



Delft University of Technology

IntervalMDP.jl

Accelerated Value Iteration for Interval Markov Decision Processes

Mathiesen, Frederik Baymler; Lahijanian, Morteza; Laurenti, Luca

DOI

[10.1016/j.ifacol.2024.07.416](https://doi.org/10.1016/j.ifacol.2024.07.416)

Publication date

2024

Document Version

Final published version

Published in

IFAC-PapersOnLine

Citation (APA)

Mathiesen, F. B., Lahijanian, M., & Laurenti, L. (2024). IntervalMDP.jl: Accelerated Value Iteration for Interval Markov Decision Processes. *IFAC-PapersOnLine*, 58(11), 1-6.
<https://doi.org/10.1016/j.ifacol.2024.07.416>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

IntervalMDP.jl: Accelerated Value Iteration for Interval Markov Decision Processes

Frederik Baymler Mathiesen* Morteza Lahijanian**
Luca Laurenti***

* Delft University of Technology, the Netherlands (e-mail: f.b.mathiesen@tudelft.nl)

** University of Colorado, Boulder, USA (e-mail: morteza.lahijanian@colorado.edu)

*** Delft University of Technology, the Netherlands (e-mail: l.laurenti@tudelft.nl)

Abstract: In this paper, we present `IntervalMDP.jl`, a Julia package for probabilistic analysis of interval Markov Decision Processes (IMDPs). `IntervalMDP.jl` facilitates the synthesis of optimal strategies and verification of IMDPs against reachability specifications and discounted reward properties. The library supports sparse matrices and is compatible with data formats from common tools for the analysis of probabilistic models, such as PRISM. A key feature of `IntervalMDP.jl` is that it presents both a multi-threaded CPU and a GPU-accelerated implementation of value iteration algorithms for IMDPs. In particular, `IntervalMDP.jl` takes advantage of the Julia type system and the inherently parallelizable nature of value iteration to improve the efficiency of performing analysis of IMDPs. On a set of examples, we show that `IntervalMDP.jl` substantially outperforms existing tools for verification and strategy synthesis for IMDPs in both computation time and memory consumption.

Copyright © 2024 The Authors. This is an open access article under the CC BY-NC-ND license (<https://creativecommons.org/licenses/by-nc-nd/4.0/>)

Keywords: Markov Decision Processes, Robust Value Iteration, Reachability, Control Synthesis, Verification.

1. INTRODUCTION

Interval Markov Decision Processes (IMDPs) (Givan et al., 2000) are a class of uncertain Markov Decision Processes (MDPs), where the transition probabilities between each pair of states are uncertain and lie in some probability interval. Due to their modeling flexibility and the existence of efficient robust value iteration algorithms (Givan et al., 2000; Haddad and Monmege, 2018), IMDPs have recently gained increasing attention both in the computer science and control communities (Haddad and Monmege, 2018; Cauchi et al., 2019; Laurenti and Lahijanian, 2023; Lavaei et al., 2022). In particular, IMDPs are popular for abstraction-based verification and control of stochastic hybrid systems Cauchi et al. (2019); Lavaei et al. (2022). However, existing tools for performing verification and strategy synthesis of IMDPs lack parallelization via multi-threading and GPU-acceleration, under-utilizing available hardware.

In this paper, we introduce `IntervalMDP.jl`, which is a Julia package for reachability analysis of IMDPs. The tool supports both reachability specifications and discounted-reward properties, and allows the user to query for both optimal strategies and quantitative values of satisfaction. The package contains both a CPU implementation and an implementation that uses General-Purpose Graphics Processing Units (GPUs) allowing one to use CUDA-capable hardware (NVIDIA Corporation, 2023) to accelerate the computation. The package is developed in Julia, which is a

modern programming language that targets the scientific community (Bezanson et al., 2012). It enables both fast prototyping and accelerated code to be written in the same language, including the ability to write CUDA kernels for custom accelerated computations. Furthermore, through Julia's parametric typing, `IntervalMDP.jl` supports single- and double-precision floating point numbers, as well as rational numbers for exact arithmetic with transition probabilities.

We evaluate `IntervalMDP.jl` on various benchmarks and compare it with PRISM (Kwiatkowska et al., 2011) and `bdmp-tool` (Lahijanian, 2021), the only tools available to perform reachability analysis for `IntervalMDP.jl` to the best of our knowledge. The benchmarks include 35 IMDPs taken from the literature, with the total number of transitions between states ranging from a few tens for the smaller models to tens of millions for the larger models. The empirical analysis shows that `IntervalMDP.jl` CPU implementation is on average 2-4× faster compared to the state of the art, while the GPU implementation can achieve speed-ups of various orders of magnitude on the larger systems. Furthermore, because of the use of sparse matrices and the Julia type system, in all cases, `IntervalMDP.jl` requires less memory compared to PRISM and `bdmp-tool`.

The paper is organized as follows. First, in Section 2, we formally introduce IMDPs and robust value iteration for IMDPs. Then, in Section 3, we give an overview of `IntervalMDP.jl` and describe how IMDPs can be created and stored in the tool and how to perform strategy

synthesis and verification. In Section 4, we detail our algorithmic approach to robust value iteration on GPUs. Finally, in Section 5, we illustrate the effectiveness of `IntervalMDP.jl` on various benchmarks.

2. THE GOAL OF INTERVALMDP.JL: ROBUST VALUE ITERATION FOR IMDPS

Interval Markov Decision Processes (IMDPs), also called bounded-parameter MDPs (Givan et al., 2000), are a generalization of MDPs, where the transition probabilities between each pair of states are not known exactly, but they are constrained to be in some independent probability intervals. Formally, an IMDP M is a tuple $M = (S, A, \underline{P}, \overline{P})$, where

- S is a finite set of states,
- A is a finite set of actions, where the set of actions available at state $s \in S$ is denoted by $\mathcal{A}(s) \subseteq A$,
- $\underline{P} : S \times A \times S \rightarrow [0, 1]$ is a function, where $\underline{P}(s, a, s')$ defines the lower bound of the transition probability from state $s \in S$ to state $s' \in S$ under action $a \in \mathcal{A}(s)$,
- $\overline{P} : S \times A \times S \rightarrow [0, 1]$ is a function, where $\overline{P}(s, a, s')$ defines the upper bound of the transition probability from state $s \in S$ to state $s' \in S$ under action $a \in \mathcal{A}(s)$.

For each state-action pair $(s, a) \in S \times A$ where $a \in \mathcal{A}(s)$, it holds that $\sum_{s' \in S} \underline{P}(s, a, s') \leq 1 \leq \sum_{s' \in S} \overline{P}(s, a, s')$ and a transition probability distribution $p_{s,a} : S \rightarrow [0, 1]$ is called *feasible* if $\underline{P}(s, a, s') \leq p_{s,a}(s') \leq \overline{P}(s, a, s')$ for all $s' \in S$. The set of all feasible distributions for the state-action pair (s, a) is denoted by $\Gamma_{s,a}$.

A *path* of an IMDP is a sequence of states and actions $\omega = (s_0, a_0), (s_1, a_1), \dots$, where $(s_i, a_i) \in S \times A$ and $a_i \in \mathcal{A}(s_i)$. We denote by $\omega(k) = s_k$ the state of the path at time $k \in \mathbb{N}_0$ and by Ω the set of all paths. A *strategy* or *policy* for an IMDP is a function π that assigns an action to a given state of an IMDP. In `IntervalMDP.jl`, we focus on *time-dependent* strategies, i.e., $\pi : S \times \mathbb{N}_0 \rightarrow A$, as they suffice for optimality for reachability and discounted cumulative reward properties (Givan et al., 2000; Delimpaltadakis et al., 2023). In the special case that π does not depend on time and solely depends on the state, it is called a *stationary* strategy, i.e., $\pi(s, k) = \pi(s, k')$ for all $k, k' \in \mathbb{N}_0$ and all $s \in S$. Similar to a strategy, an adversary η is a function that assigns a feasible distribution to a given state (Givan et al., 2000). Given a strategy and an adversary, an IMDP collapses to a finite Markov chain.

In `IntervalMDP.jl`, we consider both reachability and discounted-reward properties. In the rest of this section, we focus on the reachability problem; the discounted reward case follows similarly. For reachability on IMDPs, we define a goal region $G \subseteq S$ and a time horizon $K \in \mathbb{N}_0 \cup \{\infty\}$, then the objective is to solve the following optimization problems:

$$\text{opt}_{\pi}^{\pi} \text{opt}_{\eta}^{\eta} \mathbb{P}_{\pi, \eta}[\omega \in \Omega \mid \exists k \in [0, K], \omega(k) \in G], \quad (1)$$

where $\text{opt}^{\pi}, \text{opt}^{\eta} \in \{\min, \max\}$ and $\mathbb{P}_{\pi, \eta}$ is the probability of the Markov chain induced by strategy π and adversary η . When $\text{opt}^{\eta} = \min$, Eqn. (1) is called optimal *pessimistic* probability (or reward), and conversely Eqn. (1)

is called optimal *optimistic* probability (or reward) when $\text{opt}^{\eta} = \max$. The choice of the min/max for the action and pessimistic/optimistic probability depends on the application. For instance, in robust strategy synthesis, the interest is often in maximizing pessimistic probability, i.e., $\text{opt}^{\pi} = \max$ and $\text{opt}^{\eta} = \min$.

Eqn. (1) can be computed by solving the following value iteration:

$$\begin{aligned} V_0(s) &= \mathbf{1}_G(s) \\ V_k(s) &= \mathbf{1}_G(s) + \\ &\quad \mathbf{1}_{S \setminus G}(s) \text{opt}_{a \in \mathcal{A}(s)}^{\pi} \text{opt}_{p_{s,a} \in \Gamma_{s,a}}^{\eta} \sum_{s' \in S} V_{k-1}(s') p_{s,a}(s'), \end{aligned} \quad (2)$$

where $\mathbf{1}_G(s) = \begin{cases} 1 & \text{if } s \in G \\ 0 & \text{otherwise} \end{cases}$ is the indicator function

for set G . `IntervalMDP.jl` solves Eqn. (2) iteratively over time via linear programming. In Section 4, we show how its solution can be efficiently parallelized to take advantage of GPU hardware architectures.

3. OVERVIEW OF INTERVALMDP.JL

`IntervalMDP.jl` is a Julia package that introduces parallelization and GPU-powered processing to perform value iteration in Eqn. (2), offering efficiency for verification and strategy synthesis for IMDPs. `IntervalMDP.jl` has the following main features:

- value iteration and strategy synthesis for all combinations of opt^{π} and opt^{η} ,
- dense and sparse matrix support,
- customizable numerical precision,
- multi-threaded CPU and CUDA-accelerated value iteration, and
- data loading and writing in various formats (e.g., PRISM, `bmdp-tool`, and `IntervalMDP.jl`).

In this section, we show how to create an IMDP model in `IntervalMDP.jl` and use `IntervalMDP.jl` for verification and synthesis. The source code can be found at <https://github.com/Zinoex/IntervalMDP.jl>.

3.1 System modeling

To create an IMDP, we can programmatically construct an `IntervalMarkovDecisionProcess` object and pass it to `IntervalMDP.jl`. Here, we include an example of how to construct a 3-state IMDP with the third state being a sink state.

```
prob1 = IntervalProbabilities(;
    lower = [0.0 0.5; 0.1 0.3; 0.2 0.1],
    upper = [0.5 0.7; 0.6 0.5; 0.7 0.3],
)
prob2 = IntervalProbabilities(;
    lower = [0.1 0.2; 0.2 0.3; 0.3 0.4],
    upper = [0.6 0.6; 0.5 0.5; 0.4 0.4],
)
prob3 = IntervalProbabilities(;
    lower = [0.0; 0.0; 1.0],
    upper = [0.0; 0.0; 1.0]
)
transition_probs = [
    ["a1", "a2"] => prob1,
    ["a1", "a2"] => prob2,
```

```

["sink"] => prob3
]
imdp = IntervalMarkovDecisionProcess(transition_probs)

```

Figure 1 shows a pictorial representation of the IMDP constructed with the code above. The interval transition probabilities are specified as a list `actions => transition_probs` where the upper and lower bounds are specified using dense matrices with columns representing the actions and rows the resulting state. `IntervalMDP.jl` also supports sparse matrices encoded in the Compressed Sparse Column (CSC) format (Buluç et al., 2009) for better memory usage and computational efficiency. Note that while in the example above the actions are given as strings, namely "a1", "a2", and "sink", the package supports any arbitrary type as actions, e.g., integers and floating point numbers.

To improve compatibility with other tools, `IntervalMDP.jl` supports reading and writing IMDPs stored in various formats. Currently, there are three supported standards for storing IMDPs: PRISM, `bmdp-tool`, and `IntervalMDP.jl` data formats. The following subsections describe how to read and write in the respective formats.

PRISM `IntervalMDP.jl` supports reading and writing PRISM (Kwiatkowska et al., 2011) explicit data format. The data format is split into 4 different files, one for the states 'sta', one for the labels 'lab', one for the transition probabilities 'tra', and one for the specification 'pctl'. Therefore, our interface for reading PRISM files takes the path without file ending and adds the appropriate ending to each of the four files.

```

# Read
problem = read_prism_file(path_without_file_ending)

# Write
write_prism_file(path_without_file_ending, problem)

```

The problem structure contains both the IMDP and the specification including whether to synthesize a maximizing or minimizing strategy and whether to use an optimistic or pessimistic adversary.

bmdp-tool `bmdp-tool` (Lahijanian, 2021) data format is similar to the PRISM explicit format transition probability files, where transition probabilities are stored line-by-line with source, action, destination, and probability bounds in ASCII. Key differences include no explicit listing of states and the fact that it only supports reachability properties. This format lacks information about whether the reachability is finite or infinite time, and hence the reader only returns the set of terminal states.

```

# Read
imdp, terminal_states = read_bmdp_tool_file(path)

# Write
write_bmdp_tool_file(path, problem)

```

IntervalMDP.jl A disadvantage of both PRISM and `bmdp-tool` data formats is that they store the data in ASCII where each character requires one byte of storage space. As a result, the data format requires more




Fig. 1. An example of a 3-state IMDP where state q_3 is a sink state.

storage space than necessary for floating point data. `IntervalMDP.jl` supports NetCDF (Rew et al., 2006) to store transition probabilities efficiently. NetCDF is a multi-dimensional array-centric format with readers and writers in many languages and includes lossless compression methods for even more efficient data storage. We use JSON to store the specification, as storage space for the specification is much less a concern, and because JSON is a widely used, human-readable, file format.

```

# Read
imdp = read_intervalmdp_jl_model(model_path)
spec = read_intervalmdp_jl_spec(spec_path)
problem = Problem(imdp, spec)

problem = read_intervalmdp_jl(model_path, spec_path)

# Write
write_intervalmdp_jl_model(model_path, imdp_or_problem)
write_intervalmdp_jl_spec(spec_path, spec_or_problem)

```

Without compression, this data format requires 50% less space compared to both ASCII formats.

3.2 Specification and verification

The tool minimizes or maximizes the optimistic or pessimistic probability of reaching a given set of states, or optimizes a reward. Continuing on the example from Section 3.1, we show how to compute the maximum pessimistic probability of reaching state 3 within 100 time steps.

```

prop = FiniteTimeReachability([3], 100)
spec = Specification(prop, Pessimistic, Maximize)
problem = Problem(imdp, spec)

# Compute the value function
V, k, residual = value_iteration(problem)

```

Function `value_iteration` returns a 3-tuple containing the following: the robust optimal probability (in the above example, maximum pessimistic probability) of satisfying the reachability property, the number of iterations performed (100 in the above example), and finally, the Bellman residual for the last iteration. That is, the state-wise difference between the value function for the last and the second to last iteration. For a finite horizon property, the number of iterations is fixed, while for an infinite horizon property, a user may specify a tolerance for convergence, measured by the maximum Bellman residual.

Instead, to compute robust optimal discounted cumulative reward property, `IntervalMDP.jl` can be used as follows:

```
rewards = [1.0, 2.0, 3.0]
discount = 0.95
prop = FiniteTimeReward(rewards, discount, 100)
```

The length of the list specifying the rewards must match the number of states in the IMDP and the discount must be between 0 and 1. Then the specification and problem construction follow as with a reachability property.

Using the library, we may also synthesize the optimal policy.

```
# Compute optimal finite-time policy
prop = FiniteTimeReachability([3], 100)
spec = Specification(prop, Pessimistic, Maximize)
problem = Problem(imdp, spec)
time_dependent_policy = control_synthesis(problem)

# Compute optimal infinite-time policy
prop = InfiniteTimeReachability([3], 1e-6)
spec = Specification(prop, Pessimistic, Maximize)
problem = Problem(imdp, spec)
stationary_policy = control_synthesis(problem)
```

In the first case, with a finite time horizon property, the optimal strategy is time dependent, and thus `IntervalMDP.jl` returns a matrix of size $|S| \times K$ where K is the time horizon, whose entry in the i -th row and j -th column is the action for state i at time step j . In the second case, the optimal strategy is stationary and `IntervalMDP.jl` returns a vector of length $|S|$, where the i -th entry is the action for state i .

A core feature of `IntervalMDP.jl` is GPU-accelerated algorithms. To enable the use of CUDA, we only need to transfer the problem to the GPU. Note that GPUs generally have less memory than what is available to CPUs, and thus it is highly recommended the use of the sparse format for GPUs.

For the example above, the following code enables the use of a GPU, if available:

```
using CUDA

prop = FiniteTimeReachability([3], 100)
spec = Specification(prop, Pessimistic, Maximize)
problem = Problem(imdp, spec)

if CUDA.functional()
    problem = IntervalMDP.cu(problem)
end

V, k, residual = value_iteration(problem)
```

This trivial modification allows improvements in terms of computational time of 50-200 \times as is illustrated in Section 5. Note that `IntervalMDP.cu` is opinionated to `Float64` values and `Int32` indices to balance numerical errors and performance. However, other options for value and index types are available via the `Adapt.jl` package.




Fig. 2. Bitonic sorting network structure in $\log_2(n)$ major rounds, each of which is up to $\log_2(n)$ minor rounds. Each arrow represents a comparison and points towards the larger element. After each major round, subsets of increasing size are bitonic (increasing order, then decreasing of the same size).

4. VALUE ITERATION ON CUDA-CAPABLE GPUS

In this section, we discuss our algorithmic approach to solve Eqn. (1) on CUDA-capable hardware. For the sake of simplicity of presentation, we focus only on the maximum pessimistic probability ($\text{opt}^\pi = \max$ and $\text{opt}^\eta = \min$ in Eqn. (1)) case; the other cases follow similarly. Recall from Eqn. (1) that computing the optimal pessimistic reward reduces to iteratively solving the following problem for each $s \in S$ and for each k

$$\max_{a \in \mathcal{A}(s)} \min_{p_{s,a} \in \Gamma_{s,a}} \sum_{s' \in S} V_k(s') p_{s,a}(s'). \quad (3)$$

As Eqn. (3) is solved independently for each state s , its solution at different states can be trivially parallelized. However, in what follows we show how in `IntervalMDP.jl`, the solution of robust value iteration is parallelized further.

The inner problem in Eqn. (3) is a linear problem that must be solved for every $a \in \mathcal{A}(s)$. We solve this problem using the O-maximization procedure introduced in (Givan et al., 2000; Lahijanian et al., 2015). The O-maximization procedure is divided into two phases: ordering of states based on the value of V_k and then assigning the least probability mass possible to states with high V_k . Both phases can benefit from parallelization on a GPU. For ordering the states based on the value of the value function at the current time step V_k , we apply existing parallel algorithms that are particularly suited to run efficiently on a GPU. In particular, in `IntervalMDP.jl`, we use bitonic sort (Peters et al., 2010), which is a parallel sorting algorithm with $O(\log_2(n)^2)$ latency¹. Figure 2 shows an example of how the sorting is performed when n is a power of 2. A bitonic sorting treats the list as bitonic subsets that gets merged over $\log_2(n)$ major rounds, each of which consists of up to $\log_2(n)$ minor rounds to preserve the bitonic property. After $\log_2(n)$ major rounds, the set is the first half of a bitonic set, which implies that it is sorted.

The second phase consists of assigning probabilities within the feasible probability intervals to give the most proba-

¹ Note that when assessing parallel algorithms, the asymptotic performance is measured by the latency, which is the delay in the number of parallel operations, before the result is available. This is in contrast to traditional algorithms, which are assessed by the total number of operations.




Fig. 3. Cumulative sum as a tree-reduction. This tree reduction has a latency of $O(\log_2(n))$ as opposed to $O(n)$ of a traditional iterative algorithm.

bility mass possible to states early in the ordering. To perform this assignment step in parallel, we have to parallelize the sequential algorithm below.

```
function omaximization(ordering, lower, upper)
    p = copy(lower)
    rem = 1 - sum(lower)
    gap = upper - lower

    for o in ordering
        if gap[o] < rem
            p[o] += gap[o]
            rem -= gap[o]
        else
            p[o] += rem
            break;
        end
    end

    return p
end
```

While seemingly sequential in nature, the method implicitly implements a cumulative summation procedure where the summation is according to the ordering previously found. Consequently, we can employ the parallel algorithm based on tree reduction developed in (Ladner and Fischer, 1980) (see Figure 3), which has latency $O(\log_2(n))$. In particular, for O-maximization, the cumulative sum is computed over the gap between the upper and lower bounds according to the ordering. Thus, each element will have the sum of the gaps of the states up to and including itself in the ordering. Then each iteration in O-maximization is independent and thus can be performed in parallel, that is, with a latency of $O(1)$.

```
function omaximization(ordering, lower, upper)
    p = copy(lower)
    rem = 1 - sum(lower)
    gap = upper - lower

    # Ordered cumulative sum of gaps
    cumgap = cumulative_sum(gap[ordering])
```

```
for (i, o) in enumerate(ordering)
    rem_state = max(rem - cumgap[i] + gap[o], 0)
    if gap[o] < rem_state
        p[o] += gap[o]
    else
        p[o] += rem_state
    end
end

return p
end
```

5. COMPUTATIONAL STUDIES


In order to show the effectiveness of `IntervalMDP.jl`, we compare it against `bmdp-tool` and `PRISM`, which, to the best of our knowledge, are the only existing tools that support value iteration for IMDPs. We benchmark the tools on 35 models that are taken from the literature and include abstractions of linear and non-linear systems, including neural networks, with number of transition probabilities ranging from few tens to tens of millions.

For a fair comparison, for all models we run a (maximum pessimistic) finite time reachability query (200 time steps). All experiments were conducted on a computer with 16GB RAM, an Intel I7-6700K CPU (4 cores, 8 hyperthreads), and an NVIDIA GTX1060 6GB VRAM GPU. For each tool, we measure the computation time only, and not the time it takes to load each file.


The obtained results are shown in Figure 4 as plots of computation time as a function of the number of transitions in the IMDP. In Figure 4a, we see that `IntervalMDP.jl` substantially outperforms the other tools in terms of computation time. For instance, computing the query for the largest model takes `bmdp-tool` 6865s and `PRISM` 1235s, while `IntervalMDP.jl` takes 372s and 30s for the CPU and GPU implementation respectively. This is a speed-up of 228× and 41× of the GPU implementation of `IntervalMDP.jl` relative to `bmdp-tool` and `PRISM` respectively.

The log-log plot in Figure 4b highlights the performance differences on the smaller models. Specifically, the `bmdp-tool` and the CPU implementation of `IntervalMDP.jl` are faster than `PRISM` on the smallest models, with a speed-up of 30-60×. We conjecture that the better scaling of `PRISM` compared to `bmdp-tool` is due to using a dictionary of non-zero probability destinations for each source-action pair stored in a list, rather than a dictionary keyed by source/action/destination triplets. `IntervalMDP.jl`, on the other hand, does not store any dictionary. Instead, we track and store the indices for source, action, and destination in the CSC-format sequentially, allowing better caching when accessing probabilities. In Figure 4b, we also see that the GPU implementation of `IntervalMDP.jl` has overhead that is dominant for smaller models, while for larger models is consistently orders of magnitude faster than the other implementations. The cut-off is roughly at 25000 transitions in the IMDP.

A common bottleneck for IMDP tools is memory consumption, which is only exacerbated on a GPU, as they generally have less memory available. However, due to the CSC-format with `Float64` values and `Int32` indices,



(a) Linear scale with trendline



(b) Log-log scale

Fig. 4. Computation time of `IntervalMDP.jl` compared against `bmdp-tool` and `PRISM` on 35 IMDPs of varying sizes.

`IntervalMDP.jl` generally requires less memory compared to `PRISM` and `bmdp-tool`. For example, to run value iteration on the largest model, `IntervalMDP.jl` requires 4.88GB of memory. In contrast, `PRISM` uses 6.32GB of memory and `bmdp-tool` uses 5.38GB to run value iteration on the same problem. This is a 23% reduction relative to `PRISM` and 9% relative to `bmdp-tool`, which is including the Julia runtime.

6. CONCLUSION

We introduced `IntervalMDP.jl`, a new software package for robust value iteration of IMDPs that employs parallel algorithms and takes advantage of CUDA-accelerated hardware. Computational experiments showed speed-ups of 50-200 \times relative to state-of-the-art. `IntervalMDP.jl` is currently limited to reachability specifications along with reward optimization. We plan to extend it to support Linear Temporal Logic (Baier and Katoen, 2008), via the product construction (Abate et al., 2011), and Probabilistic Computation Tree Logic (Hansson and Jonsson, 1994), via parse tree construction (Lahijanian et al., 2015).

REFERENCES

Abate, A., Katoen, J.P., and Mereacre, A. (2011). Quantitative automata model checking of autonomous stochastic hybrid systems. In *Proceedings of the 14th international conference on Hybrid systems: computation and control*, 83–92.

Baier, C. and Katoen, J.P. (2008). *Principles of Model Checking*. The MIT Press.

Bezanson, J., Karpinski, S., Shah, V.B., and Edelman, A. (2012). Julia: A fast dynamic language for technical computing. *arXiv preprint arXiv:1209.5145*.

Buluç, A., Fineman, J.T., Frigo, M., Gilbert, J.R., and Leiserson, C.E. (2009). Parallel sparse matrix-vector and matrix-transpose-vector multiplication using compressed sparse blocks. In *Proceedings of the twenty-first annual symposium on Parallelism in algorithms and architectures*, 233–244.

Cauchi, N., Laurenti, L., Lahijanian, M., Abate, A., Kwiatkowska, M., and Cardelli, L. (2019). Efficiency through uncertainty: Scalable formal synthesis for stochastic hybrid systems. In *Proceedings of the 22nd ACM international conference on hybrid systems: computation and control*, 240–251.

Delimpaltadakis, G., Lahijanian, M., Mazo Jr, M., and Laurenti, L. (2023). Interval markov decision processes with continuous action-spaces. In *Proceedings of the 26th ACM International Conference on Hybrid Systems: Computation and Control*, 1–10.

Givan, R., Leach, S., and Dean, T. (2000). Bounded-parameter markov decision processes. *Artificial Intelligence*, 122(1-2), 71–109.

Haddad, S. and Monmege, B. (2018). Interval iteration algorithm for mdps and imdps. *Theoretical Computer Science*, 735, 111–131.

Hansson, H. and Jonsson, B. (1994). A logic for reasoning about time and reliability. *Formal aspects of computing*, 6, 512–535.

Kwiatkowska, M., Norman, G., and Parker, D. (2011). Prism 4.0: Verification of probabilistic real-time systems. In *Computer Aided Verification: 23rd International Conference*, 585–591. Springer.

Ladner, R.E. and Fischer, M.J. (1980). Parallel prefix computation. *Journal of the ACM (JACM)*, 27(4), 831–838.

Lahijanian, M. (2021). `bmdp-tool`. URL <https://github.com/aria-systems-group/bmdp-tool>. Standalone software.

Lahijanian, M., Andersson, S.B., and Belta, C. (2015). Formal verification and synthesis for discrete-time stochastic systems. *IEEE Transactions on Automatic Control*, 60(8), 2031–2045.

Laurenti, L. and Lahijanian, M. (2023). Unifying safety approaches for stochastic systems: From barrier functions to uncertain abstractions via dynamic programming. *arXiv preprint arXiv:2310.01802*.

Lavaei, A., Soudjani, S., Abate, A., and Zamani, M. (2022). Automated verification and synthesis of stochastic hybrid systems: A survey. *Automatica*, 146, 110617.

NVIDIA Corporation (2023). NVIDIA CUDA C++ programming guide. Version 12.3.

Peters, H., Schulz-Hildebrandt, O., and Luttenberger, N. (2010). Fast in-place sorting with cuda based on bitonic sort. In *Parallel Processing and Applied Mathematics: 8th International Conference*, 403–410. Springer.

Rew, R., Hartnett, E., Caron, J., et al. (2006). Netcdf-4: Software implementing an enhanced data model for the geosciences. In *22nd International Conference on Interactive Information Processing Systems for Meteorology, Oceanography, and Hydrology*.