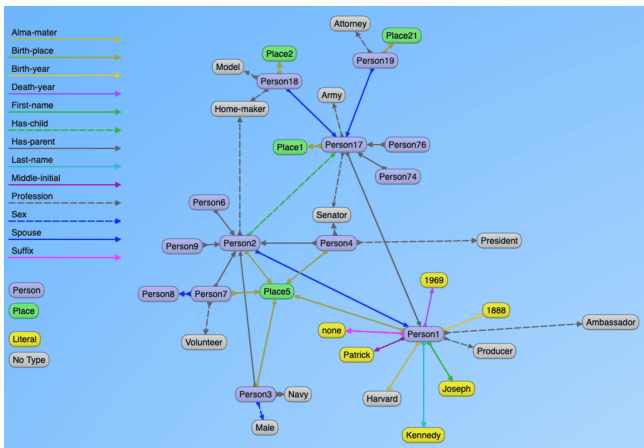


Graph Analytics with AllegroGraph and Apache Spark

AllegroGraph enables users to export data out of their Knowledge Graph and perform graph analytics with Apache Spark, one of the most popular platforms for large-scale data processing. Users immediately gain machine learning and SQL database solutions as well as GraphX and GraphFrames, two frameworks for running graph compute operations on data.

A key benefit of using Apache Spark for graph analytics within AllegroGraph is that it is built on top of Hadoop MapReduce and extends the MapReduce model to efficiently use more types of computations. Users can access interfaces (including interactive shells) for programming entire clusters with implicit data parallelism and fault-tolerance.

Apache Spark is one of the most popular platforms for large-scale data processing. In addition to machine learning, SQL database solutions, Spark also comes with GraphX and GraphFrames two frameworks for running graph compute operations on your data. In the AllegroGraph-Spark notebook, we demonstrate how to read data from AllegroGraph and then perform graph analytics with Spark.



Apache Spark was built on top of Hadoop MapReduce and it extends the MapReduce model to efficiently use more types of computations. It provides interfaces (including interactive shells) for programming entire clusters with implicit data parallelism

and fault-tolerance. For a quick start of more Spark APIs, please go to [here](#).

Visit our Github example page for more details.