INFORMS Journal on Computing
Vol. 23, No. 2, Spring 2011, pp. 189–204
ISSN 1091-9856 | EISSN 1526-5528 | 11 | 2302 | 0189

DOI 10.1287/ijoc.1100.0393

Decision Support and Optimization in Shutdown and Turnaround Scheduling

Nicole Megow
Max-Planck-Institut für Informatik, 66123 Saarbrücken, Germany, nmegow@mpi-inf.mpg.de

Rolf H. Möhring, Jens Schulz
Institut für Mathematik, Technische Universität Berlin, 10623 Berlin, Germany
{moehring@math.tu-berlin.de, jschulz@math.tu-berlin.de}

Large-scale maintenance in industrial plants requires the entire shutdown of production units for disassembly, comprehensive inspection, and renewal. We derive models and algorithms for this so-called turnaround scheduling that include different features such as time-cost trade-off, precedence constraints, external resource units, resource leveling, different working shifts, and risk analysis. We propose a framework for decision support that consists of two phases. The first phase supports the manager in finding a good makespan for the turnaround. It computes an approximate project time-cost trade-off curve together with a stochastic evaluation. Our risk measures are the expected tardiness at time t and the probability of completing the turnaround within time t. In the second phase, we solve the actual scheduling optimization problem for the makespan chosen in the first phase heuristically and compute a detailed schedule that respects all side constraints. Again, we complement this by computing upper bounds for the same two risk measures.

Our experimental results show that our methods solve large real-world instances from chemical manufacturing plants quickly and yield an excellent resource utilization. A comparison with solutions of a mixed-integer program on smaller instances proves the high quality of the schedules that our algorithms produce within a few minutes.

Key words: project management; planning; scheduling; resource constraints; risk analysis; applications; large-scale systems; chemical industries
History: Accepted by Karen Aardal, Area Editor for Design and Analysis of Algorithms; received August 2009; revised February 2010; accepted March 2010. Published online in *Articles in Advance* July 2, 2010.

1. Introduction

Large-scale maintenance activities are conducted on a regular basis in industrial settings such as chemical manufacturing, refining, or power plants. Entire production units are shut down for disassembly, comprehensive inspection, and renewal. Such a process is called *shutdown and turnaround* (or turnaround, for short). It is an essential process but causes high out-of-service cost. Therefore a good schedule for the turnaround has a high priority to the manufacturer. A good schedule is not simply a short schedule. The project execution can be speeded up at the expense of adding resource units, mostly in the form of additional workers. Thus, short projects cause high resource costs, whereas cheap projects take a long time. Moreover, in practice, task execution times typically involve uncertainty. Such uncertainty arises as a result of unforeseen repair jobs, and naturally, a short schedule is less robust against unexpected repair jobs or processing delays than a schedule with a long duration that offers more flexibility for rescheduling. Such considerations are fundamental in the decision process of a turnaround project manager. We support this process by analyzing the trade-off between project duration and project cost as well as the effects on the stability of schedules. Our main contribution is an optimization algorithm within a larger decision support framework that computes a detailed schedule of given project duration with the aim of minimizing the total resource cost.

Clearly, turnaround projects differ in size, duration, and particular specifications depending on the actual industrial site. Generally, the turnaround of a large plant may take in total up to one year and must be repeated every four to six years. Typically, such large projects are split into a sequence of shutdowns of single production units that are planned individually. The time horizon for those projects is often between two and four weeks. The complex working steps of a turnaround are planned in advance with very high granularity. They are split into many jobs that must be executed in parallel or directly after each other by workers with a particular specialization such as electricians, pipe fitters, inspectors, cranes, crane drivers, or other craftsmen. In theory, models with jobs that are executed by workers with different specializations seem plausible, but here we focus on the case where each job needs exactly one specialized type of worker.

Table 1 Comparison of Runtime and Resource Availability Cost (or Lower Bounds) of Optimal Schedules Obtained by CPLEX and Schedules Obtained by Our Heuristic (Between One and Two Resource Units per Job)

Number of jobs	Runtime heuristic	Average runtime CPLEX	Optimal solutions CPLEX (%)	Optimal solutions heuristic (%)	Gap heuristic vs. opt. (%)	Max. gap heuristic vs. opt. (%)	Gap heuristic vs. LB (%)	Max. cost gap heuristic vs. LB (%)
30	<1 sec	20 sec	100	61	8	33	—	—
40	<1 sec	20 sec	100	47	10	38	—	—
50	<1 sec	25 sec	100	59	7	29	—	—
60	<1 sec	12 min	83	40	9	33	10	33

Table 2 Comparison of Runtime and Resource Availability Cost (or Lower Bounds) of Optimal Schedules Obtained by CPLEX and Schedules Obtained by Our Heuristic (Between Six and Seven Resource Units per Job)

Number of jobs	Runtime heuristic	Average runtime CPLEX	Optimal solutions CPLEX (%)	Optimal solutions heuristic (%)	Gap heuristic vs. opt. (%)	Max. gap heuristic vs. opt. (%)	Gap heuristic vs. LB (%)	Max. cost gap heuristic vs. LB (%)
30	<1 sec	20 min	68	25	12	48	15	48
40	<1 sec	36 min	43	0	13	24	14	30
50	<1 sec	36 min	41	7	16	33	17	38
60	<1 sec	54 min	10	3	10	19	18	36

Table 3 Comparison of Runtime and Resource Availability Cost (or Lower Bounds) of Optimal Schedules Obtained by CPLEX and Schedules Obtained by Our Heuristic (Between Two and Four Resource Units per Job)

Number of jobs	Runtime heuristic	Average runtime CPLEX	Optimal solutions CPLEX (%)	Optimal solutions heuristic (%)	Gap heuristic vs. opt. (%)	Max. gap heuristic vs. opt. (%)	Gap heuristic vs. LB (%)	Max. cost gap heuristic vs. LB (%)
30	<1 sec	29 sec	100	43	9	31	—	—
40	<1 sec	19 min	73	10	11	57	14	69
50	<1 sec	45 min	15	0	11	17	15	25

column in Tables 1 and 2 reveals that with the increasing the number of jobs, CPLEX computations reach the given time limit, and thus, the computation process is aborted without having found an optimal solution. This situation occurs, in particular, if we set the resource allocation to six or seven resource units. In a few cases, our heuristic actually found an optimal solution after a few seconds, whereas CPLEX did not within one hour. The fifth column in Tables 1 and 2 quantifies how often our heuristic yields provable optimal solutions. For resource allocations between one and two units, we solve a significant number of instances to optimality. This changes when the resource requirements increase to six and seven units. In those cases, in which a provably optimal solution is found (by CPLEX), we compare its cost with those of our heuristic. The sixth column in the tables shows that we obtain solutions that are close to optimum. We leave, on average, a gap of about 7% to 10% and 38% in the worst case; see Table 1. Increasing the number of resource units yields somewhat worse results with an average gap of up to 16% and 48% in the worst case; see Table 2. The last columns in Tables 1 and 2 compare the results of our heuristic to the lower bounds obtained by CPLEX. If CPLEX has solved all instances to optimality, we omit the last entries because they are equal to the sixth and seventh columns. We compute the cost gap of our heuristic to the lower bound over all instances and to those where CPLEX found an optimal solution. In total, this does not dramatically increase the average gap, which shows that our solutions leave a gap of same size to the lower bounds as to the suboptimal solutions by CPLEX.

So far we have evaluated our heuristic on instances with up to two processing alternatives per job. Table 3 shows our computational results on instances where each job requires two, three, or four resource units that are given in advance per job. The work volume per job lies between 20 and 50. The other parameters are equal to those of the instances before. With increasing numbers of jobs as well as increasing project durations, the number of optimal solutions found by CPLEX decreases and thus the number of proven optimal solutions found by our heuristic solutions also decreases. It turns out that the average optimality gap is again about 10%. The maximum gap has been 57% for a single instance.

6. Conclusions and Research Perspectives

To the best of our knowledge, popular project management software does not support a time-cost trade-off related analysis. Resource-leveling packages do exist but seem to use very simple heuristics. Moreover, they have their limitations in the presence of working shifts, capacity bounds, or other specialized constraints such as conflicting job sets.

Motivated by applications in chemical manufacturing, we have formulated the shutdown and turnaround scheduling problem as an integrated problem that contains various optimization problems as subproblems, such as the time-cost trade-off problem, the problem of scheduling with resource capacities and working shifts, and the resource leveling problem, all of which have been considered individually previously. We reported on our successful solution approach within a more comprehensive decision support tool that additionally provides tools for risk analysis during the decision process and for the final schedule. Our optimization algorithm yields near-optimal solutions in a very short time. We hope that our work initiates more research on this general and integrated model to overcome the deficiencies of current project management tools.

Another very challenging line of research has come out of extensive discussions with practitioners about how to cope with uncertainty in turnaround scheduling. One question addresses the risk inherent in the whole planning process. So far we have only implemented tools for risk evaluation of given schedules. We applied techniques to determine an upper bound on the expected tardiness and the probability of meeting the makespan for a given project schedule. This allows the project manager to choose a schedule according to his or her risk affinity. Nevertheless, this method is only applied after the schedule optimization. We expect that an integrated approach that combines risk analysis and scheduling might yield better decision support for turnaround projects, but this is currently beyond the optimization methods available in practice.

Acknowledgments

The authors thank T.A. Cook Consultants for providing the data used to test their solution method and for explaining the details of shutdown and turnaround scheduling projects. The authors also thank Eamonn T. Coughlan, Elisabeth Günther, Ralf Hoffmann, Robert Pankrath, and Ines Spenke for their support in implementing the algorithms and developing their software tool.

References

Adel, S. M., S. E. Elmaghraby. 1984. Optimal linear approximation in project compression. *IIE Trans.* **16**(4) 339–347.

Adlakha, V. G., V. G. Kulkarni. 1989. A classified bibliography of research on stochastic PERT networks: 1966–1987. *Inform. Systems Oper. Res.* **27**(3) 272–296.

Bandelloni, M., M. Tucci, R. Rinaldi. 1994. Optimal resource leveling using non-serial dynamic programming. *J. Soc. Indust. Appl. Math.* **78**(2) 162–177.

Bourges, A., J. Killebrew. 1962. Variation in activity level in a cyclical arrow diagram. *J. Indust. Engrg.* **13**(2) 76–83.

Brucker, P., A. Drexl, R. Möhring, K. Neumann, E. Pesch. 1999. Resource-constrained project scheduling: Notation, classification, models, and methods. *Eur. J. Oper. Res.* **112**(1) 3–41.

Cieliebak, M., T. Erlebach, F. Hennecke, B. Weber, P. Widmayer. 2004. Scheduling with release times and deadlines on a minimum number of machines. J.-J. Lévy, E. W. Mayr, J. C. Mitchell, eds. *Exploring New Frontiers of Theoretical Informatics: IFIP 18th World Comput. Congress. Proc. Third IFIP Internat. Conf. Theoret. Comput. Sci.* (*TCS 2004*), Kluwer Academic Publishers, Norwell, MA, 209–222.

De, P., E. J. Dunne, J. B. Ghosh, C. E. Wells. 1995. The discrete time-cost tradeoff problem revisited. *Eur. J. Oper. Res.* **81**(2) 225–238.

De, P., E. J. Dunne, J. B. Ghosh, C. E. Wells. 1997. Complexity of the discrete time-cost tradeoff problem for project networks. *Oper. Res.* **45**(2) 302–306.

Deineko, V. G., G. J. Woeginger. 2001. Hardness of approximation of the discrete time-cost tradeoff problem. *Oper. Res. Lett.* **29**(5) 207–210.

Demeulemeester, E. L., W. S. Herroelen. 2002. *Project Scheduling: A Research Handbook*. Kluwer Academic Publishers, Norwell, MA.

Du, J., J. Y.-T. Leung. 1989. Complexity of scheduling parallel task systems. *SIAM J. Discrete Math.* **2**(4) 473–487.

Easa, S. M. 1989. Resource leveling in construction by optimization. *J. Construction Engrg. Management* **115**(2) 302–316.

Elmaghraby, S. E., J. Kamburowski. 1992. The analysis of activity networks under generalized precedence relations (GPRs). *Management Sci.* **38**(9) 1245–1263.

Falk, J. E., J. L. Horowitz. 1972. Critical path problems with concave cost-time curves. *Management Sci.* **19**(4, Part 1) 446–455.

Franck, B., K. Neumann, C. Schwindt. 2001. Project scheduling with calendars. *OR Spectrum* **23**(3) 325–334.

Fulkerson, D. R. 1961. A network flow computation for project cost curves. *Management Sci.* **7**(2) 167–178.

Grigoriev, A., M. Sviridenko, M. Uetz. 2007. Machine scheduling with resource dependent processing times. *Math. Programming* **110**(1) 209–228.

Hagstrom, J. N. 1988. Computational complexity of PERT problems. *Networks* **18**(2) 139–147.

Harris, R. B. 1990. Packing method for resource leveling (PACK). *J. Construction Engrg. Management* **116**(2) 331–350.

Jansen, K., H. Zhang. 2006. An approximation algorithm for scheduling malleable tasks under general precedence constraints. *ACM Trans. Algorithms* **2**(3) 416–434.

Kapur, K. C. 1973. An algorithm for the project cost-duration analysis problem with quadratic and convex cost functions. *IIE Trans.* **5**(4) 314–332.

Kelley, J. E., Jr. 1961. Critical-path planning and scheduling: Mathematical basis. *Oper. Res.* **9**(3) 296–320.

Lamberson, L. R., R. R. Hocking. 1970. Optimum time compression in project scheduling. *Management Sci.* **16**(10) B597–B606.

Lepère, R., D. Trystram, G. J. Woeginger. 2002. Approximation algorithms for scheduling malleable tasks under precedence constraints. *Internat. J. Found. Comput. Sci.* **13**(4) 613–627.

Ludwig, A., R. H. Möhring, F. Stork. 2001. A computational study on bounding the makespan distribution in stochastic project networks. *Ann. Oper. Res.* **102**(1–4) 49–64.

Meilijson, I., A. Nádas. 1979. Convex majorization with an application to the length of critical paths. *J. Appl. Probab.* **16**(3) 671–677.

Möhring, R. H. 2001. Scheduling under uncertainty: Bounding the makespan distribution. H. Alt, ed. *Comput. Discrete Math. Lecture Notes in Computer Science*, Vol. 2122. Springer, Berlin, 79–97.

Neumann, K., J. Zimmermann. 2000. Procedures for resource leveling and net present value problems in project scheduling with general temporal and resource constraints. *Eur. J. Oper. Res.* **127**(2) 425–443.

Neumann, K., C. Schwindt, J. Zimmermann. 2003. *Project Scheduling with Time Windows and Scarce Resources*, 2nd ed. Springer, Berlin.

Phillips, D. T., A. Garcia-Diaz. 1981. *Fundamentals of Network Analysis*. Prentice-Hall, Englewood Cliffs, NJ.

Phillips, S., Jr., M. I. Dessouky. 1977. Solving the project time/cost tradeoff problem using the minimal cut concept. *Management Sci.* **24**(4) 393–400.

Pritsker, A. A. B., L. J. Watters, P. M. Wolfe. 1969. Multiproject scheduling with limited resources: A zero-one programming approach. *Management Sci.* **16**(1) 93–108.

Schmidt, G. 2000. Scheduling with limited machine availability. *Eur. J. Oper. Res.* **121**(1) 1–15.

Shabtay, D., G. Steiner. 2007. A survey of scheduling with controllable processing times. *Discrete Appl. Math.* **155**(13) 1643–1666.

Siemens, N., C. Gooding. 1975. Reducing project duration at minimum cost: A time-cost tradeoff algorithm. *OMEGA* **3**(5) 569–581.

Skutella, M. 1998a. Approximation algorithms for the discrete time-cost tradeoff problem. *Math. Oper. Res.* **23**(4) 909–929.

Skutella, M. 1998b. Approximation and randomization in scheduling. Ph.D. thesis, Technische Universität Berlin, Berlin.

Vanhoucke, M. 2005. New computational results for the discrete time/cost trade-off problem with time-switch constraints. *Eur. J. Oper. Res.* **165**(2) 359–374.

Vanhoucke, M., D. Debels. 2007. The discrete time/cost trade-off problem: Extensions and heuristic procedures. *J. Sched.* **10**(4–5) 311–326.

Vanhoucke, M., E. Demeulemeester, W. Herroelen. 2002. Discrete time/cost trade-offs in project scheduling with time-switch constraints. *J. Oper. Res. Soc.* **53**(7) 1–11.

Weiss, G. 1986. Stochastic bounds on distributions of optimal value functions with applications to PERT, network flows and reliability. *Oper. Res.* **34**(4) 595–605.

Yang, H.-H., Y.-L. Chen. 2000. Finding the critical path in an activity network with time-switch constraints. *Eur. J. Oper. Res.* **120**(3) 603–613.

Zhan, J. 1992. Calendarization of time planning in MPM networks. *ZOR—Methods Models Oper. Res.* **36**(5) 423–438.