminus (I \cup N^-(I))$. This is the set of all vertices that are neither in I nor absorbed by I. Since I is not a kernel, U is non-empty. According to Lemma 4.2.3 applied on $D[U]$, there exists a vertex $v \in U$ such that $v \xrightarrow{r} w \implies w \to v$ holds for all vertices $w \in U$. (Again, determining such a v can be done simply by scanning the vertices of D.)

If $I \not\to v$, then adding v to I leads to $I \cup \{v\} \in \mathcal{I}$ such that $\alpha_I \prec \alpha_{I\cup\{v\}}$. Suppose now that $I \to v$. Define I' by removing from I all vertices in $N^-(v)$ and by adding v. The set I' is independent since $v \not\to I$ by definition of v and all vertices of I with an arc to v have been removed from I. Moreover, let w be any vertex such that $I' \xrightarrow{r} w$. We claim that $w \to I'$. Three cases have to be considered.

- $w \to I \setminus N^-(v)$. Then $w \to I'$ since $I \setminus N^-(v) \subseteq I'$ by definition.

- $w \to I \cap N^-(v)$. Let u be a vertex in $I \cap N^-(v)$ such that $w \to u$. We have $u \xrightarrow{b} v$ because $v \not\to I$. If $w \xrightarrow{b} u$, then (i) implies $w \xrightarrow{b} v$ because $v \not\to I$ and we are done. So, suppose that $w \xrightarrow{r} u$, and let $v' \in I'$ such that $v' \xrightarrow{r} w$. By (ii), we have $w \xrightarrow{b} v'$ (here, we use the fact that $v' \stackrel{r}{\not\to} u$, which holds either because $v' \in I$ or because $v' = v$ and $v \not\to I$). Therefore, in this case, whatever is the color of the arc (w, u), we have $w \to I'$.

- $w \not\to I$. This means that $v \xrightarrow{r} w$. Since w does not belong to I', it is distinct from v. Since $v \not\to I$, the vertex w does not belong to $I \cap N^-(v)$. Therefore, the vertex w belongs to U and by definition of v we have $w \to v$, which implies $w \to I'$.

Therefore, $I' \in \mathcal{I}$. Moreover, no vertex of $I \cap N^-(v)$ is in the same strongly connected component of D^b as v: indeed, the existence of a blue dipath from v to $I \cap N^-(v)$ would imply by Lemma 4.2.4 $v \xrightarrow{b} I$ or $I \xrightarrow{r} v$, both situations contradicting the fact that v belongs to U. Therefore, $\alpha_I \prec \alpha_{I'}$. □

4.3 Proof of Theorem 4.1.2

The proof is very similar to the proof of Theorem 4.1.1 but does not require any result about the poset of antichains.

Proof of Theorem 4.1.2. Let $\mathcal{I}$ be the set of independent sets $I \subseteq V(D)$ such that $I \xrightarrow{r} w \implies w \to I$ holds for all vertices w. We define $\preccurlyeq$ on $\mathcal{I}$ such that $I \preccurlyeq I'$ whenever for each element x in I, there exists a blue dipath from x to I'. Condition (i) ensures that $\preccurlyeq$ defines an order on $\mathcal{I}$.

$\mathcal{I}$ is not empty (just take a sink of the digraph obtained from D by keeping only the red arcs). We describe now a procedure to modify such a $I \in \mathcal{I}$ when it is not a kernel, in order to get a new element $I' \in \mathcal{I}$ such that $I \prec I'$. By finitness, this will show the existence of a kernel.

Suppose that $I \in \mathcal{I}$ is not a kernel and let $U \coloneqq V(D) \setminus (I \cup N^-(I))$. This is the set of all vertices that are neither in I nor absorbed by I. Since I is not a kernel, U is non-empty. Take v a sink of the digraph obtained from $D[U]$ by keeping only the red arcs. If $I \not\to v$, then adding v to I leads to a $I \cup \{v\} \in \mathcal{I}$ such that $I \prec I \cup \{v\}$. Suppose now that $I \to v$. Define I' by removing from I all vertices in $N^-(v)$ and by adding v. The set I' is independent since $v \not\to I$ by definition of v and all vertices of I with an arc to v have been removed from I. Moreover, let $w \in V(D)$ and $v' \in I'$ be any vertex such that $v' \xrightarrow{r} w$ (with possibly $v' = v$). We claim that $w \to I'$. Three cases have to be considered.

- $w \to I \setminus N^-(v)$. Then $w \to I'$ since $I \setminus N^-(v) \subseteq I'$ by definition.

- $w \to I \cap N^-(v)$. Let u be a vertex in $I \cap N^-(v)$ such that $w \to u$. We have $u \xrightarrow{b} v$ because $v \not\to I$. Condition (ii) implies $w \to v$ or $w \to v'$, so in every case $w \to I'$.

- $w \not\to I$. This means that $v \xrightarrow{r} w$. Since w does not belong to I', it is distinct from v. Since $v \not\to I$, the vertex w does not belong to $I \cap N^-(v)$. Therefore, the vertex w belongs to U but this is a contradiction with the definition of v.

Therefore, $I' \in \mathcal{I}$. Moreover, every vertex u in $I \setminus I'$ is in $N^-(v)$ and since $v \not\to I$, $u \xrightarrow{b} v$, therefore $I \prec I'$. □

We could not find any polynomial bound on the number I considered in the previous proof. As a result, we do not known whether any complexity result for finding a kernel in graphs satisfying conditions of Theorem 4.1.2 can be derived from the proof.

Part II

Quasi-kernels

Chapter 5

Quasi-kernels in a nutshell

This chapter is intended to be a brief but complete introduction to quasi-kernels. The first section presents two distinct proofs of the existence of a quasi-kernel in any digraph (Section 5.1). After providing the formal statement of the small quasi-kernel conjecture, we review then the current knowledge about this conjecture (Section 5.2). Then, we collect a series of results about the simultaneous existence of more than one quasi-kernel, a topic that has attracted some attention since the nineties (Section 5.3).

5.1 Existence of a quasi-kernel

We establish here in two different ways the existence of a quasi-kernel in any digraph D. These are classical proofs.

Original proof by Chvátal and Lovász [26]. By induction on the vertices of D. Consider a vertex $v \in V(D)$ and a quasi-kernel Q of $D - N^-[v]$. If $Q \cap N^+(v) \neq \varnothing$, then Q forms a quasi-kernel of D. Otherwise, $Q \cup \{v\}$ is a quasi-kernel. □

The next proof relies in an interesting way to the von Neumann–Morgenstern theorem.

Thomassé's proof [17]. Partition D into two acyclic digraphs (this is standard and can be done by locating the vertices on a line). The first one admits a kernel by Theorem 1.1.2. This kernel induces in the second digraph another acyclic digraph, which admits also a kernel, again by Theorem 1.1.2. This latter kernel is a quasi-kernel of D. □

Note that both proofs ensure that a quasi-kernel can be computed in polynomial time.

5.2 The small quasi-kernel conjecture

The main conjecture about quasi-kernels has been proposed by Erdős and Székely [41] in 1976, and can be stated as follows.

Conjecture ("Small quasi-kernel conjecture"). *Every sink-free digraph D admits a quasi-kernel of size at most $\frac{1}{2}|V(D)|$.*

So far this conjecture is still open but has been shown for a few special cases.

A *semicomplete multipartite* digraph is an orientation of a complete multipartite graph. A *quasi-transitive* digraph D is such that for every three vertices $a, b, c \in A(D)$ such that $(a, b) \in A(D)$ and $(b, c) \in A(D)$ then $(a, c) \in A(D)$ or $(c, a) \in A(D)$. A *locally semicomplete* digraph is an orientation of a graph such that for every vertex v, the underlying graph of the restriction to $N^-[v] \cup N^+[v]$ forms a complete graph.

Theorem 5.2.1 (Heard and Huang [49] 2008)**.** *The small quasi-kernel conjecture is satisfied by semicomplete multipartite digraphs, by quasi-transitive digraphs, and by locally semicomplete digraphs.*

Kostochka, Luo, and Shan have recently proved that 4-colorable digraphs—and in particular planar digraphs—satisfy the small quasi-kernel conjecture. Actually, they have established the following more general result.

Theorem 5.2.2 (Kostochka, Luo and Shan [59] 2020)**.** *The small quasi-kernel conjecture is satisfied by digraphs whose vertex set can be partitioned into two subsets, each inducing a kernel-perfect digraph.*

This theorem relates the existence of a small quasi-kernel to the existence of kernels. Such relation is even more striking in the next theorem, whose proof is actually very simple (see [3, 39] for a streamlined version of the original proof).

Theorem 5.2.3 (van Hulst [73] 2021)**.** *The small quasi-kernel conjecture is satisfied by digraphs admitting a kernel.*

We finish with very recent results, showing that there is still activity around the small quasi-kernel conjecture.

Theorem 5.2.4 (Erdős, Győri, Mezei, Salia, and Tyomkyn [39] 2023)**.** *The small quasi-kernel conjecture is satisfied by digraphs containing a kernel in the second outneighborhood of a quasi-kernel and also by orientations of unicyclic graphs.*

A *directed anti-claw* is a digraph isomorphic to D with $V(D) = \{v_1, v_2, v_3, v_4\}$ and $A(D) = \{(v_1, v_4), (v_2, v_4), (v_3, v_4)\}$. A digraph is *anti-claw-free* if it contains no induced directed anti-claw.

Theorem 5.2.5 (Ai, Gerke, Gutin, Yeo, Zhou [3] 2023)**.** *The small quasi-kernel conjecture is satisfied by anti-claw-free digraphs.*

In Chapter 6, we contribute to the conjecture by proving the conjecture for particular cases. Actually, it is not known whether every sink-free digraph has a quasi-kernel of size at most $\alpha|V|$ for any $\alpha \geqslant \frac{1}{2}$, so we start by proving that every sink-free split digraph has a quasi-kernel of size at most $\frac{3}{4}|V(D)|$. We also prove the conjecture for subclasses of split digraphs such as splits avoiding a particular structure and complete split digraphs. Then, we establish the conjecture on complete bipartite graphs before considering a reformulation of the small quasi-kernel conjecture. In the chapter, we also discuss equivalent conjectures.

5.3 Simultaneous quasi-kernels

Jacob and Meyniel [54] proved that if a digraph does not have a kernel then it must contain at least three (not necessarily disjoint) quasi-kernels. Digraphs with exactly one and two quasi-kernels have been characterized by Gutin et al. [48]. It follows from this characterization that if a digraph has precisely two quasi-kernels then these two quasi-kernels are actually disjoint.

In 2001 Gutin et al. [47] conjectured that every sink-free digraph has two disjoint quasi-kernels (this stronger conjecture implies the original small quasi-kernel conjecture). In 2004, in an update of their paper, the authors constructed a counterexample, a split digraph with 14 vertices [48].

Note that, whereas the small quasi-kernel conjecture is true for planar sink-free digraphs [59], no sink-free planar digraph without two disjoint quasi-kernels is known so far (the counterexample constructed by Gutin et al. [48] does contain a directed K_7).

5.4 Algorithmic considerations

Surprisingly enough, quasi-kernels have almost not been studied from an algorithmic point of view. The only result of this kind from the literature is the following: deciding whether there exists a quasi-kernel that contains a specified vertex is NP-complete [30].

In Chapter 7, we initiate the study of the problem of finding a quasi-kernel of minimum size problem which we call Quasi-Kernel. This problem is computationally hard, even for simple digraph classes, as shown for example by Theorem 7.2.1 in the case of acyclic orientations of cubic graphs, or by Theorem 7.2.2 in the case of acyclic orientations of bipartite graphs. Also, we investigate the problem of finding two disjoint quasi-kernels. As we shall prove in Section 7.1, not only sink-free digraphs occasionally fail to contain two disjoint quasi-kernels, but it is actually computationally hard to distinguish those that do from those that do not. In the restricted case of sink-free bounded degree planar digraph, deciding whether it has three disjoint quasi-kernels is NP-complete, as proved in Theorem 7.1.4.

For those two problems, we provide polynomial algorithms in the case of digraphs with bounded treewidth in Section 7.5.

Chapter 6

Structure

Sections 6.1 and 6.2 form a paper that is about to be submitted to a journal.

In this chapter, we study upper bounds for the size of the smallest quasi-kernel in different classes of digraphs with a particular focus on split digraphs.

We denote by $S(D)$ the sinks of a digraph D.

6.1 Split

A *split* graph is a graph whose vertices can be partitioned into a clique and an independent set. We extend this notion to digraphs by requiring that the underlying undirected graph is a split graph. For a split digraph D, we denote by $K(D)$ the set of vertices of the clique-part and by $I(D)$ the set of vertices of the independent-part. One of the motivations is the construction by Gutin et al. [48] using a split digraph for refuting their conjecture about the existence of two disjoint quasi-kernels in every sink-free digraph [47]. A split is *one-way* if if contains no arc from $K(D)$ to $I(D)$.

6.1.1 Optimality of the 1/2 ratio

We describe two infinite families of split digraphs for which the minimum size of the quasi-kernel tends to $\frac{1}{2}|V(D)|$. The first family is formed by one-way split digraphs. The second family does not contain one-way split digraphs and shows that, if the ratio 1/2 is correct for split digraphs, one-way split digraphs are not the only reason for tightness.

Consider the one-way split digraph D_n defined by

$$K(D_n) \coloneqq \{k_0, \ldots, k_{2n}\}, \quad I(D_n) \coloneqq \Big\{s_{ij} \colon 0 \leqslant i \leqslant 2n, 1 \leqslant j \leqslant n\Big\}, \quad \text{and}$$
$$A(D_n) \coloneqq \Big\{(k_i, k_{i+j}) \colon 0 \leqslant i \leqslant 2n, 1 \leqslant j \leqslant n\Big\} \cup \Big\{(s_{ij}, k_i) \colon 0 \leqslant i \leqslant 2n, 1 \leqslant j \leqslant n\Big\},$$

where the sum $i+j$ is understood modulo $2n+1$ (i.e., $i+j = i+j-2n-1$ if $i+j > 2n$).

Proposition 6.1.1. *Denote by Q_n a smallest quasi-kernel of D_n. Then*

$$\lim_{n\to+\infty} \frac{|Q_n|}{|V(D_n)|} = \frac{1}{2}.$$

Proof. Since D_n is one-way, Q_n intersects $K(D_n)$ in a single vertex. Assume that this vertex is k_0. This is without loss of generality because of the symmetry of D_n. For every $i \in \{1, \ldots, n\}$, the shortest k_i-k_0 path is of length two. On the other hand, there is an arc (k_i, k_0) for every $i \in \{n+1, \ldots, 2n\}$. Therefore $Q_n = \{k_0\} \cup \Big\{s_{ij} \colon i, j \in \{1, \ldots, n\}\Big\}$ and $|Q_n| = n^2+1$. The convergence result is then a consequence of $|V(D_n)| = 2n^2+3n+1$. □

Consider now the split digraph D'_n defined by

$$K(D'_n) \coloneqq K(D_n), \quad I(D'_n) \coloneqq I(D_n), \quad \text{and}$$
$$A(D'_n) \coloneqq A(D_n) \cup \{(k_0, s_{ij})\colon 1 \leqslant i \leqslant 2n, 1 \leqslant j \leqslant n\}.$$

It is not a one-way split digraph but a strongly connected split digraph.

Proposition 6.1.2. *Denote by Q'_n a smallest quasi-kernel of D'_n. Then*

$$\lim_{n\to+\infty} \frac{|Q'_n|}{|V(D'_n)|} = \frac{1}{2}\,.$$

Proof. Since every path from k_1 to $I(D'_n)$ is of length at least 3, the quasi-kernel Q'_n intersects necessarily $K(D'_n)$ in a single vertex. This vertex cannot be k_0 since the shortest path from s_{11} to k_0 is of length 3 and of length 4 to $I(D'_n)$. Then, $Q'_n \cap K(D'_n) = \{k_\ell\}$ with $\ell \in \{1, \ldots, 2n\}$, and $Q'_n = \{k_\ell\} \cup \{s_{(\ell+h)j}\colon 1 \leqslant h \leqslant n, 1 \leqslant j \leqslant n\} \setminus \{s_{0j}\colon 1 \leqslant j \leqslant n\}$. Therefore, $n^2 + 1 \geqslant |Q'_n| \geqslant n(n-1) + 1$ and the convergence result is a consequence of $|V(D_n)| = 2n^2 + 3n + 1$. □

6.1.2 A 3/4-bound for sink-free split digraphs

The fact that one-way split digraphs satisfy the small quasi-kernel conjecture has been noticed by several persons. We provide a short proof for sake of completeness. See [3] for a stronger version. The proof uses an easy preliminary lemma.

Lemma 6.1.3. *For every tournament T with positive weights w_v on the vertices, the following holds:*

$$\max_{v\in V(T)} w(N^-[v]) \geqslant \frac{1}{2} w(V(T))\,.$$

Proof. We have $\sum_{v\in V(T)} w_v w\Big(N^-(v)\Big) = \sum_{v\in K(D)} w_v w\Big(N^+(v)\Big)$, just because both quantities are equal to the sum $\sum_{(u,v)} w_u w_v$ taken over the arcs of T. Thus there exists a vertex $\bar{v}$ such that $w(N^-(\bar{v})) \geqslant w(N^+(\bar{v}))$, which implies $w(N^-[\bar{v}]) \geqslant \frac{1}{2} w(V(T))$. □

Proposition 6.1.4. *Every sink-free one-way split digraph has a quasi-kernel of size at most $\frac{1}{2}|V(D)|$.*

Proof. Set $w_v \coloneqq |N^-(v) \cap I(D)| + 1$ for every vertex $v \in K(D)$. Let v^* be a vertex such that $w(N^-[v^*] \cap K(D))$ is maximal. We first check that v^* is a quasi-kernel of $D[K(D)]$. Let $v \in K(D)$. By definition of v^* and since $w_v > 0$, we have $w(N^-[v^*] \cap K(D)) > w(N^-(v) \cap K(D))$, which means that $N^-[v^*] \cap N^+[v]$ is nonempty.

Define now $Q \coloneqq (\{v^*\} \cup I(D)) \setminus \Big(N^-(v^*) \cup N^{--}(v^*)\Big)$. The set Q is independent. Every vertex of $K(D)$ is at distance at most two from v^* because v^* is a quasi-kernel of $D[K(D)]$. Every other vertex is either in Q or at distance at most two of v^*. By construction, every vertex of $Q \cap I(D)$ is in the inneighborhood of $K(D) \setminus N^-[v^*]$. If $K(D) \setminus N^-[v^*]$ is empty, then the size of Q is exactly one; otherwise, it is at most $w(K(D) \setminus N^-[v^*]) = w(K(D)) - w(N^-[v^*] \cap K(D)) \leqslant \frac{1}{2} w(K(D))$, by Lemma 6.1.3. Therefore, in any case the size of Q is at most $\frac{1}{2}(|I(D)| + |K(D)|) = \frac{1}{2}|V(D)|$. □

Lemma 6.1.5. *Let X and Y be two disjoint subsets of a set V, and let Y' be a subset of Y such that $|Y'| \leqslant \frac{1}{2}|Y|$. If $|X| \leqslant \frac{1}{2}|V|$, then $|X \cup Y'| \leqslant \frac{3}{4}|V|$.*

Proof. We have $|X| + |Y| \leqslant |V|$ since X and Y are disjoint. Thus $|X| + 2|Y'| \leqslant |V|$. Adding $|X|$ on the left-hand side and $\frac{1}{2}|V|$ on the right-hand side, and then dividing by 2 leads to the desired inequality. □

Using the previous results, we prove the following theorem, thereby proving a weaker version of the small quasi-kernel conjecture for split digraphs.

Theorem 6.1.6. *Every sink-free split digraph D admits a quasi-kernel of size at most $\frac{3}{4}|V(D)|$.*

Proof. We assume that $K(D)$ and $I(D)$ are both non-empty since otherwise there is nothing to prove. Let V_i be the set of vertices at distance exactly i of $I(D)$. Notice that V_0 is equal to $I(D)$. Consistently, $V_{+\infty}$ is the set of vertices v of D for which there is no directed path from v to an element of $I(D)$. We denote by $i^{\max}$ the largest index $i \in \mathbb{Z}_+ \cup \{+\infty\}$ such that V_i is non-empty. Since $K(D)$ is non-empty, we have $i^{\max} \geqslant 1$. Note that the sets V_i that are non-empty form a partition of $V(D)$, and that if V_i is non-empty (for finite $i \geqslant 1$), then V_{i-1} is non-empty as well. The proof considers in turns the possible values of $i^{\max}$.

- *Case $i^{\max} = +\infty$.*
  The digraph $D[V_{+\infty} \cup (N^-(V_{+\infty}) \cap I(D))]$ is a one-way split digraph, which admits a small quasi-kernel Q according to Proposition 6.1.4. By definition, no vertex of $V_{+\infty}$ is the origin of an arc ending in V_i for some $i < +\infty$. Every vertex in $K(D) \setminus V_{+\infty}$ is thus at distance one of every vertex in $V_{+\infty}$, which implies that Q is a small quasi-kernel of D as well.

- *Case $i^{\max} \in \{1, 2\}$.*
  For every vertex v in V_1, pick an arbitrary vertex in $N^+(v) \cap I(D)$. This provides a subset $I^+ \subseteq I(D)$ such that $N^+(v) \cap I^+ \neq \varnothing$ for every $v \in V_1$ and $|I^+| \leqslant |V_1|$. Notice that when $i^{\max} = 1$, the set I^+ is a small quasi-kernel of D, and therefore we assume that $i^{\max} = 2$. Let $U \coloneqq \big(N^-(V_2) \cap I(D)\big) \setminus I^+$.

  – *Subcase $|U| \leqslant \frac{1}{2}|V(D)|$.*
  Let $X \coloneqq U$, $Y \coloneqq V_1 \cup I^+$, and $Y' \coloneqq I^+$.
  The two sets X and Y are disjoint. Since I^+ is disjoint from V_1, and smaller, we have $|Y'| \leqslant \frac{1}{2}|Y|$. Lemma 6.1.5 with $V \coloneqq V(D)$ shows then that $|X \cup Y'| \leqslant \frac{3}{4}|V(D)|$.
  Since X and Y' are both included in $I(D)$, the set $X \cup Y'$ is independent. Every vertex in V_1 is at distance exactly one of Y' and every vertex not in $X \cup Y'$ is at distance at most one of V_1. Therefore, $X \cup Y'$ is a quasi-kernel of D of size at most $\frac{3}{4}|V(D)|$.

  – *Subcase $|U| > \frac{1}{2}|V(D)|$.*
  Let $X \coloneqq I(D) \setminus (N^-(V_2) \cup I^+)$ and $Y \coloneqq V(D) \setminus X$. The digraph $D[V_2 \cup U]$ is a one-way split digraph and admits a small quasi-kernel Q' according to Proposition 6.1.4. Let $Y' \coloneqq Q' \cup \big(I^+ \setminus N^-(Q')\big)$. Note that Y' is disjoint from X and thus included in Y.
  We have $|I^+| \leqslant \frac{1}{2}|V_1 \cup I^+|$ and $|Q'| \leqslant \frac{1}{2}|V_2 \cup U|$. Since $V_1 \cup I^+$ and $V_2 \cup U$ are disjoint and $Y = (V_1 \cup I^+) \cup (V_2 \cup U)$, we have $|Y'| \leqslant \frac{1}{2}|Y|$. Lemma 6.1.5 with $V \coloneqq V(D)$ shows then that $|X \cup Y'| \leqslant \frac{3}{4}|V(D)|$.
  There is no arc from Q' to I^+: indeed, the only vertex of Q' in $K(D)$ (which exists because $D[V_2 \cup U]$ is a one-way split digraph) belongs to V_2 and can thus not be the origin of an arc ending in I^+. There is no arc from $I^+ \setminus N^-(Q')$ to Q' by definition. Hence, the set Y' is independent. The set X is independent by definition. The arcs leaving X all end in V_1. Since $Y' \cap K(D) \subseteq V_2$, there is no arc from X to Y' and no arc from Y' to X either. The set $X \cup Y'$ is therefore independent.
  Every vertex in I^+ is at distance at most one of Y' and every vertex in V_1 is at distance one of I^+. Thus, every vertex in V_1 is at distance at most two of Y'. Every vertex in $V_2 \cup U$ is at distance at most two of Y' as well because this latter set contains

Q'. A vertex of $I(D)$ that is neither in X nor in Y' is in $U \cup I^+$. Therefore, $X \cup Y'$ is a quasi-kernel of D of size at most $\frac{3}{4}|V(D)|$.

- *Case* $3 \leqslant i^{\max} < +\infty$.

Let $W := V_{i^{\max}-1} \cap N^+(V_{i^{\max}})$ and $P := I(D) \setminus N^-(K(D) \setminus W)$ (every arc leaving P ends in W).

– *Subcase* $|P| \leqslant \frac{1}{2}|V(D)|$.

Let $X := P$, $Y := V_{i^{\max}} \cup (N^-(V_{i^{\max}}) \cap I(D))$, and Y' be a small quasi-kernel of $D[Y]$ (which exists by Proposition 6.1.4 because the latter graph is a one-way split digraph).

The set X is included in $I(D)$ and no vertex in X is the origin of an arc ending in $V_{i^{\max}}$. Thus, X and Y are disjoint, and Lemma 6.1.5 with $V := V(D)$ shows that $|X \cup Y'| \leqslant \frac{3}{4}|V(D)|$.

Each of X and Y' is independent by definition. Moreover, there is no arc from X to Y' by definition, and no arc from Y' to X since $Y' \cap K(D) \subseteq V_{i^{\max}}$. Therefore, $X \cup Y'$ is independent.

Since $D[Y]$ is a one-way split digraph, Y' has a non-empty intersection with $V_{i^{\max}}$. Every vertex in V_i for $i \in \{1, 2, \ldots, i^{\max} - 2\}$ is the origin of an arc ending at the vertex of $Y' \cap V_{i^{\max}}$. This implies that every vertex in V_i for $i \in \{1, 2, \ldots, i^{\max} - 2\}$ is at distance at most one of Y' and that every vertex in $V_{i^{\max}-1} \cup V_{i^{\max}}$ is at distance at most two of Y'. Every vertex in $I(D) \setminus P$ is the origin of an arc ending in $K(D) \setminus W$. Since every vertex in this latter set is the origin of an arc ending at the vertex of $Q' \cap V_{i^{\max}}$, every vertex in $I(D) \setminus P$ is at distance at most two of Y'. Therefore, $X \cup Y'$ is a quasi-kernel of D of size at most $\frac{3}{4}|V(D)|$.

– *Subcase* $|P| > \frac{1}{2}|V(D)|$.

Let $X := I(D) \setminus N^-(V_{i^{\max}} \cup W)$, $Y := V_{i^{\max}} \cup W \cup (N^-(V_{i^{\max}} \cup W) \cap I(D))$, and Y' be a small quasi-kernel of $D[Y]$ (which exists by Proposition 6.1.4 because the latter graph is a one-way split digraph).

The set X is included in $I(D)$ and no vertex in X is the origin of an arc ending in $V_{i^{\max}} \cup W$. Thus, X and Y are disjoint. The sets X and P being disjoint, we have $|X| \leqslant \frac{1}{2}|V(D)|$, and Lemma 6.1.5 with $V := V(D)$ shows then that $|X \cup Y'| \leqslant \frac{3}{4}|V(D)|$.

Each of X and Y' is independent by definition. Moreover, there is no arc from X to Y' by definition, and no arc from Y' to X since $Y' \cap K(D) \subseteq V_{i^{\max}-1} \cup V_{i^{\max}}$. Therefore, $X \cup Y'$ is independent.

Since $D[Y]$ is a one-way split digraph, Y' has a non-empty intersection with $V_{i^{\max}} \cup W$. Every vertex in V_i is at distance one of every vertex in $V_{i^{\max}}$ for $i \in \{1, 2, \ldots, i^{\max} - 2\}$. Similarly, every vertex in $V_{i^{\max}-1} \setminus W$ is at distance one of every vertex in $V_{i^{\max}}$. Since every vertex in W is the end of an arc originating in $V_{i^{\max}}$, every vertex in $K(D)$ is at distance at most two of Y', whether Y' have a non-empty intersection with $V_{i^{\max}}$ or with W. The vertices in $I(D)$ are either in X or in Y. Therefore $X \cup Y'$ is a quasi-kernel of D of size at most $\frac{3}{4}|V(D)|$. □

6.1.3 Split digraphs with sinks

Kostochka, Luo, and Shan [59] proposed the following conjecture, which they proved to be equivalent to the small quasi-kernel conjecture. Recall that we denote by $S(D)$ the set of sinks of a digraph D.

Conjecture 1. *Every digraph D admits a quasi-kernel of size at most $\frac{1}{2}(|V(D)| + |S(D)| - |N^-(S(D))|)$.*

The next proposition shows that the equivalence between the small quasi-kernel conjecture and Conjecture 1 still holds when restricted to split digraphs: if we were able to prove the small quasi-kernel conjecture for all split digraphs, then Conjecture 1 would also be satisfied for all split digraphs. Moreover, this holds for weaker versions of the conjectures as well, where the 1/2-ratio is replaced by a larger ratio.

Proposition 6.1.7. *The small quasi-kernel conjecture and Conjecture 1 are equivalent for split digraphs, even with the ratio* $1/2$ *being replaced by a larger one.*

Proof. If Conjecture 1 is true for split digraphs, then the small quasi-kernel conjecture is clearly true for split digraphs as well. So, we prove just the reverse implication. We assume from now on that the small quasi-kernel conjecture is true for split digraphs, for a fixed ratio $\alpha \geqslant 1/2$. We prove by induction on the number of vertices that in every split digraph D, there exists a quasi-kernel of size at most $\alpha(|V(D)| + |S(D)| - |N^-(S(D))|)$.

If there is only one vertex, the existence of the desired quasi-kernel is obvious. Now, suppose there are at least two vertices. If the underlying undirected graph is not connected, then we apply induction on every component. If the graph has no sink, then we apply Conjecture 5.2 (version for split digraphs, and ratio α). If $S(D)$ is a quasi-kernel of D, then we are done because it is of size at most $\alpha\,(|V(D)| + |S(D)| - |N^-(S(D))|)$.

We are left with the case where the underlying undirected graph of D is onnected, $|V(D)| \geqslant 2$, $|S(D)| \geqslant 1$, and $S(D)$ is not a quasi-kernel of D, which we deal with now. Let $X_0 \coloneqq S(D)$, $X_1 \coloneqq N^-(S(D))$, $X_2 \coloneqq N^{--}(S(D)) \cap I(D)$, and $D' \coloneqq D[V(D) \setminus (X_0 \cup X_1 \cup X_2)]$. Since $S(D)$ is not a quasi-kernel, D' has at least one vertex. By induction, D' admits a quasi-kernel Q' of size at most $\alpha(|V(D')| + |S(D')| - |N^-(S(D'))|)$. Let $Q \coloneqq Q' \cup S(D)$. The set Q is an independent set by construction. Since every vertex in $X_1 \cup X_2$ is at distance at most two to $S(D)$, the set Q is a quasi-kernel of D.

We finish the proof by checking the size of Q. The construction of D' implies $|S(D')| \leqslant 1$ because any sink of D' must necessarily lie in $K(D)$. Note that if $N^-(S(D')) = \varnothing$, then D' is reduced to a single vertex, and since $S(D)$ is not a quasi-kernel of D and the vertex of D' is not a sink of D, the set X_2 is not empty. This implies that $|S(D')| - |N^-(S(D'))| \leqslant |X_2|$. We have thus

$$\begin{aligned}
|Q| &\leqslant \alpha\left(|V(D')| + |S(D')| - |N^-(S(D'))|\right) + |S(D)| \\
&\leqslant \alpha\,(|V(D)| - |X_0| - |X_1|) + |S(D)| \\
&= \alpha\left(|V(D)| - |N^-(S(D))|\right) + (1-\alpha)|S(D)| \\
&\leqslant \alpha\left(|V(D)| + |S(D)| - |N^-(S(D))|\right),
\end{aligned}$$

as required. (The last inequality comes from $1 - \alpha \leqslant \alpha$, which we have because $\alpha \geqslant 1/2$.) □

In addition to Conjecture 1, we consider another version of the conjecture about small quasi-kernels in digraphs with sinks.

Conjecture 2. *Every digraph* D *admits a quasi-kernel of size at most* $\frac{1}{2}(|V(D)| + |S(D)|)$.

The previous conjecture seems weaker than Conjecture 1. However, we prove that it is actually also equivalent to the small quasi-kernel conjecture.

Proposition 6.1.8. *Conjecture 2 and the small quasi-kernel conjecture are equivalent.*

Proof. Consider a digraph D beeing a counterexample to Conjecture 2 and D' a copy of D where each sink is replaced by a directed cycle of length two. See Figure 6.1 for an example. With this construction, D' is a counterexample to the small quasi-kernel conjecture. Indeed, suppose that D' has a quasi-kernel Q of size k. Then, contract every cycle of length two that arrive from sinks in the process of construction. Consider the

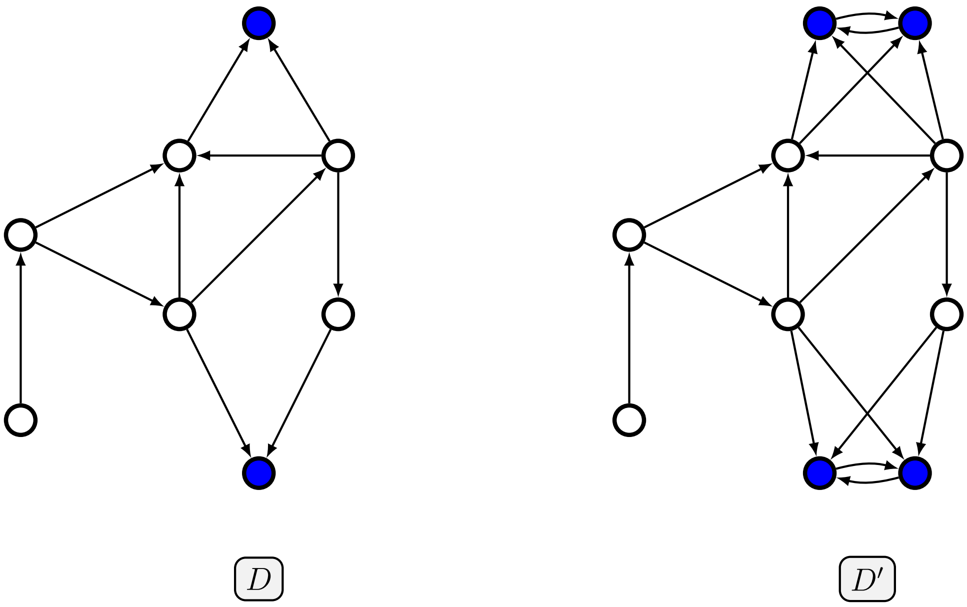

Figure 6.1: Example of the construction presented in the proof of Proposition 6.1.8.

set induced leads to a quasi-kernel of D of size k. Then, since for every quasi-kernel of D of size k we have $\frac{1}{2}(|V(D)| + |S(D)|) \leqslant k$, so every quasi-kernel of D' is of size $k \geqslant \frac{1}{2}(|V(D)| + |S(D)|) = \frac{1}{2}(|V(D')| - |S(D)| + |S(D)|) = \frac{1}{2}|V(D')|$. □

6.1.4 Split digraphs without a particular 3-cycle

Considering orientations of split digraphs forbidding a certain pattern, we obtain the small quasi-kernel conjecture for this particular class, as shown by the following proposition.

Proposition 6.1.9. *Let D be a sink-free split digraph. If there is no 3-cycle $v_1 \to v_2 \to v_3 \to v_1$ with $v_1 \in N^-(I(D)), v_2 \in I(D)$, and $v_3 \notin N^-(I(D))$, then D has a quasi-kernel of size at most $\frac{1}{2}|V(D)|$.*

Proof. Define $K_1 \coloneqq N^-(I(D))$ and $K_2 \coloneqq K(D) \setminus N^-(I(D))$. Consider the set Q_1 obtained by picking one vertex in $N^+(v) \cap I(D)$ per vertex v in K_1. Let $I_1 \coloneqq (I(D) \cap N^-(K_1)) \cup Q_1$. Then Q_1 forms a quasi-kernel of $K_1 \cup I_1$ and $|Q_1| \leqslant \frac{1}{2}(|K_1| + |I_1|)$. Consider Q_2 a quasi-kernel of $K_2 \cup (I(D) \setminus I_1)$ (from Proposition 6.1.4 since it forms a one-way split) of size $|Q_2| \leqslant \frac{1}{2}(|K_2| + |I(D) \setminus I_1|)$. Then $Q \coloneqq Q_1 \cup Q_2$ is of size $|Q| = |Q_1| + |Q_2| \leqslant \frac{1}{2}(|K(D)| + |I(D)|)$.

Now, note that if there are two vertices $x, y \in Q$ such that $x \to y$, then $x \in Q_1$ and $y \in Q_2 \cap K_2$ (because Q_1 and Q_2 are independent sets, $Q_1 \subset I(D)$, $I(D)$ is independent, and $K_2 \not\to I(D)$). Keep removing vertices x from Q_1 that have vertices in Q_2 in their outneighborhood. We eventually get an independent set Q' of size at most $\frac{1}{2}(|K(D)| + |I(D)|)$. This set Q' is a quasi-kernel of D since every vertex is at distance at most two to $Q_1 \cup Q_2$, and thus at distance two to Q': consider a vertex $x \in Q_1$ that has been removed because $x \to y$ for some $y \in Q_2 \cap K_2$; by the condition of the forbidden 3-cycle, the inneightborhood of x is included in the inneighborhood of y. □

6.2 Complete split graphs

A *complete split* graph is a split such that there is an edge between every pair of vertices in the clique-part and the independent-part.

Lemma 6.2.1. *Let D be an orientation of a complete split graph with no quasi-kernel of size one. Let x be a vertex with the maximum number of inneighbors in the clique-part. Denote by L the set of vertices with the same inneighborhood as x (including x). Then,*

- *the set L is included in the independent-part.*
- *every vertex v in L forms a quasi-kernel of $D[(V(D) \setminus L) \cup \{v\}]$.*

Proof. Suppose, aiming for a contradiction, that there is a vertex v in $K(D)$ with the maximum number of inneighbors in $K(D)$. Then $\{v\}$ is a quasi-kernel because every vertex in $N^+(v)$ has an outneighbor in $N^-(v)$, by the maximality of v; a contradiction with D having no quasi-kernel of size one. This proves the first item.

Consider a vertex v in L and a vertex u not in L. Suppose first that u is in $K(D)$. We have just seen that u has fewer inneighbors in $K(D)$ than v. Thus, u has an outneighbor in $N^-(v) \cup \{v\}$. Suppose now that u is in $I(D)$. Since u is not in L, it has an outneighbor in $N^-(v)$. In any case, there is a path of length at most two from u to v. This proves the second item. □

Theorem 6.2.2. *Let D be an orientation of a complete split graph. If D has a sink, then there is a unique minimum-size quasi-kernel, which is formed by all sinks. If D has no sink, then it has a quasi-kernel of size at most two.*

Proof. Observe that in an orientation of a complete split graph, if a vertex is a sink, then there is a path of length two from every other non-sink vertex to this sink. Thus, if D has at least one sink, then there is a unique inclusion-wise minimal quasi-kernel, which is formed by all sinks. Assume from now on that D has no sink and no quasi-kernel of size one. We are going to show that D has a quasi-kernel of size two.

Let x be a vertex maximizing $|N^-(x) \cap K(D)|$. We know from Lemma 6.2.1 that x is in $I(D)$. Suppose now, aiming for a contradiction, that every vertex v in $I(D)$ is such that $N^+(x) \subseteq N^+(v)$. Choose any vertex y in $N^+(x)$. The singleton $\{y\}$ is no quasi-kernel of $D[K(D)]$, since otherwise it would be a quasi-kernel of D of size one. A well-known consequence of the proof of Chvátal and Lovász is that in a digraph every vertex is in a quasi-kernel or has an outneighbor in a quasi-kernel. Thus, there exists a vertex z in $N^+(y) \cap K(D)$ that forms a quasi-kernel of $D[K(D)]$. The singleton $\{z\}$ is then a quasi-kernel of D as well since every vertex of $I(D)$ has y as outneighbor; a contradiction.

Hence, there is a vertex t in $I(D)$ with $N^+(x) \cap N^-(t) \neq \varnothing$. We claim that $\{x, t\}$ is a quasi-kernel of D. It is an independent set. Let L be the set of vertices having the same inneighborhood as x. Consider a vertex v in $V(D) \setminus \{x, t\}$.

If v is in L, then by definition of t there is a directed path of length two from v to t. If v is in $V(D) \setminus L$, Lemma 6.2.1 ensures that there is a directed path of length at most two from v to x. □

6.3 Complete bipartite graphs

We focus now on a class of graph which is not included in the class of split digraphs, the class of complete bipartite digraphs. A *complete bipartite digraph* is an orientation of a complete bipartite graph.

Proposition 6.3.1. *Let D be a complete bipartite digraph with no sink and denote by V_1 and V_2 its two parts. Then there exists a quasi-kernel of size at most* $\max\Big(\lceil \log_2(|V_1|) \rceil, \lceil \log_2(|V_2|) \rceil\Big) + 1$.

Proof. We prove by induction on $|V_1|$ and $|V_2|$ that D has a quasi-kernel of size at most $\lceil \log(|V_1|) \rceil + 1$ included in V_2 or a quasi-kernel of size at most $\lceil \log(|V_2|) \rceil + 1$ included in V_1.

Let $k = |V_1|$ and $\ell = |V_2|$. If $k = 2$ and $\ell = 2$, then the proposition is true. Consider now D with $k, \ell \geqslant 3$.

The digraph D has a vertex on a side whose inneighborhood is at least half of the size of the opposite side. Without loss of generality, D has a vertex v in V_2 with at least $k/2$ inneighbors. Consider $D - (v \cup N^-(v))$. If it has sinks, then the set formed by v and any sink of $D - (v \cup N^-(v))$ is a quasi-kernel of D of size two. If it has no sink, by induction, $D - (v \cup N^-(v))$ has (first case) a quasi-kernel included in $|V_1|$ of size at most $\lceil \log_2(\ell - 1) \rceil + 1$ or (second case) a quasi-kernel included in $|V_2|$ of size at most $\lceil \log_2(k/2) \rceil + 1 = \lceil \log_2(k) \rceil$.

In the first case: this quasi-kernel is included in $N^+(v)$ by hypothesis so it is a quasi-kernel of D and it size is at most $\lceil \log_2(\ell) \rceil + 1$. In the second case: we add v to the quasi-kernel and it forms a quasi-kernel of D of size at most $\lceil \log_2(k) \rceil + 1$. □

Chapter 7

Complexity

Sections 7.1, 7.2, and 7.3 form a paper that has been published at the conference WG 2022 [60]. Section 7.5 results from a collaboration with Julien Baste and Antoine Castillon.

In this chapter, we focus on the size of quasi-kernels in digraphs with an algorithmic point of view. Indeed, contrary to kernels, the existence of quasi-kernels is systematic. The question about the size of quasi-kernels is then very natural. The small quasi-kernel conjecture leads to other natural questions such as the complexity of finding a smallest quasi-kernel or of finding two disjoint quasi-kernels. We prove that the problem is computationally hard even for small classes of graphs. Finally we provide polynomial algorithms to find a smallest quasi-kernel and disjoint quasi-kernels in orientations of bounded treewidth.

7.1 Disjoint quasi-kernels

In 2011 Gutin et al. [47] conjectured that every sink-free digraph has two disjoint quasi-kernels. This conjecture implies the small quasi-kernel conjecture, since if a digraph has two disjoint quasi-kernels, at least one has less than half of the vertices of the digraph. In 2004, the same authors constructed a counterexample to this stronger conjecture with 14 vertices [48] depicted in Figure 7.1. As the following theorem proves, not only sink-free digraphs occasionally fail to contain two disjoint quasi-kernels, but it is actually computationally hard to distinguish those that do from those that do not.

Theorem 7.1.1. *Deciding if a digraph has two disjoint quasi-kernels is* NP*-complete, even for digraphs with maximum outdegree six.*

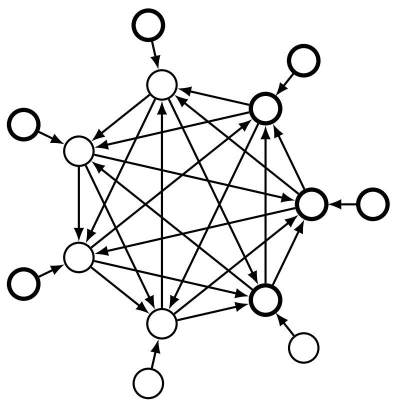

Figure 7.1: A split digraph having no two disjoint quasi-kernels.

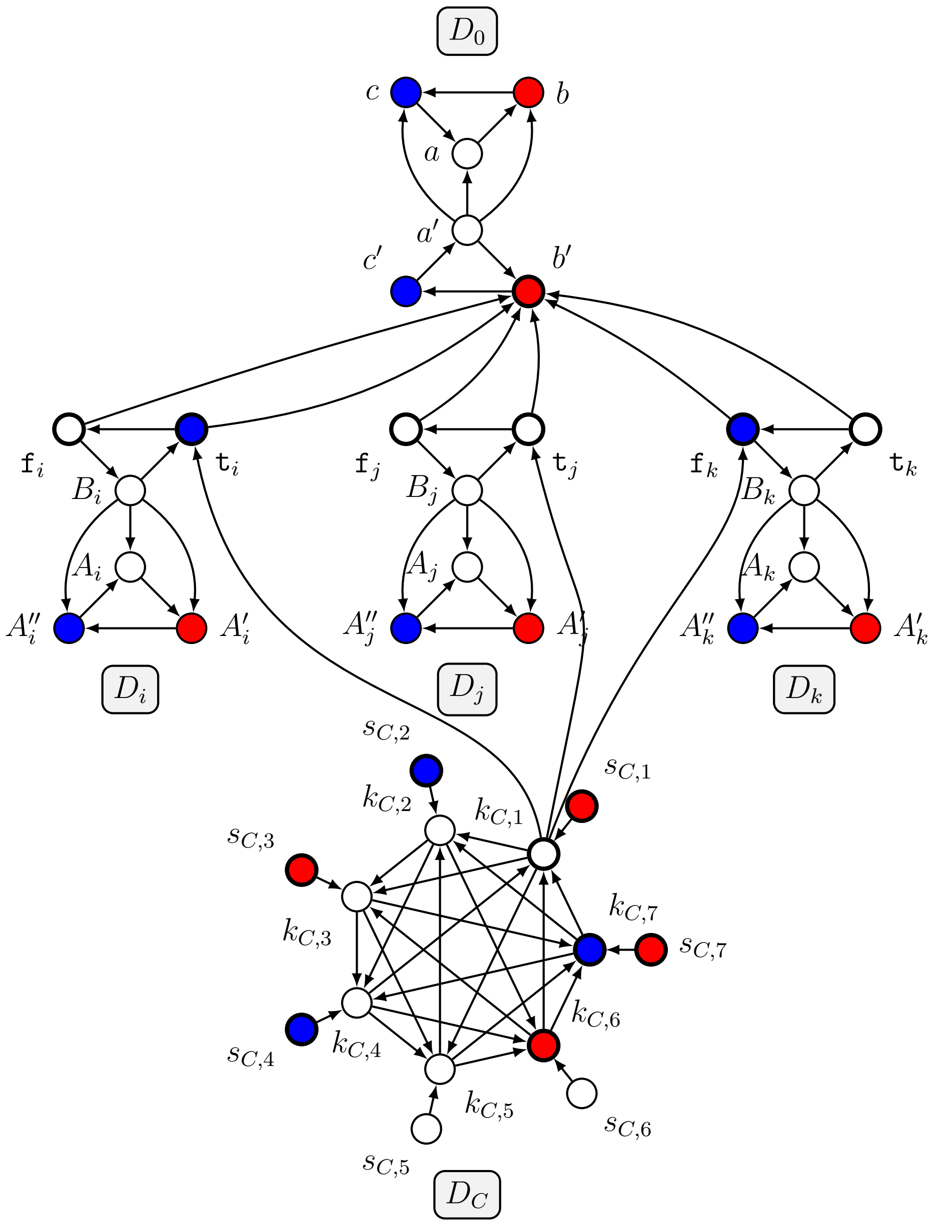

Figure 7.2: Proof of Theorem 7.1.1: Connecting the gadgets for clause $c = x_i \vee x_j \vee \neg x_k$. Red (resp. Blue) vertices denote vertices in Q_1 (resp. Q_2). Shown here is the case $\varphi(x_i) =$ $\mathtt{true}$, $\varphi(x_j) = \mathtt{false}$ and $\varphi(x_k) = \mathtt{false}$ (*i.e.*, $\mathtt{t}_i \in Q_2$, $\mathtt{f}_j \in Q_2$ and $\mathtt{f}_k \in Q_2$). Note that $\mathtt{f}_j \notin Q_2$ and $\mathtt{t}_j \notin Q_2$ implies $\varphi(x_j) = \mathtt{false}$.

Proof. Given a Boolean expression F in conjunctive normal form (CNF) where each clause is the disjunction of at most three distinct literals, 3-SAT asks to decide whether F is satisfiable. We reduce from 3-SAT, which is known to be NP-complete [55].

Consider an instance of 3-SAT. Let $X = \{x_1, x_2, \ldots, x_n\}$ be its variables, and let $F = C_1 \vee C_2 \vee \cdots \vee C_m$ be its CNF-formula. We construct a digraph D as follows.

- We start with the gadget D_0 shown on the top part of Figure 7.2 which contains the specified vertex b'.

- For every Boolean variable $x_i \in X$ we introduce the gadget D_i shown in the middle part of Figure 7.2 which contains two specified vertices $\mathtt{f}_i$ and $\mathtt{t}_i$. Furthermore, we connect D_i to D_0 with two arcs $(\mathtt{f}_i, b')$ and $(\mathtt{t}_i, b')$.

- For every clause $C = \ell_i \vee \ell_j \vee \ell_k$ of F we introduce the gadget D_C shown in the bottom part of Figure 7.2 which contains one specified vertex $k_{C,1}$. Furthermore, we connect D_C to the gadgets D_i, D_j, D_k with three arcs $(k_{C,1}, \lambda_i)$, $(k_{C,1}, \lambda_j)$ and $(k_{C,1}, \lambda_k)$, where $\lambda_i = \mathtt{t}_i$ (resp. $\lambda_j = \mathtt{t}_j$ and $\lambda_k = \mathtt{t}_k$) if ℓ_i (resp. ℓ_j and ℓ_k) is a positive literal, and $\lambda_i = \mathtt{f}_i$ (resp. $\lambda_j = \mathtt{f}_j$ & $\lambda_k = \mathtt{f}_k$) if ℓ_i (resp. ℓ_j and ℓ_k) is a negative literal.

Note that for every clause C of F, the digraph D_C is the counterexample constructed by Gutin et al. [48]. It has the important property that any two distinct vertices of $\{k_{C,i} \colon 1 \leqslant i \leqslant 7\}$ have a common outneighbor in $\{k_{C,i} \colon 1 \leqslant i \leqslant 7\}$.

It is clear that $|V(D)| = 14m + 6n + 6$ and $|A(D)| = 31m + 11n + 9$. Moreover, D has maximum outdegree six (but it has unbounded indegree; see vertex b'). We claim that the Boolean formula F is satisfiable if and only if the digraph D has two disjoint quasi-kernels.

Suppose that the Boolean formula F is satisfiable and consider any satisfying assignment φ. Construct two subsets $Q_1, Q_2 \subseteq V$ as follows.

- The elements of Q_1 are the following vertices: the vertices b and b' from D_0, the vertex A'_i from D_i for every variable $x_i \in X$, and the vertices $k_{C,6}$, $s_{C,1}$, $s_{C,3}$ and $s_{C,7}$ from D_C for every clause C of F.

- The elements of Q_2 are the following vertices: the vertices c and c' from D_0, the vertices A''_i and $\mathtt{t}_i$ from D_i for every variable $x_i \in X$ with $\varphi(x_i) = \mathtt{true}$, or the vertices A''_i and $\mathtt{f}_i$ from D_i with $\varphi(x_i) = \mathtt{false}$ and the vertices $k_{C,7}$, $s_{C,2}$ and $s_{C,4}$ from D_C for every clause C of F.

It is a simple matter to check that Q_1 and Q_2 are disjoint and that both Q_1 and Q_2 are independent subsets. Furthermore, we claim that Q_1 and Q_2 are two quasi-kernels of D. The claim is clear for Q_1. As for Q_2, it is enough to show that, for every clause C, the vertex $s_{C,1}$ is at distance at most two of some vertex in Q_2. Indeed, let $C = \ell_i \vee \ell_j \vee \ell_k$ be a clause where ℓ_i, ℓ_j and ℓ_k are positive or negative literals. Since φ is a satisfying assignment, there exists one literal, say ℓ_i, that evaluates to true in the clause C. Therefore, if $\varphi(x_i) = \mathtt{true}$ then $\mathtt{t}_i \in Q_2$ and $(k_{C,1}, \mathtt{t}_i) \in A$, and if $\varphi(x_i) = \mathtt{false}$ then $\mathtt{f}_i \in Q_2$ and $(k_{C,1}, \mathtt{f}_i) \in A$.

Conversely, suppose that there exist two disjoint quasi-kernels Q_1 and Q_2 in D. We first observe that $Q_1 \cap \{a, b, c\} \neq \varnothing$ and $Q_2 \cap \{a, b, c\} \neq \varnothing$. Then it follows that $a' \notin Q_1 \cup Q_2$ (by independence), and hence $b' \in Q_1 \cup Q_2$. Without loss of generality, suppose $b' \in Q_1$. Define an assignment φ for the Boolean formula F as follows: for $1 \leqslant i \leqslant n$, if $\mathtt{t}_i \in Q_2$ then set $\varphi(x_i) = \mathtt{true}$; otherwise set $\varphi(x_i) = \mathtt{false}$. Let us show that φ is a satisfying assignment.

By independence, we have $\mathtt{t}_i \notin Q_1$ and $\mathtt{f}_i \notin Q_1$ for $1 \leqslant i \leqslant n$.

We need the following claim.

Claim 7.1.2. *We have $\{k_{C,1}, k_{C,2}, k_{C,3}, k_{C,5}\} \cap (Q_1 \cup Q_2) = \varnothing$ for every clause C of F.*

Proof. We only prove $k_{C,1} \notin Q_1 \cup Q_2$ (the proof is similar for $k_{C,2} \notin Q_1 \cup Q_2$, $k_{C,3} \notin Q_1 \cup Q_2$ and $k_{C,5} \notin Q_1 \cup Q_2$.) Suppose, aiming at a contradiction, that $k_{C,1} \in Q_1 \cup Q_2$. Without loss of generality we may assume $k_{C,1} \in Q_1$ (the argument is symmetric if $k_{C,1} \in Q_2$). Then it follows that $\{s_{C,2}, s_{C,3}, s_{C,5}\} \subseteq Q_1$, and hence $\{s_{C,2}, s_{C,3}, s_{C,5}\} \cap Q_2 = \varnothing$. But, for any vertex $k_{C,i}$, $2 \leqslant i \leqslant 7$, we can easily check that either $d(s_{C,2}, k_{C,i}) > 2$, or $d(s_{C,3}, k_{C,i}) > 2$, or $d(s_{C,5}, k_{C,i}) > 2$. Hence Q_2 is not a quasi-kernel of D. This is the sought contradiction. □

Claim 7.1.3. *We have $s_{C,1} \in Q_1$ for every clause C of F.*

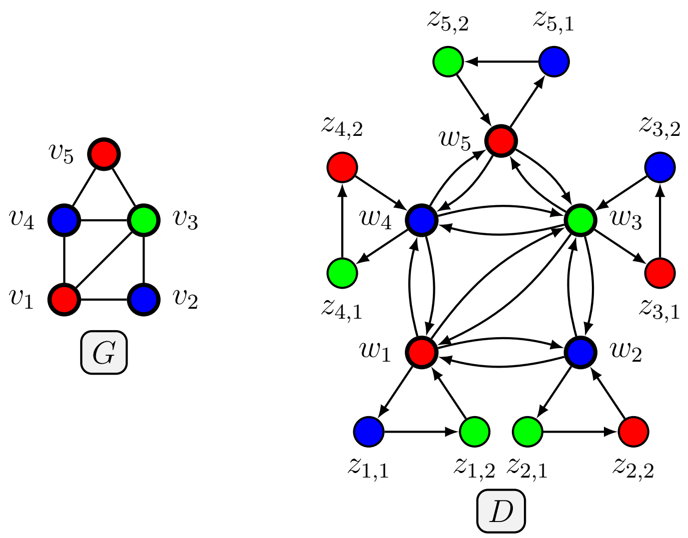

Figure 7.3: Example of the construction presented in the proof of Theorem 7.1.4.

Proof. Suppose, aiming at a contradiction, that $s_{C,1} \notin Q_1$. Combining Claim 7.1.2 with $\mathtt{t}_i \notin Q_1$ and $\mathtt{f}_i \notin Q_1$ for $1 \leqslant i \leqslant n$, we conclude that no vertex in Q_1 is at distance at most two from $s_{C,1}$. Therefore, Q_1 is not a quasi-kernel of D. This is a contradiction. □

Let $C = \ell_i \vee \ell_j \vee \ell_k$ be a clause. According to Claim 7.1.3, we have $s_{C,1} \in Q_1$. Furthermore, according to Claim 7.1.2, $\{k_{C,1}, k_{C,2}, k_{C,3}, k_{C,5}\} \cap Q_2 = \varnothing$. Then it follows that $\{\lambda_i, \lambda_j, \lambda_k\} \cap Q_2 \neq \varnothing$ where $\lambda_i = \mathtt{t}_i$ (resp. $\lambda_j = \mathtt{t}_j$ and $\lambda_k = \mathtt{t}_k$) if ℓ_i (resp. ℓ_j and ℓ_k) is a positive literal, and $\lambda_i = \mathtt{f}_i$ (resp. $\lambda_j = \mathtt{f}_j$ and $\lambda_k = \mathtt{f}_k$) if ℓ_i (resp. ℓ_j and ℓ_k) is a negative literal. Therefore φ is a satisfying assignment. □

Even though is it known that the small quasi-kernel conjecture is true for planar sink-free digraphs [59], no sink-free planar digraph without two disjoint quasi-kernels is known so far (the counterexample constructed by Gutin et al. [48] does contain a directed K_7). Whether such a planar graph exists is not known, we show that deciding whether a sink-free bounded degree planar digraph has three disjoint quasi-kernels is NP-complete.

Theorem 7.1.4. *Deciding if a digraph has three disjoint quasi-kernels is* NP*-complete, even for planar digraphs with maximum degree four.*

Proof. Given a planar graph with maximum degree four, 3-Coloring asks to decide whether there exists a proper vertex coloring of G with three colors (*i.e.*, a labeling of the vertices with three colors such that no two distinct vertices incident to a common edge have the same color). We reduce 3-Coloring which is known to be NP-complete for planar graphs with maximum degree four [31].

Let G be a planar graph with n vertices, m edges and maximum degree four. We denote by $v_1, v_2, \ldots, v_n$ the vertices of G. Without loss of generality, we assume that G has no isolated vertex. We construct a digraph D as follows. For every $1 \leqslant i \leqslant n$, we introduce C_i an oriented cycle of length three which contains a specified vertex w_i. For every edge $v_iv_j \in E(G)$, we connect the two gadgets C_i and C_j with two arcs (w_i, w_j) and (w_j, w_i) to D. More formally, we have

$$V(D) = \{w_i \colon 1 \leqslant i \leqslant n\} \cup \{z_{i,j} \colon 1 \leqslant i \leqslant n \text{ and } 1 \leqslant j \leqslant 2\},$$
$$A(D) = \{(w_i, w_j), (w_j, w_i) \colon v_iv_j \in E(G)\} \cup \{(w_i, z_{i,1}), (z_{i,1}, z_{i,2}), (z_{i,2}, w_i) \colon 1 \leqslant i \leqslant n\}.$$

It is clear that $|V'| = 3n$, $|A| = 2m+3n$ and that the digraph D is planar. Moreover, D has maximum indegree five and maximum outdegree five. See Figure 7.3 for an example.

We claim that G has a proper 3-coloring if and only if D has three distinct quasi-kernels.

Suppose first that G has a proper 3-coloring. Let $C = \{c_1, c_2, c_3\}$ be the three colors. Consider the 3-coloring of D defined as follows: if v_i is colored with color c_1 (resp. c_2 & c_3) in G, then color w_i with color c_1 (resp. c_2 & c_3), $z_{i,1}$ with color c_2 (resp. c_3 & c_1), and $z_{i,2}$ with color c_3 (resp. c_1 & c_2) in D. It is clear that the 3-coloring of D is proper. Moreover, for each color c and every vertex v there exists a directed path of length at most two from v to a vertex colored with color c. Then it follows that the 3-coloring of D induces three disjoint quasi-kernels in D. (Note that these three disjoint quasi-kernels actually form a partition of the vertices of D.)

Conversely, suppose that the digraph D has three disjoint quasi-kernels Q_1, Q_2, and Q_3. We have $w_i \in Q_1 \cup Q_2 \cup Q_3$ for every $1 \leqslant i \leqslant n$. Indeed, $\{z_{i,1}\} \cup N^+(z_{i,1}) \cup N^+(N^+(z_{i,1})) = \{z_{i,1}, z_{i,2}, w_i\}$. Thus a quasi-kernel necessarily contains one vertex in $\{z_{i,1}, z_{i,2}, w_i\}$, which implies that one of the three disjoint quasi-kernels contains w_i. Furthermore, each subset Q_1, Q_2, and Q_3 is independent because it is a quasi-kernel. Therefore, the three quasi-kernels Q_1, Q_2, and Q_3 induce a proper 3-coloring of G. □

7.2 Acyclic digraphs

In this section, we address the complexity status of Quasi-Kernel for acyclic orientations of various classes of graphs. The next two theorems show that there is not so much room for extending the positive result about orientations of graphs with bounded treewidth that will be shown in Section 7.5.

We recall that a *cubic* graph is a graph in which every vertex has degree three.

Theorem 7.2.1. Quasi-Kernel *is* NP*-complete, even for acyclic orientations of cubic graphs.*

Proof. Given a Boolean expression F in conjunctive normal form where each clause is the disjunction of three distinct literals and each literal occurs exactly twice among the clauses, (3,B2)-SAT asks to decide whether F is satisfiable. We reduce (3,B2)-SAT which is known to be NP-complete [11].

Consider an instance of (3,B2)-SAT. Denote by $x_1, x_2, \ldots, x_n$ its variables, and let $F = C_1 \vee C_2 \vee \cdots \vee C_m$ be its formula. For every Boolean variable x_i occurring in F we introduce a copy D_{x_i} of the gadget shown in Figure 7.4a which contains two specified vertices x_i and $\neg x_i$. For every clause C_j of F we introduce a copy D_{C_j} of the gadget shown in Figure 7.4b which contains three specified vertices $C_{j,1}$, $C_{j,2}$, and $C_{j,3}$. Furthermore, for every clause C_j of F and every $1 \leqslant k \leqslant 3$, we introduce a copy $D_{C_{j,k}}$ of the gadget shown in Figure 7.4c which contains one specified vertex $C''_{j,k}$.

If the literal x_i (resp. $\neg x_i$) occurs in clause C_j as the k-th literal ($k = 1, 2, 3$), we add the directed path $\{(C_{j,k}, C'_{j,k}), (C'_{j,k}, x_i)\}$ (resp. $\{(C_{j,k}, C'_{j,k}), (C'_{j,k}, \neg x_i)\}$ and the arc $(C''_{j,k}, C'_{j,k})$ as shown in Figure 7.5. Let D denote the resulting digraph. We observe that D is an orientation of a cubic graph (since every literal occurs twice) with $14n+21m+18m = 14n + 39m$ vertices.

We claim that F is satisfiable if and only if the digraph D has a quasi-kernel of size $8m + 4n$.

Suppose first that F is satisfiable and consider a satisfying assignment φ. Construct a subset Q of vertices of D as follows.

- For $1 \leqslant j \leqslant m$, if the clause C_j is satisfied by its first (resp. second and third) literal, add the vertices $C_{j,2}$ and $C_{j,3}$ (resp. $C_{j,1}$ and $C_{j,3}$, and $C_{j,1}$ and $C_{j,2}$) to Q. In case the clause C_j is satisfied by more than one literal, we choose one satisfying literal arbitrarily. Furthermore, add the three vertices t_j, t'_j, and t''_j to Q.

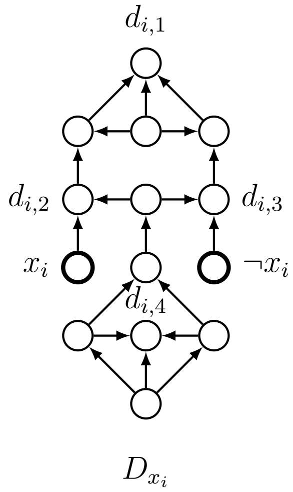

(a) Gadget D_{x_i} for Boolean variable x_i.

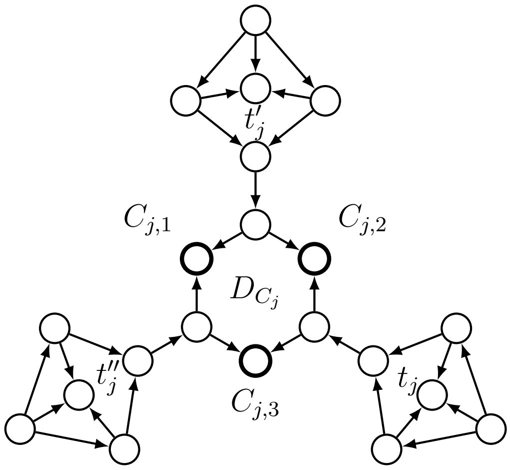

(b) Gadget D_{C_j} for clause C_j.

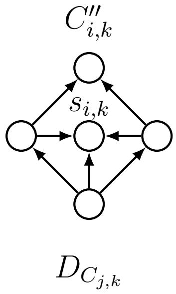

(c) Gadget $D_{C_{j,k}}$ for the k-th literal of clause C_j.

Figure 7.4: The gadgets in the proof of Theorem 7.2.1.

- For $1 \leqslant i \leqslant n$, if $\varphi(x_i) = \texttt{true}$ (resp. $\varphi(x_i) = \texttt{false}$) add the vertices x_i and $d_{i,3}$ (resp. $\neg x_i$ and $d_{i,2}$) to Q. Furthermore, add the two vertices $d_{i,1}$ and $d_{i,4}$ to Q.
- For $1 \leqslant j \leqslant m$ and $1 \leqslant k \leqslant 3$, add $s_{j,k}$ to Q.

We check at once that $|Q| = 8m + 4n$. It is now a simple matter to check that Q is a quasi-kernel of D.

Conversely, let Q be a quasi-kernel of D of size $8m + 4n$. We first observe that Q contains at least five vertices of each gadget D_{C_j} and one vertex from each gadget $D_{C_{j,k}}$. We also observe that Q contains at least four vertices of each gadget D_{x_i}. Then it follows that Q contains exactly five vertices of each gadget D_{C_j}, and exactly 4 vertices of each gadget D_{x_i}. Note that we have actually shown that no mid-vertex of a path of length two connecting some gadget D_{C_j} to some gadget D_{x_i} is in Q. Furthermore, while we have a fair degree of flexibility in the way the vertices are selected, all sinks have to be in Q (by definition). In particular, for each j, the three sinks t_j, t'_j, and t''_j are in Q, and for each i, the two sinks $d_{i,1}$ and $d_{j,4}$ are in Q as well. We now turn to defining a truth assignment φ for the variable of F. As Q contains exactly four vertices in every gadget D_{x_i}, we are left to consider the three possibilities depicted in Figure 7.6, where the red vertices denote the vertices in Q. The truth assignment φ is defined as follows: $\varphi(x_i) = \texttt{false}$ if and

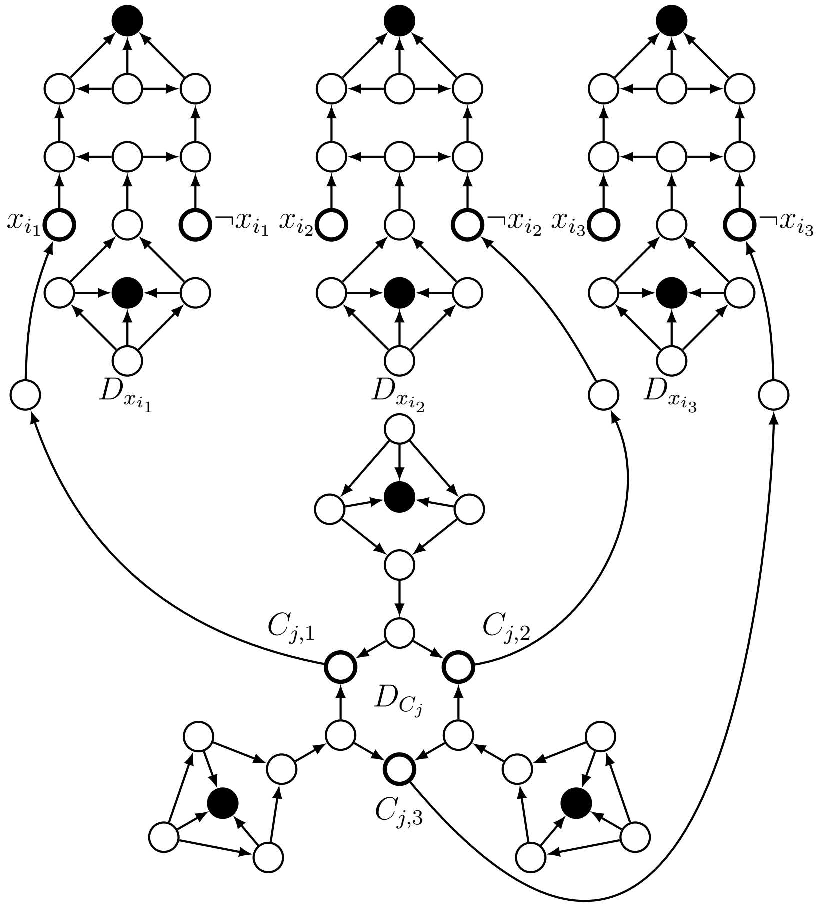

Figure 7.5: Proof of Theorem 7.2.1: connecting gadget D_{C_j} to gadgets $D_{x_{i_1}}$, $D_{x_{i_2}}$ and $D_{x_{i_3}}$ for clause $C_j = x_{i_1} \vee \neg x_{i_2} \vee \neg x_{i_3}$. Shown here is the assignment $\varphi(x_{i_1}) = \texttt{true}$, $\varphi(x_{i_2}) = \texttt{true}$ and $\varphi(x_{i_3}) = \texttt{false}$, and the clause C_j is satisfied by its first literal.

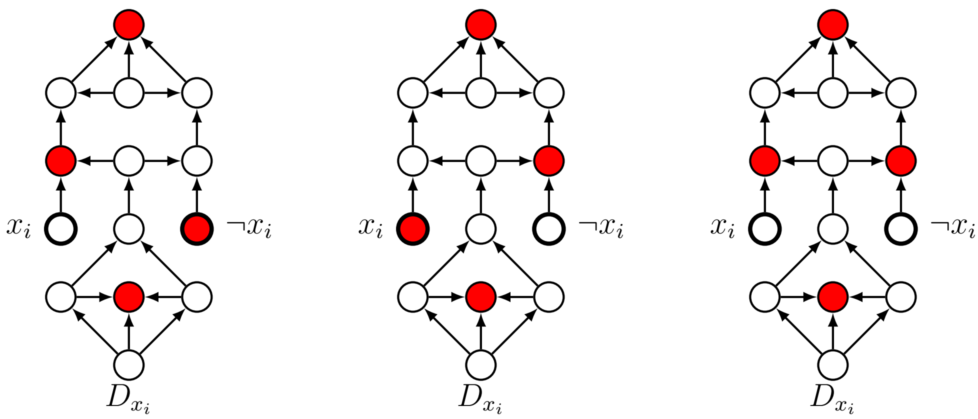

Figure 7.6: Proof of Theorem 7.2.1: truth selection.

only if $\neg x_i \in Q$. We claim that φ is a satisfying assignment. Indeed, consider any clause C_j. Combining the observation that t_j, t'_j, and t''_j are in Q together with the fact that Q contains exactly five vertices of the gadget D_{C_j}, we conclude that one of $C_{i,1}$, $C_{j,2}$, and $C_{j,3}$ is not in Q, and that there exists a directed path of length two from this vertex to some vertex x_i or $\neg x_i$ in Q. (Note that $\varphi(x_i) = \texttt{true}$ if and only if $x_i \in Q$ yields another satisfying assignment for the proposed construction.) □

Assuming $\mathsf{FPT} \neq \mathsf{W}[2]$, our next result shows that one cannot confine the seemingly inevitable combinatorial explosion of computational difficulty to an additive function of the size of the quasi-kernel, even for restricted digraph classes.

Theorem 7.2.2. Quasi-Kernel *is* W[2]*-complete when the parameter is the size of the sought quasi-kernel, even for acyclic orientations of bipartite graphs.*

Proof. Membership to W[2] is easy.

Given $\mathcal{F}$ a family of sets over a universe U and a positive integer k, Set Cover consists in deciding if there exists a subfamily $\mathcal{F}' \subseteq \mathcal{F}$ of size at most k such that $\mathcal{F}'$ covers U. We prove hardness by reducing from Set Cover which is known to be W[2]-complete [35].

Let $\mathcal{F} = \{F_1, F_2, \ldots, F_m\}$ be a family of sets over some universe $U = \{u_1, u_2, \ldots, u_n\}$ and k be a positive integer. Without loss of generality, we may assume $U = \bigcup_{1 \leqslant j \leqslant m} F_j$. We show how to produce a digraph D such that $\mathcal{F}$ has a set cover of size at most k if and only if D has a quasi-kernel of size at most $k' = k + 1$. The digraph D is defined as follows:

$$\begin{aligned} V(D) &= \mathcal{F} \cup U \cup \{s, t\}\,, \\ A(D) &= \{(u_i, F_j) \colon F_j \in \mathcal{F} \text{ and } u_i \in F_j\} \cup \{(F_j, s) \colon F_j \in \mathcal{F}\} \cup \{(s, t)\}\,. \end{aligned}$$

It is clear that $|V(D)| = m + n + 2$ and that $|A(D)| = m + 1 + \sum_{1 \leqslant j \leqslant m} |F_j|$.

Suppose that there exists a subfamily $\mathcal{F}' \subseteq \mathcal{F}$ of size at most k such that $\mathcal{F}'$ covers U. It is clear that $\mathcal{F}' \cup \{t\}$ is a quasi-kernel of D of size $k' = k + 1$.

Conversely, suppose that there exists a quasi-kernel of D of size at most $k' = k + 1$. Observe that $t \in Q$ since t is a sink, and $s \notin Q$ (by independence). Among these quasi-kernels, choose one Q that minimizes $|Q \cap U|$. We show that $Q \cap U = \varnothing$. Indeed, suppose, aiming at a contradiction, that $Q \cap U \neq \varnothing$ and let $u_i \in Q \cap U$. Furthermore, let $F_j \in N^+(u_i)$ and $U' = N^-(F_j)$. Since $s \notin Q$ we have $Q' = (Q \setminus U') \cup \{F_j\}$ is a quasi-kernel of D of size at most k' with $|Q' \cap U| < |Q \cap U|$. This contradicts our assumption, and hence $Q \cap U = \varnothing$. Then it follows $N^+(u_i) \cap Q \neq \varnothing$ for every $u_i \in U$, and hence $Q \cap \mathcal{F}$ yields a set cover of size $k' - 1 = k$. □

We finish the section with a series of propositions providing complementary evidence for the versatile hardness of computing small quasi-kernels.

Recall that a kernel is a quasi-kernel. Actually we have more: a kernel is an inclusion-wise maximal quasi-kernel. Inclusion-wise minimal quasi-kernels are easy to find with a greedy algorithm. Though, finding a minimum-size quasi-kernel included in a kernel is hard as shown by the following result, whose proof is identical to the one of Theorem 7.2.2 ($\mathcal{F} \cup \{t\}$ is actually a kernel of the digraph D).

Proposition 7.2.3. *Let D be an acyclic orientation of a bipartite graph, $K \subseteq V(D)$ be a kernel of D and k be a positive integer. Deciding whether there exists a quasi-kernel included in K of size k is* W[2]*-complete for parameter k.*

Dinur and Steuer [34] have shown that Set Cover cannot be approximated in polynomial time within a factor of $(1 - \varepsilon)\ln(|U|)$ for some constant $\varepsilon > 0$ unless P = NP. Moreover, they built an instance of Set Cover where the number of subsets is a polynomial of the universe size. Therefore, the construction used in the proof of Theorem 7.2.2 allows us to state the following inapproximability result.

Proposition 7.2.4. Quasi-Kernel *cannot be approximated in polynomial time within a factor of $(1 - \varepsilon)\ln(|V(D)|)$ for some constant $\varepsilon > 0$ unless* P = NP*, even for acyclic orientations of bipartite graphs.*

Our last result focuses on another restricted classes of digraphs, namely acyclic digraphs with bounded indegrees. We actually do not know if the problem is in APX for digraphs with unbounded indegrees. We need a preliminary lemma which we state for general digraphs.

Lemma 7.2.5. Quasi-Kernel *belongs to* APX *for digraphs with fixed maximum indegrees.*

Proof. Let D be a digraph and $Q \subseteq V(D)$ be a quasi-kernel. It is clear that $(d^2 + d + 1)|Q| \geqslant |V(D)|$, where d is the maximum indegree of D. Then it follows that any polynomial-time algorithm that computes a quasi-kernel (such as the natural algorithm from the proof Chvátal and Lovász, see Section 5.1) is a $(d^2 + d + 1)$-approximation algorithm. □

Proposition 7.2.6. Quasi-Kernel *is* APX*-complete for acyclic digraphs with maximum indegree three and maximum outdegree two.*

Proof. Membership in APX for acyclic digraphs with fixed indegrees follows from Lemma 7.2.5. Specifically, Quasi-Kernel for acyclic digraphs with maximum indegree three can be approximated in polynomial time within a factor of 13.

To prove hardness, we L-reduce from Vertex Cover in cubic graphs which is known to be APX-complete [4]. As defined in [66], letting P and P' be two optimization problems, we say that P L-reduces to P' if there are two polynomial-time algorihtms f, g, and constants $\alpha, \beta > 0$ such that for each instance I of P : algorithm f produces an instance $I' = f(I)$ of P, such that the optima of I and I', $OPT(I)$ and $OPT(I')$, respectively, satisfy $OPT(I') \leqslant \alpha OPT(I)$ and given any solution of I' with cost c', algorithm g produces a solution of I with cost c such that $|c - OPT(I)| \leqslant \beta|c' - OPT(I')|$. Let f be the following L-reduction from Vertex Cover in cubic graphs to Quasi-Kernel with maximum indegree three. Given a cubic graph G with $V(G) = [n]$ and m edges, we construct a digraph D as follows:

$$
\begin{aligned}
V(D) &= \{w_i, w'_i, w''_i \colon 1 \leqslant i \leqslant n\} \cup \{z_e, z'_e \colon e \in E(G)\}\,, \\
A(D) &= \{(w_i, w'_i), (w'_i, w''_i) \colon 1 \leqslant i \leqslant n\} \cup \{(z'_e, z_e), (z_e, w_i), (z_e, w_j) \colon e = ij \in E(G)\}\,.
\end{aligned}
$$

Note that the vertices w''_i are sinks in D. It is clear that $|V'| = 3n + 2m$, $|A| = 2n + 3m$ and, since G is a cubic graph, that every vertex has maximum indegree three in D. We also observe that the maximum outdegree is two in D. See Figure 7.7 for an example.

Consider a quasi-kernel $Q \subseteq V(D)$ of $D = f(G)$. We claim that it can be transformed in polynomial time into a vertex cover $C \subseteq V(G)$ of G such that $|C| \leqslant |Q|$. To see this, observe first that Q can be transformed in polynomial time into a quasi-kernel $Q' \subseteq V(D)$ such that (i) $|Q'| \leqslant |Q|$ and (ii) $z'_e \notin Q'$ and $z_e \notin Q'$ for every $e \in E(G)$. Indeed, repeated applications of the following two procedures enable us to achieve the claimed quasi-kernel.

- Suppose that there exists $z'_e \in Q$ for some $e = ij \in E(G)$. Then it follows that $z_e \notin Q$ (by independence). Furthermore, we have $w''_i \in Q$ and $w''_j \in Q$, and hence $w'_i \notin Q$ and $w'_j \notin Q$. Therefore, $w_i \in Q$ or $w_j \in Q$ (possibly both). On account of the above remarks, $Q' = Q \setminus \{z'_e\}$ is a quasi-kernel of D and $|Q'| < |Q|$.

- Let $Z_i \subseteq Q$ stand for the set of vertices $z_e \in Q$, where e is an edge incident to the vertex i in G. Suppose that there exists some set $Z_i \neq \varnothing$. Then it follows that $w_i \notin Q$ (by independence). Furthermore, we have $w''_i \in Q$, and hence $w'_i \notin Q$. On account of the above remarks, $Q' = (Q \setminus Z_i) \cup \{w_i\}$ is a quasi-kernel of D and $|Q'| \leqslant |Q|$.

From such a Q', construct then a vertex cover $C \subseteq V(G)$ of G as follows: for $1 \leqslant i \leqslant n$, add the vertex i to C if $w_i \in Q'$. By construction, C is a vertex cover of G of size $|C| = |Q'| - |V(G)|$

Finally, it is easy to see that from a vertex cover $C \in V(G)$ of G we can construct a quasi-kernel $Q \subseteq V(D)$ of $D = f(G)$ of size exactly $|C| + |V(G)|$: for every $1 \leqslant i \leqslant n$,

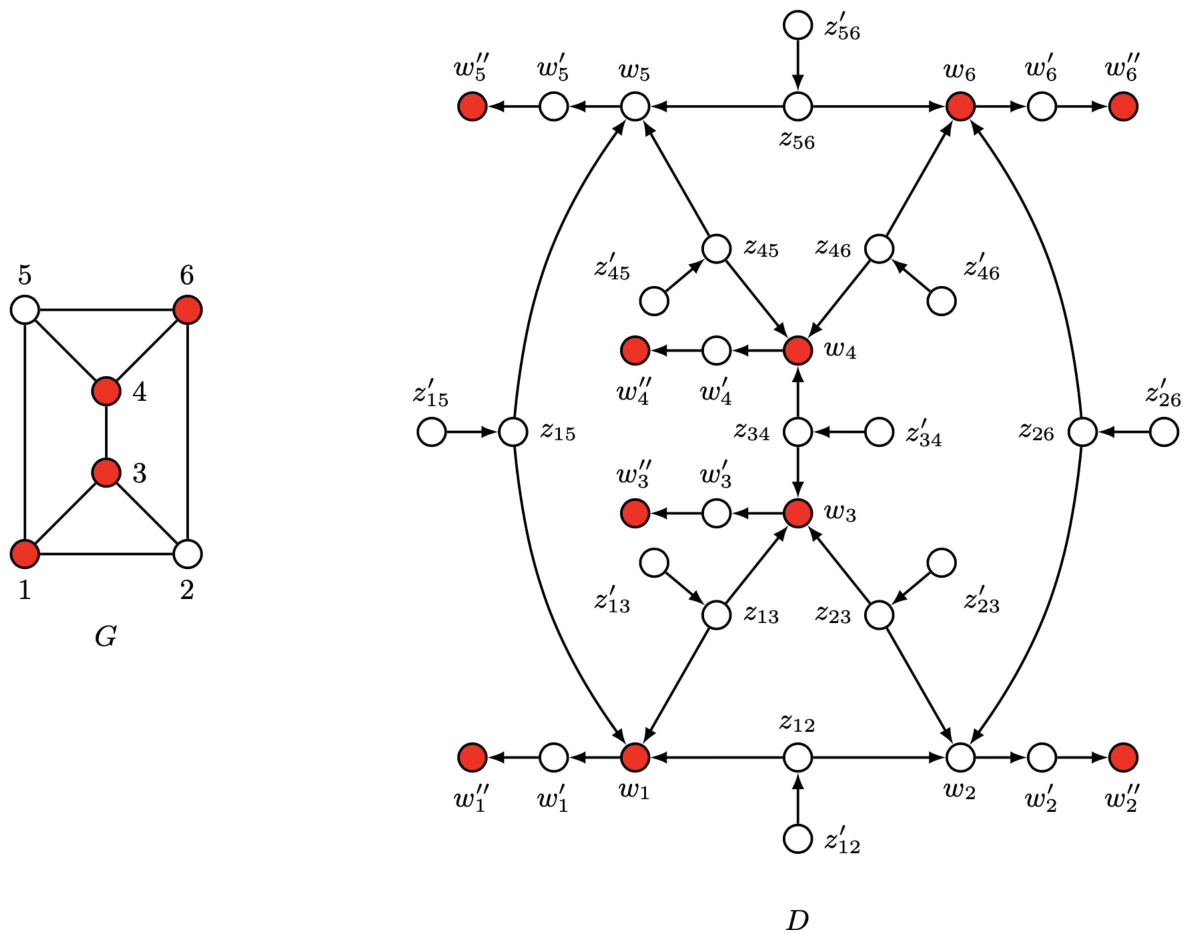

Figure 7.7: Example of the construction presented in the proof of Proposition 7.2.6.

add w''_i to Q and add w_i to Q if $i \in C$. Since G is a cubic graph, we have $|C| \geqslant |V(G)|/4$, and hence $|Q| = |C| + |V(G)| \leqslant |C| + 4|C| = 5|C|$.

Thus $\text{opt}(f(G)) \leqslant 5\,\text{opt}(G)$ and we have shown that f is an L-reduction with parameters $\alpha = 5$ and $\beta = 1$. □

7.3 Orientations of split graphs

In this section, we focus on orientations of split graphs. A split is a whose vertices can be partitioned into a clique and an independent set. This class seems to play an important role in the study of small quasi-kernels since the only examples of oriented graphs having no two disjoint quasi-kernels contain the orientation of a split graph constructed by Gutin et al. [48].

We first show that one cannot confine the seemingly inevitable combinatorial explosion of computational hardness to the size of the sought quasi-kernel.

Proposition 7.3.1. Quasi-Kernel *is* W[2]*-complete when the parameter is the size of the sought quasi-kernel even for orientations of split graphs.*

Proof. Membership in W[2] is clear. Given a digraph D and an integer q, Directed Dominating Set is the problem of deciding if there exists a dominating set of size q, i.e., a subset $L \subseteq V(D)$ of size q such that every vertex $v \in V(D)$ is either in L or has an outneighbor in L. Directed Dominating Set is W[2]-complete for parameter q [?].

We reduce Directed Dominating Set to Quasi-Kernel. Let D be a digraph and q be a positive integer. Let $n \coloneqq |V(D)|$, $m \coloneqq |A(D)|$, and $b \coloneqq 2q+3$. Consider moreover an arbitrary total order $\preccurlyeq$ on $A(D)$. Define the following split digraph D':

$$
\begin{aligned}
V(D') &\coloneqq \{s\} \cup S^1 \cup S^2 \cup K^1 \cup K^2 \\
A(D') &\coloneqq A_s \cup A_{S^1} \cup A_{S^2} \cup A_{K^1} \cup A_{K^2}
\end{aligned}
$$

where

$$
\begin{aligned}
S^1 &= \left\{s^1_v \colon v \in V(D)\right\} \\
S^2 &= \left\{s^2_i \colon 1 \leqslant i \leqslant b\right\} \\
K^1 &= \left\{k^1_a \colon a \in A(D)\right\} \\
K^2 &= \left\{k^2_i \colon 1 \leqslant i \leqslant b\right\}
\end{aligned}
$$

and

$$
\begin{aligned}
A_s &= \left\{\left(s, k^1_a\right) : a \in A(D)\right\} \cup \left\{\left(k^2_i, s\right) : 1 \leqslant i \leqslant b\right\} \\
A_{S^1} &= \left\{\left(s^1_v, k^1_{(v,v')}\right) : (v,v') \in A(D)\right\} \cup \left\{\left(k^1_{(v,v')}, s^1_{v'}\right) : (v,v') \in A(D)\right\} \\
A_{S^2} &= \left\{\left(s^2_i, k^2_i\right) : 1 \leqslant i \leqslant b\right\} \\
A_{K^1} &= \left\{\left(k^1_a, k^1_{a'}\right) : a, a' \in A(D), a \prec a'\right\} \cup \left\{(k^1_a, k^2_\ell) : a \in A(D), 1 \leqslant \ell \leqslant b\right\} \\
A_{K^2} &= \left\{(k^2_i, k^2_j) \colon 1 \leqslant i < j \leqslant b,\ i = j \bmod 2\right\} \cup \left\{(k^2_j, k^2_i) \colon 1 \leqslant i < j \leqslant b, i \neq j \bmod 2\right\}.
\end{aligned}
$$

Clearly, D' is an orientation of a split graph (i.e., $\{s\} \cup S^1 \cup S^2$ is an independent set and $K^1 \cup K^2$ induces a tournament), $|V(D')| = n+m+2b+1$, and $|A(D')| = \binom{m+b}{2} + 3m + 2b$.

We claim that there exists a dominating set of size at most q in D if and only if D' has a quasi-kernel of size at most $q+1$.

Suppose first that there exists a dominating set $L \subseteq V(D)$ of size at most q in D. Define $Q = \{s\} \cup \{s^1_v \colon v \in L\}$. We note that $Q \subseteq \{s\} \cup S^1$, and hence Q is an independent set. Furthermore, by construction, the vertex s is at distance at most two from every vertex in $S^2 \cup K^1 \cup K^2$. Since L is a dominating set, it is now clear that Q is a quasi-kernel of D' of size at most $q+1$.

Conversely, suppose that there exists a quasi-kernel $Q \subseteq V(D')$ of size at most $q+1$ in D'. By independence of Q, we have $|Q \cap (K^1 \cup K^2)| \leqslant 1$. We first claim that $s \in Q$. Indeed, suppose, aiming at a contradiction, that $s \notin Q$. Let $X = S^2 \setminus Q$. By construction, $N^+(X) = \{k^2_i \in K^2 \colon s^2_i \in X\}$. On the one hand, we have $|X| > |S^2| - |Q| \geqslant b - (q+1) = q+2$ (note that S^2 cannot contain the whole set Q), and hence $|N^+(X)| > q+2$. On the other hand, $|X|$ being positive, there exists $k^2_j \in K^2 \cap Q$ such that $N^+(X) \subseteq N^-[k^2_j] \cap K^2$. But, according to the definition of A_{K^2}, we have $|N^-[k^2_j] \cap K^2| \leqslant \lceil b/2 \rceil \leqslant q+2$ for all $k^2_i \in K^2$ and in particular for k^2_j. This is a contradiction and hence $s \in Q$. We now observe that $k^1_a \in N^+(s)$ for every $k^1_a \in K^1$ and $s \in N^+(k^2_i)$ for every $k^2_i \in K^2$. Combining this observation with $s \in Q$ and the independence of Q, we obtain $Q \cap (K^1 \cup K^2) = \varnothing$. We have thus $|S^1 \cap Q| \leqslant q$. We now turn to S^1. It is clear that s is at distance three from every vertex $s^1_v \in S_1$. Therefore, by definition of quasi-kernels, for every vertex $s^1_v \in S^1 \setminus Q$, there exists one vertex $s^1_{v'} \in S^1 \cap Q$ such that $(s^1_v, k^1_{(v,v')}) \in A(D')$ and $(k^1_{(v,v')}, s^1_{v'}) \in A(D')$. Note that, by construction, $(s^1_v, k^1_{(v,v')})$ and $(k^1_{(v,v')}, s^1_{v'})$ are two arcs of D' if and only (v,v') is an arc of D. Then it follows that $L = \{v \colon s^1_v \in Q\}$ is a dominating set in D of size at most q. □

Proposition 7.3.2. *Quasi-Kernel for orientations of split graphs is* FPT *for parameter* $|K(D)|$ *or parameter* $k + |I(D)|$*, where* k *is the size of the sought quasi-kernel.*

Proof. Let D be an orientation of a split graph, and write $n = |K(D) \cup I(D)|$. Let M be the adjacency matrix of D. It is clear that, after having computed M^2, one can decide in linear time if any given subset $Q \subseteq K(D) \cup I(D)$ is a quasi-kernel of D. This preprocessing step is $\mathcal{O}(n^3)$ time (a better running time can be achieved by fast matrix multiplication but is not relevant here). Furthermore, by independence of quasi-kernels, we have $|Q \cap K(D)| \leqslant 1$ for every quasi-kernel Q of D. This straightforward observation is the first step of the two algorithms.

Algorithm for parameter $|K(D)|$. Select (including none) a vertex of $K(D)$. Define the equivalence relation $\sim$ on I as follows: $s \sim s'$ if and only if $N^-(s) = N^-(s')$ and $N^+(s) = N^+(s')$. The key point is to observe that in any minimum cardinality quasi-kernel Q of D, for every equivalence class $I' \in I(D)/\sim$, either $I' \cap Q = \varnothing$, $I' \subseteq Q$ or any vertex of I' is in Q. For any combination, check if the selected vertices of $I(D)$ together with the selected vertex of $K(D)$ (if any) is a quasi-kernel of D of size k. The size of $I(D)/\sim$ is bounded by $4^{|K(D)|}$ since each equivalence class is determined by its out and inneighborhood. The algorithm is $\mathcal{O}(n^3 + k\,|K(D)|\,3^{|I(D)/\sim|}) = \mathcal{O}(n^3 + k\,|K(D)|\,3^{(4^{|K(D)|})})$ time.

Algorithm for parameter $k + |I(D)|$. Select (including none) a vertex of $K(D)$. For every subset $I' \subseteq I(D)$ of size $k - 1$ (or k, if no vertex of K is selected), check if I' together with the selected vertex of $K(D)$ is a quasi-kernel of D. The algorithm is $\mathcal{O}(n^3 + k\,|K(D)|\binom{|I(D)|}{k})$ time. □

7.4 Orientations of 4-partite complete graphs

In the section we provide a last hardness result about Quasi-Kernel. An immediate corollary of Theorem 6.3.1 is the existence of an algorithm to find a quasi-kernel of minimum size in any orientation of a complete bipartite graph. The next theorem ensures that the problem becomes hard for 4-partite graphs. The question about 3-partite graphs is still open.

Theorem 7.4.1. Quasi-Kernel *is* W[2]*-complete, even of orientations of complete* 4*-partite graphs.*

Proof. Given $\mathcal{S}$ a family of sets over a universe $\mathcal{X}$ and a positive integer k, Hitting Set consists in deciding if there exists a subfamily $\mathcal{X}' \subseteq \mathcal{X}$ of size at most k such that every set in $\mathcal{S}$ intersects $\mathcal{X}'$. We prove hardness of Quasi-Kernel for orientations of complete 4-partite graphs by reducing from Hitting set which is known to be W[2]-complete [35].

Let $\mathcal{S} = \{S_1, S_2, \ldots, S_m\}$ be a family of sets over some universe $\mathcal{X} = \{x_1, x_2, \ldots, x_n\}$ and k be a positive integer. Without loss of generality, we may assume $\mathcal{X} = \bigcup_{1 \leqslant j \leqslant m} S_j$ and $\bigcap_{1 \leqslant j \leqslant m} S_j = \varnothing$. We show how to produce a digraph D such that $\mathcal{S}$ has a hitting set of size at most k if and only if D has a quasi-kernel of size at most k. The digraph D is defined as follows:

$$V(D) \coloneqq \quad \mathcal{S} \cup \mathcal{X} \cup B \cup C,$$

where

$$
\begin{aligned}
B &:= \{b_i\colon 1 \leqslant i \leqslant k+1\}\\
C &:= \{c_i\colon 1 \leqslant i \leqslant k+1\}\\
A(D) &:= \{(S_j, x_i)\colon x_i \in S_j, S_j \in \mathcal{S}, x_i \in \mathcal{X}\} \cup \{(x_i, S_j)\colon x_i \notin S_j, S_j \in \mathcal{S}, x_i \in \mathcal{X}\}\\
&\cup \{(S_j, b_i)\colon S_j \in \mathcal{S}, b_i \in B\} \cup \{(S_j, c_i)\colon S_j \in \mathcal{S}, c_i \in C\}\\
&\cup \{(x_i, b_j)\colon x_i \in \mathcal{X}, b_j \in B\} \cup \{(c_i, x_j)\colon c_i \in C, x_j \in \mathcal{X})\}\\
&\cup \{(b_i, c_j)\colon b_i \in B, c_j \in C\}\,.
\end{aligned}
$$

Clearly, D is an orientation of a complete 4-partite graph.

Suppose that there exists a subfamily $\mathcal{X}' \subseteq \mathcal{X}$ of size at most k such that every set in $\mathcal{S}$ intersects $\mathcal{X}$'. It is clear that $\mathcal{X}'$ is a quasi-kernel of D of size k, since every vertex in $\mathcal{X}$ has an outneighbour in $\mathcal{S}$ and every vertex in $B \cup C$ is at distance at most two to every vertex in $\mathcal{X}$.

Conversely, suppose that there exists a quasi-kernel Q of D of size at most k. Necessarily, Q is included in $\mathcal{S}, \mathcal{X}, B$, or C. Observe that $Q \not\subseteq \mathcal{S}$ since B is a distance three to $\mathcal{S}$. Suppose $Q \subseteq B$, then $Q = B$ since every vertex $b_i \in B$ is at distance three to $B \setminus \{b_i\}$ but this contradicts the size of Q. The same contradiction occurs if we suppose $Q \subseteq C$. Then, $Q \subseteq \mathcal{X}$ and Q forms a hitting set of size k. □

7.5 Exponential algorithm

This section provides two polynomial algorithms for digraphs with bounded treewidth, one deciding the size of a minimum quasi-kernel and the other deciding if there are two disjoint quasi-kernels. Courcelle's theorem [29] ensures that those problems are polynomial-time solvable for orientations of graphs with bounded treewidth, but it is known that the complxity has a huge factor depending on the treewidth (as far as we understand from the paper of Kneis and Langerabout [58] the complexity would be of order $2^{2^{2^{\mathbf{tw}^2}}}$). This motivated the search for a direct approach of the polynomial result. In this section we provide polynomial algorithms with reasonable constant factors.

7.5.1 Preliminaries about treewidth

Let G be a graph. A *tree decomposition* of G is a pair $(T, \mathcal{X} = \{X_t\}_{t \in V(T)})$ where T is a tree and $\mathcal{X}$ is a collection of subsets of $V(G)$ such that

- $\bigcup_{t \in V(T)} X_t = V(G)$,
- for every $uv \in E(G)$, there is a t such that $\{u, v\} \in X_t$, and
- for every $v \in V(G)$, $T[\{t\colon v \in X_t\}]$ is connected.

The vertices of T are the *nodes* and the sets in $\mathcal{X}$ are the *bags* of the tree decomposition $(T, \mathcal{X})$. The *width* of $(T, \mathcal{X})$ is equal to $\max\{|X_t| - 1\colon t \in V(T)\}$ and the *treewidth* of G is the minimum width over all tree decompositions of G. We denote the treewidth of a graph G by $\mathbf{tw}(G)$ and by extension $\mathbf{tw}(D)$ the treewidth of the underlying graph of a digraph D.

Let $\mathcal{D} = (T, \mathcal{X})$ be a tree decomposition of G, r be a node of T, and $\mathcal{G} = \{G_t\colon t \in V(T)\}$ be a collection of subgraphs of G, indexed by the nodes of T. We say that the triple $(\mathcal{D}, r, \mathcal{G})$ is a *nice tree decomposition* of G if the following conditions hold:

- T is rooted at node a,

- $X_a = \varnothing$ and $G_a = G$,
- each node of $\mathcal{D}$ has at most two children in T,
- for each leaf $t \in V(T)$, $X_t = \varnothing$, and G_t is such that $V(G_t) = \varnothing$ and $E(G_t) = \varnothing$. Such t is a *leaf node*,
- if $t \in V(T)$ has exactly one child t', then either
  - $X_t = X_{t'} \cup \{v_{\text{insert}}\}$ for some $v_{\text{insert}} \notin X_{t'}$ and G_t is such that $V(G_t) = V(G_{t'}) \cup \{v_{\text{insert}}\}$ and $E(G_t) = E(G_{t'})$. The vertex t is an *introduce vertex* node and the node v_{insert} is the *insertion vertex* of X_t,
  - $X_t = X_{t'}$ and G_t is such that $V(G_t) = V(G_{t'})$ and $E(G_t) = E(G_{t'}) \cup \{e_{\text{insert}}\}$ where e_{insert} is an edge of G with endpoints in X_t. The node t is an *introduce edge* node and the edge e_{insert} is the *insertion edge* of X_t, or
  - $X_t = X_{t'} \setminus \{v_{\text{forget}}\}$ for some $v_{\text{forget}} \in X_{t'}$ and $G_t = G_{t'}$. The node t is a *forget vertex* node and v_{forget} is the *forget vertex* of X_t.
- for every $e \in E(G)$, there exists a unique node t such that e is the insertion edge of X_t, and
- if $t \in V(T)$ has exactly two children t' and t'', then $X_t = X_{t'} = X_{t''}$, $V(G_t) = V(G_{t'}) \cup V(G_{t''})$, $E(G_t) = E(G_{t'}) \cup E(G_{t''})$ and $E(G_{t'}) \cap E(G_{t''}) = \varnothing$. The node t is a *join* node.

Claim 7.5.1. *For every graph G with bounded treewidth, it is possible to find a nice tree decomposition with $\mathcal{O}(\mathbf{tw}(G)|V(G)|)$ vertices of width $\mathcal{O}(\mathbf{tw}(G))$ in polynomial time.*

Indeed, as discussed by Althaus and Ziegler [5], given a tree decomposition $(T, \mathcal{X})$ of width w, we can compute in time $\mathcal{O}(w^2(|V(T)| + |V(G)|))$ a nice tree decomposition of G of width w with at most $\mathcal{O}(w|V(G)|)$ nodes. Moreover, by Bodlaender et al. [16] we can find in time $2^{\mathcal{O}(\mathbf{tw})}n$ a tree decomposition of width $\mathcal{O}(\mathbf{tw})$ of any graph G.

In this section, we consider problems on digraphs and we use tree decompositions of their underlying graph to find polynomial-time algorithms.

For each $t \in V(T)$, we denote by V_t the set $V(G_t)$.

7.5.2 Results

Theorem 7.5.2. *If a nice tree decomposition $(T, \mathcal{X})$ of a digraph of width w is given,* Quasi-Kernel *can be solved in time $\mathcal{O}(25^w|V(T)|)$.*

Theorem 7.5.3. *If a nice tree decomposition $(T, \mathcal{X})$ of a digraph of width w is given, deciding if it has two disjoint quasi-kernels can be solved in time $\mathcal{O}(25^{2w}|V(T)|)$.*

Actually the proof of Theorem 7.5.3 can very easily be adapted to prove that deciding if a digraph has k disjoint quasi-kernels can be solved in $\mathcal{O}(25^{kw}|V(T)|)$.

7.5.3 Proof of Theorem 7.5.2

We define, for each $t \in V(T)$, the set

$$I_t = \Big\{(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N})\colon S_0 \cup S_{1V} \cup S_{1N} \cup S_{2V} \cup S_{2N} = X_t, \\ \forall \mathfrak{a} \neq \mathfrak{b} \in \{0, 1V, 1N, 2V, 2N\}, S_{\mathfrak{a}} \cap S_{\mathfrak{b}} = \varnothing\Big\}$$

and a function $r_t \colon I_t \to \mathbb{Z}_+$. For each $t \in V(T)$, an element of I_t represents a partition of X_t into the vertices in the potential quasi-kernel (S_0), vertices that we know are inneighbors of the potential quasi-kernel (S_{1V}), vertices that are intended to be inneighbors of the futur quasi-kernel but are not yet (S_{1NV}), vertices that we know are inneighbors of -potential- inneighbors of the potential quasi-kernel and the remaining vertices, which are intended to be at distance two to the futur quasi-kernel (S_{2NV}).

For each $t \in V(T)$, we define r_t recursively from the r'_t for each children t' of t. The definition depends on the type of node of t.

- Leaf: $r_t \coloneqq 0$.
- Introduce vertex: If v is the insertion vertex of X_t and t' is the child of t, then for each $(S_0, S_{1V}, S_{1NV}, S_{2V}, S_{2NV}) \in I_t$,

$$r_t(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N}) \coloneqq \begin{cases} r'_t(S_0 \setminus \{v\}, S_{1V}, S_{1N}, S_{2V}, S_{2N}) + 1 & \text{if } v \in S_0. \\ r'_t(S_0, S_{1V}, S_{1N} \setminus \{v\}, S_{2V}, S_{2N}) & \text{if } v \in S_{1N}. \\ r'_t(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N} \setminus \{v\}) & \text{if } v \in S_{2N}. \\ +\infty & \text{otherwise.} \end{cases}$$

- Forget vertex: If v is the forget vertex of X_t and t' is the child of t, then for each $(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N}) \in I_t$,

$$\begin{aligned} r_t(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N}) \coloneqq \min\{&(r_{t'}(S_0 \cup \{v\}, S_{1V}, S_{1N}, S_{2V}, S_{2N}), \\ &(r_{t'}(S_0, S_{1V} \cup \{v\}, S_{1N}, S_{2V}, S_{2N}), \\ &(r_{t'}(S_0, S_{1V}, S_{1N}, S_{2V} \cup \{v\}, S_{2N})\}. \end{aligned}$$

- Introduce edge: If $\{u, v\}$ is the introduced edge of X_t and that (u, v) in an arc in the digraph and t' is the child of t, then for each $(S_0, S_{1V}, S_{1NV}, S_{2V}, S_{2NV}) \in I_t$,

$$r_t(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N}) \coloneqq \begin{cases} \min\Big(r_{t'}(S_0, S_{1V} \setminus \{u\}, S_{1N} \cup \{u\}, S_{2V}, S_{2N}), \\ \quad r_{t'}(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N})\Big) & \text{if } u \in S_{1V} \\ & \text{ and } v \in S_0. \\ \min\Big(r_{t'}(S_0, S_{1V}, S_{1N}, S_{2V} \setminus \{u\}, S_{2N} \cup \{u\}), \\ \quad r_{t'}(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N})\Big) & \text{if } u \in S_{2V} \\ & \text{ and } v \in S_{1V} \cup S_{1N}. \\ +\infty & \text{if } u, v \in S_0. \\ r_{t'}(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N}) & \text{otherwise.} \end{cases}$$

- Join: If t' and t'' are the children of t, then for each $(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N}) \in I_t$,

$$r_t(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N}) \coloneqq \min_{\substack{S_{2N} = S'_{2N} \cap S''_{2N} \\ S_{1N} = S'_{1N} \cap S''_{1N} \\ S_0 = S'_0 = S''_0 \\ S_{2V} = S'_{2V} \cup S''_{2V} \\ S_{1V} = S'_{1V} \cup S''_{1V} \\ S'_{2V} \cup S'_{2N} = S''_{2V} \cup S''_{2N} \\ S'_{1V} \cup S'_{1N} = S''_{1V} \cup S''_{1N}}} r_{t'}(S'_0, S'_{1V}, S'_{1N}, S'_{2V}, S'_{2N}) + r_{t''}(S''_0, S''_{1V}, S''_{1N}, S''_{2V}, S''_{2N}) - |S_0|.$$

Lemma 7.5.4. *For each $t \in V(T)$ and $(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N}) \in I_t$, the quantity $r_t(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N})$ is the minimum k such that there exists a set $\hat{S} \subseteq V(G_t)$ that satisfies simultaneously:*

(i) $\hat{S}$ *is independent,*

(ii) $|\hat{S}| \leqslant k$*,*

(iii) $\hat{S} \cap X_t = S_0$*,*

(iv) *for every* $v \in V(G_t) \setminus X_t$*,* $d(x, \hat{S}) \leqslant 2$ *or* $d(v, S_{1V} \cup S_{1N}) \leqslant 1$*,*

(v) $S_{1V} \subseteq N^-_{G_t}(\hat{S})$*, and*

(vi) $S_{2V} \subseteq N^-_{G_t}(S_{1N}) \cup N^{--}_{G_t}(\hat{S})$*.*

Such a set $\hat{S}$ *is called the witness of* $(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N})$*.*

Proof of lemma 7.5.4. Consider $t \in V(T)$ and $(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N}) \in I_t$. We prove the result by induction by assuming the result true for every child of t. Let us distinguish different cases depending on the type of t.

- ***Leaf.***
  $V(G_t) = \varnothing$ and $I_t = \{(\varnothing, \varnothing, \varnothing, \varnothing, \varnothing)\}$ so the witness of $(\varnothing, \varnothing, \varnothing, \varnothing, \varnothing)$ is $\varnothing$, of size zero.

- ***Introduce vertex.***
  Let v be the insertion vertex of X_t. Consider $(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N}) \in I_t$. Let us prove that $r_t((S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N})) = k$ if and only if there exists $\hat{S}$ a witness of $(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N})$ of size k.

  First, suppose $r_t((S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N})) = k$. By definition, if $v \in S_0$ then $r'_t(S_0 \setminus \{v\}, S_{1V}, S_{1N}, S_{2V}, S_{2N}) = k-1$. By induction, there exists a witness $\hat{S}'$ of $(S_0 \setminus \{v\}, S_{1V}, S_{1N}, S_{2V}, S_{2N})$ of size $k-1$. Just adding the vertex v to $\hat{S}'$ provides a witness of $(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N})$ of size k. Indeed, since v is the insertion vertex of X_t, it is independent from ever other vertex in $V(G_t)$, so condition (i) is fulfilled, conditions (ii) and (iii) are straightforward and every other condition is true exactly by induction hypothesis.

  One can check that similar arguments works if $v \in S_{1N}$ or $v \in S_{2N}$.

  Conversely, let $(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N}) \in I_t$ and $\hat{S}$ be a witness of size k. If $v \in \hat{S}$, consider $\hat{S} \setminus \{v\}$. This forms a witness of $(S_0 \setminus \{v\}, S_{1V}, S_{1N}, S_{2V}, S_{2N})$ of size $k-1$. Indeed, conditions (i), (ii), and (iii) are clear and since v is the insertion vertex of X_t is it independent from every other vertex in $V(G_t)$, conditions (iv), (v), and (vi) are still fulfilled.

  By induction, $r_{t'}(S_0 \setminus \{v\}, S_{1V}, S_{1N}, S_{2V}, S_{2N}) = k-1$ and then, by definition of r_t for t an introduce vertex, $r_t(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N}) = k$.

  One can check that similar arguments work if $v \in S_{1N}$ or $v \in S_{2N}$.

- ***Forget vertex.***
  Let v be the forget vertex of X_t and t' be the child of t. Consider $(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N}) \in I_t$.

  First, suppose $r_t((S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N})) = k$. By definition of r_t with t a forget vertex, one quantity among $r_{t'}(S_0 \cup \{v\}, S_{1V}, S_{1N}, S_{2V}, S_{2N})$, $r_{t'}(S_0, S_{1V}, S_{1N} \cup \{v\}, S_{2V}, S_{2N})$, and $r_{t'}(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N} \cup \{v\})$ equals k. Consider the case $r_{t'}(S_0 \cup \{v\}, S_{1V}, S_{1N}, S_{2V}, S_{2N}) = k$. By induction, there exists a witness $\hat{S}$ of $(S_0 \cup \{v\}, S_{1V}, S_{1N}, S_{2V}, S_{2N})$ of size k. Then $\hat{S}$ forms a witness of $(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N})$ of size k. Indeed, conditions (i),(ii), (v), and (vi) are clear, condition (iii) is true since v is not in X_t, and condition (iv) is true because $x \in \hat{S}$ so $d(x, \hat{S}) = 0$.

  Similar arguments work for $r_{t'}(S_0, S_{1V}, S_{1N} \cup \{v\}, S_{2V}, S_{2N}) = k$ and for $r_{t'}(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N} \cup \{v\}) = k$.

Conversely, let $(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N}) \in I_t$ and $\hat{S}$ be a witness of size k.

If $v \in \hat{S}$, then $\hat{S}$ forms a witness of $(S_0 \cup \{v\}, S_{1V}, S_{1N}, S_{2V}, S_{2N})$ of size k. Indeed, condition (iii) is clear and every other condition is fulfilled because $\hat{S}$ is a witness of $(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N})$. This implies $r_{t'}(S_0 \cup \{v\}, S_{1V}, S_{1N}, S_{2V}, S_{2N}) = k$ by induction, and $r_t(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N}) \leqslant k$ by definition of r_t for a forget vertex.

Similar arguments show that if $v \in N^-(\hat{S})$ or $v \in N^{--}(\hat{S}) \cup N^-(S_{1V} \cup S_{1N})$, then $r_t(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N}) \leqslant k$.

- ***Introduce edge.***

Let $\{u, v\}$ be the introduced edge of X_t such that (u, v) is an arc of the digraph and t' the child of t. Consider $(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N}) \in I_t$.

Suppose $r_t(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N}) = k$ and consider the case where $u \in S_{1N}$ and $v \in S_0$. By definition, $r_{t'}(S_0, S_{1V} \cup \{u\}, S_{1N} \setminus \{u\}, S_{2V}, S_{2N}) = k$ and by induction, there exists a witness $\hat{S}'$ of $(S_0, S_{1V} \cup \{u\}, S_{1N} \setminus \{u\}, S_{2V}, S_{2N})$ of size k. The same set forms a witness of $(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N})$ of size k. Indeed, condition (v) is weaker for a witness of $(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N})$ than for one of $(S_0, S_{1V} \cup \{u\}, S_{1N} \setminus \{u\}, S_{2V}, S_{2N})$ and every other condition remain the same.

Similar arguments work for every other cases.

Conversely, let $(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N}) \in I_t$ and $\hat{S}$ be a witness of size k of $(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N})$. If $u \in S_{1N}$ and $v \in S_0$, then $\hat{S}$ forms a witness of $(S_0, S_{1V} \cup \{u\}, S_{1N} \setminus \{u\}, S_{2V}, S_{2N})$ of size k. Indeed, condition (v) is fulfilled because $u \in N^-(v)$ and $v \in \hat{S}$ and every other condition is clear. By induction, $r_{t'}(S_0, S_{1V} \cup \{u\}, S_{1N} \setminus \{u\}, S_{2V}, S_{2N}) = k$ and then, by definition, $r_t(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N}) = k$.

One can check that similar arguments work in every other case.

- ***Join.***

Suppose $r_t(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N}) = k$. Consider the sets $S'_0, S'_{1V}, S'_{1N}, S'_{2V}, S'_{2N}, S''_0, S''_{1V}, S''_{1N}, S''_{2V}, S''_{2N}$ reaching the minimum in the definition of $r_t(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N})$ and denote $k' = r_{t'}(S'_0, S'_{1V}, S'_{1N}, S'_{2V}, S'_{2N})$, $k'' = r_{t''}(S''_0, S''_{1V}, S''_{1N}, S''_{2V}, S''_{2N})$, $l = |S_0|$. By definition we have $k' + k'' - l = k$.

By induction, there exists $\hat{S}'$ witness of $(S'_0, S'_{1V}, S'_{1N}, S'_{2V}, S'_{2N})$ of size k' and $\hat{S}''$ witness of $(S''_0, S''_{1V}, S''_{1N}, S''_{2V}, S''_{2N})$ of size k''. $\hat{S}' \cup \hat{S}''$ forms a witness of $(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N})$ of size $k' + k'' - l$. Indeed, by definition of r_t for a join node, $S_0 = S'_0 = S''_0$ which ensures condition (i) and (iii). Condition (ii) is clear, and the other conditions are true because of conditions in the min of the definition.

Conversely, let $(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N}) \in I_t$ and $\hat{S}$ be a witness of size k. Consider $\hat{S}' = \hat{S} \cap V(G_{t'})$ and $\hat{S}'' = \hat{S} \cap V(G_{t''})$. Denote $k' = |\hat{S}'|$, $k'' = |\hat{S}''|$ and $l = |S_0|$. Notice that by construction $k' + k'' - l = |\hat{S}| = k$. Consider then $S'_0 = S''_0 = S_0$, $S'_{1V} = X_{t'} \cap N^-(\hat{S}')$ and $S''_{1V} = X_{t''} \cap N^-(\hat{S}'')$. Notice that $S_{1V} = S'_{1V} \cup S''_{1V}$. Denote $S'_{1N} = (S_{1V} \cup S_{1N}) \setminus S'_{1V}$, $S''_{1N} = (S_{1V} \cup S_{1N}) \setminus S''_{1V}$, $S'_{2V} = N^{--}(\hat{S}') \cup N^-(S'_{1V} \cup S'_{1N})$, $S''_{2V} = N^{--}(\hat{S}'') \cup N^-(S_{1V})''$. Finally, consider $S'_{2N} = (S_{2N} \cup S_{2V}) \setminus S'_{2V}$ and $S''_{2N} = (S_{2N} \cup S_{2V}) \setminus S''_{2V}$. Note that the previous sets respect the conditions of the minimum in the definition of $r_t(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N})$.

Then $\hat{S}'$ is a witness of $(S'_0, S'_{1V}, S'_{1N}, S'_{2V}, S'_{2N})$ and $\hat{S}''$ is a witness of $(S''_0, S''_{1V}, S''_{1N}, S''_{2V}, S''_{2N})$. Indeed, condition (i) and (iii) come trivially from the fact that $S_0 = S'_0 = S''_0$, condition (ii) is true by construction, and the other conditions are true by the condition in the min in the definition.

By induction we have $r_{t'}(S'_0, S'_{1V}, S'_{1N}, S'_{2V}, S'_{2N}) = k'$ and $r_{t''}(S''_0, S''_{1V}, S''_{1N}, S''_{2V}, S''_{2N}) = k''$. By definition, this leads to $r_t(S_0, S_{1V}, S_{1N}, S'_{2V}, S'_{2N}) = k' + k'' - l = k$. □

For a nice tree decomposition rooted at node a of a digraph D, we have then $r_a(\varnothing, \varnothing, \varnothing, \varnothing, \varnothing)$ the size of a smallest quasi-kernel in D. Indeed, a witness forms a quasi-kernel since conditions (iii), (v), and (vi) turn to be empty and conditions (i) and (iv) are the definition of a quasi-kernel.

Proof of Theorem 7.5.2. Using Lemma 7.5.4, we just need to analyze the running time of the algorithm computing the functions r_t recursively.

For each $t \in V(T)$, , we have $|I_t| \leqslant 5^{|X_t|}$ since $(S_0, S_{1V}, S_{1NV}, S_{2V}, S_{2NV})$ is a partition of X_t. If t is a leaf, then r_t can be computed in $\mathcal{O}(1)$. If t is an introduce vertex, a forget vertex, or an introduce edge, and t' is the child of t, then r_t can be computed in time $\mathcal{O}(|I'_t|)$. If t is a join node, and t' and $t''$are the two children of t, then r_t can be computed in time $\mathcal{O}(|I'_t||I''_t|)$. The algorithm runs in $25^{\mathbf{tw}(D)}|V(T)|$. □

7.5.4 Proof of Theorem 7.5.3

We define, for each $t \in V(T)$, the set

$$I_t = \Big\{(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N})\colon S_0 \cup S_{1V} \cup S_{1N} \cup S_{2V} \cup S_{2N} = X_t,\\ \forall \mathfrak{a} \neq \mathfrak{b} \in \{0, 1V, 1N, 2V, 2N\}, S_{\mathfrak{a}} \cap S_{\mathfrak{b}} = \varnothing\Big\}$$

and a function $r_t \colon I_t^2 \to \{0,1\}$.

For each $t \in V(T)$, an element of I_t represents a partition of X_t into the vertices in the potential quasi-kernel (S_0), vertices that we know are inneighbors of the potential quasi-kernel (S_{1V}), vertices that are intended to be inneighbors of the futur quasi-kernel but are not yet (S_{1NV}), vertices that we know are inneighbors of inneighbors or potential inneighbors of the futur quasi-kernel and the remaining vertices, which are intended to be at distance two to the futur quasi-kernel (S_{2NV}).

For each $t \in V(T)$, we define r_t recursively from the r'_t for each children t' of t. The definition depends on the type of node of t.

- Leaf : we can assume $I_t = \varnothing$, and then $r_t \equiv 1$.
- Introduce vertex : If v is the insertion vertex of X_t and t' is the child of t, then for each

  $U = (S_0, S_{1V}, S_{1NV}, S_{2V}, S_{2NV}), (T_0, T_{1V}, T_{1NV}, T_{2V}, T_{2NV})) \in I_t^2$,

$$r_t(U) \coloneqq \begin{cases} 0 & \text{if } v \in S_0 \cap T_0 \\ 0 & \text{if } v \in S_{1V} \cup S_{1N} \cup S_{2V} \cup S_{2N} \\ r_{t'}((S_0 \setminus \{v\}, S_{1V}, S_{1N} \setminus \{v\}, S_{2V}, S_{2N} \setminus \{v\}), & \\ \quad (T_0 \setminus \{v\}, T_{1V}, T_{1N} \setminus \{v\}, T_{2V}, T_{2N} \setminus \{v\})) & \text{else.} \end{cases}$$

- Forget vertex : If v is the forget vertex of X_t and t' is the child of t, then for each $U = (S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N}), (T_0, T_{1V}, T_{1N}, T_{2V}, T_{2N})) \in I_t^2$:

$$r_t(U) \coloneqq \max \Big\{ r_{t'}((S'_0, S'_{1V}, S_{1N}, S'_{2V}, S_{2N}), (T'_0, T'_{1V}, T_{1N}, T'_{2V}, T_{2N}))\colon\\ S'_0, S'_{1V}, S'_{2V} \text{ disjoints, } S'_0 \cup S'_{1V} \cup S'_{2V} = S_0 \cup S_{1V} \cup S_{2V} \cup \{v\},\\ T'_0, T'_{1V}, T'_{2V} \text{ disjoints and } T'_0 \cup T'_{1V} \cup T'_{2V} = T_0 \cup T_{1V} \cup T_{2V} \cup \{v\} \Big\}.$$

- Introduce arc: If (u, v) is the introduced arc of X_t and t' is the child of t, then for each

  $((S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N}), (T_0, T_{1V}, T_{1N}, T_{2V}, T_{2N})) \in I_t^2$,

  if $u, v \in S_0$ or $u, v \in T_0$, $r_t((S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N}), (T_0, T_{1V}, T_{1N}, T_{2V}, T_{2N})) = 0$.

Else, $r_t((S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N}), (T_0, T_{1V}, T_{1N}, T_{2V}, T_{2N}))$ is the higher value between 0,

$r_{t'}((S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N}), (T_0, T_{1V}, T_{1N}, T_{2V}, T_{2N}))$ and the following :

- if $(u, v) \in ((S_{1V} \cup S_{1N}) \times S_0) \cap ((T_{2V} \cup T_{2N}) \times (T_{1V} \cup T_{1N}))$: $r_{t'}((S_0, S_{1V} \cup \{u\}, S_{1N} \setminus \{u\}, S_{2V}, S_{2N}), (T_0, T_{1V}, T_{1N}, T_{2V} \cup \{u\}, T_{2N} \setminus \{u\}))$.
- if $(u, v) \in ((S_{2V} \cup S_{2N}) \times (S_{1V} \cup S_{1N})) \cap ((T_{1V} \cup T_{1N}) \times T_0)$: $r_{t'}((S_0, S_{1V}, S_{1N}, S_{2V} \cup \{u\}, S_{2N} \setminus \{u\}), (T_0, T_{1V} \cup \{u\}, T_{1N} \setminus \{u\}, T_{2V}, T_{2N}))\big)$.
- if $(u, v) \in ((S_{2V} \cup S_{2N}) \times (S_{1N} \cup S_{1V})) \cap ((T_{2V} \cup T_{2N}) \times (T_{1V} \cup T_{1N}))$: $r_{t'}((S_0, S_{1V}, S_{1N}, S_{2V} \cup \{u\}, S_{2N} \setminus \{u\}), (T_0, T_{1V}, T_{1N}, T_{2V} \cup \{u\}, T_{2N} \setminus \{u\}))$.
- if $(u, v) \in (T_{1V} \cup T_{1N}) \times T_0$: $r_{t'}((S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N}), (T_0, T_{1V} \cup \{u\}, T_{1N} \setminus \{u\}, T_{2V}, T_{2N}))$.
- if $(u, v) \in (T_{2V} \cup T_{2N}) \times (T_{1V} \cup T_{1N})$: $r_{t'}((S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N}), (T_0, T_{1V}, T_{1N}, T_{2V} \cup \{u\}, T_{2N} \setminus \{u\}))$.
- if $(u, v) \in (S_{1V} \cup S_{1N}) \times S_0)$: $r_{t'}((S_0, S_{1V} \cup \{u\}, S_{1N} \setminus \{u\}, S_{2V}, S_{2N}), (T_0, T_{1V}, T_{1N}, T_{2V}, T_{2N}))$
- if $(u, v) \in ((S_{2V} \cup S_{2N}) \times (S_{1N} \cup S_{1V}))$: $r_{t'}((S_0, S_{1V}, S_{1N}, S_{2V} \cup \{u\}, S_{2N} \setminus \{u\}), (T_0, T_{1V}, T_{1N}, T_{2V}, T_{2N}))$.

- Join: If t' and t'' are the children of t then for each $((S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N}), (T_0, T_{1V}, T_{1N}, T_{2V}, T_{2N})) \in I_t^2$,

  $r_t((S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N}), (T_0, T_{1V}, T_{1N}, T_{2V}, T_{2N})) =$

$$\max_{\substack{S_{2N}=S'_{2N}\cap S''_{2N}, T_{2N}=T'_{2N}\cap T''_{2N} \\ S_{1N}=S'_{1N}\cap S''_{1N}, T_{1N}=T'_{1N}\cap T''_{1N} \\ S_0=S'_0=S''_0, T_0=T'_0=T''_0 \\ S_{2V}=S'_{2V}\cup S''_{2V}, T_{2V}=T'_{2V}\cup T''_{2V} \\ S_{1V}=S'_{1V}\cup S''_{1V}, T_{1V}=T'_{1V}\cup T''_{1V} \\ S'_{2V}\cup S'_{2N}=S''_{2V}\cup S''_{2N}, T'_{2V}\cup T'_{2N}=T''_{2V}\cup T''_{2N} \\ S'_{1V}\cup S'_{1N}=S''_{1V}\cup S''_{1N}, T'_{1V}\cup T'_{1N}=T''_{1V}\cup T''_{1N}}} r_{t'}((S'_0, S'_{1V}, S'_{1N}, S'_{2V}, S'_{2N}), (T'_0, T'_{1V}, T'_{1N}, T'_{2V}, T'_{2N})).$$

Lemma 7.5.5. *For each $t \in V(T)$, $(S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N}) \in I_t$, and $(T_0, T_{1V}, T_{1N}, T_{2V}, T_{2N}) \in I_t$ the quantity $r_t((S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N}), (T_0, T_{1V}, T_{1N}, T_{2V}, T_{2N}))$ equals 1 if and only if there exist $\hat{S}, \hat{T} \subseteq V(G_t)$ that satisfies simultaneously:*

1. *$\hat{S} \cap \hat{T} = \varnothing$,*
2. *$\hat{S} \cap X_t = S_0$,*
3. *for every $v \in G_t \setminus X_t$, $d(v, \hat{S}) \leqslant 2$ or $d(v, S_{1V} \cup S_{1N}) \leqslant 1$,*
4. *$S_{1V} \subseteq N^-_{G_t}(\hat{S})$,*
5. *$S_{2V} \subseteq N^-_{G_t}(S_{1N}) \cup N^{--}_{G_t}(\hat{S})$.*
6. *$\hat{T} \cap X_t = T_0$,*
7. *for every $v \in G_t \setminus X_t$, $d(v, \hat{T}) \leqslant 2$ or $d(v, T_{1V} \cup T_{1N}) \leqslant 1$.*
8. *$T_{1V} \subseteq N^-_{G_t}(\hat{T})$, and*
9. *$T_{2V} \subseteq N^-_{G_t}(T_{1N}) \cup N^{--}_{G_t}(\hat{T})$.*

Such a pair $(\hat{S}, \hat{T})$ is called the witness of $((S_0, S_{1V}, S_{1N}, S_{2V}, S_{2N}), (T_0, T_{1V}, T_{1N}, T_{2V}, T_{2N}))$.

For a nice tree decomposition rooted at node a of a digraph D, we have then $r_a((\varnothing,\varnothing,\varnothing,\varnothing,\varnothing),(\varnothing,\varnothing,\varnothing,\varnothing,\varnothing))$ the boolean answer to the question: does D admits two disjoint quasi-kernels.

Proof of Theorem 7.5.3. Using Lemma 7.5.5, we just need to analyze the running time of the algorithm computing the functions r_t recursively. For each $t \in V(T)$, we have $|I_t| \leqslant (25)^{|X_t|}$ since $(S_0, S_{1V}, S_{1NV}, S_{2V}, S_{2NV})$ and $(T_0, T_{1V}, T_{1NV}, T_{2V}, T_{2NV})$ are partitions of X_t. If t is a leaf, then r_t can be computed in $\mathcal{O}(1)$. If t is an introduce vertex, a forget vertex, or an introduce edge, and t' is the child of t, then r_t can be computed in time $\mathcal{O}(|I'_t|)$. If t is a join node, and t' and $t''$are the two children of t, then r_t can be computed in time $\mathcal{O}(|I'_t||I''_t|)$. The algorithm runs in $25^{2\mathbf{tw}(D)}|V(T)|$. □

Appendix A

Tables - Kernels

Problem	Complexity
Deciding of the existence of a kernel	NP-complete [27]
Deciding if an orientation of a perfect graph is clique-acyclic	coNP-complete [6]
Finding a kernel in a clique-acylic orientation of a perfect graph with cliques of bounded size	PPAD [56]
Finding a kernel in simple clique-acyclic orientations of perfect claw-free graphs	Polynomial [52]
Finding a kernel in clique-acylic orientations of DE-graphs	Polynomial [32, 52]
Finding a kernel in M-clique-acylic orientations of comparability graphs	Polynomial [1]
Finding a kernel in M-clique-acylic orientations of tree-cographs	Polynomial (Corollary 2.3.2)
Finding a kernel in simple clique-acylic orientations of distance-hereditary graphs	Polynomial (Proposition 2.3.3)
Finding a kernel in a digraph with the conditions depicted in Figure 4.1	Polynomial (Theorem 4.1.1)
Deciding if a digraph has at least two crossing consecutive chords in each odd directed cycle	Polynomial (Proposition 3.3.1)
Finding a kernel in clique-acylic orientations of chordal graphs	Polynomial [52]
Finding a kernel in a clique-acylic orientation of a perfect graph	Open
Deciding if a digraph has at least two chords with consecutive heads in each odd directed cycle	Open

Conditions ensuring the existence of a kernel	Reference
Every directed cycle has at least one hence-and-forth pairs of arcs	[36]
Each odd directed cycle has at least two hence-and-forth pair of arcs	[36]
Each odd directed cycle has two chords with consecutive heads	[44]
Each odd directed cycle has two chords with consecutive heads or two non-crossing chords of odd length in the same direction or two crossing chords, one short and the other of odd length	Theorem 3.2.1
Clique-acyclic orientation of a perfect graph	[18]
Digraph respecting the conditions depicted in Figure 4.1	Theorem 4.1.1
Digraph colored in two colors with no monochromatic directed cycle avoiding structures depicted in Figure 4.2	Theorem 4.1.2

Appendix B

Tables - Quasi-kernels

Problem	Complexity
Deciding the existence of a quasi-kernel containing a specified vertex	NP-complete [30]
Finding the smallest quasi-kernel in acyclic orientations of cubic graphs	NP-complete (Theorem 7.2.1)
Deciding the existence of two disjoint quasi-kernels	NP-complete (Theorem 7.1.1)
Deciding the existence of three disjoint quasi-kernels in sink-free bounded degree planar digraphs	NP-complete (Proposition 7.1.4)
Finding the smallest quasi-kernel in orientations of split graphs	W[2]-complete (Proposition 7.3.1)
Finding the smallest quasi-kernel in orientations of complete 4-partite graphs	W[2]-complete (Theorem 7.4.1)
Finding the smallest quasi-kernel of acyclic orientations of bipartite graphs	W[2]-complete (Theorem 7.2.2)
Finding the minimum quasi-kernel for digraphs with fixed maximumm indegrees	APX-complete (Proposition 7.2.6)
Finding the smallest quasi-kernel in digraphs with bounded treewidth	Polynomial
Deciding the existence of k disjoint quasi-kernels for any k in digraphs with bounded treewidth	Polynomial

Conditions ensuring the existence of a small quasi-kernel	Reference
Semicomplete multipartite digraph	[49]
Quasi-transitive digraph	[49]
Locally semicomplete digraph	[49]
Digraph whose vertex vertex can be partitioned into two kernel-perfect digraphs	[59]
Digraph admitting a lernel	[73]
Digraph containing a kernel in the second outneighbhorhood of a quasi-kernel	[39]
Orientation of unicyclic graph	[39]
Anti-claw-free digraph	[3]

Bibliography

[1] Moncef Abbas and Youcef Saoula, *Polynomial algorithms for kernels in comparability, permutation and P_4-free graphs*, 4OR **3** (2005), no. 3, 217–225. 22, 33, 42, 79

[2] Ron Aharoni and Ron Holzman, *Fractional kernels in digraphs*, Journal of Combinatorial Theory, Series B **73** (1998), 1–6. 29

[3] Jiangdong Ai, Stefanie Gerke, Gregory Gutin, Anders Yeo, and Yacong Zhou, *Results on the small quasi-kernel conjecture*, Discrete Mathematics **346** (2023), no. 7, 113435. 49, 52, 80

[4] Paola Alimonti and Viggo Kann, *Hardness of approximating problems on cubic graphs*, Italian conference on algorithms and complexity, Springer, 1997, pp. 288–298. 67

[5] Ernst Althaus and Sarah Ziegler, *Optimal tree decompositions revisited: A simpler linear-time fpt algorithm*, Graphs and Combinatorial Optimization: from Theory to Applications: CTW2020 Proceedings (2021), 67–78. 72

[6] S. D. Andres and W. Hochstättler, *Perfect digraphs*, J. Graph Theory **79** (2014), 21–29. 79

[7] Stephan Dominique Andres and Winfried Hochstättler, *Perfect digraphs*, Journal of Graph Theory **79** (2015), no. 1, 21–29. 21, 32

[8] Jørgen Bang-Jensen and Gregory Gutin, *Digraphs: theory, algorithms and applications*, Springer Science & Business Media, 2008. 11

[9] Claude Berge, *Farbung von Graphen, deren samtliche bzw. deren ungerade Kreise starr sind*, Wissenschaftliche Zeitschrift (1961). 19

[10] Claude Berge and Pierre Duchet, *Problème*, January 1983, Publication Title: Séminaire MSH. 5, 13, 19, 34

[11] Piotr Berman, Marek Karpinski, and Alexander Scott, *Approximation hardness of short symmetric instances of MAX-3SAT*, Electron. Colloquium Comput. Complex. **049** (2003). 63

[12] Mostafa Blidia, *A parity digraph has a kernel*, Combinatorica **6** (1986), 23–27. 20

[13] Mostafa Blidia, Pierre Duchet, Henry Jacob, Frédéric Maffray, and Henry Meyniel, *Some operations preserving the existence of kernels*, Discrete Mathematics **205** (1999), 211 – 216. 5, 13, 29, 33, 34

[14] Mostafa Blidia and Konrad Engel, *Perfectly orderable graphs and almost all perfect graphs are kernel m-solvable*, Graphs and Combinatorics **8** (1992), no. 2, 103–108. 20

[15] Mostaffa Blidia, Pierre Duchet, and Frédéric Maffray, *On the orientation of Meyniel graphs*, Journal of Graph Theory **18** (1994), no. 705–711. 21

[16] Hans Bodlaender, Pål Gronås Drange, Markus Dregi, Fedor Fomin, Daniel Lokshtanov, and Michał Pilipczuk, *A c^kn 5-approximation algorithm for treewidth*, SIAM Journal on Computing **45** (2016), no. 2, 317–378. 72

[17] John Adrian Bondy, *Short proofs of classical theorems*, J. Graph Theory **44** (2003), no. 3, 159–165. 3, 11, 48

[18] Endre Boros and Vladimir Gurvich, *Perfect graphs are kernel solvable*, Discrete Math. **159** (1996), no. 1, 35–55. 5, 13, 19, 21, 28, 79

[19] ______, *Perfects graphs, kernels, and cores of cooperative games*, Discrete Math. **306** (2006), 2336–2354. 13, 29, 30

[20] Michel Burlet and Jean Fonlupt, *Polynomial algorithm to recognize a Meyniel graph*, Progress in Combinatorial Optimization, Elsevier, 1984, pp. 69–99. 21

[21] Michel Burlet and Jean Pierre Uhry, *Parity graphs*, North-Holland mathematics studies, vol. 88, Elsevier, 1984, pp. 253–277. 21

[22] Christophe Champetier, *Kernels in some orientations of comparability graphs*, Journal of Combinatorial Theory, Series B **47** (1989), no. 1, 111–113. 20, 42

[23] Maria Chudnovsky, Gérard Cornuéjols, Xinming Liu, Paul Seymour, and Kristina Vušković, *Recognizing Berge graphs*, Combinatorica **25** (2005), 143–187. 21

[24] ______, *Recognizing Berge graphs*, Combinatorica **25** (2005), no. 2, 143–186. 32

[25] Maria Chudnovsky, Neil Robertson, Paul Seymour, and Robin Thomas, *The strong perfect graph theorem*, Annals of mathematics (2006), 51–229. 19

[26] V Chvátal and L Lovász, *Every directed graph has a semi-kernel*, Hypergraph Seminar: Ohio State University 1972, Springer, 1974, pp. 175–175. 3, 11, 48

[27] Václav Chvátal, *On the computational complexity of finding a kernel*, Tech. Report CRM300, Centre de Recherches Mathématiques, Université de Montréal, 1973. 3, 11, 17, 79

[28] Economic Sciences Prize Committee et al., *Stable matching: Theory, evidence, and practical design-the sveriges riksbank prize in economic sciences in memory of alfred nobel 2012: Scientific background*, Tech. report, Tech. rep. The Nobel Foundation, Stockholm, Sweden, 2012. 24

[29] Bruno Courcelle, *The monadic second-order logic of graphs. i. recognizable sets of finite graphs*, Information and computation **85** (1990), no. 1, 12–75. 7, 14, 71

[30] Cosmina Croitoru, *A note on quasi-kernels in digraphs*, Information Processing Letters **115** (2015), no. 11, 863–865. 50, 80

[31] David Dailey, *Uniqueness of colorability and colorability of planar 4-regular graphs are NP-complete*, Discrete Math. **30** (1980), no. 3, 289–293. 62

[32] Olivier Durand de Gevigney, Frédéric Meunier, Christian Popa, Julien Reygner, and Ayrin Romero, *Solving coloring, minimum clique cover and kernel problems on arc intersection graphs of directed paths on a tree*, 4OR **9** (2011), 175–188. 33, 79

[33] Yannis Dimopoulos and Vangelis Magirou, *A graph-theoretic approach to default logic*, Information and computation **112** (1994), no. 2, 239–256. 25

[34] Irit Dinur and David Steurer, *Analytical approach to parallel repetition*, Proceedings of the forty-sixth annual ACM symposium on Theory of computing, 2014, pp. 624–633. 66

[35] Rodney Downey and Michael Fellows, *Fundamentals of parameterized complexity*, Texts in Computer Science, Springer, 2013. 66, 70

[36] Pierre Duchet, *Graphes noyau-parfaits*, Annals of Discrete Mathematics, vol. 9, Elsevier, 1980, pp. 93–101. 18, 79

[37] ———, *Parity graphs are kernel-m-solvable*, Journal of Combinatorial Theory, Series B **43** (1987), no. 2, 121–126. 19, 30

[38] Pierre Duchet and Henri Meyniel, *Une generalisation du theoreme de Richardson sur l'existence de noyaux dans les graphes orientes*, Discrete Mathematics **43** (1983), 21 – 27. 36

[39] Péter Erdős, Ervin Győri, Tamás Róbert Mezei, Nika Salia, and Mykhaylo Tyomkyn, *On the small quasi-kernel conjecture*, arXiv preprint arXiv:2307.04112 (2023). 49, 80

[40] Paul Erdős, Arthur Rubin, and Herbert Taylor, *Choosability in graphs*, Congr. Numer **26** (1979), no. 4, 125–157. 26

[41] Péter Erdős and László Székely, *Two conjectures on quasi-kernels, open problems no. 4. in fete of combinatorics and computer science*, Bolyai Society Mathematical Studies, 2010. 3, 11, 48

[42] David Gale and Lloyd Shapley, *College Admissions and the Stability of Marriage*, The American Mathematical Monthly **69** (1962), 9–15. 20, 23

[43] Hortensia Galeana-Sánchez, *A counterexample to a conjecture of Meyniel on kernel-perfect graphs*, Discrete Mathematics **41** (1982), no. 1, 105–107. 36

[44] Hortensia Galeana-Sánchez and Victor Neumann-Lara, *On kernels and semikernels of digraphs*, Discrete Mathematics **48** (1984), no. 1, 67–76. 5, 13, 19, 36, 79

[45] ———, *On kernel-perfect critical digraphs*, Discrete Mathematics **59** (1986), 257–265. 17

[46] Fred Galvin, *The list chromatic index of a bipartite multigraph*, Journal of Combinatorial Theory, Series B **63** (1995), no. 1, 153–158. 26

[47] Gregory Gutin, Khee Meng Koh, Eng Guan Tay, and Anders Yeo, *On the number of quasi-kernels in digraphs*, Report Series 01-7 (2001). 7, 14, 49, 51, 59

[48] ———, *On the number of quasi-kernels in digraphs*, J. Graph Theory **46** (2004), no. 1, 48–56. 49, 51, 59, 61, 62, 68

[49] Scott Heard and Jing Huang, *Disjoint quasi-kernels in digraphs*, J. Graph Theory **58** (2008), no. 3, 251–260 (en). 49, 80

[50] Chính T Hoàng, *On the complexity of recognizing a class of perfectly orderable graphs*, Discrete applied mathematics **66** (1996), no. 3, 219–226. 21

[51] Ayumi Igarashi, *Coalition formation in structured environments*, Proceedings of the 16th Conference on Autonomous Agents and Multiagent Systems, 2017, pp. 1836–1837. 23

[52] Ayumi Igarashi, Frédéric Meunier, and Adèle Pass-Lanneau, *Perfect graphs with polynomially computable kernels*, Discrete Applied Mathematics **272** (2020), 69–74. 27, 32, 33, 79

[53] Henry Jacob, *Kernels in graphs with a clique-cutset*, Discrete Math. **156** (1996), 265–267. 33

[54] Henry Jacob and Henry Meyniel, *About quasi-kernels in a digraph*, Discrete Math. **154** (1996), no. 1-3, 279–280. 49

[55] Richard Karp, *Reducibility among combinatorial problems*, Complexity of Computer Computations (R. Miller and J. Thatcher, eds.), Plenum Press, 1972, pp. 85–103. 60

[56] Shiva Kintali, Laura Poplawski, Rajmohan Rajaraman, Ravi Sundaram, and Shang-Hua Teng, *Reducibility among fractional stability problems*, SIAM Journal on Computing **42** (2013), no. 6, 2063–2113. 32, 79

[57] Tamás Király and Júlia Pap, *A note on kernels and Sperner's Lemma*, Discrete Applied Mathematics **157** (2009), 3327–3331. 29

[58] Joachim Kneis and Alexander Langer, *A practical approach to Courcelle's theorem*, Electronic Notes in Theoretical Computer Science **251** (2009), 65–81. 71

[59] Alexandr Kostochka, Ruth Luo, and Songling Shan, *Towards the Small Quasi-Kernel Conjecture*, arXiv:2001.04003 (2020). 49, 54, 62, 80

[60] Hélène Langlois, Frédéric Meunier, Romeo Rizzi, and Stéphane Vialette, *Algorithmic aspects of small quasi-kernels*, International Workshop on Graph-Theoretic Concepts in Computer Science, Springer, 2022, pp. 370–382. 59

[61] Min Chih Lin, Francisco Soulignac, and Jayme Szwarcfiter, *Arboricity, h-index, and dynamic algorithms*, Theoretical Computer Science **426** (2012), 75–90. 21

[62] László Lovász, *Normal hypergraphs and the perfect graph conjecture*, discrete Mathematics **2** (1972), no. 3, 253–267. 34

[63] Frédéric Maffray, *On kernels in i-triangulated graphs*, Discrete mathematics **61** (1986), no. 2-3, 247–251. 21

[64] ———, *Kernels in perfect line-graphs*, Journal of Combinatorial Theory, Series B **55** (1992), 1–8. 20, 24

[65] ———, *On the coloration of perfect graphs*, Recent Advances in Algorithms and Combinatorics (2003), 65–84. 20

[66] Christos Papadimitriou and Mihalis Yannakakis, *Optimization, approximation, and complexity classes*, J. Comput. Syst. Sci. **43** (1991), no. 3, 425–440. 67

[67] Moses Richardson, *On weakly ordered systems*, Bulletin of the American Mathematical Society **52** (1946), 113–116. 18

[68] Alvin Roth, *The evolution of the labor market for medical interns and residents: a case study in game theory*, Journal of political Economy **92** (1984), no. 6, 991–1016. 24

[69] Florian Roussel and Irena Rusu, *A linear algorithm to color i-triangulated graphs*, Information processing letters **70** (1999), no. 2, 57–62. 21

[70] Nicholas Roussopoulos, *A max* $\{m, n\}$ *algorithm for determining the graph* H *from its line graph* G, Information Processing Letters **2** (1973), no. 4, 108–112. 21

[71] Bill Sands, Norbert Sauer, and Robert Woodrow, *On monochromatic paths in edge-coloured digraphs*, Journal of Combinatorial Theory, Series B **33** (1982), no. 3, 271–275. 5, 13, 21, 42

[72] Nicolas Trotignon, *Perfect graphs: a survey*, arXiv preprint arXiv:1301.5149 (2013). 19, 20

[73] Allan van Hulst, *Kernels and small quasi-kernels in digraphs*, arXiv preprint arXiv:2110.00789 (2021). 49, 80

[74] John Von Neumann and Oskar Morgenstern, *Theory of games and economic behavior*, Princeton University Press, 1947. 3, 11, 18, 22

[75] Michał Walicki, *Resolving infinitary paradoxes*, The Journal of Symbolic Logic **82** (2017), no. 2, 709–723. 24