

A general class of combinatorial filters that can be minimized efficiently

Yulin Zhang and Dylan A. Shell

Abstract—State minimization of combinatorial filters is a fundamental problem that arises, for example, in building cheap, resource-efficient robots. But exact minimization is known to be NP-hard. This paper conducts a more nuanced analysis of this hardness than up till now, and uncovers two factors which contribute to this complexity. We show each factor is a distinct source of the problem’s hardness and are able, thereby, to shed some light on the role played by (1) structure of the graph that encodes compatibility relationships, and (2) determinism-enforcing constraints. Just as a line of prior work has sought to introduce additional assumptions and identify sub-classes that lead to practical state reduction, we next use this new, sharper understanding to explore special cases for which exact minimization is efficient. We introduce a new algorithm for constraint repair that applies to a large sub-class of filters, subsuming three distinct special cases for which the possibility of optimal minimization in polynomial time was known earlier. While the efficiency in each of these three cases previously appeared to stem from seemingly dissimilar properties, when seen through the lens of the present work, their commonality now becomes clear. We also provide entirely new families of filters that are efficiently reducible.

I. INTRODUCTION

Combinatorial filters are discrete transition systems that process streams of observations to produce outputs sequentially. They have found practical application as estimators in multi-agent tracking problems (e.g., [18]) and as representations of feedback plans/policies for robots (e.g., [11], [22]). Unlike traditional recursive Bayesian filters (the class of estimator most familiar to roboticists, see [17]), combinatorial filters allow one to ask questions regarding minimality. By reducing their size, one can design resource-efficient robots—a consideration of practical importance. More fundamentally, through filter minimization, one may discover what information is necessary to compute a particular estimate, or what needs to be tracked in order to have sufficient knowledge for a given task. Determining a task’s information requirements and limits is a basic problem with a long history in robotics [2], [9], and has begun gaining interest again (e.g., [10]). Unfortunately, given some combinatorial filter, computing the smallest equivalent filter—its *minimizer*—is an NP-hard problem.

This paper uncovers and examines two different factors which contribute to this complexity: the first has to do with the structure of the compatibility graph induced by the filter; the second involves auxiliary (or zipper) constraints added during minimization to ensure the result will be deterministic. As we show, both are distinct dimensions and

Y. Zhang is with Amazon Robotics, North Reading, MA, USA. zhangyl@amazon.com. The work was done prior to joining Amazon. D. A. Shell is with Dept. of Computer Science & Engineering, Texas A&M University, College Station, TX, USA. dshell@tamu.edu. This work was supported by the NSF through awards [IIS-1849249](#) and [IIS-2034097](#).

form independent sources of the problem’s hardness. This is the first contribution of the paper (and constitutes the subject of Section III).

Like most hard problems of practical significance, a line of research has sought specially structured sub-classes of filters that allow efficiency to be salvaged [14]. Another line has examined relaxed [16] or other restricted forms of reduction [15], [12]. In the prior work, three particular sub-classes of filter have been identified for which optimal filter minimization is known to be possible in polynomial time: (i) no-missing-edge filters [14], (ii) once-appearing-observation filters [14], and (iii) unitary filters [22].

The second portion of the paper, building upon the first, establishes a new sub-class of filters for which exact minimization is achievable in polynomial time. This sub-class strictly subsumes those of (i), (ii), and (iii), and also provides some understanding of why both factors—the compatibility graph and auxiliary/zipper constraints—are tame for these filters. Part of the answer is that it is possible to ignore the constraints because they can be repaired afterwards: Section IV introduces an algorithm for constraint repair that applies broadly, including for the new sub-class we study and, hence, the three prior ones as well. Another part of the answer, the requirement to quickly generate minimal clique covers, is feasible for the three prior sub-classes, (i)–(iii), because their compatibility graphs all turn out to be chordal. Thus, their apparent distinctiveness happens to be superficial and, in reality, their efficiency stems from some common underlying properties.

A. Context, general related work, and roadmap of the paper

As the contributions of this work are of a theoretical nature, we leave the customary motivating settings and example application domains to the work we cite next; each and every one of the following include specific problem instances, so we trust the reader will glance at those papers to allay any doubts as to practical utility. The term combinatorial filter was coined by Tovar *et al.* [18], and the minimization problem was formulated and its hardness established in [11]. The current state-of-the-art algorithm for combinatorial filter minimization was presented at ICRA’21 in [19]. The starting point for our current treatment is the authors’ paper [20], which showed that filter minimization is equivalent to the classic graph problem of determining the minimal clique cover, when augmented with auxiliary constraints.

The next section will provide necessary definitions and theoretical background. Section III first delineates important sub-families of graphs and uses them to establish our key hardness results. Section IV turns to constraints and provides an algorithm to repair constraint-violating solutions

when specific conditions are met. Thereafter, the results are consolidated into a new, efficiently minimizable sub-class of filters and this is connected with prior special sub-classes in Section V. The final section presents the conclusion. Space considerations have meant that most proofs had to be omitted; the full proofs all appear in [21].

II. PRELIMINARIES

A. Basic definitions

Definition 1 ([13]). A *deterministic filter*, or just *filter*, is a 6-tuple $\mathcal{F} = (V, v_0, Y, \tau, C, c)$, with V a non-empty finite set of states, v_0 an initial state, Y the set of observations, $\tau : V \times Y \rightarrow V$ the partial function describing transitions, C the set of outputs, and $c : V \rightarrow C$ being the output function.

Filters process finite sequences of elements from Y in order to produce a corresponding sequence of outputs (elements of C). Any filter does this by tracing from v_0 along edges (defined by the transition function) and producing outputs (via the c function) as it visits states. States v_i and v_k are understood to be connected by a directed edge bearing label y , if and only if $\tau(v_i, y) = v_k$. We will assume Y to be non-empty, and that no state is unreachable from v_0 .

This paper's central concern is the following problem:

Problem: Filter Minimization ($\text{FM}(\mathcal{F})$)

Input: A deterministic filter \mathcal{F} .

Output: A deterministic filter \mathcal{F}^* with fewest states, such that:

1. any sequence which can be traced on \mathcal{F} can also be traced on \mathcal{F}^* ;
2. the outputs they produce on any of those sequences are identical.

Solving this problem requires some minimally-sized filter \mathcal{F}^* that is functionally equivalent to \mathcal{F} , where the notion of equivalence—called *output simulation*—needs only criteria 1. and 2. to be met. For a formal definition of output simulation, see [20, Definition 5, pg. 93].¹

Lemma 2 ([11]). The problem FM is NP-hard.

B. Constrained clique covers on a graph

In giving a minimization algorithm, FM was recently connected to an equivalent graph covering problem [20]. To start, we consider this problem abstractly in isolation:

Problem: Minimum Zipped Clique Cover ($\text{MZCC}(G, Z)$)

Input: A graph $G = (V, E)$ and collection of zipper constraints $Z = \{(U_1, V_1), \dots, (U_m, V_m)\}$, with $U_i, V_i \subseteq V$.

Output: Minimum cardinality clique cover \mathbb{K} such that:

1. $\bigcup_{K_i \in \mathbb{K}} K_i = V$, with each K_i forming a clique on G ;
2. $\forall K_i \in \mathbb{K}$, if there is some ℓ such that $U_\ell \subseteq K_i$, then some $K_j \in \mathbb{K}$ must have $K_j \supseteq V_\ell$.

The constraints in Z are still rather arcane, hence the next section, in making a connection to FM , provides an explanation of how Z is used, what it is for, and why it bears the moniker *zipper*.

¹After examining filters like those here, the later sections of that paper go further by studying a generalization in which function c may be a relation. Complications arising from that generalization will not be discussed herein.

C. Filter minimization as constrained clique covering

In bridging filters and clique covers, the key idea is that certain sets of states in a filter can be identified as candidates to merge together, and such ‘mergability’ can be expressed as a graph. The process of forming covers of this graph identifies states to consolidate and, accordingly, minimal covers yield small filters. The first technical detail concerns this graph and states that are candidates to be merged:

Definition 3 (extensions and compatibility). For a state v of filter \mathcal{F} , we will use $\mathcal{L}_{\mathcal{F}}(v)$ to denote the set of observation sequences, or *extensions*, that can be traced starting from v . Two states v and w are *compatible* with each other if their outputs agree on $\mathcal{L}_{\mathcal{F}}(v) \cap \mathcal{L}_{\mathcal{F}}(w)$, their common extensions. In such cases, we will write $v \sim w$. The *compatibility graph* $G_{\mathcal{F}}$ possesses edges between states iff they are compatible.

However, simply building a minimal cover on $G_{\mathcal{F}}$ is not enough because covers may merge some elements which, when transformed into a filter, produce nondeterminism. The core obstruction is when a fork is created, as when two compatible states are merged, both of which have outgoing edges bearing identical labels, but whose destinations differ. To enforce determinism, we introduce constraints that forbid forking and require mergers to flow downwards. The following specifies such a constraint:

Definition 4 (determinism-enforcing zipper constraint). Given a set of mutually compatible states U in \mathcal{F} and the set of all its y -children $V = \{v \mid u \in U, \tau(u, y) = v\}$, then the pair (U, V) is a *determinism-enforcing zipper constraint* of \mathcal{F} .

A zipper constraint is satisfied by a clique if U is not covered in a clique, or both U and V are covered in cliques. (This is criterion 2. for MZCC.) For filters, in other words, if the states in U are to be consolidated then the downstream states, in V , must be as well.

The collection of all determinism-enforcing zipper constraints for a filter \mathcal{F} is denoted $Z_{\mathcal{F}}$. Both $G_{\mathcal{F}}$ and $Z_{\mathcal{F}}$ are clearly polynomial in the size of \mathcal{F} . Then, a minimizer of \mathcal{F} can be obtained from the solution to the minimum zipped vertex cover problem, $\text{MZCC}(G_{\mathcal{F}}, Z_{\mathcal{F}})$:

Lemma 5 ([20]). Any $\text{FM}(\mathcal{F})$ can be converted into an $\text{MZCC}(G_{\mathcal{F}}, Z_{\mathcal{F}})$ in polynomial time; hence MZCC is NP-hard.

Though we skip the details, the proof in [20] of the preceding also gives an efficient way to construct a deterministic filter from the minimum cardinality clique cover.

III. HARDNESS: REEXAMINED AND REFINED

The recasting of $\text{FM}(\mathcal{F})$ as $\text{MZCC}(G_{\mathcal{F}}, Z_{\mathcal{F}})$ leads one naturally to wonder: what precise role do the compatibility graph and zipper constraints play with regards to hardness?

A. Revisiting the original result

Firstly, examining the proof of Lemma 2, the argument in [11] proceeds by reducing the graph 3-coloring problem to filter minimization. Looking at that construction carefully, one observes that the FM instance that results from any 3-coloring problem does not have any zipper constraints.

Hence, by writing MZCC with compatibility graph of \mathcal{F} and no zipper constraints as $\text{MZCC}(G_{\mathcal{F}}, \emptyset)$, we get the following:

Lemma 6. $\text{MZCC}(G_{\mathcal{F}}, \emptyset)$ is NP-hard.

A superficial glance might cause one to think of MZCC with an empty collection of zipper constraints as the standard minimum clique cover problem, viz., № 13 of Karp's original 21 NP-complete problems [8]. Actually, Lemma 6 states that the clique cover instances arising in minimization of filters are NP-hard; note that this is neither a direct restatement of Karp's original fact nor merely entailed by it. (But see, also, Theorem 17 below.)

B. Special graphs: efficiently coverable cases

To begin to investigate problems with special structure, our starting point is to recognize that several specific sub-families of undirected graphs (some widely known, others more obscure) allow a minimal clique cover to be obtained efficiently. We formalize such cases with the following.

Definition 7. A sub-family of graphs \mathbf{G} is *efficiently coverable* if some algorithm $\mathcal{A}_{\mathbf{G}}$ exists so that, $\forall G \in \mathbf{G}$, $\mathcal{A}_{\mathbf{G}}(G)$ produces a minimal clique cover of G in polynomial time.

In filters with efficiently coverable compatibility graphs, when also $Z_{\mathcal{F}} = \emptyset$, then criterion 2. of MZCC holds vacuously and $\text{FM}(\mathcal{F})$ will be efficient. The contrast of this statement with Lemma 6, shows that the efficiently coverable sub-families carve out subsets of easy problems.

Lemmas 9, 11 and 13, and Theorem 14, which will follow, review some instances of efficiently coverable graphs:

Definition 8 ([5]). A graph is *chordal* if all cycles of four or more vertices have a chord, which is an edge not part of the cycle but which connects two vertices in the cycle.

Lemma 9 ([4]). Chordal graphs are efficiently coverable.

A strictly larger class of graphs are those that are perfect.

Definition 10 ([5]). A *perfect graph* is a graph where the chromatic number of every induced subgraph equals the order of the largest clique of that subgraph.

Lemma 11 ([6]). Perfect graphs are efficiently coverable.

As all chordal graphs are also perfect, Lemma 9 follows from Lemma 11, and the reader may wonder why then chordal graphs are worth mentioning explicitly. Three reasons: (1) the requirements of Definition 8 tend to be less demanding to check than those in Definition 10, which involve some indirectness; (2) the polynomial-time algorithm of [6] (referenced in proof of Lemma 11) is not a direct combinatorial method and, in fact, researchers continue to contribute practical methods tailored to specific sub-classes of perfect graphs (e.g., [1]); (3) chordal graphs will show up in the proofs, including in the next section.

But there are graphs, beyond only those which are perfect, that still give efficiently coverable problems:

Definition 12 ([5]). A *triangle-free* graph is a graph where no three vertices have incident edges forming a triangle.

A specific triangle-free graph that is not perfect is the Grötzsch graph.

Lemma 13 ([3]). The triangle-free graphs are efficiently coverable.

Finally, composition allows treatment of graphs with mixed properties, e.g., we might have filters with compatibility graphs where some components are perfect, and others are triangle-free. The following fact is useful in such cases.

Theorem 14 (mix-and-match). Suppose a graph $G = (V, E)$, where $E \subseteq V \times V$, is made up of components G_1, G_2, \dots, G_m , every $G_i = (V_i, E_i)$, where V is partitioned into mutually disjoint set of vertices V_1, V_2, \dots, V_m , and E_1, E_2, \dots, E_m with every $E_i \subseteq V_i \times V_i$. If each of the G_i 's is efficiently coverable, then G is efficiently coverable.

Filters naturally yield graphs comprising separate components as the output values directly partition the vertices. That is, compatibility graphs never possess edges between any vertices v and w where $c(v) \neq c(w)$.

C. Special compatibility graphs and non-empty zippers

In light of Lemma 6 showing that zippers are not needed to have hard problems, and the fact that there are sub-families of graphs for which minimal covers may be obtained efficiently, we next ask: do the zipper constraints themselves contribute enough complexity so that even with an efficiently coverable instance, we can get a hard problem?

The answer is in the affirmative and we use the sub-family of chordal graphs to establish this. We begin with a triangulation procedure that, given a general graph, constructs one which is chordal. The approach to the proof is to think about solving MZCC on the chordal version and then relate the solution back to the original problem.

A graph is non-chordal if and only if there is an m -cycle, with $m \geq 4$ and there is no edge cutting across the cycle. We can break such a cycle into smaller ones by adding edges as shortcuts. Repeating this process will triangulate such cycles and the procedure must eventually terminate as the complete graph is chordal. We call these newly introduced edges *dashed* as this is how we shall depict them visually.

Having introduced extra edges, the idea is to discourage clique covers from ever choosing to cover any of these new dashed edges via penalization. A penalty is incurred by being compelled to choose additional cliques—zipper constraints are rich enough to force such choices. This requires the introduction of a gadget we term a ‘necklace’. Suppose the original non-chordal graph G had n vertices, m_s edges, and that an additional m_d dashed edges were used to triangulate the graph. Then, as illustrated in Figure 1, we first make $\ell = m_s + m_d + 1$ copies of 2-vertex connected graphs, which we dub ‘pendants’. These are laid in a line, and between any pair of pendants, we place a single black vertex that we call a ‘bead’. The ℓ pendants and $\ell - 1$ beads are strung together via edges, each bead being connected to the two adjacent pendants. We’ll call these connecting edges ‘strings’. To each dashed edge we add $2\ell - 2 = 2(m_s + m_d)$ zipper constraints, connecting the dashed edge to the length of strings. This construction means that when a dashed edge is covered, its zipper constraints become active and then each bead

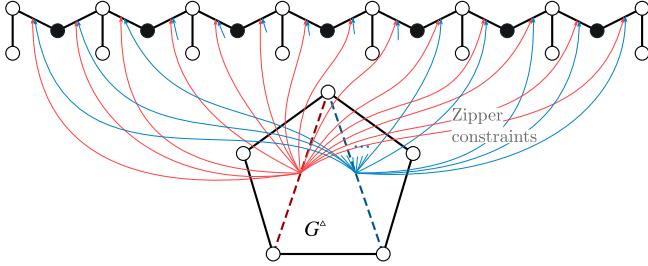


Fig. 1: A clique cover problem on a general graph is reduced to a clique cover problem on a chordal graph with extra zipper constraints. Two dashed edges are added to the pentagon in order to triangulate it, resulting in a 7-edge chordal graph. Dashed edges are made undesirable through the addition of zipper constraints that trigger the necklace string (at top). Zipper constraints are represented as arrows, shown in red and blue to associate them visually with their dashed edge. (Note: parts of some blue arrows have been elided to reduce visual clutter).

will have to appear in two separate covers, one for each neighboring pendant.

Given graph and zipper constraints (G, Z) , we will denote the result of the construction just described with (G^Δ, Z^Δ) , the first element being the chordal graph along with the necklace, and the second element being the additional constraints. Then: (1) G^Δ is chordal (dashed edges made G chordal, the necklace itself is chordal), (2) $G \subseteq G^\Delta$, as vertices/edges were added, never removed, (3) $Z \subseteq Z^\Delta$, as constraints were added, not removed. Further, notice that G^Δ and Z^Δ are no larger than some polynomial factor of n . This construction takes $O(n^4)$. The purpose of this construction is the following:

Lemma 15. Given any non-empty graph G and zipper constraints Z , a solution to $\text{MZCC}(G, Z)$ can be obtained from any solution S^Δ to $\text{MZCC}(G^\Delta, Z^\Delta)$, by restricting S^Δ to only those covers on the vertices of G .

This is one of the paper's main results: the proof calculates the penalty incurred by covering any dashed edge (i.e., the additional covers for part of the necklace) and shows it to be so large that even covering non-dashed edges in pairs would be preferable. Any optimal cover would do no worse than this and, thus, a zipper constraint can never be triggered.

Theorem 16. $\text{MZCC}(G_{\text{ee}}, Z)$, where G_{ee} is efficiently coverable and $Z \neq \emptyset$ is NP-hard.

The proof uses Lemma 6, for a known NP-hard problem, in order to obtain a reduction via Lemma 15.

The preceding leads to the following interpretation. If a filter \mathcal{F} induces an efficiently coverable $G_{\mathcal{F}}$ then:

- when $Z_{\mathcal{F}} = \emptyset$, since criterion 2. of MZCC holds vacuously, $\text{FM}(\mathcal{F})$ is in \mathbb{P} ;
- when $Z_{\mathcal{F}} \neq \emptyset$, despite $G_{\mathcal{F}}$ being efficiently coverable, $\text{FM}(\mathcal{F})$ should be suspected as intractable because it is NP-hard in the worst case.

Additional support for this link between the filter structure, FM , MZCC, and worst-case intractability is the following:

Theorem 17. A graph G can be realized as the compatibility graph of some filter if and only if G either: (1) has at least two connected components, or (2) is a complete graph.

The construction hinges on assigning unique outputs to each connected component. Thereafter, incompatibilities are easily manufactured in a filter by simply introducing new observation symbols that go to states with different outputs.

Corollary 18. Chordal graphs with zippered necklace structures, G^Δ , as in the construction involved in Lemma 15, are realizable from filters.

IV. REPAIRABLE ZIPPER CONSTRAINTS

If general zipper constraints introduce enough complexity that the problem is hard even when the graph is efficiently coverable, and yet absence of zipper constraints gives an easy problem, how do we obtain a more discriminating conception of zipper constraints and their structure? And, specifically, are there special cases of filters which give ‘nice’ zipper constraints? In this section, we first formalize sufficient conditions in order to ignore zipper constraints; in these cases once a clique cover has been obtained, we can modify it to make the zipper constraints hold. This modification step can repair, in polynomial time, any zipper-constraint-violating clique cover for which the sufficient conditions are met and, crucially, can do this without causing any increase in size.

In any graph $G = (V, E)$, we refer to the neighbors of a vertex $v \in V$ by set $\mathcal{N}_G(v) := \{w \in V \mid (w, v) \in E\} \cup \{v\}$. Note that we explicitly include v in its own neighborhood.

Through neighborhoods, the following condition now describes a type of harmony between zipper constraints and the compatibility relation.

Definition 19 (comparable zip candidates). Given a graph $G = (V, E)$, it has the *comparable zip candidates* property with respect to an associated collection of zipper constraints $Z = \{(U_1, W_1), (U_2, W_2), \dots, (U_m, W_m)\}$, where $U_k, W_k \subseteq V$, if and only if every pair of vertices $\{w_i^0, w_i^1\} = W_i$ satisfies either $\mathcal{N}_G(w_i^0) \subseteq \mathcal{N}_G(w_i^1)$ or $\mathcal{N}_G(w_i^0) \supseteq \mathcal{N}_G(w_i^1)$.

The preceding definition is a sufficient condition to yield MZCC problems whose zipper constraints may be repaired.

Lemma 20. Let filter \mathcal{F} ’s compatibility graph $G_{\mathcal{F}} = (V_{\mathcal{F}}, E_{\mathcal{F}})$ possess the comparable zip candidates property with respect to $Z_{\mathcal{F}}$. Suppose cover \mathbb{K} of $G_{\mathcal{F}}$ violates $Z_{\mathcal{F}}$. Then one may obtain, in polynomial time, a cover \mathbb{M} that will satisfy $Z_{\mathcal{F}}$ and has $|\mathbb{M}| \leq |\mathbb{K}|$.

Repair algorithm and proof. Given cover $\mathbb{K} = \{K_1, K_2, \dots, K_k\}$, the proof constructs $\mathbb{M} = \{M_1, M_2, \dots, M_m\}$ such that all zipper constraints will be satisfied through \mathbb{M} , with $m \leq k$. For $i \in \{1, \dots, k\}$, form the sets M_i :

$$M_i := K_i \cup \left\{ w_a^{\hat{n}} \mid \left(U_a, \{w_a^0, w_a^1\} \right) \in Z_{\mathcal{F}}, n \in \{0, 1\}, \right. \quad (1) \\ \left. w_a^n \in K_i, \mathcal{N}_{G_{\mathcal{F}}}(w_a^n) \subseteq \mathcal{N}_{G_{\mathcal{F}}}(w_a^{\hat{n}}) \right\},$$

where $\hat{n} := 1 - n$. Essentially, to each K_i we are adding some extra elements. When K_i contains an element w_a^n that is paired with $w_a^{\hat{n}}$, another vertex in the zipper constraints, then we include $w_a^{\hat{n}}$ if it possesses a neighborhood that is no smaller than w_a^n ’s. After doing this for all K_1, \dots, K_k ,

we obtain the collection $\mathbb{M} = \{M_1, \dots, M_m\}$ with $m \leq k$ (and $m < k$ only if some sets grew to become identical).

All vertices previously covered are still covered. Moreover, each M_i is a clique because K_i is a clique, and one can argue inductively as if the extra elements were added sequentially: As each element $w_a^{\widehat{n}}$ is added, it is compatible with element w_a^n already in the clique (owing to the pair being a zipper constraint target). Element $w_a^{\widehat{n}}$ has a compatibility neighborhood at least as large as that of w_a^n 's, so is compatible with the clique, and the composite thus forms a clique itself.

It only remains to show that all $Z_{\mathcal{F}}$ is now satisfied. Consider any constraint (U_i, W_i) , with $W_i = \{w_i^0, w_i^1\}$, then we show that both w_i^0 and w_i^1 appear together in some element in \mathbb{M} . The comparable zip candidates property means there is an $n \in \{0, 1\}$ such that $\mathcal{N}_{G_{\mathcal{F}}}(w_i^n) \subseteq \mathcal{N}_{G_{\mathcal{F}}}(w_i^{\widehat{n}})$. Since \mathbb{K} is a cover, at least one K_j exists such that $w_i^n \in K_j$. But then, via (1), M_j must include both w_i^n and $w_i^{\widehat{n}}$. \square

Notice that the scope of Definition 19 includes graphs and zipper constraints generally, while Lemma 20 concerns compatibility graphs and zipper constraints obtained specifically from filters, and the additional structure inherited from the filter shows up in the proof itself.

V. SPECIAL CASES: EFFICIENTLY REDUCIBLE FILTERS

Up to this point, sufficient conditions have been presented for favorable MZCC problem instances. Each condition concerns a separate factor: the first, in Section III, deals with structural properties of graphs; while in Section IV, the second involves the zipper constraints being accordant with neighborhoods of potentially zipped vertices. The two reflect different dimensions and, as argued above, form distinct sources of the problem's hardness. To consolidate—

A sub-class of filters that can be minimized efficiently:

Any filter with efficiently coverable compatibility graph and comparable zip candidates can be minimized efficiently.

One first constructs the compatibility graph, finds any minimum clique cover, and then repairs it using Lemma 20.

Notice that to verify membership of this sub-class involves requirements pertaining to products derived from filters; we will now give closer scrutiny to filters themselves.

A. Handy properties that yield efficiently reducible filters

We seek properties that are recognizable and verifiable on filters directly. To start with, here is a sufficient condition:

Definition 21. A filter \mathcal{F} is *globally language comparable* if, for every pair of compatible states $v \sim w$, either extensions $\mathcal{L}_{\mathcal{F}}(v) \subseteq \mathcal{L}_{\mathcal{F}}(w)$ or extensions $\mathcal{L}_{\mathcal{F}}(v) \supseteq \mathcal{L}_{\mathcal{F}}(w)$.

Filters that are globally language comparable have compatibility graphs that are efficiently coverable.

Lemma 22. If \mathcal{F} is any globally language comparable filter, then its compatibility graph $G_{\mathcal{F}}$ is chordal.

In fact, globally language comparable filters also have benign zipper constraints.

Lemma 23. If \mathcal{F} is a globally language comparable filter, then compatibility graph $G_{\mathcal{F}}$ possesses the comparable zip candidates property with respect to zipper constraints $Z_{\mathcal{F}}$.

Theorem 24. Globally language comparable filters can be minimized efficiently.

Here, the proof just puts Lemmas 22 and 23 together.

A different but also potentially useful way to identify special cases is to use the compatibility relation: for a particular problem instance, one might determine whether specific algebraic properties hold. To get there, first we start with a characteristic involving neighborhoods.

Definition 25. A filter \mathcal{F} is *neighborhood comparable* iff every pair of vertices v, w in $G_{\mathcal{F}}$ with $\mathcal{N}_{G_{\mathcal{F}}}(v) \cap \mathcal{N}_{G_{\mathcal{F}}}(w) \neq \emptyset$, satisfies either $\mathcal{N}_{G_{\mathcal{F}}}(v) \subseteq \mathcal{N}_{G_{\mathcal{F}}}(w)$ or $\mathcal{N}_{G_{\mathcal{F}}}(v) \supseteq \mathcal{N}_{G_{\mathcal{F}}}(w)$.

Next, we connect it with an algebraic property.

Property 26. A filter \mathcal{F} is neighborhood comparable iff it induces a compatibility relation \sim on states such that \sim is an equivalence relation.

Further, the neighborhood comparable property induces an efficiently coverable compatibility graph:

Lemma 27. Any neighborhood comparable filter has a compatibility graph that is chordal.

Also we can show that a neighborhood comparable filter has zipper constraints that are repairable.

Lemma 28. Any neighborhood comparable filter \mathcal{F} has compatibility graph $G_{\mathcal{F}}$ that possesses the comparable zip candidates property with respect to zipper constraints $Z_{\mathcal{F}}$.

Theorem 29. Neighborhood comparable filters can be minimized efficiently.

The properties described in the section—one on the extensions and another on neighborhoods—are useful because they are not difficult requirements to verify and they imply facts about both the compatibility graph and zipper constraints. Still, they are fairly abstract. One might wonder, for instance, whether Definitions 21 and 25 really differ essentially. (They are distinct, as we will see shortly.)

B. Prior cases in the literature

We now use the conditions just introduced to re-examine three sub-classes of filter for which polynomial-time minimization has been reported in the literature. This treatment provides a new understanding of the relationships between these special cases. To start, each sub-class must be defined.

When discussing the fact that FM is NP-hard, the authors in [11] point out that this may, at first, seem unexpected since minimization of deterministic finite automata (DFA) is efficient (e.g., via the theorem of Myhill–Nerode [7]). As an intuition for this difference, they offer the following perspective: when a sequence crashes on a DFA, that string is outside of the automaton's language—whereas, when a sequence crashes on a filter, FM allows the minimizer to

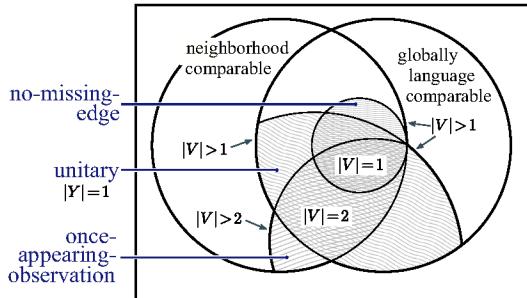


Fig. 2: The relationships between the three sub-classes of filter in terms of the globally language and neighborhood comparable properties identified. (Note: Degenerate filters are omitted.)

select any output for it, and some choices will likely give more compression than others. These degrees-of-freedom represent a combinatorial choice within the FM problem.

One way to curtail the explosion of such choices, then, is to ensure that no strings can ever crash:

Definition 30 ([14]). In a *no-missing-edge filter*, every state has an outgoing edge for every observation in Y .

For any no-missing-edge filter \mathcal{F} , we have the language $\mathcal{L}_{\mathcal{F}}(v_0) = Y^*$, i.e., the Kleene star of the set of observations.

The authors of [14] also identify a sort of obverse to the foregoing sub-class. If no-missing-edge filters fully re-use observations, every state having all of them, consider next the sub-class where observations occur at precisely one state:

Definition 31 ([14]). A *once-appearing-observation filter* is a filter where every observation in Y appears at most once.

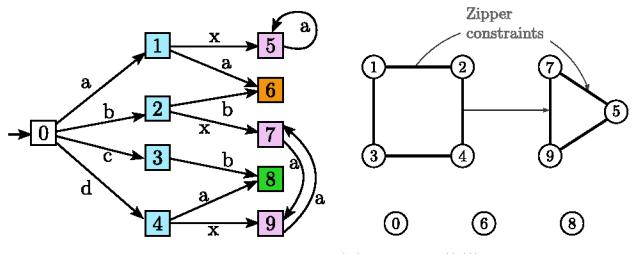
In a quite different context, the cardinality of the observation set was shown to affect complexity; constraining this gives another sub-class of filters:

Definition 32 ([22]). A *unitary filter* is a filter with a set of observations that is a singleton, i.e., $|Y| = 1$.

Despite the fact that the preceding three sub-classes impose a diverse assortment of constraints, their common trait is efficiency: any filter belonging to those sub-classes can be minimized in time that is polynomial in the input size [14, Thms. 3 and 4], [22, Thm. 3].

The sub-class of once-appearing-observation filters is disjoint from the no-missing-edge ones, with the sole exception of filters with only a single state. Any unitary filter is (exclusively) either a linear chain, or includes a cycle. In the latter case, every state will have an outgoing edge, and that filter therefore has no edges missing. A unitary filter with multiple states can only have observations that appear once when it is a chain with $|V| = 2$; longer chains are neither no-missing-edge nor once-appearing-observation filters.

Each of these three filters possess at least one of the properties introduced in Section V-A. Along with their inter-relationships, this is illustrated in Figure 2. None of the sub-classes strictly subsumes the others, yet common underlying properties explain their efficiency: (i) their compatibility graphs will be chordal and (ii) their zipper constraints are tame. The route by which these facts are established depends



(a) A simple 10-state filter $\mathcal{F}_{\text{poly}}$, with outputs visualized as colors. **(b)** Compatibility graph $G_{\mathcal{F}_{\text{poly}}}$ and zipper constraint set.

Fig. 3: A filter now recognizable as efficiently minimizable.

upon the sub-class: using the globally language comparable property, (i) comes from Lemma 22 and (ii) from Lemma 23; using the neighborhood comparable property, they come from Lemmas 27 and 28, respectively.

C. A new instance now recognizable as efficiently solvable

Thus, it has turned out that the compatibility graphs of previously known special cases are all chordal. This gives a ready means by which extra FM instances can be identified as having efficient solutions, instances that have never previously been recognized as such. One basic example is shown in Figure 3. This filter, $\mathcal{F}_{\text{poly}}$, has a compatibility graph that fails to be chordal, but is a perfect graph. Additionally, $\mathcal{F}_{\text{poly}}$ possesses the comparable zip candidates property because states 5, 7, 9 are all mutually compatible. More generally, the perfect graphs include many non-chordal members, including bipartite graphs, Turán graphs, etc. Any filters with these as compatibility graphs (and Theorem 17 says there *are* such filters) are candidates as additional special cases over and above what was previously known; one then needs to ensure their zipper constraints are benign: say, being \emptyset or repairable.

VI. CONCLUSION

This paper has identified basic properties underlying instances of filter minimization that are easy to solve. There are two aspects: the structure of the compatibility graph and the determinism-enforcing zipper constraints; they are distinct and both are shown to matter—as NP-hard problems arise if a single aspect is constricted but the other given free rein. Uncovering these facts, the key contribution of the first part of the paper, involves distinguishing and formalizing several subtle properties, proving new hardness results, and devising an efficient algorithm for repair of violated zipper constraints.

The paper then turns to more pragmatic ways the preceding insights can be leveraged: as is usual with NP-hard problems, researchers have sought conditions that identify sub-classes for which minimizers can be found efficiently. The sub-class we identify subsumes the previously known special cases. The paper further improves understanding of previously known sub-classes, detailing their differences and also drawing out commonalities. Finally, the paper gives an example filter which, prior to here, would not be recognized as possessing an efficient solution.

REFERENCES

- [1] Flavia Bonomo, Gianpaolo Oriolo, Claudia Snels, and Gautier Stauffer. Minimum Clique Cover in Claw-Free Perfect Graphs and the Weak Edmonds–Johnson Property. In *International Conference on Integer Programming and Combinatorial Optimization (IPCO)*, pages 86–97, 2013.
- [2] Bruce R. Donald. *The Compass That Steered Robotics*, pages 50–65. Springer Berlin Heidelberg, 2012. Logic and Program Semantics: Essays Dedicated to Dexter Kozen on the Occasion of His 60th Birthday.
- [3] Jack Edmonds. Paths, trees, and flowers. *Canadian Journal of Mathematics*, 17:449–467, 1965.
- [4] Fănică Gavril. Algorithms for Minimum Coloring, Maximum Clique, Minimum Covering by Cliques, and Maximum Independent Set of a Chordal Graph. *SIAM Journal on Computing*, 1(2):180–187, 1972.
- [5] Jonathan L. Gross, Jay Yellen, and Ping Zhang, editors. *Handbook of Graph Theory*. Taylor & Francis Group, second edition, 2013.
- [6] Martin Grötschel, László Lovász, and Alexander Schrijver. *Geometric Algorithms and Combinatorial Optimization*. Springer, Berlin, 1988.
- [7] J. E. Hopcroft and J. D. Ullman. *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley, Reading, MA, second edition, 1979.
- [8] Richard M Karp. Reducibility among combinatorial problems. In *Complexity of computer computations*, pages 85–103. Springer, 1972.
- [9] Steven M. LaValle. Sensing and filtering: A fresh perspective based on preimages and information spaces. *Foundations and Trends in Robotics*, 1(4):253–372, 2010.
- [10] Anirudha Majumdar and Vincent Pacelli. Fundamental Performance Limits for Sensor-Based Robot Control and Policy Learning. In *Robotics: Science and Systems*, New York City, NY, USA, June 2022.
- [11] Jason M. O’Kane and Dylan A. Shell. Concise planning and filtering: Hardness and algorithms. *IEEE Transactions on Automation Science and Engineering*, 14(4):1666–1681, 2017.
- [12] Hazhar Rahmani and Jason M. O’Kane. Equivalence notions for state-space minimization of combinatorial filters. *IEEE Transactions on Robotics*, 37(6):2117–2136, 2021.
- [13] Fatemeh Zahra Saberifar, Shervin Ghasemlou, Jason M O’Kane, and Dylan A Shell. Set-labelled filters and sensor transformations. In *Robotics: Science and Systems*, Ann Arbor, Michigan, 2016.
- [14] Fatemeh Zahra Saberifar, Ali Mohades, Mohammadreza Razzazi, and Jason M. O’Kane. Combinatorial Filter Reduction: Special Cases, Approximation, and Fixed-Parameter Tractability. *Journal of Computer and System Sciences*, 85:74–92, May 2017.
- [15] Fatemeh Zahra Saberifar, Ali Mohades, Mohammadreza Razzazi, and Jason M. O’Kane. Improper Filter Reduction. *Journal of Algorithms and Computation*, 50(1):69–99, June 2018.
- [16] Fatemeh Zahra Saberifar, Jason M. O’Kane, and Dylan Shell. Inconsequential Improprieties: Filter Reduction in Probabilistic Worlds. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and System*, 2017.
- [17] Sebastian Thrun, Wolfram Burgard, and Dieter Fox. *Probabilistic Robotics*. MIT Press, Cambridge, MA, U.S.A., 2005.
- [18] Benjamin Tovar, Fred Cohen, Leonardo Bobadilla, Justin Czarnowski, and Steven M Lavalle. Combinatorial filters: Sensor beams, obstacles, and possible paths. *ACM Transactions on Sensor Networks*, 10(3):1–32, 2014.
- [19] Yulin Zhang, Hazhar Rahmani, Dylan A Shell, and Jason M O’Kane. Accelerating combinatorial filter reduction through constraints. In *Proceedings of IEEE International Conference on Robotics and Automation*, pages 9703–9709, 2021.
- [20] Yulin Zhang and Dylan A. Shell. Cover combinatorial filters and their minimization problem. In *Algorithmic Foundations of Robotics XIV*, pages 90–106. Springer, 2021.
- [21] Yulin Zhang and Dylan A. Shell. A general class of combinatorial filters that can be minimized efficiently. *arXiv preprint arXiv:2209.04567*, 2022.
- [22] Yulin Zhang and Dylan A. Shell. Nondeterminism subject to output commitment in combinatorial filters. In *Algorithmic Foundations of Robotics XV*. Springer, 2022.