

Adaptive Subproblem Selection in Benders Decomposition for Survivable Network Design Problems

Tim Donkiewicz*

Abstract

Scenario-based optimization problems can be solved via Benders decomposition, which separates first-stage (master problem) decisions from second-stage (subproblem) recourse actions and iteratively refines the master problem with Benders cuts. In conventional Benders decomposition, all subproblems are solved at each iteration. For problems with many scenarios, solving only a selected subset can reduce computation. We quantify the potential in selecting only those subproblems that yield cuts, and develop subproblem scoring and selection strategies. The proposed multi-criteria scoring methods combine historical subproblem performance metrics with problem-specific features, trained online via logistic regression to adapt to the changing likelihood of subproblem usefulness. Multiple stopping criteria balance exploration and exploitation: cut limits, proportional solve limits, and score thresholds. We evaluate our approach using a variant of the survivable network design problem as an example.

Computational experiments on 135 test instances demonstrate the potential and practical performance of subproblem selection. Analysis reveals that 52.1% of all subproblems solved are unnecessary (they contribute no cuts and occur outside cut-free rounds). An oracle with perfect foresight reduces total solve times by 34.4%. Random selection performs significantly worse than full enumeration, showing that naive strategies can degrade performance. Our best-scoring and selection method achieves statistically significant improvements in both runtime and primal-dual integrals. We show that informed subproblem selection can improve Benders decomposition, while highlighting challenges in developing reliable prediction models.

Keywords: Integer programming, Benders decomposition, subproblem selection, survivable network design, machine learning, operations research.

*Chair of Operations Research, RWTH Aachen University, Germany. E-mail: donkiewicz@or.rwth-aachen.de. ORCID: 0000-0002-5721-3563

1 Introduction

Benders decomposition [5] is a powerful technique for structured optimization problems, separating first-stage decisions from second-stage recourse problems. In an iterative manner, the master problem is solved to propose first-stage solutions. Then, subproblems corresponding to different scenarios are solved to check feasibility or optimality of the proposed solution. If any subproblem is infeasible or more expensive than expected by the master problem’s estimation, Benders cuts are generated and added to the master problem to refine the solution space. This process repeats until the branch-and-bound algorithm terminates. Modern implementations of Benders decomposition apply Branch-and-Benders-cut [24], integrating Benders cuts as lazy constraints within a branch-and-bound framework for mixed-integer programming (see Figure 1).

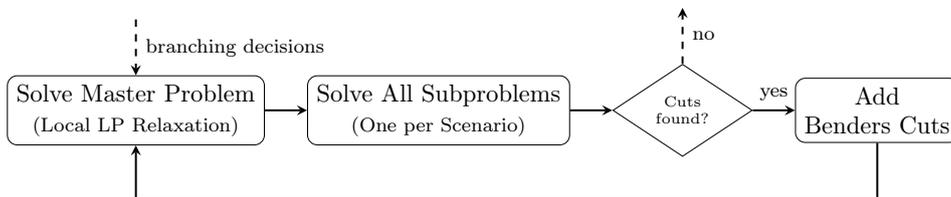


Figure 1: Basic Branch-and-Benders-cut framework.

However, when the number of scenarios grows large, solving all subproblems at each iteration becomes computationally prohibitive. Through computational experiments on survivable network design (SND) instances, we observe that 52.1% of subproblems solved under conventional Benders decomposition are unnecessary. This empirical observation motivates a data-driven computational approach to adaptive subproblem selection—we aim to dynamically learn which subproblems are most likely to contribute to the solving process by generating Benders cuts. To that end, we observe patterns in cut generation behavior, then prioritize subproblems accordingly, and stop once enough subproblems were solved (and at least one cut was generated).

We use SND as a testbed for developing adaptive subproblem selection strategies. SND problems exhibit characteristics that make them particularly suitable for this study: they naturally decompose into many subproblems (one per failure scenario), these subproblems vary substantially in their importance, and some scenarios can dominate others due to network topology. While our goal is not to advance the state-of-the-art for SND solving specifically, this problem class provides a suitable environment to demonstrate subproblem selection techniques.

To our knowledge, adaptive subproblem selection based on computational pattern learning has not been successfully demonstrated for Benders decomposition in prior work. We address this gap through four main contributions grounded in empirical analysis: First, we develop an oracle baseline

that quantifies the computational potential of perfect subproblem selection through empirical measurements. Second, we develop a multi-criteria scoring mechanism that combines observed historical performance metrics (cut generation frequency, cumulative contribution, recency) with computational features extracted from the current solution (failed edge capacity, flow, utilization, centrality). Third, we introduce an online learning approach using logistic regression that learns from computational observations during solving. We train incrementally on each subproblem outcome to predict scenario infeasibility and adapt to evolving cut generation patterns without requiring problem-specific tuning or offline training. The online training incurs a reasonable computational overhead—one update is comparable to solving a single subproblem—making it practical for real-time deployment. Fourth, we implement partial subproblem solving with multiple stopping criteria that balance computational effort between exploration and exploitation based on empirical performance indicators.

Our goal is twofold: demonstrate through computational experiments that dynamically learning from empirical observations can achieve statistically significant improvements over the baseline, and quantify via the oracle how much additional computational improvement could be attained with more sophisticated pattern recognition.

2 Related Work

Benders [5] first introduced the Benders decomposition algorithm for mixed-integer programming. Rahmaniani et al. [24] provide a comprehensive modern survey and an overview of Branch-and-Benders-cut. Benders decomposition has been applied to survivable network design problems [9, 16]. The SNDlib benchmark instances [19, 20] are widely used for evaluating network design algorithms.

2.1 Scenario Retention, Selection, and Partitioning

An approach to managing computational burden is to retain selected second-stage variables explicitly in the master problem rather than projecting them out through Benders cuts. Crainic et al. [13, 12] develop partial decomposition strategies for two-stage integer programs and stochastic multicommodity network design. Pauphilet et al. [23] show that random variable retention can achieve good performance. These methods fundamentally change the master problem structure.

Rather than modifying the master problem structure, scenario selection methods choose which subproblems to solve at each iteration. Rahmaniani et al. [24] note that solving only a subset of subproblems—especially early in the algorithm—can reduce computational burden, though such strategies remain underexplored for combinatorial optimization. Chen et al. [11] reformulate

the feasibility check for single-commodity survivable network design using an oracle based on minimum cut computations to identify violated scenarios to examine. Zhang et al. [29] propose optimization-based scenario reduction to approximate the recourse function while maintaining tight optimality gaps. Blanchot et al. [7] stop solving subproblems after the sum of the violation of obtained optimality cuts exceeds the current gap, often solving only 1% of subproblems. Mehamdi and Gendron [22] select pricing subproblems to solve in partial pricing in Branch-and-Price based on dual information. Celik and Toriello [10] develop a scenario-retention strategy that avoids resolving subproblems remaining feasible across iterations. Learning-based approaches have also emerged: Borozan et al. [8] use machine learning to predict which scenarios generate cuts binding in the final solution. These approaches adaptively select which subproblems to examine without modifying the problem structure.

2.2 Scoring Mechanisms for Algorithmic Decisions

The use of scoring functions to prioritize algorithmic decisions is common in mixed-integer programming [1, 2, 3], and multi-criteria scoring methods often outperform, e.g., pure violation-based approaches [18, 4]. Recent work has also explored machine learning for various selection decisions in MIP solvers, applying neural networks and reinforcement learning to cut selection, variable selection, and cut generation, learning to predict effectiveness from historical data and solution trajectories [25, 26, 27, 28, 30]. Li et al. [21] use machine learning to configure separator selection in branch-and-cut. Deza et al. [14] provide a comprehensive survey showing that while machine learning approaches demonstrate promise, they require substantial training data and careful feature engineering.

These principles from MIP solver design inform our multi-criteria scoring for subproblem selection, where we decide which feasibility checks to perform rather than selecting cuts from an existing pool.

2.3 Contributions in the Context of Related Works

Prior work on subproblem selection designs auxiliary optimization problems or heuristics to approximate subproblem objectives and identify scenarios to examine, often leveraging problem-specific structure. Our approach differs by scoring subproblems using multiple criteria and adaptively determining how many to solve without requiring problem reformulation. Through online learning from empirical observations during solving, we dynamically adapt prioritization to each instance and solution trajectory. We observe that for our test instances with feasibility cuts only, omitting even few cuts significantly degrades convergence, highlighting the challenge of effective subproblem selection.

Our multi-criteria scoring framework adapts principles from MIP solver decision-making (subsection 2.2) to decide which subproblems to solve before incurring costs, reducing computational effort by avoiding unnecessary solves entirely.

We do not aim to advance the state-of-the-art for survivable network design specifically, but rather use it as a representative problem class to develop and evaluate adaptive subproblem selection strategies applicable to Benders decomposition more broadly.

3 Problem Formulation

Consider a network $G = (V, E)$ with node set V and undirected edge set E . For each edge $\{i, j\} \in E$, we create two directed arcs: (i, j) and (j, i) . Let A denote the set of all such directed arcs. For any node $v \in V$, we denote by $\delta^+(v)$ the set of arcs leaving v and by $\delta^-(v)$ the set of arcs entering v . We denote by $A(e)$ the two arcs corresponding to edge $e \in E$.

For each edge $e \in E$, capacity can be installed using a set of *capacity modules* M_e . Each module type $m \in M_e$ provides capacity $u_{e,m} > 0$ at installation cost $c_{e,m} \geq 0$. Edges may have preinstalled capacity (at zero cost). The number of installed modules m on each edge e is represented by $y_{e,m} \in \mathbb{N}$.

A set of demands D must be routed, where each demand $d \in D$ specifies a source node $o_d \in V$, destination node $t_d \in V$, and demand value $b_d > 0$. For convenience, define $b_{d,v} = b_d$ if $v = o_d$, $b_{d,v} = -b_d$ if $v = t_d$, and $b_{d,v} = 0$ otherwise.

We consider a set S of *failure scenarios*, where each scenario $s \in S$ is characterized by a subset $F_s \subseteq E$ of failed edges. We also define s_0 as the *base scenario* with $F_{s_0} = \emptyset$ (no failures), and let $S_0 = \{s_0\} \cup S$ denote all scenarios including the base case. The formulation supports arbitrary failure sets (single-edge, multi-edge, or scenario-specific combinations). In this paper, we focus on *single-edge failure scenarios*: $S = \{s_e : e \in E\}$ where $F_{s_e} = \{e\}$ (edge e fails), commonly known as N-1 reliability in telecommunications and power systems. The generalization to arbitrary failure sets is straightforward and does not affect the adaptive subproblem selection methodology.

This problem can be viewed as a two-stage integer program where, in each failure scenario, the demands must be routed. In the first stage, we install capacity modules on edges (integer decisions made before failures occur) and determine flow routing for the base scenario s_0 . Then, in the second stage, for each scenario $s \in S$, we route demands on non-failed edges respecting installed capacities (recourse decisions after failures are revealed). The objective is to minimize total module installation costs.

Note that the second-stage decisions (flow routing) serve as recourse actions but have zero cost—we only consider routing feasibility, not routing

costs. The objective is purely to minimize first-stage installation costs, but the installed capacity must be sufficient to enable feasible routing under all failure scenarios.

3.1 Compact Formulation

Let $y_{e,m} \in \mathbb{N}$ be the number of modules of type m installed on link e and $f_{s,d,a} \in \mathbb{R}_+$ the flow of demand d on arc a in scenario s . We then consider the following compact formulation:

$$\min \sum_{e \in E} \sum_{m \in M_e} c_{e,m} y_{e,m} \quad (1)$$

$$\text{s.t.} \quad \sum_{a \in \delta^+(v)} f_{s,d,a} - \sum_{a \in \delta^-(v)} f_{s,d,a} = b_{d,v} \quad \forall s \in S_0, d \in D, v \in V \quad (2)$$

$$\sum_{d \in D} \sum_{a \in A(e)} f_{s,d,a} \leq \sum_{m \in M_e} u_{e,m} y_{e,m} \quad \forall s \in S_0, e \in E \setminus F_s \quad (3)$$

$$f_{s,d,a} = 0 \quad \forall s \in S_0, d \in D, e \in F_s, a \in A(e) \quad (4)$$

$$y_{e,m} \in \mathbb{N} \quad \forall e \in E, m \in M_e \quad (5)$$

$$f_{s,d,a} \geq 0 \quad \forall s \in S_0, d \in D, a \in A \quad (6)$$

The capacity constraint (3) applies only to operational links $e \in E \setminus F_s$ in each scenario. For failed links $e \in F_s$, constraint (4) explicitly forces flow to zero on both arcs $a \in A(e)$ of the failed link, prohibiting any routing through failed components.

With $|S_0| = |E| + 1$ scenarios (base case plus all single-edge failures), we have $O(|S_0| \cdot |D| \cdot |A|)$ flow variables and $O(|S_0| \cdot (|D| \cdot |V| + |E|))$ constraints. When considering multi-edge failures, the number of scenarios grows exponentially, making the compact formulation impractical. For large networks, even solving the single-edge failure case becomes computationally intractable using conventional MIP solvers.

3.2 Benders Decomposition

Benders decomposition addresses the scalability issue by separating first-stage and second-stage decisions into a master problem and subproblems. Since we only consider feasibility in the second stage, we only use Benders *feasibility* cuts to iteratively refine the master problem.

Master Problem The Benders master problem includes first-stage capacity decisions and base case flow constraints:

$$\min \sum_{e \in E} \sum_{m \in M_e} c_{e,m} y_{e,m} \quad (7)$$

$$\text{s.t.} \quad \sum_{a \in \delta^+(v)} f_{s_0,d,a} - \sum_{a \in \delta^-(v)} f_{s_0,d,a} = b_{d,v} \quad \forall d \in D, v \in V \quad (8)$$

$$\sum_{d \in D} \sum_{a \in A(e)} f_{s_0,d,a} \leq \sum_{m \in M_e} u_{e,m} y_{e,m} \quad \forall e \in E \quad (9)$$

$$\text{Benders cuts (17)} \quad (10)$$

$$y_{e,m} \in \mathbb{N} \quad \forall e \in E, m \in M_e \quad (11)$$

$$f_{s_0,d,a} \geq 0 \quad \forall d \in D, a \in A \quad (12)$$

Constraints (8) and (9) enforce flow conservation and capacity constraints for the base case scenario s_0 with no failed edges ($F_{s_0} = \emptyset$), ensuring basic feasibility before considering failure scenarios.

Subproblem For each failure scenario $s \in S$ and candidate solution \bar{y} , the subproblem checks feasibility:

$$\min \quad 0 \quad (13)$$

$$\text{s.t.} \quad \sum_{a \in \delta^+(v)} f_{s,d,a} - \sum_{a \in \delta^-(v)} f_{s,d,a} = b_{d,v} \quad \forall d \in D, v \in V \quad (14)$$

$$\sum_{d \in D} \sum_{a \in A(e)} f_{s,d,a} \leq \sum_{m \in M_e} u_{e,m} \bar{y}_{e,m} \quad \forall e \in E \setminus F_s \quad (15)$$

$$f_{s,d,a} = 0 \quad \forall d \in D, e \in F_s, a \in A(e) \quad (16)$$

Constraints (14) enforce flow conservation (corresponding to (2) in the compact formulation), and constraints (15) enforce capacity limits on operational edges (corresponding to (3)). For failed edges $e \in F_s$, constraints (16) explicitly set flow to zero on both arcs, ensuring no routing through failed components.

We observe that this is a multicommodity flow problem with fixed capacities determined by the first-stage solution \bar{y} . It is solvable in polynomial time as a linear program with $O(|D| \cdot |A|)$ variables and $O(|D| \cdot |V| + |E|)$ constraints.

If the subproblem is feasible, the capacity \bar{y} is sufficient for scenario s . If infeasible, we generate a Benders feasibility cut using the subproblem's dual solution.

Benders Feasibility Cuts When the subproblem for scenario s is infeasible, we obtain an unbounded ray (π^s, σ^s) of the dual problem, where $\pi_e^s \geq 0$ are the dual variables for capacity constraints (15) and $\sigma_{d,v}^s$ are the (unrestricted) dual variables for flow conservation constraints (14). The Benders feasibility cut is:

$$\sum_{e \notin F_s} \pi_e^s \sum_{m \in M_e} u_{e,m} y_{e,m} \geq - \sum_{d \in D} \sum_{v \in V} \sigma_{d,v}^s b_{d,v} \quad (17)$$

where we sum only over operational edges $e \notin F_s$ since failed edges have zero available capacity.

4 Adaptive Subproblem Selection

Our data-driven approach is based on two observations. Firstly, it is not required to solve all subproblems at each Benders iteration to make progress toward feasibility. Secondly, not all subproblems are equally likely to generate cuts. If we are able to predict which scenarios are unlikely to produce a cut, we can avoid solving those subproblems entirely, saving computation time.

We therefore propose Branch-and-Benders-cut with *adaptive subproblem selection*: prioritize which subproblems to solve first based on learned importance, and limit the number of subproblems solved per iteration, as illustrated in Figure 2. At each Benders iteration, we first extract features from the master solution (topology metrics like capacity, flow, and utilization). We then score all scenarios based on their predicted likelihood of generating cuts using a machine learning model. Scenarios are examined in decreasing score order, with stopping criteria determining when to terminate subproblem solving. If there exists any, at least one cut is added in each round. After each subproblem is solved, we train the regression model online using the observed outcome and updated historical performance (reliability, violation magnitude). Feasibility cuts from infeasible subproblems are added to the master, and the process repeats until convergence.

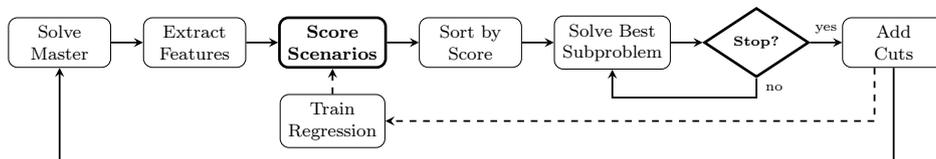


Figure 2: Adaptive subproblem selection pipeline in Benders decomposition. Scoring and subproblem selection step are highlighted.

In this section, we first describe an oracle that selects an optimal subproblem solving sequence to not miss any cuts. Then, we present our multi-criteria subproblem scoring mechanism in Section 4.2. Building upon the scoring system, we introduce stopping criteria for subproblem solving iterations in Section 4.3, aiming to approximate the oracle’s behavior.

4.1 Oracle for Evaluating Prediction Quality

To assess scoring quality, we define a *cut prediction oracle* \mathcal{O}^* with perfect foresight of all cuts generated until the algorithm converges. Let $\Pi = (P_1, P_2, \dots, P_T)$ denote a subproblem selection sequence, where $P_t \subseteq S$ are the examined scenarios in iteration t . We define the set of cut-yielding

subproblems in iteration t as $V_t \subseteq S$. The *optimal subproblem selection* as given by \mathcal{O}^* is $\Pi^* = (V_1, V_2, \dots, V_t)$, the sequence consisting of exactly those subproblems yielding a cut. Note that a sequence $\Pi = (P_1, P_2, \dots, P_T)$ with $P_t \cap V_t \subsetneq V_t$, i.e., missing some cut-yielding subproblems, may lead to lower overall runtime, but we focus on approximating the oracle’s behavior.

Definition 1 (Cut Prediction Oracle Bound). *Let τ_s^t denote the time to solve subproblem s at iteration t . Let $T(\Pi)$ denote the total subproblem solving time over all iterations for sequence Π . The cut prediction oracle achieves:*

$$T(\Pi^*) = \sum_{t=1}^T \sum_{s \in V_t} \tau_s^t \leq \sum_{t=1}^T \sum_{s \in S} \tau_s^t = T(\Pi_{all}). \quad (18)$$

Since identical cuts are added, the runtime behavior is equal otherwise, apart from effects of warmstarting.

In general, this cut prediction oracle is not computable a priori, but rather serves as a benchmark. Our goal is to design a scoring mechanism that approximates the oracle’s performance by prioritizing subproblems likely to generate cuts.

4.2 Subproblem Scoring

For each scenario $s \in S$, we maintain a score $\text{score}(s) \in [0, 1]$ that estimates its likelihood of generating a violated cut. We present a machine learning-based scoring approach that uses network topology features and historical performance metrics (Section 4.2.2).

4.2.1 Score Components

We utilize information derived from the master problem solution, in conjunction with the failure scenario and from historical scenario performance. Let (\bar{y}, \bar{f}) denote the current master solution where $\bar{y}_{e,m}$ is the number of modules of type m installed on link e , and $\bar{f}_{s_0,d,a}$ is the flow of demand d on arc a in the base case scenario s_0 (no failures). Let C_s denote the set of iterations where scenario s generated a cut, and I_s denote iterations where s was examined. Let $\nu_s^{t'}$ denote the violation magnitude of the Benders cut generated by scenario s at iteration $t' \in C_s$ (measured as the right-hand side value of the cut). Let betweenness centrality β_e be the fraction of shortest paths (in the network topology, without considering modules) between each demand pair that contain edge e . Using these definitions, we introduce eight score components as described in Table 1. Note that the violation magnitude $\nu_s^{t'}$ for feasibility cuts is not necessarily directly comparable, since the dual ray may scale arbitrarily.

Table 1: Score components for scenario prioritization.

Name	Definition	Description
Reliability	$\rho_s = C_s / I_s $	Fraction of times s produced a cut when solved
Total Share	$\tau_s = C_s /\sum_{s' \in S} C_{s'} $	Fraction of all cuts produced by s
Average Violation	$\bar{\nu}_s = \frac{1}{ C_s } \sum_{t' \in C_s} \nu_{s'}^{t'}$	Mean violation magnitude when s generated cuts
Staleness	$\zeta_s = t - \max\{t' : t' \in I_s\}$	Iterations since s was last examined
Failed Edge Capacity	$\kappa_e = \sum_{m \in M} u_{e,m} \bar{y}_{e,m}$	Total capacity on failed edge $e \in F_s$
Failed Edge Flow	$\psi_e = \sum_{d \in D} (\bar{f}_{s_0,d,a} + \bar{f}_{s_0,d,a'})$	Base case flow on failed edge $e \in F_s$
Failed Edge Utilization	$\mu_e = \psi_e/\kappa_e$	Flow-to-capacity ratio on $e \in F_s$ (when $\kappa_e > 0$)
Betweenness Centrality	β_e	Topological centrality of edge $e \in F_s$ (computed once)

A straightforward approach would combine score components via weighted linear combination $R(s) = w_\rho \rho_s + w_\tau \tau_s + w_\zeta \zeta_s + \dots$ with min-max normalization. However, determining appropriate weights is challenging: they vary across instances and solution phases. For problems with long solve times, static weights become problematic as the optimal balance shifts dynamically. This motivates the machine learning approach described next, which learns feature importance automatically.

4.2.2 Machine Learning-Based Scoring

Our regression-based approach (in the following, referred to as machine learning, or “ML”) predicts scenario infeasibility using logistic regression trained online during the Benders decomposition process. We construct an 8-dimensional feature vector $\phi(s) \in \mathbb{R}^8$ for each scenario s combining the failed edge capacity κ_e , flow ψ_e , utilization μ_e , betweenness centrality β_e , and the historical components $\bar{\nu}_s$, ρ_s , τ_s , ζ_s described above. Features are normalized to zero mean and unit variance using Welford’s online algorithm, with normalization parameters accumulated across the entire solution process (not reset during stabilization rounds).

Logistic Regression Model We use logistic regression for its low computational complexity during online training and superior interpretability compared to more sophisticated models like neural networks or ensemble methods.

Logistic regression models the probability of binary outcomes via a sigmoid transformation of a linear predictor. The model predicts infeasibility probability as:

$$p(s \text{ infeasible} \mid \phi(s); w, b) = \sigma(w^\top \phi(s) + b) \quad (19)$$

where $w \in \mathbb{R}^8$ is the weight vector, $b \in \mathbb{R}$ is the bias term, and $\sigma(z) = 1/(1 + e^{-z})$ is the logistic function mapping $\mathbb{R} \rightarrow (0, 1)$.

This formulation allows the model to learn non-uniform importance of features: features strongly predictive of infeasibility receive large absolute weights, while less informative features receive weights near zero. The sigmoid ensures output values are valid probabilities.

Online Training via Gradient Descent After solving the subproblem for scenario s at iteration t , we observe the outcome $z_s^t \in \{0, 1\}$ where $z_s^t = 1$ if the subproblem was infeasible (generated a cut) and $z_s^t = 0$ if feasible. We update the model parameters to minimize the logistic loss:

$$\ell(w, b) = -\log p(z_s^t \mid \phi(s); w, b) \quad (20)$$

Applying gradient descent with L2 regularization yields the update rules:

$$w \leftarrow w - \alpha [(\hat{z}_s^t - z_s^t)\phi(s) + \lambda w] \quad (21)$$

$$b \leftarrow b - \alpha(\hat{z}_s^t - z_s^t) \quad (22)$$

where $\hat{z}_s^t = \sigma(w^\top \phi(s) + b)$ is the predicted probability, $\alpha > 0$ is the learning rate controlling step size, and $\lambda \geq 0$ is the regularization parameter preventing overfitting by penalizing large weights.

The gradient term $(\hat{z}_s^t - z_s^t)$ represents the prediction error: positive when the model overestimates infeasibility probability, negative when underestimating. The regularization term λw in (21) shrinks weights toward zero, favoring simpler models that generalize better.

4.3 Subproblem Selection

Table 2: Stopping criteria for partial subproblem solving and naming convention. One or more can be imposed, and a Benders iteration stops subproblem solving when any criterion is met. Stabilization rounds solve all subproblems to prevent selection bias and provide additional training data.

Name	Identifier	Description
Cut limit	kC	Stop if k cuts have been generated
Relative solve limit	pP	Stop after examining a proportion $p \in [0, 1]$ of all scenarios, i.e., $\lceil p \cdot S \rceil$ subproblems
Consecutive misses	mM	Stop if no cuts found for m consecutive scenarios
Score limit	rT	Stop when the score of the next scenario falls below threshold r
Time limit	—	Stop if iteration time limit reached
Stabilization (solving all subproblems)		
Stabilization	nS	Every n iterations
Root node	nR	First n rounds at root

We solve subproblems in order of decreasing score. Using the stopping criteria introduced in Table 2, we aim to examine all cut-generating scenarios, and to stop once we are unlikely to find more cuts. Since limited examination may miss some cut-generating scenarios and, over time, introduce a selection bias, we periodically perform *stabilization rounds* where all scenarios are examined (disabling stopping criteria). This prevents systematic exclusion of subproblems and provides additional training data. Since initial training data is especially important, we always initialize the scores with at least one score initialization round, but up to n stabilization rounds at the root node (nR). At stabilization rounds, we reset the score components ρ_s , τ_s , $\bar{\nu}_s$ and ζ_s to avoid bias from outdated historical data.

5 Computational Experiments

5.1 Experimental Setup

We implement our Branch-and-Benders-cut framework with adaptive cut selection using Gurobi 12.0.3 [17], Julia 1.12.4 [6] with JuMP 1.12 [15], and lazy constraint callbacks. To improve efficiency, we maintain a single subproblem template adjusted for each scenario rather than constructing new subproblems from scratch. We reuse the Gurobi environment, allowing for LP warmstarting. We do not perform separation on fractional solutions as to isolate the impact of subproblem selection.

Table 3 presents the configuration parameters for all evaluated methods. We parameterize those configurations following an analysis of stopping criteria: For each (threshold, proportion, cut limit) individually, we examine how many cuts would have been generated under this setting. A distribution boxplot can be found in Figure 5 in Appendix A. Based on this analysis, we select more conservative (calling more subproblems, possibly solving more unrequired ones) parameters as well as more aggressive ones to evaluate the impact of stopping criteria. In particular, setting `ML-0.1T-0.6P-10C-20S-5R` uses all three main stopping criteria, and sets the limit at the upper quartile of required cuts per iteration (see Figure 5 in Appendix B). That way, we expect to capture most cut-generating scenarios while avoiding excessive unrequired subproblem solves.

Table 3: Configuration parameters for tested methods. ML methods use online logistic regression. See Table 2 for parameter descriptions.

Method	Selection	Threshold	Proportion	Cut Limit	Stabilization	Root Node Stabilization
Standard	All	—	—	—	—	—
Random-0.5P-20S-5R	Random	—	0.5	—	20	5 rounds
ML-0.1T-0.6P-10C-20S-5R	Threshold+Prop	0.1	0.6	10C	20	5 rounds
ML-0.5P-20S-5R	Proportion	—	0.5	—	20	5 rounds
ML-0.5T-5C-20S-5R	Threshold	0.5	—	5C	20	5 rounds
ML-0.7P-20S-5R	Proportion	—	0.7	—	20	5 rounds

Out of the original SNDlib instances [19, 20], only 11 solve within our time limit of three hours, which yields an insufficient sample for meaningful evaluation. We generate test instances from SNDlib benchmark networks by combining 2–5 subgraphs: (1) extract connected subgraphs via BFS from high-degree nodes using proportion $p \in [0.45, 0.65]$ to obtain $p \cdot |V|$ nodes per network; (2) merge subgraphs by connecting high-degree nodes; (3) generate $N-1$ failure scenarios ($|S| = |E|$); (4) filter scenarios that disconnect demand nodes. We generate 135 test instances with 16–229 scenarios, 8–57 nodes, 13–82 edges, and 18–857 demands. The online ML training overhead is negligible, comparable to solving a single subproblem per iteration. For stabilization frequency, we tested values of 5, 10, 20, 50, 100, and 200 iterations, with 20 providing the best balance for SND problems.

All experiments use machines with 2x Intel Xeon L5630 Quad Core processors at 2.13 GHz (8 cores, 16GB RAM). A three-hour time limit is imposed per instance. Statistical comparisons use the shifted geometric mean (shift of 10) to handle zero values.

The parameters for the regression model were determined using a grid search on a set of 14 instances, generated synthetically by the above procedure. We tried configurations with learning rates $\alpha \in \{0.02, 0.05, 0.075, 0.1, 0.15\}$ and regularization $\lambda \in \{0.05, 0.02, 0.01, 0.005, 0.001\}$. The selected parameters ($\alpha = 0.075$, $\lambda = 0.02$) provided the best runtime performance on the tuning set.

5.2 Results

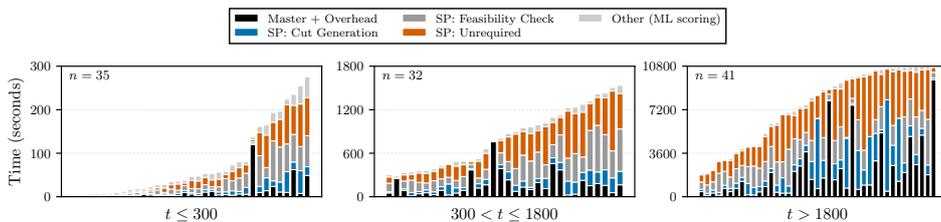


Figure 3: Time breakdown for. Master time (black), cut-generating subproblems (blue), and feasibility-proving subproblems in cut-free iterations (grey) represent necessary computation, unrequired subproblems (orange) represent wasted computation.

The potential for subproblem selection is established through the cut prediction oracle bound (see 1). As illustrated in Figure 3 (and described in detail in Table 5 in Appendix A), 34.4% of total solve time in standard Benders is spent on unrequired solving of subproblems. Since solving all subproblems that yield cuts does not always lead to the best runtime, this is not necessarily an upper bound on the achievable speedup. However, we

Table 4: Summary of runtime performance across all methods ($n = 79$ instances solved by all methods). Solved: instances solved to optimality; Sum/Mean/Sh. Geom. Mean[†]: computed over instances solved by all methods. Wilcoxon signed-rank test (one-sided): tests if each method is significantly better (faster) than Standard. Significance levels: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, n.s. = not significant.

Method	Solved	Sum [†]	Mean [†]	Sh. Geom. Mean [†]	Wilcoxon p	Significance
Standard	86	100941.2	1277.7	300.5	—	—
Random-0.5P-20S-5R	82	112118.3	1419.2	334.8	0.9996	n.s.
ML-0.5T-5C-20S-5R	86	86715.2	1097.7	270.8	0.0130	*
ML-0.5P-20S-5R	89	96097.3	1216.4	285.2	0.0136	*
ML-0.7P-20S-5R	88	100536.4	1272.6	288.3	0.2080	n.s.
0.1T-0.6P-10C-20S-5R	87	92977.6	1176.9	275.3	0.0081	**

consider it as a dimension for comparison.

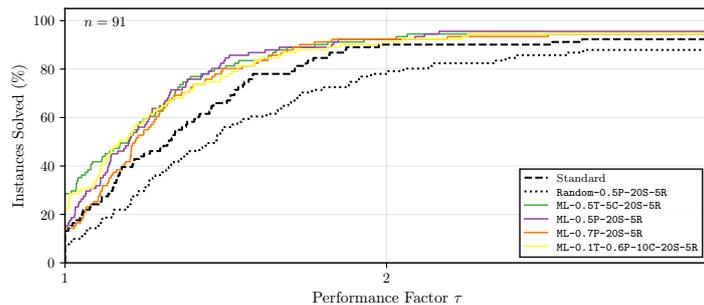


Figure 4: Performance profile on all instances solved by at least one method: Random selection performs worse than standard Benders, regression-based methods achieve modest gains.

Runtime Performance As illustrated in Figure 4 and detailed in Table 4, three out of four tested ML configurations achieve statistically significant speedups over standard Benders, the best being ML-0.1T-0.6P-10C-20S-5R. The reduction from 300.5 to 270.8 represents a speedup of 9.9%, as opposed to a “theoretical” maximum of 34.4% (see Table 5 in Appendix A). Since we aim for reducing unrequired subproblem solving times, the performance profile reduced to only those times shows a more pronounced improvement (see Figure 9 in Appendix F). The full results can be found in Table 7 in Appendix G. We observe that our scoring method indeed outperforms random selection, which is statistically significantly worse than standard Benders ($p = 0.9996$). This confirms that naive subproblem selection strategies can degrade performance, and that our proposed scoring mechanism is able to prioritize important subproblems.

Convergence Quality Beyond runtime performance, primal-dual integral analysis (detailed results in Table 8 in Appendix H) evaluates convergence quality. All four ML configurations achieve statistically significant improvements in convergence quality over standard Benders, with the best method ML-0.5T-5C-20S-5R reducing the shifted geometric mean from 1336.3 to 1145.9 billion (14.3% improvement). Random selection shows no significant difference from standard Benders ($p = 0.12$). These results demonstrate that ML-based subproblem selection not only reduces solve time but also improves the quality of the search by better exploring the solution space.

The ML model achieves an average accuracy of 84.4%, precision of 33.4%, recall of 70.9%, and F1-score of 45.2% (see Table 6 in Appendix B). The mediocre scores mean that we are “missing” cuts, which can lead to worse algorithm convergence that can also be seen in the runtime table (Table 7 in Appendix G).

The ML model’s imperfect predictions can be attributed to two main factors: first, the features used may not fully capture the complex relationships around scenario violations (as can be observed by the high variance in feature importance, see Figure 6 in Appendix C) and second, the online training approach may not obtain sufficient training data for robust learning, especially when instances solve quickly.

To counteract the former, we also tested locality-based features (e.g., utilization in the n -hop neighborhood of failed edges) but they did not receive significant weights in our preliminary experiments. Addressing the latter, we tried pre-training on similar instances, but as Figure 8 in Appendix E shows, cut generation patterns evolve substantially during solving, and so do the weights of the ML model. Additionally, we implemented root node stabilization that serves the purpose of establishing a strong initial ML model and facilitating early primal heuristics. However, further research is needed to develop more reliable prediction models that can consistently deliver robust performance improvements across diverse instances.

6 Conclusions and Future Work

This paper investigated adaptive subproblem selection in Benders decomposition, motivated by the empirical observation that 52.1% of subproblems in our variant of the survivable network design problem unnecessarily solved: The rest either yield cuts (5.6%), or verify feasibility in cut-free iterations (42.4%). We developed a multi-criteria scoring mechanism combining historical performance metrics with topology-aware features, trained online via logistic regression. Multiple stopping criteria (score thresholds, proportional limits, cut limits) adaptively determine how many subproblems to solve per iteration. The online training overhead is negligible—comparable to solving a single subproblem—making the approach practical for real-time deployment.

Computational experiments on 135 survivable network design instances reveal that in standard Benders with full subproblem enumeration, 34.4% of total solve time is spent on unrequired subproblem solving. Our best runtime configuration (ML-0.1T-0.6P-10C-20S-5R) achieves 9.9% speedup over standard Benders (270.8s vs 300.5s shifted geometric mean, $p = 0.0081$), while the best convergence configuration (ML-0.5T-5C-20S-5R) reduces primal-dual integrals by 14.3% ($p < 0.001$). Three of four ML configurations achieve statistically significant runtime improvements, and all achieve statistically significant improvements in primal-dual integrals. While the scoring performance is modest (the regression model achieves only 33.4% precision, 70.9% recall, and 45.2% F1), these results establish that adaptive subproblem selection based on learned importance can improve Benders decomposition performance. Combining multiple stopping criteria appears to be a promising direction, as our best method employs three types simultaneously. Random selection degrades performance significantly ($p = 0.9996$), confirming that informed prioritization is necessary.

Future Research Directions Multiple future research directions can be identified: First, more sophisticated prediction models—such as graph neural networks—could better capture network topology and temporal evolution of cut generation patterns, potentially improving precision and allowing for preconditioning on a set of similar instances. Second, rather than using fixed thresholds, dynamic adaptation of stopping criteria during solving could adjust exploration-exploitation tradeoffs based on observed performance, also taking into account cut usefulness. Third, analyzing feature importance across different solving stages could reveal which metrics are most predictive early versus late in the algorithm, enabling stage-specific scoring. Finally, auxiliary (problem-specific) optimization techniques based on min-cut computations could provide tighter bounds on scenario criticality, complementing the learned scoring mechanism. These directions could help close the gap between current performance and the theoretical potential identified through our analysis of subproblem selection potential.

Acknowledgements

The author sincerely thanks Oliver Gaul for his help in previous work around adaptive subproblem selection and for insightful discussions. The author also thanks Alexander Helber and Marco Lübbecke for their helpful comments.

References

- [1] Achterberg, T.: Constraint integer programming. Ph.D. thesis, Technische Universität Berlin (2007)

- [2] Achterberg, T.: SCIP: solving constraint integer programs. *Mathematical Programming Computation* **1**(1), 1–41 (2009)
- [3] Andreello, G., Caprara, A., Fischetti, M.: Embedding cuts in a branch&cut framework: a computational study with $\{0, \frac{1}{2}\}$ -cuts. *INFORMS Journal on Computing* **19**(2), 229–238 (2007)
- [4] Avella, P., Mattia, S., Sassano, A.: Metric Inequalities and the Network Loading Problem. In: *Discrete Optimization: The State of the Art*, pp. 53–84. Elsevier (2004)
- [5] Benders, J.F.: Partitioning procedures for solving mixed-variables programming problems. *Numerische mathematik* **4**(1), 238–252 (1962)
- [6] Bezanson, J., Edelman, A., Karpinski, S., Shah, V.B.: Julia: A fresh approach to numerical computing. *SIAM review* **59**(1), 65–98 (2017)
- [7] Blanchot, X., Clautiaux, F., Detienne, B., Froger, A., Ruiz, M.: The Benders by batch algorithm: Design and stabilization of an enhanced algorithm to solve multicut Benders reformulation of two-stage stochastic programs. *European Journal of Operational Research* **309**(1), 202–216 (2023)
- [8] Borozan, S., Giannelos, S., Falugi, P., Moreira, A., Strbac, G., Zhang, T.: Machine learning-enhanced Benders decomposition approach for the multi-stage stochastic transmission expansion planning problem. *Electric Power Systems Research* **226**, 109956 (2024)
- [9] Botton, Q., Fortz, B., Gouveia, L., Poss, M.: Benders decomposition for the hop-constrained survivable network design problem. *INFORMS Journal on Computing* **25**(1), 13–26 (2013)
- [10] Çelik, Ş., Martin, L., Schrottenboer, A.H., Van Woensel, T.: Exact two-step Benders decomposition for the time window assignment traveling salesperson problem. *Transportation Science* **59**(2), 210–228 (2025)
- [11] Chen, R.L., Cohn, A., Pinar, A.: An implicit optimization approach for survivable network design. In: *2011 IEEE Network Science Workshop*. pp. 159–163. IEEE (2011)
- [12] Crainic, T.G., Hewitt, M., Maggioni, F., Rei, W.: Partial Benders decomposition: general methodology and application to stochastic network design. *Transportation Science* **55**(2), 414–435 (2021)
- [13] Crainic, T.G., Hewitt, M., Rei, W.: Partial decomposition strategies for two-stage stochastic integer programs. *CIRRELT (Université de Montréal)* **88** (2014)

- [14] Deza, A., Khalil, E.B.: Machine learning for cutting planes in integer programming: A survey. arXiv preprint arXiv:2302.09166 (2023)
- [15] Dunning, I., Huchette, J., Lubin, M.: JuMP: A modeling language for mathematical optimization. *SIAM Review* **59**(2), 295–320 (2017)
- [16] Fortz, B., Poss, M.: An improved Benders decomposition applied to a multi-layer network design problem. *Operations Research Letters* **37**(5), 359–364 (2009)
- [17] Gurobi Optimization, LLC: Gurobi Optimizer Reference Manual (2024), <https://www.gurobi.com>
- [18] Hosseini, M., Turner, J.: Deepest Cuts for Benders Decomposition. *Operations Research* **73**(5), 2591–2609 (2024)
- [19] Idzikowski, F., Orłowski, S., Raack, C., Woesner, H., Wolisz, A.: Dynamic routing at different layers in IP-over-WDM networks – Maximizing energy savings. *Optical Switching and Networking, Special Issue on Green Communications* (2011)
- [20] Koster, A.M., Kutschka, M., Raack, C.: Towards Robust Network Design using Integer Linear Programming Techniques. In: *Proceedings of the NGI 2010, Paris, France. Next Generation Internet, Paris, France* (June 2010)
- [21] Li, S., Ouyang, W., Paulus, M.B., Wu, C.: Learning to Configure Separators in Branch-and-Cut. In: *Advances in Neural Information Processing Systems*. vol. 37 (2023)
- [22] Mehamdi, A.B., Lacroix, M., Martin, S.: Pricing filtering in Dantzig-Wolfe decomposition. *Operations Research Letters* **58**, 107207 (2025)
- [23] Pauphilet, J., Wu, Y.: The surprising performance of random partial Benders decomposition. *Optimization Online* (2025), 32181
- [24] Rahmaniani, R., Crainic, T.G., Gendreau, M., Rei, W.: The Benders decomposition algorithm: A literature review. *European Journal of Operational Research* **259**(3), 801–817 (2017)
- [25] Scavuzzo, L., Aardal, K., Lodi, A., Yorke-Smith, N.: Machine learning augmented branch and bound for mixed integer linear programming. *Mathematical Programming* (2024)
- [26] Tang, Y., Agrawal, S., Faenza, Y.: Reinforcement learning for integer programming: Learning to cut. In: *International Conference on Machine Learning*. pp. 9367–9376. PMLR (2020)

- [27] Turner, M., Koch, T., Serrano, F., Winkler, M.: Adaptive cut selection in mixed-integer linear programming. *Open Journal of Mathematical Optimization* **4**, 1–25 (2023)
- [28] Wang, Z., Li, X., Wang, J., Kuang, Y., Yuan, M., Zeng, J., Zhang, Y., Wu, F.: Learning cut selection for mixed-integer linear programming via hierarchical sequence model. *arXiv preprint arXiv:2302.00244* (2023)
- [29] Zhang, W., Wang, K., Jacquillat, A., Wang, S.: Optimized scenario reduction: Solving large-scale stochastic programs with quality guarantees. *INFORMS Journal on Computing* **35**(4), 886–908 (2023)
- [30] Zhang, X., Chen, L., Yang, Z., Zeng, Z.: Learning to select cutting planes in mixed-integer linear programming solving. *Expert Systems with Applications* **255**, 124654 (2025)

Appendix

A Subproblem Statistics

Table 5: Time and count breakdown for ML-train method across instances that completed within the time limit ($n = 86$). Unrequired subproblems are those that do not yield cuts in iterations where other subproblems do. Cut Generation shows subproblems that yielded cuts. Feasibility Proving shows subproblems in cut-free iterations (required to verify feasibility). Other includes ML scoring and callback overhead.

Category	Time			Subproblems	
	Time (s)	% Total	% SP	Count	% SP
Master Problem	87,626	25.8%	—	—	—
SP: Cut Generation	56,712	16.7%	23.5%	20,485	5.6%
SP: Feasibility Proving	68,367	20.2%	28.3%	155,251	42.4%
SP: Unrequired	116,594	34.4%	48.2%	190,788	52.1%
Other (ML, overhead)	9,786	2.9%	—	—	—
Total	339,086	100.0%	—	366,524	100.0%

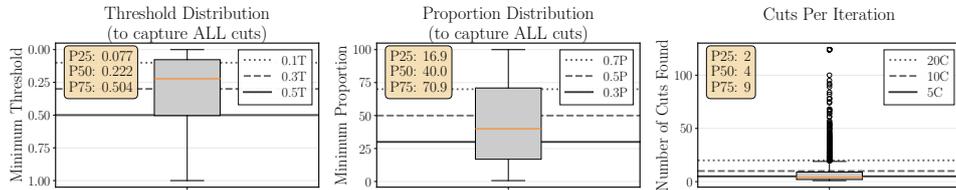


Figure 5: Distribution of minimum parameter values required to capture all cuts per iteration during ML training, with reference lines showing tested configuration values.

B Regression Training Performance Analysis

Table 6: Aggregated ML Model Classification Metrics (%) by Configuration. Values show geometric mean across all instances. Standard includes full enumeration, enabling slightly better performance due to more training data.

Configuration	Accuracy	Precision	Recall	F1-Score
Standard	86.0	34.3	71.0	46.1
ML-0.5P-20S-5R	83.8	32.8	70.4	44.6
ML-0.5T-5C-20S-5R	83.9	32.6	69.8	44.3
ML-0.7P-20S-5R	84.9	34.0	71.5	46.0
ML-0.1T-0.6P-10C-20S-5R	83.5	33.2	71.7	45.2
Average	84.4	33.4	70.9	45.2

C Regression Model Feature Importance

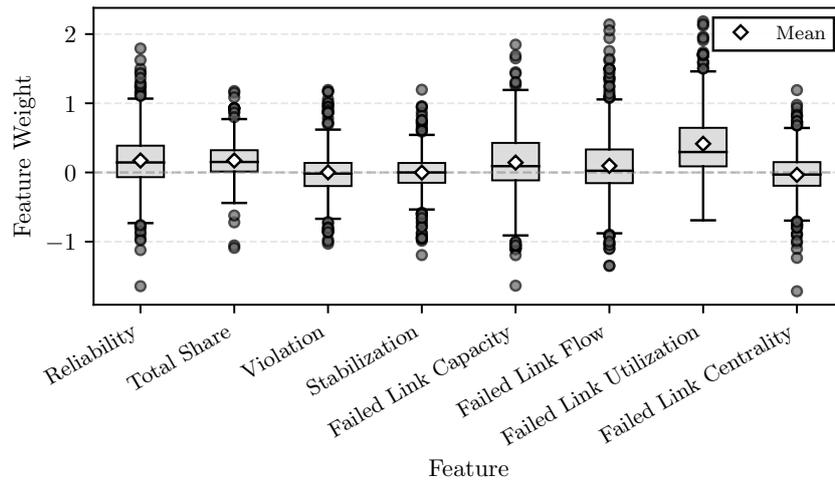


Figure 6: Distribution of learned feature weights across instances. Feature reference in Table 1. Box plots show median (black line), quartiles (box), and mean (white diamond).

D Subproblem Hit Quality Analysis

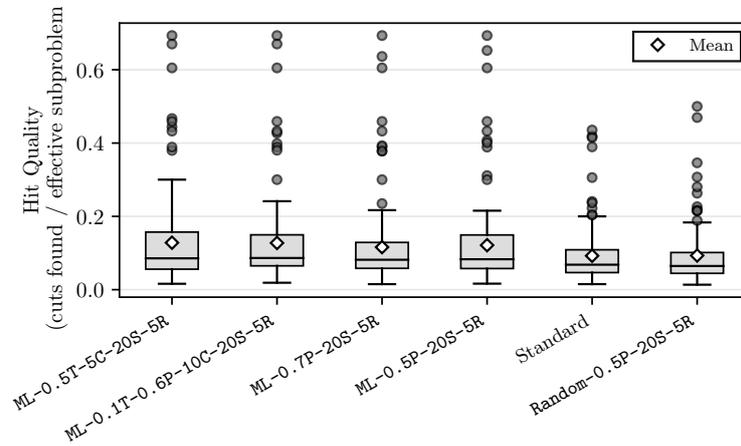


Figure 7: Distribution of hit quality (cuts found / subproblems examined) across methods. Higher values indicate more efficient subproblem selection. Box plots show median (black line), quartiles (box), and mean (white diamond). Outliers are shown as circles.

E Cut Yield Evolution

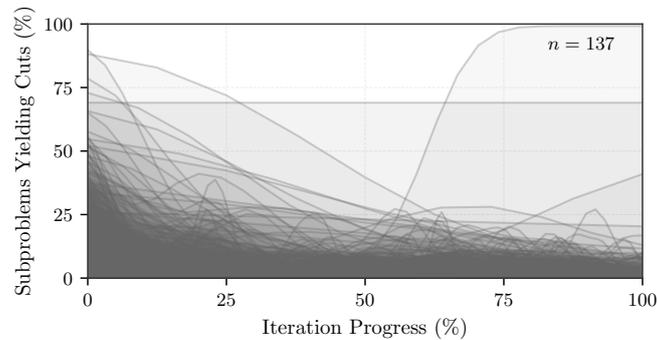


Figure 8: Cut yield evolution across iterations for full enumeration, showing declining patterns from 10–40% early on to 0–10% near convergence.

F Subproblem Performance Analysis

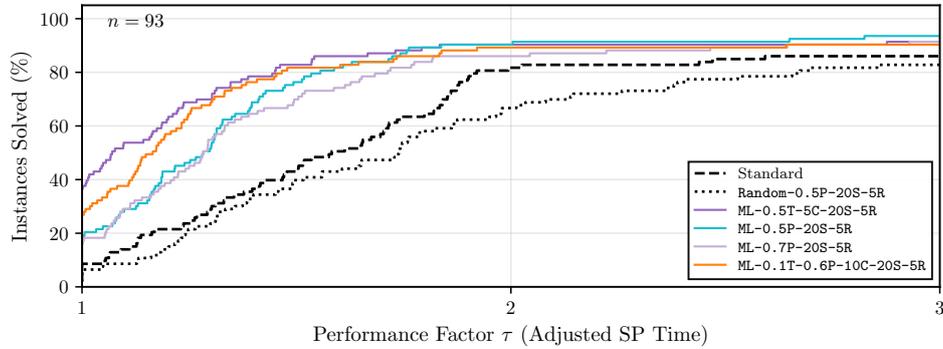


Figure 9: Performance profile based on adjusted subproblem time, which deducts the necessary subproblem time (Oracle baseline that solves all subproblems) from each method's total subproblem time. Shows computational effort spent on unrequired subproblems.

G Detailed Solve Times

Table 7: Solve times (seconds) for all instances. Timeouts: optimality gaps in brackets. Bold: methods faster than Standard (if Standard solved), or methods that solved (if Standard did not). Daggers (†): all methods optimal. Sum/Mean/Sh. Geom. Mean†: over instances solved by all. Wilcoxon test (one-sided): solved use time; unsolved use time-limit+gap%. Significance: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Instance	Baseline				ML-Based Methods			
	Standard	Sum	Mean	Sh. Geom. Mean	0-100	0-200	0-300	0-400
ab1time	6.9	8.5	8.1	5.2	5.0	5.0	5.0	4.8
act1time	36.2	89.2	55.6	68.3	53.7	53.7	53.7	53.3
di-frac	735.7	1743.3	763.6	951.3	519.1	519.1	519.1	1051.1
frac	245.7	200.1	175.9	210.8	162.8	162.8	162.8	183.6
geom†	614.7	1361.4	537.5	428.9	748.2	748.2	748.2	376.7
mean†	811.4	1151.5	1103.0	953.9	822.0	822.0	822.0	1067.7
sub1-geom†	811.4	2157.9	1181.7	1634.4	1797.9	1797.9	1797.9	1474.2
sub1	811.4	36.8	36.4	38.4	31.0	31.0	31.0	28.9
sub1-geom†	811.4	175.8	214.4	154.1	208.1	208.1	208.1	163.9
polys	6.5	8.9	7.8	7.3	7.1	7.1	7.1	6.8
sum	1419.0	2031.0	1612.4	1379.3	2021.1	2021.1	2021.1	1644.3
instance_0001_a3_026_solve01	7.4	7.7	8.3	7.3	6.7	6.7	6.7	6.1
instance_0001_a3_078_solve01	374.4	357.4	428.5	2417.3	277.1	277.1	277.1	376.2
instance_0004_a3_031_solve04	36.7	28.0	23.3	31.1	26.5	26.5	26.5	40.2
instance_0005_a3_081_solve05	1519.3	1665.5	1665.5	1665.5	1665.5	1665.5	1665.5	1511.1
instance_0007_a3_043_solve07	1396.8	1451.6	1075.8	1075.9	1201.4	1201.4	1201.4	1058.0
instance_0008_a3_016_solve08	(34.9%)	(34.9%)	(34.9%)	(34.9%)	(34.9%)	(34.9%)	(34.9%)	(34.9%)
instance_0009_a3_035_solve09	12.9	11.7	11.1	11.2	11.7	11.7	11.7	11.9
instance_0010_a3_041_solve10	201.1	295.2	296.3	296.3	296.3	296.3	296.3	284.8
instance_0011_a3_054_solve11	481.4	738.0	580.0	495.3	582.8	582.8	582.8	637.0
instance_0012_a3_078_solve12	381.3	471.7	471.7	471.7	471.7	471.7	471.7	377.8
instance_0013_a3_087_solve13	477.8	(0.0%)	(0.0%)	3621.0	4409.6	4409.6	4409.6	523.2
instance_0014_a3_108_solve14	(14.1%)	(14.1%)	(14.1%)	(14.1%)	(14.1%)	(14.1%)	(14.1%)	(14.1%)
instance_0015_a3_051_solve15	569.1	703.3	844.7	877.9	684.3	684.3	684.3	963.0
instance_0017_a3_028_solve17	727.0	5546.6	5692.5	7248.8	5209.6	5209.6	5209.6	6853.1
instance_0018_a3_081_solve18	(0.1%)	(0.1%)	(0.1%)	(0.1%)	(0.1%)	(0.1%)	(0.1%)	(0.1%)
instance_0019_a3_122_solve19	(36.7%)	(36.7%)	(36.7%)	(36.7%)	(36.7%)	(36.7%)	(36.7%)	(36.7%)
instance_0020_a3_010_solve20	280.1	264.0	137.5	206.6	237.6	237.6	237.6	258.9
instance_0021_a3_071_solve21	211.4	247.7	208.0	242.1	231.0	231.0	231.0	230.1
instance_0022_a3_021_solve22	1379.4	1060.4	879.1	884.0	1124.7	1124.7	1124.7	977.2
instance_0023_a3_037_solve23	(0.3%)	(0.3%)	(0.3%)	(0.3%)	(0.3%)	(0.3%)	(0.3%)	(0.3%)
instance_0024_a3_070_solve24	1379.4	2066.3	1208.5	2232.2	1605.6	1605.6	1605.6	1911.1
instance_0025_a3_111_solve25	(13.4%)	(14.1%)	(8.0%)	(10.0%)	(17.1%)	(17.1%)	(17.1%)	(3.2%)
instance_0026_a3_122_solve26	(22.4%)	(22.4%)	(22.4%)	(22.4%)	(22.4%)	(22.4%)	(22.4%)	(22.4%)
instance_0028_a3_159_solve28	(91.0%)	(24.9%)	(81.0%)	(21.4%)	(19.7%)	(19.7%)	(19.7%)	(91.0%)
instance_0029_a3_077_solve29	(71.6%)	(6.5%)	(68.7%)	(64.0%)	(18.7%)	(18.7%)	(18.7%)	(72.6%)
instance_0031_a3_087_solve31	(1.3%)	10488.0	(0.1%)	(0.1%)	9117.6	9117.6	9117.6	10642.5
instance_0032_a3_088_solve32	(56.1%)	(56.1%)	(56.1%)	(56.1%)	(56.1%)	(56.1%)	(56.1%)	(56.1%)
instance_0033_a3_104_solve33	158.5	721.0	658.0	604.7	561.3	561.3	561.3	632.0
instance_0035_a3_019_solve35	(9.2%)	(12.3%)	(14.0%)	(12.3%)	(9.2%)	(9.2%)	(9.2%)	(9.2%)
instance_0036_a3_033_solve36	2660.7	2961.9	2313.4	2488.8	2270.7	2270.7	2270.7	2632.2
instance_0038_a3_018_solve38	3292.1	(13.2%)	(13.2%)	(13.2%)	(13.2%)	(13.2%)	(13.2%)	(13.2%)
instance_0041_a3_029_solve41	11.7	11.3	11.3	11.1	11.3	11.3	11.3	11.9
instance_0042_a3_036_solve42	3202.1	3044.1	2766.3	2980.1	2968.7	2968.7	2968.7	2237.4
instance_0043_a3_016_solve43	1.9	1.9	1.9	1.9	2.1	2.1	2.1	2.0
instance_0044_a3_036_solve44	(0.2%)	(0.2%)	(0.2%)	(0.2%)	(0.2%)	(0.2%)	(0.2%)	(0.2%)
instance_0045_a3_107_solve45	(6.0%)	(8.4%)	(6.0%)	(7.5%)	(7.0%)	(7.0%)	(7.0%)	(5.0%)
instance_0047_a3_020_solve47	1.1	1.1	1.1	1.1	1.1	1.1	1.1	1.2
instance_0048_a3_042_solve48	(47.9%)	(45.8%)	(44.1%)	(42.6%)	(47.5%)	(47.5%)	(47.5%)	(10.3%)
instance_0049_a3_041_solve49	80.8	82.8	82.8	82.8	82.8	82.8	82.8	80.8
instance_0050_a3_011_solve50	(97.9%)	(97.9%)	(97.9%)	(97.9%)	(97.9%)	(97.9%)	(97.9%)	(97.9%)
instance_0051_a3_087_solve51	3088.2	2811.0	2611.0	2611.0	2611.0	2611.0	2611.0	487.0
instance_0052_a3_081_solve52	3088.2	2421.7	2266.0	2768.8	1937.4	1937.4	1937.4	2106.3
instance_0053_a3_087_solve53	384.0	371.4	353.7	384.4	377.5	377.5	377.5	355.2
instance_0054_a3_019_solve54	(21.1%)	(15.4%)	(15.4%)	(15.4%)	(21.1%)	(21.1%)	(21.1%)	(14.2%)
instance_0055_a3_037_solve55	3145.7	3071.1	2188.9	4113.0	3704.7	3704.7	3704.7	2976.5
instance_0056_a3_019_solve56	(98.1%)	(98.1%)	(98.1%)	(98.1%)	(98.1%)	(98.1%)	(98.1%)	(98.1%)
instance_0057_a3_108_solve57	(99.1%)	(99.1%)	(99.1%)	(99.1%)	(99.1%)	(99.1%)	(99.1%)	(99.1%)
instance_0058_a3_018_solve58	969.5	911.8	798.4	1016.9	908.4	908.4	908.4	7901.4
instance_0059_a3_043_solve59	77.7	69.7	65.9	64.5	58.3	58.3	58.3	29.8
instance_0060_a3_080_solve60	641.2	641.2	672.2	672.2	672.2	672.2	672.2	641.2
instance_0061_a3_101_solve61	(0.3%)	(0.3%)	9609.2	8925.3	7744.2	7744.2	7744.2	(0.1%)
instance_0062_a3_032_solve62	(0.6%)	(0.6%)	(0.6%)	(0.6%)	(0.6%)	(0.6%)	(0.6%)	(0.6%)
instance_0064_a3_129_solve64	(22.1%)	(100.0%)	(100.0%)	(100.0%)	(22.1%)	(22.1%)	(22.1%)	(100.0%)
instance_0065_a3_032_solve65	(8.7%)	(8.7%)	(8.7%)	(8.7%)	(8.7%)	(8.7%)	(8.7%)	(74.4%)
instance_0071_a3_013_solve71	(1.7%)	(2.2%)	(3.8%)	(1.5%)	(0.8%)	(0.8%)	(0.8%)	(2.6%)
instance_0072_a3_013_solve72	(1.9%)	(1.9%)	(1.9%)	(1.9%)	(1.9%)	(1.9%)	(1.9%)	(1.9%)
instance_0073_a3_012_solve73	(0.2%)	(5.4%)	(0.1%)	(0.1%)	(0.1%)	(0.1%)	(0.1%)	(0.1%)
instance_0074_a3_022_solve74	(9.4%)	(9.4%)	(9.4%)	(9.4%)	(9.4%)	(9.4%)	(9.4%)	(9.4%)
instance_0081_a3_016_solve81	(71.3%)	(72.5%)	(73.0%)	(69.0%)	(67.5%)	(67.5%)	(67.5%)	(72.5%)
instance_0082_a3_076_solve82	(7.6%)	(7.6%)	(7.6%)	(7.6%)	(7.6%)	(7.6%)	(7.6%)	(7.6%)
instance_0083_a3_088_solve83	4829.9	(0.1%)	(0.1%)	5184.6	(0.1%)	(0.1%)	(0.1%)	8922.0
instance_0084_a3_088_solve84	3901.9	671.4	729.1	611.1	611.1	611.1	611.1	551.2
instance_0085_a3_046_solve85	1438.4	2076.9	1710.8	1490.3	1511.3	1511.3	1511.3	1414.4
instance_0087_a3_081_solve87	(41.7%)	(41.7%)	(20.0%)	(36.9%)	(39.2%)	(39.2%)	(39.2%)	(24.7%)
instance_0088_a3_045_solve88	302.9	371.1	311.0	300.7	315.9	315.9	315.9	198.9
instance_0089_a3_081_solve89	901.9	966.3	1060.0	907.0	907.0	907.0	907.0	806.4
instance_0092_a3_026_solve92	3.0	3.6	3.7	3.7	4.2	4.2	4.2	4.2
instance_0093_a3_030_solve93	4801.0	4714.0	2413.8	3218.0	3117.1	3117.1	3117.1	1661.1
instance_0094_a3_018_solve94	561.7	438.0	408.8	462.8	291.3	291.3	291.3	361.4
instance_0095_a3_015_solve95	276.4	276.4	276.4	276.4	276.4	276.4	276.4	276.4
instance_0096_a3_048_solve96	2627.4	(2.0%)	2627.4	6206.4	9409.0	9409.0	9409.0	9409.0
instance_0099_a3_102_solve99	(34.6%)	(31.0%)	(30.4%)	(30.4%)	(36.1%)	(36.1%)	(36.1%)	(31.3%)
instance_0100_a3_032_solve100	(26.5%)	(31.0%)	(31.0%)	(31.0%)	(26.5%)	(26.5%)	(26.5%)	(26.5%)
instance_0112_a3_078_solve112	63.4	730.7	681.4	627.8	600.1	600.1	600.1	431.3
instance_0113_a3_012_solve113	(1.7%)	(1.7%)	(1.7%)	(1.7%)	(1.7%)	(1.7%)	(1.7%)	(1.4%)
instance_0115_a3_109_solve115	3807.6	1032.3	8005.6	8528.3	9501.1	9501.1	9501.1	(0.1%)
instance_0121_a3_027_solve121	91.1	43.6	60.1	32.8	58.2	58.2	58.2	69.9
instance_0122_a3_018_solve122	181.2	164.0	143.3	152.7	151.7	151.7	151.7	141.2
instance_0123_a3_010_solve123	1022.8	1292.3	1022.8	1022.8	1022.8	1022.8	1022.8	941.4
instance_0124_a3_027_solve124	19.0	32.9	21.1	20.0	20.3	20.3	20.3	39.2
instance_0125_a3_010_solve125	1112.2	1228.1	777.3	682.7	742.3	742.3	742.3	792.8
instance_0127_a3_034_solve127	35.3	37.8	28.6	11.8	32.2	32.2	32.2	30.3
instance_0128_a3_027_solve128	198.1	144.8	151.2	151.1	162.0	162.0	162.0	151.0
instance_0129_a3_045_solve129	786.4	388.7	630.1	795.2	931.5	931.5	931.5	783.7
instance_0130_a3_045_solve130	(11.3%)	(8.1%)	(8.1%)	(8.1%)	(11.3%)	(11.3%)	(11.3%)	(8.0%)
instance_0131_a3_046_solve131	(0.1%)	(0.1%)	(0.1%)	(0.1%)	(0.1%)	(0.1%)	(0.1%)	(0.1%)
instance_0132_a3_016_solve132	857.0	393.3	444.8	239.2	231.9	231.9	231.9	129.2
instance_0133_a3_076_solve133	35.4	75.4	42.6	43.3	41.7	41.7	41.7	50.7
instance_0134_a3_085_solve134	291.2	396.4	396.4	396.4	396.4	396.4	396.4	291.8
instance_0139_a3_015_solve139	113.7	257.9	312.1	398.1	297.2	297.2	297.2	233.6
instance_0140_a3_045_solve140	1140.0	1140.0	1140.0	1140.0	1140.0	1140.0	1140.0	768.9
instance_0137_a3_027_solve137	6381.7	6014.9	5440.0	6014.0	7491.7	7491.7	7491.7	6213.3
instance_0139_a3_144_solve139	(34.4%)	(17.1%)	(81.1%)	(61.0%)	(61.0%)	(61.0%)	(61.0%)	(30.2%)
instance_0142_a3_011_solve142	7.9	7.2	7.2	8.4	7.9	7.9	7.9	7.9
instance_0143_a3_011_solve143	151.1	206.6	1697.2	1216.6	1277.3			

H Detailed Convergence Quality

Table 8: Primal-dual integral values (billions, \int_0^T (primal(t) - dual(t)) dt , cost-seconds). Lower is better. Bold: better than Standard. Wilcoxon test (one-sided). Significance: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Instance	Baseline			ML-Based Methods			
	Standard	0, 50, 100, 200, 300	0, 50, 100, 200, 300	0, 50, 100, 200, 300	0, 50, 100, 200, 300	0, 50, 100, 200, 300	0, 50, 100, 200, 300
airlines	50.0	40.0	50.0	50.0	50.0	50.0	40.0
airtrans	257	165	117	201	377	317	317
airtrans	101	61.1	101	101	61.1	101	20.1
france	1020	1260	1050	1070	1120	1120	1060
great	2320	520	335	335	290	290	230
newyork	61.2	01.6	61.7	61.4	61.3	61.3	71.6
schal-garmay	1115	1251	1191	1191	1161	1161	1171
job	40.0	40.0	40.0	40.0	40.0	40.0	40.0
job	10.1	20.1	10.1	10.1	10.1	10.1	10.1
potkins	10.0	20.0	20.0	20.0	20.0	20.0	20.0
potkins	110	150	200	200	220	220	220
instance_0001_a3_276_eeed001	20.0	20.0	20.0	20.0	20.0	20.0	20.0
instance_0001_a3_276_eeed002	1000	1200	670	690	690	690	670
instance_0004_a3_233_eeed004	90.0	50.0	60.0	40.0	40.0	40.0	40.0
instance_0005_a3_211_eeed005	300	300	270	290	290	290	270
instance_0007_a3_243_eeed007	2000	5200	1900	2270	1770	1770	2000
instance_0009_a3_195_eeed009	15100	13100	11500	11500	12000	12000	12000
instance_0009_a3_235_eeed009	30.0	30.0	30.0	30.0	30.0	30.0	30.0
instance_0010_a3_275_eeed010	1.00	1.00	1.00	1.00	1.00	1.00	1.00
instance_0011_a3_284_eeed011	30.0	40.0	40.0	40.0	40.0	40.0	40.0
instance_0012_a3_278_eeed012	1020	1020	1020	1020	1020	1020	1020
instance_0013_a3_287_eeed013	60.0	60.0	60.0	60.0	60.0	60.0	60.0
instance_0014_a3_210_eeed014	2070	2070	4121	4121	4121	4121	2070
instance_0015_a3_251_eeed015	6.0	6.0	6.0	6.0	6.0	6.0	6.0
instance_0017_a3_288_eeed017	1200	1010	7200	6900	7300	7300	6900
instance_0018_a3_261_eeed018	200	200	200	200	200	200	200
instance_0020_a3_101_eeed020	150	150	150	150	150	150	150
instance_0021_a3_271_eeed021	1050	1410	1100	1100	1070	1070	1050
instance_0022_a3_210_eeed022	40.0	51.0	2400	3800	3010	3200	3700
instance_0023_a3_287_eeed023	2200	2181	1500	1500	1500	1500	1500
instance_0025_a3_251_eeed025	3020	4700	4900	5270	4230	4230	5200
instance_0026_a3_211_eeed026	18122	14002	11002	11002	12012	12012	12001
instance_0027_a3_212_eeed027	7000	10300	10000	9000	9000	9000	9000
instance_0028_a3_105_eeed028	11170	10300	12100	11170	11100	11100	10700
instance_0029_a3_177_eeed029	10700	10700	10430	10470	10400	10400	10400
instance_0030_a3_287_eeed030	11372	6922	6922	6102	6102	6102	7712
instance_0032_a3_208_eeed032	526	486	527	527	526	526	527
instance_0033_a3_214_eeed033	80.0	81.9	650	829	649	649	800
instance_0035_a3_219_eeed035	2400	2500	2395	2415	2335	2335	2301
instance_0036_a3_283_eeed036	4300	3800	2900	2900	2900	2900	2900
instance_0042_a3_218_eeed042	2255	3205	1011	1204	1275	1275	1274
instance_0043_a3_230_eeed043	20.0	10.0	20.0	20.0	20.0	20.0	20.0
instance_0042_a3_217_eeed042	2100	20.0	1820	1700	2200	2200	1810
instance_0043_a3_238_eeed043	10100	10100	10100	10100	10100	10100	10100
instance_0044_a3_238_eeed044	10.0	20.0	270	270	200	200	200
instance_0045_a3_107_eeed045	10120	10120	9500	9500	10100	10100	10100
instance_0047_a3_220_eeed047	10.0	10.0	10.0	10.0	10.0	10.0	10.0
instance_0048_a3_242_eeed048	90.0	90.0	90.0	90.0	90.0	90.0	90.0
instance_0049_a3_213_eeed049	110	80.0	70.0	70.0	70.0	70.0	70.0
instance_0051_a3_211_eeed051	10700	10700	10700	10700	10700	10700	10700
instance_0051_a3_211_eeed051	1411	1381	1711	1141	1111	1111	1111
instance_0052_a3_282_eeed052	700	600	600	670	670	670	670
instance_0053_a3_205_eeed053	30.0	30.0	30.0	30.0	30.0	30.0	30.0
instance_0054_a3_124_eeed054	10700	10700	10000	9000	9000	9000	7000
instance_0054_a3_287_eeed054	10000	20000	18000	13700	11720	11720	11310
instance_0056_a3_124_eeed056	11020	11200	10670	11270	10820	10700	10700
instance_0056_a3_278_eeed056	11070	11070	11070	11070	11070	11070	11070
instance_0058_a3_205_eeed058	32.3	31.9	31.9	32.1	32.5	32.5	32.4
instance_0059_a3_282_eeed059	300	350	420	420	420	420	320
instance_0059_a3_210_eeed059	4100	27020	17400	32378	22770	22770	18420
instance_0059_a3_213_eeed059	10700	10700	10700	10700	10700	10700	10700
instance_0060_a3_219_eeed060	10370	11018	11008	11538	94200	11018	11018
instance_0071_a3_211_eeed071	10300	47100	10370	10370	20020	20020	40200
instance_0073_a3_210_eeed073	10370	51821	40822	40812	40812	40812	30747
instance_0073_a3_210_eeed073	10370	50166	50166	50166	30066	30066	30747
instance_0075_a3_212_eeed075	10370	21011	13110	12832	13272	13272	12702
instance_0081_a3_101_eeed081	10370	42100	40100	40100	40100	40100	30400
instance_0082_a3_276_eeed082	22100	10207	31170	30100	35100	35100	30700
instance_0083_a3_230_eeed083	50.0	40.0	50.0	50.0	40.0	40.0	40.0
instance_0083_a3_238_eeed083	150	150	150	150	150	150	150
instance_0085_a3_205_eeed085	4000	11700	4700	4700	6100	6100	6100
instance_0087_a3_208_eeed087	1301	12043	70107	62128	67328	67328	72047
instance_0088_a3_285_eeed088	300	200	200	200	200	200	200
instance_0091_a3_282_eeed091	190	200	210	210	210	210	210
instance_0092_a3_285_eeed092	20.0	10.0	10.0	10.0	10.0	10.0	10.0
instance_0093_a3_290_eeed093	7300	6119	3040	4000	6700	6700	6020
instance_0094_a3_285_eeed094	200	200	200	200	200	200	200
instance_0096_a3_215_eeed096	40.0	40.0	50.0	50.0	40.0	40.0	50.0
instance_0096_a3_284_eeed096	40.0	20.1	41.1	41.1	40.1	40.1	40.1
instance_0099_a3_102_eeed099	11411	114103	67740	69756	71804	71804	67047
instance_0100_a3_102_eeed100	4710	4710	4710	4710	4710	4710	4710
instance_0110_a3_110_eeed110	2000	3001	16443	17008	17725	17725	17002
instance_0110_a3_110_eeed110	2007	1706	1802	1706	1706	1706	1706
instance_0121_a3_227_eeed121	80.0	70.0	80.0	80.0	80.0	80.0	80.0
instance_0122_a3_204_eeed122	400	620	500	500	500	500	500
instance_0122_a3_205_eeed122	500	1150	700	900	400	400	470
instance_0122_a3_227_eeed122	100	100	100	100	100	100	100
instance_0125_a3_100_eeed125	40.0	77.0	3000	2700	2300	2300	2400
instance_0125_a3_100_eeed125	100.0	90.0	90.0	101	120	120	110
instance_0126_a3_227_eeed126	20.0	20.0	20.0	20.0	20.0	20.0	20.0
instance_0126_a3_227_eeed126	2020	1200	1500	1420	1500	1500	1400
instance_0129_a3_205_eeed129	7000	11020	49200	72100	91000	91000	52010
instance_0130_a3_210_eeed130	2000	1800	2100	2100	2200	2200	2100
instance_0132_a3_216_eeed132	100	100	110	110	110	110	110
instance_0133_a3_278_eeed133	100	600	300	300	300	300	300
instance_0134_a3_282_eeed134	80.0	70.0	90.0	90.0	90.0	90.0	90.0
instance_0135_a3_285_eeed135	200	200	200	200	200	200	200
instance_0136_a3_283_eeed136	400	454	453	453	464	464	463
instance_0137_a3_287_eeed137	6021	10506	7500	8752	7000	7000	7000
instance_0139_a3_214_eeed139	10300	20000	20700	31170	31120	31120	30500
instance_0140_a3_276_eeed140	20.0	20.0	20.0	20.0	20.0	20.0	20.0
instance_0142_a3_211_eeed142	1100	1400	900	970	900	900	900
instance_0146_a3_283_eeed146	110	100	80.0	80.0	80.0	80.0	80.0
instance_0151_a3_278_eeed151	3400	3000	3030	3030	3000	3000	3000
instance_0152_a3_215_eeed152	50000	40000	30000	30000	25000	25000	20000
instance_0160_a3_210_eeed160	21800	22750	14100	14111	14401	14401	13811
instance_0162_a3_238_eeed162	30.1	30.1	30.1	30.1	30.1	30.1	30.1
instance_0164_a3_271_eeed164	330	270	210	200	230	230	200
instance_0166_a3_285_eeed166	60.2	71.4	77	77	70.7	70.7	70.1
instance_0167_a3_250_eeed167	200	200	180	180	180	180	180
instance_0168_a3_285_eeed168	32100	12472	7330	9234	12273	12273	8255
instance_0169_a3_287_eeed169	100	100	100	100	100	100	100
instance_0171_a3_274_eeed171	1800	4510	3300	3372	14071	14071	3300
instance_0173_a3_230_eeed173	20.0	20.0	20.0	20.0	20.0	20.0	20.0
instance_0174_a3_294_eeed174	9500	11100	10000	8031	8000	8000	7070
instance_0175_a3_270_eeed175	200	200	200	200	200	200	200
instance_0177_a3_110_eeed177	6000	5201	5200	3000	3821	3821	3871
instance_0178_a3_110_eeed178	4902	7300	6200	6214	6204	6204	6204
instance_0180_a3_242_eeed180	150	140	150	150	150	150	150
instance_0181_a3_277_eeed181	6200	6000	3200	3200	4800	4800	4900
instance_0182_a3_287_eeed182	1271	12227	7152	8243	8043	8043	8163
instance_0183_a3_213_eeed183	12320	12320	12320	12320	12320	12320	12320
instance_0184_a3_212_eeed184	10872	40052	24433	30034	31015	31015	29114
instance_0185_a3_294_eeed185	300	200	200	200	200	200	200
instance_0189_a3_210_eeed189	37212	1114	63757	63077	63707	63707	63037
instance_0190_a3_210_eeed190	1207	1114	900	1000	1014	1014	1004
Mean	180000	170000	160500	160300	160300	160300	160000
Mean	1800	1513	1240	1240	1201	1205	1173
St. Dev. Mean	1184.4	1271.4	1106.4	1106.4	1221.0	1221.7	1151.7
Wilcoxon p		0.002	0.004	0.000	0.001	0.001	0.001
Significance				**	**	**	**