

On Scheduling Complex Dags for Internet-Based Computing

(extended abstract)

Grzegorz Malewicz
Department of Computer Science
University of Alabama
Tuscaloosa, AL 35487, USA
greg@cs.ua.edu

Arnold L. Rosenberg and Matthew Yurkewych
Department of Computer Science
University of Massachusetts Amherst
Amherst, MA 01003, USA
{rsnbrg,yurk}@cs.umass.edu

Conceptual tools are developed to aid in crafting a theory of scheduling complex computation-dags for Internet-based computing. The goal of the schedules produced is to render tasks eligible for allocation to remote clients (hence for execution) at the maximum possible rate. This allows one to utilize remote clients well, and also lessen the likelihood of the “gridlock” that ensues when a computation stalls for lack of eligible tasks. Earlier work has introduced a formalism for studying this optimization problem and has identified optimal schedules for several significant families of structurally uniform dags. The current paper extends this work via a methodology for devising optimal schedules for a much broader class of complex dags. These dags are obtained via composition from a prespecified collection of simple building-block dags. The paper introduces a suite of algorithms that decompose a given dag to expose its building-blocks, and a priority relation on building-blocks. When the building-blocks are appropriately interrelated, the dag can be scheduled optimally.