



Konrad-Zuse-Zentrum
für Informationstechnik Berlin

Takustraße 7
D-14195 Berlin-Dahlem
Germany

TIMO BERTHOLD

Measuring the impact of primal heuristics

Supported by the DFG Research Center MATHEON *Mathematics for key technologies* in Berlin.

ZIB-Report 13-17 (March 2013)

Herausgegeben vom
Konrad-Zuse-Zentrum für Informationstechnik Berlin
Takustraße 7
D-14195 Berlin-Dahlem

Telefon: 030-84185-0
Telefax: 030-84185-125

e-mail: bibliothek@zib.de
URL: <http://www.zib.de>

ZIB-Report (Print) ISSN 1438-0064
ZIB-Report (Internet) ISSN 2192-7782

Abstract

In modern MIP solvers, primal heuristics play a major role in finding and improving feasible solutions early in the solution process. However, classical performance measures such as time to optimality or number of branch-and-bound nodes reflect the impact of primal heuristics on the overall solving process badly. This article discusses the question of how to evaluate the effect of primal heuristics.

Therefore, we introduce a new performance measure, the “primal integral” which depends on the quality of solutions found during the solving process as well as on the points in time when they are found. Our computational results reveal that heuristics improve the performance of MIP solvers in terms of the primal bound by around 80%. Further, we compare five state-of-the-art MIP solvers w.r.t. the newly proposed measure.

Keywords: primal heuristic, mixed integer programming, performance measure

Mathematics Subject Classification: 90C11, 90C59

1 Introduction

When implementing optimization software, two questions naturally arise: how does the new code perform with respect to existent codes and which are the best settings for a particular algorithm? This goes back to the early days of operations research: Hoffman et al. reported a first computational experiment to compare different implementations of linear programming algorithms in 1953 [4]. Just as researchers and software vendors want to distinguish their code on general test sets, a user wants to tune an optimization software for a particular set of problems. However, all parties require suitable criteria for measuring the performance of a software implementation.

With the rise of computational research, standards and guidelines for conducting computational experiments were proposed [2, 6, 9]. One key issue of the cited articles is the choice of suitable performance indicators. In mathematical programming, the running time to optimality is the “gold standard” for performance comparisons. For branch-and-bound based algorithms, the number of branch-and-bound nodes is another typical measure. Similarly, when using a simplex or an interior point based solver, the number of iterations is commonly used. Both, iterations and nodes attempt to estimate the running time by a measure that is less dependant on the hardware and at the same time better reflects the sheer computational complexity.

The setting In this article, we will use mixed integer (linear) programming as a showcase for our computational experiments. *Mixed integer programming* (MIP) is to solve the optimization problem

$$\tilde{x}_{\text{opt}} = \operatorname{argmin}\{c^T x \mid Ax \leq b, x_j \in \mathbb{Z} \text{ for all } j \in J\},$$

with $A \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$, $c \in \mathbb{R}^n$, and $J \subseteq \{1, \dots, n\}$. Mixed integer programs are a typical example for optimization problems that can be efficiently solved by branch-and-bound.

One advantage of branch-and-bound based algorithms, as opposed to, e.g., pure cutting plane algorithms, is that sub-optimal (though feasible) solutions, so called *incumbents*, often are available early during the course of the algorithm. *Primal heuristics*¹ boost this characteristic even further. As a consequence, sufficiently good solutions might be available long before the branch-and-bound search terminates.

¹A primal heuristic is, roughly speaking, an incomplete algorithm that aims at finding high-quality feasible solutions quickly. In general, it is neither guaranteed to be successful, nor does it provide dual information on the quality of the solution by itself.

Existing measures The time needed to find a first feasible solution, an optimal solution, or a solution within a certain gap to optimality (see, e.g., [5]) are performance criteria that concentrate on the primal part of the solution process. Each of these has its individual strengths and weaknesses. The time to first solution entirely disregards the solution quality: for about one quarter of the MIPLIB 2010 [8] benchmark instances, a trivial solution of all variables set to their lower bound (or all to their upper bound) is feasible, but most of the times useless. Particularly when analyzing heuristics embedded in a complete solver, the time to the first solution mainly measures the time needed for preprocessing and solving the root node relaxation: the MIP solvers CPLEX, GUROBI and XPRESS find solutions for 72, 70, and 64 of the 84 feasible MIPLIB 2010 benchmark instances during root node processing. The time to optimal solution, however, ignores that slightly sub-optimal but practically sufficient solutions might have been found long before. Finally, taking the time to a certain gap is an attempt to balance this, but the choice of the threshold is arbitrary by design.

Altogether, the important consideration for primal heuristics should be the trade-off between speed and solution quality. None of the above measures entirely meets this requirement. In marked contrast, two of the named measures, time to first solution and time to optimality, rather represent extreme points. It is the goal of this paper to introduce a new performance measure that reflects the development of the solution quality over the complete optimization process.

The challenge We argue that standard performance criteria are not well-suited to describe the impact that primal heuristics have within a solver. Take the following observation: On the one hand, for state-of-the-art MIP solvers, the impact of primal heuristics on the overall running time and the number of branch-and-bound nodes is typically minor to negligible² whereas other components such as cutting planes or branching rules change these numbers by a factor of two or three.

On the other hand, the solver vendors, such as CPLEX, GUROBI or XPRESS, seem to consider primal heuristics to be a “trade secret”. It stays unrevealed, which heuristics they use, when those are called, or just how many of them a solver features – whereas for other components, there are plenty of user parameters and statistical output available. One interpretation of this discrepancy might be that primal heuristics are considered a – if not the – crucial part of the software and their value is simply not reflected by the performance measures that we commonly use.

2 The primal integral

In this article, we introduce a new performance measure, in particular for benchmarking primal heuristics, that takes into account the whole solution process. The goal is to measure the progress of the primal bound’s convergence towards the optimal solution over the entire solving time. Therefore, we make use of the *primal gap* of a feasible solution, consider this as a function over time and compute the integral of that function.

Definition 2.1. *Let \tilde{x} be a solution for a MIP, and \tilde{x}_{opt} be an optimal (or best known) solution*

²Recent presentations by software vendors mention values in the range of five to fifteen percent.

for that MIP. We define the primal gap $\gamma \in [0, 1]$ of \tilde{x} as:

$$\gamma(\tilde{x}) := \begin{cases} 0, & \text{if } |c^T \tilde{x}_{opt}| = |c^T \tilde{x}| = 0, \\ 1, & \text{if } c^T \tilde{x}_{opt} \cdot c^T \tilde{x} < 0, \\ \frac{|c^T \tilde{x}_{opt} - c^T \tilde{x}|}{\max\{|c^T \tilde{x}_{opt}|, |c^T \tilde{x}|\}}, & \text{else.} \end{cases}$$

Note that for two feasible MIP solutions \tilde{x}_1, \tilde{x}_2 with $c^T \tilde{x}_1 < c^T \tilde{x}_2$ and $\text{sgn}(c^T \tilde{x}_2) = \text{sgn}(c^T \tilde{x}_{opt})$ it holds that $\gamma(\tilde{x}_1) < \gamma(\tilde{x}_2)$. Now assume that we have available the objective function values of intermediate incumbent solutions and the points in time when those have been found – for a given MIP solver, a certain problem instance and a fixed computational environment.

Definition 2.2. Let $t_{max} \in \mathbb{R}_{\geq 0}$ be a limit on the solution time of a MIP solver. Its primal gap function $p: [0, t_{max}] \mapsto [0, 1]$ is defined as:

$$p(t) := \begin{cases} 1, & \text{if no incumbent until point } t, \\ \gamma(\tilde{x}(t)), & \text{with } \tilde{x}(t) \text{ being the} \\ & \text{incumbent solution at point} \\ & t, \text{ else.} \end{cases}$$

The primal gap function $p(t)$ is a step function that changes whenever a new incumbent is found. It is monotonically decreasing and zero from the point on at which the optimal solution is found.

Definition 2.3. Let $T \in [0, t_{max}]$ and let $t_i \in [0, T]$ for $i \in 1, \dots, I-1$ be the points in time when a new incumbent solution is found, $t_0 = 0, t_I = T$. We define the primal integral $P(T)$ of a run as:

$$P(T) := \int_{t=0}^T p(t) dt = \sum_{i=1}^I p(t_{i-1}) \cdot (t_i - t_{i-1}).$$

We suggest to use $P(t_{max})$ for measuring the quality of primal heuristics. It features two simple, but important attributes: First, whenever a better solution is found at the same point in time, $P(t_{max})$ decreases. Second, whenever the same solution is found at an earlier point in time, $P(t_{max})$ decreases. Briefly: the primal integral favors finding good solutions early. For the performance measures discussed in the introduction, at most one of these two attributes holds in general.

The fraction $P(t_{max})/t_{max}$ can be seen as the average solution quality during the search process. Spoken differently, the smaller $P(t_{max})$ is, the better is the expected quality of the incumbent solution if we stop the solver at an arbitrary point in time. The primal integral is an absolute measure in the sense that it is only defined by a single solver – unlike, for instance, a performance profile³ which compares relative performance and is defined by a set of solvers.

Note that for pure feasibility problems (instances with a zero objective function), the primal integral and the overall running time will give the same measure, up to a constant scaling factor. This follows from the simple observation, that for feasibility instances, the primal gap is one, before the solution is found (and the solution process thereby terminates) and zero afterwards.

³ Performance profiles [3] represent the relative performance of a set of algorithms compared to a virtually best solver as a curve in a graph. A performance profile shows how many percent of the instances, which is the ordinate, in a given test set a given solver could solve within a time factor, which is the abscissa, of the best solver.

Hence when for two different runs the running time differs by a factor of k , the primal integral $p(t)$ will also differ by a factor of k .

In a recent work on rounding heuristics for MINLP, Nannicini and Belotti [11] used the percentage of total running time for which one given algorithm gave rise to a strictly better solution than another one to compare two solution processes. This can be formulated in terms of our notation as

$$\frac{1}{T} \int_{t=0}^T \chi_{\{p_1 > p_2\}}(t) dt$$

with p_1 and p_2 the primal gap functions of two runs of a solver (with different settings) and $\chi_{\{p_1 > p_2\}}(t)$ a characteristic function, being one if $p_1(t) > p_2(t)$ and zero otherwise. Main differences to the primal integral are that this does neither take the actual objective function values of the solutions into account, nor is it an absolute measure since it compares the relative performance of two algorithms.

3 Computational experiments

We conducted experiments with five state-of-the-art MIP solvers: CBC 2.7.0 [12], FICO XPRESS 23.01.06 [13], GUROBI 5.1.0 [14], IBM ILOG CPLEX 12.5.0 [15], and SCIP 3.0.1 [1] compiled with SOPLEX 1.7.1 [16] as LP solver. CPLEX, GUROBI, and XPRESS are among the fastest commercial MIP solvers, CBC and SCIP among the fastest open-source MIP solvers [10]. Further, these are the five codes that have been used to compile the MIPLIB 2010 [8] benchmark set which we chose as a test set for our computational experiments. We excluded the three infeasible instances `ash608gpia-3col`, `enlight14`, and `ns1766074`. Additionally, we excluded the instance `mspp16` because SCIP and CBC ran out of memory. This leaves 83 instances in the test set.

The results for running the solvers in default mode are taken from the benchmarks for optimization software webpage of Hans Mittelmann [10], as of 20. February 2013. Additional results with disabled primal heuristics were obtained on the same computer, a 64bit Intel Xeon X5680 CPU at 3.20GHz with 12 MB cache and 32 GB main memory, running an OPENSUSE 12.1 with a GCC 4.6.2 compiler. Turboboost was disabled. In all experiments, there was only one job at a time to avoid random noise in the measured running time that might be caused by cache-misses if multiple processes share common resources.

As a first test, we compare the performance of SCIP and CPLEX when running with and without primal heuristics. We show the evolution of the primal gap in Figures 1 and 2. The dashed line corresponds to the average (taken over 83 instances) primal gap function, when running the solver in default mode. The dark area corresponds to the average primal integral. Accordingly, the dotted line and the light shaded (plus the dark shaded) area correspond to the average primal gap function and the average primal integral when running the solver without heuristics.

For SCIP, the average value of $P(t_{\max})/t_{\max}$ was 9.05% when using primal heuristics and 16.18% without. For CPLEX, it was 1.92% and 3.58%, respectively. This indicates, that on this test set, for these two solvers, primal heuristics lead to an improvement of 78.8% and 86.5% in the primal bound, on average. As a comparison, the running time to optimality was 819.0 seconds with and 910.5 seconds without primal heuristics for SCIP (+11%); 181.1 seconds and 239.9 seconds for CPLEX (+32%). Both solvers solved four instances less when not using primal heuristics. Note that for both solvers the “default” function is strictly smaller than the “no heur” function. This implies that, independent of the chosen time limit, using primal heuristics is superior in terms of the primal integral.

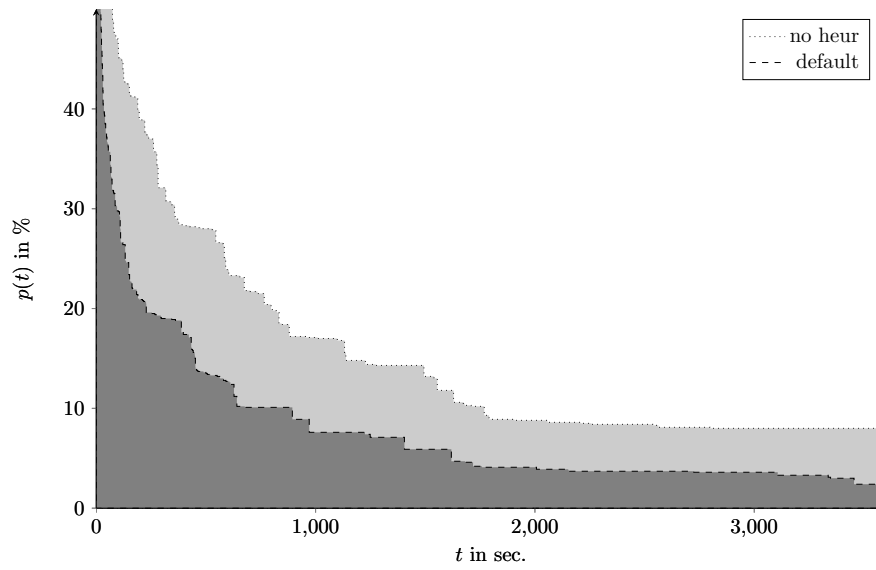


Figure 1: Course of the primal gap when running SCIP with and without primal heuristics

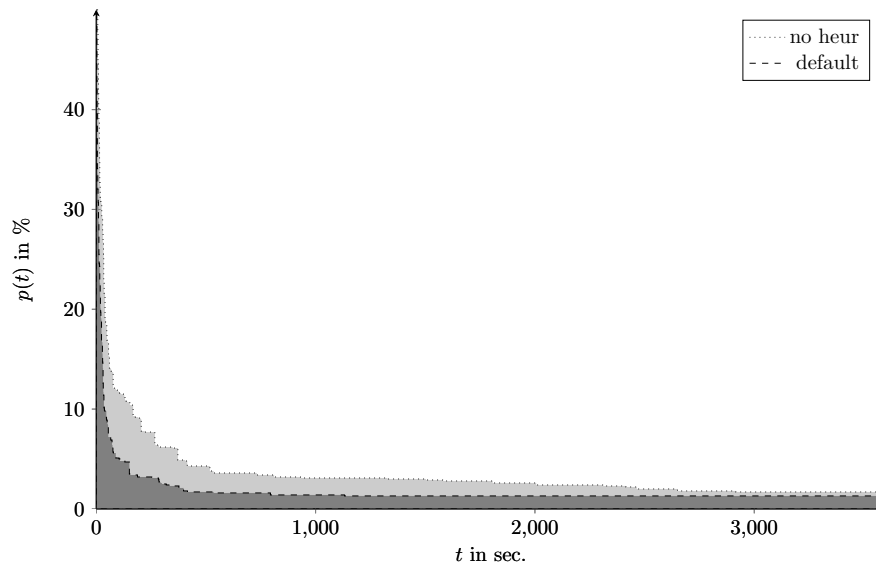


Figure 2: Course of the primal gap when running CPLEX with and without primal heuristics

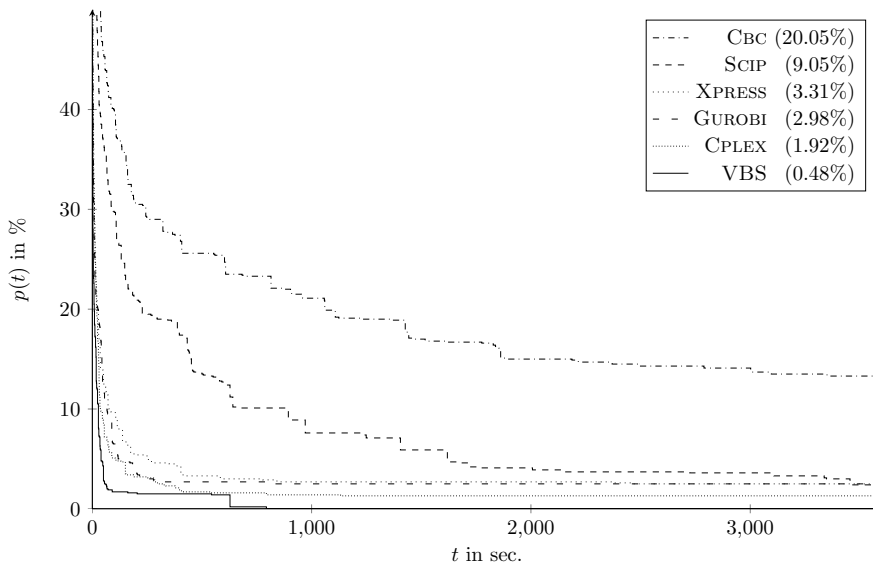


Figure 3: Course of the primal gap for five different MIP solvers plus a virtual best solver

Of course, we can use the primal integral as another metric to compare the performance of solvers against each other. It has been argued in [8] (and previously in [6]) that more than one performance indicator should be used for evaluating mathematical programming software, ideally based on different aspects of the optimization process. Figure 3 shows the course of the average primal gap function for CBC, SCIP, XPRESS, GUROBI, CPLEX and a *virtual best solver* (VBS). VBS takes for each instance the minimum of the primal bounds of the five solvers at each point in time. The results for the average primal integral, in particular the order of the solvers, are clearly different from those for the mean time to optimality as it is shown at [10], which shows that different solvers have different strengths.

Even more important, averages tell you little about a single instance. The primal integral of the virtual best solver in Figure 3 is a factor 3.8 smaller than the primal integral of the best individual solver, meaning that the portfolio of solvers is significantly better than any single solver. For 82 of the 83 instances, at least one solver found an optimal⁴ solution within 800 seconds. In contrast, for each individual solver, there are at least four instances for which the solver did not find an optimal solution after one hour. Each solver contributes to VBS, meaning that each solver has the single best primal bound for some instances for some time. Altogether this shows once again (compare [8]), that having a portfolio of MIP solvers is beneficial.

4 Variants and Extensions

Two main directions for modifications of the primal integral are (i) using a different base measure $p(t)$ and (ii) extending the integral function $P(T)$.

Concerning (i), this article suggested to take the integral over a primal gap function defined by the current incumbent and a best known feasible solution. Analogously, we can define a *dual integral* by considering a dual gap function between the current and a best known dual

⁴Here, we consider a solution \tilde{x} optimal when $\gamma(\tilde{x}) \leq 10^{-6}$.

bound. Using the “gap closed” (one of the first appearances was in 1985 [7]) as a measure for the performance of a cutting plane algorithm has the same pitfall as using the objective function of the best found solution for primal heuristics: it only considers the final state, ignoring the path that led there.

Taking the integral over the primal-dual gap as reported by most general purpose MIP solvers can serve as a measure of their convergence speed. Again, this might be particularly worthwhile when a testset contains many instances which hit an imposed time limit. Reporting the gap at termination is prone to variations caused by bound changes around the time limit. Using a *primal-dual integral* instead reduces the impact of events that happen around the time limit: if as an extreme example one solver improves the bound after 3599 seconds and the other after 3601, the primal-dual integral will differ by less than 0.1%, the gap after one hour can be arbitrarily different. Unlike the primal integral, a primal-dual integral does not even require the value of an optimal or best known solution as an input.

Concerning (ii), logarithmic scales are often used to put the focus on the factor between measured values rather than on their absolute difference. It can be argued for the time axis as well as for the gap axis that it makes sense to rather use a logarithmic scale – and incorporate the logarithm into the definition of $P(T)$.

We used primal heuristics for mixed integer linear programming as a showcase in this article, but never exploited specific structures of this problem class. The suggested measures can be used for any class of optimization problem which features a meaningful primal or dual bound and any algorithm that produces a monotonic sequence of bound values. Similarly, instead of a gap function, any performance measure which evolves monotonic over time could be used.

5 Conclusion

In this article, we introduced a new performance measure to evaluate mathematical programming software. The *primal integral* takes the development of the incumbent solution over time into account, thereby favoring finding good solutions early. It is less prone, though not immune, to common weaknesses of standard performance measures, notably the dependence on an (arbitrarily chosen) time limit.

We argued that the primal integral is particularly useful to measure the impact of primal heuristics. To this end, we showed that for two state-of-the-art MIP solvers, the primal integral changes by a factor of nearly two when disabling primal heuristics. The running time to optimality only increased moderately in the same experiment.

We conclude that the primal integral and its variants (e.g., a primal-dual integral) are a valuable extension of the portfolio of available performance measures. We suggest to take it into account whenever analyzing algorithms that work particularly on the primal bound within an optimization procedure.

Acknowledgments

We thank Gregor Hendel, who provided the python scripts that compute the primal integral and plot the corresponding figures. We are further indebted to Hans Mittelmann who made the log files available to us. Many thanks to Jonas Schweiger for proof-reading a first draft of this paper. This research has been supported by the DFG Research Center MATHEON *Mathematics for key technologies* in Berlin.

References

- [1] T. Achterberg. SCIP: Solving Constraint Integer Programs. *Mathematical Programming Computation*, 1(1):1–41, 2009.
- [2] H. P. Crowder, R. S. Dembo, and J. M. Mulvey. Reporting computational experiments in mathematical programming. *Mathematical Programming*, 15:316–329, 1978.
- [3] E. D. Dolan and J. J. Moré. Benchmarking optimization software with performance profiles. *Mathematical Programming*, 91:201–213, 2002.
- [4] A. Hoffman, M. Mannon, D. Sokolowsky, and N. Wiegmann. Computational experience in solving linear programs. *Journal of the Society for Industrial and Applied Mathematics*, 1(1):17–33, 1953.
- [5] K. L. Hoffman and M. Padberg. Solving airline crew scheduling problems by branch-and-cut. *Management Science*, 39(6), 1993.
- [6] R. H. F. Jackson, P. T. Boggs, S. G. Nash, and S. Powell. Guidelines for reporting results of computational experiments. report of the ad hoc committee. *Mathematical Programming*, 49:413–425, 1991.
- [7] E. Johnson, M. Kostreva, and U. Suhl. Solving 0-1 integer programming problems arising from large scale planning models. *Operations Research*, 33:803–819, 1985.
- [8] T. Koch, T. Achterberg, E. Andersen, O. Bastert, T. Berthold, R. E. Bixby, E. Danna, G. Gamrath, A. M. Gleixner, S. Heinz, A. Lodi, H. Mittelman, T. Ralphs, D. Salvagnin, D. E. Steffy, and K. Wolter. MIPLIB 2010. *Mathematical Programming Computation*, 3(2):103–163, 2011.
- [9] C. C. McGeoch. Toward an experimental method for algorithm simulation. *INFORMS Journal on Computing*, 8(1):1–15, 1996.
- [10] H. Mittelman. Decision tree for optimization software: Benchmarks for optimization software. <http://plato.asu.edu/bench.html>, accessed 5. March 2013.
- [11] G. Nannicini and P. Belotti. Rounding-based heuristics for nonconvex MINLPs. In P. Bonami, L. Liberti, A. J. Miller, and A. Sartenaer, editors, *Proceedings of the EWMINLP*, pages 159–167, April 2010.
- [12] CBC user guide – COIN-OR. <http://www.coin-or.org/Cbc>.
- [13] FICO Xpress Optimizer. <http://www.fico.com/en/Products/DMTools/xpress-overview/Pages/Xpress-Optimizer.aspx>.
- [14] Gurobi optimization. <http://www.gurobi.com/>.
- [15] IBM ILOG CPLEX Optimizer. <http://www-01.ibm.com/software/integration/optimization/cplex-optimizer/>.
- [16] SoPlex. An open source LP solver implementing the revised simplex algorithm. <http://soplex.zib.de/>.