
Data Races and the Discrete Resource-time Tradeoff Problem with Resource Reuse over Paths

Rathish Das^{*1}, Shih-Yu Tsai¹, Sharmila Duppala¹, Jayson Lynch², Esther Arkin¹,
Rezaul Chowdhury¹, Joseph Mitchell¹, and Steven Skiena¹

¹Stony Brook University – United States

²MIT – United States

Abstract

A determinacy race occurs if two or more logically parallel instructions access the same memory location and at least one of them tries to modify its content. Races often lead to nondeterministic and incorrect program behavior. A data race is a special case of a determinacy race which can be eliminated by associating a mutual-exclusion lock or allowing atomic accesses to the memory location. However, such solutions can reduce parallelism by serializing all accesses to that location. For associative and commutative updates, reducers allow parallel race-free updates at the expense of using some extra space. We ask the following question. Given a fixed budget of extra space to mitigate the cost of races in a parallel program, which memory locations should be assigned reducers and how should the space be distributed among the reducers in order to minimize the overall running time? We argue that the races can be captured by a directed acyclic graph (DAG), with nodes representing memory cells and arcs representing read-write dependencies between cells. We then formulate our optimization problem on DAGs. We concentrate on a variation of this problem where space reuse among reducers is allowed by routing extra space along a source to sink path of the DAG and using it in the construction of reducers along the path. We consider two reducers and the corresponding duration functions (i.e., reduction time as a function of space budget). We generalize our race-avoiding space-time tradeoff problem to a discrete resource-time tradeoff problem with general non-increasing duration functions and resource reuse over paths. For general DAGs, the offline problem is strongly NP-hard under all three duration functions, and we give approximation algorithms. We also prove hardness of approximation for the general resource-time tradeoff problem and give a pseudo-polynomial time algorithm for series-parallel DAGs.

Keywords: Parallel and Distributed Algorithms, Multi, Core Architectures, Resource Management and Awareness, Scheduling Problems

^{*}Speaker