Knowledge Compilation with Empowerment

Lucas Bordeaux[1] and Joao Marques-Silva[1,2]

[1] Microsoft Research, Cambridge, UK
[2] University College Dublin, Ireland

Abstract. When we encode constraints as Boolean formulas, a natural question is whether the encoding satisfies a "propagation completeness" property: is the basic unit propagation mechanism able to deduce all the literals that are logically entailed? We consider the problem of automatically finding encodings with this property, i.e. of compiling a naïve definition of a constraint into a good, propagation-complete encoding. Well-known Knowledge Compilation techniques from AI can be used for this purpose, but the constraints for which they can produce a polynomial size encoding are few. We show that the notion of *empowerment*, recently introduced in the SAT literature, allows producing encodings that are shorter than with previous techniques, sometimes exponentially.

1 Introduction

Modeling problems with constraints is as much an art as it is a science. Very often the same relations between variables can be encoded in a wide range of possible ways. Although semantically equivalent these various encodings may present dramatic differences in terms of complexity: some encodings are "well-posed" in some sense, which guarantees a tractable reasoning on the constraints, while some others are formulated in a way that hinders the deduction mechanisms of constraint solvers.

The Question we investigate is how to automatize the search for such "well-posed" encodings: given a naïve definition of a constraint in propositional logic, how do we produce an equivalent, but "well-posed", encoding of the constraint? Both of the words *naïve* and *well-posed* require clarification. The specific notion of "well-posed" we focus on is defined precisely later in the paper by the property of *propagation completeness*, which states that the simple unit propagation rule that is at the core of SAT solvers effectively deduces all the literals that are logically entailed. The notion of "naïve" encoding is unavoidably more subjective: by this we essentially mean a concise encoding that logically captures the semantics of the constraint, but ignores any of the redundancies or "modeling tricks" that experts will usually add for performance reasons.

Producing propagation-complete encodings is an important question on which significant prior work can be found. In particular the Knowledge Compilation literature proposes a rich set of techniques for the preprocessing of logical formulas into a "compiled" form that allows certain types of reasoning, including literal and clausal entailment, to be done in polynomial time [9]. Clausal Knowledge Compilation formalisms, initiated in [18], are particularly relevant. In parallel, recent Constraint Programming

(CP) literature proposes propagation-complete SAT encodings of specific constraints [2, 6, 17, 13, 4, 11, 4]. Bacchus [2] makes the connection between the two areas and explicitly suggests the use of Knowledge Compilation techniques to automatically obtain propagation-complete encodings of complex constraints.

Our Contribution in this paper is to relate clausal Knowledge Compilation to the recently introduced notion of *empowerment* [1, 16]. We show that generating empowering redundant clauses can be viewed as the key rewriting technique that can be used to generate "useful" redundant constraints that ultimately produce a propagation-complete formula. This observation is simple and, in retrospect, natural, but we prove that restricting the clause generation to empowering clauses can in some cases reduce exponentially the size of compiled formulas, compared to previous CNF knowledge compilation techniques.

Overview of the Paper. We start by reviewing the notation and required material in Section 2. Section 3 presents the notion of empowerment and studies the question of how to find empowering clauses. Section 4 introduces and studies the technique of compilation using empowerment, which is compared to prior CNF compilation techniques in Section 5. Finally, Section 6 outlines a number of research directions.

2 Preliminaries

This section overviews the needed preliminary material and notation.

2.1 Resolution

The paper will mostly consider logical formulas in CNF (Conjunctive Normal Form). Recall that these formulas are conjunctions of clauses, each of which is a disjunction of literals (variable or its negation). We write $\varphi \models c$ to indicate that clause c is a logical consequence of φ, i.e. all models of φ also satisfy c. We say that c is an *implicate* of φ. We denote by $\bot$ the empty clause; by $\mathit{vars}(\varphi)$ the set of variables that have an occurrence in a formula φ; by $|c|$ the length (number of literals) of a clause c.

Propositional *resolution* is the well-known deduction rule that infers from two clauses of the form $(A \vee x)$ and $(\neg x \vee B)$ the consequence $(A \vee B)$. *Unit resolution*, also known as unit propagation, is the special case where A or B is empty. We use the symbol $\vdash$ for deduction using resolution, with subscripts corresponding to restricted cases of resolution: in particular given a CNF formula φ we write $\varphi \vdash_1 c$ if the clause c can be deduced from φ using unit resolution.

2.2 Propagation-Completeness

We call *propagation-completeness* the property of CNF encodings where unit propagation computes all the valid unit consequences of the formula.

Definition 1 (Propagation-Completeness). *A CNF formula φ is propagation-complete if, for any set of literals $\{l_1, \cdots, l_k\}$, any literal d that is logically entailed can be obtained by unit propagation, i.e.,*

$$\textit{if} \quad \varphi \wedge l_1 \wedge \cdots \wedge l_k \models d \quad \textit{then} \quad \varphi \wedge l_1 \wedge \cdots \wedge l_k \vdash_1 d$$

In particular, if $\varphi \wedge l_1 \wedge \cdots \wedge l_k$ is inconsistent, and φ is propagation-complete, unit propagation deduces $\bot$. Propagation completeness has been implicitly considered in recent CP papers [12, 2, 6, 17, 13, 4]. because of its connection to *Domain Consistency* (also known as *Generalized Arc-Consistency*). In these papers, constraints are encoded in SAT using a "natural" domain encoding: each variable ranges over a finite domain and is encoded in SAT using one Boolean per value in the domain. (This stands in contrast to "logarithmic" encodings where b Booleans encode 2^b possible values). In this setting, a key observation is that if the encoding of the constraint is propagation-complete, then unit propagation on the SAT encoding effectively finds the same implications as a dedicated Domain Consistency propagator.

2.3 Knowledge Compilation.

We can now define more formally the problem we are considering throughout the paper:

Problem 1. Given a CNF formula φ, produce a "compiled formula" that is equivalent to φ and propagation-complete.

This question has been extensively studied in the *Knowledge Compilation* literature; more precisely what is studied is usually a variant called *clausal entailment* (see, e.g. [9]): given a clause $c \equiv (l_1 \vee \cdots \vee l_k)$, do we have $\varphi \models c$? It is easy to see that if a formula is propagation-complete then clausal entailment on the formula can be done in polynomial time as we can simply check whether $\varphi \wedge \neg l_1 \wedge \cdots \wedge \neg l_k \vdash_1 \bot$. Conversely, if a polynomial-time algorithm (not necessarily based on unit propagation) exists for clausal entailment then we can check efficiently whether $\varphi \wedge l_1 \wedge \cdots \wedge l_k \models d$ for any literal d by checking whether $(\neg l_1 \vee \cdots \vee \neg l_k \vee d)$ is entailed. Clausal entailment and propagation-completeness are therefore essentially equivalent, but tractable clausal entailment can resort to any type of polynomial-time algorithm, while propagation-completeness specifically focusses on unit propagation.

In general, if the formula φ to be compiled is arbitrary, the compilation process hits fundamental complexity limits: clausal entailment, like other intractable problems, is "non-compilable" in general, unless $P = NP/poly$ [7]; hence reaching propagation-completeness requires in the worst case an exponentially large encoding. There are nevertheless many specific constraints whose satisfiability and unit implication problems are polynomial-time solvable, and for many of them (but *not all*, as shown in [5]!), concise propagation-complete CNF encodings have been proposed. [6, 17, 13, 4, 11, 4]. Problem 1 asks how we automatically find such encodings.

3 Empowering Implicates

In this section we review the notion of *empowering clause* and relate it to knowledge compilation. We next address the question: how do we compute empowering clauses for a CNF formula?

3.1 Empowerment

It is well-known that adding logically redundant constraints (in SAT: implicates) to a problem can *in some cases* improve propagation [8], but that not all redundant clauses are useful. Consider for instance $\varphi = \{(\neg x \vee y), (\neg y \vee z)\}$. The clause $(\neg x \vee z)$ is an implicate, however this clause does not add anything that really benefits propagation: from x we can deduce z and from $\neg z$ we can deduce $\neg x$, with or without this clause. The property that this clause is missing is *empowerment* [16]. This notion has been introduced in a different setting, but we argue that it is the *exact* characterization of "useful" clause. Implicates such as $(\neg x \vee z)$ that are useless are said to be *absorbed* by the CNF φ.

Definition 2 (Empowerment [16]; Absorption [1]). *Let φ be a CNF formula, and $c = (l_1 \vee \cdots \vee l_k)$ be an implicate of φ. The clause c is empowering*[3] *w.r.t. φ if one of its literals l_i, called* empowered literal*, is such that:*

$$\varphi \wedge \bigwedge_{j \in 1..k,\ j \neq i} \neg l_j \quad \not\vdash_1 \quad l_i.$$

c is absorbed *by φ if it has no empowered literal.*

Our first observation is that empowerment is intimately related to compilation:

Definition 3 (Closure under Empowerment; Completion). *A CNF formula ψ is* closed under empowerment *if it absorbs any implicate not in ψ. Given a formula φ, let ψ be a set of implicates of φ; if the formula $\varphi \cup \psi$ is closed under empowerment, then we say that it is a* completion *of φ.*

Proposition 1. *A formula is closed under empowerment iff it is propagation-complete.*

Proof. If φ has en empowering clause $(l_1 \vee \cdots \vee l_k \vee d)$ with d the empowering literal, then we have $\varphi \wedge \neg l_1 \wedge \cdots \wedge \neg l_k \models d$ while $\varphi \wedge \neg l_1 \wedge \cdots \wedge \neg l_k \not\vdash_1 d$, and so φ is not propagation-complete. Conversely if φ is not propagation-complete and there is a set of literals $l_1 \cdots l_k, d$ such that $\varphi \wedge l_1 \wedge \cdots \wedge l_k \models d$ while $\varphi \wedge l_1 \wedge \cdots \wedge l_k \not\vdash_1 d$, then $(\neg l_1 \vee \cdots \vee \neg l_k \vee d)$ is empowering.

[3] Reference [16] uses the term 1-empowering, but we drop the 1- prefix for simplicitiy. Also we use in fact the original formulation from a AAAI conference paper that precedes [16]: for technical reasons [16] adds the extra condition $\varphi \wedge \bigwedge_{j \neq i} \neg l_j \not\vdash_1 \bot$, which we do not need.

Given a CNF formula φ of size s (size here means the sum of clause lengths), and a candidate clause $c = (l_1 \vee \cdots \vee l_k)$, there are two contexts in which it may be useful to check whether c is an empowering implicate. If we do not know whether c is an implicate, then checking whether it is is coNP-complete. In some other cases, c may be *known to be an implicate*, for instance if it is obtained by application of steps of propositional resolution. Checking whether it is empowering is in this case easier: we simply have to check whether any of the l_is ($i \in 1..k$) is obtained by propagation when we assert the conjunction $\bigwedge_{j \in 1..k, j \neq i} \neg l_j$. This can be done in time $\mathcal{O}(k \cdot s)$ since it is sufficient to run (and undo) k propagations.

3.2 Finding Empowering Implicates using QBF

Empowerment was introduced in the context of SAT solvers, and it was noted that learnt clauses in DPLL solvers are in fact empowering [16, 1]. However the literature has not, to our knowledge, proposed any *complete* method for finding empowering clauses. Such a method should fail to return an empowering implicate only when the formula is proved to be propagation-complete. We propose such a method based on an encoding to Quantified Boolean Formulae (QBF). The generated Quantified Boolean Formula asks whether there exists a clause that is a valid implicate and that is empowering. This QBF encoding is described below.

A set of variables X is assumed, with $|X| = n$, $X = \{x_1, \ldots, x_n\}$. Literals are represented as x_i^p, where $x_i \in X$, $p \in \{0, 1\}$, $x_i \equiv x_i^1$ and $\neg x_i \equiv x_i^0$. Essentially, p denotes the truth value of x_i that satisfies the literal associated with x_i^p. Furthermore, let $T(x_i^p)$ denote the clauses satisfied by x_i when assigned value $p \in \{0, 1\}$.

Consider a CNF formula φ and a clause c. Define $\varphi^a = \varphi \cup \{\neg l_c : l_c \in c\}$ and $\varphi^c = \varphi^a \setminus \{\neg l\}$, for a given literal $l \in c$. Recall that a clause c is empowering for φ if there exists $l \in c$ such that [4]: (i) $\varphi \models c$; (ii) $\varphi^c \not\models_1 \bot$; and (iii) $\varphi^c \not\models_1 l$.

The identification of an empowering clause c consists of the following main steps: (i) Select the literals of clause c, among all possible sets of literals; and (ii) Validate the conditions of Definition 2 for clause c. The configuration of clause c is done through a set of auxiliary variables $S = \{s_i^p \,|\, x_i \in X \wedge p \in \{0, 1\}\}$. Clause c is defined as follows: $c = \{l_1^0, l_1^1, l_2^0, l_2^1, \ldots, l_n^0, l_n^1\}$, where $l_i^p \leftrightarrow x_i^p \wedge s_i^p$. Moreover, $(\neg s_i^0 \vee \neg s_i^1)$ holds for $i = 1, \ldots, n$. A complete assignment to the variables in set S decides which literals actualy integrate clause c.

Given the definition of set S, the model outlined above can be refined as follows:

$$\exists S.(\varphi \models c) \wedge (\varphi^c \not\models_1 \bot) \wedge (\varphi^c \not\models_1 l) \tag{1}$$

which captures the conditions of Definition 2. In the remainder of this section, the complete QBF is derived by specifying the following predicates (each associated with one of the conditions of empowering clause):

$$\exists S.\text{Unsat}(\varphi^a) \wedge \text{ProperUPConfig}(\varphi^c) \wedge \bigvee_{l \in c} \text{NonUPImplied}(\varphi^c, l) \tag{2}$$

[4] Although this section assumes the most recent definition of empowering clause [16], the proposed QBF can easily be adapted to the simpler definition, i.e. without condition (ii) above.

where Unsat(φ^a) holds if φ^a is unsatisfiable, ProperUPConfig(φ^c) holds if unit propagation (UP) on φ^c is non-inconsistent, and NonUPImplied(φ^c, l) holds if l is not implied by unit propagation. Predicate Unsat(φ^a) is represented by $\forall X.\neg\varphi^a(X)$. The other two predicates require modeling the proper outcomes of unit propagation (UP) as a propositional formula.

The selection of the set of variables to be declared assigned by unit propagation is achieved through a few sets of auxiliary variables. The first set U, is defined as follows: $U = \{u_i^p \mid x_i \in X \wedge p \in \{0,1\}\}$, where u_i^p is true iff $x_i \in X$ takes value $p \in \{0,1\}$ by unit propagation. Clearly, $(\neg u_i^0 \vee \neg u_i^1)$ holds for all $1 \leq i \leq n$. Also, if $u_i^0 = u_i^1 = 0$ then x_i is unassigned by unit propagation. Moreover, $u_{i,d}^p$ denotes whether x_i takes value p in no more than d unit propagation steps. The consistent assignments to $u_{i,d}^p$ are defined recursively as follows: (i) $u_{i,0}^p = 1$ iff (x_i^p) is a unit clause; and (ii) for $d > 0$, $u_{i,d}^p = 1$ if x_i^p was already set to 1 at earlier propagation stages than d, or if there exists a clause $c_t \in \varphi^c$ with all literals but x_i^p assigned value 0 in no more than $d-1$ unit propagation steps. Formally,

$$\begin{array}{ll} u_{i,0}^p = 1 & \text{if } (x_i^p) \in \varphi^c \\ u_{i,0}^p = 0 & \text{if } (x_i^p) \notin \varphi^c \\ u_{i,d}^p = u_{i,d-1}^p \vee \bigvee\limits_{c_t \in T(x_i^p)} \bigwedge\limits_{\substack{x_j^q \in c_t \\ x_j^q \neq x_i^p}} u_{j,d-1}^{1-q} & \text{for } d > 0 \end{array} \tag{3}$$

For simplicity of notation, in the rest of the section we simply denote by u_i^p the variables $u_{i,n}^p$, that represent the state of each x_i at the end of propagation, i.e. when $d = n$ (last "propagation level").

Let $\varphi^{u,d}(u_i^p)$ denote the set of clauses from (3). Then, the predicate ProperUP(u_i^p) is given by the conjunction of $\varphi^{u,d}(u_i^p)$, $(\neg u_i^0 \vee \neg u_i^1)$, $(u_{i,d-1}^p \rightarrow u_{i,d}^p)$, with $1 \leq d \leq n$, and $(u_i^p \leftrightarrow u_{i,n}^p)$.

For a given clause c_t, the predicate SAT(c_t) is given by $(\vee_{x_i^p \in c_t} u_i^p)$. After consistent unit propagation, all clauses *must* either be satisfied or non-unit. Define v_i to be true iff x_i is unassigned, i.e. $v_i \leftrightarrow (\neg u_i^0 \wedge \neg u_i^1)$. Then the NonUnit($c_t$) predicate is defined by $(\sum_{x_i^p \in c_t} v_i \geq 2)$.

The predicate ProperUPConfig(φ^c, U) can be defined as follows:

$$\bigwedge_{\substack{1 \leq i \leq n \\ p \in \{0,1\}}} \text{ProperUP}(u_i^p) \wedge \bigwedge_{c_t \in \varphi^c} (\text{SAT}(c_t) \vee \text{NonUnit}(c_t)) \tag{4}$$

The existence of a proper UP configuration is then given by: $\exists U.$ProperUPConfig(φ^c, U). Finally, assuming literal l corresponds to x_k^p, predicate NonUPImplied(φ^c, x_k^p, U) is defined as follows: ProperUPConfig(φ^c, U) $\rightarrow \neg u_k^p$, i.e. for any non-inconsistent unit propagation x_k^p remains unassigned. Thus, the condition that x_k is not implied, is given by $\forall U.$NonUPImplied(φ^c, l, U).

Proposition 2. *Given a formula φ, the question:* is there a clause c that is an empowering implicate of φ? *is polynomial-time reducible to a QBF formula.*

Proof. See QBF derivation above.

We also note that the QBF encoding can easily be tuned to express the existence of empowering clauses *of bounded length*: this amounts to constraining the length of the desired clause in the encoding.

4 Compilation by Iterative Empowerment

The most natural approach to knowledge compilation using empowering clauses is to iteratively add empowering clauses to the formula until it ultimately becomes propagation-complete. We call this approach *compilation by iterative empowerment*. Here we study this approach and focus on a disciplined approach where empowering clauses are introduced by increasing length.

4.1 Compilation Sequences

Knowledge compilation by empowering implicate generation is highly dependent on the order in which empowering clauses are generated. The notion of *compilation sequence* captures this ordering:

Definition 4. *Let φ be a CNF formula. A* compilation sequence *is a sequence of clauses $[\varphi; c_1; \cdots; c_k]$ such that:*

- *Each clause c_i for $1 \leq i \leq k$ is an implicate of φ that is empowering wr.t. the partially compiled formula $\varphi \wedge c_1 \wedge \cdots \wedge c_{i-1}$;*
- *$\varphi \wedge c_1 \wedge \cdots \wedge c_k$ is ultimately propagation-complete.*

Example 1. Consider the formula: $\varphi \equiv \{(a \vee b), (\neg a \vee x), (\neg a \vee y), (\neg b \vee x), (\neg b \vee y), (\neg x \vee y), (\neg y \vee x)\}$. There are two possible compilation sequences for this formula: $[\varphi; (x)]$ and $[\varphi; (y)]$. In other words we may start by adding the empowering implicate (x), in which case (y) becomes deducible by unit propagation and therefore is not an empowering clause; or we may start by adding (y) in which case (x) is not empowering.

4.2 Clause Deprecation

Being empowering w.r.t. $\varphi \wedge c_1 \wedge \cdots \wedge c_{i-1}$ does not, in general, guarantee that c_i will remain empowering w.r.t. to the fully compiled formula. One issue is that as we generate empowering clauses sequentially, some clauses created at an earlier stage may become non-empowering later as more clauses are added. We call this *clause deprecation*:

Definition 5 (Clause Deprecation). *In a compilation sequence $[\varphi; c_1; \cdots; c_k]$ we say that a clause c_i is* deprecated *by the addition of clause c_j $(k \geq j > i \geq 1)$ if c_i is empowering w.r.t $\varphi \wedge \bigwedge_{h \in 1..j-1, h \neq i} c_h$ and non-empowering w.r.t. $\varphi \wedge \bigwedge_{h \in 1..j, h \neq i} c_h$.*

Example 2. Consider the formula: $\varphi \equiv \{(a \vee b), (\neg a \vee x), (\neg a \vee y), (\neg b \vee x), (\neg b \vee y), (\neg x \vee y)\}$. The two possible compilation sequences are: $[\varphi; (y); (x)]$ and $[\varphi; (x)]$. In the first sequence, (y) is empowering w.r.t. φ but is not empowering w.r.t. $\varphi \wedge (x)$, i.e. becomes deprecated by the addition of clause (x).

In extreme examples, some poorly selected sequences can ultimately contain exponentially many deprecated clauses, as shown in the following example.

Example 3. Consider the following formulas parameterized by a size m:

$$\bigwedge_p \left(y \vee \bigvee_h x_{ph} \right) \quad \wedge \quad \bigwedge_h \bigwedge_p \bigwedge_{p'>p} (y \vee \neg x_{ph} \vee \neg x_{p'h})$$

where p ranges over $1..m$ and h ranges over $1..m-1$. (These formulas are a variant of the well-known Pigeon-Hole Principle (PHP) formulas encoding $y \vee PHP_m$.) It is clear that for this formula we can generate exponentially many clauses, and in particular many empowering ones, whereas the unit clause (y) is indeed the only meaningful implicate. Specifically, a possible compilation sequence starts by φ, then adds all clauses of the form $\left(y \vee \bigvee_{p \in 1..m-1} \neg x_{p,\delta(p)}\right)$, for all bijections δ from $1 \cdots m-1$ onto itself (i.e. permutations); then adds (y). All clauses are empowering at the time where they are added. Yet adding the final clause (y) deprecates all the previous clauses.

Example 3 shows that generating a short (here, unit) empowering clause can sometimes prevent the creation of many more empowering clauses. This suggests a strategy of *length-increasing* iterative empowerment, detailed next.

4.3 Length-Increasing Iterative Empowerment

We now consider what happens if we generate clauses *by increasing length*: we first saturate the formula under empowering clauses of length 1; only then do we consider length 2; and so forth. This is shown as Algorithm *length_increasing_empower* in Fig. 1. The algorithm is similar to the simple width-increasing algorithm resolution proposed in e.g. [3], but (1) only generates clauses that are empowering, and (2) exhaustively checks that no implicate of length L exists before incrementing L. The latter test can be done by a width-bounded QBF encoding as shown in the previous section. This algorithm is non-deterministic in that there are many possible choices in the selection of the empowering clause at any stage. The sequences of implicates it generates can be characterized as follows:

Definition 6 (Length-Increasing Compilation Sequence). *A compilation sequence* $[\varphi; c_1; \cdots; c_k]$ *is* length-increasing *if for every* $i \in 1..k$ *we have that all implicates of length strictly less than* $|c_i|$ *are absorbed by* $\varphi \wedge c_1 \wedge \cdots \wedge c_{i-1}$.

A key property of compilation by length-increasing iterative empowerment is that it limits the effects of deprecation, in the following sense:

Proposition 3. *In a length-increasing sequence, a clause of length* b *never deprecates a clause of length* $a < b$.

Proof. Consider a sequence $[\varphi; c_1; \cdots; c_k]$, and clauses c_i and c_j with $1 \leq i < j \leq k$. Denote by a the length of c_i and by b the length of c_j, with $a < b$ We assume

function *length_increasing_empower*(φ):
    let $\psi := \emptyset$
    for L **from** 1 **to** *width*(φ)
        while there exists a clause c of length L that is empowering w.r.t. $\varphi \cup \psi$
            $\psi := \psi \cup \{c\}$
        % minimization code optionally goes here
    return $\varphi \cup \psi$

function *minimize*(φ):
    let $\psi := \varphi$
    foreach clause c **in** ψ *% arbitrary order*
        if c is absorbed by $\psi \setminus \{c\}$
            $\psi := \psi \setminus \{c\}$
    return ψ

Fig. 1. Algorithms: *length_increasing_empower* takes a CNF φ and returns a completion of it; *minimize* takes a propagation-complete CNF φ and returns a subset of it that is equivalent, propagation-complete, and minimal.

that c_j deprecates c_i and show a contradiction. We denote by Φ the formula: $\varphi \wedge \bigwedge_{h \in 1..j-1,\ h \neq i} c_h$. Furthermore let

$$c_i \text{ be of the form: } (l_1 \vee \cdots \vee l_a), \text{ and}$$
$$c_j \text{ be of the form: } (\lambda_1 \vee \cdots \vee \lambda_b).$$

At step i in the sequence, clause c_i is empowering w.r.t. at least one of its literals, say l_a. By definition of deprecation the clause remains empowering until step $j-1$ and becomes non-empowering at step j, i.e. we have:

$$\Phi \wedge \neg l_1 \wedge \cdots \wedge \neg l_{a-1} \not\vdash_1 l_a$$

$$\Phi \wedge c_j \wedge \neg l_1 \wedge \cdots \wedge \neg l_{a-1} \vdash_1 l_a$$

This means that clause c_j becomes unit when added to $\Phi \wedge \neg l_1 \wedge \cdots \wedge \neg l_{a-1}$ (if it did not the propagation would be unchanged and we still would not get l_a). One of its literals is implied, say λ_b. But then the clause $(l_1 \vee \cdots \vee l_{a-1} \vee \lambda_b)$ is an empowering implicate of length a, and the sequence is not length-increasing.

4.4 Minimality

It may be desirable in some cases to compute propagation-complete formulas that are *minimal*, where removing any clause would cause the formula not to be propagation-complete anymore.

Definition 7 (Minimal propagation-complete formula). *A propagation-complete CNF formula* $\{c_1 \cdots c_m\}$ *is* minimal *if no* c_i *is absorbed by the set of remaining clauses, i.e.* $\{c_j \ : \ j \neq i\}$.

Example 2 shows that length-increasing iterative empowerment does not necessarily lead to formulas that are minimal: deprecation can happen *between generated implicates of the same length*. Minimizing a propagation-complete formula can be done by simply checking one by one, in some arbitrary order, whether any clause is absorbed by the rest of the formula, as shown in Algorithm *minimize* of Fig. 1. If a clause c is verified to be empowering during the execution of the algorithm, it is clear that it will remain empowering at the end of the execution, where more clauses have been removed; therefore considering each clause once is enough. Minimizing a formula of length s (counted in sum of clause lengths) takes time $\mathcal{O}(s^2)$ since we need to do/undo one propagation for each literal of every clause.

To construct a minimal propagation-complete formula using the length-increasing compilation approach, it is possible to interleave the generation of empowering clauses of every length with minimization steps. In Algorithm *length_increasing_empower* of Fig. 1 the minimization code can be added in the commented area. Because of Proposition 3, it is sufficient, once the generation of empowering implicates of a certain length L has completed, to verify the empowerment of clauses of length L, i.e. the **foreach** loop of Algorithm *minimize* can be restricted to clauses whose length is L.

5 Iterative Empowerment versus Prime Implicate Saturation

We now study compilation by iterative empowerment and compare it against prior implicate-based approaches to knowledge compilation. For most of our results, we focus on the length-increasing compilation scheme, but do not assume minimality. Our main point is that it provides a strictly more "succinct" compilation language than prior compilation methods by prime implicates, where succinctness is defined as in [9].

5.1 Previous Compilation Schemes.

A well-known way to obtain a propagation-complete formula is to saturate it by prime implicates, as was proposed by, e.g. [18, 15] and suggested for the encoding of constraints by [2]. An implicate of a formula φ is a clause c that is a valid consequence, i.e. $\varphi \models c$. An implicate is *prime* if it is not subsumed by another implicate, i.e., there is no implicate c' that contains a strict subset of the literals of c. We denote by $\mathit{prime}(\varphi)$ the set of prime implicates of a formula φ. (This set is uniquely defined.)

It is known that prime implicate generation can yield clauses that are useless for propagation-completeness; for instance from the formula $\varphi = (\neg x \vee y), (\neg y \vee z)$ the absorbed clause $(\neg x \vee z)$ is a prime implicate. Heuristics have been proposed to avoid generating such absorbed clauses. In particular, when resolving two clauses $(A \vee x)$ and $(\neg x \vee B)$, it is easy to see that if $\mathit{vars}(A) \cap \mathit{vars}(B) = \emptyset$, then the generated clause $A \vee B$ is absorbed by $(A \vee x), (\neg x \vee B)$. The resolution step is in this case called *non-merge* (it is called *merge* if $\mathit{vars}(A) \cap \mathit{vars}(B) \neq \emptyset$). [18] proposes several methods that essentially remove from a formula saturated under prime implicates clauses that can be obtained using non-merge resolution. This technique is a particular case of the idea that we advocate of removing absorbed clauses, but this particular case is limited: we

prove in Proposition 7 that in some cases the simplifications obtained using non-merge resolution are exponentially weaker than those obtained with iterative empowerment.

We compare iterative empowerment to compilation using prime implicates and its improvement based on merge resolution. Later approaches to prime implicate generation have been proposed, for instance [14], but these approaches depart from explicit CNF generation, while it is our aim to produce CNF encodings amenable to SAT solving. ([14] uses, specifically, a compact clause representation with BDDs).

5.2 Iterative Empowerment versus Prime Implicates.

We first note that length-increasing compilation sequences can never generate more clauses than prime implicate generation, as they only contain prime implicates. We then show that compilation by saturation under empowering clauses can in some cases be exponentially more compact than approaches based on prime implicate generation.

Proposition 4. *All empowering implicates generated in a length-increasing compilation sequence are prime.*

Proof. Consider a sequence $[\varphi; c_1; \cdots; c_k]$, and clauses c_i and c_j. We assume that c_i subsumes c_j and show a contradition. Since the literals of c_i are a strict subset of those of c_j, we have $|c_i| < |c_j|$ and $i < j$, i.e. clause c_j is generated after c_i in the sequence. But c_j is not empowering w.r.t. the formula that already includes c_i: whenever c_j becomes unit, c_i either also becomes unit, or becomes inconsistent; in both cases no new useful new literal can be inferred from c_j.

A class of formulas that exhibit an exponential separation between prime implicate saturation and iterative empowerment is the so-called EVEN formulas, introduced next. (Comments on the right-hand side of the formula explain the meaning of each block of clauses in this formula.) These are CNF encodings of the formula $x_1 \oplus \cdots \oplus x_n = 0$, true when the number of 1s is even.

Definition 8 (EVEN Formulas). *We denote by* EVEN_n *the following formula over the sets of variables* $X = \{x_1 \cdots x_n\}$ *and* $Y = \{y_1 \cdots y_n\}$*:*

$$
\begin{array}{ll}
(\neg x_1 \vee y_1) \wedge (\neg y_1 \vee x_1) & \text{"}y_1 = x_1\text{"} \\
\wedge \bigwedge_{i \in 2..n} \left(\begin{array}{l} (\neg y_i \vee y_{i-1} \vee x_i) \wedge \\ (\neg y_i \vee \neg y_{i-1} \vee \neg x_i) \wedge \\ (y_i \vee \neg y_{i-1} \vee x_i) \wedge \\ (y_i \vee y_{i-1} \vee \neg x_i) \end{array} \right) & \text{"}y_i = y_{i-1} \oplus x_i\text{"} \\
\wedge (y_n) & \text{"}\textit{output gate } y_n \textit{ is true}\text{"}
\end{array}
$$

Proposition 5. *The* EVEN_n *formulas are closed under empowerment yet have exponentially many prime implicates.*

Proof. We first note that the constraint hyper-graph of these formulas is Berge-acyclic. It is well-known that for acyclic constraint networks achieving arc consistency for each

constraint of the hyper-graph is enough to achieve arc-consistency for the whole network [10]. It follows that the formula is propagation-complete; In other words any implicate is absorbed. Now there exist an exponential number of prime implicates: (1) any clause over the variables $\{x_1 \cdots x_n\}$ that has an odd number of negative literals is an implicate. (2) these implicates are prime: if we remove any of their literals we obtain an invalid clause. (3) there are 2^{n-1} such clauses.

5.3 Iterative Empowerment versus Merge Resolution.

We next consider the optimizations to prime implicate generation from [18], in which we discard clauses that are produced by *non-merge* resolution steps. This method allows to remove some absorbed clauses, however we show that it can still generate exponentially many non-empowering clauses. This is exhibited by the following class of formulas.

Definition 9 (MERGE-EVEN formulas). *We denote by* MERGE-EVEN$_n$ *the variant of* EVEN *in which every clause receives an additional positive literal f, such that $f \notin X \cup Y$:*

$$\begin{array}{l} (f \vee \neg x_1 \vee y_1) \wedge (f \vee \neg y_1 \vee x_1) \\ \wedge \bigwedge_{i \in 2..n} \left(\begin{array}{l} (f \vee \neg y_i \vee y_{i-1} \vee x_i) \wedge \\ (f \vee \neg y_i \vee \neg y_{i-1} \vee \neg x_i) \wedge \\ (f \vee y_i \vee \neg y_{i-1} \vee x_i) \wedge \\ (f \vee y_i \vee y_{i-1} \vee \neg x_i) \end{array} \right) \\ \wedge (f \vee y_n) \end{array}$$

The set of solutions of MERGE-EVEN$_n$ is exactly the union of: (1) all assignments where f is true (with any value assigned elsewhere); (2) all assignments where f is false and EVEN$_n$ holds.

Proposition 6. *The set of prime implicates of* MERGE-EVEN$_n$ *is:* $\{(f \vee A) \ : \ A \in$ *prime(*EVEN$_n$*)*$\}$

Proof. Given any clause $A \in$ *prime(*EVEN$_n$*)*, the clause $(\neg f \vee A)$ is an implicate of MERGE-EVEN. To show that every clause $(f \vee A)$ is prime, we have two cases:

- If we remove literal f from the clause then A is not an implicate because there are solutions that assign f to true and contradict A.
- Any clause of the form $(f \vee B)$, where B is a strict subset of A, is not an implicate because B would be an implicate of EVEN and A would not be a prime implicate of EVEN.

Conversely given any prime implicate of MERGE-EVEN note the following. First the implicate contains literal f since the assignments where f is true can assign arbitrary values to the variables in X, Y. Now given that the implicate is of the form $(f \vee A)$, clause A is a prime implicate of MERGE-EVEN since the values assigned to X, Y when f is false are exactly the solutions to MERGE-EVEN.

Proposition 7. *The formula* MERGE-EVEN$_n$ *is closed under empowerment, but has exponentially many prime implicates; furthermore all of these implicates are produced by merge-resolution.*

Proof. Assume the existence of a non-empowering prime implicate of MERGE-EVEN. It is a clause of the form $(f \vee A \vee l)$, not subsumed by any clause of MERGE-EVEN, and where literal l is not obtained from MERGE-EVEN by unit propagation when f is false and all literals in A are false. When f is set to false the formula simplifies to EVEN, therefore it is also the case that l is not obtained from EVEN when all literals in A are false. This implies that $(A \vee l)$ is an empowering clause of EVEN not subsumed by any clause in it, which contradicts Proposition 5. We now show that all the resolvents applied in the compilation are Merge. All clauses of EVEN include a positive occurrence of f. Any resolution on two such clauses is a merge resolution operation and produces in a clause that also includes a positive occurrence of f. Therefore all clauses in any resolution proof will include only clauses with positive occurrence of f, and make use of merge resolution exclusively.

Note also that in formula MERGE-EVEN the merge literal is always f. This means that the optimization of algorithm *FPI*$_1$ of [18] does not apply and that this algorithm also generates exponentially many absorbed implicates.

6 Conclusion and Perspectives

Clausal (CNF) formalisms played an important role in early Knowledge Compilation work, but more recent work has favoured non-clausal formalisms that are in many respects more expressive [9]. In our view clausal Knowledge Compilation remains important because of its connections to propagation-complete encodings, and to our initial Question (formalized as Problem 1). Our main findings in this paper are the following:

- The notion of *empowerment* of [1, 16] sheds a new light on clausal Knowledge Compilation: several heuristics had previously been proposed to limit the "useless" redundant constraints generated by classical prime implicate methods; but they did not, to our knowledge, explicitly define "useless". Our belief is that empowerment characterizes "the desired property" of the implicates used in compilation.
- In particular we showed that length-increasing empowerment is a knowledge compilation scheme that has many appealing features in terms of limited deprecation and easier minimization, and comparison with other prime implicate generation schemes.

References

1. Atserias, A., Fichte, J.K., Thurley, M.: Clause-learning algorithms with many restarts and bounded-width resolution. J. of Artif. Intel. Research (JAIR) 40, 353–373 (2011), preliminary version in SAT 2009
2. Bacchus, F.: GAC via unit propagation. In: Princ. and Practice of Constraint Programming (CP). pp. 133–147 (2007)

3. Ben-Sasson, E., Wigderson, A.: Short proofs are narrow - resolution made simple. J. of the ACM 48(2), 149–169 (2001)
4. Bessiere, C., Katsirelos, G., Narodytska, N., Quimper, C.G., Walsh, T.: Decompositions of all different, global cardinality and related constraints. In: Int. Joint. Conf. on A.I. (IJCAI). pp. 419–424 (2009)
5. Bessiere, C., Katsirelos, G., Narodytska, N., Walsh, T.: Circuit complexity and decompositions of global constraints. In: Int. Joint. Conf. on A.I. (IJCAI). pp. 412–418 (2009)
6. Brand, S., Narodytska, N., Quimper, C.G., Stuckey, P.J., Walsh, T.: Encodings of the sequence constraint. In: Princ. and Practice of Constraint Programming (CP) (2007)
7. Cadoli, M., Donini, F., Liberatore, P., Schaerf, M.: Preprocessing of intractable problems. Information and Computation 176, 89–120 (2002)
8. Choi, C.W., Lee, J.H.M., Stuckey, P.J.: Removing propagation redundant constraints in redundant modeling. ACM Trans. on Computational Logic (TOCL) 8(4) (2007)
9. Darwiche, D., Marquis, P.: A knowledge compilation map. J. of Artif. Intel. Research (JAIR) 17, 229–264 (2002)
10. Dechter, R.: Tractable structures for constraint satisfaction problems. In: Handbook of Constraint Programming, chap. 7, pp. 209–244. Elsevier (2006)
11. Feydy, T., Stuckey, P.J.: Lazy clause generation reengineered. In: Princ. and Practice of Constraint Programming (CP). pp. 352–366 (2009)
12. Gent, I.P.: Arc consistency in SAT. In: Euro. Conf. on A.I. (ECAI). pp. 121–125 (2002)
13. Huang, J.: Universal Booleanization of constraint models. In: Princ. and Practice of Constraint Programming (CP). pp. 144–158 (2008)
14. Marquis, P., Sadaoui, S.: A new algorithm for computing theory prime implicates compilations. In: Conf. on Artif. Intel. (AAAI). pp. 504–509 (1996)
15. Marquis, P.: Knowledge compilation using theory prime implicates. In: Int. Joint. Conf. on A.I. (IJCAI). pp. 837–845 (1995)
16. Pipatsrisawat, K., Darwiche, A.: On the power of clause-learning SAT solvers as resolution engines. Artif. Intel. 175(2), 512–525 (2011), preliminary works in AAAI 2008 (notion of *empowerment*) and CP 2009.
17. Quimper, C.G., Walsh, T.: Decompositions of grammar constraints. In: Conf. on Artif. Intel. (AAAI). pp. 1567–1570 (2008)
18. del Val, A.: Tractable databases: How to make propositional unit resolution complete through compilation. In: Knowledge Representation and Reasoning (KR). pp. 551–561 (1994)