

1 MASP-Reduce: a proposal for distributed 2 computation of stable models

3 **Federico Igne**

4 University of Udine and New Mexico State University

5 ignefederico@gmail.com

6 **Agostino Dovier**

7 University of Udine, Udine, Italy

8 agostino.dovier@uniud.it

9 **Enrico Pontelli**

10 New Mexico State University, NM, USA

11 epontell@cs.nmsu.edu

12 — Abstract —

13 The availability of efficient answer set solvers (e.g., CLASP and its descendants [8, 2]) gave Answer
14 set programming (ASP) a leading role in languages for knowledge representation and reasoning.
15 The simple syntax is surely one of the main strengths of the paradigm; moreover the stable models
16 semantics intuitively resembles the human reasoning process in a clean and *logical* way. ASP is
17 regarded as the computational embodiment of non-monotonic reasoning because of its simple
18 syntax and elegant non-monotonic semantics. The popularity of ASP is demonstrated by the
19 increasing number of authors publishing ASP-based research work in artificial intelligence as well
20 as non-logic programming venues, and its use as a natural alternative to other paradigms (e.g.,
21 SAT solving). Most of the answer set solvers are currently developed as two-phases procedures
22 (save some exceptions — e.g., [3, 11]) . The first stage is called *grounding* and computes the
23 equivalent propositional logic program of an input logic program, instantiating each rule over
24 the domain of its variables. Modern solvers also apply some simplifications and heuristics to
25 the program, in order to ease the computation during the second stage. The computation of
26 the answer sets of a logic program is carried out by the *solving* stage, which also deals with the
27 non-deterministic reasoning involved in the model.

28 ASP encoding of sophisticated applications in real-world domains (e.g., planning, phylogenetic
29 inference) highlighted the strengths and weaknesses of this paradigm. Most of the times, the
30 technology underlying the ASP solvers, lacks the ability to keep up with the demand of complex
31 applications. This has been, for example, highlighted in a study on the use of ASP to address
32 complex planning problems [13, 5, 6]. With respect to these studies, it is clear that one of the
33 main limitations of this paradigm resides in the grounding process and the ability to compute the
34 stable models of large ground programs. This limitation is even more obvious when the whole
35 computation is performed in-memory.

36 This work tries to partially solve the problem of processing large ground programs that can
37 exceed capabilities for in-memory computation—using parallelism and distributed computing.
38 We aim to study, analyse and develop a fully distributed answer set solver and use a distributed
39 environment to efficiently represent and reason over large programs whose grounding would be
40 prohibitive for a single general-purpose machine. We propose a solver that uses *MapReduce*, a
41 distributed programming paradigm, mainly used to work with huge volumes of data on structured
42 networks of computers (*workers*) [4]. Implementations of the MapReduce model (e.g., [4]) are
43 usually executed on clusters to take full advantage of the parallel nature of the architecture. The
44 paradigm provides a basic interface consisting of two methods: `map(·)` that maps a function over
45 a collection of objects and outputs a collection of “key-value” tuples; `reduce(·)` that takes as input
46 a collection of key-value pairs and merges the values of all entries with the same key; the merging
47 operation is user-defined.



© Federico Igne, Agostino Dovier, Enrico Pontelli;
licensed under Creative Commons License CC-BY

42nd Conference on Very Important Topics (CVIT 2016).

Editors: Alessandro Dal Palù and Paul Tarau; Article No. 23; pp. 23:1–23:3

OpenAccess Series in Informatics

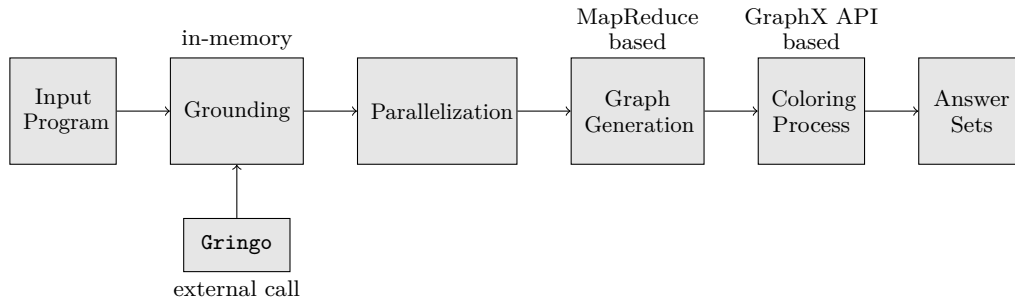


OASICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

48 An inspiration for the work proposed here comes from the proposal by Konczak et al. [9, 10],
 49 which addresses the problem of finding the answer sets of a ground normal logic program by
 50 means of computing the admissible colorings of the relative Rule Dependency Graph (RDG).
 51 This is done by defining a set of operators on the RDG of a program. These operators deal with
 52 the non-deterministic coloring of nodes and the deterministic propagation of colors. [1] used this
 53 technique in the development of the NoMoRe (Non-monotonic Reasoning with Logic Programs)
 54 solver. This implementation is purely sequential and in-memory.

55 In this research we investigate the above-mentioned graph coloring approach and extended
 56 it so as to include weight constraint rules. We investigate its mapping to MapReduce and other
 57 distributed programming paradigms that build over MapReduce. The solver we are developing,
 58 called MASP-Reduce, is written in Scala [12, 7], it uses Apache Spark [14] as a library for
 59 distributed computation, and it natively works on the Hadoop Distributed File System (HDFS).
 60 The library gives access to a complete set of primitives for the *MapReduce* programming paradigm,
 61 and on top of this, it implements GraphX, a distributed direct multigraph with a complete and
 62 easy-to-use interface [14].

63 The development of *MASP-Reduce* is heavily based on the concept of rule dependency graph
 64 of ASP programs. Graphs turn out to be a good data structure for distributed programming,
 65 since they can directly exploit the underlying network configuration. Up to now, the software
 66 is comprehensive of a solver and of a graph generator that converts a ground program in a rule
 67 dependency graph (see Figure below). As a future work, we plan to implement a distributed
 68 grounder taking full advantage of the MapReduce paradigm, so that the Grounding block is
 69 incorporated into the Parallelization block.



70
 71 We tested the solver both in a local environment (a notebook) and in a distributed environ-
 72 ment, namely four nodes of a cluster, where each node is a 12-core Intel CPUs, with each core dual
 73 hyperthreaded for a total of 48 OS-visible processors per node; each node has 256GB of RAM,
 74 ~3TB of hard disk local storage and ~512GB solid state local storage. The implementation
 75 works on simple examples. However, during the development we encountered a few challenges
 76 that prevented us from providing a full testing phase report. Spark is presented as an easy and
 77 ready-to-use tool for distributed programming; this might be true in a few cases, but most of the
 78 times one needs to fine-tune the system in order to reach an optimal configuration; this tuning
 79 process takes into account a vast number of parameters, and is mostly program-specific—and it
 80 is work in progress for our project.

81 As far as we know, this is the first work addressing the implementation of a distributed
 82 answer set solver using MapReduce paradigm and non-standard graph coloring as answer set
 83 characterization. This deeply influenced own roadmap, which couldn't take advices from previous
 84 works, leading to an incremental approach to development.

85 The system is still far from complete; we are planning on working on the development of a
 86 distributed grounder in the next few months. We are also considering the implementation of a
 87 few coloring heuristics and learning techniques to improve the performances of the solver.

88 **2012 ACM Subject Classification** Computing methodologies → Logic programming and answer
 89 set programming. Software and its engineering → Massively parallel systems

90 **Keywords and phrases** ASP solving, Parallelism, Map-reduce

91 **Digital Object Identifier** 10.4230/OASIS.ICLP.2018.2018.23

92 **Funding** A. Dovier is partially supported by UNIUD PRID “Encase” and by INdAM GNCS
 93 projects. F. Igne benefited of a scholarship from the Scuola Superiore of the University of Udine.

94 **Acknowledgements** We thank Huiping Cao for the availability of the BigDat cluster (KDD lab).

95 ——— References ———

- 96 **1** C. Anger, M. Gebser, T. Linke, A. Neumann, and T. Schaub. The NoMoRe++ System.
 97 In C. Baral, G. Greco, N. Leone, and G. Terracina, editors, *Proc of LPNMR 2005* volume
 98 3662 of *LNCS*, pages 422–426. Springer, 2005.
- 99 **2** M. Banbara, B. Kaufmann, M. Ostrowski, and T. Schaub. Clingcon: The next generation.
 100 *TPLP*, 17(4):408–461, 2017.
- 101 **3** A. Dal Palù, A. Dovier, E. Pontelli, and G. Rossi. GASP: Answer set programming with
 102 lazy grounding. *Fundam. Inform.*, 96(3):297–322, 2009.
- 103 **4** J. Dean and S. Ghemawat. *MapReduce: Simplified Data Processing on Large Clusters*,
 104 volume 51, pages 107–113. ACM, Jan. 2008.
- 105 **5** A. Dovier, A. Formisano, and E. Pontelli. An empirical study of constraint logic program-
 106 ming and answer set programming solutions of combinatorial problems. *J. Exp. Theor.*
 107 *Artif. Intell.*, 21(2):79–121, 2009.
- 108 **6** A. Dovier, A. Formisano, and E. Pontelli. Perspectives on logic-based approaches for rea-
 109 soning about actions and change. In M. Balduccini and T. C. Son, editors, *Logic Program-*
 110 *ming, Knowledge Representation, and Nonmonotonic Reasoning*, volume 6565 of *LNCS*,
 111 pages 259–279. Springer, 2011.
- 112 **7** École Polytechnique Fédérale. Scala — Object-Oriented Meets Functional (website), 2018.
 113 [last accessed Feb. 2018] <http://www.scala-lang.org/>.
- 114 **8** M. Gebser, R. Kaminski, B. Kaufmann, and T. Schaub. Clingo = ASP + control: Prelim-
 115 inary report. *CoRR*, abs/1405.3694, 2014.
- 116 **9** K. Konczak, T. Linke, and T. Schaub. Graphs and colorings for answer set programming:
 117 Abridged report. In V. Lifschitz and I. Niemelä, editors, *Proc of LPNMR2004* volume 2923
 118 of *LNCS*, pages 127–140. Springer, 2004.
- 119 **10** K. Konczak, T. Linke, and T. Schaub. Graphs and colorings for answer set programming.
 120 *TPLP*, 6(1-2):61–106, 2006.
- 121 **11** C. Lefèvre, C. Béatrix, I. Stéphane, and L. Garcia. Asperix, a first-order forward chaining
 122 approach for answer set computing. *TPLP*, 17(3):266–310, 2017.
- 123 **12** M. Odersky, L. Spoon, and B. Venners. *Programming in Scala: Updated for Scala 2.12*.
 124 Artima Incorporation, USA, 3rd edition, 2016.
- 125 **13** T. Son and E. Pontelli. Planning for biochemical pathways: A case study of answer set
 126 planning in large planning problem instances. In M. De Vos and T. Schaub, editors, *Proc*
 127 *of the First International SEA’07 Workshop, Tempe, Arizona, USA*, volume 281 of *CEUR*
 128 *Workshop Proceedings*, pages 116–130, 01 2007.
- 129 **14** The Apache Software Foundation. Apache Hadoop, Spark, and Graphx (websites), 2018.
 130 [last accessed Feb. 2018] <http://hadoop.apache.org/>, <https://spark.apache.org/>,
 131 <https://spark.apache.org/docs/latest/graphx-programming-guide.html>.