

PASS: A Parallel Activity-Search System

Abstract:

Given a setmathcal of activities expressed via temporal stochastic automata, and a set of observations detections of low level events, we study the problem of identifying instances of activities from mathcal. While past work has developed algorithms to solve this problem, in this paper, we develop methods to significantly scale these algorithms. Our PASS architecture consists of three parts: leveraging past work to represent all activities in Mathcal via a single “merged” graph, partitioning the graph into a set of subgraphs, where is the number of compute nodes in a cluster, and (iii) developing a parallel activity detection algorithm that uses a different compute node in the cluster to intensively process each subgraph. We propose three possible partitioning methods and a parallel activity-search detection algorithm that coordinates computations across nodes in the cluster. We report on experiments showing that our algorithms enable us to handle both large numbers of observations per second as well as large merged graphs.