

Sharing is Caring: Combination of Theories^{*}

Dejan Jovanović and Clark Barrett

New York University

New York University Technical Report: TR2011-940

Abstract. One of the main shortcomings of the traditional methods for combining theories is the complexity of guessing the arrangement of the variables shared by the individual theories. This paper presents a reformulation of the Nelson-Oppen method that takes into account explicit equality propagation and can ignore pairs of shared variables that the theories do not care about. We show the correctness of the new approach and present care functions for the theory of uninterpreted functions and the theory of arrays. The effectiveness of the new method is illustrated by experimental results demonstrating a dramatic performance improvement on benchmarks combining arrays and bit-vectors.

1 Introduction

The seminal paper of Nelson and Oppen [14] introduced a general framework for combining quantifier-free first-order theories in a modular fashion. Using the Nelson-Oppen framework, decision procedures for two individual theories can be used as black boxes to create a decision procedure for the combined theory. Although the Nelson-Oppen combination method as originally formulated requires stably-infinite theories, it can be extended to handle non-stably-infinite theories using an approach based on polite theories [12, 13, 17].

The core idea driving the method (and ensuring its correctness) is the exchange of equalities and disequalities over the interface variables between the theories involved in the combination. Interface variables are the problem variables that are shared by both theories (or an extended set of variables in the polite combination framework), and both theories must agree on an arrangement over these variables. Most modern satisfiability modulo theories (SMT) solvers perform the search for such an arrangement by first using aggressive theory propagation to determine as much of the arrangement as possible and then relying on an efficient SAT solver to guess the rest of the arrangement, backtracking and learning lemmas as necessary [1, 3, 6].

In some cases, if the theories that are being combined have additional properties, such as convexity and/or complete and efficient equality propagation, there are more efficient ways of obtaining a suitable arrangement. But, in general, since the number of shared variables can be substantial, guessing an arrangement over the shared variables can have an exponential impact¹ on the running time [15]. Trying to minimize the burden of non-deterministic guessing is thus of the utmost importance for a practical and efficient combination mechanism. For example, a recent model-based theory combination approach [7], in which the solver keeps a model for each theory, takes the optimistic stance of eagerly propagating all equalities that hold in the model (whether or not they are truly implied), obtaining impressive performance improvements.

In this paper we tackle the problem of minimizing the amount of non-deterministic guessing by equipping the theories with an *equality propagator* and a *care function*. The role of the theory-specific equality propagator is, given a context, to propagate entailed equalities and disequalities over the interface variables. The care function, on the other hand, provides information about which variable pairs among the interface variables are important for maintaining the satisfiability of a given formula. With the information provided by these two functions we can, in many cases, drastically reduce the search space for finding a suitable arrangement. We present a reformulation of the Nelson-Oppen method that uses these two functions to decide a combination of two theories. The method can easily be adapted to the combination method for polite theories, where reducing the number of shared variables is even more important (the polite theory combination method requires extending the set of interface variables significantly).

^{*} This work was funded in part by SRC contract 2008-TJ-1850.

¹ If the two theories can be decided in time $O(\mathcal{T}_1(n))$ and $O(\mathcal{T}_2(n))$, the combination can be decided in $O(2^{n^2} \times (\mathcal{T}_1(n) + \mathcal{T}_2(n)))$.

2 Preliminaries

We start with a brief overview of the syntax and semantics of many-sorted first-order logic. For a more detailed exposition, we refer the reader to [11, 20].

A *signature* Σ is a triple (S, F, P) where S is a set of *sorts*, F is a set of *function symbols*, and P is a set of *predicate symbols*. For a signature $\Sigma = (S, F, P)$, we write Σ^S for the set S of sorts, Σ^F for the set F of function symbols, and Σ^P for the set P of predicates. Each predicate and function symbol is associated with an *arity*, a tuple constructed from the sorts in S . Functions whose arity is a single sort are called *constants*. We write $\Sigma_1 \cup \Sigma_2 = (S_1 \cup S_2, F_1 \cup F_2, P_1 \cup P_2)$ for the union² of signatures $\Sigma_1 = (S_1, F_1, P_1)$ and $\Sigma_2 = (S_2, F_2, P_2)$. Additionally, we write $\Sigma_1 \subseteq \Sigma_2$ if $S_1 \subseteq S_2$, $F_1 \subseteq F_2$, $P_1 \subseteq P_2$, and the symbols of Σ_1 have the same arity as those in Σ_2 . We assume the standard notions of a Σ -*term*, Σ -*literal*, and Σ -*formula*. In the following, we assume that all formulas are quantifier-free, if not explicitly stated otherwise. A literal is called *flat* if it is of the form $x = y$, $x \neq y$, $x = f(y_1, \dots, y_n)$, $p(y_1, \dots, y_n)$, or $\neg p(y_1, \dots, y_n)$, where x, y, y_1, \dots, y_n are variables, f is a function symbol, and p is a predicate symbol. If ϕ is a term or a formula, we will denote by $\text{vars}_\sigma(\phi)$ the set of variables of sort σ that occur (free) in ϕ . We overload this function in the usual way, $\text{vars}_S(\phi)$ denoting variables in ϕ of the sorts in S , and $\text{vars}(\phi)$ denoting all variables in ϕ . We also sometimes refer to a set Φ of formulas as if it were a single formula, in which case the intended meaning is the conjunction $\bigwedge \Phi$ of the formulas in the set.

Let Σ be a signature, and let X be a set of variables whose sorts are in Σ^S . A Σ -*interpretation* \mathcal{A} over X is a map that interprets each sort $\sigma \in \Sigma^S$ as a non-empty domain A_σ ,³ each variable $x \in X$ of sort σ as an element $x^{\mathcal{A}} \in A_\sigma$, each function symbol $f \in \Sigma^F$ of arity $\sigma_1 \times \dots \times \sigma_n \times \tau$ as a function $f^{\mathcal{A}} : A_{\sigma_1} \times \dots \times A_{\sigma_n} \rightarrow A_\tau$, and each predicate symbol $p \in \Sigma^P$ of arity $\sigma_1 \times \dots \times \sigma_n$ as a subset $p^{\mathcal{A}}$ of $A_{\sigma_1} \times \dots \times A_{\sigma_n}$. A Σ -*structure* is a Σ -interpretation over an empty set of variables. As usual, the interpretations of terms and formulas in an interpretation \mathcal{A} are defined inductively over their structure. For a term t , we denote with $t^{\mathcal{A}}$ the evaluation of t under the interpretation \mathcal{A} . Likewise, for a formula ϕ , we denote with $\phi^{\mathcal{A}}$ the truth-value (true or false) of ϕ under interpretation \mathcal{A} . A Σ -formula ϕ is *satisfiable* iff it evaluates to true in some Σ -interpretation over (at least) $\text{vars}(\phi)$. Let \mathcal{A} be an Ω -interpretation over some set V of variables. For a signature $\Sigma \subseteq \Omega$, and a set of variables $U \subseteq V$, we denote with $\mathcal{A}^{\Sigma, U}$ the interpretation obtained from \mathcal{A} by restricting it to interpret only the symbols in Σ and the variables in U .

We will use the definition of theories as classes of structures, rather than sets of sentences. We define a theory formally as follows (see e.g. [19] and Definition 2 in [17]).

Definition 1 (Theory). *Given a set of Σ -sentences \mathbf{Ax} a Σ -theory $T_{\mathbf{Ax}}$ is a pair (Σ, \mathbf{A}) where Σ is a signature and \mathbf{A} is the class of Σ -structures that satisfy \mathbf{Ax} .*

Given a theory $T = (\Sigma, \mathbf{A})$, a T -*interpretation* is a Σ -interpretation \mathcal{A} such that $\mathcal{A}^{\Sigma, \emptyset} \in \mathbf{A}$. A Σ -formula ϕ is T -satisfiable iff it is satisfiable in some T -interpretation \mathcal{A} . This is denoted as $\mathcal{A} \models_T \phi$, or just $\mathcal{A} \models \phi$ if the theory is clear from the context.

As theories in our formalism are represented by classes of structures, a combination of two theories is represented by those structures that can interpret both theories (Definition 3 in [17]).

Definition 2 (Combination). *Let $T_1 = (\Sigma_1, \mathbf{A}_1)$ and $T_2 = (\Sigma_2, \mathbf{A}_2)$ be two theories. The combination of T_1 and T_2 is the theory $T_1 \oplus T_2 = (\Sigma, \mathbf{A})$ where $\Sigma = \Sigma_1 \cup \Sigma_2$ and $\mathbf{A} = \{\Sigma\text{-structures } \mathcal{A} \mid \mathcal{A}^{\Sigma_1, \emptyset} \in \mathbf{A}_1 \text{ and } \mathcal{A}^{\Sigma_2, \emptyset} \in \mathbf{A}_2\}$.*

The set of Σ -structures resulting from the combination of two theories is indeed a theory in the sense of Definition 1. If \mathbf{Ax}_1 is the set of sentences defining theory T_1 , and \mathbf{Ax}_2 is the set of sentences defining theory T_2 , then \mathbf{A} is the set of Σ -structures that satisfy the set $\mathbf{Ax} = \mathbf{Ax}_1 \cup \mathbf{Ax}_2$ (see Proposition 4 in [17]).

Given decision procedures for the satisfiability of formulas in theories T_1 and T_2 , we are interested in constructing a decision procedure for satisfiability in $T_1 \oplus T_2$ using these procedures as black boxes. The

² In this paper, we always assume that function and predicate symbols from different theories do not overlap, so that the union operation is well-defined. On the other hand, two different theories are allowed to have non-disjoint sets of sorts.

³ In the rest of the paper we will use the calligraphic letters $\mathcal{A}, \mathcal{B}, \dots$ to denote interpretations, and the corresponding subscripted Roman letters $A_\sigma, B_\sigma, \dots$ to denote the domains of the interpretations.

Nelson-Oppen combination method [14, 19, 20] gives a general mechanism for doing this. Given a formula ϕ over the combined signature $\Sigma_1 \cup \Sigma_2$, the first step is to *purify* ϕ by constructing an equisatisfiable set of formulas $\phi_1 \cup \phi_2$ such that each ϕ_i consists of only Σ_i -formulas. This can easily be done by finding a pure (i.e. Σ_i - for some i) subterm t , replacing it with a new variable v , adding the equation $v = t$, and then repeating this process until all formulas are pure. The next step is to force the decision procedures for the individual theories to agree on whether variables appearing in both ϕ_1 and ϕ_2 (called *shared* or *interface* variables) are equal. This is done by introducing an *arrangement* over the shared variables [17, 19].

Here we will use a more general definition of an arrangement that allows us to restrict the pairs of variables that we are interested in. We do so by introducing the notion of a *care graph*. Given a set of variables V , we will call any graph $\mathbf{G} = \langle V, E \rangle$ a care graph, where $E \subseteq V \times V$ is the set of care graph edges. If an edge $(x, y) \in E$ is present in the care graph, it means that we are interested in the relationship between the variables x and y .

Definition 3 (Arrangement). *Given a care graph $\mathbf{G} = \langle V, E \rangle$ where sorts of variables in V range over a set of sorts S , with $V_\sigma = \text{vars}_\sigma(V)$, we call $\delta_{\mathbf{G}}$ an arrangement over \mathbf{G} if there exists a family of equivalence relations*

$$\mathcal{E} = \{ E_\sigma \subseteq V_\sigma \times V_\sigma \mid \sigma \in S \} ,$$

such that the equivalence relations restricted to E induce $\delta_{\mathbf{G}}$, i.e. $\delta_{\mathbf{G}} = \bigcup_{\sigma \in S} \delta_\sigma$, where each δ_σ is an individual arrangement of V_σ (restricted to E):

$$\delta_\sigma = \{ x = y \mid (x, y) \in E_\sigma \cap E \} \cup \{ x \neq y \mid (x, y) \in \overline{E}_\sigma \cap E \} ,$$

where \overline{E}_σ denotes the complement of E_σ (i.e. $V_\sigma \times V_\sigma \setminus E_\sigma$). If the care graph \mathbf{G} is a complete graph over V , we will denote the arrangement simply as δ_V .

The Nelson-Oppen combination theorem states that ϕ is satisfiable in $T_1 \oplus T_2$ iff there exists an arrangement δ_V of the shared variables $V = \text{vars}(\phi_1) \cap \text{vars}(\phi_2)$ such that $\phi_i \cup \delta_V$ is satisfiable in T_i . However, as mentioned earlier, some restrictions on the theories are necessary in order for the method to be complete. Sufficient conditions for completeness are: the two signatures have no function or predicate symbols in common; and the two theories are *stably-infinite* over (at least) the set of common sorts $\Sigma_1^S \cap \Sigma_2^S$. Stable-infiniteness was originally introduced in a single-sorted setting [15]. In the many-sorted setting stable-infiniteness is defined with respect to a subset of the signature sorts (see Definition 6 from [20]).

3 New Combination Method

In this section we present a new method for combining two signature-disjoint theories. The method is based on Nelson-Oppen, but it makes equality propagation explicit and also includes a *care function* for each theory, enabling a more efficient mechanism for determining equalities and dis-equalities among the shared variables. Another notable difference from the original method is that we depart from viewing the combination problem as symmetric. Instead, as in the method for combining polite theories [12, 13, 17], one of the theories is designated to take the lead in selecting which variable pairs are going to be part of the final arrangement.

We first define the equality propagator and the care function, and then proceed to presenting and proving the combination method.

Definition 4 (Equality Propagator). *For a Σ -theory T we call a function $\mathfrak{P}_T^{\overline{[\cdot]}}$ an equality propagator for T if, for every set V of variables, it maps every set ϕ of flat Σ -literals into a set of equalities and dis-equalities between variables:*

$$\mathfrak{P}_T^{\overline{[\cdot]}}[V](\phi) = \{x_1 = y_1, \dots, x_m = y_m\} \cup \{z_1 \neq w_1, \dots, z_n \neq w_n\} ,$$

where $\text{vars}(\mathfrak{P}_T^{\overline{[\cdot]}}[V](\phi)) \subseteq V$ and

1. *for each equality $x_i = y_i \in \mathfrak{P}_T^{\overline{[\cdot]}}[V](\phi)$ it holds that $\phi \models_T x_i = y_i$;*
2. *for each dis-equality $z_i \neq w_i \in \mathfrak{P}_T^{\overline{[\cdot]}}[V](\phi)$ it holds that $\phi \models_T z_i \neq w_i$;*
3. *$\mathfrak{P}_T^{\overline{[\cdot]}}[V]$ is monotone, i.e. $\phi \subseteq \psi \implies \mathfrak{P}_T^{\overline{[\cdot]}}[V](\phi) \subseteq \mathfrak{P}_T^{\overline{[\cdot]}}[V](\psi)$; and*

4. $\mathfrak{P}_T^{\equiv}[V]$ contains at least those equalities and dis-equalities, over variables in V , that appear in ϕ .

An equality propagator, given a set of theory literals, returns a set of entailed equalities and dis-equalities between the variables in V . It does not need to be complete (i.e. it does not need to return *all* entailed equalities and dis-equalities), but the more complete it is, the more helpful it is in reducing the arrangement search space.

When combining two theories, the combined theory can provide more equality propagation than just the union of the individual propagators. The following construction defines an equality propagator that reuses the individual propagators in order to obtain a propagator for the combined theory. This is achieved by allowing the propagators to incrementally exchange literals until a fix-point is reached.

Definition 5 (Combined Propagator). Let T_1 and T_2 be two theories over the signatures Σ_1 and Σ_2 , equipped with equality propagators $\mathfrak{P}_{T_1}^{\equiv}[\cdot]$ and $\mathfrak{P}_{T_2}^{\equiv}[\cdot]$, respectively. Let $T = T_1 \oplus T_2$ and $\Sigma = \Sigma_1 \cup \Sigma_2$. Let V be a set of variables and ϕ a set of flat Σ -literals partitioned into a set ϕ_1 of Σ_1 -literals and a set ϕ_2 of Σ_2 -literals. We define the combined propagator $\mathfrak{P}_T^{\equiv}[\cdot]$ for the theory T as

$$\mathfrak{P}_T^{\equiv}[V](\phi) = (\mathfrak{P}_{T_1}^{\equiv} \oplus \mathfrak{P}_{T_2}^{\equiv})[V](\phi) = \psi_1^* \cup \psi_2^* ,$$

where $\langle \psi_1^*, \psi_2^* \rangle$ is the least fix-point of the following operator \mathcal{F}

$$\mathcal{F}\langle \psi_1, \psi_2 \rangle = \langle \mathfrak{P}_{T_1}^{\equiv}[V](\phi_1 \cup \psi_2), \mathfrak{P}_{T_2}^{\equiv}[V](\phi_2 \cup \psi_1) \rangle .$$

The fix-point exists as the propagators are monotone and the set V is finite. Moreover, the value of the fix-point is easily computable by iteration from $\langle \emptyset, \emptyset \rangle$. Also, it's clear from the definition that the combined propagator is at least as strong as the individual propagators, i.e. $\mathfrak{P}_{T_1}^{\equiv}[V](\phi_1) \subseteq \mathfrak{P}_T^{\equiv}[V](\phi_1) \subseteq \mathfrak{P}_T^{\equiv}[V](\phi)$, $\mathfrak{P}_{T_2}^{\equiv}[V](\phi_2) \subseteq \mathfrak{P}_T^{\equiv}[V](\phi_2) \subseteq \mathfrak{P}_T^{\equiv}[V](\phi)$.

Definition 6 (Care Function). For a Σ -theory T we call a function $\mathfrak{C}[\cdot]$ a care function for T with respect to a T -equality propagator $\mathfrak{P}_T^{\equiv}[\cdot]$ when for every set V of variables and every set ϕ of flat Σ -literals

1. $\mathfrak{C}[V]$ maps ϕ to a care graph $\mathbf{G} = \langle V, E \rangle$;
2. if $x = y$ or $x \neq y$ are in $\mathfrak{P}_T^{\equiv}[V](\phi)$ then $(x, y) \notin E$;
3. if $\mathbf{G} = \langle V, \emptyset \rangle$ and ϕ is T -satisfiable then, for any arrangement δ_V such that $\mathfrak{P}_T^{\equiv}[V](\phi) \subseteq \delta_V$, it holds that $\phi \cup \delta_V$ is also T -satisfiable.

For any Σ -theory T and a set of variables V , the *trivial care function* $\mathfrak{C}_0[\cdot]$ is the one that maps a set of variables to a complete graph over the pairs of variables that are not yet decided. i.e.

$$\mathfrak{C}_0[V](\phi) = \langle V, \{(x, y) \in V \times V \mid x = y, x \neq y \notin \mathfrak{P}_T^{\equiv}[V](\phi)\} \rangle .$$

Notice that $\mathfrak{C}_0[\cdot]$ trivially satisfies the conditions of Definition 6 with respect to any equality propagator. To see this, the only case to consider is when the care graph returned has no edges and ϕ is satisfiable. But in this case, if $\mathfrak{P}_T^{\equiv}[V](\phi) \subseteq \delta_V$, then we must have $\mathfrak{P}_T^{\equiv}[V](\phi) = \delta_V$, and so clearly $\phi \cup \delta_V$ is satisfiable.

3.1 Combination Method

Let T_i be a Σ_i -theory, for $i = 1, 2$ and let $S = \Sigma_1^{\mathbb{S}} \cap \Sigma_2^{\mathbb{S}}$. Further, assume that each T_i is stably-infinite with respect to S_i , decidable, and equipped with a theory propagator $\mathfrak{P}_{T_i}^{\equiv}[\cdot]$. Additionally, let T_2 be equipped with a care function $\mathfrak{C}_{T_2}[\cdot]$ operating with respect to $\mathfrak{P}_{T_2}^{\equiv}[\cdot]$. We are interested in deciding the combination theory $T = T_1 \oplus T_2$ over the signature $\Sigma = \Sigma_1 \cup \Sigma_2$. We denote the combined theory propagator with $\mathfrak{P}_T^{\equiv}[\cdot]$. The combination method takes as input a set ϕ of Σ -literals and consists of the following steps:

Purify: The output of the purification phase is two new sets of literals, ϕ_1 and ϕ_2 such that $\phi_1 \cup \phi_2$ is equisatisfiable (in T) with ϕ and each literal in ϕ_i is a flat Σ_i -literal, for $i = 1, 2$. This step is identical to the first step in the standard Nelson-Oppen combination method.

Arrange: Let $V = \text{vars}(\phi_1) \cap \text{vars}(\phi_2)$ be the set of all variables shared by ϕ_1 and ϕ_2 . Let the care graph \mathbf{G}_2 be a fix-point of the following operator

$$\mathcal{G}(\mathbf{G}) = \mathbf{G} \cup \mathfrak{C}_{T_2}[[V]](\phi_2 \cup \mathfrak{P}_T^{\bar{=}}[[V]](\phi_1 \cup \phi_2 \cup \delta_{\mathbf{G}})) , \quad (1)$$

where we choose the arrangement $\delta_{\mathbf{G}}$ non-deterministically.

Check: Check the following formulas for satisfiability in T_1 and T_2 respectively

$$\phi_1 \cup \mathfrak{P}_T^{\bar{=}}[[V]](\phi_1 \cup \phi_2 \cup \delta_{\mathbf{G}_2}) , \quad \phi_2 \cup \mathfrak{P}_T^{\bar{=}}[[V]](\phi_1 \cup \phi_2 \cup \delta_{\mathbf{G}_2}) .$$

If both are satisfiable, output satisfiable, otherwise output unsatisfiable.

Notice that above, since the graph is finite, and the the operator \mathcal{G} is increasing, the fix-point always exists. Moreover, it is in our interest to choose the minimal such fix-point, which we can obtain by doing a fix-point iteration starting from $\mathbf{G}_0 = \langle V, \emptyset \rangle$. Another important fact is that for any fix-point \mathbf{G} , with respect to the $\delta_{\mathbf{G}}$ we have chosen, of the operator \mathcal{G} above, we must have that the care function from (1) returns an empty graph. This follows from the fact that the propagator must return all the equalities and dis-equalities from $\delta_{\mathbf{G}}$, by definition, and the care function then must ignore them, also by definition.

Example 1. Consider the case of combining two theories T_1 and T_2 equipped with trivial care functions and propagators $\mathfrak{P}_{T_i}^{\bar{=}}[[V]]$ that simply return those input literals that are either equalities or dis-equalities over variables in V . Assume that ϕ_1 and ϕ_2 are the output of the purification phase, and let V be the set of variables shared by ϕ_1 and ϕ_2 . Since $\mathfrak{C}_{T_2}[[\cdot]]$ is a trivial care function, we will choose a arrangement $\delta_{\mathbf{G}_2}$ over the set V of shared variables that completes the set of equalities and dis-equalities over V . Since equality propagators simply keep the input equalities and dis-equalities over V , and all other relationships between variables in V are determined by $\delta_{\mathbf{G}_2}$, the combined propagator will return a complete arrangement δ_V and we will then check $\phi_1 \cup \delta_V$ and $\phi_2 \cup \delta_V$ for satisfiability. This shows that our method can effectively simulate the standard Nelson-Oppen combination method. We now show the correctness of the method.

Theorem 1. *Let T_i be a Σ_i -theory, stably-infinite with respect to the set of sorts S_i , and equipped with equality propagators $\mathfrak{P}_{T_i}^{\bar{=}}[[\cdot]]$, for $i = 1, 2$. Additionally, let T_2 be equipped with a care function $\mathfrak{C}_{T_2}[[\cdot]]$ operating with respect to $\mathfrak{P}_{T_2}^{\bar{=}}[[\cdot]]$. Let $\Sigma = \Sigma_1 \cup \Sigma_2$, $T = T_1 \oplus T_2$ and let ϕ be a set of flat Σ -literals, which can be partitioned into a set ϕ_1 of Σ_1 -literals and a set ϕ_2 of Σ_2 -literals, with $V = \text{vars}(\phi_1) \cap \text{vars}(\phi_2)$. If $\Sigma_1^{\mathbb{S}} \cap \Sigma_2^{\mathbb{S}} = S_1 \cap S_2$, then following are equivalent*

1. ϕ is T -satisfiable;
2. there exists some care-graph \mathbf{G}_2 , and a corresponding arrangement $\delta_{\mathbf{G}_2}$, that are fix-point solutions of (1), such that the following sets are T_1 - and T_2 -satisfiable respectively

$$\phi_1 \cup \mathfrak{P}_T^{\bar{=}}[[V]](\phi_1 \cup \phi_2 \cup \delta_{\mathbf{G}_2}) , \quad \phi_2 \cup \mathfrak{P}_T^{\bar{=}}[[V]](\phi_1 \cup \phi_2 \cup \delta_{\mathbf{G}_2}) .$$

Moreover, T is stably-infinite with respect to $S_1 \cup S_2$.

Proof. (1) \Rightarrow (2) : Suppose $\phi = \phi_1 \cup \phi_2$ is T -satisfiable in a T -interpretation \mathcal{A} . Let δ_V be the full arrangement over V satisfied by \mathcal{A} . Since δ_V trivially is a fix-point of (1), \mathcal{A} satisfies δ_V , and the propagator only adds formulas that are entailed, it is clear that \mathcal{A} satisfies both sets of formulas, which proves one direction.

(2) \Leftarrow (1) : Assume that there is a T_1 -interpretation \mathcal{A}_1 and a T_2 -interpretation \mathcal{A}_2 (and assume wlog that both interpret all the variables in V) such that $\mathcal{A}_1 \models_{T_1} \phi_1 \cup \mathfrak{P}_T^{\bar{=}}[[V]](\phi_1 \cup \phi_2 \cup \delta_{\mathbf{G}_2})$ and $\mathcal{A}_2 \models_{T_2} \phi_2 \cup \mathfrak{P}_T^{\bar{=}}[[V]](\phi_1 \cup \phi_2 \cup \delta_{\mathbf{G}_2})$. Let δ_V be the arrangement over the complete graph on V satisfied by \mathcal{A}_1 , so

$$\delta_{\mathbf{G}_2} \subseteq \mathfrak{P}_{T_2}^{\bar{=}}[[V]](\phi_2 \cup \delta_{\mathbf{G}_2}) \subseteq \mathfrak{P}_T^{\bar{=}}[[V]](\phi_1 \cup \phi_2 \cup \delta_{\mathbf{G}_2}) \subseteq \delta_V .$$

Because \mathbf{G}_2 is a fix-point, we know that $\mathfrak{C}_{T_2}[[V]](\phi_2 \cup \mathfrak{P}_T^{\bar{=}}[[V]](\phi_1 \cup \phi_2 \cup \delta_{\mathbf{G}_2})) = \langle V, \emptyset \rangle$. We then know, by property 3 of the care function, that there is a T_2 -interpretation \mathcal{B}_2 such that $\mathcal{B}_2 \models_{T_2} \phi_2 \cup \delta_V$. Since δ_V is a complete arrangement over all the shared variables and we also have that $\mathcal{A}_1 \models_{T_1} \phi_1 \cup \delta_V$, we can now appeal to the correctness of the standard Nelson-Oppen combination method to obtain a T -interpretation \mathcal{C} that satisfies $\phi_1 \cup \phi_2 = \phi$. The proof that the combined theory is stably-infinite can be found in [12]. \square

3.2 Extension to Polite Combination

The method described in Section 3 relies on the correctness argument for the standard Nelson-Oppen method, meaning that the theories involved should be stably-infinite for completeness. A more general combination method based on the notion of *polite* theories (and not requiring that both theories be stably-infinite) was introduced in [17] and clarified in [12, 13]. Here, we assume familiarity with the concepts appearing in those papers, and show how they can be integrated into the combination method of this paper.

Assume that the theory T_2 is polite with respect to the set of sorts S_2 such that $\Sigma_1 \cap \Sigma_2 \subseteq S_2$, and is equipped with a witness function $witness_2$. We modify the combination method of Section 3.1 as follows:

1. In the **Arrange** and **Check** phases, instead of using ϕ_2 , we use the formula produced by the witness function, i.e. $\phi'_2 = witness_2(\phi_2)$.
2. We define $V = vars_S(\phi'_2)$ instead of $V = vars(\phi_1) \cap vars(\phi_2)$.

Theorem 2. *Let T_i be a Σ_i -theory polite with respect to the set of sorts S_i , and equipped with equality propagators $\mathfrak{P}_{T_i}^{\equiv}[\cdot]$, for $i = 1, 2$. Additionally, let T_2 be equipped with a care function $\mathfrak{C}_{T_2}[\cdot]$ operating with respect to $\mathfrak{P}_{T_2}^{\equiv}[\cdot]$. Let $\Sigma = \Sigma_1 \cup \Sigma_2$, $T = T_1 \oplus T_2$ and let ϕ be a set of flat Σ -literals, which can be partitioned into a set ϕ_1 of Σ_1 -literals and a set ϕ_2 of Σ_2 -literals. Let $\phi'_2 = witness_{T_2}(\phi_2)$ and $V = vars_S(\phi'_2)$. If $S \subseteq S_1 \cap S_2$, then following are equivalent*

1. ϕ is T -satisfiable;
2. there exists a care-graph \mathbf{G}_2 and arrangement $\delta_{\mathbf{G}_2}$, fix-point solutions of (1), such that the following sets are T_1 - and T_2 -satisfiable respectively

$$\phi_1 \cup \mathfrak{P}_T^{\equiv}[V](\phi_1 \cup \phi'_2 \cup \delta_{\mathbf{G}_2}) , \quad \phi'_2 \cup \mathfrak{P}_T^{\equiv}[V](\phi_1 \cup \phi'_2 \cup \delta_{\mathbf{G}_2}) .$$

Moreover, T is polite with respect to $S_1 \cup (S_2 \setminus \Sigma_1)$.⁴

4 Theory of Uninterpreted Functions

The theory of uninterpreted functions over a signature Σ_{euf} is the theory $T_{\text{euf}} = (\Sigma_{\text{euf}}, \mathbf{A})$, where \mathbf{A} is simply the class of all Σ_{euf} -structures. Conjunctions of literals in this theory can be decided in polynomial time by congruence closure algorithms (e.g. [18]). We make use of insights from these algorithms in defining both the equality propagator and the care function. For simplicity, we assume Σ_{euf} contains no predicate symbols, but the extension to the case with predicate symbols is straightforward.

Equality Propagator. Let ϕ be a set of flat literals, let V be a set of variables, and let \sim_c be the smallest congruence relation⁵ over terms in ϕ containing $\{(x, t) \mid x = t \in \phi\}$. We define a dis-equality relation \neq_c as the smallest relation satisfying

$$x \sim_c x' \wedge y \sim_c y' \wedge x' \neq y' \in \phi \implies x \neq_c y .$$

Now, we define the equality propagator as

$$\mathfrak{P}_{\text{euf}}^{\equiv}[V](\phi) = \{x = y \mid x, y \in V, x \sim_c y\} \cup \{x \neq y \mid x, y \in V, x \neq_c y\}.$$

It is easy to see that $\mathfrak{P}_{\text{euf}}^{\equiv}[\cdot]$ is indeed an equality propagator. Moreover, it can easily be implemented as part of a decision procedure based on congruence closure.

Example 2. Given the set $\phi = \{x = z, y = f(a), z \neq f(a)\}$, the equality propagator would return $\mathfrak{P}_{\text{euf}}^{\equiv}[\{x, y\}](\phi) = \{x = x, y = y, x \neq y, y \neq x\}$.

⁴ The remaining proofs are relegated to an appendix.

⁵ In this context, a congruence relation is an equivalence relation that also satisfies the congruence property: if $f(x_1, \dots, x_n)$ and $f(y_1, \dots, y_n)$ are terms in ϕ , and if for each $1 \leq i \leq n$, $x_i \sim_c y_i$, then $f(x_1, \dots, x_n) \sim_c f(y_1, \dots, y_n)$.

Care Function. The definition of the care function is based on the fact that during congruence closure, we only care about equalities between pairs of variables that occur as arguments in the same position of the same function symbol. Again, let V be a set of variables and let ϕ be a set of flat literals, such that ϕ only contains function symbols from $F = \{f_1, f_2, \dots, f_n\}$.

For a set of formulas ϕ , let $\mathcal{E}(\phi)$ denote the smallest equivalence relation over the terms occurring in ϕ containing $\{(x, t) \mid x = t \in \phi\}$. For an equivalence relation E , let E^* denote the *congruence closure* of E (i.e. the smallest congruence relation containing E). In order to make our care-function more precise, we will first approximate the implications that possible equalities over variables in V could trigger. We do so by taking all possible equalities over V , i.e. let $\delta_{\overline{V}}$ be the full arrangement over the shared variables where all variables of the same sort are equal. Now, to see what these equalities could imply, we let $E_{\phi}^{\overline{V}} = \mathcal{E}(\phi \cup \delta_{\overline{V}})^*$.

For each function symbol $f \in F$ of arity $\sigma_1 \times \sigma_2 \times \dots \times \sigma_k \mapsto \sigma$, let E_f be a set containing pairs of variables that could trigger an application of congruence because of two terms that are applications of f . More precisely, let $E_f \subseteq V \times V$ be a maximal set of pairs $(x, y) \in V \times V$, that are not already decided by the propagator ($x \approx_c y$ and $\neg x \neq_c y$), such that for each $(x, y) \in E_f$ we have:

1. there are x_i and y_i such that $x \sim_c x_i$ and $y \sim_c y_i$;
2. there are terms $f(x_1, \dots, x_i, \dots, x_k) \approx_c f(y_1, \dots, y_i, \dots, y_k)$ in ϕ ;
3. for $1 \leq j \leq k$, variables x_j and y_j could become equal, $(x_j, y_j) \in E_{\phi}^{\overline{V}}$;
4. for $1 \leq j \leq k$, variables x_j and y_j are not known to be disequal, $\neg(x_j \neq_c y_j)$.

Now, we let $E = \bigcup_{f \in F} E_f$, and define the care function mapping ϕ to the care graph \mathbf{G} as $\mathbf{C}_{\text{euf}}[\![V]\!](\phi) = \mathbf{G} = \langle V, E \rangle$.

Example 3. Consider the following sets of literals

$$\begin{aligned} \phi_1 &= \{f(x_1) \neq f(y_1), y_1 = x_2\} \quad , \\ \phi_2 &= \{z_1 = f(x_1), z_2 = f(y_1), g(z_1, x_2) \neq g(z_2, y_2)\} \quad , \\ \phi_3 &= \{y_1 = f(x_1), y_2 = f(x_2), z_1 = g(x_1), z_2 = g(x_2), h(y_1) \neq h(z_1)\} \quad . \end{aligned}$$

and corresponding sets of shared variables $V_1 = \{x_1, x_2\}$, $V_2 = \{x_1, x_2, y_1, y_2\}$, $V_3 = \{x_1, x_2, y_2, z_2\}$. The care function above would return the care graphs $\mathbf{G}_1 = \langle V, \{(x_1, x_2)\} \rangle$, $\mathbf{G}_2 = \langle V, \{(x_1, y_1), (x_2, y_2)\} \rangle$, and $\mathbf{G}_3 = \langle V, \{(x_1, x_2)\} \rangle$.

Note that the the care function for ϕ_3 does not return the pair (y_2, z_2) , which is important in case x_1 and x_2 become equal. This is remedied in the procedure itself, by computing the fix-point, which, in case we choose $x_1 = x_2$, will add the pair (y_2, z_2) to the care graph in the second step.

Theorem 3. *Let T_{euf} be the theory of uninterpreted functions with equality over the signature Σ_{euf} . $\mathbf{C}_{\text{euf}}[\![\cdot]\!]$ is a care function for T_{euf} with respect to the equality propagator $\mathfrak{P}_{\text{euf}}^{\overline{V}}[\![\cdot]\!]$.*

5 Theory of Arrays

The extensional theory of arrays T_{arr} operates over the signature Σ_{arr} that contains the sorts $\{\text{array}, \text{index}, \text{elem}\}$ and function symbols

$$\text{read} : \text{array} \times \text{index} \mapsto \text{elem} \quad , \quad \text{write} : \text{array} \times \text{index} \times \text{elem} \mapsto \text{array} \quad ,$$

where read represents reading from an array at a given index, and write represents writing a given value to an array at an index. The semantics of the theory are given by the three axioms:

1. $\forall a:\text{array}. \forall i:\text{index}. \forall v:\text{elem}. \text{read}(\text{write}(a, i, v), i) = v$,
2. $\forall a:\text{array}. \forall i, j:\text{index}. \forall v:\text{elem}. i \neq j \rightarrow \text{read}(\text{write}(a, i, v), j) = \text{read}(a, j)$,
3. $\forall a, b:\text{array}. (\forall i:\text{index}. \text{read}(a, i) = \text{read}(b, i)) \rightarrow a = b$.

The flat literals of the theory are of the form $x = \text{read}(a, i)$, $a = \text{write}(b, i, x)$, $i = j$, $i \neq j$, $x = y$, $x \neq y$, $a = b$, $a \neq b$, where here and below we use the convention that x, y, v are variables of sort elem , i, j are variables of sort index , a, b , and c are variables of sort array , and w, z are variables of any sort. For a set ϕ of flat T_{arr} -literals, we also define $\alpha(\phi)$ to be the subset of ϕ that does not contain literals of the form $a = \text{write}(b, i, v)$.

Decision Procedure. Before presenting the equality propagator and care function, it will be helpful to present a simple rule-based decision procedure for T_{arr} based on [9].⁶ Given a set Γ of flat T_{arr} -literals, we define \approx_a^Γ as $\mathcal{E}(\alpha(\Gamma))^*$ and the corresponding disequality relation \neq_a^Γ as the smallest relation satisfying:

$$w \approx_a^\Gamma w' \wedge z \approx_a^\Gamma z' \wedge w \neq z \in \Gamma \implies w' \neq_a^\Gamma z' .$$

Additionally, let $\Gamma[l_1, \dots, l_n]$ denote that literals $l_i, 1 \leq i \leq n$ appear in Γ , and for every pair (a, b) of variables in $\text{vars}_{\text{array}}(\Gamma)$, let $k_{a,b}$ be a distinguished fresh variable of sort index. Let \mathcal{D}_{arr} be the following set of inference rules.

$$\begin{aligned} \text{RIntro1} & \frac{\Gamma[a = \text{write}(b, i, v)]}{\Gamma, v = \text{read}(a, i)} \quad \text{if } v \not\approx_a^\Gamma \text{read}(a, i) \\ \text{RIntro2} & \frac{\Gamma[a = \text{write}(b, i, v), x = \text{read}(c, j)]}{\Gamma, i = j \quad \Gamma, \text{read}(a, j) = \text{read}(b, j)} \quad \text{if } \begin{array}{l} a \approx_a^\Gamma c \text{ or } b \approx_a^\Gamma c, \\ i \not\approx_a^\Gamma j, \\ \text{read}(a, j) \not\approx_a^\Gamma \text{read}(b, j) \end{array} \\ \text{ArrDiseq} & \frac{\Gamma[a \neq b]}{\Gamma, \text{read}(a, k_{a,b}) \neq \text{read}(b, k_{a,b})} \quad \text{if } \neg(\text{read}(a, k_{a,b}) \neq_a^\Gamma \text{read}(b, k_{a,b})) \end{aligned}$$

Note that non-flat literals appear in the conclusions of rules RIntro2 and ArrDiseq. We use this as a shorthand for the flattened version of these literals. For example, $\text{read}(a, j) = \text{read}(b, j)$ is shorthand for $x = \text{read}(a, j) \wedge y = \text{read}(b, j) \wedge x = y$, where x and y are fresh variables (there are other possible flattenings, especially if one or more of the terms appears already in Γ , but any of them will do). We say that a set Γ of literals is \mathcal{D}_{arr} -saturated if no rules from \mathcal{D}_{arr} can be applied.

Theorem 4. *The inference rules of \mathcal{D}_{arr} are sound and terminating.*

Theorem 5. *Let Γ be a \mathcal{D}_{arr} -saturated set of flat T_{arr} -literals. Then Γ is T_{arr} -satisfiable iff $\alpha(\Gamma)$ is T_{euf} -satisfiable.*

Equality Propagator. Let ϕ be a set of flat literals and V a set of variables. Consider the following modified versions of RIntro2 that are enabled only if one of the branches can be ruled out as unsatisfiable:

$$\begin{aligned} \text{RIntro2a} & \frac{\Gamma[a = \text{write}(b, i, v), x = \text{read}(c, j)]}{\Gamma, i = j} \quad \text{if } \begin{array}{l} a \approx_a^\Gamma c \text{ or } b \approx_a^\Gamma c, \\ i \not\approx_a^\Gamma j, \\ \text{read}(a, j) \neq_a^\Gamma \text{read}(b, j) \end{array} \\ \text{RIntro2b} & \frac{\Gamma[a = \text{write}(b, i, v), x = \text{read}(c, j)]}{\Gamma, \text{read}(a, j) = \text{read}(b, j)} \quad \text{if } \begin{array}{l} a \approx_a^\Gamma c \text{ or } b \approx_a^\Gamma c, \\ i \neq_a^\Gamma j, \\ \text{read}(a, j) \not\approx_a^\Gamma \text{read}(b, j) \end{array} \end{aligned}$$

Let $\mathcal{D}'_{\text{arr}}$ be obtained from \mathcal{D}_{arr} by replacing RIntro2 with the above rules. Since these rules mimic RIntro2 when they are enabled, but are enabled less often, it is clear that $\mathcal{D}'_{\text{arr}}$ remains sound and terminating. Let Γ' be the result of applying $\mathcal{D}'_{\text{arr}}$ until no more apply (we say that Γ' is $\mathcal{D}'_{\text{arr}}$ -saturated). We define the equality propagator as:

$$\mathfrak{P}_{\text{arr}}^{\text{=}}[\![V]\!] (\phi) = \{w = z \mid w, z \in V, w \approx_a^{\Gamma'} z\} \cup \{w \neq z \mid w, z \in V, w \neq_a^{\Gamma'} z\}.$$

It is easy to see that $\mathfrak{P}_{\text{arr}}^{\text{=}}[\![\cdot]\!]$ satisfies the requirements for a propagator. Though not necessary for the care function we present here, a more powerful propagator can be obtained by additionally performing congruence closure over write terms.

Care Function. Let ϕ be a set of flat literals and V a set of variables. First, since a simple propagator cannot compute all equalities between array variables, we will ensure that the relationships between all pairs of array variables in V have been determined. To do so we define the set E_a^ϕ of pairs of array variables in V that are not yet known equal or dis-equal

$$E_a^\phi = \{(a, b) \in V \times V \mid a \not\approx_a^\phi b \wedge \neg(a \neq_a^\phi b)\} .$$

⁶ The main difference is that in our procedure, we exclude literals containing write from the T_{euf} -satisfiability check as they are not needed and this allows us to have a simpler care function.

Next, since the inference rules can introduce new `read` terms, we compute the smallest set R^ϕ with possible such terms, i.e

- if $x = \text{read}(a, i) \in \phi$ or $a = \text{write}(b, i, v) \in \phi$, then $\text{read}(a, i) \in R^\phi$,
- if $a = \text{write}(b, i, v) \in \phi$, $\text{read}(c, j) \in R^\phi$, $i \not\approx_a^\phi j$, and $a \approx_a^\phi c \vee b \approx_a^\phi c$, then both $\text{read}(a, j) \in R^\phi$ and $\text{read}(b, j) \in R^\phi$,
- if $a \neq b \in \phi$, then both $\text{read}(a, k_{a,b}) \in R^\phi$ and $\text{read}(b, k_{a,b}) \in R^\phi$.

Crucial in the introduction of the above `read` terms, is the set of index variables whose equality could affect the application of the `RIntro2` rule. We capture these variables by defining the set E_i^ϕ as the set of all pairs (i, j) such that:

- $i \not\approx_a^\phi j$ and $\neg(i \neq_a^\phi j)$
- $\exists a, b, c, v. a = \text{write}(b, i, v) \in \phi, \text{read}(c, j) \in R^\phi$, and $a \approx_a^\phi c \vee b \approx_a^\phi c$.

Finally, we claim that with the variables in E_a^ϕ and E_i^ϕ decided, we can essentially use the same care function as for T_{euf} , treating `read` as uninterpreted. We therefore define the third set E_r^ϕ to be the set of all pairs $(i, j) \in V \times V$ of undecided indices, $i \not\approx_a^\phi j$ and $\neg(i \neq_a^\phi j)$, such that there are a, b, i', j' with $a \approx_a^\phi b$, $i \approx_a^\phi i', j \approx_a^\phi j', \text{read}(a, i') \in R^\phi, \text{read}(b, j') \in R^\phi, \text{read}(a, i') \not\approx_a^\phi \text{read}(b, j')$.

With the definitions above, we can define the care graph as $\mathbf{C}_{\text{arr}}[[V]](\phi) = \mathbf{G} = \langle V, E \rangle$, where the set of edges is defined as

$$E = \begin{cases} E_a^\phi & \text{if } E_a^\phi \neq \emptyset, \\ E_i^\phi & \text{if } E_i^\phi \neq \emptyset, \text{ and} \\ E_r^\phi & \text{otherwise.} \end{cases}$$

Note that as defined, E_i^ϕ may include pairs of index variables, one or more of which are not in V . Unfortunately, the care function fails if E_i^ϕ is not a subset of $V \times V$. We can ensure that it is either by expanding the set V until it includes all variables in E_i^ϕ or doing additional case-splitting up front on pairs in E_i^ϕ , adding formulas to ϕ , until $E_i^\phi \subseteq V \times V$.

Theorem 6. *Let T_{arr} be the theory of arrays. $\mathbf{C}_{\text{arr}}[[\cdot]]$ is a care function for T_{arr} with respect to the equality propagator $\mathfrak{P}_{\text{arr}}[[\cdot]]$ for all sets ϕ of literals and V of variables such that $E_i^\phi \subseteq V \times V$.*

Example 4. Consider the following constraints involving arrays and bit-vectors of size m , where \times_m denotes unsigned bit-vector multiplication:

$$\bigwedge_{k=1}^n (\text{read}(a_k, i_k) = \text{read}(a_{k+1}, i_{k+1}) \wedge i_k = x_k \times_m x_{k+1}) \quad . \quad (2)$$

Assume that only the index variables are shared, i.e. $V = \{i_1, \dots, i_{n+1}\}$. In this case, both E_a^ϕ and E_i^ϕ will be empty and the only `read` terms in R^ϕ will be those appearing in the formula. Since none of these are reading from equivalent arrays, the empty care graph is a fix-point for our care function, and we do not need to guess an arrangement.

Note that in the case when V contains `array` variables, the care graph requires us to split on all pairs of these variables (i.e. the care function is trivial over these variables). Fortunately, in practice it appears that index and element variables are typically shared, and only rarely are array variables shared.

6 Experimental Evaluation

We implemented the new method in the `Cvc3` solver [2], and in the discussion below, we denote the new implementation as `Cvc3+C`. We focused our attention on the combination of the theory of arrays and the theory of fixed-size bit-vectors (`QF_AUFBV`). This seemed like a good place to start because there are many benchmarks which generate a non-trivial number of shared variables, and additional splits on shared bit-vector variables can be quite expensive. This allowed us to truly examine the merits of the new combination

Table 1. Experimental results.

	Boolector		Yices		MathSAT		Z3		Cvc3		Cvc3+C	
crafted (40)	2100.13	40	6253.32	34	468.73	30	112.88	40	388.29	9	14.22	40
matrix (11)	1208.16	10	683.84	6	474.89	4	927.12	11	831.29	11	45.08	11
unconstr (10)	3.00	10		0	706.02	3	54.60	2	185.00	5	340.27	8
copy (19)	11.76	19	1.39	19	1103.13	19	4.79	19	432.72	17	44.75	19
sort (6)	691.06	6	557.23	4	82.21	4	248.94	3	44.89	6	44.87	6
delete (29)	3407.68	18	1170.93	10	2626.20	14	1504.46	10	1766.91	17	1302.32	17
member (24)	2807.78	24	185.54	24	217.35	24	112.23	24	355.41	24	320.80	24
	10229.57	127	8852.25	97	5678.53	98	2965.02	109	4004.51	89	2112.31	125

method. In order to evaluate our method against the current state-of-the-art, we compared to Boolector [4], Yices [10], Cvc3, and MathSAT [5], the top solvers in the QF_AUFBV category from the 2009 SMT-COMP competition (in order). Additionally, we included the Z3 solver [8] so as to compare to the model-based theory combination method [7]. All tests were conducted on a dedicated Intel Pentium E2220 2.4 GHz processor with 4GB of memory. Individual runs were limited to 15 minutes.

We crafted a set of new benchmarks based on Example 4 from Section 5, taking $n = 10, \dots, 100$, with increments of 10, and $m = 32, \dots, 128$, with increments of 32. We also included a selection of problems from the QF_AUFBV division of the SMT-LIB library. Since most of the benchmarks in the library come from model-checking of software and use a flat memory model, they mostly operate over a single array representing the heap. Our method is essentially equivalent to the standard Nelson-Oppen approach for such benchmarks, so we selected only the benchmarks that involved constraints over at least two arrays. We anticipate that such problems will become increasingly important as static-analysis tools become more precise and are able to infer separation of the heap (in the style of Burstall, e.g. [16]). All the benchmarks and the Cvc3 binaries used in the experiments are available from the authors' website.⁷

The combined results of our experiments are presented in Table 1, with columns reporting the total time (in seconds) that a solver used on the problem instances it solved (not including time spent on problem instances it was unable to solve), and the number of solved instances. Compared to Cvc3, the new implementation Cvc3+C performs uniformly better. On the first four classes of problems, Cvc3+C greatly outperforms Cvc3. On the last three classes of problems, the difference is less significant. After examining the benchmarks, we concluded that the multitude of arrays in these examples is artificial – the many array variables are just used for temporary storage of sequential updates on the same starting array – so there is not a great capacity for improvement using the care function that we described. A scatter-plot comparison of Cvc3 vs Cvc3+C is shown in Figure 1(a). Because the only difference between the two implementations is the inclusion of the method described in this paper, this graph best illustrates the performance impact this optimization can have.

When compared to the other solvers, we find that whereas Cvc3 is not particularly competitive, Cvc3+C is very competitive and in fact, for several sets of benchmarks, performs better than all of the others. This again emphasizes the strength of our results and suggests that combination methods can be of great importance for performance and scalability of modern solvers. Overall, on this set of benchmarks, Boolector solves the most (solving 2 more than Cvc3+C). However, Cvc3+C is significantly faster on the benchmarks it solves. Figure 1(b) shows a scatter-plot comparison of Cvc3+C against Boolector.

7 Conclusion

We presented a reformulation of the classic Nelson-Oppen method for combining theories. The most notable novel feature of the new method is the ability to leverage the structure of the individual problems in order to reduce the complexity of finding a common arrangement over the interface variables. We do this by defining theory-specific care functions that determine the variable pairs that are relevant in a specific problem. We proved the method correct, and presented care functions for the theories of uninterpreted functions and arrays. The new method is asymmetric as only one of the theories takes charge of creating the arrangement graph over the interface variables. Since many theories we combine in practice are parametrized by other theories,

⁷ <http://cs.nyu.edu/~dejan/sharing-is-caring/>




Fig. 1. Comparison of Cvc3, Cvc3+C and Boolector. Both axes use a logarithmic scale and each point represents the time needed to solve an individual problem.

the non-symmetric approach has an intuitive appeal. We draw intuition for the care functions and correctness proofs directly from the decision procedures for specific theories, leaving room for new care functions backed by better decision algorithms. Another benefit of the presented method is that it is orthogonal to the previous research on combinations of theories. For example, it would be easy to combine our method with a model-based combination approach—instead of propagating all equalities between shared variables implied by the model, one could restrict propagation to only the equalities that correspond to edges in the care graph, gaining advantages from both methods.

We also presented an experimental evaluation of the method, comparing the new method to a standard Nelson-Oppen implementation and several state-of-the-art solvers. Compared to the other solvers on a selected set of benchmarks, the new method performs competitively, and shows a robust performance increase over the standard Nelson-Oppen implementation.

References

1. Clark Barrett, Robert Nieuwenhuis, Albert Oliveras, and Cesare Tinelli. Splitting on demand in SAT Modulo Theories. In *Logic for Programming, Artificial Intelligence, and Reasoning*, volume 4246 of *LNCS*, pages 512–526. Springer, 2006.
2. Clark Barrett and Cesare Tinelli. CVC3. In *Computer Aided Verification*, volume 4590 of *LNCS*, pages 298–302. Springer-Verlag, 2007.
3. Marco Bozzano, Roberto Bruttomesso, Alessandro Cimatti, Tommi Junttila, Silvio Ranise, Peter van Rossumd, and Roberto Sebastiani. Efficient theory combination via Boolean search. *Information and Computation*, 204(10):1493–1525, 2006.
4. Robert Brummayer and Armin Biere. Boolector: An efficient SMT solver for bit-vectors and arrays. In *Tools and Algorithms for the Construction and Analysis of Systems*, volume 5505 of *LNCS*, pages 174–177. Springer, 2009.
5. Roberto Bruttomesso, Alessandro Cimatti, Anders Franzén, Alberto Griggio, and Roberto Sebastiani. The MathSAT 4 SMT solver. In *Computer Aided Verification*, volume 5123 of *LNCS*, pages 299–303. Springer, 2008.
6. Roberto Bruttomesso, Alessandro Cimatti, Anders Franzen, Alberto Griggio, and Roberto Sebastiani. Delayed theory combination vs. Nelson-Oppen for satisfiability modulo theories: A comparative analysis. *Annals of Mathematics and Artificial Intelligence*, 55(1):63–99, 2009.
7. Leonardo de Moura and Nikolaj Bjørner. Model-based Theory Combination. In *5th International Workshop on Satisfiability Modulo Theories*, volume 198 of *Electronic Notes in Theoretical Computer Science*, pages 37–49. Elsevier, 2008.
8. Leonardo de Moura and Nikolaj Bjørner. Z3: An Efficient SMT Solver. In *Tools and Algorithms for the Construction and Analysis of Systems*, volume 4963 of *LNCS*, page 337. Springer, 2008.
9. Leonardo de Moura and Nikolaj Bjørner. Generalized, efficient array decision procedures. In *Formal Methods in Computer-Aided Design, 2009*, pages 45–52. IEEE, November 2009.

10. Bruno Dutertre and Leonardo de Moura. The YICES SMT Solver. *Tool paper at <http://yices.csl.sri.com/tool-paper.pdf>*, 2006.
11. Herbert B. Enderton. *A mathematical introduction to logic*. Academic press New York, 1972.
12. Dejan Jovanović and Clark Barrett. Polite theories revisited. Technical Report TR2010-922, Department of Computer Science, New York University, January 2010.
13. Dejan Jovanović and Clark Barrett. Polite theories revisited. In *Logic for Programming, Artificial Intelligence, and Reasoning*, volume 6397 of *LNCS*, pages 402–416. Springer Berlin / Heidelberg, 2010.
14. Greg Nelson and Derek C. Oppen. Simplification by cooperating decision procedures. *ACM Transactions on Programming Languages and Systems*, 1(2):245–257, October 1979.
15. Derek C. Oppen. Complexity, convexity and combinations of theories. *Theoretical Computer Science*, 12(3):291–302, 1980.
16. Zvonimir Rakamarić and Alan J. Hu. A Scalable Memory Model for Low-Level Code. In *Verification, Model Checking, and Abstract Interpretation*, *LNCS*, page 304. Springer-Verlag, 2009.
17. Silvio Ranise, Christophe Ringeissen, and Calogero G. Zarba. Combining Data Structures with Nonstably Infinite Theories Using Many-Sorted Logic. In *Frontiers of Combining Systems*, volume 3717 of *LNCS*, pages 48–64. Springer, 2005.
18. Robert E. Shostak. An algorithm for reasoning about equality. In *5th international joint conference on Artificial intelligence*, pages 526–527. Morgan Kaufmann Publishers Inc., 1977.
19. Cesare Tinelli and Mehdi T. Harandi. A new correctness proof of the Nelson–Oppen combination procedure. In *Frontiers of Combining Systems*, *Applied Logic*, pages 103–120. Kluwer Academic Publishers, March 1996.
20. Cesare Tinelli and Calogero Zarba. Combining decision procedures for sorted theories. In *Logic in Artificial Intelligence*, volume 3229 of *LNAI*, pages 641–653. Springer, 2004.

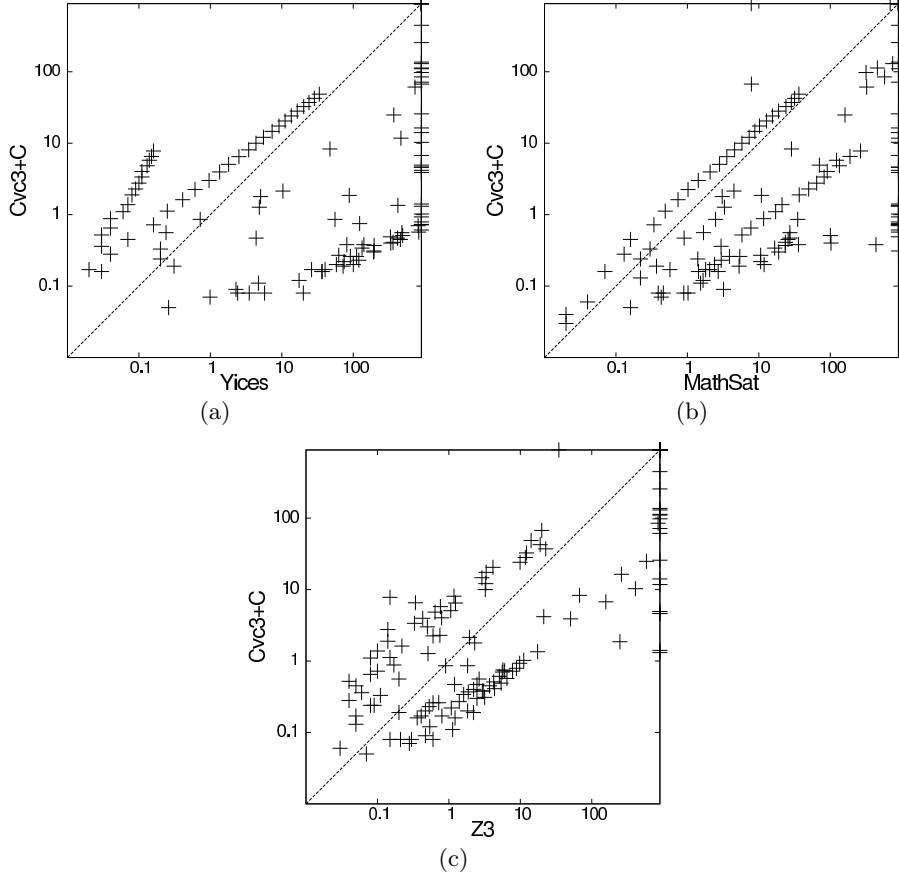


Fig. 2. Additional scatter-plot comparisons. Both axes are in logarithmic scale and each point represents the times needed to solve an individual problem.

A Appendix

A.1 Proof of Correctness for Polite theories

Theorem 2. Let T_i be a Σ_i -theory polite with respect to the set of sorts S_i , and equipped with equality propagators $\mathfrak{P}_{T_i}^{\equiv}[\cdot]$, for $i = 1, 2$. Additionally, let T_2 be equipped with a care function $\mathfrak{C}_{T_2}[\cdot]$ operating with respect to $\mathfrak{P}_{T_2}^{\equiv}[\cdot]$. Let $\Sigma = \Sigma_1 \cup \Sigma_2$, $T = T_1 \oplus T_2$ and let ϕ be a set of flat Σ -literals, which can be partitioned into a set ϕ_1 of Σ_1 -literals and a set ϕ_2 of Σ_2 -literals. Let $\phi'_2 = \text{witness}_{T_2}(\phi_2)$ and $V = \text{vars}_S(\phi'_2)$. If $S \subseteq S_1 \cap S_2$, then following are equivalent

1. ϕ is T -satisfiable;
2. there exists a care-graph \mathbf{G}_2 and arrangement $\delta_{\mathbf{G}_2}$, fix-point solutions of (1), such that the following sets are T_1 - and T_2 -satisfiable respectively

$$\phi_1 \cup \mathfrak{P}_T^{\equiv}[V](\phi_1 \cup \phi'_2 \cup \delta_{\mathbf{G}_2}) , \quad \phi'_2 \cup \mathfrak{P}_T^{\equiv}[V](\phi_1 \cup \phi'_2 \cup \delta_{\mathbf{G}_2}) .$$

Moreover, T is polite with respect to $S_1 \cup (S_2 \setminus S_1)$.

Proof. The proof is identical to the one given in Theorem 1 for the case of stably-infinite theories, except that in the last step, instead of relying on the correctness of the standard Nelson-Oppen method, we rely on the correctness of the method for combination of polite theories as described in [12, 13, 17].

A.2 Proof of Correctness for $\mathfrak{C}_{\text{euf}}[\cdot]$

We prove correctness of the care function for $\mathfrak{C}_{\text{euf}}[\cdot]$ by relying on the following well-known proposition.

Proposition 1. *A set ϕ of flat literals is T_{euf} -satisfiable iff for each dis-equality $x \neq y \in \phi$, $(x, y) \notin \mathcal{E}(\phi)^*$.*

Theorem 3. *Let T_{euf} be the theory of uninterpreted functions with equality over the signature Σ_{euf} . $\mathfrak{C}_{\text{euf}}[\cdot]$ is a care function for T_{euf} with respect to the equality propagator $\mathfrak{P}_{\text{euf}}^{\leftarrow}[\cdot]$.*

Proof. Let ϕ be a satisfiable set of flat literals, V a set of variables, let $\mathbf{G} = \mathfrak{C}_{\text{euf}}[\mathbb{V}](\phi) = \langle V, \emptyset \rangle$. Suppose we are given an arrangement δ_V (corresponding to equivalence relation E_V) with $\mathfrak{P}_{\text{euf}}^{\leftarrow}[\mathbb{V}](\phi) \subseteq \delta_V$. We must show that $\phi \wedge \delta_V$ is T_{euf} -satisfiable. Let $E_1 = \mathcal{E}(\phi)^*$, which corresponds to \sim_c by definition; we know by Proposition 1, above, that

$$x \neq y \in \phi \implies (x, y) \notin E_1 .$$

Now, let $E_2 = \mathcal{E}(\phi \cup \delta_V)^*$. In order to prove the theorem it suffices to show that

$$x \neq y \in \phi \cup \delta_V \implies (x, y) \notin E_2 .$$

We start by showing that

$$E_2 = \mathcal{E}(E_1 \cup E_V) . \quad (3)$$

To see this, note first that by basic properties of equivalence and congruence closures we have that

$$E_2 = \mathcal{E}(\phi \cup \delta_V)^* = (\mathcal{E}(\phi)^* \cup \mathcal{E}(\delta_V))^* = \mathcal{E}(E_1 \cup E_V)^* .$$

To see that $\mathcal{E}(E_1 \cup E_V)^* = \mathcal{E}(E_1 \cup E_V)$, suppose that this is not the case. Then there must be some pair of function applications $f(x_1, x_2, \dots, x_n)$ and $f(y_1, y_2, \dots, y_n)$, appearing in ϕ , such that

$$f(x_1, x_2, \dots, x_n) \sim_c f(y_1, y_2, \dots, y_n) ,$$

but for each $1 \leq i \leq n$, $(x_i, y_i) \in \mathcal{E}(E_1 \cup E_V)$. Since E_1 is congruence closed, there must be some i such that $(x_i, y_i) \notin E_1$ and $(x_i, y_i) \in \mathcal{E}(E_1 \cup E_V)$.

But then we must have a chain of equalities from x_i to y_i , such that at least one equality comes from E_V , and hence this chain must contain at least two variables from V . Let's pick x'_i , the first variable from V , and y'_i , the last variable from V , in this equality chain.

Since the chains from x_i to x'_i , and y_i to y'_i do not contain any variables from V , all these equalities must come from E_1 so it must be that $x_i \sim_c x'_i$ and $y_i \sim_c y'_i$. We can conclude that $x'_i \sim_c y'_i$ since otherwise we could deduce that $x_i \sim_c y_i$, and we would have the pair $(x_i, y_i) \in E_1$. We additionally must have $\neg x'_i \neq_c y'_i$, as otherwise this dis-equality would be in the propagated set, and we chose δ_V to be compatible with it. Finally, notice that for each $1 \leq i \leq n$, we have that $(x_i, y_i) \in \mathcal{E}(E_1 \cup E_V) \subseteq E_V^{\leftarrow}$.

But, now we can see that x'_i and y'_i satisfy all the requirements of the care function, so the edge (x'_i, y'_i) then must be in the graph \mathbf{G} , contradicting the fact that it should be empty, which then establishes (3).

Next, it is clear that, since δ_V is compatible with the propagated equalities from ϕ , it must be that

$$v, w \in V \text{ and } (v, w) \in E_1 \implies (v, w) \in E_V . \quad (4)$$

We can now show that

$$v, w \in V \text{ and } (v, w) \in E_2 \implies (v, w) \in E_V . \quad (5)$$

Suppose $v, w \in V$ and $(v, w) \in E_2$ but $(v, w) \notin E_V$. We know $(v, w) \notin E_1$ by (4). By (3), the only other possibility is that there is some transitive chain from v to w using pairs from both E_1 and E_V . Let $(t_0, t_1), (t_1, t_2), \dots, (t_{n-1}, t_n)$ be the smallest such chain (with $v = t_0$ and $w = t_n$). Let (t_i, t_{i+1}) be the first pair such that $t_i \in V$ but $t_{i+1} \notin V$ (there must be such a pair since, by (5), every pair in $E_1 \cap (V \times V)$ is also in E_V so that if $t_k \in V$ for all k , we would have $(t_0, t_n) \in E_V$). Then, let (t_j, t_{j+1}) be first pair such that $j > i$ and $t_j \notin V$ and $t_{j+1} \in V$ (there must be such a pair since $t_n \in V$). Notice that every pair from (t_i, t_{i+1}) to (t_j, t_{j+1}) must be in E_1 since each contains a term not in V . But then $(t_i, t_{j+1}) \in E_1$ which contradicts the assumption that $(t_0, t_1), (t_1, t_2), \dots, (t_{n-1}, t_n)$ is the smallest chain from t_0 to t_n . This establishes (5).

Finally, we return to the main proof and show that if $x \neq y \in \phi \cup \delta_V$, then $(x, y) \notin E_2$. We consider two cases.

- First, suppose $x \neq y \in \delta_V$ (and thus $x, y \in V$). Clearly, we cannot also have $x = y \in \delta_V$, so $(x, y) \notin E_V$. It follows by (5) that $(x, y) \notin E_2$.
- On the other hand, suppose that $x \neq y \in \phi$. We know that $(x, y) \notin E_1$, because $\phi \cup \delta_{\mathbf{G}}$ is T_{euf} -satisfiable. Thus, by (3), if $(x, y) \in E_2$, there must be some transitive chain from x to y using at least one pair from E_V . Let $(t_0, t_1), (t_1, t_2), \dots, (t_{n-1}, t_n)$ be such a chain (with $x = t_0$ and $y = t_n$). Let i be the first index such that $t_i \in V$, and let j be the last index such that $t_j \in V$. We know that $i < j$ since at least one pair must be from E_V . Clearly, $(t_i, t_j) \in E_2$, but since both $t_i \in V$ and $t_j \in V$, we also know by (5) that $(t_i, t_j) \in E_V$. Now, notice that (t_0, t_i) and (t_j, t_n) are both in E_1 since E_V only contains pairs that are both in V . It follows by the definition of the propagator that $t_i \neq t_j \in \mathfrak{P}_{\text{euf}}^{\neq} \llbracket V \rrbracket (\phi \wedge \delta_{\mathbf{G}})$. But $\delta_V \supseteq \mathfrak{P}_{\text{euf}}^{\neq} \llbracket V \rrbracket (\phi \wedge \delta_{\mathbf{G}})$, so we must have $t_i \neq t_j \in \delta_V$. But this is impossible since $(t_i, t_j) \in E_V$.

A.3 Proof of Correctness for the Decision Procedure for T_{arr}

While the procedure we described in Section 5 is based on the one described in [9], it differs enough that for the reader's convenience, we provide a self-contained proof of correctness for our procedure here.

Theorem 4. *The inference rules of \mathcal{D}_{arr} are sound and terminating.*

Proof. By sound, we mean that for each rule, the set of literals in the premise is T_{arr} -satisfiable iff one of the conclusion sets is T_{arr} -satisfiable. It is not hard to see that the soundness of the RIntro1 rule follows from axiom 1 of T_{arr} , the soundness of RIntro2 follows from axiom 2, and the soundness of ArrDiseq from axiom 3.

To see that the rules are terminating, first notice that applying a rule results in a new set Γ which no longer satisfies the side conditions of the rule just applied, so every applicaiton of a rule along a derivation branch must involve different “trigger” formulas (the ones in square brackets). Now, no rule introduces array disequalities, so it is clear that ArrDiseq can only be applied a finite number of times. Similarly, no rule introduces a new literal containing `write`, so RIntro1 can only be applied a finite number of times. Now, suppose we have a set Γ in which both the RIntro1 rule and the ArrDiseq rule no longer apply, and consider rule RIntro2 which may introduce new `read` terms. The rule cannot, however, introduce new `array` or `index` variables, so there are only a finite number of `read` terms that can be generated. Each application of RIntro2 merges the equivalence classes of either two `index` terms or two `read` terms. Since there are a finite number of both, eventually, no more merges will be possible.

Theorem 5. *Let Γ be a \mathcal{D}_{arr} -saturated set of flat T_{arr} -literals. Then Γ is T_{arr} -satisfiable iff $\alpha(\Gamma)$ is T_{euf} -satisfiable.*

Proof. Since T_{euf} includes all structures and $\alpha(\Gamma) \subseteq \Gamma$, the only-if direction is trivial. For the other direction, suppose Γ is a \mathcal{D}_{arr} -saturated set of flat T_{arr} -literals and denote with \approx_a^Γ the equivalence relation (from Section 5) on terms in Γ . Let \mathcal{A} be a maximally diverse T_{euf} model of $\alpha(\Gamma)$ (for instance, the \approx_a^Γ -quotient of the term model) and note that it has the property that for any two terms s and t of the same sort, $s \approx_a^\Gamma t$ iff $s^{\mathcal{A}} = t^{\mathcal{A}}$. We must show that Γ is T_{arr} -satisfiable. We will construct a T_{arr} interpretation \mathcal{B} that satisfies Γ . We define the domains of \mathcal{B} as:

$$\begin{aligned} B_{\text{index}} &= A_{\text{index}} \text{ ,} \\ B_{\text{elem}} &= A_{\text{elem}} \text{ ,} \\ B_{\text{array}} &= \{ f \mid f : B_{\text{index}} \mapsto B_{\text{elem}} \} \text{ .} \end{aligned}$$

We further define:

$$\begin{aligned} \text{read}^{\mathcal{B}} &= \lambda a : B_{\text{array}}. \lambda i : B_{\text{index}}. a(i) \text{ ,} \\ \text{write}^{\mathcal{B}} &= \lambda a : B_{\text{array}}. \lambda i : B_{\text{index}}. \lambda x : B_{\text{elem}}. (\lambda j : B_{\text{index}}. \text{if } i = j \text{ then } x \text{ else } a(j)) \text{ ,} \\ i^{\mathcal{B}} &= i^{\mathcal{A}} \text{ ,} \\ x^{\mathcal{B}} &= x^{\mathcal{A}} \text{ .} \end{aligned}$$

We interpret each array $a \in \text{vars}_{\text{array}}(\Gamma)$ as the corresponding function from \mathcal{A} in a restricted manner. Let e_0 be some distinguished element of B_{elem} . Then

$$a^{\mathcal{B}} = \lambda e : B_{\text{index}}. \begin{cases} x^{\mathcal{A}} & \text{if } x = \text{read}(b, i) \in \Gamma \text{ and } a \approx_a^{\Gamma} b \text{ and } i^{\mathcal{A}} = e \\ e_0 & \text{otherwise.} \end{cases}$$

To see that this definition is well-defined, suppose that for some variable a , we have both $x = \text{read}(b, i) \in \Gamma$ and $y = \text{read}(c, j) \in \Gamma$, with $a \approx_a^{\Gamma} b \approx_a^{\Gamma} c$ and $i^{\mathcal{A}} = j^{\mathcal{A}} = e$. Clearly, $b^{\mathcal{A}} = c^{\mathcal{A}}$, but then it must be the case that $\text{read}(b, i)^{\mathcal{A}} = \text{read}(c, j)^{\mathcal{A}}$ and so $x^{\mathcal{A}} = y^{\mathcal{A}}$.

It is easy to see that the definitions of `read` and `write` satisfy the axioms of T_{arr} . Now, we proceed to show that $\mathcal{B} \models \Gamma$. First, note that by definition, equalities and disequalities between variables of sort `index` or `elem` are trivially satisfied. Next, consider an equality of the form $x = \text{read}(a, i)$. Since this equality is in Γ , we know by the definition of $a^{\mathcal{B}}$ that $a^{\mathcal{B}}(i^{\mathcal{B}}) = x^{\mathcal{B}}$, so by the definition of `read`, such equalities must be satisfied. This shows that for terms t of sort `index` or `elem`, $t^{\mathcal{B}} = t^{\mathcal{A}}$ and thus if s is a term of the same sort as t , $s \approx_a^{\Gamma} t$ iff $s^{\mathcal{B}} = t^{\mathcal{B}}$. Similarly, it is not hard to see that since $\alpha(\Gamma)$ is satisfiable, we must have that if $s \not\approx_a^{\Gamma} t$ then $s^{\mathcal{B}} \neq t^{\mathcal{B}}$.

Next, consider equalities and disequalities between `array`-variables. For every disequality $a \neq b \in \Gamma$, we know that (because Γ is saturated), $\text{read}(a, k_{a,b}) \neq_a^{\Gamma} \text{read}(b, k_{a,b})$, and thus $\text{read}(a, k_{a,b})^{\mathcal{B}} \neq \text{read}(b, k_{a,b})^{\mathcal{B}}$, from which it is clear that $a^{\mathcal{B}} \neq b^{\mathcal{B}}$. To see that equalities $a = b$ are satisfied, note that we have $a \approx_a^{\Gamma} b$, and thus the definitions of $a^{\mathcal{B}}$ and $b^{\mathcal{B}}$ will yield the same function.

Finally, consider an equality of the form $a = \text{write}(b, i, v)$. Let $f_a = a^{\mathcal{B}}$ and $f_{\text{write}} = \text{write}(b, i, v)^{\mathcal{B}}$. We will show that for all `index`-elements ι , $f_a(\iota) = f_{\text{write}}(\iota)$. First, suppose that $\iota = i^{\mathcal{B}}$. In this case, it is clear that $f_{\text{write}}(\iota) = v^{\mathcal{B}}$ by the definition of `write` ^{\mathcal{B}} . Also, by the `RIntro1` rule (and saturation of Γ), we know that $v \approx_a^{\Gamma} \text{read}(a, i)$ and so $\text{read}(a, i)^{\mathcal{B}} = v^{\mathcal{B}}$. Then, by the definition of $a^{\mathcal{B}}$, we must have $f_a(\iota) = v^{\mathcal{B}}$.

Suppose, on the other hand that $\iota \neq i^{\mathcal{B}}$. Note that by the definition of `write` ^{\mathcal{B}} , this implies that $f_{\text{write}}(\iota) = b^{\mathcal{B}}(\iota)$. Suppose now that we have $x = \text{read}(c, j) \in \Gamma$ with $a \approx_a^{\Gamma} c$ and $j^{\mathcal{B}} = \iota$. In this case, the definition of $a^{\mathcal{B}}$ ensures that $f_a(\iota) = \text{read}(c, j)^{\mathcal{B}}$. Looking at rule `RIntro2`, we can see that because Γ is saturated and $i^{\mathcal{B}} \neq j^{\mathcal{B}}$, we must have $\text{read}(a, j)^{\mathcal{B}} = \text{read}(b, j)^{\mathcal{B}}$. But the first is equal to $\text{read}(c, j)^{\mathcal{B}}$ by saturation and rule `RIntro1`, and the second is equal to $b^{\mathcal{B}}(\iota)$ by the definition of `read` ^{\mathcal{B}} . Thus, $f_{\text{write}}(\iota) = f_a(\iota)$. A similar case is when $x = \text{read}(c, j) \in \Gamma$ with $b \approx_a^{\Gamma} c$ and $j^{\mathcal{B}} = \iota$. Here, we have $f_{\text{write}}(\iota) = \text{read}(c, j)^{\mathcal{B}}$ by definition, and we can again conclude that $\text{read}(a, j)^{\mathcal{B}} = \text{read}(b, j)^{\mathcal{B}}$ by rule `RIntro2`. But the first is equal to $f_a(\iota)$ by the definition of `read` and the second is equal to $\text{read}(c, j)^{\mathcal{B}}$ and thus to $f_{\text{write}}(\iota)$. In the final case, when neither of the previous cases hold, the definitions of $a^{\mathcal{B}}$ and $b^{\mathcal{B}}$ ensure that $f_a(\iota) = a^{\mathcal{B}}(\iota) = e_0 = b^{\mathcal{B}}(\iota) = f_{\text{write}}(\iota)$.

Since \mathcal{B} satisfies the axioms and each of the literals in Γ , this shows that Γ is T_{arr} -satisfiable.

A.4 Proof of Correctness for $\mathfrak{C}_{\text{arr}}[\cdot]$

In the proofs below, we make use of the definitions from Section 5. We start with a few simple lemmas. The first lemma states that both \approx_a^{ϕ} and \neq_a^{ϕ} are monotonic.

Lemma 1. *Suppose ϕ and Γ are sets of flat T_{arr} -literals with $\phi \subseteq \Gamma$. Then:*

- $s \approx_a^{\phi} t \rightarrow s \approx_a^{\Gamma} t$, and
- $s \neq_a^{\phi} t \rightarrow s \neq_a^{\Gamma} t$.

Proof. This is a straightforward consequence of the fact that adding additional information can only increase the set of consequences of a set of formulas.

The next lemma shows that if some derivation path from ϕ could introduce $\text{read}(a, i)$, then $\text{read}(a, i) \in \mathbb{R}^{\phi}$.

Lemma 2. *Let ϕ be a set of flat T_{arr} -literals, and \mathbb{R}^{ϕ} the set defined as in Section 5, and suppose that $\mathfrak{C}_{\text{arr}}[\mathbb{V}](\phi) = \langle V, \emptyset \rangle$. If Γ is a satisfiable set of literals obtained from ϕ via a sequence of \mathcal{D}_{arr} -inferences, then if $\text{read}(a, i)$ appears in Γ , $\text{read}(a, i) \in \mathbb{R}^{\phi}$.*

Proof. The proof is by induction on inference rule applications. For the base case, suppose $\Gamma = \phi$. The first rule defining \mathbf{R}^ϕ ensures that $\text{read}(a, i) \in \mathbf{R}^\phi$.

For the inductive case, suppose every term $\text{read}(a, i)$ appearing in Γ is in \mathbf{R}^ϕ and let Γ' be obtained by applying an inference rule to Γ . Suppose the inference rule is **RIntro1**. This introduces a term of the form $\text{read}(a, i)$. It also requires that we have an equality $a = \text{write}(b, i, v) \in \Gamma$. But no rule introduces such equalities, so it must have been in ϕ originally. The second rule defining \mathbf{R}^ϕ then ensures that $\text{read}(a, i) \in \mathbf{R}^\phi$.

Next, suppose the inference rule is **RIntro2**. The right branch of this rule may introduce $\text{read}(a, j)$ and $\text{read}(b, j)$. In this case, we know there are equalities $a = \text{write}(b, i, v)$ and $x = \text{read}(c, j)$ in Γ with $i \not\approx_a^\Gamma j$, and either $a \approx_a^\Gamma c$ or $b \approx_a^\Gamma c$. As before, we must have $a = \text{write}(b, i, v) \in \phi$, and by the inductive hypothesis, we know that $\text{read}(c, j) \in \mathbf{R}^\phi$. Furthermore, because $E_a^\phi = \emptyset$, we know that all relationships between **array** variables are already determined by ϕ , so either $a \approx_a^\phi c$ or $b \approx_a^\phi c$; and we know from Lemma 1 that $i \not\approx_a^\phi j$. We can then see that the third rule defining \mathbf{R}^ϕ ensures that $\text{read}(a, j)$ and $\text{read}(b, j)$ are in \mathbf{R}^ϕ .

Finally, suppose the inference rule is **ArrDiseq**. This rule may introduce $\text{read}(a, k_{a,b})$ and $\text{read}(b, k_{a,b})$. This can only happen if $a \neq b \in \Gamma$. Since no rules introduce disequalities between **array** variables, this implies that $a \neq b \in \phi$, and so the last rule defining \mathbf{R}^ϕ ensures that $\text{read}(a, k_{a,b}) \in \mathbf{R}^\phi$ and $\text{read}(b, k_{a,b}) \in \mathbf{R}^\phi$.

Theorem 6. *Let T_{arr} be the theory of arrays. $\mathbf{C}_{\text{arr}}[\cdot]$ is a care function for T_{arr} with respect to the equality propagator $\mathfrak{P}_{\text{arr}}^\pm[\cdot]$ for all sets ϕ of literals and V of variables such that $E_i^\phi \subseteq V \times V$.*

Proof. Assume that we are given a set ϕ of flat Σ_{arr} -literals and a set V of variables with $E_i^\phi \subseteq V \times V$. Let $\mathbf{C}_{\text{arr}}[V](\phi) = \langle V, \emptyset \rangle$, and assume that ϕ is T_{arr} -satisfiable and δ_V is a variable arrangement such that $\delta_V \supseteq \mathfrak{P}_{\text{arr}}^\pm[V](\phi)$. Because ϕ is T_{arr} -satisfiable, there must exist some set $\Gamma \supseteq \phi$ such that:

- Γ is derivable from ϕ using the rules of \mathcal{D}_{arr} .
- Γ is \mathcal{D}_{arr} -saturated,
- $\alpha(\Gamma)$ is T_{euf} -satisfiable.

Let \sim_c and \neq_c be the relations defined in Section 4, but with respect to $\alpha(\Gamma)$. Notice that by definition we have $s \sim_c t$ iff $s \approx_a^\Gamma t$ and $s \neq_c t$ iff $s \not\approx_a^\Gamma t$.

We claim that $\delta_V \supseteq \mathfrak{P}_{\text{euf}}^\pm[V](\alpha(\Gamma))$. We know that $\delta_V \supseteq \mathfrak{P}_{\text{arr}}^\pm[V](\phi)$, so it suffices to show that $\mathfrak{P}_{\text{euf}}^\pm[V](\alpha(\Gamma)) = \mathfrak{P}_{\text{arr}}^\pm[V](\phi)$. Notice that, if Γ' is the $\mathcal{D}'_{\text{arr}}$ -saturated set obtained starting from ϕ (the set is unique because $\mathcal{D}'_{\text{arr}}$ is completely deterministic), then by matching up the definitions, it is clear that $\mathfrak{P}_{\text{arr}}^\pm[V](\phi) = \mathfrak{P}_{\text{euf}}^\pm[V](\alpha(\Gamma'))$. Thus, it suffices to show that $\mathfrak{P}_{\text{euf}}^\pm[V](\alpha(\Gamma)) = \mathfrak{P}_{\text{euf}}^\pm[V](\alpha(\Gamma'))$. In fact, we can show that $\Gamma = \Gamma'$. Suppose not. The only way this could happen is if there is some \mathcal{D}_{arr} -derivation starting from ϕ in which rule **RIntro2** applies but rules **RIntro2a** and **RIntro2b** do not. Let Γ'' be the first set in the \mathcal{D}_{arr} -derivation from $\phi \cup \delta_{\mathbf{G}}$ to Γ in which this is the case. In order for the rule to be enabled, we must have $a = \text{write}(b, i, v), x = \text{read}(c, j) \in \Gamma''$. As we have noted before, derivations do not introduce equalities containing applications of **write**, so we must have $a = \text{write}(b, i, v) \in \phi$. We also know by Lemma 2 that $\text{read}(c, j) \in \mathbf{R}^\phi$. We also have $a \approx_a^{\Gamma''} c$ or $b \approx_a^{\Gamma''} c$, so it follows from the fact that $E_a^\phi = \emptyset$ that $a \approx_a^\phi c$ or $b \approx_a^\phi c$. But now, since $E_i^\phi = \emptyset$, clearly we must have $i \approx_a^\phi j$ or $i \neq_a^\phi j$. In the first case, we know that $i \approx_a^{\Gamma''} j$, so rule **RIntro2** is not applicable, contradicting our assumption. In the second case, we know that $i \neq_a^{\Gamma''} j$ which means that **RIntro2b** is applicable, which also contradicts our assumption.

Now, let $\mathbf{G}' = \langle V, E' \rangle = \mathbf{C}_{\text{euf}}[V](\alpha(\Gamma))$ be the T_{euf} care graph based on $\alpha(\Gamma)$ (with **read** treated as an uninterpreted function). We claim that $E' = \emptyset$. First note that because $E_a^\phi = \emptyset$, we know that for variables $a, b \in \text{vars}_{\text{array}}(V)$, either $a \approx_a^\phi b$ or $a \neq_a^\phi b$ (and thus $a \sim_c b$ or $a \neq_c b$), so it is impossible to have $(a, b) \in E'$. Next, notice that since variables of sort **elem** cannot appear as arguments to functions in $\alpha(\Gamma)$, there are no pairs $(x, y) \in E'$. Finally, suppose we have a pair of variables (i, j) of sort **index** such that $(i, j) \in E'$. Let $\delta_{\bar{V}}$ be the arrangement over V that sets all variables of the same sort whose relationship is not yet determined in ϕ to be equal, and let $E^\pm = \mathcal{E}(\alpha(\Gamma) \cup \delta_{\bar{V}})^*$. By the definition of $\mathbf{C}_{\text{euf}}[\cdot]$, we know that there exist a, b, i', j' such that

1. $\text{read}(a, i')$ and $\text{read}(b, j')$ appear in $\alpha(\Gamma)$,
2. $\text{read}(a, i') \not\approx_a^\Gamma \text{read}(b, j')$,
3. $(a, b) \in E^\pm$,
4. $\neg(a \neq_a^\Gamma b)$, and
5. $i \approx_a^\Gamma i'$, and $j \approx_a^\Gamma j'$,

We can immediately conclude from (1) that $\text{read}(a, i') \in R^\phi$ and $\text{read}(b, j') \in R^\phi$ by Lemma 2. Also, (2) implies that $\text{read}(a, i') \not\approx_a^\phi \text{read}(b, j')$ by Lemma 1.

We next consider the implications of (5). Notice that the only equalities between index variables that could have been introduced during the derivation from ϕ to Γ are those introduced by rule **RIntro2**. But, as we argued above, if this rule is enabled and could introduce $i = j$, then $(i, j) \in E_i^\phi$. But we know that $E_i^\phi = \emptyset$. It follows that no equalities between index variables are introduced in the derivation. So, if $i \approx_a^\Gamma i'$, then $i \approx_a^\phi i'$. The only way this can hold is if there is some chain of equalities from ϕ linking i to i' . Let i'' be the first variable in the chain that is also in V . Clearly, $i \approx_a^\phi i''$. Similarly, there is a j'' on a chain from j to j' such that $j \approx_a^\phi j''$. But now notice that if we can establish $a \approx_a^\phi b$ (see below), it will follow from the definition of E_r^ϕ (and the fact that $E_r^\phi = \emptyset$) that either $i'' \approx_a^\phi j''$ or $i'' \neq_a^\phi j''$. But $i'' \approx_a^\Gamma i$ and $j'' \approx_a^\Gamma j$, so in the first case, we clearly have $i \approx_a^\Gamma j$ and thus $i \sim_c j$, and in the second case, we have $i \neq_a^\Gamma j$ and thus $i \neq_c j$. This contradicts our assumption that $(i, j) \in E'$. It remains to show $a \approx_a^\phi b$. By (4), we know that $\neg(a \neq_a^\Gamma b)$, and thus by Lemma 1, $\neg(a \neq_a^\phi b)$. But since $E_a^\phi = \emptyset$, we must then have $a \approx_a^\phi b$.

We have thus established that $E' = \emptyset$. Now, because $\delta_V \supseteq \mathfrak{P}_{\text{euf}}^{\leftarrow} \llbracket V \rrbracket (\alpha(\Gamma))$, by Theorem 3, $\alpha(\Gamma) \cup \delta_V$ must be T_{euf} -satisfiable. But $\alpha(\Gamma) \cup \delta_V = \alpha(\Gamma \cup \delta_V)$, so $\alpha(\Gamma \cup \delta_V)$ is T_{euf} -satisfiable. Finally, since Γ is \mathcal{D}_{arr} -saturated, and δ_V can only add new equalities and disequalities between variables of sort `index` or `elem`, it is clear that $\Gamma \cup \delta_V$ must also be \mathcal{D}_{arr} -saturated, so by Theorem 5, $\Gamma \cup \delta_V$ is T_{arr} -satisfiable, from which we can conclude that $\phi \cup \delta_V$ is T_{arr} -satisfiable.