

Efficient Computation of the Well-Founded Semantics over Big Data

Ilias Tachmazidis, Grigoris Antoniou, and Wolfgang Faber

University of Hudderseld, UK

Abstract. Data originating from the Web, sensor readings and social media result in increasingly huge datasets. The so called Big Data comes with new scientific and technological challenges while creating new opportunities, hence the increasing interest in academia and industry. Traditionally, logic programming has focused on complex knowledge structures/programs, so the question arises whether and how it can work in the face of Big Data. In this paper, we examine how the well-founded semantics can process huge amounts of data through mass parallelization. More specifically, we propose and evaluate a parallel approach using the MapReduce framework. Our experimental results indicate that our approach is scalable and that well-founded semantics can be applied to billions of facts. To the best of our knowledge, this is the first work that addresses large scale nonmonotonic reasoning without the restriction of stratification for predicates of arbitrary arity.

This paper appears in:

Proceedings of the 30th International Conference on Logic Programming (ICLP 2014).
To appear.

Proceedings of the 21st RCRA workshop on *Experimental Evaluation of Algorithms for Solving Problems with Combinatorial Explosion* (RCRA 2014).
Vienna, Austria, July 17–18, 2014.