

Models and algorithms for the reconfiguration of distributed wireless switching systems

Renaud Sirdey

Received: 13 April 2007 / Published online: 5 June 2007
© Springer-Verlag 2007

Abstract This paper is a summary of the author's PhD thesis entitled "Models and algorithms for the reconfiguration of wireless switching systems". The thesis deals with the study of a strongly *NP*-hard resource-constrained scheduling problem arising from the telecommunication industry. This work was supervised by Jacques Carlier and Dritan Nace, both from Université de Technologie de Compiègne, and carried out while the author was a System Architect within Nortel GSM Access R&D organization. The thesis, which is written in both French and English, has been defended on 29 March 2007 and is available by email request to the author.

Keywords Scheduling · Polyhedral combinatorics · Branch-and-bound · Branch-and-cut · Simulated annealing · Distributed systems

MSC classification (2000) 68M20 · 90C57 · 90C59 · 68M14

1 Introduction

This PhD thesis is devoted to the study of a strongly *NP*-hard resource-constrained scheduling problem: the *Process Move Programming (PMP) problem*. This problem

This research was supported in part by Association Nationale de la Recherche Technique grant CIFRE-121/2004.

R. Sirdey (✉)
Service d'architecture BSC (PC 12A7), Nortel GSM Access R&D,
Parc d'activités de Magny-Châteaufort, 78928 Yvelines Cedex 09, France
e-mail: renauds@nortel.com

R. Sirdey
Laboratoire Heudiasyc, Université de Technologie de Compiègne,
Compiègne Cedex, France

arises from the telecommunication industry, in relation to the operability of certain high availability real-time distributed systems. It is one of the many combinatorial problems that the author has encountered in the design of wireless switching systems (see [Sirdey 2007](#); [Sirdey et al. 2003](#)).

In a few words, the PMP problem consists, starting from an arbitrary admissible initial distribution of *processes* on the *processors* of a distributed system, in finding a least disruptive sequence of operations (non-impacting process migrations or temporary process interruptions) at the end of which the system ends up in another predefined arbitrary state. The main constraint lies in the fact that the capacity of the processors must not be exceeded during the reconfiguration.

After establishing the strong *NP*-hardness of the problem ([Sirdey et al. 2005](#)), we have approached it from different angles: exact resolution of special cases, exact resolution in the general case and efficient approximate resolution. Each of these approaches achieves different software engineering trade-offs.

2 Polynomial cases

We have first focused on the study of polynomially solvable special cases ([Sirdey et al. 2005](#)). To an instance of the PMP problem we associate a directed multigraph, denoted by D and called the *transfer multigraph*, whose vertices are associated to the processors and such that an arc is associated to each of the process moves.

It is then easy to prove that if the transfer multigraph is acyclic, then the problem can be solved in linear time, without any interruption, by considering a topological ordering of the vertices of D . Furthermore, we have shown that in the case in which this graph contains some directed cycles, its strongly connected components can be processed independently in the reverse order of a topological ordering on the reduced transfer multigraph.

We have also established that the case in which the consumption of each of the processes is constant can be solved in polynomial time by an algorithm proceeding by successive extractions of maximal eulerian subgraphs and applications of the above strong connectivity-based decomposition rule.

3 A branch-and-bound algorithm

Although it was clear that only efficient approximate resolution algorithms were industrially viable, due to the real time context of our application, such algorithms need to be validated. Hence, we have first focused on an exact resolution branch-and-bound algorithm for the PMP problem ([Sirdey et al. 2005](#)).

The algorithm initially starts with the worst possible solution, consisting in interrupting all the moves. Then an admissible solution (a set of interruptions and an admissible ordering) is built, each branching decision consisting in choosing an interrupted process to concatenate to the solution ordering, among those for which doing so preserves the solution admissibility. A leaf is obtained when no such process exists. This scheme is complemented by a pseudopolynomial lower bound, requiring the resolution of one knapsack problem per processor, as well as by a dominance rule

which generalizes several other intuitive rules. In particular, this rule allows to filter out solutions which are locally equivalent on a per processor basis.

Although this algorithm was able to solve “hard” instances with more than 100 moves in less than 20 min, it has reasonable chances of success for instances with up to around 50 moves within a 20 min time limit. Note that 50 moves is the order of magnitude of the biggest instance in the context of our industrial application. This algorithm, which is reasonably easy to implement, allowed us to build a database of instances with known optimal solution for the validation of approximate resolution algorithms.

4 A simulated annealing algorithm

In order to practically solve the PMP problem we have then focused on the design of a simulated annealing-based approximate resolution algorithm (Sirdey et al. 2006).

At the basis of this algorithm lies the notion of (α, β) -acceptable solution which we have introduced. Given α and $\beta \in [0, 1]$, an (α, β) -acceptable solution to a minimization problem has, with probability at least α , a value less than or equal to $\text{OPT} + \beta(\text{WRT} - \text{OPT})$, where OPT is the value of an optimal solution and WRT is the value of a worst possible one. We have then used the markovian theory underlying the homogeneous simulated annealing algorithm to derive conditions under which such solutions may be produced.

Applying these (widely applicable) results leads to a pseudopolynomial time approximation algorithm for the PMP problem which explores the permutations of the moves in a way such that the optimum can always be generated. Although our reasoning is, at the end of day, of heuristic nature, we have performed extensive computational experiments which strongly suggest that our algorithm meets its design intent for $\alpha = 0.95$ and $\beta = 0.05$.

5 Polyhedral combinatorics

Lastly, in an attempt to push the exact resolution of the PMP problem in the realm of bigger instances, we have applied the powerful techniques of polyhedral combinatorics to that problem (Sirdey and Kerivin 2006, 2007; Kerivin and Sirdey 2007). This work was done in collaboration with Hervé Kerivin from Université de Clermont-Ferrand II.

From a theoretical standpoint we have studied two polytopes, namely the *partial linear ordering* and the *process move program* polytopes, which we have introduced. In particular, we have given several classes of facet-defining inequalities for these polytopes and studied the associated separation problems. These results were used to design a branch-and-cut algorithm based on a linear relaxation theoretically solvable in pseudopolynomial time using the ellipsoid algorithm.

Lastly, extensive computational experiments demonstrated the practical relevance of the algorithm. On top of being able to solve exactly “hard” instances with up to 119 moves, our algorithm was able to *provably* solve all the instances, with up to 180 moves, on which it was tried with an optimality gap of less than 5% in less than 4 hours.

6 Perspectives

It should be emphasized, as initially hinted by a fellow PhD student, Olivier Klopfenstein from France Telecom R&D, that there is a close connection between the PMP problem and certain flavours of the *Reroute Sequence Planning (RSP) problem* in MPLS networks (Józsa and Makai 2003). It thus turns out that the models and algorithms presented in this thesis naturally generalize to that latter problem and can be expected to significantly improve on the resolution methods currently available for the RSP problem.

References

- Józsa BG, Makai M (2003) On the solution of reroute sequence planning problem in MPLS networks. *Comput Networks* 42:199–210
- Kerivin H, Sirdey R (2007) Polyhedral combinatorics of a resource-constrained ordering problem part II: on the process move program polytope. Technical Report PE/BSC/INF/017913 V01/EN, Service d'architecture BSC, Nortel GSM Access R&D, France (submitted for publication)
- Sirdey R (2007) Combinatorial optimization problems in wireless switch design. *4OR* (in press)
- Sirdey R, Kerivin H (2006) A branch-and-cut algorithm for a resource-constrained scheduling problem. *RAIRO—Oper Res* (in press)
- Sirdey R, Kerivin H (2007) Polyhedral combinatorics of a resource-constrained ordering problem part I: on the partial linear ordering polytope. Technical Report PE/BSC/INF/017912 V01/EN, Service d'architecture BSC, Nortel GSM Access R&D, France (submitted for publication)
- Sirdey R, Plainfossé D, Gauthier J-P (2003) A practical approach to combinatorial optimization problems encountered in the design of a high availability distributed system. In: *Proceedings of INOC*, pp 532–539
- Sirdey R, Carlier J, Kerivin H, Nace D (2005) On a resource-constrained scheduling problem with application to distributed systems reconfiguration. *Eur J Oper Res* (in press)
- Sirdey R, Carlier J, Nace D (2006) Approximate resolution of a resource-constrained scheduling problem. *J Heuristics* (in press)