

State Space *c*-Reductions of Concurrent Systems in Rewriting Logic

Alberto Lluch Lafuente¹, José Meseguer², and Andrea Vandin¹

¹ IMT Institute for Advanced Studies Lucca, Italy

² University of Illinois in Urbana-Champaign, USA

Abstract. We present *c-reductions*, a state space reduction technique. The rough idea is to exploit some equivalence relation on states (possibly capturing system regularities) that preserves behavioral properties, and explore the induced quotient system. This is done by means of a canonizer function, which maps each state into a (non necessarily unique) canonical representative of its equivalence class. The approach exploits the expressiveness of rewriting logic and its realization in Maude to enjoy several advantages over similar approaches: flexibility and simplicity in the definition of the reductions (supporting not only traditional symmetry reductions, but also name reuse and name abstraction); reasoning support for checking and proving correctness of the reductions; and automatization of the reduction infrastructure via Maude’s meta-programming features. The approach has been validated over a set of representative case studies, exhibiting comparable results with respect to other tools.

1 Introduction

State space reduction techniques have been extensively investigated since the birth of automated verification and have contributed to their success by enhancing the performance of tools and allowing for the analysis of larger and larger systems. The most prominent case is probably that of model checking [1], where techniques such as abstract interpretation, partial order reduction, and symmetry reduction allow to consider smaller though equivalent systems and thus save precious space and time resources.

Symmetry reduction, for instance, enjoys a vast literature [2] but, while most of the authors agree on the effectiveness of the technique, symmetry reduction has not established itself as standard feature of verification tools as is the case of other state space reduction techniques. Notable examples are model checkers such as SPIN [3], where symmetry reduction and abstract interpretation are not built-in while other techniques like partial order reduction are.

There are several reasons for this. (i) automatic detection of system regularities is a hard task, very often delegated to the system designer since there are only few solutions available (e.g. [4]); (ii) their exploitation is done by complementing (or enriching) the system description language with some meta-data in a different language, so that the user is hence required to use this new language; (iii) the

implementation of state space reduction techniques has to be combined (both theoretically and practically) with the rest of the techniques and algorithms implemented in the model checker, and often this integration effort has to be repeated for every new version, improvement or technique of the model checker; and (iv) checking correctness of the defined reductions is not easy and requires reasoning techniques (e.g. theorem proving) that are not integrated in the model checking framework, or that are not part of the user’s skills.

Contributions. This paper proposes an approach called *c-reductions* that substantially mitigates the above problems (i)–(iv). Its key idea is the reduction of each state into a (non necessarily unique) *canonical* representative of an equivalence class induced by a relation on states that preserves the system properties under study (i.e. a bisimulation). The approach is fully general: it subsumes many reduction techniques (symmetry reduction, name reuse and name abstraction) and can be applied to any Kripke structure. Since any computable Kripke structure can be formally specified by a finitely rewrite theory [5] we use the setting of rewriting logic [6], and its realization in Maude [7], in order to exploit its expressiveness and support of formal verification techniques like model checking and theorem proving.

The *c*-reduction approach tackles the above mentioned issues in the following way: (i) reductions are defined using ordinary ingredients of the system description language, namely, equations; (ii) the implementation of reductions does not interfere with the theory of rewriting logic or with the Maude engine and its commands, since it is essentially based on equational simplification, a standard feature of the setting; (iii) the approach is highly automatizable (thanks to Maude’s metaprogramming facilities based on logical reflection) and, at the same time, flexible enough to allow customization (e.g. by allowing the user to define the canonization functions or parts of them); (iv) correctness checks are semi-automatized by techniques well supported by the Maude formal environment, e.g. critical pair analysis to check that canonizer functions do not “interfere” with other functions and are “coherent” with respect to behavioural rules.

We have evaluated our approach over an ample set of examples by considering the ease of defining reduction strategies, the effectiveness of the correctness checks, and the performance of the resulting reductions. With respect to previous works we have observed performance gains in some cases (including previous implementations of symmetry reductions in Maude [8]), more flexibility in the definition of reductions, which allow us to subsume a wide range of reductions including permutation and rotation symmetries, name reuse and name abstraction.

Synopsis. § 2 offers the necessary background. § 3 presents the *c*-reduction technique in a generic way, focusing on Kripke structures. § 4 describes the realization of *c*-reductions in rewriting logic, highlighting the theoretical results, and the reasoning and verification mechanisms and tools, that can be exploited to analyze and develop and *c*-reductions. § 5 evaluates the performance of *c*-reductions through some representative examples. Finally, § 6 concludes the paper, describes subjects of current work and outlines future research directions.

2 Preliminaries

We focus on system specifications in the form of theories of rewriting logic [6], which can be realized in practice in terms of Maude modules [7].

Definition 1 (Rewrite theory). *A rewrite theory \mathcal{R} is a tuple $\mathcal{R} = (\Sigma, E \cup A, R, \phi)$ where Σ is a signature, specifying the basic syntax (function symbols) and type infrastructure (sorts, kinds and subsorting) for terms, i.e., state descriptions; E is a set of (possibly conditional) equations, which induce equivalence classes of terms, and (possibly conditional) membership predicates, which refine the typing information; A is a set of axioms which also induce equivalence classes of terms, i.e., equational axioms describing structural equivalences between terms, like associativity and commutativity; R is a set of (possibly conditional) non-equational rules, which specify the local concurrent transitions in a system whose states are $E \cup A$ -equivalence classes of ground Σ -terms; and where $\phi : \Sigma \rightarrow \mathcal{P}_{fin}(\mathbb{N})$ is a frozenness map, assigning to each function symbol f of arity n a subset $\phi(f) \subseteq \{1..n\}$ of its frozen argument positions, i.e. positions under which rewriting with rules in R is forbidden.*

For the sake of simplicity, we assume that the system under study is described by a rewrite theory $\mathcal{R} = (\Sigma, E \cup A, R, \phi)$ whose rules are “topmost” for a designated kind `[State]` of states. We also assume that an operator `{_}` is used to enclose states, so that all transitions in R have that operator as their top operator in their lefthand sides. What follows can be of course generalized but the former assumptions facilitate the presentation. Finally, we shall assume that \mathcal{R} has good executability properties, i.e., that E is sufficient complete, confluent and terminating modulo A (that is that the equational part correctly defines functions), and R is coherent with E modulo A [7] (that is that applying equations to evaluate functions does not interfere with the application of the rules that dictate the system transitions). All such properties can be checked using the standard Maude tools.

We consider a well-known semantical domain for rewrite theories, namely, Kripke structures, which are suitable to formulate state space exploration problems like model checking.

Definition 2 (Kripke structure). *A Kripke structure K is a tuple $K = (S, \rightarrow, L, AP)$ such that S is a denumerable set of states, $\rightarrow \subseteq S \times S$ is a transition relation between states, and $L : S \rightarrow 2^{AP}$ is a labelling function mapping states into sets of atomic propositions AP (i.e. observations on states).*

The Kripke semantics of a rewrite theory is defined as expected, with `State`-sorted terms as states and one-step rewrites between `State`-sorted terms as transitions. The labelling function is defined by Boolean predicates specified equationally in the rewrite theory. As proved in [5], any computable Kripke structure, even an infinite-state one, can be obtained this way from an executable rewrite theory using only a finite signature Σ , and finite sets E of equations, A of axioms and R of rules.

Definition 3 (Rewrite theory semantics). Let $\mathcal{R} = (\Sigma, E \cup A, R, \phi)$ be a rewrite theory with a designated state sort **State**, symbols and equations for a labelling function L , and atomic propositions AP . The Kripke structure associated to \mathcal{R} is $K_{\mathcal{R}} = (T_{\text{State}/E \cup A}, \rightarrow, L_{E/A}, AP)$ such that $T_{\text{State}/E \cup A}$ are all **State**-sorted states, \rightarrow is defined as $\{[t] \rightarrow [t'] \mid \mathcal{R} \vdash u \rightarrow_{R, \text{State}}^1 v\}$ (i.e. transitions are one-step rewrites between $E \cup A$ equivalence classes of **State**-terms in \mathcal{R}), and $L_{E/A}$ is the interpretation of L the initial algebra $\mathcal{T}_{\Sigma/E \cup A}$.

We will consider bisimulation as the key semantic equivalence.

Definition 4 (bisimulation). Let $K = (S, \rightarrow_K, L, AP)$, $H = (S, \rightarrow_H, L, AP)$ be two Kripke structures, and let $\sim \subseteq S \times S$ be an equivalence relation on S . We say that \sim is a bisimulation between K and H iff (i) $s \sim s'$ implies $L(s) = L(s')$; (ii) $s \rightarrow_H r$ and $s \sim s'$ implies that there is a r' s.t. $s' \rightarrow_K r'$ and $r \sim r'$; and (iii) $s \rightarrow_K r$ and $s \sim s'$ implies that there is an r' s.t. $s' \rightarrow_H r'$ and $r' \sim r$.

We recall here some basic notions of group theory, since we shall often instantiate our approach to the case of equivalence classes of (bisimilar) states which can be defined as the orbits of a group action, which yields the cases of symmetry reductions as special instances of our approach.

Definition 5 (groups, generators and actions). A group is a tuple (G, \circ) such that G is a set of elements and $\circ : G \times G \rightarrow G$ is a binary operation on G where \circ is associative; there is an identity element $e \in G$ such that $\forall f \in G. f \circ e = e \circ f = f$; and each element $f \in G$ has an inverse $f^{-1} \in G$ such that $f \circ f^{-1} = f^{-1} \circ f = e$.

Let (G, \circ) be a group and $H \subseteq G$ be a subset of G . The group generated by H denoted $\langle H, \circ \rangle$ is defined as the closure of H under the inverse and product operators $(-)^{-1}$ and \circ of G . In general $\langle H, \circ \rangle$ will be a subgroup of (G, \circ) , but if $\langle H, \circ \rangle$ coincides with (G, \circ) , then H is said to generate G and its elements are called generators.

Let (G, \circ) be a group and A be a set. A group action of G on A is a monoid homomorphism $[[\cdot]] : G \rightarrow (A \rightarrow A)$, that is, $[[f \circ g]] = [[f]] \circ [[g]]$, where $f; g$ denotes function composition in $(A \rightarrow A)$, and $[[e]] = id_A$, with id_A the identity function on A .

Notable examples are full symmetric and rotation groups, which capture typical symmetries introduced by replication (e.g. of processes with identical behavior) in concurrent systems. Generators define groups in a finite and concise manner. Well known examples of generators for full symmetric and rotation groups are transpositions and single rotations, respectively. Group actions (which allow us to apply group operations to concrete domains like state descriptions) can be defined by means of equations of the form $[[[f]]](\mathbf{t}) = \mathbf{t}'$ where f denotes a group element (possibly a generator) and \mathbf{t}, \mathbf{t}' are **State**-sorted terms. Since rewriting logic (and its realization in Maude) allows to handle the axioms of sets (associativity, commutativity, identity), a group action can be described very concisely. Such conciseness is crucial for our approach since most of the correctness checks rely on these equations: the fewer equations, the fewer proofs and checks to be carried out.

3 C-Reductions for Kripke Structures

This section introduces the idea of *canonical reductions*, abbreviated c-reductions as a generic means to reduce a Kripke structure by exploiting some equivalence relation \sim on states which is also a bisimulation. In the next section we will explain how c-reductions are realized and analyzed in rewriting logic, where Kripke structures are specified by rewrite theories.

We start defining canonizer functions, the means to compute (non necessarily) unique canonical representatives for given equivalence classes of states (e.g. a canonical permutation of the identifiers of processes with identical behavior).

Definition 6 (canonizer functions). *Let $K = (S, \rightarrow, L, AP)$ be a Kripke structure, and let \sim be an equivalence relation which is a bisimulation on K . A function $c : S \rightarrow S$ is a \sim -canonizer (resp. strong \sim -canonizer) iff for each $s \in S$ we have $s \sim c(s)$ (resp. $s \sim c(s)$, and $s \sim s' \rightarrow c(s) = c(s')$).*

Canonizer functions are used to compute smaller but semantically equivalent (i.e. bisimilar) Kripke structures by applying canonizers after each transition.

Definition 7 (c-reduction of a Kripke structure). *Let $K = (S, \rightarrow, L, AP)$ be a Kripke structure, and let $c : S \rightarrow S$ be a \sim -canonizer function for some equivalence relation \sim which is a bisimulation on K . We call the Kripke structure $K_c = (S, (\rightarrow; c), L, AP)$ the c-reduction of K , where the composed transition relation $\rightarrow; c$ is defined as $\{(s, c(r)) \in S^2 \mid s \rightarrow r\}$.*


We sometimes decompose transitions $s \rightarrow c(s')$ in K_c to make explicit the canonization step from a state s to its canonical representative $c(s)$ with $s \rightarrow s' \rightarrow_c c(s')$.

An important result is then that c-reductions are bisimulation preserving.

Theorem 1 (\sim -preservation). *Let $K = (S, \rightarrow, L, AP)$ be a Kripke structure, let \sim be a bisimulation on K , and let c be a \sim -canonizer function. Then \sim is a bisimulation relation between K and K_c .*

Proof. The relation \sim is obviously a simulation from K to K_c , i.e. it lets the latter simulate the former (c.f. beside figure (i)). Indeed if we have $s \rightarrow r$, we know that any $s' \sim s$ can simulate the transition $s \rightarrow r$ by a transition $s' \rightarrow r'$ since \sim is a bisimulation for K . But then we have $s \sim c(r')$ (since c preserves \sim , and \sim is symmetric and transitive). Therefore, \sim is a simulation relation from K to K_c .

It is also easy to show that \sim is a simulation from K_c to K as well (c.f. beside figure (ii)). Indeed, a transition $s \rightarrow r \rightarrow_c c(r)$, can be simulated by some transition $s' \rightarrow r'$ such that $s \sim$



$s' \sim r'$ and $r \sim r'$ since \sim is a bisimulation for K . But, as in the above case, we also have that $r \sim c(r')$ (since c preserves \sim and \sim is transitive). Therefore, \sim is a simulation relation from K_c to K .

Since we have already shown that \sim is a simulation from K to K_c , we conclude that \sim is a bisimulation between K and K_c . \square

4 C-Reductions in Rewriting Logic

We now describe a methodology for defining and analyzing c -reductions in rewriting logic. We shall often mention how such formalization can benefit from the tool support offered by the Maude framework such as Maude’s LTL Model Checker, Invariant Analyzer [9], Inductive Theorem Prover [10, 11] and Church Rosser and Coherence Checker [12].

We assume there is some regularity in \mathcal{R} (e.g. replicated process) that implicitly defines a bisimulation \sim on states (e.g. a permutation of processes) to be exploited to ease the analysis of \mathcal{R} . Our methodology is then based on the following steps: (i) implement the group action that induces \sim and verify that it is indeed a group action (§4.1) which ensures \sim to be an equivalence relation; (ii) verify that \sim preserves the state predicates (§4.2); (iii) verify that \sim is a bisimulation (§4.3); (iv) define a c -canonizer and show it to be a (possibly strong) \sim -canonizer (§4.4); (v) build the c -reduction of \mathcal{R} (§4.5); and (vi) evaluate its performance (§5).

We remark that the steps of this methodology proceed by tackling different levels of abstraction so that they act as building blocks to be re-used whenever needed. For instance verifying a c -reduction strategy does not require performing all the verification steps if it is based on a state equivalence that has been already proven to be correct. In practice, bisimulation relations and their canonizers need not to be defined and proven to be correct for every system, as there will be classes of systems for which they can be specified once and for all. In such cases, one can define reductions as theory transformations for wide classes of examples corresponding, for instance, to certain permutation groups, or to other useful equivalence relations besides the symmetry reduction case. In Maude this could be done by exploiting reflection so that the c -reduction is automatized as a function at the metalevel, possibly after checking some proof obligations. In addition, in some cases it might be more convenient to bypass some steps, for instance to check the correctness of a canonizer when it is not clear which is the underlying group action (which may not exist at all, since even if we focus here on group-theoretic reductions, c -reductions are more general).

This paper puts the emphasis on steps (iv) to (vi) which constitute our main contribution here; while steps (i) to (iii) enjoy less attention due to lack of space.

4.1 Checking group actions

In the case of equivalences yielded by group actions, one can start defining the group, the generators and their group actions and provide a formal proof of their properties. Many symmetries such as full and rotation symmetries have been identified and thoroughly studied in the past so we won’t pay much attention to this aspect in this work. In what follows we will sometimes instantiate our methodology on the particular case of symmetries assuming correctly defined groups, generators and actions.

We assume that the relation on states we have to analyze is equationally defined by topmost equations relating two **State**-sorted terms. In the case of group actions with generators, we can just give equations of the form $[[g]](\{t\})$

$= \{\mathfrak{t}'\}$, for \mathfrak{g} being a generator (plus the equations defining the inverse operation, unless inverses of generators can be obtained by composition of generators as in the case of transpositions and rotations). Such equations define the actual application of a generator.

Of course, we need to check that for each generator g and each state u we have $\llbracket g \circ g^{-1} \rrbracket(u) = u$, but this can be done easily because we can devise inductive proofs exploiting generators. For instance, in the case of full symmetries all we have to do is to prove the equality $\llbracket i \leftrightarrow j \rrbracket(\llbracket i \leftrightarrow j \rrbracket(\{t\})) = \{t\}$, i.e. that applying the same transposition $i \leftrightarrow j$ twice amounts to applying the identity so that rewrite steps are reversible.

If one has correctly specified a group action, and in particular checked (or given as explicit equations) the monoid homomorphism equations $\llbracket [\text{id}] \rrbracket(\{\mathfrak{t}\}) = \{\mathfrak{t}\}$ and $\llbracket [\mathfrak{g} \circ \mathfrak{g}'] \rrbracket(\{\mathfrak{t}\}) = \llbracket [\mathfrak{g}] \rrbracket(\llbracket [\mathfrak{g}'] \rrbracket(\{\mathfrak{t}\}))$ plus the above-mentioned check of the inverses for generators, one has already proved that \sim is an equivalence relation. Indeed a group action implicitly defines an equivalence relation between states as follows $\{t\} \sim \{t'\} \Leftrightarrow \exists f \in G \text{ s.t. } \llbracket f \rrbracket(\{t\}) = \{t'\}$.

4.2 Checking that \sim preserves atomic predicates

To ensure that \sim preserves the atomic predicates AP under consideration we can proceed as follows.

We first define a rewrite theory \mathcal{R}_\sim for the sole purpose of this analysis defined as $\mathcal{R}_\sim = (\Sigma \cup \Sigma_\sim, E \cup E_\sim \cup A, \{\{\mathfrak{t}\} \Rightarrow \llbracket [\mathfrak{g}] \rrbracket(\{\mathfrak{t}\})\} \mid \mathfrak{g} \text{ is a generator or the inverse of a generator}\}, \phi)$. In words, we substitute the rules of \mathcal{R} by rules that allows us to go from a state u to another state that v that results from applying a generator (or the inverse of a generator) to u .

As we mentioned above, it is easy to see that in this rewrite theory two states are reachable if and only if they are in the same orbit, i.e. that the relation \sim is defined by the rewrite relation $\rightarrow_{\mathcal{R}_\sim}^*$. That is, that $u \sim v$ exactly when the above defined rewrite theory can prove $u \rightarrow_{\mathcal{R}_\sim}^* v$.

It is easy to see that proving that a predicate p is preserved by \sim is exactly proving that p is invariant under \mathcal{R}_\sim . Such a proof can be greatly facilitated if p is defined by pattern matching, i.e. by topmost unconditional equations of the form $\mathfrak{p}(\{\mathfrak{t}\}) = \text{true}$.

More formally, if we let \mathcal{R}_\sim be the rewrite theory defined above (and we check that the lefthand sides of all its rules are free constructor terms modulo A), all we have to do is to prove that for each rule $\{\mathfrak{t}'\} \Rightarrow \{\mathfrak{t}''\} \in \mathcal{R}_\sim$, equation $\mathfrak{p}(\{\mathfrak{t}'\}) = \text{true} \in E$, and A -unifier θ between \mathfrak{t}' and \mathfrak{t} (i.e. a mapping θ of variables into non-necessarily ground terms such that $\theta(\mathfrak{t}') =_A \theta(\mathfrak{t})$), then we have $\mathfrak{p}(\{\theta(\mathfrak{t}'')\}) = \text{true}$. If such is the case, indeed, we can conclude that for each pair of states $u, v \in T_{\text{State}_{E/A}}$ it holds that $u \sim v$ implies $\mathfrak{p}(u) = \mathfrak{p}(v)$, i.e. that \mathfrak{p} is invariant under \sim .

Fortunately, Maude has a tool called the Invariant Analyzer [13, 9] (InvA) which can automate a good part of the effort of proving such invariants, leaving the remaining proof obligations for Maude's inductive theorem prover [11].

As a matter of fact, as part of our experimentation we were able to use the tool to check that a simple neighborhood property of a model of the dining philosophers problem (“no two consecutive philosophers eating”) is invariant under rotations but not under arbitrary permutations. The first result was obtained in fully automatic way, while the second follows from an analysis of the few proof obligations provided by InvA.

4.3 Checking that \sim is a bisimulation

Once \sim has been shown to preserve the atomic predicates of interest, all there is left to check is that \sim is a bisimulation. In our setting this amounts to checking the joinability of suitable “critical pairs” between the state transition rules $\{\mathfrak{t}\} \Rightarrow \{\mathfrak{t}'\}$ if `cond` in our original set R , and the rules $\{\mathfrak{t}'\} \Rightarrow \{\mathfrak{t}''\}$ of \mathcal{R}_\sim .

For A -unifiers θ between \mathfrak{t} and \mathfrak{t}'' bisimulation is ensured if we prove that for each critical pair denoted with ordinary arrows in the diagram on the right, there is a rule in R giving us a one-step rewrite $\{\theta(\mathfrak{t}'')\} \rightarrow_R \{w\}$ for which we can prove: $\{\theta(\mathfrak{t}')\} \rightarrow_{R_\sim}^* \{w\}$.

Part of this can be automatized by using Maude’s automated reachability analysis capabilities, possibly including narrowing-based ones since some of the above terms might not be ground (i.e. they might contain variables). Of course, it is not as easy as it sounds, since the above proofs are inductive, so that sometimes it may not be straightforward to show these properties for all ground terms just by rewriting terms with variables: some inductive arguments may also be needed.

$$\begin{array}{ccc} \{\theta(\mathfrak{t})\} & \xrightarrow{R} & \{\theta(\mathfrak{t}')\} \\ R_\sim \downarrow & & R_\sim \downarrow^* \\ \{\theta(\mathfrak{t}'')\} & \xrightarrow{R} & \{w\} \end{array}$$

4.4 Defining and analyzing canonizer functions

The next step is to define canonizers functions $c : [\text{State}] \rightarrow [\text{State}]$ in a protecting extension of the rewrite theory \mathcal{R} under study. Note that in order to define the function $c : [\text{State}] \rightarrow [\text{State}]$ we may need to define some auxiliary functions (for example, the ordering relations used in symmetry reductions to determine a representative for each orbit).

Definition 8 (c-extension of a rewrite theory). *Let $\mathcal{R} = (\Sigma, E \cup A, R, \phi)$ be a rewrite theory and c a canonizer function. The c -extension of \mathcal{R} is the rewrite theory $\mathcal{R}^c = (\Sigma \uplus \Sigma_c, E \cup G_c \cup A, R, \phi_c)$ where $c \in \Sigma_c$, and $\Sigma_c \setminus \{c\}$ are other auxiliary functions; $(\Sigma \uplus \Sigma_c, E \cup G_c \cup A)$ is a protecting extension of $(\Sigma, E \cup A)$, and is ground confluent, ground terminating, and sufficiently complete wrt. the same signature of constructors; and ϕ_c extends ϕ with frozenness maps for all arguments of all functions in Σ_c .*

For a given bisimulation \sim many candidate canonizers may exist, each leading to different results in terms of the size of the reduced state space and computational requirements.

In any case, all canonizer functions must enjoy correctness properties such as the extended equational theory being confluent, terminating, sufficiently complete

with respect to constructors (which may be the same constructors as those of $(\Sigma, E \cup A)$). More importantly (for what regards our contribution), canonizer functions must preserve \sim , i.e. they must be \sim -canonizers. This might require some theorem proving but it can be relatively easy in most cases, since we can use the equations G_c and show that each one preserves \sim .

For example, in the case of *local* reduction strategies [4] the equations G_c defining c are of the form $c(\{t\}) = c(\llbracket g \rrbracket(\{t\}))$ if $\llbracket g \rrbracket(\{t\}) < \{t\}$ for each possible generator (or inverse generator) g with $<$ defining an ordering relation on states, plus an equation $c(\{t\}) = \{t\}$ [owise] to deal with the case when none of the previous equations is applicable (denoted by the keyword [owise]), that is when there is no way to transform a state into a *smaller* equivalent one by applying a generator. Since such equations define c in terms of group actions or of the identity function when all conditions fail, preservation of \sim is immediate by the very definition of c . A very similar situation is that of *enumeration* strategies [4] where canonizers are defined as $c(\{t\}) = \min\{\llbracket f \rrbracket(\{t\})\}$ where f ranges over all elements of the group under consideration. Again, preservation of \sim by c follows from the very definition of c . Indeed, for all states u , $c(u)$ will be necessarily of the form $\llbracket g_1 \circ g_2 \circ \dots \circ g_n \rrbracket(u)$, with each g_i being a generator. All the c -reduction strategies evaluated in §5 are essentially of these forms.

Proposition 1. *Let \mathcal{R} be a rewrite theory, \mathcal{R}_\sim be defined as usual, and \mathcal{R}^c be the c extension of \mathcal{R} such that the equations of E_c defining c are of the above described form. Then, for all states $u \in T_{State/E \cup A}$ we have that $c(u) \sim u$.*

In practice, when implementing c -strategies in the above form, all we have to check are the ordinary well-definedness properties of the equations of c : termination and sufficient completeness with respect to constructors, for which one can rely on the standard Maude tools.

Note that proving ground confluence of c is not sufficient to show that c is a strong canonizer. It may still be the case that for some two states u, v such that $u \sim v$ we have that $c(u) \neq c(v)$. Indeed, what we need to prove to show that c is a strong canonizer is that for all generators g and states s we have $c(s) = c(\llbracket g \rrbracket(s))$. It is easy to see that if the previous property holds, an inductive argument allow us to conclude $c(s) = c(\llbracket f \rrbracket(s))$ for any possible group operation f and hence for any two equivalent states $s \sim s' = \llbracket f \rrbracket(s)$.

But the above check can be reduced to an ordinary (automatic) confluence check if we extend E_c with the set of equations $\mathbf{test}(g, c(\{t\})) = c(\llbracket g \rrbracket(\{t\}))$ and $\mathbf{test}(g, c(\{t\})) = c(\{t\})$, for all generators g (and inverse generators) of the group under consideration, where \mathbf{test} is an auxiliary function symbol that ensure that the generator is applied only once. We denote this extension of E_c with $E_{c\sim}$. In words, the idea is that applying or not a group generator and then canonizing should result in the very same state.

Proposition 2. *Let \mathcal{R} be a rewrite theory, \mathcal{R}_\sim be defined as usual, \mathcal{R}^c be the c extension of \mathcal{R} , and let $E_{c\sim}$ be defined as above. If $(\Sigma \cup \Sigma_c, E \cup E_{c\sim}, A)$ is confluent, then c is a strong canonizer.*

This result allows us to rely on Maude’s confluence checker. Of course, there are cases in which no check is needed. For instance, it is well-known that enumeration strategies yield strong canonizers, while local strategies are not strong in general.

4.5 Building c-reductions

The next step is to build the c-reduction. Recall that we have defined the c-reduction \mathcal{K}_c of a Kripke structure \mathcal{K} in Def. 7. Starting from a system specification given as a rewrite theory \mathcal{R} , our goal is to build $\mathcal{K}_c(\mathcal{R})$, but this can be easily done by applying a theory transformation mapping \mathcal{R} into its c-reduction \mathcal{R}_c , which is defined as follows.

Definition 9 (c-reduction of a rewrite theory). *Let $\mathcal{R}^c = (\Sigma \cup \Sigma_c, E \cup G_c \cup A, R, \phi_c)$ be the c-extension of a rewrite theory $\mathcal{R} = (\Sigma, E \cup A, R, \phi)$. We call the rewrite theory $\mathcal{R}_c = (\Sigma \cup \Sigma_c, E \cup G_c \cup A, R_c, \phi_c)$ where $R_c = \{t \Rightarrow c(t') \text{ if } \text{cond} \mid (t \Rightarrow t' \text{ if } \text{cond}) \in R\}$ a c-reduction of \mathcal{R} .*

In words, \mathcal{R}_c is very much like \mathcal{R} , but each rule $t \Rightarrow t' \text{ if } \text{cond}$ in R , is transformed into a rule $t \Rightarrow c(t') \text{ if } \text{cond}$, i.e., into a rule where the canonizer function c is applied to the right hand side to ensure that canonization is performed after each system transition.

Note that in some cases it might be convenient not to apply the canonizer after each step. For instance if we know that the rule will always result in a canonical state we can save the time of applying the function.


It is trivial to show that \mathcal{R}_c is a c-reduction by construction, and in particular that $\mathcal{K}_c(\mathcal{R}) = \mathcal{K}(\mathcal{R}_c)$. It can also be shown that it has good executability properties. By the properties required for G_c , it inherits all the properties of the equational part of \mathcal{R} , namely sufficient completeness, confluence and termination modulo A . What is left to show is that \mathcal{R}_c is coherent modulo A .

Lemma 1 (coherence of \mathcal{R}_c). *Let $\mathcal{R} = (\Sigma, E \cup A, R, \phi)$ be ground confluent, terminating, sufficiently complete, and with R ground coherent with E modulo A . Then \mathcal{R}_c defined as above has R_c ground coherent with $E \cup G_c$ modulo A .*

Proof. Since \mathcal{R} is topmost and c is a frozen operator, any ground $\Sigma \cup \Sigma_c$ -term u of

sort **State** such that $\mathcal{R} \vdash u \rightarrow^1$

v must be a Σ term. Therefore, its $E \cup G_c/A$ -canonical form u' is exactly its E/A -canonical form. By the ground coherence of R with E modulo A we obtain the inner pentagon in the diagram on the right which trivially yields the outer pentagon, proving $E \cup G_c \vdash c(t') = c(u')$ as desired. \square



Theorem 2 (executability of \mathcal{R}_c). *Let $\mathcal{R} = (\Sigma \cup c, E \cup G_c \cup A, R, \phi)$ be a rewrite theory with good executability properties, then \mathcal{R}_c has good executability properties.*

Proof. The proof immediately follows from the fact that \mathcal{R}_c has the same equational part as \mathcal{R}_c and from Lemma 1. \square

Experiment	n	SymmSpin						c-reductions			
		ideal	weak		strong		weak		strong		
			%S	%T	%S	%T	%S	%T	%S	%T	
Peterson	2	-50.00%	-47.90%	+0.00%	-49.05%	+0.00%	-45.46%	+0.00%	-45.46%	+0.00%	
	3	-83.30%	-74.90%	-7.93%	-82.54%	-7.93%	-76.60%	+0.00%	-82.98%	+150.00%	
	4	95.80%	-91.38%	-76.96%	-95.50%	-83.62%	-92.85%	-75.00%	-94.81%	+50.00%	
DBM	7	-99.98%	-98.86%	-99.99%	-99.39%	-80.00%	-98.83%	-87.00%	-99.70%	+20.00%	
	8	-99.99%	-99.48%	-99.99%	-99.78%	-47.37%	-99.45%	-95.33%	-99.91%	-64.76%	

Table 1. SymmSpin vs c-reductions in Maude.

5 Performance experiments

We present here a subset of the experiments that we are carrying out, with the main purpose of validating the effectiveness of the implementation of the c-reduction approach in Maude.

Our main hypothesis to be checked is that the relative performance gain (in terms of runtime and state space reductions) is comparable to the one obtained by state-of-the-art model checkers. The second hypothesis is that c-reductions are more efficient than the previous approach to symmetry reductions in Maude described in [8]. Finally, we enrich the experiments where we combine various c-reductions not supported by the previously mentioned tools.

Comparison with SymmSPIN We have chosen SymmSPIN [14, 15] as a representative model checker with which to compare our approach. SymmSPIN extends the SPIN [3] model checker with support for symmetry reductions. It is worth to remark that we do not perform absolute comparison as we aim at checking the usefulness of our approach (experiments with SymmSPIN are anyway not reproducible since the tool is not available for download).

We have implemented in Maude two of the benchmark models tested in [14, 15], namely Peterson’s mutual exclusion protocol [16], and a database management system [17]. Both examples exhibit a full symmetry due to the presence of families of replicated concurrent processes with identical behavior.

We have considered various c-reduction strategies. For the sake of simplicity, we consider only the two best strategies of SymmSPIN and our implementation, for which regards time and state space reduction. We call them *strong* and *weak* as they are actually strong and non-strong canonizers.

Table 1 presents our results, for instances of the models with increasing number of components (n). We offer only results for those instances for which it has been possible to generate the unreduced state spaces, so to compare the relative gain of the reductions. The table also includes the “ideal” gain, which in the case of full symmetries is a factor of 1 divided by the factorial of the number n of participants, since the size of each orbit is at most $n!$. Even if the experiment is not a competition we highlight cells corresponding to the best results in each category (state space and run-time gain) for each model instance.

From the table we see that the two approaches provide reductions near to the ideal gain. The two strategies based on **weak** provide very similar outcomes, while the ones based on **strong** reduces similarly. SymmSPIN is more time-efficient, which is not a surprise, since the reduction algorithms are implemented in a procedural language (C) and efficiently compile, while our implementation is

Params		Not reduced		Reflection-based			c-reductions					
				States	%S	%T	weak			strong		
N	M	States	Ideal				States	%S	%T	States	%S	%T
2	5	38,029	-50%	19,295	-49%	+962%	21,630	-43%	+114%	19,025	-50%	+463%
3	2	72,063	-83%	13,280	-82%	+730%	29,534	-59%	+101%	12,235	-83%	+344%
3	3	952,747	-83%	174,428	-81%	+565%	307,532	-68%	-50%	160,121	-83%	+30%

Table 2. Reflection-based symmetry reduction vs. c-reductions in Maude.

based on a declarative language (Maude) running over an interpreter (the Maude engine).

Comparison with reflection-based symmetry reduction in Maude Our second set of experiments aims at checking whether c-reductions offers better performances than the symmetry reduction implementation in Maude described in [8]. Very briefly, the main idea of [8] is to select the canonical representative of a state on the basis of the lexicographical order of the meta-representation of the state, which is achieved by exploiting Maude’s reflection capabilities.

The comparison is performed over the Chain-Replication protocol used in [8]. As in the previous case, the replication of identical processes yields a full symmetry. Table 2 presents our results in the same format as Table 1 with the only exception that the model is instantiated with two parameters: the number n of replicated components and the number m of queries they perform. The table shows that the reductions of the reflection-based approach of [8] stand in the ranges between the ones obtained by our strategies. In particular our weak strategy offers worse space reductions, while the strong one offers better space reductions. More interestingly, our reduction strategies introduce much less time overhead, differing often by an order of magnitude. This is not a surprise, since resorting to Maude’s meta-level involves a considerable overhead.

Exploiting permutations, rotations, reuse and abstraction Our last set of experiments regards the joint application of a number of c-reductions of different nature, not supported by the previously mentioned tools. As a test case we have considered a message-passing solution to the Dining Philosophers problem along the lines of the case study used in [18], where newly generated messages (representing forks) receive fresh identities (as the original purpose of [18] was to reason about individual messages). There are a couple of regularities that can be exploited in the form of c-reductions and that happen to yield bisimulations (for an empty set of atomic predicates): the rotational symmetry of philosophers, the full symmetry of messages, the reuse of message identifiers and their abstraction. Of course, the situation is different when one considers state predicates that involve the identity of philosophers or messages. However, our goal here is to validate the effectiveness of the mentioned c-reductions.

Table 3 reports the results. The table presents the size of the state space and the time (in ms) to generate it, for instances of the model with increasing number of philosophers. The table considers the state spaces generated in the following cases: reuse of message identifiers (NR), reuse together with (rotational and full) symmetry reduction (NR+RS+FS), abstraction of name identifiers (NA), and abstraction of name identifiers together with (rotational) symmetry reduction

N	NR		NR+RS+FS		NA		NA+RS	
	States	Time	States	Time	States	Time	States	Time
2	21	0	10	0	18	0	10	0
3	115	8	27	12	76	0	27	8
4	801	100	86	60	322	20	86	48
5	6,251	1,456	275	320	1,364	124	275	248
6	54,869	20,765	982	1,732	5,778	740	982	1,412
7	541,731	463,080	3,499	11,828	24,476	6,624	3499	8,124
8	O.T.	O.T.	13,016	49,651	103,682	29,329	13,006	35,594
9	O.T.	O.T.	48,828	247,987	439,204	192,072	48,819	186,329

Table 3. c -reductions for the dining philosophers.

(NA+RS). The sizes of the unreduced state spaces are not shown since they are infinite (due to the creation of messages with fresh identifiers).

The first clear advantage is that reuse of message identifiers yields finite state spaces (since the number of messages in each state is bounded by n). Besides this, we see how combining various c -reductions results in better and more performant reductions. To be noticed is the fact that name reusing alone ran out of time (more than 5 hours) for models instantiated with more than 7 philosophers, while combining it with symmetry reductions (for messages and philosophers) allows us to manage greater instances. In particular, the best reductions are obtained with the combination “name abstraction + rotational symmetry”, while “name abstraction” alone offers the fastest explorations from 2 to 8 philosophers, and is outperformed by the combination “name abstraction + rotational symmetry” for greater instances.

6 Conclusion

We have presented c -reductions, a state space reduction technique, whose main idea is to use canonizer functions mapping each state into a (not necessarily unique) canonical representative of its equivalence class modulo a bisimulation equivalence relation, capturing some specific system regularities. The main differentiating features with respect to other state space reduction techniques are (i) reductions are defined using the system description language; (ii) (re)use of the standard techniques of the verification setting to check correctness of the reduction; (iii) automatization: both for applying the reduction and for checking its correctness; and (iv) generality: it subsumes in an uniform way symmetry reduction as well as other kinds of reductions (name reuse, abstraction).

We also presented a representative subset of the performance experiments we are carrying out. We have observed a comparable performance with mature tools such as SPIN and performance gains wrt. to a previous implementations of symmetry reductions in Maude [8]. Moreover, the flexibility of our approach has allowed us to define a wide range of reductions. Beyond the classical permutation and rotation symmetries, we have considered some simple cases of name reuse and name abstraction, which are crucial to deal with the infinite state spaces of systems with dynamic allocation of resources. Compared to the approach presented in [14, 15] we are able to treat a wider class of systems, where identifiers

of symmetric objects can appear as pointers in attributes of other objects, and with wider classes of symmetries such as rotational ones. Similar remarks can be made about the approach presented in [8], with respect to which we offer a wider class of reduction strategies and better performance.

A closely related work is the automatic symmetry detection of [4] which also provides an alternative symmetry reduction extension (TopSPIN) of the SPIN model checker. Our approach does not consider automatic detection of symmetries but, instead, user-definable ones, together with a methodology to check their correctness, with the main advantage being that we rely on tools and techniques used to perform the verification of the system itself. Both lines of research are essentially complementary and can converge in a synergistic way.

We are also investigating how to apply our approach to improve the efficiency of rewriting-logic based interpreters of programming languages with primitives for dynamic memory allocation, and to further explore the combination of *c*-reductions with other state space reduction techniques, with a particular attention to those already proposed in the setting of rewriting logic and Maude, such as the language-generic partial order reduction proposed in [19], and not forgetting abstraction techniques [20] where some promising efforts have been carried out [21].

Current efforts are also devoted to a deeper investigation of state space reduction techniques based on name-reuse. Such techniques have a limited impact on Kripke models, where object identifiers are “global”, in the sense that they represent an entity through all the states of a model. These techniques become instead fundamental using more sophisticated models (together with non-propositional property specification languages), where object identifiers are local to single states, and “trans-states identities” are explicitly represented through suitable mappings associated to state transitions. Notable examples are History Dependent Automata [22], High-level Allocational Büchi Automata [18], Graph Transition Systems [23] and Counterpart models [24].

We are also working further (and experimenting) on the verification steps that we could only sketch in this paper (steps §4.1-§4.3). Moreover we would like to further explore reductions not based on group actions, extend the approach to deal with equivalence preorder like simulations, and further develop the automatization of our approach and the use of the integrated Maude formal environment [25].

References

1. Clarke, E.M., Grumberg, O., Peled, D.A.: Model Checking. The MIT Press (1999)
2. Wahl, T., Donaldson, A.F.: Replication and abstraction: Symmetry in automated formal verification. *Symmetry* **2** (2010) 799–847
3. Holzmann, G.: The SPIN model checker: primer and reference manual. Addison-Wesley Professional (2003)
4. Donaldson, A.F., Miller, A.: A computational group theoretic symmetry reduction package for the SPIN model checker. In: Algebraic Methodology and Software Technology. (2006) 374–380

5. Meseguer, J., Palomino, M., Martí-Oliet, N.: Algebraic simulations. *Journal of Logic and Algebraic Programming* **79** (2010) 103–143
6. Meseguer, J.: Conditional rewriting logic as a united model of concurrency. *Theoretical Computer Science (TCS)* **96** (1992) 73–155
7. Clavel, M., Durán, F., Eker, S., Lincoln, P., Martí-Oliet, N., Meseguer, J., Talcott, C.L.: *All About Maude*. Volume 4350 of LNCS. Springer (2007)
8. Rodríguez, D.E.: Combining techniques to reduce state space and prove strong properties. In: *Proceedings of the 7th International Workshop on Rewriting Logic and its Applications (WRLA 2008)*. Volume 238(3) of ENTCS. (2009) 267 – 280
9. The Maude Invariant Analyzer Tool (InvA), camilorocha.info/software/inva.
10. Clavel, M., Palomino, M., Riesco, A.: Introducing the ITP tool: a tutorial. *Journal of Universal Computer Science* **12** (2006) 1618–1650
11. Maude Interactive Theorem Prover, maude.cs.uiuc.edu/tools/itp/.
12. Durán, F., Meseguer, J.: A church-rosser checker tool for conditional order-sorted equational maude specifications. In Ölveczky, P.C., ed.: *Proceedings of the 8th International Workshop on Rewriting Logic and its Applications (WRLA'10)*. Volume 6381 of LNCS., Springer (2010) 69–85
13. Rocha, C., Meseguer, J.: Proving safety properties of rewrite theories. In Corradini, A., Klin, B., Crstea, C., eds.: *Algebra and Coalgebra in Computer Science*. Volume 6859 of LNCS. Springer (2011) 314–328
14. Bosnacki, D., Dams, D., Holenderski, L.: A heuristic for symmetry reductions with scalarsets. In: *International Symposium of Formal Methods Europe on Formal Methods for Increasing Software Productivity (FME'01)*, Springer (2001)
15. Bosnacki, D., Dams, D., Holenderski, L.: Symmetric SPIN. *International Journal on Software Tools for Technology Transfer (STTT)* **4** (2002) 92–106
16. Lynch, N.A.: *Distributed Algorithms*. Morgan Kaufmann Publishers Inc. (1996)
17. Valmari, A.: Stubborn sets for reduced state generation. In: *Proceedings on Advances in Petri nets 1990*, Springer (1991)
18. Distefano, D., Rensink, A., Katoen, J.P.: Model checking birth and death. In Baeza-Yates, R.A., Montanari, U., Santoro, N., eds.: *2nd International Conference on Theoretical Computer Science (TCS'02)*. Volume 223., Kluwer (2002)
19. Farzan, A., Meseguer, J.: Partial order reduction for rewriting semantics of programming languages. In: *Proceedings of the 6th International Workshop on Rewriting Logic and its Applications (WRLA'06)*. Volume 176 of ENTCS. (2007) 61–78
20. Meseguer, J., Palomino, M., Martí-Oliet, N.: Equational abstractions. *Theoretical Computer Science* **403** (2008) 239–264
21. Donaldson, A.F., Kaiser, A., Kroening, D., Wahl, T.: Symmetry-aware predicate abstraction for shared-variable concurrent programs. In Gopalakrishnan, G., Qadeer, S., eds.: *Proceedings of the 23rd International Conference on Computer Aided Verification (CAV'11)*. Volume 6806 of LNCS., Springer (2011) 356–371
22. Montanari, U., Pistore, M.: Structured coalgebras and minimal HD-automata for the π -calculus. *Theoretical Computer Science* **340** (2005) 539–576
23. Rensink, A.: Isomorphism checking in GROOVE. *ECEASST* **1** (2006)
24. Gadducci, F., Lluch Lafuente, A., Vandin, A.: Counterpart semantics for a second-order μ -calculus. In Ehrig, H., Rensink, A., Rozenberg, G., Schürr, A., eds.: *5th International Conference on Graph Transformation (ICGT'10)*. Volume 6372 of LNCS. Springer (2010) 282–297
25. The Maude Formal Environment, <http://maude.lcc.uma.es/MFE/>.