Scalable Multi-Agent Simulation based on MapReduce

Tobias Ahlbrecht, Jürgen Dix, and Niklas Fiekas

Department of Informatics
Clausthal University of Technology
Julius-Albert-Str. 4 D-38678 Clausthal-Zellerfeld, Germany
{tobias.ahlbrecht✉,dix,niklas.fiekas}@tu-clausthal.de

Abstract. Jason is perhaps the most advanced multi-agent programming language based on *AgentSpeak*. Unfortunately, its current Java-based implementation does not scale up and is seriously limited for simulating systems of hundreds of thousands of agents.
We are presenting a scalable simulation platform for running huge numbers of agents in a Jason style simulation framework. Our idea is (1) to identify independent parts of the simulation in order to parallelize as much as possible, and (2) to use and apply existing technology for parallel processing of large datasets (e.g. MapReduce).
We evaluate our approach on a simple benchmark and show that it scales up linearly (in the number of agents).

1 Introduction

This is the final report for a research project started in the winter semester 2016/17. The project was mostly completed in that term and a paper with the results [Ahlbrecht et al., 2016a] was presented at the EUMAS 2016 conference. The developed platform has since been used in the SWZ project SimSE for agent-based simulation of software development processes [Ahlbrecht et al., 2016b], published as an open source project [1] and successfully used in the Multi-Agent Programming Contest [Pieper, 2017]. The present paper is based on the EUMAS 2016 version and extends it with further evaluation results.

The goal is to simulate a huge number of agents, (hundreds of thousands or even more). Current approaches implemented in Java often do not scale up (see [Ahlbrecht et al., 2016c] for a detailed discussion). Similarly, declarative approaches (e.g. those based on *AgentSpeak*) are well suited for modeling simulations, but do not support efficient implementation.

Here we focus on a new approach for implementing scalable multi-agent simulation platforms with MapReduce. The main idea is to identify parts of the simulated environment that are completely independent from each other and can thus be processed in parallel. This is particularly useful in scenarios based

[1] Source code available at https://github.com/niklasf/pyson

on large existing datasets, but can also be applied to multi-agent simulation in general.

In the following we give a very brief introduction to Jason and MapReduce and comment on related work. The main part is Section 3, where we show how Jason can be interpreted in a way that is compatible with MapReduce. While previous approaches have used limited agent models [Radenski, 2013] or restricted languages [Wang et al., 2010] our approach supports full Jason-style *AgentSpeak*. We believe that similar agent languages can be translated accordingly.

Key points of any simulation are (1) modeling and (2) implementing the environment: we elaborate on both in Section 4. Finally we evaluate our approach in Section 5 using a benchmark for our early proof of concept implementation[2] and conclude with Section 6.

2 Related work

2.1 Jason

Jason is a Java based platform for multi-agent simulation with an extended version of *AgentSpeak* [Bordini et al., 2007]. *AgentSpeak* is a language to describe BDI agents that mixes a declarative approach to reasoning (Prolog) and an imperative way of stating plans [Rao, 1996]. Jason extends the language with useful functionality such as agent communication. Jason is widely used [Bordini and Dix, 2013] but does not scale well when the simulation size is increased beyond thousands of agents, even when the agents are very simple.

2.2 MapReduce

MapReduce is a programming paradigm designed to simplify the parallel processing of large datasets [Dean and Ghemawat, 2008] by abstracting away low level architecture (single thread, multi-core computer, grid of commodity computers), synchronization, error recovery, locking and distribution of work among the nodes of a cluster. The algorithm is defined in terms of *map* and *reduce* functions that operate on key value pairs. Map functions operate independently on key value pairs $\langle k, v \rangle$. After a shuffling step that groups items by their keys, *reduce* functions operate on sequences of values in each group:

$$\text{Map} : (K, V) \to (K, V)^* \quad ; \quad \text{Reduce} : (K, V^*) \to (K, V)^*$$

Algorithms in terms of these functions can be executed using a MapReduce framework like Spark, Hadoop, MR4C, MapReduce-MPI or Disco, which automatically partition the dataset for parallel execution.

[2] Source code available at `https://github.com/niklasf/pyson`

2.3 Simulation with MapReduce

There are several design patterns for MapReduce that have been used outside of agent simulation. Lin and Schatz [2010] reinterpret classical graph algorithms as communication along the edges of graphs. Combined with some optimizations this yields significant speedups over previous implementations.

Many simulations deal with objects in a simulated physical world. If objects can only influence their neighborhood the simulation can be interpreted as a series of spatial joins. Zhang et al. [2009] provide a technique for parallelizing spatial joins using MapReduce.

These approaches have then been used data analysis and classical simulation, but more recently also in agent system simulation with agent models that have been restricted accordingly: Radenski [2013] uses graph algorithms to simulate cellular automatons. Wang et al. [2010] use spatial joins for behavioral simulations, where agent actions are restricted to associative operations on the environment. Here we try to lift these restrictions by simulating arbitrary Jason agents.

3 Translating Jason to MapReduce

When agents deliberate but do not communicate or execute actions in the environment they can be executed independently in Map steps. In this section we discuss key requirements for a Jason interpreter that allows doing that. Then the state of agents and the state of the environment can be represented in *key value pairs* such that *actions that advance the simulation can be performed efficiently with Map and Reduce steps.*

Most MapReduce platforms commit datasets to disk after each MapReduce step. However this overhead can be avoided for multi-agent simulation: In case of data loss computation steps can simply be repeated. We therefore choose Apache Spark as our underlying platform. Spark features the concept of *Resilient Distributed Datasets* with configurable levels or persistence. Additionally, Spark uses the scripting language Python as one of the primary supported languages. This allows us to use Python as a single language for the platform as well as for scripting the simulated environment and available actions. There are three key requirements for the Jason interpreter:

- **Serializability**: The state of agents must be serializable at any given time to allow Spark to serialize and transmit them to other nodes of the cluster.
- **Ability to pause and resume individual agents**: In distributed computing local operations are near-instant while network operations take orders of magnitudes more time. An agent waiting for data from the network needs to be paused in order not to block the execution of other agents.
- **Memory efficiency**: The interpreter must have a low memory footprint so that hundreds of thousands of agents can fit into main memory.

For memory efficiency we embed native Python data types directly into Jason (`bool`, `int` and `float` for numerics, `tuple` for lists). Variables and belief literals are defined as classes in Python (`Var()` and `Literal(functor, args)`). All other Python objects are treated as atoms. To avoid making copies of objects, all substitutions (mappings of variables to terms) are kept in a separate dictionary. Additionally, agents have a stack of substitutions and choice points that allows them to undo failed partial unifications.[3]

To allow pausing and resuming individual agents (even while they are executing a Prolog query) we use Python generators to iterate over alternatives, with a technique similar to YieldProlog[4]. Finally choosing the Python implementation PyPy guarantees serializability of Python objects including functions, closures and generators.

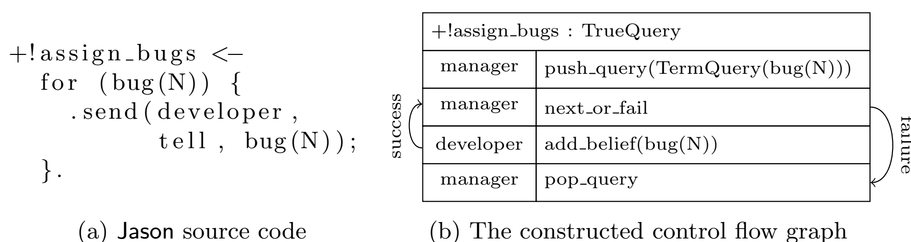

(a) Jason source code (b) The constructed control flow graph

Fig. 1: Example: A `manager` agent sends bug details to a `developer` agent

For AgentSpeak(L) the control flow in a plan is linear. Jason defines additional control structures such as branches and loops. To capture both we represent plans as a control flow graph where nodes are high level instructions. Each node has at most two outgoing edges labeled `success` or `failure` that are followed depending on the result of the current instruction. If a node does not have the corresponding edge this is interpreted as plan achievement (success) or plan failure respectively.

Intentions in *AgentSpeak* are defined as a stack of partially instantiated plans [Rao, 1996]. To avoid copying plans for each instantiation we use a separate intention data structure instead. The data structure contains (i) the instantiated plan head from the point of view of the caller, (ii) a pointer to the current instruction in the control flow graph, (iii) the current substitution `scope` (mapping of variables to terms), (iv) stacks to undo unifications and continue with a different choice (`stack`, `query_stack`, `choicepoint_stack`).

These instructions are used as an intermediate representation of Jason programs:

noop(agent, intention) Does nothing and succeeds always.

[3] This technique is well known in Prolog interpreters [Wielemaker et al., 2012, Winikoff, 1996].

[4] http://yieldprolog.sourceforge.net/

add_belief(term, agent, intention) Applies the current substitution to `term` and adds it to the belief base. Triggers a belief addition event.

remove_belief(term, agent, intention) Unifies `term` with the first matching belief and removes it from the belief base. Triggers a belief removal event.

test_belief(term, agent, intention) Tries to find a substitution such that `term` is a logical consequence of the belief base. Triggers a belief test event.

call(trigger, goal_type, term, agent, intention) Tries to find a plan matching `trigger`, `goal_type` and `term` and adds it as a subplan to the current intention.

call_delayed(trigger, goal_type, term, agent, intention) Tries to find a plan matching `trigger`, `goal_type` and `term` and creates a new intention with it.

push_query(query, agent, intention) Starts the Prolog query `query` and adds the resulting Python generator to the query stack. This is also used for actions that can yield multiple results.

next_or_fail(agent, intention) Tries to advance the topmost generator.

pop_query(agent, intention) Removes the topmost generator from the stack.

Observation 1 (Correct-, and Completeness) *The described interpreter satisfies the hard requirements outlined above. In addition, all Jason programs can be transformed to programs in our instruction set.*

4 Handling the Environment

To simulate the environment, a number of different object types have to be modeled. Possible *actions* and *percepts* make up a major part, as they imply the environment's behavior and thus determine the computational effort. Environments need a notion for each "thing" that is not an agent: we call it *artifact*.

The entire state of the simulation is stored in key value pairs. It comprises the agents $\langle uuid, agent\rangle$ and artifacts from the environment. A cycle of the simulation starts with a map step where each agent state is mapped to the next. Messages to other agents are emitted as key value pairs using a Jason-style belief annotation for the sender: $\langle recipientUuid, message[source(senderUuid)]\rangle$. Actions selected by the agent emit additional key values pairs (usually of the form $\langle affectedArtifactUuid, action\rangle$).

The actual effects of the actions are computed in a reduce phase where key value pairs are grouped by recipient or affected artifact. Reduce operations in Spark must be associative. Additionally commutativity is a reasonable requirement to achieve deterministic results even when the order of the values is nondeterministic. Actions that return results must include the UUID of the agent so that results can be emitted as a key value pair $\langle uuid, resultMessage\rangle$.

Values for distinct keys are reduced in parallel. This leads directly to the following observation.

Observation 2 *The environment needs to be designed such that potentially conflicting actions always affect the same key.*

While this can be trivially achieved by using a monolithic environment with a single key, it is likely that the reduction for that key will be a bottleneck. Thus, to allow parallel execution, we need the following complementary goal.

Observation 3 *Independent actions* should *affect distinct keys.*

For many scenarios there is a natural way to decompose the environment into key value pairs. For example Wang et al. [2010] partition a spatial environment into overlapping areas to simulate social force. Since areas overlap, the same action (effects) may be sent to multiple keys. Summation of forces is used as an associative and commutative reduce operation. However, as not all simulations decompose spatially (see the *Simulating Software Evolution* scenario) we propose the following additions:

- Instead of hardcoding the concept of spatial location we introduce groups that agents can subscribe to and send multicast messages to. This mechanism will also be exploited for percept generation and distribution.
- Deterministic reservoir sampling [Vitter, 1985] as an associative and commutative operation to fairly select one of multiple conflicting actions. This works for arbitrary actions since they no longer have to be associative and/or commutative themselves.

Currently, the whole environment has to be hand-coded as a Python script. The next step is to provide a thin wrapper around Spark to abstract away from its concrete functionality so as not to burden the user with having to learn everything about MapReduce in order to use the platform. In a later step, the final environment metamodel will be combined with our already existing Jason metamodel to provide the user with schematic modeling facilities (i.e. diagramming) to enable kick-starting new projects.

5 Evaluation

5.1 Application: Simulating Software Evolution

As mentioned before, the platform is part of a bigger project on simulating software development processes using agent-based technology to gain insights on (specific) software evolution. In this scenario, agents can perform abstract modifications on the software project, i.e. "fix bugs" or "refactor methods". Representing the developers with simple agents in *Repast Simphony* already proved a viable solution [Honsel et al., 2015].

To evaluate the modelling capabilities of the platform we reimplemented scenarios from this project (Bug lifespan from [Honsel et al., 2015] and refactorings from [Honsel et al., 2017]). Implementation effort is similar and our graph transformations map easily to Prolog queries. However our platform does not yet have built-in tools to visualize data from running simulations, which is useful for debugging.

It should be noted that we are not yet taking advantage of the advanced reasoning and planning capabilities of Jason. We believe that implementing this will enable the agents to adopt goal-oriented behavior, e.g. based on code change patterns, and ultimately yield more detailed simulation results. Furthermore, beliefs will be crucial to simulate how the agents gain experience in the process (see [Ahlbrecht et al., 2016b]).

5.2 Throughput of the interpreter

The running time of complex simulations depends heavily on the specific scenario, how well it can be parallelized in principle and even on the concrete implementation on any given agent platform. We resort to to a simple benchmark developed by Ahlbrecht et al. [2016c] and compare the performance of our platform running on different Python interpreters (Python 2, Python 3, PyPy) with the performance of other platforms (Jason, *Maserati*). While it only really measures the throughput of the interpreter on a single node (relating to the implementation described in Section 3), it already differentiates our platform from previous approaches.

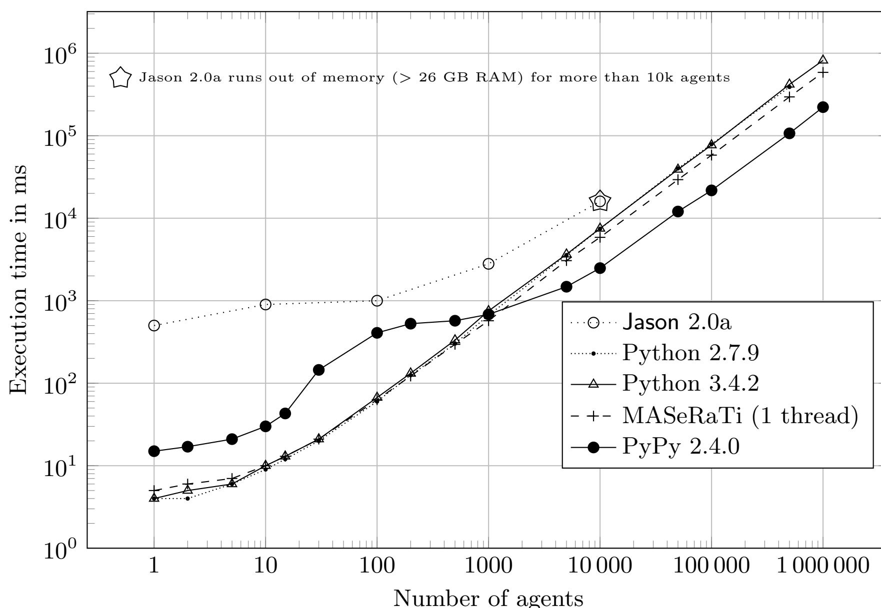

Fig. 2: Execution times of the counting scenario for increasing numbers of agents

The test environment is a freshly installed Debian Jessie using an Intel Xeon CPU @ 4 x 2.30 GHz and 26 GB RAM (n1-highmem-4 on Google Compute engine). Besides a terminal session via SSH and standard background processes

no other programs were running. Graphed times are averages over 10 runs for each point.

Jason 2.0a runs out of memory for 50 000 agents, but could potentially complete the simulation on a machine with even more RAM. The other platforms scale roughly linearly as expected for this simple scenario. We achieve the best performance with PyPy which uses *Just-In-Time compilation* and *hotspot optimization* (see the disproportional speedup for a medium number of agents).

Albeit beeing JIT the optimizations are reliable. The highest standard deviation encountered in the benchmarks was 2.14s relative to an average execution time of 106.821s (for 500 000 agents).

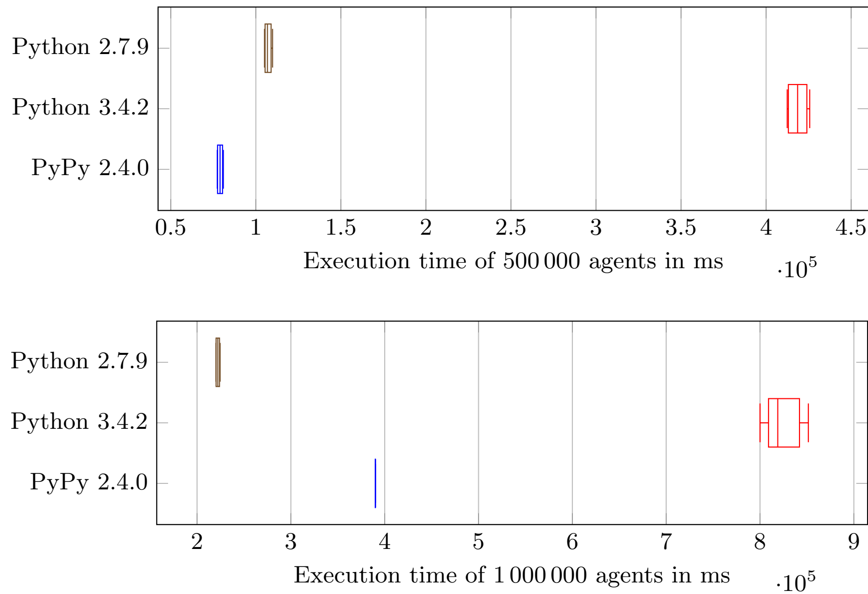

Fig. 3: Range of measurements with Python 2, Python 3 and PyPI with JIT compilation

6 Conclusion

We have presented a scalable Jason interpreter that is part of a bigger project on quality control of software development processes (see [Ahlbrecht et al., 2016b]). However, we believe our approach is rather general and can be applied to similar agent languages based roughly on *AgentSpeak* (which allows us to use the built-in modelling constructs). All that needs to be done is to find a suitable translation of this language into MapReduce (as described in Section 3). An advantage of our approach is the possibility to use off-the-shelf professional tools to deal with MapReduce.

Our evaluation shows linear scalability (in the number of agents) in a simple benchmark, even for a reimplementation of Jason. In more realistic simulations this will only be the best case and it remains to test other benchmarks and to tailor our system for the application in the planned project. But we are planning to apply our approach also to other areas, where parallelization in the simulation of an environment pays off.

References

Tobias Ahlbrecht, Jürgen Dix, and Niklas Fiekas. Scalable multi-agent simulation based on mapreduce. In *Multi-Agent Systems and Agreement Technologies*, pages 364–371. Springer, 2016a.

Tobias Ahlbrecht, Jürgen Dix, Niklas Fiekas, Jens Grabowski, Verena Herbold, Daniel Honsel, Stephan Waack, and Marlon Welter. Agent-based simulation for software development processes. Technical Report IfI-16-02, TU Clausthal, September 2016b. URL `http://www.in.tu-clausthal.de/fileadmin/homes/techreports/ifi1602ahlbrecht.pdf`.

Tobias Ahlbrecht, Jürgen Dix, Niklas Fiekas Philipp Kraus, and Jörg P. Müller. An architecture for scalable simulation of systems of cognitive agents. *International Journal of Agent-Oriented Software Engineering*, 2016c.

Rafael Bordini and Jürgen Dix. Chapter 13: Programming multi-agent systems. In Gerhard Weiss, editor, *Multiagent systems*, pages 587–639. MIT-Press, 2013.

Rafael H. Bordini, Jomi Fred Hübner, and Michael Wooldridge. *Programming Multi-Agent Systems in AgentSpeak Using Jason (Wiley Series in Agent Technology)*. John Wiley & Sons, 2007. ISBN 0470029005.

Jeffrey Dean and Sanjay Ghemawat. MapReduce: Simplified data processing on large clusters. *Commun. ACM*, 51(1):107–113, January 2008. ISSN 0001-0782. doi: 10.1145/1327452.1327492. URL `http://doi.acm.org/10.1145/1327452.1327492`.

Daniel Honsel, Niklas Fiekas, Verena Herbold, Marlon Welter, Tobias Ahlbrecht, Stephan Waak, Jürgen Dix, and Jens Grabowski. Simulating software refactorings based on graph transformations. 2017. (To appear in SimScience).

Verena Honsel, Daniel Honsel, Steffen Herbold, Jens Grabowski, and Stephan Waack. Mining software dependency networks for agent-based simulation of software evolution. In *2015 30th IEEE/ACM International Conference on Automated Software Engineering Workshop (ASEW)*, pages 102–108. IEEE, 2015.

Jimmy Lin and Michael Schatz. Design patterns for efficient graph algorithms in MapReduce. In *Proceedings of the Eighth Workshop on Mining and Learning with Graphs*, MLG '10, pages 78–85, New York, NY, USA, 2010. ACM. ISBN 978-1-4503-0214-2. doi: 10.1145/1830252.1830263. URL `http://doi.acm.org/10.1145/1830252.1830263`.

Jonathan Pieper. Entwicklung eines BDI basierten Agententeams für den Multi-Agent Programming Contest, September 2017. Bachelorarbeit.

Atanas Radenski. Using MapReduce streaming for distributed life simulation on the cloud. ECAL, 284-291(2013), 2013.

Anand S. Rao. AgentSpeak(L): BDI agents speak out in a logical computable language. In *Proceedings of the 7th European Workshop on Modelling Autonomous Agents in a Multi-agent World: Agents Breaking Away*, MAAMAW '96, pages 42–55, Secaucus, NJ, USA, 1996. Springer-Verlag New York, Inc. ISBN 3-540-60852-4. URL `http://dl.acm.org/citation.cfm?id=237945.237953`.

Jeffrey S Vitter. Random sampling with a reservoir. *ACM Transactions on Mathematical Software (TOMS)*, 11(1):37–57, 1985.

Guozhang Wang, Marcos Antonio Vaz Salles, Benjamin Sowell, Xun Wang, Tuan Cao, Alan J. Demers, Johannes Gehrke, and Walker M. White. Behavioral simulations in MapReduce. *CoRR*, abs/1005.3773, 2010. URL `http://arxiv.org/abs/1005.3773`.

Jan Wielemaker, Tom Schrijvers, Markus Triska, and Torbjörn Lager. SWI-Prolog. *Theory and Practice of Logic Programming*, 12(1-2):67–96, 2012.

Michael Winikoff. W-Prolog, 1996. URL `http://waitaki.otago.ac.nz/~michael/wp/index.html`.

Shubin Zhang, Jizhong Han, Zhiyong Liu, Kai Wang, and Zhiyong Xu. SJMR: parallelizing spatial join with MapReduce on clusters. In *Proceedings of the 2009 IEEE International Conference on Cluster Computing, August 31 - September 4, 2009, New Orleans, Louisiana, USA*, pages 1–8, 2009. doi: 10.1109/CLUSTR.2009.5289178. URL `http://dx.doi.org/10.1109/CLUSTR.2009.5289178`.