

ABSTRACT: Finding the optimal selection of an OWL reasoner and service interface for a specific ontology-based application is challenging. Matching application requirements with service offerings from available reasoning engines has become more and more difficult over time with recent optimizations for certain reasoning services and new reasoning algorithms for different fragments of OWL. This work is largely motivated by real-world experiences and will report about interesting findings in the course of developing an ontology-based application. Benchmarking outcomes of several reasoning engines are discussed - especially with respect to accompanying sound and completeness tests. Furthermore we give an account of issues and a performance comparison with respect to various service and communication protocols which show that this largely neglected component may have an enormous impact on the overall performance.