Estimating the limiting value of optimality for very large NP problems

George M. White

Abstract The search for better solutions to large NP-hard problems such as timetabling, personnel scheduling, resource allocation, etc., often requires approximation methods. These methods can often yield solutions that are often "very good", although it is generally impossible to say just how good these solutions are. Not only do we not know what the best solutions are, we also don't know how far away we are from the optimum solution - we may be very close but we may also be quite far. This paper describes a method of using historical data to estimate the limiting optimality of the solution to a problem if the problem arises from a situation taken from the real world and there is sufficient historical data about proposed solutions. Knowledge of the limiting optimality can provide guidance in estimating just how far a "very good" solution lies from the best solution, even when we don't know what the best solution is.

Keywords NP-hard problems · examination scheduling · penalty estimation

1 Introduction

If you are reading this text, you are likely very familiar with the difficulties of solving large scale problems. These are problems that are commonly known by names such as the travelling salesman problem, the graph (or vertex) colouring problem, examination scheduling, staff scheduling and the like.

Our challenge is nearly always to find a solution to these problems that is in some sense "the best", a concept that is more easily stated than defined. Most attempts to define just what is meant by "the best" are very context-sensitive but one that seems satisfactory in many circumstances is based upon the principle of utility proposed by Jeremy Bentham, that the right way to act is the way that causes "the greatest good for the greatest number of people". Thus the best solution is usually the one that

G.M. White
School of Information Technology and Engineering
University of Ottawa
Ottawa K1N 6N5 Canada
Tel.: +613-562-5800 x6677
Fax: +613-562-5664
E-mail: white@site.uottawa.ca

causes the least expense to travelling salesmen, uses the fewest colours to colour a map, minimizes the misery of exam writing students, or minimizes the complaints from staff when the teaching schedule, nursing schedule, employee roster ,..., is published.

When such problems are solved on our computers, this principle is used to formulate a mathematical expression strongly dependent on the exact nature of the problem, as is the method used to extract it. Our interest here is focussed on a class of problems known as non-deterministric, polynomial time hard problems (NP-hard), sometimes defined informally as those problems as hard as the hardest problems in NP. A much more detailed and precise definition and discussion is found in the classic book (Garey and Johnson (1979)).

Dispite their complexity some NP-hard problems have been solved to completeness. The travelling salesman problem, TSP, has been studied since at least 1832. Provably optimal solutions to this problem can be obtained by a variety of techniques entailing a prodidgeous amounts of computer time.

Examination scheduling problems have been in existance ever since there were examinations. Casting these schedules has evolved from a chore to be done by hand with pencils, paper and large erasers to a programming exercise to be solved by computer.

Sports scheduling has resisted computerization for a long time but is now slowly and with some reluctance being increasingly done by machine (Easton et al (2003)).

A discussion of the possibility of predicting future results based on the analysis of past results appears in the next section along with two examples of its potential in elite speed sports. The logistic equation is a potential candidate for obtaining this prediction and appears in section 3. The TSP and examination scheduling are developed in sections 4 and 5. Conclusions are found in section 6.

2 Is prediction possible?

Investigators working on the difficult problems discussed earlier continue to make progress year after year because of improvements to both hardware and software:

– processor speed
– algorithm used
– initial conditions
– parameter tuning
– manpower available

Similar conditions are true in elite sporting events. The goal is to produce a superior outcome and to achieve this goal choices must be made from a wide number of variables such as training, genetics, health, equipment, weather conditions, altitude, environment, diet, etc.

One example that can be cited is the men's 100 metre sprint. Whoever is the current record holder is often referred to as "The World's Fastest Man". A plot of the progressive world record in this event and the year in which it was achieved is shown in figure 1. The curved line is an attempt to fit the points to an equation and the dashed, straight line near the bottom of the graph is the asymptotic value of the curve.

A second example is the men's 5000 metre speed skate. The records for the fastest man on ice are shown in figure 2. For sporting events it is evident that although the times required to establish a new record are always being reduced, they will never be reduced to zero. In the case of sprinting, the dashed line represents an asymptotic

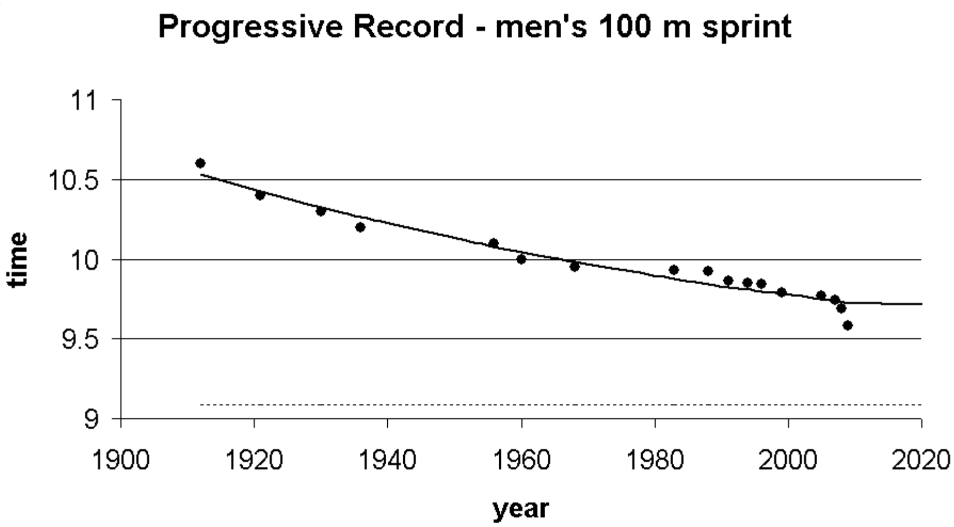

Fig. 1 Progressive record for men's 100 metre sprint

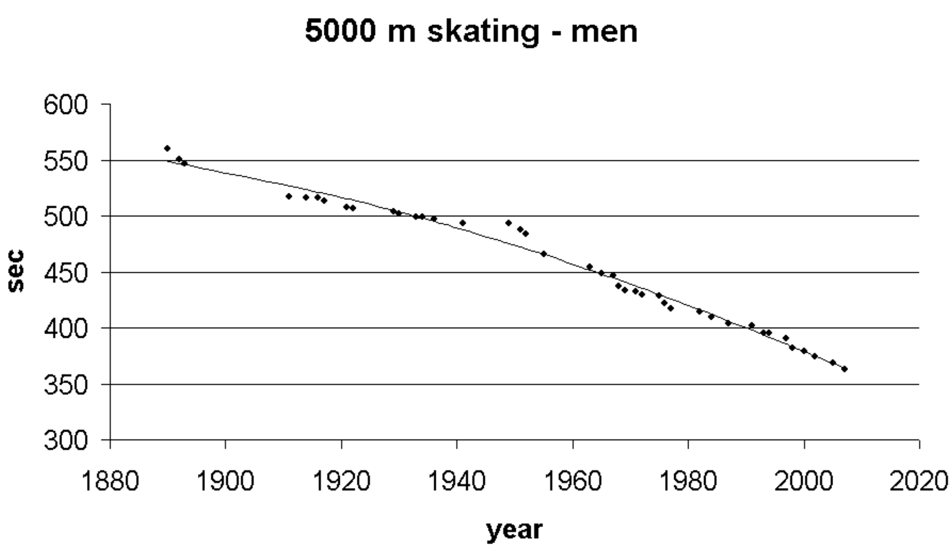

Fig. 2 Progressive record for men's 5000 metre skate

limiting value of the "ultimate" speed record and is suggested by the flattening of the curve at time goes on. For the skating event, the curve is steepening rather than flattening and suggests that there is some distance to go before the records converge on some limiting value. The analysis of curves such as this will be discussed in the next section.

This approach to records can analgously be applied to similar records in the area of large NP-hard problems touched on earlier. Perhaps a quantitative analysis can give some insight into the TSP problem and the examination problem. We may be able to predict future records, given a date, and perhaps also to estimate limiting values.

There is some evidence that other problems arising from the real world exhibit a similar behaviour.

3 The Logistic Curve

The shape of the curve that describes the experimental running best value leads to the conclusion that, at some time in the future, the best values will reach a limit. *i.e.*

$$\lim_{t\to\infty} \frac{dP(t)}{dt} = 0 \tag{1}$$

where $P(t)$ is the penalty obtained at time t. Since the exam scheduling problem is known to be NP-hard, the form of the derivative dP/dt is unknown, but it is reasonable to assume that it is some function of the current best penalty.

$$\frac{dP}{dt} = f(P) \tag{2}$$

Expanding this as a Maclaurin series yields

$$\frac{dP}{dt} = f(P) = a_0 + a_1P + a_2P^2 + a_3P^3 + ... \tag{3}$$

To simplify the form of the equation we might first try to approximate it as $dP/dt = a_0$. Then when P attains its limiting value, we have $dP/dt = 0$ and therefore $a_0 = 0$. This cannot possibly be the case. The next form to consider is $dP/dt = a_1P$ which equals 0 only if $P = 0$; this is probably not the case. The next simplest form is $\frac{dP}{dt} = a_1P + a_2P^2 = P(a_1 + a_2P)$ This has the desired properties. Recall that when P takes its limiting value, $dP/dt = 0$. It follows that $a_1+a_2P_{limit} = 0$ or $P_{limit} = -a_1/a_2$.

In the literature, this equation often appears in the form

$$\frac{dP}{dt} = rP - \frac{rP^2}{k} = rP(1 - \frac{P}{k}) \tag{4}$$

This is a differential equation whose solution is

$$P(t) = \frac{kP_0e^{rt}}{k + P_0(e^{rt} - 1)} \tag{5}$$

where $P(t)$ is the penalty obtained at time t and r, k and P_0 are adjustable parameters. When $t = 0, P(t) = P(0) = P_0$. As $t \to \infty$, $P(t) \to k$.

A sketch of this equation in the form $P(rt)$ with k set equal to $5 * P_0$ and $r = 1$ is shown in figure 3. This equation is used to describe birth-death processes and race results among other applications. Here we will use it to analyse some published NP-hard results.

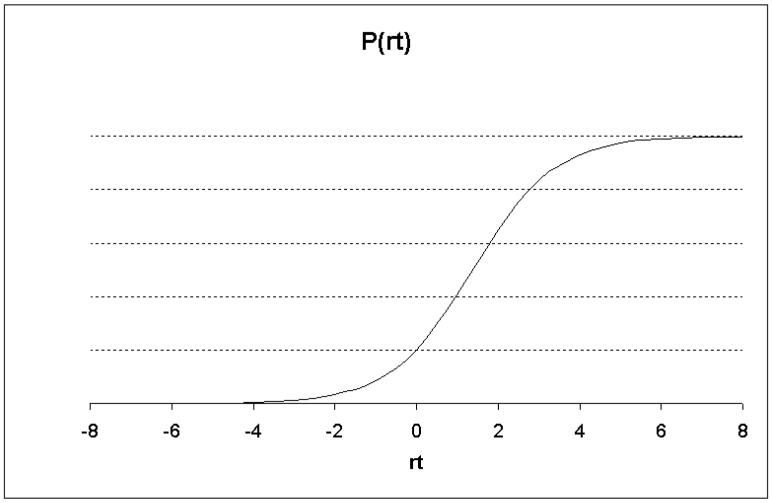

Fig. 3 Sketch of $P(rt)$ vs rt

4 The Travelling Salesman Problem (TSP)

The TSP is a real-world problem that has many practical applications and has therefore been intensively studied ever since it was introduced. The volumous history of the problem and its literature has been reviewed in several books (see for example the books by Lawler et al (1985); Applegate et al (2006)).

A book published in Germany in the 1830s described the problem in the context of actual travelling salemen who were wanting to cover their territory in the shortest possible time, but did no mathematical analysis. The first investigator to treat the problem in a mathematical setting was an Irishman, Sir William Rowan Hamilton, who studied the problem in the 1850s. Exact solutions to non-trivial problems were slow to appear because of the amount of calculation that had to be done but, by 1954, Danzig, Fulkerson and Johnson pubished the results for a 49 city instance, the capitals of the lower 48 United States plus Washington.

Since then the size of successfully solved TSP problems has grown steadily with the present record holder being D. Applegate and 6 colleagues, (Applegate et al (2006)) who solved a 85900 city instance in 2006. A graph of the sizes of solved instances and the year they were obtained is shown in figure 4.

The plot on the right shows the size of the successful instances vs. the year published (large points) and the best fitted curve of $P(t)$ to these points. Because of the large range of values, the same data has been plotted on a semi-log scale on the right. The smooth curve is a plot of equation 5 fitted by the Solver tool of Microsoft Excel to the data points. The fit is remarkable considering the passage of time between the first and the last points (about 36 years) and the variety of computers and software used to perform the calculations.

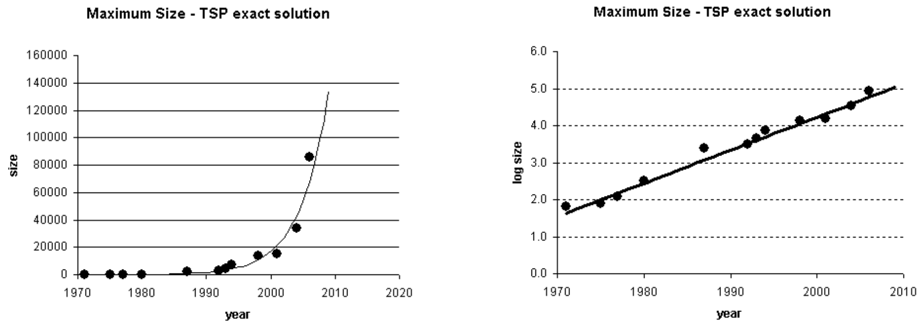

Fig. 4 Progressive record for TSP solved problem instances

5 The Examination Scheduling Problem

Examination scheduling is one of the earliest applications of computer technology to an academic problem. The possibility of finding new methods to study the problem, the promise of a useful application and the availability of real data from real sources, gave a strong impetus for academics to study the problem. Accordingly computer researchers started investigating this problem (see Broder (1964); Peck and Williams (1966)) and a few programs were written and used in practice (White and Chan (1979); Carter (1983)).

Researchers have employed many techniques in order to find better solutions. Discussions of methods and summaries of progress have been well treated in Qu et al (2009).

A *feasible* exam timetable is one in which no student is required to sit for more than one exam at a time. Although any feasible timetable will work, some of these timetables are worse than others. Several measures of the "badness" of a timetable have been proposed, such as

- the total number of consecutive exams a student must write
- the total number of consecutive exams plus the total number of exams separated by exactly one free timeslot

In 1996, a seminal paper (Carter et al (1996, 1997)) proposed a penalty, based on some earlier work, that is equal to the weighted sum of course pair penalties (Laporte and Desroches (1984)). Two exams taken by one student separated by n timeslots incurs a penalty p_n. The number of such penalties incurred by all the students is w_n. The penalty of the entire timetable is then defined to be

$$\sum_{i=1}^{5} p_i w_i \tag{6}$$

where $p_1 = 16, p_2 = 8, p_3 = 4, p_4 = 2, p_5 = 1$, and the summation is calculated over all students involved. The penalty so obtained is then divided by the number of students involved to get a *standard penalty*. The authors also referenced a depository of 13 data sets taken from real institutions that they used to test their algorithms. The benchmarks that resulted have been used ever since as a basis of comparison.

Some problems in the original data sets have been detected and corrected (see Qu et al (2009)).

Progress during this time has been made by many researchers who have employed many different approaches with a view to lowering the standard penalty when using new algorithms with the same data. The results are not unlike those obtained in track and field events where athletes attempt to lower the time required to cover a specified distance, say 100 metres, where basic conditions are unchanged. The outcomes of a foot race depend on a large number of variables: the individual racer, training, genetics, health, equipment, weather conditions, altitude, environment, diet, and the use of banned substances, among other things. The outcomes of experiments that cast examination timeables likewise depend on a large set of variables such as the algorithm used, initial conditions, parameter tuning, manpower available, processor speed, etc.

An examination of the recent timetabling literature shows that a wide variery of techniques has indeed been used and many researchers have published tables of their best results for the Toronto data base. A plot of the published values of the standard penalty against the year in which this result was published for the data set *yor-f-83* is shown in figure 5 (left). The running best penalty of a schedule for a given year is just

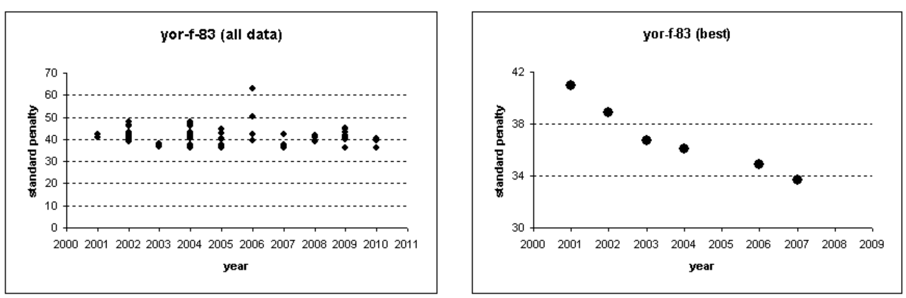

Fig. 5 Data points available for the data set *yor-f-83*

that value, obtained in that year, that had the lowest value. The set of running best penalties obtained yet for that data set at a given year is formed by choosing only those values that are better than any preceding lowest value. A plot of these best values is shown in figure 1 (right). Note that the axes of this graph have been rescaled.

Most of the data sets for which sufficient data is available show the same general tendency exhibited by the *yor-f-83* set. The data is sparse as of this writing but the behaviour of the best points for each year appears to indicate a trend. The improvement in later years is smaller than it was in earlier years and it is very unlikely that the value of the penalty will ever reach zero. This suggests that the best penalty points for each year may fall along a smooth curve that starts with some initial value, decreases slowly over the years and approaches some asymptotic non-zero positive value. This raises the question as to whether past performance can be used to forecast future results. If this is true, then perhaps an analytical study of past attempts to obtain lower standard penalties can be used to predict future lower standard penalties.

The problem of finding the best schedule arises from the sheer size of the solution space and the fact that the problem itself is NP-hard. For the *yor-f-83* or *yor83 I* data, the problem involves (a) partitioning the 181 exams into 21 timeslots and then

(b) permuting the order of these timeslots in order to find the resulting schedule having the lowest penalty.

The number of ways of partitioning n exams into k timeslots is given by a Stirling Number of the second kind

$$\left\{ \begin{matrix} n \\ k \end{matrix} \right\} = \left\{ \begin{matrix} n-1 \\ k-1 \end{matrix} \right\} + k \left\{ \begin{matrix} n-1 \\ k \end{matrix} \right\} \tag{7}$$

with

$$\left\{ \begin{matrix} n \\ 1 \end{matrix} \right\} = \left\{ \begin{matrix} n \\ n \end{matrix} \right\} = 1 \tag{8}$$

Each set of partitions can be arranged in $k!$ ways.

Thus for the *yor83 I* data, there are $\left\{ \begin{matrix} 181 \\ 21 \end{matrix} \right\} = 4.09 \times 10^{219}$ ways to partition and $21! = 5.11 \times 10^{19}$ permutations of these partitions giving a total solution space of 2.09×10^{239} entries. As a basis of comparison, the total number of protons in the universe has been estimated to be roughly 10^{80}.

For the data set *yor-f-83*, the progressive "world record" was tabulated along with the year in which the work was published (see figure 5 right). The records correspond to the best result (if any) published during the corresponding calendar year. If the record was broken more than once during the year, the best result was taken.

The original data points obtained by Carter et al (1996) were not used in the analysis because it was the first time that the data and the penalty used were presented to the research world. The long time before the next result was published is not representative of the interval separating the next improvements. The logistic curve (6) was fitted to the remaining data points using the Solver tool in Microsoft Excel. The goal of the solver was to minimize the squared deviations between the published results and the fitted equation while adjusting the constants k, P_0 and r. When this is done, the limiting value of the standard penalty is calculated to be 32.44. A graph of the best fitted logistics curve, the data points and the limiting value P_{lim} is shown in figure 6. The same procedure was followed for some of the other data sets. The results obtained by this analysis is shown in table 1.

Table 1 Limiting penalty values

data set	no. points	P_{lim}	std.dev.	notes
car-f-92	5	3.82	0.08	1
tre-s-92	5	7.30	0.20	
ute-s-92	5	23.08	0.57	
yor-f-83	6	32.44	0.30	

The meaning of the first two columns of this table is obvious. Column three lists the limiting value of the penalty, $P_{lim} = k$, as calculated by the least squares fit. Column four, labelled *std.dev*, is a measure of the expected deviation of this limit and is calculated as:

$$\sqrt{\frac{1}{n-1} \sum (P_i - y_i)^2} \tag{9}$$

where

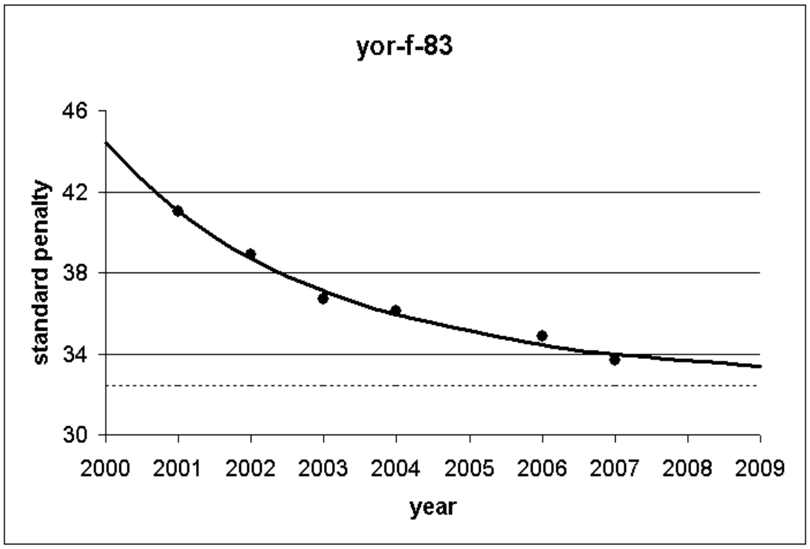

Fig. 6 The points, the best fit and the limiting value for the data set *yor-f-83*

– P_i is the value of the fitted curve for the year in question (referred to as $P(t)$ in the continuous domain of equation (6)).
– y_i is the value of the corresponding best penalty found in that year (if any).
– n is the number of points.

The column marked *notes* refers to the explanatory caveat listed below.

1. One paper, published in 2001, reports solution values that could be interpreted as being too good, too soon (see example in figure 1 (left)). This may be because of an error in the authors' calculations or it may be because of the superiority of their algorithms. In our calculation the reported value was omitted and the best value for that year was taken from the remaining values.

6 Conclusions

The data available for the TSP is far larger than the data for examination scheduling. The goodness of fit as defined by equation (9) equals 32.9, a value that is very small in comparison to the size of the problems now being successfully solved.

The numerical values of the limiting penalty given in table 1 should be used with caution. They are based on small amounts of data and their values will change as better solutions with smaller penalties emerge. The largest number of points available for any of the data sets is 6. The smallest number is 5; only just large enough for a curve to be fitted but not large enough to inspire great confidence in the result. Most of the data sets have few values in their running "best value" sets. One data set, *pur-s-93*, has so few results available that we were unable to even begin an analysis.

It must be realized that the values are asymptotic. Solutions having the limiting penalty will never be realised in finite time.

Not all the data sets yield a curve fit as well as does *yor-s-83* (see figure 6). The standard deviation, calculated as 0.030, is very good in this case. Values of accuracy that are not so good demonstrate either that the logistics curve is not a good description of the data or that the numerical parameters were badly fitted. This situation may be improved as more data becomes available.

There is no underlying theory developed as yet that supports the behaviour of the running best penalties observed over time. The basic examination scheduling problem is NP-hard and the best solution can be obtained only by exhaustive search.

The values observed and the limiting values calculated cannot be used in reverse to determine the schedule that produced them, *i.e.* there is no bijective relation between the solution space and the penalty space.

The x axis of the graphs corresponds to the year in which the data was published, not the year in which it was first obtained. This results in an uncertainty in these values. Also various publications use different accuracies in their reporting of the best penalties obtained. The difference in these accuracies may result in too many or too few data points selected in the list of best current values for a given year.

The limiting value cannot be used to rank solution methods. The fact that a certain algorithm yielded the latest best solution does not imply that a solution having that limiting penalty could be generated using that algorithm, or any other known algorithm.

However, the calculations can be used to estimate how close a given penalty is to a plausible minimum. This knowledge may then be used to calculate if the solution in question is "good enough" to use in practice. The best estimates of limiting values can be incorporated into algorithms and used as part of a stopping criterion.

They may also be used to estimate whether a penalty instance is reasonable, given the other penalties and the dates when they were obtained.

References

Applegate DL, Bixby RE, Chvatal V, Cook WJ (2006) The Travelling Salesman Problem: A Computational Study. Princeton University Press

Broder S (1964) Final examination scheduling. Communications of the ACM 7:494–498

Carter M, Laporte G, Lee S (1996) Examination timetabling: Algorithmic strategies and applications. Journal of the Operational Research Society 47:373–383

Carter M, Laporte G, Lee S (1997) Corrigendum. Journal of the Operational Research Society 48:225

Carter MW (1983) A decomposition algorithm for practical timetabling problems. Working Paper 83-06, Industrial Engineering, University of Toronto

Easton K, Nemhauser G, Trick M (2003) Solving the traveling tournament problem: A combined integer and constraint programming approach. In: Burke E, De Causmaecker P (eds) PATAT IV: Lecture Notes in Computer Science, Springer, Berlin, vol 2740, pp 63–77

Garey MR, Johnson DS (1979) Computers and Intractability: A Guide to the Theory of NP-completeness. Freeman

Laporte G, Desroches S (1984) Examination timetabling by computer. Computers and Operations Research 11:351–360

Lawler E, Lenstra J, Kan AR, Shmoys D (1985) The Travelling Salesman Problem: A Guided Tour of Combinatorial Optimization. John Wiley and Sons

Peck JEL, Williams MR (1966) Algorithm 286: Examination scheduling. Communications of the ACM 9:433–434

Qu R, Burke EK, McCollum B, Merlot LTG, Lee SY (2009) A survey of search methodologies and automated system development for examination timetabling. J Scheduling 12(1):55–89

White GM, Chan PW (1979) Towards the construction of optimal examination schedules. INFOR 17(3):219–229