

Kernelization, Proof Complexity and Social Choice

Gabriel Istrate ✉

West University of Timișoara, Romania

Cosmin Bonchiș

West University of Timișoara, Romania

Adrian Crăciun

West University of Timișoara, Romania

Abstract

We display an application of the notions of kernelization and data reduction from parameterized complexity to proof complexity: Specifically, we show that the existence of data reduction rules for a parameterized problem having (a). a small-length reduction chain, and (b). small-size (extended) Frege proofs certifying the soundness of reduction steps implies the existence of subexponential size (extended) Frege proofs for propositional formalizations of the given problem.

We apply our result to infer the existence of subexponential Frege and extended Frege proofs for a variety of problems. Improving earlier results of Aisenberg et al. (ICALP 2015), we show that propositional formulas expressing (a stronger form of) the Kneser-Lovász Theorem have quasipolynomial size Frege proofs for each constant value of the parameter k .

Another notable application of our framework is to impossibility results in computational social choice: we show that, for any fixed number of agents, propositional translations of the Arrow and Gibbard-Satterthwaite impossibility theorems have subexponential size Frege proofs.

2012 ACM Subject Classification Theory of computation → Proof complexity

Keywords and phrases Kernelization, Frege proofs, Kneser-Lovász Theorem, Arrow’s theorem, Gibbard-Satterthwaite theorem

Digital Object Identifier 10.4230/LIPIcs.ICALP.2021.135

Category Track B: Automata, Logic, Semantics, and Theory of Programming

Related Version *Full Version:* <https://arxiv.org/abs/2104.13681>

1 Introduction

The central task of *proof complexity* [10, 30] is that of understanding (and distinguishing) the relative power of various propositional proof systems. Proving lower bounds for stronger and stronger proof systems might (in principle) be a way to eventually confirm the various conjectures of computational complexity. Yet we are far from being able to prove exponential lower bounds for some concrete problems in strong proof systems.

One of the most important open problems in this area, explicitly raised by Bonet, Buss and Pitassi [6] is that of separating the complexity of Frege proof systems (“textbook style propositional proofs”) from that of extended Frege proof systems (which in addition can introduce new variables as substitutes for arbitrary propositional formulas). That is, we would like to find explicit classes of propositional formulas that have extended Frege proofs of polynomial size but have exponential lower bounds on the size of the shortest Frege proofs.

Many classes of problems that are candidates for separating the two systems have been proposed, e.g. statements based on linear algebra [24, 36], propositional encodings of the Paris-Harrington independence results [11], Ramsey’s theorem [31], central theorems from extremal combinatorics [1, 34], or the Kneser-Lovász formula from combinatorial topology [2, 27]).



© Gabriel Istrate, Cosmin Bonchiș, and Adrian Crăciun;
licensed under Creative Commons License CC-BY 4.0

48th International Colloquium on Automata, Languages, and Programming (ICALP 2021).

Editors: Nikhil Bansal, Emanuela Merelli, and James Worrell; Article No. 135; pp. 135:1–135:21

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



So far most of the proposed examples have turned out to have sub-exponential Frege proofs (only a couple of candidate formula classes for the purported separation have been advanced, such as local improvement principles [28], or truncations of the octahedral Tucker lemma [2]). On the other hand many of the tractability results listed above have been obtained using techniques that are highly problem-specific, with relatively little transferability to more general classes of formulas. The existence of such general methods would be highly desirable: such general results could guide the search for examples witnessing the desired separation by pointing to structural properties one needs to avoid in order to construct them.

The purpose of this paper is to present a more general approach for proving sub-exponential upper bounds for the Frege and extended Frege proof complexity of some classes of propositional formulas. We point out that concepts from the theory of parameterized complexity [18], specifically those of *data reduction* and *kernelization* [19] may be relevant to proof complexity as well¹. We give a metatheorem that translates a data reduction for the original problem whose soundness can be witnessed by polynomial size (extended) Frege proofs into subexponential proofs for the corresponding propositional translation of the original problem. The exact size of these proofs is controlled by three factors: the length of the data reduction chain, the nature and size of the proofs witnessing the soundness of the reduction rules, and the size of proofs of unsatisfiability for the formulas in the kernel.

We give several applications of our metatheorem. The use of kernelization techniques does not only allow to tackle the complexity of new problems, but also to improve existing results: In [2] it was shown that propositional formulas $Kneser_{n,k}$ expressing a principle from topological combinatorics known as *the Kneser-Lovász theorem* have, for every fixed value of parameter k , quasipolynomial size Frege proofs. We improve this result by showing quasipolynomial upper bounds for a principle stronger than the Kneser-Lovász theorem, known as Schrijver's theorem. Other applications of our metatheorem concern several (mostly graph-theoretic) problems whose kernelization had previously been studied in the theory of parameterized algorithms. The problems we study are well-known examples satisfying two conditions: First, their negative instances have natural formulations as unsatisfiable CNF formulas. Second, they have efficient kernelizations, often with a small kernel. The problems we have chosen illustrate an important point: *several techniques used in the literature to prove the existence of a kernelization can often be efficiently simulated by (extended) Frege proofs*. Perhaps the most interesting set of applications of our general metatheorem comes, however, from the theory of computational social choice [7]: as it was recently observed, various propositional formalizations of impossibility principles in the theory of computational social choice have computer-assisted proofs that reduce the task of mathematically proving these theorems to the verification of a finite number of cases of unsatisfiability of propositional formulas (using SAT solvers; see [22] for a fairly recent overview)². **We obtain subexponential upper bounds on the complexity of Frege proofs for propositional formulations of Arrow's theorem and the Gibbard-Satterthwaite theorem: quasipolynomial in general, polynomial for a fixed number of agents.**

¹ This is not the first time a connection between parameterized complexity and proof complexity was made; see e.g. [5, 16]. However, our concerns are rather different.

² A similar phenomenon had been independently uncovered for the Kneser-Lovász theorem [2].

2 Preliminaries

We assume basic familiarity with concepts from three distinct areas: proof complexity, parameterized algorithms and computational social choice. We refer the reader to [7, 19, 30] for book-length treatments of these topics. Nevertheless, for purposes of readability we review a couple of relevant notions in the sequel:

► **Definition 1.** A Frege proof system is a sound and complete propositional proof systems having a finite number of axioms and inference rules. An extended Frege proof augments Frege proofs by allowing new variables to substitute complex formulas.

All Frege proof systems are equivalent up to polynomial transformations [15]. Therefore, for concreteness, we will employ a standard “textbook proof style” system having *modus ponens* as the unique inference rule. In both cases we measure the length of a proof by the number of steps in it. Thus the effect of the extension rule in extended Frege proofs is reducing proof length.

We use the shorthand $[m]$ for the set $\{1, 2, \dots, m\}$, $[i : j]$ for $\{i, i + 1, \dots, j\}$, and write $A \cong B$ when sets A, B have the same cardinality. Function $f(\cdot)$ is called *quasipolynomial* if there exists $k > 0$ such that $f(n) = O(2^{O(\log^k(n))})$. We will need the following simple

► **Lemma 2.** Suppose C is a CNF formula and Z_1, \dots, Z_m are literals s.t. $C \wedge (Z_1 \wedge Z_2 \wedge \dots \wedge Z_m)$ is unsatisfiable, as witnessed by a resolution (Frege) proof of length k . Then one can derive from C clause $\overline{Z_1} \vee \overline{Z_2} \vee \dots \vee \overline{Z_m}$ via a resolution (Frege) proof of size at most k .

► **Definition 3 (Parameterized problem).** Let Σ be an alphabet. L is a parametrized problem over Σ^* iff $L \subseteq \Sigma^* \times \mathbb{N}$. Define the support of L , by $\text{supp}(L) = \{x \in \Sigma^* \mid (\exists k \in \mathbb{N}) : (x, k) \in L\}$.

Let L be a parameterized problem in co-NP. Let ϕ be a “canonical” reduction of L to $\overline{\text{SAT}}$. When ϕ is clear from the context, we identify L with the set of pairs $\phi(L) := \{(\phi(x, k), k) : (x, k) \in L\}$, slightly abusing notation, and writing L instead of $\phi(L)$.

► **Example 4 (Graph colorability).** Let $COL = \{(G, i) \mid \chi(G) \leq i\}$. We can encode instances (G, k) of COL as SAT instances $(\phi(G, k), k)$ by the reduction ϕ informally defined by:

- For $v \in V(G)$ and $1 \leq i \leq k$ define boolean $X_{v,i} = \text{TRUE}$ iff v is colored with color i .
- For every pair of distinct vertices $v, w \in G$ we define variable $Y_{v,w}$. The semantics is that $Y_{v,w} = \text{TRUE}$ means that v and w are connected by an edge. Thus, for all sets $\{v, w\}$ that correspond to an edge we add to $\phi(G, k)$ the unit clause $Y_{v,w}$. On the other hand, for sets $\{v, w\}$ that correspond to non-edges we add to $\phi(G, k)$ the unit clause $\overline{Y_{v,w}}$.
- For every $v \in V$ add $X_{v,1} \vee X_{v,2} \vee \dots \vee X_{v,k}$. (“ v must be colored with one of colors 1 to k ”)
- For $v \in V$ and $1 \leq i < j \leq k$ add $\overline{X_{v,i}} \vee \overline{X_{v,j}}$. (“ v cannot be colored with both i and j ”)
- For every set $v, w \in V$ and $i \in 1 \dots k$, add clause $\overline{Y_{v,w}} \vee \overline{X_{v,i}} \vee \overline{X_{w,i}}$. (“if v and w are connected then they cannot both be colored with color i ”)

► **Definition 5.** Define graph, $Kn_{n,k}$, parameterized by an integer $k \geq 1$: The vertex set of $Kn_{n,k}$ is $\binom{n}{k}$, the set of subsets of $\{1, 2, \dots, n\}$ with k elements. Two sets A, B represent adjacent vertices iff $A \cap B = \emptyset$.

The Kneser-Lovász theorem (see e.g. [17]) is a statement about the chromatic number of $Kn_{n,k}$, equivalently restated as $\chi(Kn_{n,k}) > n - 2k + 1^3$. It is expressed as a parameterized problem as follows: $L_{Kn} = \{(Kn_n^k, i) : n \geq 2k > 1, i \leq n - 2k + 1\}$. Note that $L_{Kn} \subseteq$

³ actually $\chi(Kn_{n,k}) = n - 2k + 2$. However, the existence of a $(n - 2k + 2)$ -coloring is easy [17].

\overline{COL} , hence we can use the translation from Example 4 to canonically translate L_{K_n} as a set of unsatisfiable propositional formulas. Similar propositional translations of other constraint satisfaction problems appear e.g. in [25, 33].

The next problem is just the graph coloring problem, but with a different parameterization:

► **Definition 6.** *An instance of the Dual Coloring problem is a pair (G, k) , where G is a graph with n vertices and k is an integer. To decide: is $\chi(G) \leq n - k$? That is, let $DualCol = \{(G, k) : \chi(G) \leq n - k\}$. We have $(G, k) \in DualCol \Leftrightarrow (G, n - k) \in COL$. For this reason the translation of $\overline{DualCOL}$ into SAT modifies the one from \overline{COL} to SAT in Example 4 in an obvious way.*

Given graph G , a *vertex cover* in G is a set $S \subseteq V(G)$ such that for every edge $e = (v, w)$, $v \in S$ or $w \in S$. We denote by $vc(G)$ the size of the smallest vertex cover of G .

► **Example 7 (Vertex Cover).** Let $\overline{VC} = \{(G, i) \mid i < vc(G)\}$ be the set of unsatisfiable instances of Vertex Cover. We can encode (negative) instances (G, k) of VC as instances $\phi(G, k)$ of SAT by the reduction ϕ informally defined as follows:

- For every $v \neq w \in V$, $(v, w) \in E$ add new unit clause $Y_{v,w}$ to the formula. For $(v, w) \notin E$ add new unit clause $\overline{Y_{v,w}}$, to the formula.
- For $v \in V(G)$ and $i \in 1 \dots k$ define boolean variable $X_{v,i}$ with the informal semantics $X_{v,i}$ is TRUE when vertex v is the i 'th vertex in a vertex cover of size k . To encode this semantics add to the formula, for every $v \in V$ and $i \in 1, \dots, k$, clause $\overline{X_{v,i}} \vee (\bigvee_{w \neq v} Y_{v,w})$. This ensures that if v is chosen in the vertex cover then it covers some edge (v, w) . With some extra technical complications one can do away with adding these clauses.
- For every $i = 1, \dots, k$ we add to the formula clause $\bigvee_{v \in V} X_{v,i}$.
- For every $v \neq w \in V$ and $1 \leq i \leq k$ we add to the formula clause $\overline{X_{v,i}} \vee \overline{X_{w,i}}$.
- For every $v \in V$ and $1 \leq i < j \leq k$ we add to the formula clause $\overline{X_{v,i}} \vee \overline{X_{v,j}}$.
- For $v \neq w \in V$ add to the formula clause $\overline{Y_{v,w}} \vee X_{v,1} \vee \dots \vee X_{v,k} \vee X_{w,1} \vee \dots \vee X_{w,k}$.

► **Definition 8 (Kernelization).** *Let L be a parametrized problem. A kernelization algorithm (or, shortly, kernelization) Ker for the problem L is an algorithm that works as follows: on input (x, k) , Ker outputs (in time polynomial in $|(x, k)|$) a pair (x', k') , such that the following are true: $(x, k) \in L$ iff $(x', k') \in L$, and $|x'|, k' \leq g(k)$, where g is a computable function. Pair (x', k') is called the kernel of (x, k) , while $g(k)$ is called the size of the kernel.*

One can convert a kernelization into an algorithm by solving kernel instances by other means (e.g. brute force). A kernelization is often the reflexive, transitive closure of a finite set of *data reduction rules*: we apply the rules as long as possible, until we are left with an instance, the kernel, to which no rule can be applied anymore.

► **Definition 9 (Data reduction rule).** *Let L be a parameterized problem. A **data reduction rule for L** is an algorithm A that maps (in time polynomial in $|x| + k$) an instance (x, k) of L to an instance (x', k') such that $(x, k) \in L$ iff $(x', k') \in L$ (we say that the two instances are equivalent, or that the reduction rule is safe), and $|x'| \leq |x|$. In practice, a data reduction rule may be well-defined only for $|x| \geq f(k)$, for some function $f(\cdot)$, as we can simply extend it to smaller instances (x, k) by defining $A(x, k) = (x, k)$. All kernelizations in this paper have this nature, and we will assume this to be true for all the results we give in the sequel.*


► **Definition 10 (Data reduction chain).** *Given parameterized problem L kernelizable via data reductions (A_1, A_2, \dots, A_r) , a **data reduction chain** for instance (x, k) of L is a sequence $(x_0, k_0), (x_1, k_1), \dots, (x_m, k_m)$, where $(x_0, k_0) = (x, k)$, $A_t(x_m, k_m) = (x_m, k_m)$, for all $t = 1, \dots, r$ and, for all $i = 1, \dots, m$ there exists $j \in 1, \dots, r$ such that $(x_i, k_i) = A_j(x_{i-1}, k_{i-1})$.*

► **Example 11** (Data reduction for Kneser instances:). Reductions $(Kn_n^2, a) \rightarrow (Kn_{n-1}^2, a - 1)$ and $(Kn_n^3, a) \rightarrow (Kn_{n-1}^3, a - 1)$ were used in [27] to give polynomial size extended Frege upper bounds for Kneser formulas for $k = 2, 3$.

- For $k \geq 2$ there exists $N(k) \leq k^4$ such that for $n > N(k)$ $(\mathcal{K}_n^k, a) \rightarrow (\mathcal{K}_{n-1}^k, a - 1)$. This was used in [2] to give polynomial size extended Frege upper bounds for Kneser formulas.
- For $k \geq 2$ there exists $N(k) \leq k^4$ such that for $n > N(k)$ $(Kn_n^k, a) \rightarrow (Kn_{n-\frac{n}{2k}}^k, a - \frac{n}{2k})$. This was used in [2] to give quasipolynomial size Frege proofs for Kneser formulas.

Our results will work by formalizing *data reductions* which solve the decision for an instance X_1 of a parameterized problem by reducing it to a smaller instance X_2 of the same problem. This will map, propositionally, to witnessing the soundness of an implication $\Phi_1 \vdash \Phi_2$. However, to make Φ_2 a propositional translation of an instance of the same problem, we will need to use variable substitutions $Y = \Xi[X]$. That is, the reduction we witness is actually $\Phi_1[X] \vdash \Phi_2[Y]$. When we will refer to the proof complexity of witnessing an implication $\Phi_1 \vdash \Phi_2$ using (extended) Frege proofs, what we mean is that one can derive the clauses of $\Phi_2[Y]$ using the clauses of $\Phi_1[X]$ as axioms. Since they use substitutions, these are extended Frege proofs. To convert them into Frege proofs one needs to unwind the definitions of newly defined variables.

► **Definition 12.** A *crown decomposition of a graph G* (see e.g. Fig. 1 (b).) is a decomposition of $V(G)$ into three subsets C, H, R , $C \neq \emptyset$ such that (1). C is an independent set. (2). No vertex in C is adjacent to a vertex in R . (3). There exists a matching of H in C , i.e. a set of disjoint edges covering H with the other endpoint in C .



■ **Figure 1** (a). The Kneser graph $Kn_{5,2}$. (b). A crown decomposition of a graph.

Given a set S and $T \subseteq S$, we will denote, as in [37,38], by S_{-T} the set $S \setminus T$. We will also write S_{-a} instead of $S_{-\{a\}}$. When $S = [m]$, of course $[m]_{-T} \cong [m - |T|]$ for every $T \subseteq [m]$.

► **Definition 13.** Given a set of m objects, identified with the set $[m]$, a preference profile is a linear ordering of $[m]$, i.e. a permutation $\pi \in S_m$. Given $a, b \in [m]$ we say that a is preferred to b (written $a <_{\pi} b$) iff $\pi^{-1}(a) < \pi^{-1}(b)$. **Note that preferred objects are lower in the ordering.** We denote by $top(\pi)$ the object $\pi^{-1}(1)$, i.e. the object that is preferred in π to all others. Given a preference profile π and $T \subseteq [m]$, denote by π_{-T} the restriction of π to $[m]_{-T}$, and by π^{+T} the preference profile derived from π by making all elements $a \in T$ less preferred than any other $b \in [m]$ (with an arbitrary fixed order among them, e.g. the order induced on $[m]$ by the identical permutation).

► **Definition 14.** Given a set of m objects, identified with the set $[m]$ and a set of n agents, a **social choice function (SCF)** is a mapping $s : S_m^n \rightarrow Z$. Z is a set equal to S_m (for Arrow’s theorem) and to $[m]$ (for the Gibbard-Satterthwaite theorem). A SCF is **dictatorial**

if there exists $i \in [m]$ such that for all $R_1, R_2, \dots, R_n \in S_m$, $s(R_1, R_2, \dots, R_n) = R_i$ ($s(R_1, R_2, \dots, R_n) = \text{top}(R_i)$ for the Gibbard-Satterthwaite theorem). A SCF is **unanimous** if whenever a is preferred to b in all profiles R_1, R_2, \dots, R_n then a is preferred to b in profile $s(R_1, R_2, \dots, R_n)$. SCF s satisfies the **independence of irrelevant alternatives (IIA) axiom** (for Arrow's theorem) if whenever $a, b \in [m]$ are two different objects and $(R_1, R_2, \dots, R_n) \in S_m^n$ and $(R'_1, R'_2, \dots, R'_n) \in S_m^n$ are two vectors of preference profiles such that, for all $i = 1, \dots, m$, R_i and R'_i agree in their relative preference of a or b , then $s(R_1, R_2, \dots, R_n)$ and $s(R'_1, R'_2, \dots, R'_n)$ agree in their relative preference of a or b . A SCF is **onto** iff it is onto as a function. Finally, for every pair (R, o) , $R = (R_1, \dots, R_n)$ and player $1 \leq i \leq m$, denote by $\text{pr}(i, o, R)$ the set of objects o' s.t. $R_i^{-1}(o) \leq R_i^{-1}(o')$ (i.e. i weakly prefers o to o' in R_i). A SCF s is **strategyproof** (for the Gibbard-Satterthwaite theorem) iff, for every strategy profile R , if o is the outcome of preference profile R then i cannot misrepresent its preferences as $\pi \in S_m, \pi \neq R_i$ so that the social choice for the resulting profile $s(i, R, \pi)$ is an o' that i strictly prefers to o .

Given an SCF $W : [m]^n \rightarrow Z$ and $B \subseteq [m]$ we define function $W_{-B} : [m]_{-B}^n \rightarrow Z$ to be defined as follows: $W_{-B}(R_1, R_2, \dots, R_n) = W(R_1^{+B}, R_2^{+B}, \dots, R_n^{+B})_{-B}$. In other words, we extend profiles R_1, R_2, \dots, R_n by making objects in B less preferred than all other objects, apply W on the resulting profiles, then drop objects from B from the result.

3 Main (Meta)Theorem and Applications

In the next definition we formalize the complexity of simulating data reduction steps by (extended) Frege proofs. Clearly, we want to encode the scenario where each such step can be simulated by efficient proofs. Our main result will allow a slightly more general setting, where the safety of each reduction step can be established by a “case by case argument with a limited number of cases”. This will lead not to a chain but to a *tree* of logical reductions:

► **Definition 15.** Given reduction rule \mathcal{A} for problem L and function $h(\cdot)$, **the soundness of \mathcal{A} has (extended) Frege proofs of size $h(\cdot)$** iff there is an integer $R \geq 1$ s.t. for every $(x, k) \notin L$ and every step $(x_i, k_i) \rightarrow (x_{i+1}, k_{i+1})$ in the reduction chain the following are true:

- There exists $r'_i \leq R$, tautology $\Xi_i := \bigvee_{t=1}^{r'_i} \Xi_{i,t}$ and formulas $\eta_{i,1}, \dots, \eta_{i,r'_i}$ isomorphic (up to a variable renaming) to $\Phi(x_{i+1}, k_{i+1})$ s.t. for $t = 1, \dots, r'_i$, $\Phi(x_i, k_i) \wedge \Xi_{i,t} \vdash \eta_{i,t}$.
- Proving the soundness of Ξ_i and of all reductions $\Phi(x_i, k_i) \wedge \Xi_{i,t} \vdash \eta_{i,t}$ can be accomplished by (extended) Frege proofs of **total** size at most $h(|\Phi(x, k)|)$.

In other words, the kernelization may encode a case-by-case construction with a constant number, at most R , of cases. Each case will lead to a different reduced formula. The consistency of all those reductions is witnessed by proofs with a common polynomial upper bound. Given this definition, our main (meta)theorem is:

► **Theorem 16.** Let L be a parameterized problem that is kernelizable via a finite number of data reduction rules (A_1, A_2, \dots, A_r) with kernel size $g(\cdot)$.

1. Assume that negative instance (x, k) of L has a data reduction chains of length $C(x, k)$, and that the soundness of each reduction rule A_1, A_2, \dots, A_r can be witnessed using extended Frege proofs of size at most $h(|\Phi(x, k)|)$, for some function $h(\cdot)$. Then L has extended Frege proofs of size $O((\sum_{i=0}^{C(x,k)} R^i)[h(|\Phi(x, k)|) + 2^{O(\text{poly}(g(k)))}])$. In particular, if $R = 1$ and for every fixed k we have $C(x, k) = O(\text{poly}(|\Phi(x, k)|))$ then, for every fixed k , negative instances $\Phi(x, k)$ of L have extended Frege proofs of size polynomial in $|\Phi(x, k)|$.

2. Assume that negative instances (x, k) of L have data reduction chains of length $C(x, k) = O(1)$ ($O(\log(|\Phi(x, k)|))$), respectively, where the constant may depend on k , and that the safety of each reduction $\Phi(x_i, k_i) \vdash \Phi(x_{i+1}, k_{i+1})$ is witnessed by Frege proofs of size $\leq p(|\Phi(x, k)|)$, for some fixed polynomial $p(\cdot)$. Then for every fixed k , negative instances $(\Phi(x, k), k)$ of L have Frege proofs of size polynomial (quasipolynomial) in $|\Phi(x, k)|$.

Proof.

1. Given an instance (x, k) of L , and data reduction chain $(x, k) = (x_0, k_0), (x_1, k_1), \dots, (x_m, k_m)$, an extended Frege proof for $\Phi(x, k)$ is obtained by concatenating the proofs for statements $\Phi(x_{i-1}, k_{i-1}) \vdash \Phi(x_i, k_i)$ with an extended Frege proof of the kernel instance. There is one complication, though, induced by the fact that we allow at most R cases in the reduction: the reduction chain maps to a **tree** of propositional proofs, since for each node $\Phi(x_i, k_i)$ we have $r'_i \leq R$ children $\Xi_{i,t}$, all isomorphic to $\Phi(x_{i+1}, k_{i+1})$ (but different). The total number of nodes in this tree is at most $\sum_{t=0}^{C(x,k)} R^t$.

Then the whole chain of reductions from $\Phi(x, k)$ to $\Phi(x_m, k_m)$ can be proved to be sound by proofs of length $(\sum_{t=0}^{C(x,k)} R^t) \cdot h(|\Phi(x, k)|)$. There are at most $R^{C(x,k)}$ copies of the kernel instance. Each of them can be proved (in brute force) in size $O(2^{|\Phi(x_m, k_m)|}) = O(2^{\text{poly}(g(k))})$, since $(x_m, k_m) \in \ker(L)$ and any unsatisfiable formula Ξ with n variables has Frege proofs of size $O(2^n)$.

The length of the total proof is thus $O(\sum_{t=0}^{C(x,k)} R^t \cdot [h(|\Phi(x, k)|) + 2^{\text{poly}(g(k))}])$. We infer the desired result when $C(x, k) = O(\text{poly}(|\Phi(x, k)|))$.

2. We unwind the substitutions implicit in the extended Frege proofs. For $R = 1$ (i.e. a reduction chain), arguing that the blow-up due to making substitutions is quasipolynomial as long as the chain length is logarithmic is identical to similar arguments made in [9], [2] for other problems, and we omit further details.

In our case the complication arises since we no longer have a chain but a tree. However, we can upper bound the complexity of Frege proofs by $R^{C(x,k)}$ times the complexity of a single chain (a root-to-leaf path in this tree). As long as $C(x, k) = O(\log(|\Phi(x, k)|))$, the term $R^{C(x,k)}$ has a magnitude polynomial in $|\Phi(x, k)|$. Multiplying this polynomial by the quasipolynomial complexity of each chain still yields a proof of complexity quasipolynomial in $|\Phi(x, k)|$. ◀

► **Observation 17.** *There is an important uniformity aspect of kernelization that we haven't used in the preceding proof: the fact that data reductions are specified by polynomial time algorithms. This issue will be important in applying the above result: often the existence of a data reduction is proved by an algorithm whose soundness (for all instances) would be rather cumbersome to simulate in propositional proofs. This is the case when results involve general techniques for developing kernelizations, such as the Crown Decomposition Lemma or the Sunflower Lemma. As long as we do not insist, however, on actually generating the proof, but merely on proving its existence, we can get away with proving the soundness of individual instances. That is, if we can prove the soundness of an individual application of a propositional reduction rule, $\Phi(x, k) \vdash \Phi(x', k')$, taking for granted the existence/definition of (x', k') , we can prove the existence of efficient proofs, **without actually having to generate a propositional proof of the soundness of the reduction techniques.***

Next we highlight some applications of our main (meta)theorem:

3.1 Proof Complexity of (Dual) Coloring

► **Theorem 18.** *There exists a kernelization that reduces instances (G, k) of DUALCOL to a kernel of size at most $3k - 2$. The length of reduction chain in this kernelization is $O(k)$. The soundness of each reduction step can be witnessed by polynomial size Frege proofs. Hence, for every fixed k , negative instances (G, k) of DUALCOL have Frege proofs of size polynomial in $|\Phi_{G,k}|$.*

Proof. The kernelization is a variant of the classical one from the parameterized complexity literature, based on *crown decompositions* (Definition 12). It consists of three data reductions:

- (a). Let $All(G)$ be the set of vertices v adjacent to all other vertices in G . If $All(G) \neq \emptyset$ then $(G, k) \in DualCol \Leftrightarrow (G \setminus All(G), k - |All(G)|) \in DualCol$.
- (b). If $All(G) = \emptyset$ but \overline{G} has a matching of size k , $x_1, y_1, \dots, x_k, y_k$, with x_i being matched to y_i for $i = 1, \dots, k$, then $(G, k) \in DualCol$ (so reduce it to an arbitrary positive instance).
- (c). Assume that rules (a),(b) do not apply. Let (C, H, R) be a crown decomposition of the graph \overline{G} . Reduce (G, k) to (G', k') , by deleting $H \cup C$ from G and $k' = k - |H|$.

Without loss of generality, we will only apply rule (c). to crown decompositions where $|H| \neq \emptyset$ and all nodes in C are matched to some node in H . This is possible for the following reason: if $|H| \neq \emptyset$ and the original crown decomposition had other vertices in C , just move them to R . If, on the other hand $|H| = \emptyset$ then all vertices in C would be connected to all vertices in $C \cup R$, hence to all vertices of G . But this cannot happen, since the case $All(G) \neq \emptyset$ is covered by the first data reduction rule.

The Crown Decomposition Lemma (Lemma 4.5 of [19]) makes sure that at least one of reduction rules (a),(b),(c) applies to every graph with more than $3k - 2$ vertices.

The safety of reduction rule (c) can be informally justified as follows: since vertices in a crown decomposition of \overline{G} are matched in a matching m , vertices $v \in H$ and $m(v) \in C$ are not connected in G , hence they can be colored with the same color. At the same time, $m(v)$ is connected in G to all the vertices of G' , hence must assume a color different from all the colors of vertices of G' . Also $m(v_1)$ and $m(v_2)$ are connected, so must assume distinct colors. In conclusion, vertices of C must use $|C|$ different colors, and G is $n - k$ colorable if and only if G' is $n - k - |C|$ colorable. But $|G'| = n - 2|C|$, so G is $n - k$ colorable if and only if G' is $|G'| - (k - |C|)$ colorable.

Rule (b). does not apply to unsatisfiable instances of DualCol. Hence we have to argue about the size of Frege proofs witnessing the soundness of rules (a) and (c), namely: Let $\Phi_{v,1}[\overline{Y}]$ be the formula $\bigwedge_{w \neq v} Y_{v,w}$ (informally, $v \in All(G)$). We need to provide proofs that witness that

$$\Phi(G, k) \wedge \Phi_{v,1}[\overline{Y}] \vdash \Phi(G \setminus \{v\}, k - 1), \text{ and}$$

$$\Phi(G, k) \vdash \Phi(G', k').$$

For the first implication, define new variables $Z_{w,i}$ via the substitution, for $w \neq v \in V(G)$, $Z_{w,i'} \leftrightarrow X_{w,i} \wedge X_{v,l}$, where $i' = i$ for $i < l$, $i' = i - 1$ for $i > l$.

We start by deriving, by resolving unit literals $Y_{v,w}$ (which are part of the formula), for all $w \neq v \in V(G)$ and i , clauses $\overline{X_{v,i}} \vee \overline{X_{w,i}}$. Then we derive, for every $w \neq v \in V(G)$ and i , clauses $\overline{X_{v,i}} \vee (\bigvee_{j \neq i} X_{w,j})$. This is done by resolving $(\bigvee_{j=1}^k X_{w,j})$ and $\overline{X_{v,i}} \vee \overline{X_{w,i}}$. We then derive clauses $\overline{X_{v,i}} \vee (\bigvee_{i'=1}^{k-1} Z_{w,i'})$. By resolving all these clauses against $\bigvee_{i=1}^k X_{v,i}$ we derive $(\bigvee_{i'=1}^{k-1} Z_{w,i'})$. Similar tricks allow deriving clauses $\overline{Z_{w,i}} \vee \overline{Z_{w,j}}$ and $\overline{Y_{v,w}} \vee \overline{Z_{v,i}} \vee \overline{Z_{w,i}}$ from the corresponding clauses in the X variables.

As for the second reduction rule, intuitively we want to encode the fact that if a vertex $v \in G'$ is colored with color i in G , then coloring it with color $i - less(v)$, where $less(v)$ is the number of nodes in H colored with a color smaller than i , yields a legal coloring of G' . This is true since all nodes in H must get colors (in G) different from all colors in G' .

For an arbitrary vertex $v \in G'$, let $less(v)$ be the number of nodes $w \in C$ (here C refers to the class of the crown decomposition of \overline{G}) such that $col(w) < col(v)$. One can compute the binary representation of number $less(v)$ using Frege proofs as follows: we create a boolean variable $T_{v,w}$ which will be true iff $col(w) < col(v)$. One can compute $T_{v,w}$ as

$$T_{v,w} := \bigvee_{i < j} X_{w,i} \wedge X_{v,j}.$$

Now we simply use the predicate $COUNT(T_{v,w})_{w \in C}$ to compute the binary representation of $less(v)$. Here COUNT is the Buss counting predicate [8]. We will also derive the following formulas:

$$\overline{X_{v,i}} \vee \bigvee_{t=1}^i [less[v] = t] \quad (1)$$

To accomplish that, we use the pigeonhole principle PHP_i^{i+1} to prove that

$$\overline{X_{v,i}} \vee [COUNT((X_{w,j})_{w \in C, j < i}) \leq i] \quad (2)$$

Indeed, assuming $\overline{X_{v,i}} = TRUE$ we can derive any disjunction of length $i + 1$ consisting of literals of type $\overline{X_{w,j}}$, with $w \in C$, $j < i$. This is because for all $w_1 \neq w_2 \in C$, $k_1 \neq k_2$ $\overline{X_{w_1,k_1}} \vee \overline{X_{w_2,k_1}}$ and $\overline{X_{w_1,k_1}} \vee \overline{X_{w_1,k_2}}$ are clauses of $\Phi(G, k)$. By Lemma 2 we can derive equation (2). Next, simple arguments along the lines of [8] establishes the equivalence between formulas $U \leq i$ and $\bigvee_{k=1}^i [U = k]$. Here U is a bit vector of appropriate length to represent i . We use (2) and this to derive (1).

Now, for every $v \in V(G') = R$ we define a new variable $Z_{v,i}$, designed to be true iff the color of v in the induced coloring on G' is j . We will enforce this by making the substitutions

$$Z_{v,j} := \bigvee_{i=1}^j X_{v,i} \wedge [less(v) = i - j] \quad (3)$$

First note that Z respects the color classes of G : if v_1, v_2 have the same color in G then they have the same color in G' . Furthermore, the substitution does not collapse two different color classes of G into a single color class in G' : it simply relabels the colors of vertices in G' with elements of $1, 2, \dots, k'$. Therefore, if $Z_{v_1,j} = Z_{v_2,j} = TRUE$ then there exists a unique i_0 such that $X_{v_1,i_0} = X_{v_2,i_0} = TRUE$.

We need to derive clauses $\bigvee_{t=1}^{k'} Z_{v,t}$ as well as, for $vw \in E(G')$, $\overline{Z_{v,j}} \vee \overline{Z_{w,j}}$. Deriving the first type of clauses is easy: we use formulas (1) and $X_{v,1} \vee X_{v,2} \vee \dots \vee X_{v,k}$.

As for the second one, note that all clauses $\overline{X_{v,i}} \vee \overline{X_{w,i}}$ are part of $\Phi(G, k)$. Given the observation we made above and this fact, assuming $Z_{v,j} = Z_{w,j} = TRUE$ we can derive a contradiction. By Lemma 2 we can, therefore, derive (with the same complexity) clause $\overline{Z_{v,j}} \vee \overline{Z_{w,j}}$.

◀

3.2 Proof Complexity of Schrijver's Theorem

In this section we deal with the proof complexity of a stronger version of the Kneser-Lovász theorem known as Schrijver's Theorem [35]. This is a statement about the chromatic number of the so-called *stable Kneser graph* $SK_{n,k}$, defined as follows:

► **Definition 19.** Call a set $A \subseteq \binom{[n]}{k}$ stable if A does not contain two elements that are consecutive (we also consider n and 1 as consecutive). Denote the set of stable sets by $\binom{[n]}{k}_{st}$. The stable Kneser graph $SKn_{n,k}$ is the subgraph of $Kn_{n,k}$ induced by the set $\binom{[n]}{k}_{st}$.

Schrijver's theorem asserts that the chromatic number of the stable Kneser graph $SKn_{n,k}$ is $n - 2k + 2$. We are, of course, interested mainly in the harder part of this result, the lower bound $\chi(SKn_{n,k}) > n - 2k + 1$. Since $SKn_{n,k}$ is the subset of the Kneser graph (see e.g. Figure 1, where the central star is the stable Kneser graph $SKn_{5,2}$), this strengthens the (harder part of the) Kneser-Lovász theorem. The propositional translation of Schrijver's theorem is immediate, and the resulting unsatisfiable formulas, that we will denote by $Schrijver_{n,k}$ are subformulas of formulas $Kneser_{n,k}$. We have the following:

► **Theorem 20.** For every fixed k , formulas $Schrijver_{n,k}$ have Frege proofs of size quasipolynomial in n .

For the (somewhat more involved) proof of this theorem, we refer the reader to [26].

3.3 Buss Meets Buss: the Proof Complexity of Vertex Cover

In this subsection we study the proof complexity of Vertex Cover, the “drosophila of parameterized complexity” [20]. We apply our result to a variation of the standard kernelization of VC (called in [20] *the Buss reduction*, hence the title of this subsection) to prove:

► **Theorem 21.** Instances (G, k) of VC have a kernelization with a data reduction chain of length $O(k)$ to a kernel with at most k^2 vertices. The soundness of each step in this data reduction can be witnessed by Frege proofs of size polynomial in $|\Phi(G, k)|$. Hence, for every fixed k negative instances $\Phi(G, k)$ of VC have Frege proofs of size polynomial in $|\Phi(G, k)|$.

Proof. Informally, we will use the following two data reduction rules:

- (a). if G has a vertex v of degree larger than k then G has a VC of size $\leq k$ if and only if $G \setminus \{v\}$ has a VC of size $\leq k - 1$. Indeed, v must be part of any VC of G of size $\leq k$.
- (b). if $Isolated(G)$ denotes the set vertices v in G that are isolated then G has a VC of size $\leq k$ iff $G \setminus Isolated(G)$ has a VC of size $\leq k$.

The kernel of these two reduction rules, the set of instances (G, k) of VC such that none of the two rules applies is composed of graphs of at most k^2 vertices only [19].

To encode the soundness of these rules by polynomial-size Frege proofs we use the predicate $COUNT_k^n(x_1, x_2, \dots, x_n)$ from [8]. Formula $COUNT_k^n(x_1, x_2, \dots, x_n)$ is TRUE if and only if at least k of the variables x_1, x_2, \dots, x_n are true. For every fixed k , $COUNT_k^n$ can be computed by polynomial size Frege proofs.

We will define a sequence of formulas:

1. For $v \in V$, $\Phi_{v,1}(\bar{Y}) = COUNT_k^{n-1}((Y_{v,w})_{w \neq v \in V})$. Informally, $\Phi_{v,1}$ is true in graph G iff the degree of v is at least k .
2. For $v \in V$, $\Phi_{v,2}(\bar{X}, \bar{Y}) = (\bigwedge_{i=1}^k \overline{X_{v,i}}) \wedge \Phi_{v,1}(\bar{Y}) \wedge \Phi_{VC}(G, k)[\bar{X}, \bar{Y}]$.

For every neighbor w of v , by resolving $Y_{v,w}$ with clause $\overline{Y_{v,w}} \vee X_{v,1} \vee \dots \vee X_{v,k} \vee X_{w,1} \vee \dots \vee X_{w,k}$ of $\Phi_{v,2}$ we derive clause $X_{v,1} \vee \dots \vee X_{v,k} \vee X_{w,1} \vee \dots \vee X_{w,k}$. By resolving successively with $\overline{X_{v,1}}, \dots, \overline{X_{v,k}}$ we derive clause $X_{w,1} \vee \dots \vee X_{w,k}$.

Formula $\bigwedge_{w \in N(v)} (X_{w,1} \vee \dots \vee X_{w,k})$ is isomorphic to the Pigeonhole Principle $PHP_{|N(v)|}^k$ which has polynomial-size Frege refutations [8]. Plugging in this proof of this statement into our argument, we conclude that the implication $\Phi_{v,2}(X, \bar{Y}) \vdash \square$ can be witnessed by polynomial size Frege proofs, hence, by Lemma 2, so does the implication $\Phi_{v,1}(\bar{Y}) \wedge \Phi_{VC}(G, k)[\bar{X}, \bar{Y}] \vdash \bigvee_{i=1}^k X_{v,i}$.

As for the second reduction rule, it is just as easy: for every vertex $v \in V$ which is isolated and every $i = 1, \dots, k$, we first derive by resolution (using negative clauses $\overline{Y_{v,w}}$ and clause $\overline{X_{v,i}} \vee (\bigvee_{w \neq v} Y_{v,w})$) unit clauses $\overline{X_{v,i}}$. We then use these clauses to resolve away every other occurrence of $X_{v,i}$ from the formula, obtaining a formula isomorphic to $\Phi(G \setminus \text{Isolated}(G), k)$. \blacktriangleleft

3.4 Proof Complexity of Edge Clique Cover

In this section we study the proof complexity of the following problem:

► **Definition 22** (Edge Clique Cover). *Given graph G and integer k , to decide is whether one can find sets of vertices $V_1, V_2, \dots, V_k \subseteq V$ s.t. each V_i induces a clique, and for every edge $e = (v, w) \in E$ there exists $1 \leq i \leq k$ s.t. $v, w \in V_i$ (“each edge is covered by some clique”). We represent instance (G, k) of Edge Clique Cover by propositional formula $\Phi_{G,k}$ as follows:*

- For every pair of distinct vertices $v, w \in V$ define a variable $Y_{v,w}$. For every edge $(v, w) \in E(G)$ add unit clause $Y_{v,w}$. For $(v, w) \notin E(G)$ add unit clause $\overline{Y_{v,w}}$.
- For $v \in V$ and $1 \leq i \leq k$ define boolean variable $X_{v,i} = \text{TRUE}$ iff $v \in V_i$.
- For $v, w \in V$ and $1 \leq i \leq k$ add $\overline{X_{v,i}} \vee \overline{X_{w,i}} \vee Y_{v,w}$ (“if $v, w \in V_i$ then $vw \in E(G)$ ”) and $\overline{Y_{v,w}} \vee (\bigvee_{j=1}^k (X_{v,j} \wedge X_{w,j}))$. Of course, as written above the latter formula is not CNF, but it can be converted easily by expanding the last disjunction.

The following is our result for the Edge Clique Cover problem. The main technical novelty is reducing the length of the data reduction chain (compared to the usual kernelization) from linear to logarithmic, so that we can get quasipolynomial-size Frege proofs:

► **Theorem 23.** *There exists a kernelization that reduces instances (G, k) of problem EDGE CLIQUE COVER with graph G having n vertices to a kernel with at most 2^k nodes. The length of the data reduction chain is $O(\log_{1+\frac{1}{2^k-1}}(n))$. The soundness of each reduction step can be witnessed by polynomial size Frege proofs. Consequently, for fixed k , negative instances (G, k) of EDGE CLIQUE COVER have extended Frege proofs of polynomial size and Frege proofs of quasipolynomial size in $|\Phi_{G,k}|$.*

Proof. We use the following data reduction rules:

- (a). If $|\text{Isolated}(G)| \geq \frac{n}{2^k}$ then reduce (G, k) to $(G \setminus \text{Isolated}(G), k)$.
- (b). If there exists a set $S \subseteq V$, $|S| \geq \frac{n}{2^k}$ such that vertices in S induce a clique in G , and for all $v, w \in S$ we have $N[v] = N[w]$, where $N[v]$ stands for the closed neighborhood of v , then reduce G to (G', k') , where G' is the graph obtained by identifying vertices v, w , and $k' = k$ whenever $N[v] = N[w] \neq \emptyset$, $k' = k - 1$, otherwise.

The soundness of the first reduction rule, $\Phi(G, k) \vdash \Phi(G \setminus \text{Isolated}(G), k)$ can be witnessed by efficient Frege proofs similar to those for the vertex cover problem.

As for the second rule, the formula

$$\Xi_S(G) := \bigwedge_{w, v \in S} \bigwedge_{r \in V} (Y_{v,r} \leftrightarrow Y_{w,r})$$

135:12 Kernelization, Proof Complexity and Social Choice

(where, of course, $A \leftrightarrow B$ can be equivalently rewritten as $(\bar{A} \vee B) \wedge (A \vee \bar{B})$) expresses the fact that $N[v] = N[w]$ for all $v, w \in S$. So we need to prove the soundness of the rule

$$\Phi(G, k) \wedge \Xi_S(G) \vdash \Phi(G', k'). \quad (4)$$

Without loss of generality we will only deal with the case $N[v] \neq \emptyset$ for all $v \in S$, (the other case, $N[v] = \emptyset$ for all $v \in S$, can be handled with minor modifications to this argument). By slightly abusing notation, we will denote by S the vertex of G' obtained by contraction. Let $s \in S$ be an arbitrary vertex.

We define substitutions: $Y'_{v,w} := Y_{v,w}$ for all $v, w \in G'$, $v, w \neq S$. If, say, $v = S$ we define $Y'_{S,w} := Y_{s,w}$. Also define $X'_{v,i} := X_{v,i}$ for $v \neq S$, $X'_{S,i} := X_{s,i}$. The substitution yields a formula isomorphic to $\Phi(G', k')$, and the proof of the safety is basically trivial.

To obtain the result note that the number of vertices goes down geometrically, by a ratio of $1 - \frac{1}{2^k}$ at each step.

► **Lemma 24.** *Rules (a). (b). are safe. Also, for every graph G with $n > 2^k$ vertices one of rules (a). (b). applies.*

Proof. Let G be a graph to which rules (a). (b). do not apply and which has an edge clique cover of size k . Consider an encoding $b(v)$ of every vertex v on k bits such that for every $v \in V$, $b(v)$ is a bit vector whose i 'th bit is one iff v is a part of the i 'th clique.

There must be a set of vertices $S \subset V$, $|S| \geq \frac{n}{2^k}$ such that for all $u, v \in S$, $b(u) = b(v) = b$, for some $b \in \{0, 1\}^k$.

If $b = 0^k$ then, since every edge in G must be covered by one of the k cliques, it follows that every $v \in S$ is an isolated vertex. Hence rule (a). applies.

If, on the other hand $b \neq 0^k$, say $b_i \neq 0$, then every $v \in S$ must belong to the i 'th clique. Hence S induces a clique in G . ◀

◀

3.5 Proof Complexity of the Hitting set problem

In the d -Hitting Set problem we are given an universe U and a family \mathcal{A} of subsets of U , all of cardinality at most d , as well as an integer k . To decide is whether there exists a set $H \subseteq U$ containing at most k elements, such that H intersects every $P \in \mathcal{A}$.

A formalization of the d -Hitting set problem as an instance of SAT is obtained as follows:

- **Example 25.** Let $P = (U, \mathcal{A}, k)$ be an instance of d -Hitting set. Define formula Φ_P by:
 - For $i \in U$, $j = 1, \dots, k$ add variable $X_{i,j}$, TRUE iff i is the j 'th chosen element.
 - For $i \neq i' \in U$, $1 \leq j \leq k$ add clauses $\bigvee_{i \in U} X_{i,j}$ ("some i is the j 'th chosen element") and $\overline{X_{i,j}} \vee \overline{X_{i',j}}$ ("at most one i can be the j 'th chosen element").
 - For $A \in \mathcal{A}$ add $(\bigvee_{i \in A} (\bigvee_{j=1, \dots, k} X_{i,j}))$ ("some element of A is among the k chosen elements")

Our result, which only guarantees polynomial size *extended* Frege proofs, is:

► **Theorem 26.** *There exists a kernelization mapping instances (U, \mathcal{A}, k) of d -HittingSet with $|U| = n$ elements to a kernel with at most $d \cdot d! \cdot k^d$ sets (hence at most $d^2 \cdot d! \cdot k^d$ elements). The data reductions chains in this kernelization have length $O(n^d/k)$, and their soundness can be witnessed by polynomial-size Frege proofs. Hence for every fixed k, d , unsatisfiable instances $\Phi_{(U, \mathcal{A}, k)}$ of d -HittingSet have extended Frege proofs of size $O(\text{poly}(|\Phi_{(U, \mathcal{A}, k)}|))$.*

Proof. We employ the standard kernelization of d -Hitting set based on sunflowers [19]:

► **Definition 27.** *A sunflower with k -petals and core Y is a collection of sets S_1, \dots, S_k , all different from Y , such that for $1 \leq i < j \leq k$, $S_i \cap S_j = Y$.*

We are going to propositionally encode the following informally stated data reduction rule: let (U, \mathcal{A}, k) be an instance of the d -Hitting set such that \mathcal{A} contains a sunflower $\mathcal{S} = \{S_1, S_2, \dots, S_{k+1}\}$ of cardinality $k + 1$ with core Y . We reduce (U, \mathcal{A}, k) to the instance (U', \mathcal{A}', k) , where $\mathcal{A}' = (\mathcal{A} \setminus \mathcal{S}) \cup \{Y\}$ and $U' = \cup_{X \in \mathcal{A}'} X$. Indeed, consider a hitting set H for (U, \mathcal{A}, k) . By definition, H meets every element of $\mathcal{A} \setminus \mathcal{S}$. If H did not meet Y then it would have to meet each of the $k + 1$ disjoint petals $S_j \setminus Y$. Hence $|H| \leq k$ iff H meets Y .

To simulate this argument propositionally, define for $i \in U'$ substitutions $X'_{i,j} := X_{i,j}$. We need to (a). derive clause $(\bigvee_{i \in Y} (\bigvee_{j=1, \dots, k} X_{i,j}))$ (b). for every $j = 1, \dots, k$, derive clauses $\bigvee_{i \in U'} X_{i,j}$.

For the first clause we show that $\Phi(U, \mathcal{A}, k) \wedge (\bigwedge_{i \in Y} (\bigwedge_{j=1, \dots, k} \overline{X_{i,j}})) \vdash \emptyset$ and then we invoke Lemma 2. To do that we first derive, for $l = 1, \dots, k + 1$ (by resolving literals $\overline{X_{i,j}}$) clauses $(\bigvee_{i \in S_l \setminus Y} (\bigvee_{j=1, \dots, k} X_{i,j}))$.

Substituting each variable $X_{i,j}$, where $i \in S_j \setminus Y$ to a new variable $Y_{i,j}$ yields a formula isomorphic to PHP_{k+1}^k which has polynomial size Frege proofs [8]. Putting all these things together we get polynomial-size Frege proofs witnessing the soundness of one step of the data reduction.

The number of clauses drops at every reduction step by k , so the length of the data reduction chain is $O(n^d/k)$. \blacktriangleleft

4 Proof Complexity of principles in Computational Social Choice

A great number of applications come from the theory of Social Choice [7]: motivated by pioneering work of [38], a significant amount of research in Artificial Intelligence has investigated the provability of such results in logical settings (see [22] for a recent survey). We show that the most interesting of these results (Arrow's theorem and the Gibbard-Satterthwaite theorem) have proof complexity counterparts: the unsatisfiability of formulas encoding them can be certified by Frege proofs of subexponential length. A first example of application is Arrow's Theorem. The formulas encoding the nonexistence of a social welfare function satisfying the conditions of Arrow's theorem are rather large. Nevertheless, such an encoding exists, and was used explicitly in [38] to give a computer-assisted proof of Arrow's theorem⁴:

► **Definition 28.** Consider an instance with n agents and m objects to rank. There are $(m!)^n$ possible profiles for the complete rankings of the m objects, and $m!$ possible aggregate orderings of the m objects. Formula $\text{Arrow}_{m,n}$ (unsatisfiable for $m, n \geq 3$) has $(m!)^{n+1}$ variables $X_{R,\pi}$, one for each possible pair (R, π) consisting of ranking profile R , and an aggregate ordering $\pi \in S_m$. The constraints are the following:

- For every $R \in \mathcal{R}$ and $\pi_1 \neq \pi_2 \in S_m$ add clauses $\bigvee_{\pi \in S_m} X_{R,\pi}$ ("every profile is aggregated to some ordering") and $\overline{X_{R,\pi_1}} \vee \overline{X_{R,\pi_2}}$ ("no profile is aggregated to more than one ordering")
- For $i = 1, \dots, n$ we add to $\text{Arrow}_{m,n}$ clauses $\bigvee_{R \in \mathcal{R}} \overline{X_{R,R_i}}$. These forbid aggregations that always output the ordering given by the i 'th agent, i.e. dictatorial rank aggregations.
- For every two objects a, b let $S_{a,b}^m$ be the set of orderings π where for all $i = 1, \dots, n$, $\pi^{-1}(a) < \pi^{-1}(b)$ (i.e. a is preferred to b in ordering π). Let $\mathcal{R}_{a,b}$ be the set of profiles such that for every $i = 1, \dots, n$, $R_i \in S_{a,b}^m$ (i.e. all agents prefer a to b). For every $R \in \mathcal{R}_{a,b}$ add to $\text{Arrow}_{m,n}$ clauses $\bigvee_{\pi \in S_{a,b}^m} X_{R,\pi}$. These constraints encode unanimity (if all agents prefer object a to b then a is preferred to b in the aggregated ranking).

⁴ For encodings of Arrow's Theorem in more powerful logical frameworks see [14, 23].

- For all profiles $R, R' \in \mathcal{R}$ and objects a, b such that all players rank a, b in the same way in both R, R' and all pairs $\pi_1, \pi_2 \in S^m$ that rank a, b in a different way (i.e. $\pi_1^{-1}(a) < \pi_1^{-1}(b)$ but $\pi_2^{-1}(a) > \pi_2^{-1}(b)$ or viceversa) we add to $\text{Arrow}_{m,n}$ clauses $\overline{X_{R,\pi_1}} \vee \overline{X_{R',\pi_2}}$. These encode independence of irrelevant alternatives (if R, R' coincide with respect to the relative ordering of a, b then their aggregate orderings also rank a, b in the same way).

Results in [38] yield a kernelization for $\text{Arrow}_{m,n}$ with a reduction chain of length $O(m+n)$. We improve them by providing a kernelization with reduction chains whose length only depends on n , implying the existence of polynomial size Frege proofs for constant values of n :

- **Theorem 29.** *Formulas $\text{Arrow}_{m,n}$ have a kernelization with data reduction chains of length $\leq C(n+1)$, with constant C independent from m, n , whose safety is witnessed by polynomial-size Frege proofs. Hence (a) formulas $\text{Arrow}_{m,n}$ have Frege proofs of size quasipolynomial in $|\text{Arrow}_{m,n}|$. (b). For every fixed $n \geq 3$ there exists a polynomial $p_n(\cdot)$ such that for all $m \geq 3$ formulas $\text{Arrow}_{m,n}$ have Frege proofs of size at most $p_n(|\text{Arrow}_{m,n}|)$.*

Proof. The kernelization has two data reduction rules, described informally as follows:

- (a). If $n \geq 2, m \geq 6$ and $W : [S_m]^n \rightarrow [S_m]$ is a function that is non-dictatorial, IIA and unanimous then there exists an $T \subseteq [m]$, $|T| = m - 5$ such that $W_{-T} : [S_{[m]-T}]^n \rightarrow [S_{[m]-T}]$ (see the end of Section 2) has the same properties. In other words, one can reduce in one step the set of alternatives from $[m]$ (which has m elements) to $[m]_{-T}$ (which has 5).
- (b). See [38]: If $n, m \geq 3$ and $W : [S_m]^n \rightarrow [S_m]$ is non-dictatorial, IIA and unanimous then at least one of the functions $W_{1,2}, W_{1,3}, W_{2,3} : [S_m]^{n-1} \rightarrow [S_m]$ defined as follows:
 $W_{i,j}(R_1, R_2, \dots, \widehat{R}_i, \dots, R_n) = W(R'_1, \dots, R'_n)$ is non-dictatorial, IIA and unanimous. Here $R'_i = R_j$, $R'(k) = R_k, k \neq i$. In other words, one can reduce in one step the number of agents by one.

- **Lemma 30.** *If W is unanimous, IIA and non-dictatorial then for every $B \subseteq [m]$, function W_{-B} is unanimous and IIA.*

Proof. Suppose that $a, a' \in [m]_{-B}$ and $a <_{R_i} a'$ for all $i \in [m]_{-B}$. Then $a <_{R_i^{+B}} a'$. By unanimity of W , $a <_{W(R_1^{+B}, \dots, R_n^{+B})} a'$. Since a, a' were arbitrary, it follows that W_{-B} is unanimous. As for IIA, let $a, a' \in [m]_{-B}$ and (R_1, R_2, \dots, R_n) and $(R'_1, R'_2, \dots, R'_n)$ be preference profiles such that, for every $i = 1, \dots, n$, R_i and R'_i agree with respect to the relative ordering of a, a' . Then for every $i = 1, \dots, n$, R_i^{+B} and R'_i^{+B} agree with respect to the relative ordering of a, a' . By the IIA axiom for W , $W(R_1^{+B}, R_2^{+B}, \dots, R_n^{+B})$ and $W(R'_1^{+B}, R'_2^{+B}, \dots, R'_n^{+B})$ agree with respect to the relative ranking of a, a' . Hence so do $W_B(R_1, R_2, \dots, R_n)$ and $W_B(R'_1, R'_2, \dots, R'_n)$. ◀

- **Lemma 31.** *Reduction (a). is safe.*

Proof. Consider an arbitrary set $T \subseteq [m]$ of cardinality $m - 6$, e.g. $T = \{7, \dots, m\}$. Let $x \notin T$, e.g. $x = 6$ and $U = T \cup \{x\}$. If W_{-U} is non-dictatorial we are done. Otherwise, assume w.l.o.g. that agent 1 is a dictator for W_{-U} . Since 1 is not a dictator for W , there must exist indices $c \neq d \in [m]$ and preference profiles $<_1, \dots, <_n$ on $[m]$ such that $c <_1 d$ but $d <_{W(<_1, \dots, <_n)} c$. Let $y, a, b \notin T$, different from x, c, d , and let $V \subseteq [m]$, $|V| = m - 5$, $a, b, c, d, x \notin V$, $y \in V$. Such a V exists, since $m \geq 6$. Clearly $V \neq U$, since $y \in V \setminus U$. We claim that function W_{-V} is not dictatorial.

First note that 1 cannot be a dictator for W_{-V} . Indeed, consider $<_{i,-V}$ the restriction of $<_i$ to $[m]_{-V}$. We have $c <_{1,-V} d$ but $d <_{W_{-V}(<_{1,-V}, \dots, <_{n,-V})} c$. The first relation holds because $c <_1 d$ and $c, d \notin V$. The second relation holds because to compare c, d according to

$W_{-V}(\langle_{1,-V}, \dots, \langle_{n,-V})$ we apply function W on $(\langle_{1,-V}^{+V}, \dots, \langle_{n,-V}^{+V})$. But since $c, d \notin V$, $\langle_{i,-V}^{+V}$ coincides with \langle_i with respect to the ordering of c, d for $i = 1, \dots, n$. Invoking the IIA property of W for tuples $(\langle_1, \dots, \langle_n)$ and $(\langle_{1,-V}^{+V}, \dots, \langle_{n,-V}^{+V})$ justifies the second relation.

Assume now that another agent, say 2, were a dictator for W_{-V} . Let \langle_1, \langle_2 be preference profiles on $[m]$ s.t. $a \langle_1 b$ but $b \langle_2 a$, and $\langle_3, \dots, \langle_n$ be arbitrary preference profiles. First, invoking the first relation, the fact that W_{-U} is computed using W , and that 1 is a dictator for W_{-U} we get that $a \langle_{W_{-U}(\langle_{1,-U}, \langle_{2,-U}, \dots, \langle_{n,-U})} b$. Invoking the IIA property of W on tuples $(\langle_{1,-U}, \langle_{2,-U}, \dots, \langle_{n,-U})$ and $(\langle_1, \langle_2, \dots, \langle_n)$ we get that $a \langle_{W(\langle_1, \langle_2, \dots, \langle_n)} b$. Using a similar reasoning for the function W_{-V} we get that $b \langle_{W(\langle_1, \langle_2, \dots, \langle_n)} a$, a contradiction. ◀

The safety of reduction (b) was (mathematically) proved in [38]. Next we outline how to simulate these mathematical arguments using polynomial-size Frege proofs. First, note that formula $Arrow_{m,n}$ has $(m!)^{n+1}$ variables, all of them appearing explicitly in the formula. But $n = O(\log((m!)^{n+1}))$, so indeed a reduction chain of length $O(n)$ has length logarithmic in $|Arrow_{m,n}|$. Invoking our metatheorem yields a proof of point (a) of Theorem 29. Point (b) follows by invoking point 2 of the same metatheorem.

First reduction rule: Define, for $Q \subseteq [m]$, $|Q| = m - 5$, and $i = 1, \dots, n$ formulas

$$Nondict_{-Q,i} := \bigvee_{R \in \mathcal{R}_{-Q}} \overline{X_{R+Q, R_i^{+Q}}}$$

(informally, formula $Nondict_{-Q,i}$ is true iff i is not a dictator for W_{-Q}).

$$Unanimous_{-Q} := \bigwedge_{a,b \in [m]_{-Q}} \bigvee_{\substack{R \in \mathcal{R}_{a,b,-Q} \\ \pi \in S_{a,b}^m}} X_{R+Q, \pi+Q}.$$

$$IIA_{-Q} := \bigwedge_{(R, R', \pi, \pi') \in \mathcal{R}_{-Q}} (\overline{X_{R+Q, \pi_1^{+Q}}} \vee \overline{X_{R'+Q, \pi_2^{+Q}}})$$

(where, for simplicity, we have omitted the IIA restrictions on R, R', π_1, π_2 , see Definition 28). Note that $Arrow_{m,n} = Unanimous_{-\emptyset} \wedge IIA_{-\emptyset} \wedge \bigwedge_{i=1}^n Nondict_{-\emptyset,i}$.

We will prove that

$$Arrow_{m,n} \wedge \bigwedge_{i=1}^n Nondict_{-[6:m],i} \vdash Arrow_{5,n} \quad (5)$$

and, for $i = 1, \dots, n$

$$Arrow_{m,n} \wedge \bigvee_{i=1}^n \bigwedge_{R \in \mathcal{R}_{-[6:m]}} X_{R+[6:m], R_i^{+[6:m]}} \vdash Arrow_{5,n} \quad (6)$$

Employing tautology $\bigwedge_{i=1}^n Nondict_{-[6:m],i} \vee \bigvee_{i=1}^n (\bigwedge_{R \in \mathcal{R}_{-[6:m]}} X_{R+[6:m], top(R_i^{+[6:m])})}$ and substitutions implicit in (5) and (6) we conclude that $Arrow_{m,n} \vdash P_1 \vee P_2$, where both P_1, P_2 are formulas isomorphic to $Arrow_{5,n}$. Thus we are in the framework of our metatheorem with $R = 2$.

We need to specify the substitutions implicit in (5) and (6). First, formalizing the mathematical argument in Lemma 30 we show that $Arrow_{m,n} \vdash IIA_{-Q} \wedge Unanimous_{-Q}$. The propositional content of this implication is trivial: the clauses of $Unanimous_{-Q}$ and IA_{-Q} are simply subclauses of $Arrow_{m,n}$.

Because of this, the substitution witnessing implication (5) is quite simple: it replaces a restricted variable $X'_{R,\pi}$ of $Arrow_{5,n}$ with the variable $X_{R+[6:m], \pi+[6:m]}$ of $Arrow_{m,n}$.

135:16 Kernelization, Proof Complexity and Social Choice

As for (6), define, for all $c < d \in [5]$ and $i = 1, \dots, n$, formulas

$$Witness_{c,d,i} := \bigvee_{\substack{R \in \mathcal{R}: R_i \in S_{c,d}^m \\ \pi \in S_{d,c}^m}} X_{R,\pi} \quad (7)$$

Informally, formula $Witness_{c,d,i}$ is true when pair (c, d) acts as a witness that agent i is not a dictator for W , since $c <_{R_i} d$ but $d <_{W(R)} c$.

A next step is to prove that $Arrow_{m,n} \wedge \bigwedge_{R \in \mathcal{R}_{-[6:m]}} X_{R^{+[6:m]}, R_i^{+[6:m]}} \vdash \bigvee_{c < d \in [5]} Witness_{c,d,i}$.

This is easy: we use literals $X_{R^{+[6:m]}, R_i^{+[6:m]}}$ to prove (by resolution)

$$Arrow_{m,n} \wedge \bigwedge_{R \in \mathcal{R}_{-[6:m]}} X_{R^{+[6:m]}, R_i^{+[6:m]}} \wedge \bigwedge_{c < d \in [5]} \bigwedge_{\substack{R \in \mathcal{R}: R_i \in S_{c,d}^m \\ \pi \in S_{d,c}^m}} \overline{X_{R,\pi}} \vdash \square \quad (8)$$

(the last conjunction negates formulas $Witness_{c,d,i}$) and then invoke Lemma 2 to get a proof of the same length of the implication we claimed.

Now we prove, for all $j = 1, \dots, n$ that $Arrow_{m,n} \wedge Witness(c, d, i) \vdash Nondict_{-V,j}$, where V is defined as in the proof of Lemma 31.

The proof of the implication for $j = i$ uses unit literal X_{R,R_j} (negation of one from $Nondict_{-V,i}$) and $\overline{X_{R,\pi}} \vee \overline{X_{R^{+V}, \pi_2}}$ (part of the IIA part of $Arrow_{m,n}$) to derive clauses $\overline{X_{R^{+V}, \pi_2}}$ for all π_2 that rank c, d in a different way than π . Resolving away these literals from clause $\bigvee_{\pi \in S_m} X_{R^{+V}, \pi}$ (part of $Arrow_{m,n}$) derives clause $Nondict_{-V,i}$.

As for the case $j \neq i$, we want to show that

$$Arrow_{m,n} \wedge Witness(c, d, i) \wedge \bigwedge_{R \in \mathcal{R}_{[5]}} X_{R^{+[6:m]}, R_j^{+[6:m]}} \vdash \square \quad (9)$$

By Lemma 2 this will imply $Nondict_{-V,j}$. Let $R = (R_1, R_2, \dots, R_n)$ be a profile on $[5]$ such that $a <_{R_1} b$ but $b <_{R_2} a$ ($a, b \in [5]$ are defined as in Lemma 31). Combining the derivations of $Nondict_{-V,j}$ in (9) with the ones of $Unanimous_{-V}$ and IIA_{-V} (outlined before), plus a bijective identification of $[m]_{-V}$ and $[5]$ yields a substitution that proves $Arrow_{5,n}$, completing the proof of (6).

Second reduction rule: We refer to Lemma 2 of [38] for (mathematical) details of the reduction. What is important is that the soundness of the statement that at least one of $W_{1,2}, W_{2,3}, W_{1,3}$ is non-dictatorial is established by a case-by-case analysis. It is first proved that it cannot be that all these functions have the same dictator. Then it is established that if i is the dictator of $W_{1,2}$, j the dictator of $W_{1,3}$, k the dictator of $W_{2,3}$ then $i \in \{2, 3\}$, $j \in \{2, 3\}$, $k \in \{1, 3\}$. For all eight possible cases for triplets (i, j, k) we obtain a contradiction: either we explicitly provide a profile R showing that triplet (i, j, k) cannot represent the set of dictators for the three function, or we employ an argument similar to the one in the case $i = j = k$. ◀

As for the Gibbard-Satterthwaite theorem, we use the following formalization:

► **Definition 32.** Consider an instance with n agents and m objects. There are $(m!)^n$ possible profiles for the complete rankings of the m objects, and m possible outcomes. Formula $GS_{m,n}$ has $(m!)^n \times m$ variables $X_{R,o}$, one for each possible pair consisting of a strategy profile R and a value $o \in [m]$, the value of the SCF on profile R . The constraints are the following:

- For $R \in \mathcal{R}$ add clauses $\bigvee_{o \in [m]} X_{R,o}$ (“every joint profile is aggregated to some object”) and $\overline{X_{R,o_1}} \vee \overline{X_{R,o_2}}$ (“no joint profile is aggregated to more than one object”).
- For $i = 1, \dots, n$ we add to $GS_{m,n}$ clauses $\bigvee_{R \in \mathcal{R}} \overline{X_{R, \text{top}(R_i)}}$. They forbid social choice functions that always output the top preference of the i 'th agent, i.e. dictatorial aggregations.
- For $i = 1, \dots, n$ add $\bigvee_{R \in \mathcal{R}} X_{R,o}$. This eliminates social choice functions that are not onto.
- Add, for every pair (R, o) , player $1 \leq i \leq m$ and $\pi \in S_m$, $\overline{X_{R,o}} \vee (\bigvee_{o' \in \text{pr}(i,o,R)} X_{s(i,R,\pi),o'})$. These clauses state that the social choice function is strategyproof.

► **Theorem 33.** For every fixed m formulas $GS_{n,m}$ have a kernelization of length $O(n)$ whose soundness has polynomial time Frege proofs. Hence, formulas $GS_{n,m}$ expressing the Gibbard-Satterthwaite theorem have (a). Frege proofs of size quasipolynomial in $|GS_{n,m}|$, (b). for every fixed n , Frege proofs of size polynomial in $|GS_{n,m}|$.

Proof. The argument is very similar to that of Arrow's theorem: We use the following two data reductions: (a). for $n \geq 2, m \geq 4$, if $W : [m]^n \rightarrow [m]$ is a social choice function that is onto, non-dictatorial and strategy-proof then there exists $T \subseteq [m]$, $|T| = m - 3$ such that function $W_{-T} : ([m]_{-T})^n \rightarrow [m]_{-T}$ is onto, non-dictatorial and strategyproof, and (b). [37] if $W : [m]^n \rightarrow [m]$ is a social choice function that is onto, non-dictatorial and strategy-proof then one of functions $W_{1,2}, W_{1,3}, W_{2,3} : [S_m]^{n-1} \rightarrow [m]$ defined by $W_{a,b}(R_1, R_2, \dots, \widehat{R_a}, \dots, R_n) = W(R'_1, \dots, R'_n)$ is non-dictatorial, onto and strategy-proof. Here $R'_a = R_b$, $R'_k = R_k$ for $k \neq a$.

First, it is not obvious that, as formulated in the paragraph, W_{-B} is well-defined. The reason is that $W_{-B}(R)$ invokes W on profile R^{+B} , and it is not obvious that if R is a profile on $[m]_{-B}$ the the outcome of W is an element of $[m]_{-B}$, as needed by the definition.

Suppose that $W_{-B}(R) = W(R^{+B}) \in B$ for some profile R on $[m]_{-B}$. We claim that $W(S) \in B$ for every profile $S = (S_1, S_2, \dots, S_n)$, contradicting the hypothesis that W is onto. Indeed, if $W_{-B}(R) = W(R^{+B}) \in B$ then $W_{-B}(R_{-1}, S_1) = W(R_{-1}^{+B}, S_1^{+B}) \in B$, otherwise agent 1 would have an opportunity to manipulate at profile R^{+B} by misrepresenting its preference as S_1^{+B} . Applying this argument inductively for agents $2, 3, \dots, n$ (replacing R_i by S_i) we infer that $W(S) \in B$, which is what we claimed.

► **Lemma 34.** If $W : S_m^n \rightarrow [m]$ is onto, strategyproof and non-dictatorial then for every $B \subseteq [m]$, function W_{-B} is onto and strategyproof.

Proof. Suppose there exists some profile R and agent $i \in [n]_{-B}$ such that i could manipulate $W_{-B}(R)$ by misrepresenting its profile as R'_i . This means that i could manipulate W on profile R^{+B} by misrepresenting its profile as $R'_i{}^{+B}$, contradicting the fact that W is strategy-proof. Hence W_{-B} is strategy-proof.

Suppose now that $a \in [m]_{-B}$. Since W is onto, there must exist a profile R such that $W(R) = a$. Consider the profile R'_i that modifies R by moving a to the top of preference profile R_i . We claim that $W(R'_i) = a$. Indeed, if this was not the case then i could manipulate on profile R'_i by misrepresenting its preferences. Continuing the argument inductively for all agents we infer that if \overline{R} is the profile that modifies R by moving a to the top of all profile preferences then $W(\overline{R}) = a$. This means that W is unanimous, hence W_{-B} also is. But then there is a profile \underline{R} such that $W_{-B}(\underline{R}) = a$: simply make a the top of all profiles. Since a was arbitrary, it follows that W_{-B} is onto. ◀

135:18 Kernelization, Proof Complexity and Social Choice

We next prove the safety of reduction rule (a) (for (b) see [37]). We show that there exists $T \subseteq [m]$, $|T| = m - 3$ such that W_{-T} is non-dictatorial. Together with Lemma 34 this establishes the safety of rule a.

Step 1. Given two different sets T_1, T_2 of size $m - 3$, $|T_1 \cap T_2| = m - 4$ (in other words, $|\overline{T_1}| = |\overline{T_2}| = 3$, $|\overline{T_1} \cap \overline{T_2}| = 2$), we show that functions W_{-T_1}, W_{-T_2} must have the same dictator, if they have one.

Suppose, indeed, that W_{-T_1} has dictator i , W_{-T_2} has dictator $j \neq i$. Let $d \in T_2 \setminus T_1$, $c \in T_1 \setminus T_2$, $a, b \in \overline{T_1} \cap \overline{T_2}$. Define profiles

$$R_s = a < b < \dots < \text{sorted}(T_2) < c, \text{ for } s \neq j$$

$$R_j = b < a < \dots < \text{sorted}(T_2) < c.$$

$W(R) = a$, since i is a dictator for W_{-T_1} . Now, if we replace R_i by

$$R'_i = a < b < \dots < c < \text{sorted}(T_2).$$

obtaining profile R' , then $W(R') = a$, otherwise agent i could manipulate by reporting profile R_i instead. We continue changing iteratively profiles $R_s, s \neq j$ to $R'_s = a < b < \dots < c < \text{sorted}(T_2)$, one profile at a time, until all profiles R_s except R_j have been replaced by R'_s . Call this profile \underline{R} . That is, $\underline{R}_s = a < b < \dots < c < \text{sorted}(T_2)$ for $s \neq j$, while $\underline{R}_j = b < a < \dots < \text{sorted}(T_2) < c$.

Since no agent $s \neq j$ had an opportunity to manipulate, it must be that $W(\underline{R}) = a$. Consider now profile $\underline{R}'_j = b < a < \dots < c < \text{sorted}(T_2)$. Since W_{-T_2} has agent j as a dictator, $W(\underline{R}_1, \dots, \underline{R}'_j, \dots, \underline{R}_n) = W_{-T_2}(\underline{R}_{1,-T_2}, \dots, \underline{R}'_{j,-T_2}, \underline{R}_{n,-T_2}) = b$. So agent j has an opportunity to manipulate at profile $(\underline{R}_1, \dots, \underline{R}'_j, \dots, \underline{R}_n)$ by reporting instead \underline{R}'_j .

Step 2. Either there exists a set T of size $m - 3$ such that W_{-T} is not dictatorial, or all functions $W_{-T}, |T| = m - 3$ must have the same dictator. Indeed, we can “interpolate” between any two sets of cardinality $m - 3$ by a sequence of sets falling under step 1.

Step 3. We show that it is not possible that all functions $W_{-T}, |T| = m - 3$ have the same dictator i . Since W is not dictatorial, there exists a profile R such that $b = W(R)$ is different from $a = \text{top}(R_i)$. Let $T \subseteq [m]$, $|T| = m - 3$, $a, b \notin T$ and consider a profile R' that modifies R by moving $b, \text{sorted}(T)$ to the bottom of all preferences (in this order), that is $R' = (R_{-b}^{+b})_{- \text{sorted}(T)}^{+ \text{sorted}(T)}$.

We have $W(R') = W_{-T}(R'_{-T}) = \text{top}(R'_i) = a$, since W_{-T} has i as dictator and $\text{top}(R'_i) = a$. Let us create a path between R and R' by changing one profile R_s at a time to R'_s , the last move being R_i . The value of W does not change at any step s , since W_{-T} has i as dictator, and the relative orders of elements in $[m]_{-T}$ does not change at any profile as a result of a change $R_s \rightarrow R'_s, s \neq i$, or agent s would have an opportunity of manipulation at one of the two profiles, using R_s, R'_s , whichever yields a result ranked lower in R_s, R'_s . But this yields a contradiction, since $W(R) = b$ and $W(R') = a \neq b$.

For some details on the propositional simulations we refer to [26]. ◀

5 Conclusions and open problems

We believe that the most important contribution of our paper is to show that several techniques for proving kernelization can sometimes be simulated by efficient extended Frege proofs. The proof techniques in this list includes: crown decomposition, the sunflower lemma, ad-hoc methods. It is an interesting challenge to enlarge the list of methods and problems that have such a simulation or, conversely, show limits of some of these methods.

It is rewarding to note that methods from algebraic topology can be used to reframe and extend results in both Topological Combinatorics and Computational Social Choice: Arrow's theorem has topological proofs [4, 12, 13]. On the other hand the original results on Kneser's conjecture [32] have been strengthened using more advanced topological methods, e.g. [3] (see [29] for a book-length treatment). The results in [2, 38] can be interpreted as stating that in both cases one can bypass topological arguments by purely combinatorial arguments (plus computer-assisted verification of finitely many cases). It is an interesting question whether this is still true for the results requiring more sophisticated topological methods as well.

In Theorem 26 we have only obtained polynomial size extended Frege proofs. Similarly, in Theorems 29, 33 we have only obtained polynomial-size Frege proofs when the number of agents is fixed. We leave open the issue of improving these results. On the other hand, our results have only touched on the most basic topics on the proof complexity of statements in computational social choice. There has been significant progress in this area (see e.g. [22]) and we believe that our framework may be applicable to some of this work (e.g. to the Preservation Theorem of [21]). It would be interesting to see if this is really the case.

References

- 1 James Aisenberg, Maria Luisa Bonet, and Sam Buss. Quasipolynomial size Frege proofs of Frankl's theorem on the trace of sets. *Journal of Symbolic Logic*, 81(2):687–710, 2016.
- 2 James Aisenberg, Maria Luisa Bonet, Sam Buss, Adrian Crăciun, and Gabriel Istrate. Short proofs of the Kneser–Lovász coloring principle. *Information and Computation*, 261:296–310, 2018.
- 3 E. Babson and D.N. Kozlov. Proof of the Lovász conjecture. *Annals of Mathematics*, 165:965–1007, 2007.
- 4 Yuliy M Baryshnikov. Topological and discrete social choice: in a search of a theory. In *Topological social choice*, pages 53–63. Springer, 1997.
- 5 Olaf Beyersdorff, Nicola Galesi, Massimo Lauria, and Alexander A Razborov. Parameterized bounded-depth Frege is not optimal. *ACM Transactions on Computation Theory (TOCT)*, 4(3):1–16, 2012.
- 6 Maria Luisa Bonet, Samuel R Buss, and Toniann Pitassi. Are there hard examples for Frege systems? In *Feasible Mathematics II*, pages 30–56. Birkhäuser Boston, 1995.
- 7 Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D Procaccia. *Handbook of computational social choice*. Cambridge University Press, 2016.
- 8 S. Buss. Polynomial size proofs of the propositional pigeonhole principle. *Journal of Symbolic Logic*, 52(4):916–927, 1987.
- 9 Sam Buss. Quasipolynomial size proofs of the propositional pigeonhole principle. *Theoretical Computer Science*, 576:77–84, 2015.
- 10 Sam Buss and Jakob Nordström. Proof complexity and SAT solving. *Chapter to appear in the 2nd edition of Handbook of Satisfiability*, 2021. Draft version available at <https://www.math.ucsd.edu/~sbuss/ResearchWeb/ProofComplexitySAT>.
- 11 Lorenzo Carlucci, Nicola Galesi, and Massimo Lauria. Paris-Harrington tautologies. In *2011 IEEE 26th Annual Conference on Computational Complexity*, pages 93–103, 2011.

- 12 Graciela Chichilnisky. Social choice and the topology of spaces of preferences. *Advances in Mathematics*, 37(2):165–176, 1980.
- 13 Graciela Chichilnisky. The topological equivalence of the Pareto condition and the existence of a dictator. *Journal of Mathematical Economics*, 9(3):223–233, 1982.
- 14 Giovanni Ciná and Ulle Endriss. A syntactic proof of Arrow’s theorem in a modal logic of social choice functions. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pages 1009–1017, 2015.
- 15 Stephen A Cook and Robert A Reckhow. The relative efficiency of propositional proof systems. *The Journal of Symbolic Logic*, 44(1):36–50, 1979.
- 16 Stefan Dantchev, Barnaby Martin, and Stefan Szeider. Parameterized proof complexity. *Computational Complexity*, 20(1):51, 2011.
- 17 M. de Longueville. *A Course in Topological Combinatorics*. Springer, 2012.
- 18 Rodney G Downey and Michael R Fellows. *Fundamentals of parameterized complexity*. Springer, 2013.
- 19 Fedor V Fomin, Daniel Lokshtanov, Saket Saurabh, and Meirav Zehavi. *Kernelization: theory of parameterized preprocessing*. Cambridge University Press, 2019.
- 20 Fedor V Fomin, Daniel Lokshtanov, Saket Saurabh, and Meirav Zehavi. Parameterized algorithms. *Beyond the Worst-Case Analysis of Algorithms*, page 27, 2020.
- 21 Christian Geist and Ulrich Endriss. Automated search for impossibility theorems in social choice theory: Ranking sets of objects. *Journal of Artificial Intelligence Research*, 40:143–174, 2011.
- 22 Christian Geist and Dominik Peters. Computer-aided methods for social choice theory. *Trends in Computational Social Choice*, pages 249–267, 2017.
- 23 Umberto Grandi and Ulle Endriss. First-order logic formalisation of Arrow’s theorem. In *International Workshop on Logic, Rationality and Interaction*, pages 133–146. Springer, 2009.
- 24 Pavel Hrubes and Iddo Zameret. Short proofs for the determinant identities. *SIAM Journal on Computing*, 44(2):340–383, 2015.
- 25 G. Istrate, A. Percus, and S. Boettcher. Spines of random constraint satisfaction problems: Definition and connection with computational complexity. *Annals of Mathematics and Artificial Intelligence*, 44(4):353–372, 2005.
- 26 Gabriel Istrate, Cosmin Bonchis, and Adrian Craciun. Kernelization, proof complexity and social choice, 2021. [arXiv:2104.13681](https://arxiv.org/abs/2104.13681).
- 27 Gabriel Istrate and Adrian Craciun. Proof complexity and the kneser-lovász theorem. In *International Conference on Theory and Applications of Satisfiability Testing*, pages 138–153. Springer, 2014.
- 28 Leszek Aleksander Kołodziejczyk, Phuong Nguyen, and Neil Thapen. The provably total NP search problems of weak second order bounded arithmetic. *Annals of Pure and Applied Logic*, 162(6):419–446, 2011.
- 29 D. Kozlov. *Combinatorial Algebraic Topology*. Springer Verlag, 2008.
- 30 Jan Krajíček. *Proof complexity*, volume 170. Cambridge University Press, 2019.
- 31 Massimo Lauria, Pavel Pudlák, Vojtěch Rödl, and Neil Thapen. The complexity of proving that a graph is Ramsey. *Combinatorica*, 37(2):253–268, 2017.
- 32 L. Lovász. Kneser’s conjecture, chromatic number, and homotopy. *Journal of Combinatorial Theory, Series A*, 25:319–324, 1978.
- 33 C. Moore, G. Istrate, D. Demopoulos, and M. Vardi. A continuous-discontinuous second-order transition in the satisfiability of a class of Horn formulas. *Random Structures and Algorithms*, 31(2):173–185, 2007.
- 34 Akihiro Nozaki, Toshiyasu Arai, Noriko H Arai, et al. Polynomial-size Frege proofs of Bollobás’ theorem on the trace of sets. *Proceedings of the Japan Academy, Series A, Mathematical Sciences*, 84(8):159–161, 2008.
- 35 A. Schrijver. Vertex-critical subgraphs of Kneser graphs. *Nieuw Arch. Wiskd., III. Ser.*, 26:454–461, 1978.

- 36 Michael Soltys and Stephen Cook. The proof complexity of linear algebra. *Annals of Pure and Applied Logic*, 130(1-3):277–323, 2004.
- 37 Pingzhong Tang and Fangzhen Lin. A computer-aided proof to Gibbard–Satterthwaite theorem. Technical report, Technical report, 2008. <http://iis.tsinghua.edu.cn/~kenshin>, 2008.
- 38 Pingzhong Tang and Fangzhen Lin. Computer-aided proofs of Arrow’s and other impossibility theorems. *Artificial Intelligence*, 173(11):1041–1053, 2009.