

ABSTRACT: Many applications operate on time-“sensitive” data. Some of these data are only valid for certain intervals (e.g., job-assignments, versions of software code), others describe temporal events that happened at certain points in time (e.g., a person’s birthday). Until recently, the only way to incorporate time into Semantic Web models was as a data type property. Temporal RDF, however, considers time as an additional dimension in data preserving the semantics of time. In this paper we present a syntax and storage format based on named graphs to express temporal RDF. Given the restriction to preexisting RDF-syntax, our approach can perform any temporal query using standard SPARQL syntax only. For convenience, we introduce a shorthand format called T-SPARQL for temporal queries and show how T-SPARQL queries can be translated to standard SPARQL. Additionally, we show that, depending on the underlying data’s nature, the temporal RDF approach vastly reduces the number of triples by eliminating redundancies resulting in an increased performance for processing and querying. Last but not least, we introduce a new indexing approach method that can significantly reduce the time needed to execute time point queries (e.g., what happened on January 1st).