

Polynomial Kernels and User Reductions for the Workflow Satisfiability Problem

Gregory Gutin · Stefan Kratsch ·
Magnus Wahlström

Received: date / Accepted: date

Abstract The WORKFLOW SATISFIABILITY PROBLEM (WSP) is a problem of practical interest that arises whenever tasks need to be performed by authorized users, subject to constraints defined by business rules. We are required to decide whether there exists a *plan* – an assignment of tasks to authorized users – such that all constraints are satisfied.

The WSP is, in fact, the CONSERVATIVE CONSTRAINT SATISFACTION PROBLEM (i.e., for each variable, here called *task*, we have a unary authorization constraint) and is, thus, NP-complete. It was observed by Wang and Li (2010) that the number k of tasks is often quite small and so can be used as a parameter, and several subsequent works have studied the parameterized complexity of WSP regarding parameter k .

We take a more detailed look at the kernelization complexity of $\text{WSP}(\Gamma)$ when Γ denotes a finite or infinite set of allowed constraints. Our main result is a dichotomy for the case that all constraints in Γ are regular: (1) We are able to reduce the number n of users to $n' \leq k$. This entails a kernelization to size $\text{poly}(k)$ for finite Γ , and, under mild technical conditions, to size $\text{poly}(k+m)$ for infinite Γ , where m denotes the number of constraints. (2) Already $\text{WSP}(R)$ for some $R \in \Gamma$ allows no polynomial kernelization in $k+m$ unless the polynomial hierarchy collapses.

Keywords Workflow satisfiability problem · parameterized complexity · kernelization

G. Gutin
Royal Holloway, University of London, UK
E-mail: G.Gutin@rhul.ac.uk

S. Kratsch
TU Berlin, Germany
E-mail: stefan.kratsch@tu-berlin.de

M. Wahlström
Royal Holloway, University of London, UK
E-mail: Magnus.Wahlstrom@rhul.ac.uk

1 Introduction

A business process is a collection of interrelated tasks that are performed by users in order to achieve some objective. In many situations, a task can be performed only by certain *authorized* users; formally, every task is accompanied by an *authorization list* of all users who are authorized to perform the task. Additionally, either because of the particular requirements of the business logic or security requirements, we may require that certain sets of tasks cannot be performed by some sets of users [7]. Such restrictions are referred to as *constraints*, and may include rules such as *separation-of-duty* (also known as the “two-man” rule), which may be used to prevent sensitive combinations of tasks being performed by a single user, and *binding-of-duty*, which requires that a particular combination of tasks is performed by the same user. The use of constraints in workflow management systems to enforce security policies has been studied extensively in the last fifteen years; see, e.g., [3, 7, 17].

It is possible that the combination of constraints and authorization lists is “unsatisfiable”, in the sense that there does not exist an assignment of users to tasks (called a *plan*) such that all constraints are satisfied and every task is performed by an authorized user. A plan that satisfies all constraints and allocates an authorized user to each task is called *valid*. The WORKFLOW SATISFIABILITY PROBLEM (WSP) takes a workflow specification as input and returns a valid plan if one exists and NO otherwise. It is important to determine whether a business process is satisfiable or not, since an unsatisfiable one can never be completed without violating the security policy encoded by the constraints and authorization lists.

Let us illustrate the above notions by the following simple instance W^* of WSP. Let us be given tasks s_1, s_2, s_3 , authorization lists $A(s_1) = \{u_1, u_2, u_3\}$, $A(s_2) = \{u_1, u_4, u_5\}$, $A(s_3) = \{u_1, u_6\}$, a binding-of-duty constraint $s_1 = s_2$ (meaning that s_1 and s_2 must be assigned the same user) and two separation-of-duty constraints $s_1 \neq s_3$ and $s_2 \neq s_3$. Note that the only valid plan is the assignment of s_1 and s_2 to u_1 and s_3 to u_6 .

It is worth noting that WSP is a special class of constraint satisfaction problems where for each variable s (called a task in the WSP language) we have an arbitrary unary constraint (called an authorization) that assigns possible values (called users) for s ; this is called the CONSERVATIVE CONSTRAINT SATISFACTION problem. Note, however, that while usually in constraint satisfactions problems the number of variables is much larger than the number of values, for WSP the number of tasks is usually much smaller than the number of users. It is important to remember that for WSP we do not use the term ‘constraint’ for authorizations and so when we define special types of constraints, we do not extend these types to authorizations, which remain arbitrary.

Wang and Li [17] were the first to observe that the number k of tasks is often quite small and so can be considered as a parameter. As a result, WSP can be studied as a parameterized problem. Wang and Li [17] proved that, in general, WSP is $W[1]$ -hard, but WSP is fixed-parameter tractable (FPT) if we

consider some special types of practical constraints which include separation-of-duty and binding-of-duty constraints.

Crampton et al. [9] found a faster fixed-parameter algorithm to solve the special cases of WSP studied in [17] and showed that the algorithm can be used for a wide family of constraints called regular (in fact, regular constraints include all constraints studied in [17]). Subsequent research has demonstrated the existence of fixed-parameter algorithms for WSP in the presence of other constraint types [6,8]. In particular, Cohen et al. [6] showed that WSP with only so-called *user-independent* constraints is FPT. A constraint c on tasks t_1, \dots, t_r is user-independent when for any tuples u_{i_1}, \dots, u_{i_r} and u_{j_1}, \dots, u_{j_r} of users such that $u_{i_p} = u_{i_q}$ if and only if $u_{j_p} = u_{j_q}$ for all $1 \leq p < q \leq r$, c is satisfied by the assignment of t_ℓ to u_{i_ℓ} for each $\ell \in [r] := \{1, \dots, r\}$ if and only if c is satisfied by the assignment of t_ℓ to u_{j_ℓ} for each $\ell \in [r]$. Intuitively, given a satisfying assignment we may arbitrarily swap users for other users that are not presently assigned (and this may be iterated, giving arbitrary bijections). As an example, separation-of-duty and binding-of-duty are user-independent constraints. Crampton et al. [9] also launched the study of polynomial kernels and polynomial partial kernels¹ (in the latter only the number of users is required to be bounded by a polynomial in k), but obtained results only for concrete families of constraints.

In this work, we explore the kernelization properties of WSP in more detail. We focus on regular constraints, which are a special family of user-independent constraints, but since their definition is quite technical, we will defer it to Section 2.1. We study both the possibility of polynomial kernels and of simplifying WSP instances by reducing the set of users (i.e., partial kernels).² Our goal is to determine for which types of constraints such user-limiting reductions are possible, i.e., for which sets Γ does the problem $\text{WSP}(\Gamma)$ of WSP restricted to using constraint types (i.e., relations) from Γ admit a reduction to $\text{poly}(k)$ users? We study this question for both finite and infinite sets Γ of regular constraints, and show a strong separation: Essentially, either every instance with k tasks can be reduced to at most k users, or there is no polynomial-time reduction to $\text{poly}(k)$ users unless the polynomial hierarchy collapses. (However, some technical issues arise for the infinite case.)

Our results. Our main result is a dichotomy for the $\text{WSP}(\Gamma)$ problem when Γ contains only regular relations. We show two results. On the one hand, if every relation $R \in \Gamma$ is *intersection-closed* (see Section 4), then we give a polynomial-time reduction which reduces the number of users in an instance to $n' \leq k$, without increasing the number of tasks k or constraints m . This applies even if Γ is infinite, given a natural assumption on computable properties of the relations. On the other hand, we show that given even a single relation R which is regular but not intersection-closed, the problem $\text{WSP}(R)$ admits no polynomial kernel, and hence no reduction to $\text{poly}(k)$ users, unless the

¹ The term *partial kernel* was first used in [4].

² Such reductions are of interest by themselves as some practical WSP algorithms iterate over users in search for a valid plan [5].

polynomial hierarchy collapses. For finite sets Γ , this gives a dichotomy in a straight-forward manner: For every finite set Γ of regular relations, $\text{WSP}(\Gamma)$ admits a polynomial kernel if every $R \in \Gamma$ is intersection-closed, and otherwise not unless the polynomial hierarchy collapses.

However, for infinite sets Γ things get slightly more technical, for two reasons: (1) An instance with k tasks and few users could still be exponentially large due to the number of constraints, analogously to the result that HITTING SET admits no polynomial kernel parameterized by the size of the ground set [11] (cf. [13]). (2) More degenerately, without any restriction on Γ , an instance could be exponentially large simply due to the encoding size of a single constraint (e.g., one could interpret a complete WSP instance on k tasks as a single constraint on these k tasks). Both these points represent circumstances that are unlikely to be relevant for practical WSP instances. We make two restrictions to cope with this: (1) We allow the number m of constraints as an extra parameter, since it could be argued that $m \leq \text{poly}(k)$ in practice. (2) We require that each constraint of arity $r \leq k$ can be expressed by $\text{poly}(r)$ bits. E.g., this allows unbounded arity forms of all standard constraints. Using this, we obtain a more general dichotomy: For any (possibly infinite) set Γ of regular relations, $\text{WSP}(\Gamma)$ admits a kernel of size $\text{poly}(k + m)$ if every $R \in \Gamma$ is intersection-closed, otherwise not, unless the polynomial hierarchy collapses.

Note that prior to our work there was no conjecture on how a polynomial kernel dichotomy for all regular constraints may look like (we cannot offer such a conjecture for the more general case of user-independent constraints). The positive part follows by generalizing ideas of Crampton et al. [9]; the negative part is more challenging, and requires more involved arguments, especially to show the completeness of the dichotomy (i.e., that every relation R which is regular but not intersection-closed can be used in our lower bounds proof; see Section 4.3).

Organization. We define WSP formally and introduce a number of different constraint types, including regular constraints, in Section 2. In Section 3 we give several lower bounds for the kernelization of $\text{WSP}(\Gamma)$. In Section 4 we prove our main result, namely the dichotomy for regular constraints. We conclude in Section 5.

2 Preliminaries

We define a *workflow schema* to be a tuple (S, U, A, C) , where S is the set of tasks in the workflow, U is the set of users, $A: S \rightarrow 2^U$ assigns each task $s \in S$ an *authorization list* $A(s) \subseteq U$, and C is a set of workflow constraints. For the instance W^* of WSP of the previous section, $S = \{s_1, s_2, s_3\}$, $U = \{u_1, \dots, u_6\}$, $C = \{s_1 = s_2, s_1 \neq s_3, s_2 \neq s_3\}$, and $A(s_1) = \{u_1, u_2, u_3\}$, $A(s_2) = \{u_1, u_4, u_5\}$, $A(s_3) = \{u_1, u_6\}$. A *workflow constraint* is a pair $c = (L, \Theta)$, where $L \subseteq S$ is the *scope* of the constraint and Θ is a set of functions from L to U that specifies those assignments of elements of U to elements of

L that *satisfy* the constraint c . For the constraint $s_1 = s_2$ above, $L = \{s_1, s_2\}$ and $\Theta = \{L \rightarrow \{u_i\} : i \in \{1, \dots, 6\}\}$.

Given $T \subseteq S$ and $X \subseteq U$, a *plan* is a function $\pi: T \rightarrow X$; a plan $\pi: S \rightarrow U$ is called a *complete plan*. Given a workflow constraint (L, Θ) , $T \subseteq S$, and $X \subseteq U$, a plan $\pi: T \rightarrow X$ *satisfies* (L, Θ) if either $L \setminus T \neq \emptyset$, or π_L (π restricted to L) is contained in Θ , where $\pi_L: L \rightarrow U$ such that $\pi_L(s) = \pi(s)$ for all $s \in L$. A plan $\pi: T \rightarrow X$ is *eligible* if π satisfies every constraint in C . A plan $\pi: T \rightarrow X$ is *authorized* if $\pi(s) \in A(s)$ for all $s \in T$. A plan is *valid* if it is complete, authorized and eligible. For an algorithm that runs on an instance (S, U, A, C) of WSP, we will measure the running time in terms of $n = |U|$, $k = |S|$, and $m = |C|$.

2.1 WSP constraints and further notation

Let us first recall some concrete constraints that are of interest for this work:

$(=, T, T')$, (\neq, T, T') : These generalize the binary *binding-of-duty* and *separation-of-duty* constraints and were previously studied in [9,17]. They demand that there exist $s \in T$ and $s' \in T'$ which are assigned to the same (resp. different) users. We shorthand $(s = s')$ and $(s \neq s')$ if $T = \{s\}$ and $T' = \{s'\}$.

(t_ℓ, t_r, T) : A plan π satisfies (t_ℓ, t_r, T) , also called a *tasks-per-user counting constraint*, if a user performs either no tasks in T or between t_ℓ and t_r tasks.

Tasks-per-user counting constraints generalize the cardinality constraints which have been widely adopted by the WSP community [1,2,14,16].

$(\leq t, T)$, $(\geq t, T)$: These demand that the tasks in T are assigned to at most t (resp. at least t) different users. They generalize *binding-of-duty* and *separation-of-duty*, respectively, and enforce security and diversity [5].

All these constraints share the property that satisfying them depends only on the partition of tasks that is induced by the plan. This property is referred to as *user-independence*; see below.

Regular and user-independent constraints. Formally, a constraint (L, Θ) is *user-independent* if for any $\theta \in \Theta$ and any permutation $\psi: U \rightarrow U$, we have $\psi \circ \theta \in \Theta$. Note that this definition of user-independent constraints is equivalent to the (more informal) definition of these constraints given in the previous section.

For $T \subseteq S$ and $u \in U$ let $\pi: T \rightarrow u$ denote the plan that assigns every task of T to u . A constraint $c = (L, \Theta)$ is *regular* if it satisfies the following condition: For any partition L_1, \dots, L_p of L such that for every $i \in [p]$ there exists an eligible plan $\pi: L \rightarrow U$ and user u such that $\pi^{-1}(u) = L_i$, the plan $\bigcup_{i=1}^p (L_i \rightarrow u_i)$, where all u_i 's are distinct, is eligible. Consider, as an example, a tasks-per-user counting constraint (t_ℓ, t_r, L) . Let L_1, \dots, L_p be a partition of L such that for every $i \in [p]$ there exists an eligible plan $\pi: L \rightarrow U$ and user u such that $\pi^{-1}(u) = L_i$. Observe that for each $i \in [p]$, we have $t_\ell \leq |L_i| \leq t_r$

and so the plan $\bigcup_{i=1}^p (L_i \rightarrow u_i)$, where all u_i 's are distinct, is eligible. Thus, any tasks-per-user counting constraint (t_ℓ, t_r, L) is regular. Regular constraints are a special class of user-independent constraints, but not every user-independent constraint is regular. Crampton et al. [9] show that constraints of the type (\neq, T, T') are regular and so are constraints of the type $(=, T, T')$, where at least one of the sets T, T' is a singleton. In general, $(=, T, T')$ is not regular [9]: Consider, e.g., $(=, \{s_1, s_2\}, \{s_3, s_4\})$, where we have eligible plans for every choice of assigning some s_i to a private user, but assigning all four tasks to private users is ineligible.

Since regular constraints are of central importance to this paper, we introduce some further notation and terminology. Below, we generally follow Crampton et al. [9]. Let $W = (S, U, A, C)$ be a workflow schema, and π an eligible complete plan for W . Then \sim_π is the equivalence relation on S defined by π , where $s \sim_\pi s'$ if and only if $\pi(s) = \pi(s')$. We let S/π be the set of equivalence classes of \sim_π , and for a task $s \in S$ we let $[s]_\pi$ denote the equivalence class containing s .

For a constraint $c = (L, \Theta)$, a set $T \subseteq L$ of tasks is *c-eligible* if there is a plan $\pi: L \rightarrow U$ that satisfies c , such that $T \in L/\pi$. It is evident from the definition that c is regular if and only if the following holds: For every plan $\pi: L \rightarrow U$, π satisfies c if and only if every equivalence class $T \in L/\pi$ is *c-eligible*. In this sense, a regular constraint c is entirely defined by the set of *c-eligible* sets of tasks. It is clear that regular constraints are closed under conjunction, i.e., if C is a set of regular constraints on a set T of tasks, then the set of plans $\pi: T \rightarrow U$ which satisfy every $c \in C$ defines a new regular constraint (T, Θ) .

In a similar sense, if $c = (L, \Theta)$ is user-independent but not necessarily regular, then c can be characterized on the level of partitions of L : Let $\pi, \pi': L \rightarrow U$ be two plans such that $L/\pi = L/\pi'$. Then either both π and π' are eligible for c , or neither is. Overloading the above terminology, if c is a user-independent constraint, then we say that a partition L/π is *c-eligible* if a plan π generating the partition would satisfy the constraint. We may thus refer to the partition L/π itself as either eligible or ineligible. As with regular constraints, user-independent constraints are closed under conjunction.

Describing constraints via relations. We will frequently describe constraint types in terms of relations. In the following, we restrict ourselves to user-independent constraints. Let $R \subseteq \mathbb{N}^r$ be an r -ary relation, and $(s_1, \dots, s_r) \in S^r$ a tuple of tasks, with repetitions allowed (i.e., we may have $s_i = s_j$ for some $i \neq j$, $i, j \in [r]$). An *application* $R(s_1, \dots, s_r)$ (of R) is a constraint (L, Θ) where $L = \{s_i : i \in [r]\}$ and $\Theta = \{(\pi: L \rightarrow \mathbb{N}) : (\pi(s_1), \dots, \pi(s_r)) \in R\}$. Here, we identify users $U = \{u_1, \dots, u_n\}$ with integers $[n] = \{1, \dots, n\}$. We say that R is *user-independent (regular)* if every constraint $R(s_1, \dots, s_n)$ resulting from an application of R is user-independent (regular). In particular, a user-independent relation R can be defined on the level of partitions, in terms of whether each partition L/π of its arguments is eligible or not, and a regular relation can be defined in terms of eligible sets, as above.

Given a (possibly infinite) set Γ of relations as above, a *workflow schema over Γ* is one where every constraint is an application of a relation $R \in \Gamma$, and $\text{WSP}(\Gamma)$ denotes the WSP problem restricted to workflow schemata over Γ . To cover cases of constraints of unbounded arity, we allow Γ to be infinite.

For example, in the workflow schema W^* of the previous section every constraint is an application of binary relations $=$ and \neq .

Well-behaved constraint sets. To avoid several degenerate cases associated with infinite sets Γ we make some standard assumptions on our constraints. We say that a set Γ of user-independent relations is *well-behaved* if the following hold: (1) Every relation $R \in \Gamma$ can be encoded using $\text{poly}(r)$ bits, where r is the arity of R ; note that this does not include the space needed to specify the scope of an application of R . (2) For every application $c = (L, \Theta)$ of a relation $R \in \Gamma$, we can test in polynomial time whether a partition of L is c -eligible; we can also test in polynomial time whether a set $S \subseteq L$ is c -eligible, and if not, then we can (if possible) find a c -eligible set S' with $S \subset S' \subseteq L$. All relations corresponding to the concrete constraints mentioned above, are well-behaved.

2.2 Parameterized complexity and kernelization

A *parameterized problem* \mathcal{Q} is a subset of $\Sigma^* \times \mathbb{N}$ for some finite alphabet Σ . A parameterized problem \mathcal{Q} is *fixed-parameter tractable* (FPT) if there is a computable function $f: \mathbb{N} \rightarrow \mathbb{N}$ and an algorithm that, given (x, k) , takes time $f(k)|x|^{\mathcal{O}(1)}$ and correctly decides whether $(x, k) \in \mathcal{Q}$. A *kernelization* of \mathcal{Q} is a polynomial-time computable function $K: (x, k) \mapsto (x', k')$ such that $(x, k) \in \mathcal{Q}$ if and only if $(x', k') \in \mathcal{Q}$, and such that $|x'|, k' \leq h(k)$ for some $h(k)$. Here, (x, k) is an *instance* of \mathcal{Q} , and $h(k)$ is the *size* of the kernel. We say that K is a *polynomial kernelization* if $h(k) = k^{\mathcal{O}(1)}$. For an introduction to parameterized complexity we refer to, e.g., [12].

Our main tool for studying existence of polynomial kernels is kernelization-preserving reductions. Given two parameterized problems \mathcal{Q}_1 and \mathcal{Q}_2 , a *polynomial parametric transformation* (PPT) from \mathcal{Q}_1 to \mathcal{Q}_2 is a polynomial time computable function $\Psi: (x, k) \mapsto (x', k')$ such that for every input (x, k) of \mathcal{Q}_1 we have $(x', k') \in \mathcal{Q}_2$ if and only if $(x, k) \in \mathcal{Q}_1$, and such that $k' \leq p(k)$ for some $p(k) = k^{\mathcal{O}(1)}$. Note that if \mathcal{Q}_2 has a polynomial kernel and if there is a PPT from \mathcal{Q}_1 to \mathcal{Q}_2 , then \mathcal{Q}_1 has a *polynomial compression*, i.e., a kernel-like reduction to an instance of (possibly) a different problem with total output size $k^{\mathcal{O}(1)}$. Furthermore, for many natural problems (including all considered in this paper), we are able to complete these reductions using NP-completeness of \mathcal{Q}_1 to produce a polynomial kernel for \mathcal{Q}_1 . Conversely, by giving PPTs from problems that are already known not to admit polynomial compressions (under some assumption) we rule out polynomial kernels for the target problems. For more background on kernelization we refer the reader to the recent survey by Lokshtanov et al. [15].

3 Lower bounds for kernelization

In this section we begin our investigation of the preprocessing properties of the WORKFLOW SATISFIABILITY PROBLEM. We establish lower bounds against polynomial kernels for WSP for several widely-used constraint types. Like for many other problems, e.g., HITTING SET(n) or CNF SAT(n), there is little hope to get polynomial kernels for WSP when we allow an unbounded number of constraints of arbitrary arity, cf. [10, 11, 13]. As an example, we give Lemma 1, whose proof uses a PPT from CNF SAT(n) to WSP(≥ 2) with only two users.

Lemma 1 *Let WSP(≥ 2) be the WSP problem with constraints ($\geq 2, L$) for task sets L of arbitrary arity. Then WSP(≥ 2) admits no polynomial kernelization with respect to the number k of tasks unless the polynomial hierarchy collapses, even if the number of users is restricted to $n = 2$.*

Proof We give a PPT from SAT(n), i.e., SAT parameterized by the number n of variables. The fact that this problem admits no polynomial kernelization or compression is due to work of Dell and van Melkebeek [10]. Let an instance ϕ of SAT(n) be given, and let n denote the number of variables in the CNF-formula ϕ . For ease of presentation, let the variables of ϕ be x_1, \dots, x_n . (To recall, a CNF-formula is a conjunction of clauses, each of which is a disjunction of literals x_i or $\neg x_i$.)

We construct a WSP(≥ 2) instance with two users t and f , which intuitively represent true and false assignment to literals of ϕ . For ease of reading, we state all constraints ($\geq 2, L$) by simply declaring the sets L to which they are applied. To begin, for each variable x_i we create two tasks s_i and \bar{s}_i and a constraint $\{s_i, \bar{s}_i\}$; both users are authorized for all these tasks. Intuitively, assigning users to s_i and \bar{s}_i corresponds to setting x_i and $\neg x_i$ to true or false; the constraint $\{s_i, \bar{s}_i\}$ ensures that exactly one is true and one is false. Additionally, add one more task d for which only user f is authorized. For each clause c of ϕ we create the following set L_c : Add to L_c the task s_i if c contains a literal x_i , and the task \bar{s}_i if c contains a literal $\neg x_i$. Additionally add the task d .

If ϕ has a satisfying assignment, then we get a valid plan π by $\pi(d) = f$ and

$$\pi(s_i) = \begin{cases} t & \text{if } x_i \text{ is true,} \\ f & \text{if } x_i \text{ is false,} \end{cases} \quad \pi(\bar{s}_i) = \begin{cases} f & \text{if } x_i \text{ is true,} \\ t & \text{if } x_i \text{ is false.} \end{cases}$$

For each constraint L_c , any satisfying assignment for c must assign true to some literal $\ell_j \in \{x_i, \neg x_i\}$ for some x_i . This corresponds directly to a task s_i or \bar{s}_i in L_c which is assigned user t . Since d is always assigned user f , this fulfills ($\geq 2, L_c$).

Conversely, let π be a valid plan for the created WSP(≥ 2) instance. We create an assignment for $\{x_1, \dots, x_n\}$ by setting x_i to true if $\pi(s_i) = t$ and to false otherwise. We already argued earlier that $\pi(s_i) \neq \pi(\bar{s}_i)$ due to constraint $\{s_i, \bar{s}_i\}$. Now let c be any clause of ϕ . In the corresponding set L_c we

have the task d which must be assigned user f . Due to the constraint $(\geq 2, L_c)$ at least one other task in L_c must be assigned user t . If this is a task s_i , then c contains x_i and our defined assignment sets x_i to true, satisfying c . If this is a task \bar{s}_i then c contains $\neg x_i$. We know that $\pi(\bar{s}_i) = t$, so $\pi(s_i) = f$, implying that our assignment sets x_i to false, satisfying $\neg x_i$ and c .

Thus, our reduction is correct. It is easy to see that the construction can be performed in polynomial-time and that the number of tasks is $2n+1 \leq \text{poly}(n)$. Thus, the PPT from $\text{SAT}(n)$ to $\text{WSP}(\geq 2)$ proves that the latter problem has no polynomial kernelization unless the polynomial hierarchy collapses. \square

In our further considerations we will avoid cases like the above, by either taking m as an additional parameter or by restricting Γ to be finite, which implies bounded arity (namely the maximum arity over the finitely many $R \in \Gamma$). We also assume that all constraints are well-behaved (cf. Section 2.1). We then have the following, showing that bounding the number of users implies a polynomial kernel.

Proposition 1 *Let Γ be a set of relations. If Γ is finite, then $\text{WSP}(\Gamma)$ has a polynomial kernel under parameter $(k+n)$; if Γ is infinite but Γ is well-behaved, then $\text{WSP}(\Gamma)$ has a polynomial kernel under parameter $(k+m+n)$.*

Proof An instance of $\text{WSP}(\Gamma)$ is defined by describing its tasks, users, authorization lists and constraints. Note that the former three can be written down in space $\mathcal{O}(kn)$, hence it remains to describe the constraints. If Γ is finite, then there is a maximum arity r of any relation in Γ , and at most $|\Gamma| \cdot k^r = k^{\mathcal{O}(1)}$ possible constraints can be defined as applications of relations $R \in \Gamma$, and each constraint can be defined in short space (e.g., by an index into Γ and $\mathcal{O}(r \log k)$ bits giving the scope of the constraint). Hence all constraints can be described in space polynomial in $k+n$.

If Γ is infinite but well-behaved, then under the parameter $k+m+n$ it suffices to be able to describe each constraint in space $\text{poly}(n+k)$; this is possible by assumption. Hence in both cases, there is a simple encoding of an instance in space $\text{poly}(k+n)$ resp. $\text{poly}(k+n+m)$. \square

The following lemma addresses a special case of ternary constraint $R(a, b, c)$ and proves that already $\text{WSP}(R)$ admits no polynomial kernelization in terms of $k+m$. This lemma will be a cornerstone of the dichotomy in the following section. We also get immediate corollaries ruling out polynomial kernels in $k+m$ for constraints $(=, S, S')$ and $(\leq t, S)$, since $(=, \{a\}, \{b, c\})$ and $(\leq 2, \{a, b, c\})$ fulfill the requirement of the lemma.

The proof will be by a PPT from the problem MULTI-COLORED HITTING SET(m) (MCHS), which was considered in [11, 13]. The input is a vertex set V , a collection $\mathcal{H} = \{E_1, \dots, E_m\}$ of subsets of V , an integer ℓ , and a function $\phi: V \rightarrow [\ell]$ which *colors* each vertex of V in one of ℓ colors. The task is to find a set $Q \subseteq V$ containing exactly one vertex of each color such that $Q \cap E_i \neq \emptyset$ for every $E_i \in \mathcal{H}$. It follows from work of Dom et al. [11] that this problem admits no polynomial kernel or compression under the parameter

m unless the polynomial hierarchy collapses: A simple PPT from HITTING SET(m) works by making ℓ copies of each element and giving each copy a different color. Since instances with $\ell > m$ are trivial for HITTING SET(m) we may restrict MCHS to $\ell \leq m$ without harming this lower bound. Furthermore, the problem is complete for a kernelization hardness class known as WK[1], which is conjectured to imply further lower bounds [13].

We now proceed with the PPT. Since we will refer to this type of relation several times, let us give it a name. Let a *switching relation* be a ternary user-independent constraint $R(a, b, c)$ which is satisfied by plans with induced partition $\{\{a, b\}, \{c\}\}$ or $\{\{a, c\}, b\}$, but not by plans with partition $\{\{a\}, \{b\}, \{c\}\}$. Our result is the following.

Lemma 2 *Let $R(a, b, c)$ be a switching relation. Then $\text{WSP}(R)$ does not admit a polynomial kernel with respect to parameter $k + m$ unless the polynomial hierarchy collapses.*

Proof We give a PPT from MULTI-COLORED HITTING SET(m), described above. Let an instance $(V, \mathcal{H}, \ell, \phi)$ of MCHS be given, where $\mathcal{H} = \{E_1, \dots, E_m\}$ with $E_i \subseteq V$, $\ell \leq m$, and $\phi: V \rightarrow [\ell]$. Let $V_j = \phi^{-1}(j)$, $j \in [\ell]$. We may assume that $m \geq 2$, or else solve the instance in polynomial time and return a corresponding dummy yes- or no-instance, and that $V_j \neq \emptyset$ for each $j \in [\ell]$. A solution to MCHS is now a multi-colored hitting set $Q = \{v_1, \dots, v_\ell\}$ where $v_i \in V_i$ for each $i \in [\ell]$ and $Q \cap E_i \neq \emptyset$ for every $i \in [m]$.

We start by letting the set of users be $U := V$. We make $(\ell - 1) \cdot m$ tasks

$$e_{1,2}, \dots, e_{1,\ell}, \quad e_{2,2}, \dots, e_{2,\ell}, \quad \dots, \quad e_{m,2}, \dots, e_{m,\ell},$$

i.e., $(\ell - 1)$ tasks $e_{i,2}, \dots, e_{i,\ell}$ for each set E_i . For every $i \in [m]$, let the authorization of $e_{i,\ell}$ be E_i ; all remaining tasks $e_{i,j}$ get authorization U . Furthermore, we introduce tasks s_1, \dots, s_ℓ that are intended for choosing a hitting set of size ℓ and that have authorizations $A(s_j) := V_j$.

We introduce the following constraints for each $i \in \{1, \dots, m\}$:

1. Introduce $R(e_{i,2}, s_1, s_2)$.
2. For all $j \in \{3, \dots, \ell\}$ introduce $R(e_{i,j}, e_{i,j-1}, s_j)$.

Clearly, this construction can be performed in polynomial time and the parameter values are number of tasks $k = (\ell - 1) \cdot m + \ell = \mathcal{O}(m^2)$ and number of constraints $m' = (\ell - 1) \cdot m = \mathcal{O}(m^2)$. Thus, to prove that this constitutes a PPT it remains to prove that the created $\text{WSP}(R)$ instance has a valid plan π if and only if (V, \mathcal{H}) has a multi-colored hitting set.

Let $S \subseteq V$ be a multi-colored hitting set for (V, \mathcal{H}) . For each color j let v_j be the unique element in $S \cap V_j$. We now create π and begin with $\pi(s_j) := v_j$, consistent with the authorization of s_j . We now define the user assignment for tasks associated with some set $E_i \in \mathcal{H}$. First of all, we note that $S \cap E_i \neq \emptyset$ and thus we can arbitrarily select $v_t \in S \cap E_i$. We now assign users as follows for $j \in \{2, \dots, \ell\}$:

$$\pi(e_{i,j}) = \begin{cases} v_1 & \text{if } j < t, \\ v_t & \text{if } j \geq t. \end{cases}$$

We note that $v_t \in E_i$ and thus v_t is authorized for task $e_{i,\ell}$; all other tasks are authorized for all users anyway. It can be easily verified that $R(e_{i,2}, s_1, s_2)$ and $R(e_{i,j}, e_{i,j-1}, s_j)$ for $j \in \{3, \dots, \ell\}$ hold for this assignment of users:

- $R(e_{i,2}, s_1, s_2)$ is satisfied because π gives partition $\{\{e_{i,2}, s_2\}, \{s_1\}\}$ if $t = 2$ and partition $\{\{e_{i,2}, s_1\}, \{s_2\}\}$ otherwise.

It remains to check $R(e_{i,j}, e_{i,j-1}, s_j)$ for $j \in \{3, \dots, \ell\}$.

- If $j \in \{3, \dots, t-1\}$ then $\pi(e_{i,j-1}) = v_1$, $\pi(e_{i,j}) = v_1$, and $\pi(s_j) = v_j$, satisfying $R(e_{i,j}, e_{i,j-1}, s_j)$ with partition $\{\{e_{i,j}, e_{i,j-1}\}, \{s_j\}\}$.
- If $j = t$ then $\pi(e_{i,j-1}) = v_1$, $\pi(e_{i,j}) = v_t$, and $\pi(s_j) = v_j = v_t$, satisfying $R(e_{i,j}, e_{i,j-1}, s_j)$ with partition $\{\{e_{i,j}, s_j\}, \{e_{i,j-1}\}\}$.
- If $j \in \{t+1, \dots, \ell\}$ then $\pi(e_{i,j-1}) = v_t$, $\pi(e_{i,j}) = v_t$, and $\pi(s_j) = v_j \neq v_t$, satisfying $R(e_{i,j}, e_{i,j-1}, s_j)$ with partition $\{\{e_{i,j}, e_{i,j-1}\}, \{s_j\}\}$.

For the converse, assume that π is a valid plan for the created $\text{WSP}(R)$ instance. For $j \in \{1, \dots, \ell\}$ let $v_j := \pi(s_j)$ and note that $v_j \in A(s_j) = V_j$. We claim that $S := \{v_1, \dots, v_\ell\}$ is a multi-colored hitting set for (V, \mathcal{H}) according to ϕ . Consider any set $E_i \in \mathcal{H}$ and recall that $A(e_{i,\ell}) = E_i$. We claim $\pi(e_{i,\ell}) \in S$, which would imply that $S \cap E_i \neq \emptyset$. To prove this, we prove inductively that $\pi(e_{i,j}) \in \{v_1, \dots, v_j\}$ for $j \in \{2, \dots, \ell\}$. (Recall that authorizations for all these other tasks associated with E_i are simply $U = V$, i.e., we really need the property for $j = \ell$ and $e_{i,\ell}$.)

1. For $j = 2$ consider the constraint $R(e_{i,2}, s_1, s_2)$. Since $A(s_1) \cap A(s_2) = V_1 \cap V_2 = \emptyset$ we know that $\pi(s_1) \neq \pi(s_2)$. Furthermore, by assumption of the lemma $R(e_{i,2}, s_1, s_2)$ is not satisfied if π induces partition $\{\{e_{i,2}\}, \{s_1\}, \{s_2\}\}$. Since $\pi(s_1) \neq \pi(s_2)$ we must have $\pi(e_{i,2}) \in \{\pi(s_1), \pi(s_2)\} = \{v_1, v_2\}$. This proves our claim for $j = 2$.
2. Now consider some $j \geq 3$ such that the claim holds for all smaller j . In particular $\pi(e_{i,j-1}) \in \{v_1, \dots, v_{j-1}\}$. Note that $\{v_1, \dots, v_{j-1}\} \subseteq V_1 \cup \dots \cup V_{j-1}$ which has an empty intersection with V_j . Thus, considering the constraint $R(e_{i,j}, e_{i,j-1}, s_j)$ we find that $\pi(e_{i,j-1}) \neq \pi(s_j)$. Thus, by the same argument as in the previous item we find that $\pi(e_{i,j}) \in \{\pi(e_{i,j-1}), \pi(s_j)\} \subseteq \{v_1, \dots, v_j\}$. This completes our claim.

It follows that $\pi(e_{i,\ell}) \in S$, which implies that $S \cap E_i = S \cap A(e_{i,\ell}) \supseteq \{\pi(e_{i,\ell})\} \neq \emptyset$. Thus, S is indeed a multi-colored hitting set for (V, \mathcal{H}) according to ϕ . Since $\text{MULTI-COLORED HITTING SET}(m)$ is known not to admit a polynomial compression unless the polynomial hierarchy collapses, this completes the proof. \square

Since both $(=, \{a\}, \{b, c\})$ and $(\leq 2, \{a, b, c\})$ fulfill the requirement of the lemma, we get the following corollary.

Corollary 1 $\text{WSP}(=, S, S')$ and $\text{WSP}(\leq t, S)$ do not admit a kernelization to size polynomial in $k + m$ unless the polynomial hierarchy collapses.

4 A dichotomy for regular constraints

In this section, we present a dichotomy for the kernelization properties of $\text{WSP}(\Gamma)$ when Γ is a well-behaved set of regular relations.

Let us describe the dichotomy condition. Let $c = (L, \Theta)$ be a regular constraint, and $E \subseteq 2^L$ the set of c -eligible subsets of L ; for ease of notation, we let $\emptyset \in E$. Note that by regularity, E defines c . We say that c is *intersection-closed* if for any $T_1, T_2 \in E$ it holds that $T_1 \cap T_2 \in E$. Similarly, we say that a regular relation $R \in \Gamma$ is intersection-closed if every application $R(s_1, \dots, s_r)$ of R is. Note that the conjunction of intersection-closed constraints again defines an intersection-closed constraint (this is readily verified from the definitions). Since $(s = s')$ is regular and intersection-closed, this implies that R is intersection-closed if and only if every application $R(s_1, \dots, s_r)$ of R with r distinct tasks s_i is intersection-closed. Finally, a set Γ of relations is intersection-closed if every relation $R \in \Gamma$ is. Our dichotomy results will essentially say that $\text{WSP}(\Gamma)$ admits a polynomial kernel if and only if Γ is intersection-closed; see Theorem 1 below.

The rest of the section is laid out as follows. In Section 4.1 we show that if Γ is regular, intersection-closed, and well-behaved, then $\text{WSP}(\Gamma)$ admits a reduction to $n' \leq k$ users; by Prop. 1, this implies a polynomial kernel under parameter $(k + m)$, and under parameter (k) if Γ is finite. Section 4.2 is a short technical section, where we introduce a notion of user-independent relation *implemented* by a set of relations and prove Lemma 7, both important for Section 4.3. In Section 4.3 we show that for any single relation R that is not intersection-closed, the problem $\text{WSP}(R)$ admits no polynomial kernel, by application of Lemma 2. In Section 4.4 we consider the implications of these results for the existence of efficient user-reductions.

In summary, we will show the following result for kernelization. Again, a discussion of the consequences for user-reductions is deferred until Section 4.4.

Theorem 1 *Let Γ be a possibly infinite set of well-behaved regular relations. If every relation in Γ is intersection-closed, then $\text{WSP}(\Gamma)$ admits a polynomial-time many-one reduction down to $n' \leq k$ users, implying a polynomial kernel under parameter $k + m$ (and a polynomial kernel under parameter k if Γ is finite). Otherwise, $\text{WSP}(\Gamma)$ admits no kernel of size $\text{poly}(k + m)$ unless the polynomial hierarchy collapses (even if Γ consists of a single such relation R).*

4.1 A user reduction for intersection-closed constraints

We now give a procedure that reduces a WSP instance $W = (S, U, A, C)$ with n users, k tasks and m constraints to one with $k' \leq k$ tasks, $n' \leq k'$ users and $m' \leq m$ constraints, under the assumption that every constraint $c \in C$ occurring in the instance is intersection-closed and that our language is well-behaved. The approach is as in Crampton et al. [9, Theorem 6.5], but becomes more involved due to having to work in full generality; we also use a more

refined marking step that allows us to decrease the number of users from k^2 to k , a significant improvement. As noted (Prop. 1), under the appropriate further assumption on the constraints, this gives a polynomial kernel under parameter $k + m$ or k .

We begin by noting a consequence of sets closed under intersection.

Lemma 3 *Let $c = (L, \Theta)$ be an intersection-closed constraint, and let $T \subseteq L$ be c -ineligible. If there is a superset T' of T which is c -eligible, then there is a task $s \in L \setminus T$ such that every c -eligible superset T' of T contains s .*

Proof Let T_\cap be the intersection of all c -eligible supersets $T' \supset T$. Then T_\cap must itself be c -eligible. We clearly have $T \subseteq T_\cap$, and since T is c -ineligible the containment must be strict. Hence there is some $s \in T_\cap \setminus T$; this task s must be contained in all c -eligible supersets of T . \square

We refer to the task s guaranteed by the lemma as a *required addition* to T by c . Note that assuming well-behavedness, we can make this lemma constructive, i.e., in polynomial time we can test whether a set T is eligible for a constraint, whether it has an eligible superset, and find all required additions if it does. This can be done by first asking for an eligible superset T' of T , then greedily finding a minimal set $T \subset T'' \subseteq T'$. Then every $s \in T'' \setminus T$ is a required addition.

Our reduction proceeds in three phases. First, we detect all binary equalities implied by the constraints, i.e., all explicit or implicit constraints ($s = s'$), and handle them separately by merging tasks, intersecting their authorization lists. The output of this phase is an instance where any plan which assigns to every task a unique user is eligible (though such a plan may not be authorized); in particular, since our constraints are regular, we have that all singleton sets of tasks are c -eligible for every constraint c of the instance.

The second phase of the kernel is a user-marking process, similar to the kernels in [9] but with a stronger bound on the number of users. This procedure is based around attempting to produce a system of *distinct representatives* for $\{A(s) : s \in S\}$, i.e., to find a plan $\pi: S \rightarrow U$ such that π is authorized and $\pi(s) \neq \pi(s')$ for every $s \neq s'$. Via Hall's theorem, this procedure either succeeds, or produces a set T of tasks such that fewer than $|T|$ users are authorized to perform any task in T . In the latter case, we mark all these users, discard the tasks T , and repeat the procedure. Eventually, we end up with a (possibly empty) set of tasks S' which allows for a set of distinct representatives, and mark these representatives as well. Refer to a task s as *easy* if it was appointed a representative in this procedure, and *hard* if it was not (i.e., if it was a member of a set T of discarded tasks). We discard every non-marked user, resulting in a partially polynomial kernel with $k' \leq k$ tasks and $n' \leq k' \leq k$ users.

Finally, to establish the correctness of the kernelization, we give a procedure that, given a partial plan for the set of hard tasks, either extends the plan to a valid plan or derives that no such extended plan exists.

Lemma 4 *Let (S, U, A, C) be a workflow schema with k tasks, n users, and m constraints, with at least one equality constraint $(s = s')$, $s \neq s'$. In polynomial time, we can produce an equivalent instance (S', U', A', C') with at most $k - 1$ tasks, n users, and m constraints. Furthermore, if the constraints in C were given as applications $R(\dots)$ of some relations R , $R \in \Gamma$, then the constraints in C' can be given the same way.*

Proof Drop the constraint $(s = s')$. For every other constraint $c = (L, \Theta)$ with s' in the scope, replace c by the corresponding constraint produced by replacing s' by s . (If $c = R(s_1, \dots, s_r)$ for some relation R , then this produces a new application $R(\dots)$ of the same relation R . This application may contain the task s in more than one position, however, this is allowed by our model of constraints.) Update the authorization list so that $A(s) := A(s) \cap A(s')$. Finally, discard the task s' . The new instance has a valid plan if and only if the old instance does. \square

We now show the detection of equalities.

Lemma 5 *Let (S, U, A, C) be a workflow schema where every constraint is regular and intersection-closed. Then we can in polynomial time reduce the instance to the case where every singleton $\{s\}$, $s \in S$, is eligible.*

Proof We provide a procedure that, using calls to Lemma 3, detects all equality constraints $(s = s')$ implied by the schema, and applies Lemma 4 for every such constraint found. The procedure goes as follows:

1. For every task $s \in S$, check whether the singleton set $\{s\}$ is c -eligible for every constraint c in the instance.
2. If all such singleton task sets are eligible, then we are done.
3. Otherwise, let $\{s\}$ be ineligible for some constraint $c = (L, \Theta)$, and let s' be a required addition to $\{s\}$ by c . Apply Lemma 4 to the constraint $(s = s')$ and restart from Step 1.

We now show the correctness of the procedure. On the one hand, it is clear that by the termination of the procedure, no equality constraints can remain, implicit or explicit, since a constraint $(s = s')$ contradicts that the singletons $\{s\}$ and $\{s'\}$ are eligible (or, more formally, that the partition $\{\{s\} : s \in S'\}$ is eligible for the remaining tasks S'). On the other hand, if s' is a required addition to $\{s\}$ by some constraint c , then c must imply the constraint $(s = s')$, since in this case, by Lemma 3, every c -eligible set containing s also contains s' . \square

Next, we describe the user-marking procedure in detail. We assume that Lemma 5 has been applied, i.e., that all singleton sets are eligible.

1. Let $M = \emptyset$, let S be the set of all tasks, and U the set of all users.
2. While $\{A(s) \cap U : s \in S\}$ does not admit a system of distinct representatives: Let $T \subseteq S$ such that $|\bigcup_{s \in T} A(s)| < |T|$. Let $U_T = \bigcup_{s \in T} A(s)$. Add U_T to M , remove U_T from U , and remove from S every task s such that $A(s) \subseteq M$.

3. Add to M the distinct representatives of the remaining tasks S , if any.
4. Discard all users not occurring in M from the instance.

We refer to the set M of users produced above as the *marked* users, and let $S_{\text{hard}} \subseteq [k]$ be the set of hard tasks, i.e., the set of tasks removed in Step 2 of the procedure. Finally, we show the correctness of the above procedure.

Lemma 6 *Let (S, U, A, C) be a workflow schema where all constraints are regular and intersection-closed, and where all singleton sets are eligible. There is a valid plan for the instance if and only if there is a valid plan only using marked users.*

Proof We describe a procedure that, for any eligible and authorized partial plan assigning users to the hard tasks, either produces a valid plan using only marked users or proves that the plan cannot be extended to a complete plan. Note that necessarily, the partial plan we begin with can use only marked users.

The algorithm works as follows. Let π be the partial plan.

1. Let $\mathcal{P} = S_{\text{hard}}/\pi$ be the partition induced by the partial plan π .
2. Repeat the following until all sets $T \in \mathcal{P}$ are c -eligible for every $c \in C$:
 - (a) Let T be an ineligible set in \mathcal{P} , and let s be a required addition to T .
 - (b) If s is already assigned by π , or if $A(s) \not\subseteq \pi(s')$, $s' \in T$, then reject.
 - (c) Otherwise, add s to T in \mathcal{P} and update π (i.e., $\pi(s) = \pi(s')$, $s' \in T$).
3. Pad the partition \mathcal{P} with singleton task sets, i.e., let every unassigned task be performed by a user which has no other duties. Update π accordingly to a complete plan using the distinct representatives for the remaining tasks.

We show correctness of this procedure. First, by Lemma 3, clearly the modification in Steps 2a–2c are necessary for any valid plan. It also follows that any rejection performed during these steps will be correct. Second, the padding in Step 3 is possible, since the users appointed as distinct representatives will be used for no other task. Hence, we can create a complete plan π corresponding to the resulting padded partition \mathcal{P} , and this plan will be valid, as every task set $T \in \mathcal{P}$ is either a singleton set, and hence c -eligible for every c , or the output of the loop 2a–2c and thus c -eligible for every c by assumption; furthermore, authorization was checked at every step. \square

Putting the above pieces together yields the following theorem.

Theorem 2 *The WORKFLOW SATISFIABILITY PROBLEM, restricted to well-behaved constraint languages where every constraint is regular and intersection-closed admits a kernel with $m' \leq m$ constraints, $k' \leq k$ tasks, and $n' \leq k'$ users.*

Proof The polynomial running time and the bound $k' \leq k$ are immediate from the above, and the correctness has already been argued. It remains to show that the number of marked users is at most equal to the number of tasks. This follows inductively, since the final addition to M , in Step 3 of the marking

procedure, adds exactly $|S|$ users, and every previous addition, in an iteration of Step 2, adds fewer users to M than the number of tasks in T , all of which will be removed from the set S . Finally, the result is a kernel by Prop. 1. \square

4.2 Implementations and implications

Let $W = (S, U, A, C)$ be a workflow schema and $T \subseteq S$ a set of tasks. The *projection of W onto T* is a constraint $c = (T, \Theta)$ where $\theta: T \rightarrow U$ is contained in Θ if and only if there is a valid plan π for W that extends θ . Consider the WSP instance W^* introduced in Section 1 and let $T = \{s_2, s_3\}$. Since the only valid plan is π with $\pi(s_1) = \pi(s_2) = u_1$ and $\pi(s_3) = u_6$, the projection of W onto T is $(T, \{\theta\})$, where θ is defined by $\theta(s_2) = u_1, \theta(s_3) = u_6$.

Further, let $R \subseteq \mathbb{N}^r$ be a user-independent relation, $Q = \{q_1, \dots, q_r\}$ a set of r distinct tasks, and Γ a set of relations. We say that Γ *implements* R if, for any r -tuple $\mathcal{A} = (A(q_1), \dots, A(q_r))$ of authorization lists, there is a workflow schema $W = (S, U, A, C)$ over Γ that can be computed in polynomial time, such that $Q \subseteq S$ and for every plan $\pi: Q \rightarrow U$, either π is not authorized with respect to \mathcal{A} , or π is eligible for the projection of W onto Q if and only if it is eligible for $R(q_1, \dots, q_r)$. Furthermore, $|S| + |C|$ does not depend on \mathcal{A} , and U equals $\bigcup_{i \in [r]} A(q_i)$ plus a constant number of local users, i.e., new users who will not be authorized to perform any task outside of $S \setminus Q$.

Lemma 7 *Let Γ and Γ' be finite workflow constraint languages such that Γ' implements R for every $R \in \Gamma$. Then there is a PPT from $\text{WSP}(\Gamma)$ to $\text{WSP}(\Gamma')$, both with respect to parameter k and $k + m$.*

Proof Let $W = (S, U, A, C)$ be an instance of $\text{WSP}(\Gamma)$. We will create an equivalent instance W' of $\text{WSP}(\Gamma')$, which we will refer to as the *output* of the reduction. For every constraint $c \in C$ which is an application of a relation $R \in \Gamma \setminus \Gamma'$, let W_c be an implementation of R (using the relevant authorization lists $A(s)$ from A). Add all constraints, tasks and users of W_c to the output such that all tasks respectively users local to W_c are chosen distinct from existing task respectively users. Also add to the output every constraint $c \in C$ which is an application of a relation $R \in \Gamma \cap \Gamma'$. Clearly, this creates an output instance W' of $\text{WSP}(\Gamma')$ which is equivalent to W , and which is computed in polynomial time. Furthermore, if W contains k tasks, n users and m constraints, then W' contains $\mathcal{O}(k + m)$ tasks, $\mathcal{O}(m)$ constraints, and $\mathcal{O}(n + m)$ users. Hence it is a PPT under parameter $k + m$. Finally, since Γ is finite, we have $m = k^{\mathcal{O}(1)}$ and the reduction is a PPT under parameter k . \square

4.3 Kernel lower bounds for not intersection-closed constraints

We now give the other side of the dichotomy by showing that within the setting of regular constraints, even a single relation R which is not intersection-closed can be used to construct a kernelization lower bound, following one of the

constructions in Section 3. First, we need an auxiliary lemma. A relation R is *satisfiable* if it is satisfied by at least one plan.

Lemma 8 *Let R be a (satisfiable) regular relation R which is not intersection-closed, let $=$ be the binary equality relation, and let $\Gamma = \{R, =\}$. Then either Γ implements a switching relation, or Γ implements the binary disequality relation \neq .*

Proof Assume that R is r -ary, and let $L = \{s_1, \dots, s_r\}$ be a set of r distinct tasks. We consider an application $c = R(s_1, \dots, s_r)$ of R . Let E_R be the c -eligible subsets of L . First, if possible, let $T, T' \in E_R$ be disjoint sets such that $T \cup T' \notin E_R$, and let $\pi: L \rightarrow U$ be a c -eligible plan with $T, T' \in L/\pi$. Merge the tasks T into a single task s and the tasks T' into a single task t by applications of $=$ (e.g., add constraints $(s_i = s_j)$ for every pair $s_i, s_j \in T$ and every pair $s_i, s_j \in T'$). Let $A(s)$ and $A(t)$ be arbitrary but non-empty, and let $A(s') = U_\perp$ for any other task s' , where U_\perp is a sufficiently large supply of dummy users (i.e., a set of users who are only authorized to perform these tasks $s' \in L$). Then the resulting workflow schema W has a valid plan π where $\pi(s) = u_s$ and $\pi(t) = u_t$ for any $u_s \in A(s)$ and $u_t \in A(t)$ with $u_s \neq u_t$, but no such plan where $u_s = u_t$. Thus, the projection of W onto $\{s, t\}$ implements \neq and we are done.

Second, if not, then we have that in any c -eligible partition L/π , we may freely merge parts to create a coarser c -eligible partition. In particular, for any $T \in E_R$ we have $(L \setminus T) \in E_R$. If there is any pair of sets $T, T' \in E_R$ such that $T \cup T' \notin E_R$, then we proceed as follows. Let $P = L \setminus T$ and $Q = L \setminus T'$; then $P, Q \in E_R$ but $P \cap Q \notin E_R$. Merge each of the sets $P \cap Q, P \setminus Q, Q \setminus P$, and $L \setminus (P \cup Q)$ into single tasks, respectively, a, b, c and d . Let $R(a, b, c, d)$ be the resulting relation. Then R defines a regular relation where $\{a\}$ is ineligible and where partitions $\{\{a, b\}, \{c, d\}\}$ and $\{\{a, c\}, \{b, d\}\}$ are both eligible (by the eligible partitions $\{P, L \setminus P\}$ and $\{Q, L \setminus Q\}$). We will implement a switching relation.

Consider first the relation $R'(a, b, c) = \exists d : R(a, b, c, d)$. If $\{\{a\}, \{b\}, \{c\}\}$ is an ineligible partition for R' , then R' is a switching relation, and we get that Γ implements R' by constructing a workflow schema as above (using equality constraints to merge tasks) and projecting onto $\{a, b, c\}$. Otherwise, we must have $\{a, d\} \in E_R$, hence also $\{b\}, \{c\} \in E_R$. If $\{d\} \in E_R$, then we may restrict $A(d)$ to a supply of dummy users and again implement a switching relation $R'(a, b, c)$ by projecting onto $\{a, b, c\}$; note that this is allowed in our definition of implementation, and that this choice of $A(d)$ eliminates the partition $\{\{a\}, \{b\}, \{c\}\}$ from R' . In the remaining case, using that E_R is closed under complementation, we find that R is characterized precisely by excluding $\{a\}$ and $\{d\}$ as eligible sets. In this case, $R'(a, b, c) = \exists d : R(b, a, d, c)$ is a switching relation (using $A(d) = A(b) \cup A(c)$), which again can be implemented as above using equality constraints and projection onto $\{a, b, c\}$. Hence Γ implements R' .

Lastly, we have the case that E_R is closed under arbitrary union and under complementation. Then E_R is also under intersection, as $T \cap T' = \overline{\overline{T} \cup \overline{T'}}$. This contradicts our assumptions about R . \square

Using \neq , we can more easily construct a switching relation $R(a, b, c)$.

Lemma 9 *Let R be a regular relation which is not intersection-closed, and let $=$ and \neq denote the binary equality and disequality relations. Then $\Gamma = \{R, =, \neq\}$ implements a switching relation $R'(a, b, c)$.*

Proof Let $c = R(L)$ be an application of R on task set L , and let $E_R \subseteq 2^L$ be the set of c -eligible subsets of L . Say that R is a *counterexample relation* with respect to tasks $a, b, c \in L$ if there is a set $P \in E_R$ with $P \cap \{a, b, c\} = \{a, b\}$ and a set $Q \in E_R$ with $Q \cap \{a, b, c\} = \{a, c\}$, but $P \cap Q \notin E_R$. We will modify R into a simpler counterexample relation, by imposing further regular constraints using binary relations $=$ and \neq , to finally produce the sought-after switching relation $R'(a, b, c)$.

The first modification we consider is to merge tasks $s, s' \in L$. If there is any pair of tasks s, s' such that merging s and s' (i.e., adding an equality constraint ($s = s'$)) still yields a counterexample relation, then merge s and s' ; repeat this exhaustively. For the rest of the proof, we will treat merged tasks as a single task, and assume that R is a minimal counterexample with respect to merging operations. In particular, this implies $|P \cap Q| = 1$ for any sets P and Q as above.

Next, we impose a set of binary disequality constraints ($s \neq s'$). Let P and Q be sets as above, let $P \cap Q = \{a\}$, and let $P' \subseteq P$ be a minimal eligible set with $a \in P'$. Similarly let $Q' \subseteq Q$ be a minimal eligible set with $a \in Q'$. Impose a binary disequality constraint ($a \neq s$) for any $s \in L \setminus (P' \cup Q')$, and ($b' \neq c'$) for any $b' \in P' - a$ and $c' \in Q' - a$. We argue that the resulting workflow, projected down to tasks a, b'', c'' for $b'' \in P' - a$ and $c'' \in Q' - a$, implements a switching relation. It is clear by P' and Q' that the partitions $\{\{a, b''\}, \{c''\}\}$ and $\{\{a, c''\}, \{b''\}\}$ are eligible; the added disequality constraints have no effect on this. On the other hand, there is no eligible partition where a, b'', c'' are all contained in different sets, as due to the disequality constraints, a can only be contained in a set contained in either P' or in Q' , and both P' and Q' were chosen as minimal. Hence the projection of the constructed schema onto $\{a, b'', c''\}$ implements a switching relation $R(a, b'', c'')$. \square

Theorem 3 *Let R be a regular relation which is not intersection-closed. Then $\text{WSP}(R)$ admits no kernel of size $\text{poly}(k + m)$ unless the polynomial hierarchy collapses.*

Proof By Lemma 7, it suffices to show that $\Gamma = \{R\}$ implements a switching relation R' . First consider the set $\Gamma' = \{R, =\}$. By Lemma 8, Γ' implements either a switching relation R' or the binary disequality relation \neq ; in the latter case, we have that $\Gamma'' = \{R, =, \neq\}$ implements a switching relation by Lemma 9, hence so does Γ' by the transitivity of implementations. We find

that $\text{WSP}(R, =)$ admits no polynomial kernel unless the polynomial hierarchy collapses.

Finally, there is a trivial PPT from $\text{WSP}(R, =)$ to $\text{WSP}(R)$ by merging tasks: Let $(s = s')$ be an equality constraint in an instance of $\text{WSP}(R, =)$. We may then apply Lemma 4 to s and s' , producing an equivalent instance with modified authorization lists and fewer tasks. Repeating this until no equality constraints remain in the instance yields an instance of $\text{WSP}(R)$. Hence the same lower bound applies to $\text{WSP}(R)$. \square

This finishes the proof of Theorem 1.

Corollary 2 *Let Γ be a set of regular relations. If Γ is well-behaved, then $\text{WSP}(\Gamma)$ admits a polynomial kernel in parameter $k + m$ if Γ is intersection-closed, otherwise not, unless the polynomial hierarchy collapses. If Γ is finite, then the same dichotomy holds for parameter k instead of $k + m$.*

4.4 On user bounds for WSP

In this section we return to the question of preprocessing WSP down to a number of users that is polynomial in the number k of tasks. As seen above, the positive side of our kernel dichotomy relies directly on a procedure that reduces the number of users in an instance, while the lower bounds refer entirely to the *total size* of the instance. Could there be a loophole here, allowing the number of users to be bounded without directly resulting in a polynomial kernel? Alas, it seems that while such a result cannot be excluded, it might not be very useful.

Corollary 3 *Let Γ be a set of user-independent relations containing at least one relation which is regular but not intersection-closed. Unless the polynomial hierarchy collapses, any polynomial-time procedure that reduces the number of users in an instance down to $\text{poly}(k)$ must in some cases increase either the number k of tasks, the number m of constraints, or the coding length of individual constraints superpolynomially in $k + m$.*

Proof Let $R \in \Gamma$ be a constraint that is not intersection-closed. By Theorem 3, under our assumption there is no polynomial-time procedure that reduces every instance of $\text{WSP}(R)$ down to size $p(k + m)$ for any $p(t) = t^{\mathcal{O}(1)}$. Also note that a natural encoding of WSP instances has coding length $\mathcal{O}(kn + km\ell)$, where ℓ is the largest coding length of an individual constraint. Hence, a procedure which bounds $n = k^{\mathcal{O}(1)}$ must sometimes increase one of the parameters k , m and ℓ to $(k + m)^{\omega(1)}$. \square

5 Conclusion

In this paper, we have considered kernelization properties of the workflow satisfiability problem $\text{WSP}(\Gamma)$ restricted to use only certain types $R \in \Gamma$ of

Table 1 Overview of results for typical user-independent constraints. We recall that the WSP problem is FPT with respect to k when all constraints are user-independent.

	regular	\cap -closed	poly(k) user reduction		bounded arity resp. finite Γ		well-behaved infinite Γ	
(\neq, T, T') $(\geq 2, T)$	yes	yes	yes	[9]	PK(k)	[9]	PK($k+m$)	Cor.2
$(1, t_u, T)$ (t_l, t_u, T)	yes	yes no	yes no	[9] Cor.3	PK(k) no PK($k+m$)	[9] Thm.1	PK($k+m$) no PK($k+m$)	Cor.2
$(=, s, T')$ $(=, T, T')$	yes no	no n.a.	no	Cor.3	no PK($k+m$)	Thm.2	no PK($k+m$)	Cor.2
$(\geq t, T)$ $(\leq t, T)$	no no	n.a. n.a.	yes no	[9] Cor.3	PK(k) no PK($k+m$)	[9] Cor.1	PK($k+m$) no PK($k+m$)	Prop.1 Cor.1
reg.+ \cap -cl. regular	yes yes	yes no	yes no	Thm.2 Cor.3	PK(k) no PK($k+m$)	Thm.1 Thm.1	PK($k+m$) no PK($k+m$)	Cor.2 Cor.2

constraints. We have focused on the case that all relations $R \in \Gamma$ are regular. For this case, we showed that $\text{WSP}(\Gamma)$ admits a reduction down to $n' \leq k$ users if every $R \in \Gamma$ is *intersection-closed* (and obeys some natural assumptions on efficiently computable properties), otherwise (under natural restrictions) no such reduction is possible unless the polynomial hierarchy collapses. In particular, this implies a dichotomy on the kernelizability of WSP under the parameters k for finite Γ , and $k+m$ for infinite languages Γ (subject to the aforementioned computability assumptions). This extends kernelization results of Crampton et al. [9], and represents the first kernelization lower bounds for regular constraint languages. Some results are summarized in Table 1.

An interesting open problem is to extend this result beyond regular constraints, e.g., to general user-independent constraints.

Acknowledgements Our research was partially supported by EPSRC grant EP/K005162/1. We are grateful to the referees for their useful comments and suggestions.

References

1. American National Standards Institute. ANSI INCITS 359-2004 for role based access control, 2004.
2. Elisa Bertino, Piero A. Bonatti, and Elena Ferrari. TRBAC: A temporal role-based access control model. *ACM Trans. Inf. Syst. Secur.*, 4(3):191–233, 2001.
3. Elisa Bertino, Elena Ferrari, and Vijayalakshmi Atluri. The specification and enforcement of authorization constraints in workflow management systems. *ACM Trans. Inf. Syst. Secur.*, 2(1):65–104, 1999.
4. Nadja Betzler, Robert Brederbeck, and Rolf Niedermeier. Partial kernelization for rank aggregation: Theory and experiments. In *IPEC*, volume 6478 of *LNCS*, pages 26–37. Springer, 2010.
5. D. Cohen, J. Crampton, A. Gagarin, G. Gutin, and M. Jones. Engineering algorithms for workflow satisfiability problem with user-independent constraints. In *FAW 2014*, volume 8497 of *LNCS*, pages 48–59. Springer, 2014.

6. David Cohen, Jason Crampton, Andrei Gagarin, Gregory Gutin, and Mark Jones. Iterative plan construction for the workflow satisfiability problem. *J. Artif. Intel. Res.*, 51:555–577, 2014.
7. Jason Crampton. A reference monitor for workflow systems with constrained task execution. In *SACMAT*, pages 38–47. ACM, 2005.
8. Jason Crampton, Robert Crowston, Gregory Gutin, Mark Jones, and M. S. Ramanujan. Fixed-parameter tractability of workflow satisfiability in the presence of seniority constraints. In *FAW-AAIM*, volume 7924 of *LNCS*, pages 198–209. Springer, 2013.
9. Jason Crampton, Gregory Gutin, and Anders Yeo. On the parameterized complexity and kernelization of the workflow satisfiability problem. *ACM Trans. Inf. Syst. Secur.*, 16(1):4, 2013.
10. Holger Dell and Dieter van Melkebeek. Satisfiability allows no nontrivial sparsification unless the polynomial-time hierarchy collapses. *J. ACM*, 61(4):23, 2014.
11. Michael Dom, Daniel Lokshtanov, and Saket Saurabh. Kernelization lower bounds through colors and ids. *ACM Transactions on Algorithms*, 11(2):13, 2014.
12. Rodney G. Downey and Michael R. Fellows. *Fundamentals of Parameterized Complexity*. Texts in Computer Science. Springer, 2013.
13. Danny Hermelin, Stefan Kratsch, Karolina Soltys, Magnus Wahlström, and Xi Wu. A completeness theory for polynomial (Turing) kernelization. In *IPEC*, volume 8246 of *LNCS*, pages 202–215. Springer, 2013.
14. James Joshi, Elisa Bertino, Usman Latif, and Arif Ghafoor. A generalized temporal role-based access control model. *IEEE Trans. Knowl. Data Eng.*, 17(1):4–23, 2005.
15. Daniel Lokshtanov, Neeldhara Misra, and Saket Saurabh. Kernelization - preprocessing with a guarantee. In *The Multivariate Algorithmic Revolution and Beyond*, volume 7370 of *LNCS*, pages 129–161. Springer, 2012.
16. Ravi S. Sandhu, Edward J. Coyne, Hal L. Feinstein, and Charles E. Youman. Role-based access control models. *IEEE Computer*, 29(2):38–47, 1996.
17. Qihua Wang and Ninghui Li. Satisfiability and resiliency in workflow authorization systems. *ACM Trans. Inf. Syst. Secur.*, 13(4):40, 2010.