

Multi-Dimensional Models Facilitate Automatic Reformulation: The Case Study of the SET Game

Amanda Swearngin¹, Berthe Y. Choueiry², Eugene C. Freuder³

¹ ESQuaReD Laboratory University of Nebraska-Lincoln
aswearng@cse.unl.edu

² Constraint Systems Laboratory, University of Nebraska-Lincoln
choueiry@cse.unl.edu

³ Cork Constraint Computation Centre, University College Cork, Ireland
e.freuder@4c.ucc.ie

Abstract. In this paper we describe a reformulation strategy for solving multi-dimensional Constraint Satisfaction Problems (CSPs). This strategy operates by iteratively considering, in isolation, each one of the uni-dimensional constraints in the problem, and exploits the approximate symmetries induced by the selected constraint on the domains in order to enforce this constraint on the simplified problem. We use the game of SET, a combinatorial card game, as a toy problem to motivate our strategy and to explain and illustrate its operation. However, we believe that our approach is applicable to more complex domains of scientific and industrial importance, and deserves more thorough investigations in the future. Our approach sheds a new light on the dynamic reformulation of multi-dimensional CSPs. Importantly, it advocates that modeling tools for Constraint Programming should allow the user to specify the constraints directly on the attributes of the domain objects (i.e., variables and values) so that their multi-dimensionality can be exploited during problem solving.

1 Introduction

Yoshikawa & Wada [1992] introduced multi-dimensional Constraint Satisfaction Problems (CSPs) and used them to model many applications of practical importance such as resource allocation and configuration. In a multi-dimensional CSP, the domains of all variables are identical and the domain values are specified according to a set of domain dimensions (i.e., attributes). Constraints that apply to a single domain dimension are said to be uni-dimensional, otherwise they are multi-dimensional. In this paper, we propose a general reformulation strategy for solving multi-dimensional CSPs that reduces the cost of problem solving by facilitating the discovery of approximate symmetries. Our strategy, shown in Figure 1, operates on a multi-dimensional CSP by iteratively enforcing each uni-dimensional constraint on the corresponding domain dimension while ignoring all other constraints and domain dimensions. Ignoring all but one constraint allows one to identify, in the relaxed problem, symmetries that do not hold in the original problem [Freuder & Sabin, 1995; 1997]. Such approximate symmetries can be exploited to reduce the computational cost of enforcing the constraint on the relaxed problem. Each step in Figure 1 may discover unsolvability or produce one or more simplified subproblems where the considered constraint holds. At the end of the process,

any constraint solver can be used to solve the resulting problem(s) by enforcing the remaining constraints. In [1995; 1997], Freuder & Sabin propose a similar approach that exploits of symmetry known as neighborhood value interchangeability [Freuder, 1991]. Their strategy differs from ours in that it has a single abstraction step (ref. ‘2. *Reduce*’ in [Freuder & Sabin, 1997]). Indeed, the first simplified CSP is solved and its solution is used to solve the original CSP. In contrast, in our approach, we foresee a *sequence of reformulation steps*, each enforcing a single uni-dimensional constraint.

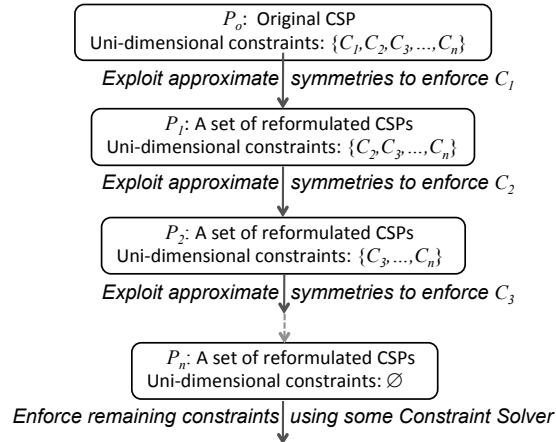


Fig. 1. A general reformulation strategy for multi-dimensional CSPs.

In this paper, we introduce a reformulation algorithm that ‘instantiates’ the general strategy of Figure 1 to solve the game of SET,⁴ a combinatorial card game. This game was invented in 1974 by Marsha Jean Falco,⁵ a population geneticist. She was inspired to create the game by her work on determining whether epilepsy in German shepherd dogs is inherited [Davis & McLagan, 2003]. We propose for SET a multi-dimensional constraint model with four uni-dimensional constraints. We describe a basic constraint solver for solving an instance of the game. Our solver is a simple backtrack search procedure with symmetry breaking that finds all the solutions of the instance. Our reformulation algorithm for SET features the following components: (1) Heuristics for selecting the uni-dimensional constraint to consider at each step; (2) The use of meta-interchangeability [Freuder, 1991] as an approximate symmetry; and (3) A disjunctive decomposition of an intermediate CSP into subproblems with non-overlapping solutions as a result of enforcing the selected uni-dimensional constraint. We show that our reformulation significantly reduces the problem-solving effort.

We have implemented our approach in an online interactive system for a single player and for two players. Naturally, the value of the system is *not* in solving the game, which, given its size, can be played by humans without the help of a computer and is widely enjoyed by children and mathematicians alike. However, we believe that our system, when completed,⁶ will be useful to teach and demonstrate problem-solving

⁴ [http://en.wikipedia.org/wiki/Set_\(game\)](http://en.wikipedia.org/wiki/Set_(game))

⁵ <http://www.setgame.com/set/history.htm>

⁶ A tool for displaying and explaining the reformulation steps has yet to be developed.

strategies to students in Computer Science and to the general public. Beyond SET, our approach sheds a new light on the dynamic reformulation of CSPs, leading the way to new strategies for effective problem solving. We believe that our approach is applicable to more complex domains of practical industrial importance beyond the toy problem considered in this paper, which calls for future investigations. In particular, we advocate that modeling tools for Constraint Programming should allow the user to specify the constraints directly on the attributes of the domain objects (i.e., variables and values) so that their multi-dimensionality can be exploited during problem solving.

This paper is structured as follows. Section 2 gives background information about the game and multi-dimensional CSPs. Section 3 describes our model and the search procedure for solving it. Section 4 introduces our reformulation of the model. Section 5 discusses our reformulation algorithm for SET. Section 6 discusses our results. Section 7 presents our interactive interface, available online on <http://gameofset.unl.edu>. Section 8 relates our work to previous research. Finally, Section 9 concludes this paper drawing directions for future research.

2 Background

We provide background information about the game and the modeling techniques.

2.1 The Game of SET

SET is a combinatorial card game consisting of a deck of 81 playing cards. Each card is uniquely determined by the values of four *attributes*, namely, the *number* of objects drawn on the card and their *color*, *filling*, and *shape*. We denote these attributes by N , C , F , and S , respectively. Each attribute takes one of three possible values as follows: $\{1,2,3\}$ for the dimension *number*, $\{red,green,purple\}$ for *color*, $\{striped,full,empty\}$ for *filling*, and $\{squiggle,oval,diamond\}$ for *shape*, see as shown in Figure 2.







Number	1	2	3
Color	red	green	purple
Filling			
Shape			

Fig. 2. The four attributes and their values.




Fig. 3. A solution set.

To play the game, twelve cards are dealt and placed, face up, on the table, visible to all players. The players compete to find a collection of exactly three cards that constitute what we call a *solution set*. For each of the above listed attributes, the three cards of a solution set must all have either the same value or different values for the attribute. Figure 3 shows three cards forming a solution set. In this example, the three cards differ on all four attributes. The first player to identify a solution set picks it up, and the cards are replaced with new ones from the deck. If all players agree that a solution set cannot be found among the twelve cards, three more are dealt and the game resumes. The operation is repeated until a set is found. Usually, the number of cards on the table quickly returns to twelve. The maximum number of cards on the table at any one time is 21, because any combination of 21 cards is guaranteed to have at least a solution set.

The proof was performed by exhaustive computation [Davis & McLagan, 2003]. The game continues until all cards have been picked up or there are no more sets among the remaining cards. The winner is the player with the largest number of sets at the end of the game. For the sake of space, our examples will show game instances with only nine cards although the original game considers twelve cards.

Obviously, a simple nested *for*-loop can ‘easily’ generate all combinations of three cards. Each combination can then be tested to check whether or not it satisfies the constraints. Such an approach is shortsighted. Indeed, human beings are unlikely to play the game by examining $\binom{12}{3} = 220$ combinations. Instead, it is fair to assume that they use various modeling and reformulation strategies. It would be equally ridiculous to use the simple nested *for*-loop to solve industrial-size problems because the number of combinations grows exponentially with the size of the problem and few of the generated combinations satisfy the constraints. For example, $\binom{12}{3} = 220$ combinations have on average 2.77 solutions and $\binom{81}{3} = 85320$ combinations have only 1080 solutions. Thus, the simple nested *for*-loop is a strategy that is not interesting to teach a human player or to use for automation in an industrial setting.

2.2 Constraint Satisfaction Problems (CSPs)

A Constraint Satisfaction Problem (CSP) is defined by $\mathcal{P} = (\mathcal{V}, \mathcal{D}, \mathcal{C})$ where \mathcal{V} is a set of variables, \mathcal{D} a set of domains, and \mathcal{C} a set of constraints. Each variable $V_i \in \mathcal{V}$ has a finite domain $D_i \in \mathcal{D}$. Each constraint in \mathcal{C} applies to a subset of the variables (the scope of the constraint), and restricts the combination of values that those variables can take at the same time. A solution to a CSP is an assignment of a value to each variable such that all the constraints are satisfied. Solving a CSP requires finding one or all solutions. In [1992], Yoshikawa & Wada defined of a *multi-dimensional* CSP as a CSP where (1) All variables have the same domains,⁷ and (2) The domain can be specified by a multi-dimensional array of values, where each value is described by a combination of domain dimensions. They called *one-dimensional* constraint a constraint defined over a single domain dimension, otherwise the constraint is said to be *multi-dimensional*.

3 Solving SET as a CSP

Now, we describe our constraint model and a search procedure for finding all solutions.

3.1 A Constraint Model for SET

Our constraint model has (only!) three variables corresponding to the three cards of a solution set: $\mathcal{V} = \{V_1, V_2, V_3\}$. All three variables have the same domain, which are the cards c_i placed on the table: $D_1 = D_2 = D_3 = \{c_1, c_2, \dots, c_i\}$, where $i \in [3, 21]$. We model the domain of a SET variable as a multi-dimensional array indexed by the following attributes of the playing cards: *number*, *color*, *filling*, and *shape* (see Figure 3). We also include the unique identifier *id* of a playing card as a fifth pseudo-attribute in order to uniquely identify each card. The two-dimensional array shown in Figure 4, which we call the *domain table*, is a flattened representation of a multi-dimensional domain.

Attribute/val	c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9
Number	1	2	3	1	2	3	1	2	3
Color	r	g	p	r	g	p	r	g	p
Filling	f	e	s	f	e	s	f	e	s
Shape	o	d	o	o	d	o	o	d	o

Fig. 4. The domain table of a nine-cards example.




Fig. 5. The constraint network of the CSP of a SET instance.

To represent the game of SET, we use a set of five ternary uni-dimensional constraints on $\{V_1, V_2, V_3\}$. The constraints, for an attribute A where $A \in \{N, C, F, S\}$ and the card id , are defined according to the following templates: $A^=$, All three variables must have the same value for the attribute A ; and A^\neq , All three variables must have all different values for the attribute A . The constraints are:

1. $N^= \oplus N^\neq$: All three cards have the same number or three different numbers.
2. $C^= \oplus C^\neq$: All three cards have the same color or three different colors.
3. $F^= \oplus F^\neq$: All three cards have the same filling or three different fillings.
4. $S^= \oplus S^\neq$: All three cards have the same shape or three different shapes.
5. id^\neq : A solution set consists of three (distinct) cards.

The constraints set is $C_o = \{(N^\neq \oplus N^=), (C^\neq \oplus C^=), (F^\neq \oplus F^=), (S^\neq \oplus S^=), id^\neq\}$. Figure 5 shows the constraint network of our model. Constraints can be specified either in extension or in intension. Given the number of variables of the constraint model (i.e., three), implementing the constraint in intension seems to be the simplest choice. Indeed, if the constraints were to be implemented in extension, the constraints can be built once and for all (domains size equal to 81) or dynamically every time a CSP is formed to be solved (domains size in [12,21]). The number of allowed tuples to generate and store for a domain of 81 cards is $P(81, 3) = \frac{81!}{(81-3)!} = 511920$. Otherwise, it is between: $P(12, 3) = \frac{12!}{(12-3)!} = 1320$ and $P(21, 3) = \frac{21!}{(21-3)!} = 7980$. Although the remaining constraints are significantly smaller (less than 3^3 tuples), it is obviously more cost effective to implement all constraints in intension than in extension.

3.2 Solving the Constraint Model

We implemented a simple backtrack search to find all the solution sets of any number of playing cards. In order to avoid generating solutions that can be obtained from other solutions by simple permutation of the cards over the three variables, we have added to the model a symmetry-breaking constraint based on lexicographical ordering of the cards unique identifiers. Given that the constraint on an attribute A has two mutually exclusive components (i.e., $A^=$ and A^\neq), our backtrack search implements a convenient combination of forward checking and back-checking [Prosser, 1993]. The symmetry

⁷ We will relax this condition, sometimes, during reformulation.

breaking constraint, the id^\neq constraint, and all equality constraints (A^\neq) are enforced by forward checking. The four all-different constraints (A^\neq) other than id^\neq are enforced by back-checking. The depth of the search tree is at most three. When the first variable is instantiated, the only checkable constraints are id^\neq and the symmetry breaking one. They are enforced by forward checking. After instantiating the second variable, we can determine by back-checking, for each of the remaining four attributes, which of the two constraints (equality A^\neq or inequality A^\neq) holds between the first two variables. The constraint that applies is selected and the other one is “switched off.” At this point, if any equality constraint on the current variable is “switched on,” it is enforced by forward checking. The domains of the third variable is consistent with all applicable equality constraints. If any all-different constraint A^\neq is applicable (i.e., switched on), then back-checking is applied to consider only consistent instantiations. In general, and depending on the constraints, one may be able to enforce higher consistency levels.

4 Constraint Model Reformulation

In this section, we reformulate our constraint model for SET by exploiting the multi-dimensionality of the domains. First, we describe how we partition the variables’ domains based on the values of a given domain dimension. We show that this process yields, in the case of the game of SET, a disjunctive decomposition of the CSP producing a set of CSPs at each reformulation step, which we illustrate with two examples.

4.1 Model Reformulation

Each reformulation step in Figure 1 partitions the variables’ domains according to a domain dimension and enforces the constraint relative to that dimension. We illustrate this process on an example. Consider the six-card game of Figure 6. All cards are red and have an empty filling. The only constraints left to enforce are $\{N^\neq \oplus N^\neq, S^\neq \oplus S^\neq, id^\neq\}$. Considering first the dimension ‘number,’ we partition the variables’ domains into equivalences classes based on the values of the chosen dimension, *number* (see Figure 7). This operation corresponds to domain partitioning by value meta-




Fig. 6. Simple example.




Fig. 7. CSP model.




Fig. 8. The domains of the variables are identical.




Fig. 9. Domains are all different.

interchangeability [Freuder, 1991]. All subproblems can be generated by a Cartesian product of the domain partitions. However, enforcing the constraint $N^= \oplus N^\neq$ and ignoring symmetrical subproblems that can be obtained by domains permutation yield a decomposition of the CSP into only four CSPs: one for each domain partition enforcing $N^=$ (Figure 8) and one where the domains have all-different partitions enforcing N^\neq (Figure 9). The solutions of the CSP are also partitioned over the four generated problems. This process can be repeated for the remaining domain dimension.

4.2 Reformulation strategy for SET

The strategy of Figure 1 exploits the one-dimensionality of the constraints of the CSP model where each one-dimensional constraint is applied to the CSP in sequence, one at a time. For SET, we showed how exploiting value meta-interchangeability for a given dimension and enforcing the constraint relative to that dimension decomposes the CSP into four subproblems whose solution sets do not overlap. Combining the reformulation strategy of Figure 1 and domain partitioning yields the reformulation tree of Figure 10.

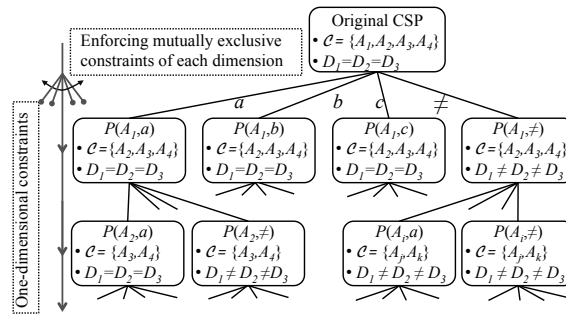


Fig. 10. One-dimensional constraints are applied in sequence, mutually exclusive in parallel.

In this figure, A_i denotes a domain dimension, a, b, c domain partitions, and $P(A_i, a)$ the subproblem resulting from enforcing the constraint relative to A_i and a . Two important questions arise in general:

1. *Which dimension (i.e., attribute) to choose at each step?* Naturally, one should reduce the branching factor by choosing, for example, the dimension that yields the smallest problem, the most symmetries, the largest domain partitions, etc.
2. *Which subproblems are generated at each branching step?* The decomposition of Section 4.1 depends on the interchangeability and constraint types that hold for the considered dimension. Generalizing this decomposition for all types of symmetries and constraints requires further investigation.

In Section 5.2, we answer the above two questions for the game of SET. Our approach is based on the analysis of the domain table shown in Figure 4. Below, we motivate that approach with two examples, see Sections 4.3 and 4.4. All generated subproblems have the same set of variables. They differ in the constraints and the variables' domains. To generate the children of a given problem in the tree of Figure 10, we need to specify the set of constraints and the domain set of each child. The constraints set of a child subproblem is that of the parent minus the constraints of the attribute used in the branching.

The domains of a child is smaller than those of its parent, which is the main incentive for the decomposition. As for the set of domains of a child, we distinguish the case where the enforced constraint at the branching step is an equality or an inequality constraint. The former keeps all domains equal whereas the latter yields a new problem where variables have different domains. Below, we illustrate the above on two examples.

4.3 Branching on equality constraints

Figure 11 shows the example of Figure 4 with a compacted the domain table. Focusing on the dimension *filling*, we notice that *filling* takes only two values *e* and *f*.

			$D_1 = D_2 = D_3$	c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9
c_4	c_5	c_6	Number	3	1	1	2	3	3	3	3	1
c_7	c_8	c_9	Color	r	r	g	p	g	r	p	r	r
			Filling	f	e	e	f	f	f	f	e	e
			Shape	o	d	s	d	d	s	d	o	o

Fig. 11. The compacted domain table of the example of Figure 4.

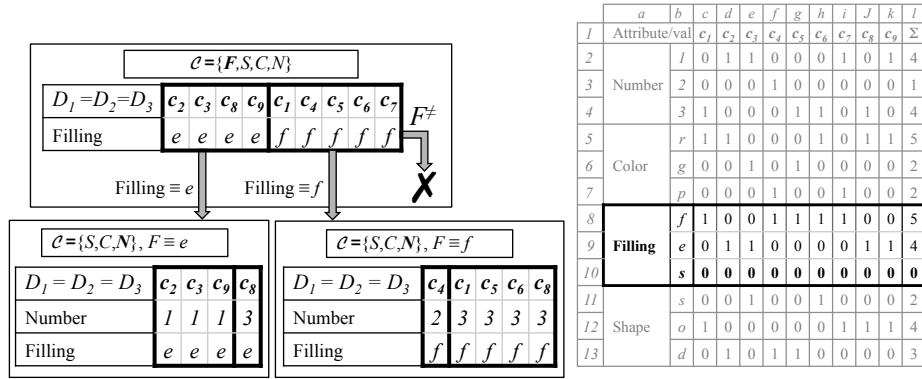


Fig. 12. Branching on equality constraints $F \equiv$.

Fig. 13. Cell 10*l* shows that no card with a *striped* filling exists.

We conclude that we should form only two CSPs: one for the cards with an *empty* filling (i.e., $\{c_2, c_3, c_8, c_9\}$) and the other for the cards with a *full* filling (i.e., $\{c_1, c_4, c_5, c_6, c_7\}$). Because there is no card with a *striped* filling, we should not generate (1) a subproblem for this third possible value or (2) a subproblem where $F \neq$ holds. The two generated subproblems have one fewer constraint than their parent and the domains of the variables in each subproblem are the same. The tree of Figure 12 illustrates the partitioning of the cards c_1, \dots, c_9 as described above, then the partitioning of $\{c_2, c_3, c_8, c_9\}$ and $\{c_1, c_4, c_5, c_6, c_7\}$ given the domain dimension *number*. To detect the above-described situation, our algorithm examines the (detailed) domain table shown in Figure 13. Column *l* sums up the number of cards in the domain for a given attribute value (represented as a row in the table). The null entry in Cell 10*l* indicates that there is no card with a *striped* filling (Row 10) in the domain, and that selecting *filling* as a dimension for reformulation would yield only two subproblems and not four.

4.4 Branching on an inequality constraint

Figure 14 shows the compacted domain table of the six-cards example of Figure 6.

c_1	c_2	c_3	$D_1 = D_2 = D_3$	c_1	c_2	c_3	c_4	c_5	c_6
R	R	R	Number	2	2	1	1	3	1
			Color	r	r	r	r	r	r
			Filling	e	e	e	e	e	e
			Shape	o	d	s	d	o	o

Fig. 14. The compacted domain table of the example of Figure 6.




Fig. 15. $D_1 \neq D_2 \neq D_3$ when enforcing $N \neq$.

	a	b	c	d	e	f	g	h	l
1	Attribute/val	c_1	c_2	c_3	c_4	c_5	c_6	Σ	
2		1	0	0	1	1	0	1	3
3	Number	2	1	1	0	0	0	0	2
4		3	0	0	0	0	1	0	1
11		s	0	0	1	0	0	0	1
12	Shape	o	1	0	0	0	1	1	3
13		d	0	1	0	1	0	0	2

Fig. 16. Cells 3l and 4l indicate that the subproblems for values 2 and 3 need not be generated.

The remaining one-dimensional constraints are over *shape* and *number*. The reformulation step can choose either dimension. Choosing the dimension *number* partitions the variables' domains into three sets of cards: $\{c_3, c_4, c_6\}$, $\{c_1, c_2\}$, and $\{c_5\}$ with *number* values 1, 2, and 3, respectively, see Figure 15. The subproblems with the domains $\{c_1, c_2\}$ and $\{c_5\}$ have less than three cards and are thus unsolvable. Our algorithm realizes that fact by checking on the values in cells 3l and 4l of Figure 16.⁸ Now, in the subproblem where $N \neq$ holds, the three CSP variables have all-different domains, one for each value of the attribute *number* (i.e., 1, 2, and 3). In summary, because the entries in 3l and 4l of Figure 16 are in $[1, 3)$ we generate the the subproblem for $N \neq$ but not the one with domains $\{c_1, c_2\}$ or $\{c_5\}$.

5 Reformulation Algorithm

Below we describe our reformulation algorithm for the game of SET. It is motivated by the examples of Figures 12 and 14 and implementing the general strategy of Figure 1.

5.1 Decomposition Tree for SET

The tree nodes maintain the following information about the subproblem at the node:

⁸ Only two domain dimensions are shown in Figure 16, reflecting the constraints applicable on the problem on top of Figure 14.

1. *Set of variables*: The set of variables and their domains.
2. *Set of constraints \mathcal{C}* : The set of applicable constraints. The one-dimensional constraints are reduced by one from a parent to a child.
3. *Domain table*: The table describing the multi-dimensional domain such as the one shown in Figure 4 to which we add the column indexed l as shown in Figures 13, 16, and Table 1. The l column sums up the values of the c_i 's in each row in the corresponding table. It indicates the number of cards with the corresponding attribute value. This domain table is useful for choosing the attribute to branch on as illustrated in Sections 4.3 and 4.4.

Table 1. Domain table.

l	a	b	c	d	e	f	g	h	i	j	k	l
1	Attribute/val	c_1	c_2	c_3	c_4	c_5	c_6	c_7	c_8	c_9	Σ	
2		l	1	0	0	0	0	1	0	0	0	2
3	Number	2	0	1	0	0	1	0	1	0	1	4
4		3	0	0	1	1	0	0	0	1	0	3
5		r	1	0	0	0	0	1	0	0	0	2
6	Color	g	0	1	0	1	1	0	0	0	1	4
7		p	0	0	1	0	0	0	1	1	0	3
8		f	1	0	0	0	0	0	0	1	1	3
9	Filling	e	0	1	0	0	0	0	0	0	0	1
10		s	0	0	1	1	1	1	1	0	0	5
11		s	1	0	0	0	0	0	0	0	1	2
12	Shape	o	0	1	0	1	1	1	0	0	0	4
13		d	0	0	1	0	0	0	1	1	0	3

Table 2. Summary of domain table.

l	a	b	c	d	e	f	g	h
1	Attribute/val	Σ_{D_1}	Σ_{D_2}	Σ_{D_3}	Π_{D_1}	Σ_{D_1}	$\Pi\Sigma_{D_1}$	
2		l	1	1	0	0	2	
3	Number	2	1	1	2	2	4	24
4		3	1	1	1	1	3	
5		r	1	0	0	0	1	
6	Color	g	1	2	1	2	4	16
7		p	1	1	2	2	4	
8		f	1	0	2	0	3	
9	Filling	e	1	0	0	0	1	12
10		s	1	3	1	3	4	
11		s	1	0	1	0	2	
12	Shape	o	1	3	0	0	4	24
13		d	1	0	2	0	3	

Note that only the entries corresponding to the applicable constraints need to be represented and updated. Entries corresponding to constraints enforced in an ancestor node in the decomposition tree are omitted as in Figure 16.

4. *Summary of domain table*: When the variables' domains are all different (i.e., $D_1 \neq D_2 \neq D_3$), we generate an additional table that is the *summary of the domain table*, see Table 2. Here again, we keep and maintain only entries (i.e., rows) corresponding to applicable constraints. Columns c , d , and e sum up the number of cards with a given attribute value in the corresponding domain. Columns f and g are the product and sum, respectively, of the entries in columns c , d , and e in the same row. Finally, column h is the product of the values in column g for the same attribute. The heuristics introduced below justify the use of this summary information.

5.2 Flow Chart

Our algorithm, see flow chart in Figure 17, operates under the following assumptions. An agenda is used to keep track of all 'open problems.' At the start, the original problem is placed on the agenda. Solving with BT search indicates finding all the solutions to a subproblem using the search procedure described in Section 3.2. The solution sets found are stored in some unspecified data structure. One may choose to enforce some level of consistency on each generated subproblem before placing it the agenda. Subproblems that are deemed unsolvable are not added to the agenda.

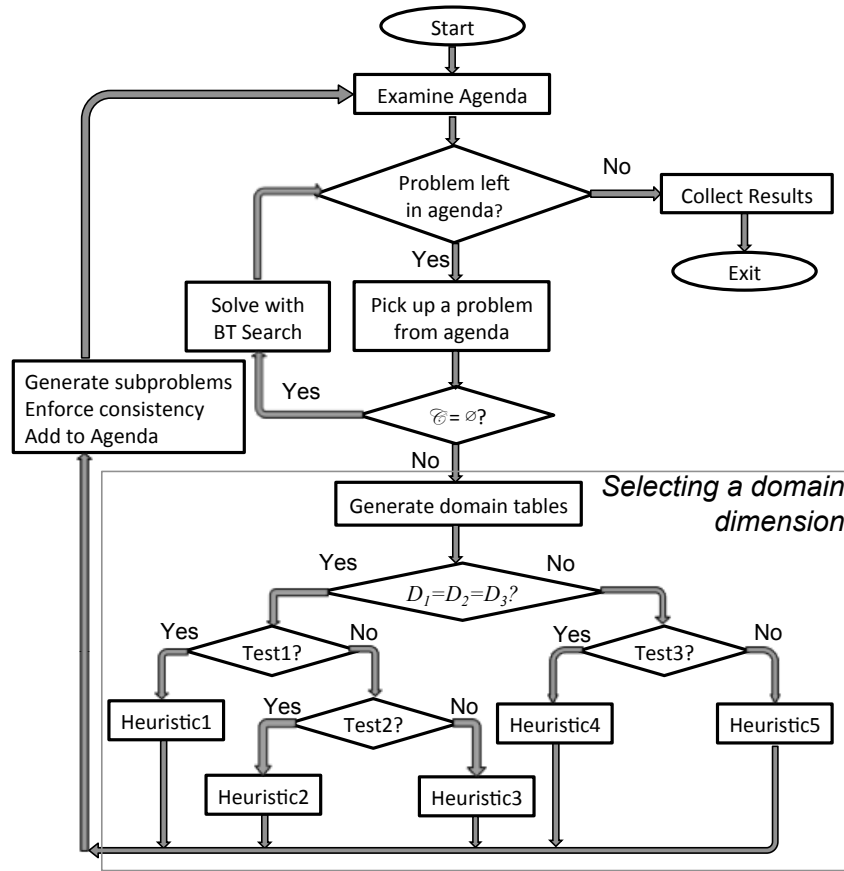


Fig. 17. Our reformulation algorithm.

The algorithm uses the following test conditions and heuristics. Below, we distinguish problems with identical domains (i.e., $D_1 = D_2 = D_3$) and the others (i.e., $D_1 \neq D_2 \neq D_3$). When $D_1 = D_2 = D_3$, we examine only the domain table:

- *Test1*: When the column l in the domain table of the subproblem has a *null* entry, the two implications are enforced:
 1. No card has that dimension value, the corresponding subproblem need not be generated. For example, *striped* for dimension filling in Figure 12.
 2. For the same reason, we cannot form a subproblem where domains are all different. For example, F^\neq in Figure 12 and N^\neq Figure 15.
- *Heuristic1*: We branch on the attribute with the largest number of zeros in the column l because branching on this attribute yields the smallest number of new subproblems, and will not require creating an A^\neq subproblem.
- *Test2*: When the value in column l for any value i for a dimension A is in $[1, 3)$, the A^i subproblem need not be generated. For example, values 2 and 3 for the dimension *number* in Figure 15.

- *Heuristic2*: We branch on the dimension that has the largest number of entries in column l that are less than three. When there are no null entries in column l for any dimension, all four subproblems must be generated.
- *Heuristic3*: We randomly choose any of the ‘active’ dimensions. This step requires generating all four subproblems.

When $D_1 \neq D_2 \neq D_3$, we consider only the table called ‘summary of domain table’ (see example in Table 2). The columns c , d , and e indicate the size of domain D_1 , D_2 , and D_3 . Their product is available in column f of the table: it is zero if some domain is empty. Column g computes the number of cards that have the same value for the same dimension, the product for all the products for the same dimension is recorded in column h . A null value in column h indicates that the dimension needs to be chosen in priority because it would yield the generation of at most two subproblems.

- *Test3*: The entry in column h for an attribute A of the summary of domain table is null. In this case, there is at least one value i for A that appears in no domain.
- *Heuristic4*: We branch on the attribute where column h is null, breaking ties in favor of the attribute with the largest number of null entries in column f . Note that when an entry in column f is null (for an attribute value i), the subproblem A^i and that where A^{\neq} holds will have at least one empty domain, and need not be generated (or will be pruned).
- *Heuristic5*: We branch on the attribute where column f has the largest number of null entries. We generate all four subproblems, those for the attribute value i with a null entry in column f will have an empty domain and will be pruned.

The above three tests cover all cases, and yield a complete and sound algorithm for selecting a dimension for reformulation. Generating and storing the various tables is linear in the number of attribute values. Also, once a problem is decomposed, the corresponding tables can be discarded. The number of generated subproblems is $\mathcal{O}(|A|^{i+1})$ where $|A|$ is the number of attributes and i is the maximum number of attribute values.

6 Results

To evaluate the effectiveness of our reformulation strategy, we run the following three algorithms on a set of 1000 random SET instances of domain size $\{3, 4, \dots, 81\}$: Brute force, search (Section 3.2), and reformulation followed by search (Section 5). ‘Brute force’ is the algorithm with the three nested *for*-loops that generates all combinations of three cards then tests whether or not they satisfy the constraints. Figure 18 shows the number of constraint checks for increasing domain size. Figure 18 shows the number of constraint checks (#CC) for increasing domain size; Figure 19 shows the number of nodes visited (#NV) for increasing domain size; and Table 3 compares the performance of all three algorithms on two domain sizes (12 and 81 cards) displaying the average number of solutions, #CC, #NV, and CPU time. The goal is more to compare the trends than the raw numbers. From those results, it is apparent that reformulation dramatically reduces the rapid growth, with increasing the domain size, of the number of constraint checks (Figure 18) and significantly that of nodes visited (Figure 19), thus demonstrating the benefits of our reformulations. The time measurements are not significant given

the size of the problem, the precision of the clock (i.e., 10 ms), and the time necessary for setting up the data structures for search. While it is true that the game of SET is a simple problem, we do believe that our techniques are widely applicable and will have significant impact on industrial size applications.




Fig. 18. Comparing the numbers of constraint checks for increasing problem size.




Fig. 19. Comparing the number of nodes visited for increasing problem size. Brute force was not included given its large values.

Table 3. Comparing the performance of three algorithms.

Algorithm	#Cards	#Sol	#CC	#NV	Time [msec]
Brute force	12	2.77	1956.80	220	0
Search			1726.65	80.77	62.46
Reformulation			85.08	12.65	5.85
Brute force	81	1080	758808	85320	0
Search			553365.00	4401.00	101.04
Reformulation			31158.00	2565.00	39.44

7 An Interactive Interface for SET

We have built a graphical user interface for SET (see Figure 20) that uses the two approaches for finding all solution sets for a given set of cards on the board. The interface allows us to compare our solvers performance in the context of a real game of SET. The game can be played in two modes: “Single Player” and “Two Players.” The interface features the two automated solvers: by Backtrack Search (Section 3.2) and by Reformulation (Section 5.2). The users can play the game in CSP mode or non-CSP mode. If they are in CSP mode, they can see statistics about the current board (number of constraint check, CPU time, etc.) as well as switch between the two solvers. We have also provided a “2 Card Hint” button, which highlights two cards appearing in the same solution set, letting the user find the third. We also have a “1 Card Hint” button, which highlights one card that appears in a solution set, letting the user find the remaining two cards. If at any time there is no solution found for the twelve cards on the board, three more cards will flip open creating a board of up to 21 cards. If there are no solutions found in the cards displayed, the game is over, and the user can start again. Our applet is available online on <http://gameofset.unl.edu>.

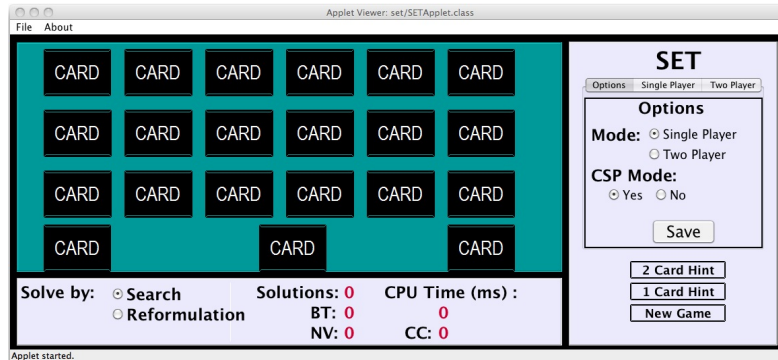


Fig. 20. The graphical user interface.

8 Discussion & Related Work

Like most popular puzzles, the rules of SET are rather simple yet the game is quite addictive. We find combinatorial puzzles to be effective vehicles to introduce the general public to Constraint Processing (CP). They are also amazingly suitable tools to attract Computer Science students to study CP and train them in modeling, search, and propagation techniques. For example, past students have developed constraint models and various propagation algorithms for Minesweeper⁹ [Bayer, Snyder, & Choueiry, 2006] and Sudoku¹⁰ [Reeson *et al.*, 2007]. In addition to the educational benefits, our initiatives have inspired new research directions as documented in [Karakashian *et al.*, 2010b; Woodward *et al.*, 2011]. Combinatorial puzzles have thus allowed us to serve all three tenants of the academic mission, namely, research, education, and outreach.

In [1995; 1997], Freuder & Sabin described an abstraction procedure for solving multi-dimensional CSPs, considering both multi-dimensional constraints and multi-dimensional domains. They evaluated them on n -queens problems in [1995; 1997], and on randomly generated problems with a controlled level of interchangeability in [1997]. Our motivation and procedures significantly overlaps with theirs: (1) Their reformulation considers a single reformulation step, while ours is designed to accommodate any number of one-dimensional constraints for reformulation; (2) After reformulation and search, their technique includes a refinement step, which is not needed in our approach; and (3) Finally, their experiments, conducted on random problems, do not provide the strikingly convincing results that our approach does.

9 Conclusions & Future Work

In conclusion, we have created a reformulation strategy for multi-dimensional CSPs that shows a significant reduction in the search effort. We believe that the techniques we have proposed here can be applied to problems with more complex domains and with more

⁹ <http://minesweeper.unl.edu>

¹⁰ <http://sudoku.unl.edu>

of a real-world significance. In addition to building a graphical tool to explain problem solving by reformulation to students and the general public, there are several directions for future work that we would like to explore. The first is to build a general theory of reformulation for multi-dimensional CSPs that unifies the one proposed in [Freuder & Sabin, 1997] and the one we have proposed above. The second is to investigate the applicability and usefulness of such a theory for general CSPs especially in light of the advances in the study of interchangeability and symmetry in CSPs [Freuder, 1991; Gent, Petrie, & Puget, 2006; Karakashian *et al.*, 2010a]. Finally, we would like to investigate whether the definition of a multi-dimensional CSP provided in [Yoshikawa & Wada, 1992] needs to be revised to allow variable domains to be different.

Acknowledgments This material is based in part upon works supported by the National Science Foundation under Grant No. CCF-0747009, Grant No. CNS-0855139, and Grant No. RI-1117956, and by Science Foundation Ireland under Grant No. 05/IN/1886.

References

- [Bayer, Snyder, & Choueiry, 2006] Bayer, K.; Snyder, J.; and Choueiry, B. Y. 2006. An Interactive Constraint-Based Approach to Minesweeper. In *Proc. of AAAI-2006*, 1933–1934.
- [Davis & McLagan, 2003] Davis, B. L., and McLagan, D. 2003. The Card Game SET. *The Mathematical Intelligencer* 25 (3):33–40.
- [Freuder & Sabin, 1995] Freuder, E. C., and Sabin, D. 1995. Interchangeability Supports Abstraction and Reformulation for Constraint Satisfaction. In *Symposium on Abstraction, Reformulation and Approximation, SARA'95*, 62–68.
- [Freuder & Sabin, 1997] Freuder, E. C., and Sabin, D. 1997. Interchangeability Supports Abstraction and Reformulation for Multi-Dimensional Constraint Satisfaction. In *Proc. of AAAI-97*, 191–196.
- [Freuder, 1991] Freuder, E. C. 1991. Eliminating Interchangeable Values in Constraint Satisfaction Problems. In *Proc. of AAAI-91*, 227–233.
- [Gent, Petrie, & Puget, 2006] Gent, I.; Petrie, K.; and Puget, J.-F. 2006. *Handbook of Constraint Programming*. Elsevier. chapter 10, 329–376.
- [Karakashian *et al.*, 2010a] Karakashian, S.; Woodward, R.; Choueiry, B. Y.; Prestwich, S.; and Freuder, E. C. 2010a. A Partial Taxonomy of Substitutability and Interchangeability. In *International Workshop on Symmetry in Constraint Satisfaction Problems (SymCon 10)*, 1–18.
- [Karakashian *et al.*, 2010b] Karakashian, S.; Woodward, R.; Reeson, C.; Choueiry, B. Y.; and Bessiere, C. 2010b. A First Practical Algorithm for High Levels of Relational Consistency. In *Proc. of AAAI-2010*, 101–107.
- [Prosser, 1993] Prosser, P. 1993. Hybrid Algorithms for the Constraint Satisfaction Problem. *Computational Intelligence* 9 (3):268–299.
- [Reeson *et al.*, 2007] Reeson, C. G.; Huang, K.-C.; Bayer, K. M.; and Choueiry, B. Y. 2007. An Interactive Constraint-Based Approach to Sudoku. In *Proc. of AAAI-2007*, 1976–1977.
- [Woodward *et al.*, 2011] Woodward, R.; Karakashian, S.; Choueiry, B. Y.; and Bessiere, C. 2011. Solving Difficult CSPs with Relational Neighborhood Inverse Consistency. In *Proc. of AAAI-2011*, 112–119.
- [Yoshikawa & Wada, 1992] Yoshikawa, M., and Wada, S. 1992. Constraint Satisfaction in A Multi-Dimensional Domain. In *First International Conference on Artificial Intelligence Planning Systems (AIPS 92)*, 252–259.