An Efficient Tabu Search Heuristic for the School Timetabling Problem

Haroldo G. Santos[1], Luiz S. Ochi[1], and Marcone J.F. Souza[2]

[1] Computing Institute, Fluminense Federal University, Niterói, Brazil
{hsantos,satoru}@ic.uff.br
[2] Computing Department, Ouro Preto Federal University, Ouro Preto, Brazil
marcone@iceb.ufop.br

Abstract. The School Timetabling Problem (STP) regards the weekly scheduling of encounters between teachers and classes. Since this scheduling must satisfy organizational, pedagogical and personal costs, this problem is recognized as a very difficult combinatorial optimization problem. This work presents a new Tabu Search (TS) heuristic for STP. Two different memory based diversification strategies are presented. Computational experiments with real world instances, comparing with a previously proposed TS found in the literature, show that the proposed method produces better solutions for all instances, as well faster times are observed in the production of good quality solutions.

1 Introduction

The School Timetabling Problem (STP) embraces the scheduling of sequential encounters between teachers and students so as to insure that requirements and constraints are satisfied. Typically, the manual solution of this problem extends for various days or weeks and normally produces unsatisfactory results due to the fact that lesson periods could be generated which are inconsistent with pedagogical needs or could even serve as impediments for certain teachers or students. STP is considered a NP-hard problem [5] for nearly all of its variants, justifying the usage of heuristic methods for its resolution. In this manner, various heuristic and metaheuristic approaches have been applied with success in the solution of this problem, such as: Tabu Search (TS) [10,4,8], Genetic Algorithms [11] and Simulated Annealing (SA) [2].

The application of TS to the STP is specially interesting, since this method is, as local search methods in general, very well suited for the interactive building of timetables, a much recognized quality in timetable building systems. Furthermore, TS based methods often offer the best know solutions to many timetabling problems, when compared to other metaheuristics [3,9]. The diversification strategy is an important aspect in the design of a TS algorithm. Since the use of a tabu list is not enough to prevent the search process from becoming trapped in certain regions of the search space, other mechanisms have been proposed. In particular, for the STP, two main approaches have been used: adaptive relaxation [8,4] and random restart [10]. In adaptive relaxation the costs involved

in the objective function are dynamically changed to bias the search process to newly, unvisited, regions of the search space. In random restart a new solution is generated and no previous information is utilized.

This work employs a TS algorithm that uses an informed diversification strategy, which takes into account the history of the search process to bias the selection of diversification movements. Although it uses only standard Tabu Search components, it provides better results than more complex previous proposals [10].

The article is organized as follows: section 2 presents related works; section 3 introduces the problem to be treated; section 4 presents the proposed algorithm; section 5 describes the computational experiments and their results; and finally, section 6 formulates conclusions and future research proposals.

2   Related Works

Although the STP is a classical combinatorial optimization problem, no widely accepted model is used in the literature. The reason is that the characteristics of the problem are highly dependent on the educational system of the country and the type of institution involved. As such, although the basic search problem is the same, variations are introduced in different works [3,4,8,10]. Described afterwards, the problem considered in this paper derives from [10] and considers the timetabling problem encountered in typical Brazilian high schools. In [10], a GRASP-Tabu Search (`GTS-II`) metaheuristic was developed to tackle this problem. The `GTS-II` method incorporates a specialized improvement procedure named "Intraclasses-Interclasses", which uses a shortest-path graph algorithm. At first, the procedure is activated aiming to attain the feasibility of the constructed solution, after which, it then aims to improve the feasible solution. The movements made in the "Intraclasses-Interclasses" also remain with the tabu status for a given number of iterations. Diversification is implemented through the generation of new solutions, in the GRASP constructive phase. In [9] three different metaheuristics that incorporate the "Intraclasses-Interclasses" were proposed: Simulated Annealing, Microcanonical Optimization (MO) and Tabu Search. The TS proposal outperformed significantly both SA and MO.

3   The Problem Considered

The problem considered deals with the scheduling of encounters with teachers and students over a weekly period. The schedule is made up of d days of the week with h daily periods, defining $p = d \times h$ distinct periods. There is a set T with t teachers that teach a set S of s subjects to a set C of c classes, which are disjoint sets of students with the same curriculum. The association of teachers to subjects in certain classes is previously fixed and the workload is informed in a matrix of requirements $R_{t \times c}$, where r_{ij} indicates the number of lessons that teacher i shall teach for class j. Classes are always available, and must have their time schedules, of size p, completely filled out, while teachers indicate a

set of available periods. Also, teachers may request a number of double lessons per class. These lessons are lessons which must be allocated in two consecutive periods on the same day. This way a solution to the STP problem must satisfy the following constraints:

1. no class or teacher can be allocated for two lessons in the same period;
2. teachers can only be allocated respecting their availabilities;
3. each teacher must fulfill his/her weekly number of lessons;
4. no class can have more than two lessons of a given subject per day.

Also, there are the following desirable features that a timetable should present:

1. the time schedule for each teacher should include the least number possible of days;
2. double lessons requests must be satisfied whenever possible;
3. "gaps" in the time schedule of teachers should be avoided, that is: periods of no activity between two lesson periods.

3.1 Solution Representation

A timetable is represented as a matrix $Q_{t\times p}$, in a such way that each row represents the complete weekly timetable for a given teacher. As such, the value $q_{ik} \in \{0, 1, \cdots, c\}$, indicates the class for which the teacher i is teaching during period k ($q_{ik} \in \{1, \cdots, c\}$), or if the teacher is available for allocation ($q_{ik} = 0$). The advantage of this representation is that it eliminates the possibility for the occurrence of conflicts in the timetable for teachers. The occurrence of conflicts in classes happens when in a given period k more than one teacher is allocated to that class. Allocations are only allowed in periods with availability of teachers. A partial sample of a timetable with 5 teachers can be found in Figure 1, with value "X" indicating the unavailabilities of teachers.

Teacher \ Period	1	2	3	4	5	$\cdots d \times h$
1	1	0	0	2	2	$\cdots$
2	0	X	X	0	1	$\cdots$
3	X	X	1	0	3	$\cdots$
4	0	1	0	1	0	$\cdots$
5	0	0	2	3	X	$\cdots$

Fig. 1. Fragment of generated timetable

3.2 Objective Function

In order to treat STP as an optimization problem, it is necessary to define an objective function that determines the degree of infeasibility and satisfaction of

requirements; that is, pretends to generate feasible solutions with minimal number of unsatisfied requisites. Thus, a timetable Q is evaluated with the following objective function, which should be minimized:

$$f(Q) = \omega \times f_1(Q) + \delta \times f_2(Q) + \rho \times f_3(Q) \tag{1}$$

where f_1 counts, for each period k, the number of times that more than one teacher teaches the same class in period k and the number of times that a class has no activity in k. The f_2 portion measures the number of allocations that disregard the daily limits of lessons of subjects in classes (constraint 4). As such, the timetable can only be considered feasible if $f_1(Q) = f_2(Q) = 0$. The importance of the costs involved defines a hierarchy so that: $\omega > \delta \gg \rho$. The f_3 component in the objective function measures the satisfaction of personal requests from teachers, namely: double lessons, non existence of "gaps" and timetable compactness, as follows:

$$f_3(Q) = \sum_{i=1}^{t} \alpha_i \times b_i + \beta_i \times v_i + \gamma_i \times c_i \tag{2}$$

where α_i, β_i, and γ_i are weights that reflect, respectively, the relative importance of the number of "gaps" b_i, the number of week days v_i each teacher is involved in any teaching activity during the same shift, and the non negative difference c_i between the minimum required number of double lessons and the effective number of double lessons in the current agenda of teacher i.

4 The Proposed Algorithm

Tabu Search (TS) methods [6] are adaptive procedures that make use of search history information to guide the improvement heuristic so that it is not confounded by the absence of improvement movements. The design of a TS algorithm involves basically the following definitions: the construction procedure that will provide an initial solution, the type of movement that will be applied in the search process and which memory structures will be used. These components will be described as follows:

4.1 Constructive Algorithm

The constructive algorithm basically consists of a greedy randomized constructive procedure [7]. While in other works the option for a randomized construction is to allow diversification, through the re-start of the process, in this case the purpose is only to have control of the randomization degree of the initial solution. To build a solution, step-by-step, the principle of allocating the *most urgent* lessons in the *most appropriate* periods is used. In this case, the urgency degree θ_{ij} of allocating a lesson from teacher i for class j is computed considering the available periods V_i from teacher i, the available periods W_j from class j and the number of unscheduled lessons $unscheduled_{ij}$ of teacher

i for class j, as follows: $\theta_{ij} = \frac{unscheduled_{ij}}{|V_i \cap W_j|+1}$. The algorithm then builds a restricted candidate list (RCL) with the most urgent lessons, in a such a way that: $RCL = \{i,j\} \mid \theta_{ij} \geq \overline{\theta} - (\overline{\theta} - \underline{\theta}) \times \alpha$, where $\overline{\theta} = \max\{\theta_{ij} \mid i \in T,\ j \in C\}$ and $\underline{\theta} = \min\{\theta_{ij} \mid i \in T,\ j \in C\}$. The α parameter allows tuning the randomization degree of the algorithm, varying from the pure greedy algorithm ($\alpha = 0$) to a completely random ($\alpha = 1$) selection of the teacher and class to allocation. At each step, the urgency degrees are recomputed. The selected lesson is allocated attempting to maintain the timetable free of conflicts and giving priority to periods with less teacher availability.

4.2 Tabu Search Components

The TS procedure starts from the initial solution provided by the constructive algorithm and, at each iteration, fully explores the neighborhood $N(Q)$ to select the next movement. The movement, in this case, consists in the swap of two values in the timetable of a teacher $i \in \{1, \cdots, t\}$, and can be defined as $\langle i, p_1, p_2 \rangle$, such that $q_{ip_1} \neq q_{ip_2}$, $p_1 < p_2$ and $p_1, p_2 \in \{1, \cdots, p\}$. The best movement is chosen at each iteration, even if it does not improves the best solution found. In order to try to prevent cycles a short term memory is updated, containing the last reverse movements performed. Once a movement m enters the tabu list, it will remain with the tabu status for a random number of iterations $tt(m)$, such that $tt(m) \in \{minTabuTenure, \cdots, maxTabuTenure\}$. The tabu list defines a modified neighborhood, in a way that the algorithm will select the best non tabu movement from this neighborhood. In order to not exclude good solutions from the allowed set of movements, an aspiration criteria is defined. In this case, if the movement improves the best solution found so far, it will loose its tabu status.

Since short-term memory is not enough to prevent the search process from become entrenched in certain regions of the search space, long-term memory is also employed. In this case a transition based long-term memory, that considers the frequency of moves involving a given teacher and class was used. These frequencies are computed using a matrix $Z_{t \times c}$, whose counts z_{ij}, represent the number of movements that were done involving teacher i and class j. The counts are zeroed whenever the best solution found is updated. The transition ratio of frequency of movements for teacher i and class j is computed in the following way, where $\overline{z} = \max\{z_{ij} \mid i \in T,\ j \in C\}$:

$$transitionRatio_{ij} = \frac{z_{ij}}{\overline{z}} \tag{3}$$

These values will be used in the diversification strategy, in a way that the execution of few explored movements will be stimulated. This is done through the incorporation of penalties in the evaluation of movements. The penalty for a movement considers the cost of the best solution found so far $f(Q^*)$, and involves two allocations, say q_{ip_1} and q_{ip_1}. These allocations can consist in lessons for two classes, or a lesson to one class and a free period. This way, the penalty for a movement involving teacher i and allocations of periods p_1, $a_1 = q_{ip_1}$ and p_2, $a_2 = q_{ip_2}$, can be calculated as:

$$penalty_{ia_1a_2} = \begin{cases} transitionRatio_{ia_1} \times f(Q^*) & \text{if } a_1 \neq 0 \text{ and } a_2 = 0 \\ transitionRatio_{ia_2} \times f(Q^*) & \text{if } a_1 = 0 \text{ and } a_2 \neq 0 \\ (transitionRatio_{ia_1} + transitionRatio_{ia_2})/2 \times f(Q^*) & \text{if } a_1 \neq 0 \text{ and } a_2 \neq 0 \end{cases}$$

Another penalty function also considers the teacher workload to promote diversification. In this case, the objective is to favor movements involving teachers whose timetable changes would probably produce bigger modifications in the solution structure. This penalty function ($penaltyTL$), is computed as follows:

$$penaltyTL_{ia_1a_2} = \frac{penalty_{ia_1a_2}}{\sum_{j=1}^{c} r_{ij} / \max_{l=1}^{t} \left(\sum_{j=1}^{c} r_{lj}\right)} \tag{4}$$

The diversification strategy is applied whenever signals that regional entrenchment may be in action are detected. In this case, the number of non-improvement iterations is evaluated before starting the diversification strategy. Movements performed in this phase can be viewed as *influential movements* [6], in a way that these movements try to modify the solution structure in a influential (non-random) way.

The pseudo-code in Figure 2 presents the proposed TSDS heuristic. The parameters $minTT$ and $maxTT$ inform, respectively, the lower bound and the upper bound of the tabu tenure values that can be randomly selected at each iteration. The parameters $activationDiv$ and $iterationsDiv$ inform the number of non improvement iterations necessary to start the diversification process, and the number of iterations that the process will remain active, respectively. The binary operator $\oplus$ means the application of a movement in the current solution. The function *computePenalty* can use one of the penalty functions previously presented. In the following sections, the implementation that considers the penalty function that only takes into account the frequency ratio of transitions will be referred as TSDS, while the implementation that use the penalty function that takes into account also the workload of teachers will be referred as TSDSTL. For comparison purposes, an implementation without the diversification strategy (TS), also will be considered.

```
procedure TSDS(Q, minTT, maxTT, activationDiv, iterationsDiv)
begin
  f* = f(Q); Q* = Q; TabuList = ∅;
  noImprovementIterations = 0; iteration = 0;
  repeat
    deltaCost(bestMov) = ∞; iteration + +;
    for each movement m ∈ N(Q)
      penalty = 0;
      if ((noImprovementIterations mod activationDiv) < iterationsDiv) and
        (noImprovementIterations > iterationsDiv) then
        penalty = computePenalty(m);
      endif
      if (f(Q) − f(Q ⊕ m) + penalty < deltaCost(bestMov) and (bestMov ∉ TabuList))
        or (f(Q ⊕ m) < f*)
        bestMov = m;
        deltaCost(bestMov) = f(Q) − f(Q ⊕ m);
      endif
    end for
    Q = Q ⊕ m;
    tabuTenure(m) = random(minTT, maxTT);
    UpdateTabuList(m, iteration);
    ComputeMovementFrequency(m);
    if (f(Q) < f*) then
      noImprovementIterations = 0; Q* = Q; f* = f(Q*);
    else
      noImprovementIterations + +;
    endif
  until (stoppingCriterionReached())
  return Q*;
end.
```

Fig. 2. Pseudo-code for TSDS algorithm

5   Computational Experiments and Discussion

Experiments were done in the set of instances originated from [10], and the data referred to Brazilian high schools, with 25 lesson periods per week for each class, in different shifts. In Table 1 some of the characteristics of the instances can be verified, such as dimension and sparseness ratio (sr), which can be computed considering the total number of lessons $(\#lessons)$ and the total number of unavailable periods (u): $sr = \frac{t \times p - (\#lessons + u)}{t \times p}$. Lower sparseness values indicate more restrictive problems and likewise, more difficult resolution.

Instance	Teachers	Classes	Total Lessons	Double Lessons	Sparseness Ratio (sr)
1	8	3	75	21	0.43
2	14	6	150	29	0.50
3	16	8	200	4	0.30
4	23	12	300	66	0.18
5	31	13	325	71	0.58
6	30	14	350	63	0.52
7	33	20	500	84	0.39

Table 1. Characteristics of problem instances

The algorithms were coded in C++. The implementation of `GTS-II` was the same presented in [10], and was implemented in C. The compiler used was GCC 3.2.3 using flag `-O2`. The experiments were performed in a micro-computer with an AMD Athlon XP 1533 MHz processor, 512 megabytes of RAM running the Linux operating system.

The weights in the objective function were defined as in [10]: $\omega = 100$, $\delta = 30$, $\rho = 1$, $\alpha_i = 3$, $\beta_i = 3$ and $\gamma_i = 1$, $\forall i = 1, \cdots, t$.

In the first set of experiments, the objective was to verify the average solution cost produced by each algorithm, within a given time limit. The results (Table 2) consider the average best solution found in 20 independent executions, with the following time limits to instances $1, \cdots, 7$, respectively: $\{90, 280, 380, 870, 1930, 1650, 2650\}$. The parameters for `GTS-II` and the time limits are the same proposed in [10]. The parameters for `TSDS` and its variations are: $\alpha = 0.1$ (constructive algorithm), $minTT = 20$, $maxTT = 25$, $activationDiv = 500$ and $iterationsDiv = 10$. Best results are shown in bold.

As it can be seen in Table 2, although only minor differences can be observed among the two implementations that use different penalty functions in the diversification strategy, results show that versions that use the informed diversification strategy perform significantly better than `GTS-II` and `TS`.

In other set of experiments, the objective was to verify the empirical probability distribution of reaching a given sub-optimal target value (i.e. find a solution

Instance	GTS-II	TSDSTL	TSDS	TS
1	204.80	203.42	**203.37**	207.05
2	350.10	**344.84**	345.36	349.26
3	455.70	439.94	**439.05**	455.58
4	686.30	**669.69**	672.15	670.92
5	796.30	782.74	**780.74**	782.84
6	799.10	783.38	**781.77**	787.85
7	1,076.20	1,060.84	**1,059.05**	1,071.21

Table 2. Average results with fixed time limits

with cost at least as good as the target value) in function of time in different instances. The sub-optimal values were chosen in a way that the slowest algorithm could terminate in a reasonable amount of time. In these experiments, TSDSTL and GTS-II were evaluated and the execution times of 150 independent runs for each instance were computed. The experiment design follows the proposal of [1]. The results of each algorithm were plotted associating with the i-th smallest running time t_i a probability $p_i = (i - \frac{1}{2})/150$, which generates points $z_i = (t_i, p_i)$, for $i = 1, \cdots, 150$. As it can be be seen in Figures 3 to 6 the TSDSTL heuristic achieves high probability values ($\geq 50\%$) of reaching the target values in significantly smaller times than GTS-II. This difference is enhanced mainly in instance 4, which presents a very low sparseness ratio. This result may be related to the fact that the "Intraclasses-Interclasses" procedure of GTS-II works with movements that use free periods, which are hard to find in this instance. Another analysis show that at the time when 95% of TSDSTL runs have achieved the target value, in average, only 64% of GTS-II runs have achieved the target value. Considering the time when 50% of TSDSTL runs have achieved the target value, only 11%, in average, of GTS-II runs have achieved the target value.

6 Concluding Remarks

This paper presented a new tabu search heuristic to solve the school timetabling problem. Experiments in real world instances showed that the proposed method outperforms significantly a previously developed hybrid tabu search algorithm, and it has the advantage of a simpler design.

Contributions of this paper include the empirical verification that although informed diversification strategies are not commonly employed in tabu search implementations for the school timetabling problem, its incorporation can significantly improve the method robustness. The proposed method not only produced better solutions for all test instances but also performed faster than a hybrid tabu search approach.

Although the proposed method offers quite an improvement, future researches may combine the "Intraclasses-Interclasses" procedure with an informed diversification strategy, which could lend to even better results .

Acknowledgements

References

1. Aiex, R. M., Resende, M. G. C., Ribeiro, C. C.: Probability distribuition of solution time in GRASP: an experimental investigation, Journal of Heuristics, **8** (2002), 343–373
2. Abramson, D.: Constructing school timetables using simulated annealing: sequential and parallel algorithms. Management Science. **37** (1991) 98–113.
3. Colorni, A., Dorigo, M., Maniezzo, V.: Metaheuristics for High-School Timetabling. Computational Optimization and Applications. **9** (1998) 277–298.
4. Costa, D.: A Tabu Search algorithm for computing an operational timetable. European Journal of Operational Research Society. **76** (1994) 98–110.
5. Even, S., Itai, A., Shamir, A.: On the complexity of timetabling and multicommodity flow problems. SIAM Journal of Computation. **5** (1976) 691–703.
6. Glover, F., Laguna, M.: Tabu Search. Kluwer Academic Publishers, Boston Dordrecht London (1997)
7. Resende, M.G.C., Ribeiro. C.C.: Greedy randomized adaptive search procedures. Handbook of Metaheuristics. Kluwer. (2003) 219–249
8. Schaerf, A.: Tabu search techniques for large high-school timetabling problems. Report CS-R9611. Centrum voor Wiskunde en Informatica, Amsterdam (1996)
9. Souza, M.J.F.: Programação de Horários em Escolas: Uma Aproximação por Metaheurísticas, D.Sc. Thesis (in Portuguese), Universidade Federal do Rio de Janeiro - Rio de Janeiro (2000)
10. Souza, M.J.F., Ochi, L.S., Maculan, N.: A GRASP-Tabu search algorithm for solving school timetabling problems. In: Resende, M.G.C., Souza, J.P. (eds.): Metaheuristics: Computer Decision-Making. Kluwer Academic Publishers, Boston (2003) 659–672
11. Wilke, P, Gröbner, M., Oster, N.: A hybrid genetic algorithm for school timetabling. In: AI 2002: McKay B. and Slaney J. (eds.): Advances in Artificial Intelligence. Springer Lecture Notes in Computer Science, Vol. 2557. Springer-Verlag, New York (2002) 455–464

Fig. 3. Empirical probability distribution of finding target values in function of time for instances 1 and 2
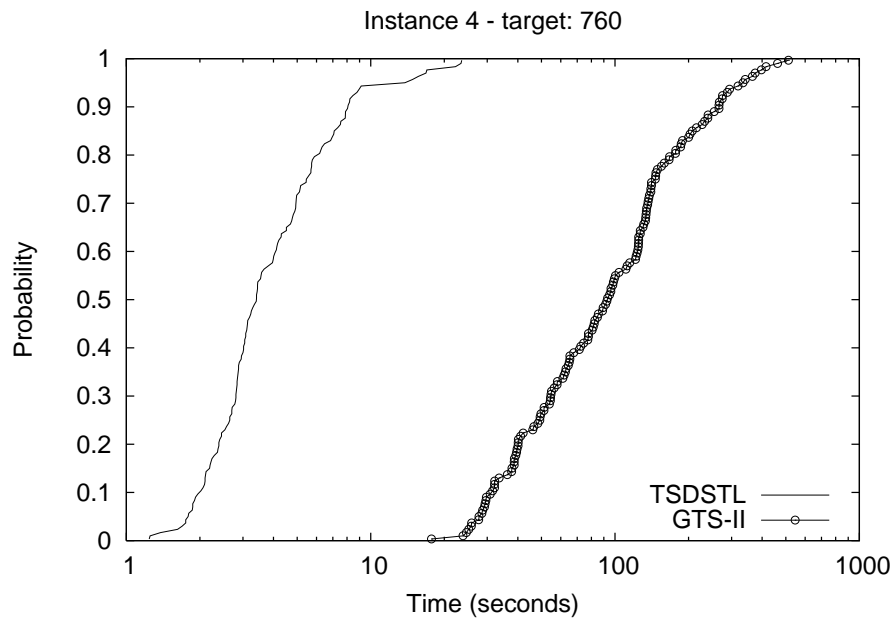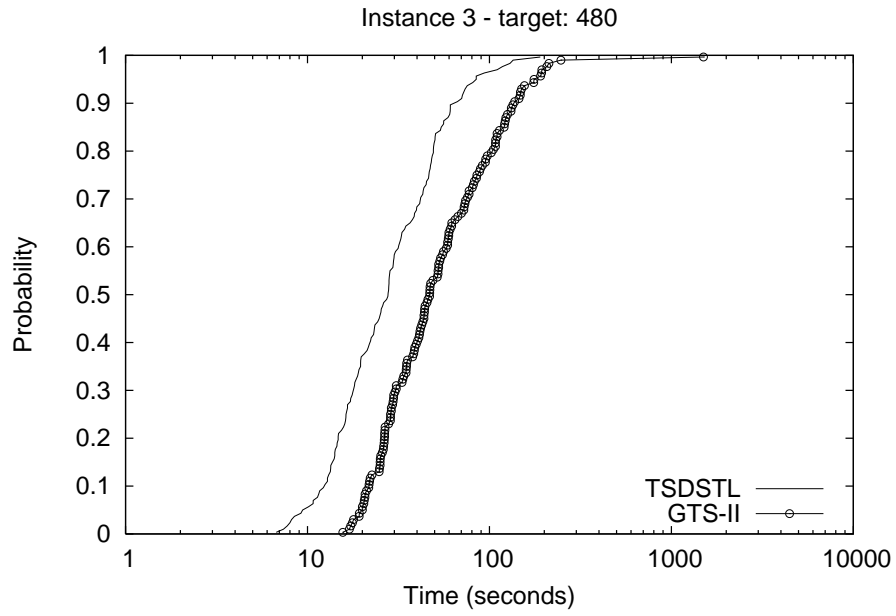
Fig. 4. Empirical probability distribution of finding target values in function of time for instances 3 and 4

Fig. 5. Empirical probability distribution of finding target values in function of time for instances 5 and 6

Fig. 6. Empirical probability distribution of finding target value in function of time for instance 7