Compositional Value Iteration with Pareto Caching

Kazuki Watanabe, Marck van der Vegt, Sebastian Junges, and Ichiro Hasuo

The de-facto standard approach in MDP verification is based on value iteration (VI). We propose compositional VI, a framework for model checking compositional MDPs, that addresses efficiency while maintaining soundness. Concretely, compositional MDPs naturally arise from the combination of individual components, and their structure can be expressed using, e.g., string diagrams. Towards efficiency, we observe that compositional VI repeatedly verifies individual components. We propose a technique called Pareto caching that allows to reuse verification results, even for previously unseen queries. Towards soundness, we present two stopping criteria: one generalizes the optimistic value iteration paradigm and the other uses Pareto caches in conjunction with recent baseline algorithms. Our experimental evaluations shows the promise of the novel algorithm and its variations, and identifies challenges for future work.

1 Introduction

1.1 MDP Model Checking and Value Iteration

Markov decision processes (MDPs) are the standard model for sequential decision making in stochastic settings. A standard question in the verification of MDPs is: *what is the maximal probability that an error state is reached.* MDP model checking is an active topic in the formal verification community.

Value iteration (VI) [17] is an iterative and approximate method whose performance in MDP model checking is well-established [10] [3] [9]. Several extensions with *soundness* have been proposed; they provide, in addition to under-approximations, also over-approximations with a desired precision [19] [10] [7] [1] [16], so that an approximate answer comes with an error bound. These sound algorithms are implemented in mature model checkers such as Prism [14], Modest [8], and Storm [12].

1.2 Compositional Model Checking

Even with these state-of-the-art algorithms, it is a challenge to model check large MDPs efficiently with high precision. Experiments observe that MDPs with more than 10^8 states are too large for those algorithms [13] [22] [23]—they simply do not fit in memory. However, such large MDPs often arise as models of complicated stochastic systems, e.g. in the domains of network and robotics. Furthermore, even small models may be numerically challenging to solve due to their structure [7] [9] [1].

Compositional model checking is a promising approach to tackle this scalability challenge. Given a compositional structure of a target system, compositional model checking executes a divide-and-conquer algorithm that avoids loading the entire state space at once, often solving the above mem-

パレートキャッシュを用いた Compositional な Value Iteration

渡邉知樹, 蓮尾一郎, 国立情報学研究所, National Institute of Informatics.

渡邉知樹, 蓮尾一郎, 総合研究大学院大学, The Graduate University for Advanced Studies (SOKENDAI).

Marck van der Vegt, Sebastian Junges, ラドバウド大学, Radboud University.

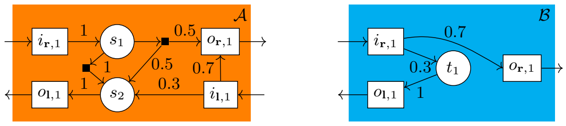

図 1 open MDPs $\mathcal{A}$ and $\mathcal{B}$.

ory problem. Moreover, reusing the model checking results for components can lead to speed-up by magnitudes. Although finding a suitable compositional structure for a given "monolithic" MDP is still open, many systems come with such an *a priori* compositional structure. For example, such compositional structures are often assumed in robotics and referred to as *hierarchical models* [11] [2] [13] [6] [20] [15] [21].

Recently, *string diagrams of MDPs* are introduced for compositional model checking [22] [23]; the current paper adopts this formalism. There, MDPs are extended with (open) entrances and exits (Fig. 1), and they get composed by *sequential composition* $⨟$ and *sum* $\oplus$. See Fig. 2, where the right-hand sides are simple juxtapositions of graphs (wires get connected in $⨟$). This makes the formalism focused on sequential (as opposed to parallel) composition. This restriction eases the design of compositional algorithms; yet, the formalism is rich enough to capture the compositional structures of many system models.

1.3 Current Work: Compositional Value Iteration

In this paper, we present a *compositional value iteration (CVI)* algorithm that solves reachability probabilities of string diagrams of MDPs, operating in a divide-and-conquer manner along compositional structures. Our approximate VI algorithm comes with *soundness*—it produces error bounds—and exploits compositionality for *efficiency*.

Specifically, for soundness, we lift the recent paradigm of *optimistic value iteration (OVI)* [10] to the current compositional setting. We use it both for local (component-level) model checking and—in one of the two global VI stopping criteria that we present—for providing a global over-approximation.

For efficiency, firstly, we adopt a *top-down* compositional approach where each component is model-checked repeatedly, each time on a different weight **w**, in a *by-need* manner. Secondly, in order to suppress repetitive computation on similar weights, we introduce a novel technique of *Pareto caching* that allows "approximate reuse" of model checking results. This closely relates to multi-objective probabilistic model checking [4] [5] [18], without the explicit goal of building Pareto curves. Our Pareto caching also leads to another (*sound*) global VI stopping criterion that is based on the approximate bottom-up approach [23].

Our algorithm is approximate (unlike the exact one in [22]), and top-down (unlike the *bottom-up* approximate one in [23]). Experimental evaluation demonstrates its performance thanks to the combination of these two features. See [24] for details.

参 考 文 献

[1] Baier, C., Klein, J., Leuschner, L., Parker, D., and Wunderlich, S.: Ensuring the Reliability of Your Model Checker: Interval Iteration for Markov Decision Processes, *CAV (1)*, LNCS, Vol. 10426, Springer, 2017, pp. 160–180.
[2] Barry, J. L., Kaelbling, L. P., and Lozano-Pérez, T.: DetH*: Approximate Hierarchical Solution of Large Markov Decision Processes, *IJCAI*, IJCAI/AAAI, 2011, pp. 1928–1935.
[3] Budde, C. E., Hartmanns, A., Klauck, M., Kretínský, J., Parker, D., Quatmann, T., Turrini, A., and Zhang, Z.: On Correctness, Precision, and Performance in Quantitative Verification - QComp 2020 Competition Report, *ISoLA (4)*, LNCS, Vol. 12479, Springer, 2020, pp. 216–241.
[4] Etessami, K., Kwiatkowska, M. Z., Vardi, M. Y., and Yannakakis, M.: Multi-Objective Model Checking of Markov Decision Processes, *Log. Methods Comput. Sci.*, Vol. 4, No. 4(2008).
[5] Forejt, V., Kwiatkowska, M. Z., and Parker, D.: Pareto Curves for Probabilistic Model Checking, *ATVA*, LNCS, Vol. 7561, Springer, 2012, pp. 317–332.

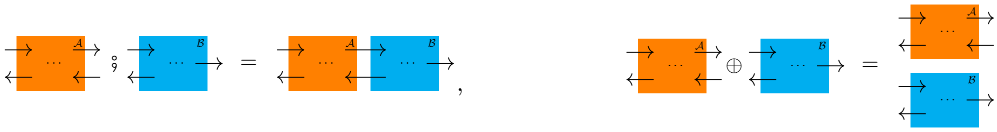

図 2 sequential composition $\mathcal{A} \fatsemi \mathcal{B}$ and sum $\mathcal{A} \oplus \mathcal{B}$ of open MDPs. The framework is *bidirectional* (edges can be left- and right-ward); thus loops can arise in $\mathcal{A} \fatsemi \mathcal{B}$.

[6] Gopalan, N., desJardins, M., Littman, M. L., MacGlashan, J., Squire, S., Tellex, S., Winder, J., and Wong, L. L. S.: Planning with Abstract Markov Decision Processes, *ICAPS*, AAAI Press, 2017, pp. 480–488.

[7] Haddad, S. and Monmege, B.: Interval iteration algorithm for MDPs and IMDPs, *Theor. Comput. Sci.*, Vol. 735(2018), pp. 111–131.

[8] Hartmanns, A. and Hermanns, H.: The Modest Toolset: An Integrated Environment for Quantitative Modelling and Verification, *TACAS*, LNCS, Vol. 8413, Springer, 2014, pp. 593–598.

[9] Hartmanns, A., Junges, S., Quatmann, T., and Weininger, M.: A Practitioner's Guide to MDP Model Checking Algorithms, *TACAS (1)*, LNCS, Vol. 13993, Springer, 2023, pp. 469–488.

[10] Hartmanns, A. and Kaminski, B. L.: Optimistic Value Iteration, *CAV (2)*, LNCS, Vol. 12225, Springer, 2020, pp. 488–511.

[11] Hauskrecht, M., Meuleau, N., Kaelbling, L. P., Dean, T. L., and Boutilier, C.: Hierarchical Solution of Markov Decision Processes using Macro-actions, *UAI*, Morgan Kaufmann, 1998, pp. 220–229.

[12] Hensel, C., Junges, S., Katoen, J., Quatmann, T., and Volk, M.: The probabilistic model checker Storm, *Int. J. Softw. Tools Technol. Transf.*, Vol. 24, No. 4(2022), pp. 589–610.

[13] Junges, S. and Spaan, M. T. J.: Abstraction-Refinement for Hierarchical Probabilistic Models, *CAV (1)*, LNCS, Vol. 13371, Springer, 2022, pp. 102–123.

[14] Kwiatkowska, M. Z., Norman, G., and Parker, D.: PRISM 4.0: Verification of Probabilistic Real-Time Systems, *CAV*, LNCS, Vol. 6806, Springer, 2011, pp. 585–591.

[15] Neary, C., Verginis, C. K., Cubuktepe, M., and Topcu, U.: Verifiable and Compositional Reinforcement Learning Systems, *ICAPS*, AAAI Press, 2022, pp. 615–623.

[16] Phalakarn, K., Takisaka, T., Haas, T., and Hasuo, I.: Widest Paths and Global Propagation in Bounded Value Iteration for Stochastic Games, *CAV (2)*, LNCS, Vol. 12225, Springer, 2020, pp. 349–371.

[17] Puterman, M. L.: *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, Wiley Series in Probability and Statistics, Wiley, 1994.

[18] Quatmann, T.: *Verification of multi-objective Markov models*, Phd thesis, 2023.

[19] Quatmann, T. and Katoen, J.: Sound Value Iteration, *CAV (1)*, LNCS, Vol. 10981, Springer, 2018, pp. 643–661.

[20] Saxe, A. M., Earle, A. C., and Rosman, B.: Hierarchy Through Composition with Multitask LMDPs, *ICML*, Proceedings of Machine Learning Research, Vol. 70, PMLR, 2017, pp. 3017–3026.

[21] Vien, N. A. and Toussaint, M.: Hierarchical Monte-Carlo Planning, *AAAI*, AAAI Press, 2015, pp. 3613–3619.

[22] Watanabe, K., Eberhart, C., Asada, K., and Hasuo, I.: Compositional Probabilistic Model Checking with String Diagrams of MDPs, *CAV (3)*, LNCS, Vol. 13966, Springer, 2023, pp. 40–61.

[23] Watanabe, K., van der Vegt, M., Hasuo, I., Rot, J., and Junges, S.: Pareto Curves for Compositionally Model Checking String Diagrams of MDPs, *TACAS (2)*, LNCS, Vol. 14571, Springer, 2024, pp. 279–298.

[24] Watanabe, K., van der Vegt, M., Junges, S., and Hasuo, I.: Compositional Value Iteration with Pareto Caching, 2024. a longer version, to appear in CAV'24.