

Soufflé in the Cloud

Lyndon Henry

October 10, 2017

Abstract

With the rise the in popularity of cloud computing, there has been strong movement towards software and hardware being provided as a service. Service-based solutions have the motivating advantage of simplifying the construction of distributed systems, absolving the user from the difficult task of designing, implementing, and managing scalable and fault-tolerant network infrastructure. However, there remains a great deal of complexity for the user, as software must be designed with accommodation of the specifics of their chosen platform – the original problem being solved becomes a secondary concern to the requirements imposed by the platform itself. This often necessitates an in-depth understanding of the particular platform, and a plethora of nuanced, domain specific knowledge that may change in potentially any way at any time, and must be constantly updated and reviewed to maintain compatibility.

We extend Soufflé, a compiler and interpreter for the declarative logic programming language Datalog, to synthesize services themselves from logic specifications. Statements of Datalog programs are partitioned into independent subprograms, and allocated to individual nodes of a network. These subprograms may then be executed concurrently, achieving a high degree of parallelism by elimination of unnecessary sequential behaviours. The system enables a total automation of both program and data distribution, that is kept entirely hidden from the user. A Datalog program for Soufflé is written no differently whether it is to be run locally on a single machine, or in a distributed setting of multiple machines; regardless of the size of data, or number of users it must handle.