

R-CHECK: A Model Checker for Verifying Reconfigurable MAS

Extended Abstract

Yehia Abd Alrahman
University of Gothenburg
Gothenburg, Sweden

Shaun Azzopardi
University of Gothenburg
Gothenburg, Sweden

Nir Piterman
University of Gothenburg
Gothenburg, Sweden

ABSTRACT

Reconfigurable multi-agent systems consist of a set of autonomous agents, with integrated interaction capabilities that feature opportunistic interaction. Agents seemingly reconfigure their interactions interfaces by forming collectives, and interact based on mutual interests. Finding ways to design and analyse the behaviour of these systems is a vigorously pursued research goal. We propose a model checker, named R-CHECK, to allow reasoning about these systems both from an individual- and a system- level. R-CHECK also permits reasoning about interaction protocols and joint missions. R-CHECK supports a high-level input language with symbolic semantics, and provides a modelling convenience for interaction features such as reconfiguration, coalition formation, self-organisation, etc.

KEYWORDS

Model-checking; Agent Theories and Models; Verification of Multi-Agent Systems

ACM Reference Format:

Yehia Abd Alrahman, Shaun Azzopardi, and Nir Piterman. 2022. R-CHECK: A Model Checker for Verifying Reconfigurable MAS: Extended Abstract. In *Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2022)*, Online, May 9–13, 2022, IFAAMAS, 3 pages.

1 OVERVIEW OF R-CHECK

R-CHECK accepts a high-level language that is based on the symbolic RECIPE formalism [2, 3]. We will present the syntax of R-CHECK language and informally describe its semantics. For a full exposition of the formal definition of R-CHECK and its usage through sizeable case studies, we refer the reader to [4].

We first introduce the class `agent`, its structure, and how to instantiate it; we introduce the syntax of its behaviour and how to create a system of agents. The class `agent` is reported in Fig. 1.

```
1  agent agent_name
2  local:
3  var_name:type , ... , var_name:type
4  init:  $\theta_T$ 
5  relabel:
6  common_var <- Exp
7  .
8  common_var <- Exp
9  receive-guard:  $g^r(V_T, ch)$ 
10 repeat: P
```

Figure 1: An agent class

Each agent class has a name that identifies a specific type of behaviour; and uses a set of channels to interact with others. We permit creating multiple instances/copies with the same class of behaviour. An agent has a local state “local” represented by a set of local variables V_T , and a relabelling function to interact with other agents anonymously. The initial state of an agent `init`: θ_T is a predicate characterising the initial assignments to the agent local variables. The section `receive-guard`: $g^r(V_T, ch)$ specifies the connectedness of the agent to channels given a current assignment to its local variables. The non-terminating behaviour of an agent is represented by `repeat`: P , executing the process P indefinitely.

An agent type of name “ A ” can be instantiated as follows $A(id, \theta)$. That is, we create an instance of “ A ” with identity id and an additional initial restriction θ . Here, we take the conjunction of θ with the predicate in the `init` section of the type “ A ” as the initial condition of this instance. We use the parallel composition operator \parallel to inductively define a system as in the following production rule:

$$\text{(System)} \quad S ::= A(id, \theta) \mid S_1 \parallel S_2$$

That is, a system is either an instance of agent type or a parallel composition (with reconfigurable multicast and broadcast semantics as in [2, 3]) of set of instances of (possibly) different types. Agents interact by *state-parametric* message exchange.

The syntax of an R-CHECK process is inductively defined as:

$$\begin{aligned} \text{(Process)} \quad P &::= P;P \mid P+P \mid \text{rep } P \mid C \\ \text{(Command)} \quad C &::= l:C \mid \langle \Phi \rangle ch! \pi d U \mid \langle \Phi \rangle ch? U \end{aligned}$$

A process P is either a sequential composition of two processes $P;P$, a non-deterministic choice between two processes $P+P$, a loop $\text{rep } P$, or a command C . There are three types of commands corresponding to either a labelled command, a message-send or a message-receive. A command of the form $l:C$ is a syntactic labelling and is used to allow the model checker to reason about syntactic elements as we will see later. A command of the form $\langle \Phi \rangle ch! \pi d U$ corresponds to a message-send. The predicate Φ is an assertion over the current local state of an agent, i.e., is a precondition that must hold before the transition can be taken. As the names suggest, ch , π and (respectively) d are the communication channel, the sender predicate (specifying the targeted receivers), and the assignment to data variables (i.e., the actual content of the message). Lastly, U is an update to local variables after taking the transition. We use “!” to distinguish send transitions. A command of the form $\langle \Phi \rangle ch? U$ corresponds to a message-receive. Differently from message-send, the predicate Φ can also predicate on the received values from the incoming message, i.e., the assignment d .

We can easily create an R-CHECK system as in Equation 1 below. There, a set of identical clients anonymously coordinate with a resource manager to get virtual machines (VM).

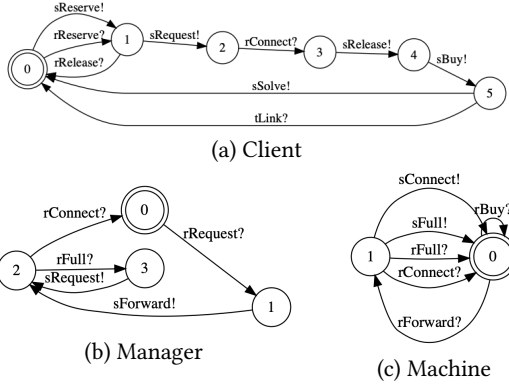


Figure 2: symbolic structure automata

$$\begin{aligned}
\text{system} = & \text{Client}(\text{client1}, \text{TRUE}) \parallel \text{Client}(\text{client2}, \text{TRUE}) \\
& \parallel \text{Client}(\text{client3}, \text{TRUE}) \parallel \text{Manager}(\text{manager}, \text{TRUE}) \\
& \parallel \text{Machine}(\text{machine1}, \text{gLink} = \text{g1} \wedge \text{pLink} = \text{vmm1}) \\
& \parallel \text{Machine}(\text{machine2}, \text{gLink} = \text{g1} \wedge \text{pLink} = \text{vmm2}) \\
& \parallel \text{Machine}(\text{machine3}, \text{gLink} = \text{g2} \wedge \text{pLink} = \text{vmm3})
\end{aligned} \quad (1)$$

Notice that the machines $\{\text{machine}_1, \dots, \text{machine}_3\}$, each belongs to a specific group and uses a private link to interact. For instance, machine_1 belongs to group “g1” (the high performance machines) and has a private link named “vmm1”.

The symbolic behaviour of this system is reported in Fig. 2, where send transitions “!” synchronise with receive ones “?” having corresponding labels. The full example is reported in [4].

2 NUXMV AND MODEL-CHECKING

We integrate R-CHECK with the nuXmv model checker [6] to enable an enhanced symbolic LTL model-checking. We also demonstrate our developments using examples. We will show how the combined features of R-CHECK, the symbolic LTL model-checking, and nuXmv provides a powerful tool to verify high-level features of reconfigurable and interactive systems.

R-CHECK provides an interactive simulator that allows the user to simulate the system either randomly or based on predicates that the user supplies. For instance, one can refer to message -send and -receive using command labels. A constraint on a send transition like “client1-sReserve”, to denote the sending of the message labelled with “sReserve”, means that this transition is feasible in the current state of simulation. R-CHECK is also supported with a web editor, syntax highlighting and a visualising tool. For instance, once the model of Equation 1 is compiled, R-CHECK produces the corresponding labelled and symbolic structure automata in Fig. 2, and thus the user may use these automata to reason about interactions.

Symbolic Model Checking. R-CHECK supports both symbolic LTL model checking and bounded LTL model checking. We illustrate the capabilities of R-CHECK by several examples based on Equation 1, the automata in Fig. 2 as the system under consideration.

We show how to verify properties about agents both from individual and interaction protocols level by predicating on message exchange rather than on atomic propositions. It should be noted that the transition labels in Fig. 2 are not mere labels, but rather

predicates with truth values changing dynamically at run-time, introducing opportunistic interaction. For instance, we can reason about a client and its connection to the system as follows:

$$G(\text{client1-sReserve} \rightarrow F \text{client1-sRelease}) \quad (1)$$

$$G(\text{client1-sRequest} \rightarrow F \text{client1-rConnect}) \quad (2)$$

The liveness condition (1) specifies that the client does not hold a live lock on a shared link. Namely, the client releases the shared link eventually. The liveness condition (2) specifies that the *system* is responsive, i.e., after the client’s request, other agents collaborate to eventually supply a connection.

We can also reason about synchronisation and reconfiguration in relation to local state as in the following:

$$G(\text{manager-sForward} \rightarrow X \text{machine1-rForward}) \quad (3)$$

$$F(\text{client1-sRelease} \wedge G(\neg \text{client1-rConnect})) \quad (4)$$

In (3), we refer to synchronisation, i.e., the manager has to forward the request before the machine can receive it. We can refer to reconfiguration in (4), i.e., eventually the client disconnects from the common link, and it can never be able to receive connection on that link.

We can also specify channel mobility and joint missions from a declarative and centralised point of view.

$$\left(\begin{array}{l}
F(\text{client1-mLink} \neq \text{empty}) \ \& \\
\quad F(\text{client2-mLink} \neq \text{empty}) \ \& \\
\quad \quad F(\text{client3-mLink} \neq \text{empty}) \\
F(\text{client1-sSolve} \mid \text{client2-sSolve} \mid \text{client3-sSolve})
\end{array} \right) \longrightarrow$$

That is, every client will eventually receive a mobile link (i.e., its $\text{mLink} \neq \text{empty}$) where it will use this private link to get a VM, and eventually one client will initiate the termination of the mission by synchronising with the other clients to solve the joint problem.

We are unaware of a model-checker that enables reasoning at such a high-level.

3 CONCLUDING REMARKS

We introduced the R-CHECK model checking toolkit for verifying and simulating reconfigurable multi-agent system. R-CHECK is supported with a command line tool, a web editor with syntax highlighting and visualisation. We integrated R-CHECK with nuXmv to enable LTL symbolic (bounded) model checking. We showed that this specialised integration provides a powerful tool that permits verifying high-level features such as interaction protocols, joint missions, channel mobility, reconfiguration, self-organisation, etc.

R-CHECK combines the lessons learnt from communication models like *AbC* [1, 5] and *RECIPE* [2, 3], and mainstreams model checkers like *MCMAS* [12] which is based on *Interpreted Systems* [8], *MTSA* toolkit [7] (based on Hoare’s CSP calculus [10] and *Fluent Linear Temporal logic (FLTL)* [9]), *SPIN* [11] (for protocol design). Furthermore, R-CHECK strives for expressiveness while preserving minimality and simplicity.

Future works. We plan to integrate LTOIL to R-CHECK from [3]. Indeed, the authors in [3] provide a PSPACE algorithm for LTOIL model checking (improved from EXPSpace in [2]). This way, we would not only be able to refer to message exchange in logical formulas, but also to identify the intentions of agents in the interaction and characterise potential interacting partners.

REFERENCES

- [1] Yehia Abd Alrahman, Rocco De Nicola, and Michele Loreti. 2019. A calculus for collective-adaptive systems and its behavioural theory. *Inf. Comput.* 268 (2019). <https://doi.org/10.1016/j.ic.2019.104457>
- [2] Yehia Abd Alrahman, Giuseppe Perelli, and Nir Piterman. 2020. Reconfigurable Interaction for MAS Modelling. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems, AAMAS '20, Auckland, New Zealand, May 9-13, 2020*, Amal El Fallah Seghrouchni, Gita Sukthankar, Bo An, and Neil Yorke-Smith (Eds.). International Foundation for Autonomous Agents and Multiagent Systems, 7–15.
- [3] Yehia Abd Alrahman and Nir Piterman. 2021. Modelling and verification of reconfigurable multi-agent systems. *Auton. Agents Multi Agent Syst.* 35, 2 (2021), 47. <https://doi.org/10.1007/s10458-021-09521-x>
- [4] Yehia Abd Alrahman, Shaun Azzopardi, and Nir Piterman. 2022. R-CHECK: A Model Checker for Verifying Reconfigurable MAS. arXiv:2201.06312 [cs.LO]
- [5] Yehia Abd Alrahman, Rocco De Nicola, and Michele Loreti. 2020. Programming interactions in collective adaptive systems by relying on attribute-based communication. *Sci. Comput. Program.* 192 (2020), 102428. <https://doi.org/10.1016/j.scico.2020.102428>
- [6] Alessandro Cimatti and Alberto Griggio. 2012. Software Model Checking via IC3. In *Computer Aided Verification - 24th International Conference, CAV 2012, Berkeley, CA, USA, July 7-13, 2012 Proceedings (Lecture Notes in Computer Science, Vol. 7358)*, P. Madhusudan and Sanjit A. Seshia (Eds.). Springer, 277–293. https://doi.org/10.1007/978-3-642-31424-7_23
- [7] Nicolás D'Ippolito, Dario Fischbein, Marsha Chechik, and Sebastián Uchitel. 2008. MTSa: The Modal Transition System Analyser. In *23rd IEEE/ACM International Conference on Automated Software Engineering (ASE 2008), 15-19 September 2008, L'Aquila, Italy*. IEEE Computer Society, 475–476. <https://doi.org/10.1109/ASE.2008.78>
- [8] Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. 1995. *Reasoning about Knowledge*. MIT Press.
- [9] Dimitra Giannakopoulou and Jeff Magee. 2003. Fluent model checking for event-based systems. In *Proceedings of the 9th European software engineering and 11th ACM SIGSOFT international symposium on Foundations of software engineering (Helsinki, Finland)*. ACM, 257–266.
- [10] C. A. R. Hoare. 2021. Communicating Sequential Processes. In *Theories of Programming: The Life and Works of Tony Hoare*, Cliff B. Jones and Jayadev Misra (Eds.). ACM / Morgan & Claypool, 157–186. <https://doi.org/10.1145/3477355.3477364>
- [11] Gerard J. Holzmann. 1997. The model checker SPIN. *IEEE Transactions on software engineering* 23, 5 (1997), 279–295.
- [12] Alessio Lomuscio, Hongyang Qu, and Franco Raimondi. 2017. MCMAS: an open-source model checker for the verification of multi-agent systems. *STTT* 19, 1 (2017), 9–30.