

A Multistart Approach to Near-optimal Concurrency Dynamics in Neighborhood-constrained Systems

Felipe M. G. França[✉] Carlos A. Matinhon[†] Gladstone M. Arantes Jr[✉]
Luerbio Faria[‡]

Abstract

This paper introduces a multistart approach to the problem of finding near-optimal "concurrency" dynamics, i.e., a metrics for the amount of concurrently operating elements in distributed systems in which the operation of each single node is constrained by its neighborhood. The determination of initial acyclic orientations upon anonymous (no labeled nodes) distributed systems of arbitrary topology are of particular interest to this approach. This practical form of symmetry breaking will be considered as a priming strategy to the application of Scheduling by Edge Reversal (SER), a simple and powerful distributed scheduling algorithm. SER requires an initial acyclic orientation on the graph G representing the target resource sharing system in order to work correctly. Such initial acyclic orientation determines an amount of concurrency ρ associated to the scheduling dynamics. Other interesting problems, such as coloring, maximal independent set and clique covering can be easily associated to SER dynamics.

SER is based on the idea of a population of processes executing upon access to shared atomic resources. Shared-resources systems are represented as a finite oriented graph G where processes are the nodes of G and an oriented edge exists between any two nodes whenever they share a resource. A process is allowed to operate if it has all its edges directed to itself, a sink in G . SER works starting from any acyclic orientation ω on G . This means that there is at least one sink node exists. Consider, just for the purpose of this explanation, an ideal synchronous environment where according to the beginning of a clock pulse, only sinks are allowed to operate while other nodes remain idle. After operating and before the clock pulse ends, sinks reverse the orientation of their edges by sending messages to all their neighbours, each one becoming a source. It is easy to see that, again, another acyclic orientation ω^0 , is formed and a new set of sinks can operate at the next clock pulse. All subsequent orientations are also acyclic and the scheduling mechanism consists basically of consecutive sets of sinks being defined in G through time. This simple dynamics ensures that no deadlock or starvation will ever occur since at every acyclic orientation there is at least one sink, i.e., one node allowed to operate.

We use two new probabilistic distributed algorithms for the generation of random acyclic orientations at the construction phase of our multistart procedure. A local search procedure is then applied, looking for improvement over ρ , through transformations over the length of the sink decomposition ω associated to the initial acyclic orientation ω . In this sense, some possible local search strategies are considered, as: (i) immediate application of SER from ω (this does not change the ρ associated to ω); (ii) selective transformation of arbitrary nodes into sinks or sources, and (iii) combining (i) and (ii).

Keywords: anonymous systems, mutual exclusion, randomized distributed algorithms, symmetry breaking.

[✉]COPPE, Universidade Federal do Rio de Janeiro, Brazil, felipe@cos.ufrj.br

[†]Institute of Computing, Universidade Federal Fluminense, Brazil, mart@dcc.ic.uff.br

[‡]Faculdade de Formação de Professores, Universidade Estadual do Rio de Janeiro, Brazil, luerbio@cos.ufrj.br