
Scheduling Task Graphs on Bounded Shared-Memory Platforms

Loris Marchal^{*1}, Bertrand Simon¹, and Frédéric Vivien¹

¹Laboratoire de l'Informatique du Parallélisme – École Normale Supérieure - Lyon, Université Claude Bernard Lyon 1, Institut National de Recherche en Informatique et en Automatique, Centre National de la Recherche Scientifique – France

Abstract

Scientific workflows are frequently modeled as Directed Acyclic Graphs (DAG) of tasks, which represent computational modules and their dependencies in the form of data produced by a task and used by another one. This formulation allows the use of runtime systems which dynamically allocate tasks onto the resources of increasingly complex computing platforms. However, for some workflows, such a dynamic schedule may run out of memory by exposing too much parallelism. This paper focuses on the problem of transforming such a DAG to prevent memory shortage, and concentrates on shared memory platforms. We first propose a simple model of DAGs which is expressive enough to emulate complex memory behaviors. We exhibit a polynomial-time algorithm that computes the maximum peak memory of a DAG, that is, the maximum memory needed by any parallel schedule. We then consider the problem of reducing this maximum peak memory to make it smaller than a given bound by adding new fictitious edges, while trying to minimize the critical path of the graph.

Keywords: task graphs, memory

^{*}Speaker