Really Big Data: Analytics on Graphs with Trillions of Edges

Willy Zwaenepoel

School of Computer and Communication Sciences, EPFL, Lausanne, Switzerland
willy.zwaenepoel@epfl.ch

Abstract

Big graphs occur naturally in many applications, most obviously in social networks, but also in many other areas such as biology and forensics. Current approaches to processing large graphs use either supercomputers or very large clusters. In both cases the entire graph must reside in memory before it can be processed. We are pursuing an alternative approach, processing graphs from secondary storage. While this comes with a performance penalty, it makes analytics on very large graphs feasible on a small number of commodity machines. We have developed two systems, one for a single machine and one for a cluster of machines. X-Stream, the single machine solution, aims to make all secondary storage access sequential. It uses two techniques to achieve this goal, edge-centric processing and streaming partitions. Chaos, the cluster solution, starts from the observation that there is little benefit to locality when accessing data from secondary storage over a high-speed network. As a result, Chaos spreads graph data uniformly randomly over storage devices, and uses randomized access to achieve I/O balance. Chaos furthermore uses work stealing to achieve computational load balance. By using these techniques, it avoids the need for expensive partitioning during pre-processing, while still achieving good scaling behavior. With Chaos we have been able to process an 8-trillion-edge graph on 32 machines, a new milestone for graph size on a small cluster. I will describe both systems and their performance on a number of benchmarks and in comparison to state-of-the-art alternatives. This is joint work with Laurent Bindschaedler (EPFL), Jasmina Malicevic (EPFL) and Amitabha Roy (Intel Labs).

1998 ACM Subject Classification C.2.4 Distributed Systems

Keywords and phrases large graphs

Digital Object Identifier 10.4230/LIPIcs.OPODIS.2016.3

Category Keynote

20th International Conference on Principles of Distributed Systems (OPODIS 2016).
Editors: Panagiota Fatourou, Ernesto Jiménez, and Fernando Pedone; Article No. 3; pp. 3:1–3:1

Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany