THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Efficient Storage of Pareto Points in Biobjective Mixed Integer Programming

Citation for published version:

Adelgren, N, Belotti, P & Gupte, A 2018, 'Efficient Storage of Pareto Points in Biobjective Mixed Integer Programming', *INFORMS Journal on Computing*, vol. 30, no. 2, pp. 324-338.
<https://doi.org/10.1287/ijoc.2017.0783>

Digital Object Identifier (DOI):

[10.1287/ijoc.2017.0783](https://doi.org/10.1287/ijoc.2017.0783)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

INFORMS Journal on Computing

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

INFORMS Journal on Computing

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Efficient Storage of Pareto Points in Biobjective Mixed Integer Programming

Nathan Adeltgren, Pietro Belotti, Akshay Gupte

To cite this article:

Nathan Adeltgren, Pietro Belotti, Akshay Gupte (2018) Efficient Storage of Pareto Points in Biobjective Mixed Integer Programming. INFORMS Journal on Computing 30(2):324-338. <https://doi.org/10.1287/ijoc.2017.0783>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2018, INFORMS

Please scroll down for article—it is on subsequent pages

With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Efficient Storage of Pareto Points in Biobjective Mixed Integer Programming

Nathan Adeligren,^a Pietro Belotti,^b Akshay Gupte^c

^a Department of Mathematics and Computer Science, Edinboro University of Pennsylvania, Edinboro, Pennsylvania 16444;

^b FICO, Birmingham B37 7GN, United Kingdom; ^c Department of Mathematical Sciences, Clemson University, Clemson, South Carolina 29634

Contact: nadelgren@edinboro.edu, <http://orcid.org/0000-0003-3836-9324> (NA); pietrobelotti@fico.com, <http://orcid.org/0000-0001-6591-6886> (PB); agupte@clemson.edu, <http://orcid.org/0000-0002-7839-165X> (AG)

Received: September 30, 2014

Revised: April 15, 2016; May 1, 2017; August 15, 2017

Accepted: September 3, 2017

Published Online: April 30, 2018

<https://doi.org/10.1287/ijoc.2017.0783>

Copyright: © 2018 INFORMS

Abstract. In biobjective mixed integer linear programs (BOMILPs), two linear objectives are minimized over a polyhedron while restricting some of the variables to be integer. Since many of the techniques for finding or approximating the Pareto set of a BOMILP use and update a subset of nondominated solutions, it is highly desirable to efficiently store this subset. We present a new data structure, a variant of a binary tree that takes as input points and line segments in $\mathbb{R}^2$ and stores the nondominated subset of this input. When used within an exact solution procedure, such as branch and bound (BB), at termination this structure contains the set of Pareto optimal solutions.

We compare the efficiency of our structure in storing solutions to that of a dynamic list, which updates via pairwise comparison. Then we use our data structure in two biobjective BB techniques available in the literature and solve three classes of instances of BOMILP, one of which is generated by us. The first experiment shows that our data structure handles up to 10^7 points or segments much more efficiently than a dynamic list. The second experiment shows that our data structure handles points and segments much more efficiently than a list when used in a BB.

History: Accepted by David Woodruff, editor-in-chief.

Funding: Partial support by Office of Naval Research [Grant N00014-16-1-2725] is acknowledged.

Keywords: biobjective mixed integer optimization • Pareto set • data structure

1. Introduction

Biobjective mixed integer linear programs (BOMILP) have the following form:

$$\begin{aligned} \underset{x,y}{\text{minimize}} \quad & f(x,y) := \begin{cases} f_1(x,y) := c_1^\top x + d_1^\top y, \\ f_2(x,y) := c_2^\top x + d_2^\top y \end{cases} \\ \text{s.t.} \quad & (x,y) \in P_I := \{(x,y) \in \mathbb{R}^m \times \mathbb{Z}^n : Ax + By \leq b\}, \end{aligned} \quad (1)$$

where P_I is a bounded set. Define as $\Omega := \{\omega \in \mathbb{R}^2 : \omega = f(x,y) \forall (x,y) \in P_I\}$ the collection of all points in $\mathbb{R}^2$ which can be obtained using the objective function values of feasible solutions to (1). We refer to the space $\mathbb{R}^2$ containing Ω as the *objective space*.

For any two vectors $v^1, v^2 \in \mathbb{R}^2$ we use the following notation: $v^1 \leq v^2$ if $v_i^1 \leq v_i^2$ for $i = 1, 2$; $v^1 < v^2$ if $v^1 \leq v^2$ and $v^1 \neq v^2$; and $v^1 < v^2$ if $v_i^1 < v_i^2$ for $i = 1, 2$. Given distinct $(\bar{x}, \bar{y}), (x', y') \in P_I$, we say that $f(\bar{x}, \bar{y})$ *dominates* $f(x', y')$ if $f(\bar{x}, \bar{y}) \leq f(x', y')$. This dominance is *strong* if $f(\bar{x}, \bar{y}) < f(x', y')$; otherwise it is *weak*. A point $(\bar{x}, \bar{y}) \in P_I$ is *(weakly) efficient* if $\nexists (x', y') \in P_I$ such that $f(x', y')$ (strongly) dominates $f(\bar{x}, \bar{y})$. The set of all efficient solutions in P_I is denoted by X_E . A point $\bar{\omega} = f(\bar{x}, \bar{y})$ is called *Pareto optimal* if $(\bar{x}, \bar{y}) \in X_E$. Given

$\Omega' \subseteq \Omega$ we say that $\omega' \in \Omega'$ is *nondominated* in Ω' if $\nexists \omega'' \in \Omega'$ such that ω'' dominates ω' . Note that Pareto optimal points are nondominated in P_I . We consider a BOMILP solved when the set of Pareto optimal points $\Omega_p := \{\omega \in \mathbb{R}^2 : \omega = f(x,y) \forall (x,y) \in X_E\}$ is found.

It is known (Ehrgott 2005) that a biobjective LP (BOLP) can be solved by taking convex combinations of $f_1(\cdot)$ and $f_2(\cdot)$ and solving a finite number of LPs. Thus, for BOLP, the set of Pareto points can be characterized as $\Omega_p = \{(\xi_1, \xi_2) \in \mathbb{R}^2 : \xi_2 = \psi(\xi_1)\}$ where $\psi(\cdot)$ is a continuous, convex, piecewise linear function obtained using extreme points of the feasible region. For biobjective integer programs (BOIPs), it is known that Ω_p is a finite set of points in $\mathbb{R}^2$. Now consider the case of BOMILP. Let $Y = \text{Proj}_y P_I$ be the set of integer feasible subvectors to (1). Since P_I is bounded, we have $Y = \{y^1, \dots, y^k\}$ for some finite k . Then for each $y^i \in Y$, there is an associated BOLP, referred to as a *slice problem* and denoted $\mathbb{P}(y^i)$, obtained by fixing $y = y^i$ in (1):

$$\begin{aligned} \mathbb{P}(y^i) \quad & \min_x \quad \begin{cases} f_1(x) = c_1^\top x + d_1^\top y^i, \\ f_2(x) = c_2^\top x + d_2^\top y^i \end{cases} \\ \text{s.t.} \quad & Ax \leq b - By^i. \end{aligned} \quad (2)$$

Figure 1. (Color online) Example of Solutions Generated When Solving an Instance of BOMILP

Problem $\mathbb{P}(y^i)$ has a set of Pareto solutions $\Omega_i := \{(\xi_1, \xi_2) \in \mathbb{R}^2: \xi_2 = \psi_i(\xi_1)\}$, where $\psi_i(\cdot)$ is a continuous convex piecewise linear function as explained before. Then $\Omega_p \subseteq \bigcup_{i=1}^k \Omega_i$ and this inclusion is strict in general. In particular, we have

$$\Omega_p = \bigcup_{i=1}^k \left(\Omega_i \setminus \bigcup_{j \neq i} (\Omega_j + \mathbb{R}_+^2 \setminus \{0\}) \right). \quad (3)$$

Such union of sets is not, in general, represented by a convex piecewise linear function. Figure 1 shows an example with $k = 5$.

It should be noted that finding Ω_p is not a trivial task in general. In the worst case, $\Omega_p = \bigcup_{i=1}^k \Omega_i$ and one may have to solve every slice problem to termination, which can have exponential complexity. For multiobjective IPs (i.e., $m = 0$), De Loera et al. (2009) prove that Ω_p can be enumerated in polynomial-time for fixed n , which extends the well-known result that single-objective IPs can be solved in polynomial-time for fixed n (Lenstra 1983). We are unaware of any similar results for BOMILP.

Exact procedures for solving BOMILP with general integers have been recently the subject of intense research. Exact methods have been presented by Belotti et al. (2012, 2016), Boland et al. (2014), and more recently by Soylu and Yıldız (2016). Özpeynirci and Köksalan (2010) give an exact method for finding *supported* solutions of BOMILP. Most other techniques in the literature tackle specific cases. Vincent et al. (2013) improved upon the method of Mavrotas and Diakoulaki (2005) for mixed 0-1 problems. Stidsen et al.

(2014) propose a method for solving mixed 0-1 problems in which only one of the objectives contains continuous variables. Belotti et al. (2012), Mavrotas and Diakoulaki (2005), Stidsen et al. (2014), and Vincent et al. (2013) are based on branch-and-bound (BB) procedures in which the Pareto set is determined by solving several BOLPs. Instead, Boland et al. (2014), Özpeynirci and Köksalan (2010), and Soylu and Yıldız (2016) determine the Pareto set by solving several MIPs, albeit in different ways. Although, for example, Boland et al. (2014) recursively partition the objective space to circumscribe subsets of Pareto points and segments, Soylu and Yıldız (2016) find the Pareto frontier incrementally, starting from a solution to a single-objective problem. The pure integer case has been studied for binary variables (Kiziltan and Yucaoglu 1983), general integers (Ralphs et al. 2006) and specific classes of combinatorial problems (Sourd and Spanjaard 2008, Przybylski et al. 2010, Jozefowiec et al. 2012).

We present a data structure for efficiently storing a nondominated set of feasible solutions to a BOMILP. For lack of a better name, we call it *biobjective tree*, or BoT. The BoT can be used in exact and heuristic solution procedures that aim at finding or approximating the Pareto set. Data structures such as *quad-trees* have been used for storing Pareto points in the past (Sun and Steuer 1996, Sun 2006), although only in the pure integer case. Sun and Steuer (1996) stored nondominated solutions using both quad-trees and dynamic lists which were updated via pairwise comparison. They showed that in the pure integer, biobjective case, dynamic lists store nondominated solutions more efficiently than quad-trees.

In Section 2, we describe the BoT; and in Section 3, we describe the insertion function, prove its correctness, and provide an example of use. In Section 4, we present the results of two experiments. The first experiment shows that a BoT is able to store nondominated points and segments more efficiently than a dynamic list, and can insert up to 10^7 solutions in reasonable time. In the second experiment, we utilize a BoT in the BB procedures of Belotti et al. (2012) and Adelgren and Gupte (2016) to solve specific instances of BOMILP. The results show that BoT handles points and segments much more efficiently than a list when used in a BB.

2. Biobjective Tree (BoT)

Figure 1(a) shows an example of solutions that might be generated when solving an instance of BOMILP. We would like to store the nondominated portion of these points and segments, as shown in Figure 1(b). A BoT stores only the nondominated subset of the solutions regardless of the order in which they are inserted. Consider the set of solutions in Figures 1(a) and suppose that the segments connecting (1, 17), (2, 15), (4, 14), and (9, 13) are currently stored. When inserting the point (5, 11), which dominates a portion of the segment connecting (4, 14) and (9, 13), the portion must be removed before the point is added. Similarly, when the segment connecting (6, 16) and (7, 10) is inserted, because a portion of this segment is dominated by (5, 11), only the nondominated portion of the segment should be added.

A BoT is a binary search tree (BST) in which each node represents either a singleton or a line segment associated with a Pareto point or set of Pareto points in the objective space. Denote as Π the set of nodes in the tree. Each $\pi \in \Pi$ is defined as a triplet $\pi = (S, l, r)$, its components defined as follows:

- $S = (x_1, x_2, y_1, y_2)$ represents a segment $[(x_1, y_1), (x_2, y_2)]$ defined in the objective space; the segment collapses to a point if $x_1 = x_2$ and $y_1 = y_2$.
- l and r are the left and the right child node of π , respectively.

We identify each element of a triplet π as $\pi.S$, $\pi.l$, and $\pi.r$. Let us define operators that are used in the remainder: $size(\pi)$ is the size of the subtree rooted at π , defined as $size(\emptyset) = 0$ and $size(\pi) = 1 + size(\pi.l) + size(\pi.r)$; also, $parent(\pi)$ is a node π' such that $\pi'.l = \pi$ or $\pi'.r = \pi$, if any exists, or $\emptyset$ otherwise. Hence, π is a *leaf node* if $\pi.l = \pi.r = \emptyset$, or the *root node* if $parent(\pi) = \emptyset$. Finally, $subtree(\pi)$ is the subtree rooted at π . In the remainder, even though a node π is defined by a triplet, we sometimes use π to refer to $\pi.S$ for ease of notation (especially with set operations), but only where we believe this does not lead to confusion. In particular, for $\pi = (S, l, r)$ the operation $\pi \cap A$, where $A \subseteq \mathbb{R}^2$, returns a node $\pi' = (S', l, r)$ such that $S' = S \cap A$.

A BoT contains nodes that correspond each to a segment $S = [(x_1, y_1), (x_2, y_2)]$ such that $x_1 \leq x_2$ and $y_1 \geq y_2$, as otherwise S can be reduced to a single point. The BoT maintains a nonstrict total order $\leq$ between nodes: for two nodes $\pi' = (S', l', r')$, $\pi'' = (S'', l'', r'') \in \Pi$, with $S' = [(x'_1, y'_1), (x'_2, y'_2)]$ and $S'' = [(x''_1, y''_1), (x''_2, y''_2)]$, $\pi' \leq \pi''$ if either π' and π'' are the same, or $x'_2 \leq x''_1$ and $y'_2 \geq y''_1$. Our notation extends to node segments as well, i.e., $\pi' \leq \pi'' \equiv \pi'.S \leq \pi''.S$. Also, for any $\pi \in \Pi$, we have $\pi.l \leq \pi \leq \pi.r$.

All operations on the BoT must conserve this total order; as for any BST, enumerating its sorted elements amounts to an *in-order* parse of the tree. Removal of a subtree from the BST and rebalancing subtrees preserves the order (Knuth 1998). Most tree operations carry over to the BoT without change, but insertion, discussed in Section 3, is radically different: a BST insertion increases the tree size by one, while in a BoT, inserting a single node might have a large-scale effect. For instance, an entire subtree might be deleted if the inserted segment S dominates all nodes of the subtree. Even if S does not dominate any of the current nodes, the nondominated portion of S could be as many as $t + 1$ disjoint segments (if t nodes are stored).

This is the main point of departure with classical BSTs and other data structures (red-black trees, for instance) that have fast insertion and lookup. This also explains why other data structures such as a list or a hash table might be less suited for this purpose: lookup and insertion in the list are rather expensive at $O(t)$ if t is the number of elements in the list; insertion of a point in a hash table, and the possible deletion of a large number of elements is inefficient in a hash table for many reasons. One such reason is lack of *locality* of the data: a good hash function would guarantee large separation, in the hash table, between two adjacent points or segments in the objective space. Finding all segments that are dominated by one point or segment would prove inefficient, as it might require visiting the entire data structure.

The following notation is useful in the remainder of the paper: we partition $\mathbb{R}^2$ relative to the segment $\pi.S = [(x_1, y_1), (x_2, y_2)]$ of a node π into four regions $R_\alpha(\pi)$ for $\alpha \in \{\text{up}, \text{dn}, \text{left}, \text{right}\}$. We define $R_{\text{up}}(\pi) = \pi.S + \{(x, y) \in \mathbb{R}^2: x > 0, y > 0\}$, i.e., the set of points dominated by π ; $R_{\text{dn}}(\pi) = \pi.S + \{(x, y) \in \mathbb{R}^2: x \leq 0, y \leq 0\}$; $R_{\text{left}}(\pi) = \{(x, y) \in \mathbb{R}^2: x \leq x_1, y > y_1\}$; and finally $R_{\text{right}}(\pi) = \{(x, y) \in \mathbb{R}^2: x > x_2, y \leq y_2\}$ (see Figure 2).

Example 1. Suppose $\pi \in \Pi$ is defined by the segment between (2, 5) and (3, 3). Further suppose that π_1 is the point (1, 5), π_2 is the point (2, 6), and both π_1 and π_2 are inserted into the BoT. Observe Figure 2(c). The point associated with node π_1 weakly dominates the left-most point of the segment associated with π , and

Figure 2. (Color online) Partition of $\mathbb{R}^2$ w.r.t. a Segment or Point $\pi.S$ (a, b); Example of Weak Domination (c)

thus π_1 should be stored. However, the point associated with node π_2 is weakly dominated by the segment associated with π , and so π_2 should not be stored.

The previous example leads us to our next choice. Consider π and π_1 in Figure 2(c): we should mark the segment as open at $(2,5)$, i.e., $\pi.S =](2,5), (3,3)]$, assuming no other points dominate $(3,3)$. However, for simplicity, we do not memorize whether a segment $\pi.S$ is open or closed on either extreme. Instead, upon termination of the BB, for any node π we check the nodes immediately to the left and to the right to determine if $\pi.S$ is (partially) open.

3. Insertion

The INSERT function is described in Algorithm 1. It takes two inputs: a node π^* to be inserted and a node π , which is the root of the tree. Because a node π does not hold any information of all other nodes of the subtree rooted at π , the insertion of π^* might have to be propagated both to $\pi.l$ and to $\pi.r$. For this reason, the INSERT function is recursive. If the recursive call inserts π^* in $\pi.l = \emptyset$ (or $\pi.r = \emptyset$), then $\pi.l$ becomes π^* . The main call to INSERT is done by passing the root node of the tree, denoted as π_0 , as the second argument.

The function REPLACE(π', π'') replaces $\pi' \in \Pi$ with $\pi'' \in \Pi$, leaving the BoT otherwise unchanged. We instead denote with $\pi' \leftarrow \pi''$ the process of replacing π' and its entire subtree with π'' and its entire subtree. REMOVE_NODE(π), described in Algorithm 2, deletes π if it is a leaf node, otherwise it replaces π with a node π' of its subtree that retains the total order. For this, π' must be adjacent to π to the left or right, i.e., it is the left-most node of the right subtree of π or right-most node of the left subtree of π .

Algorithm 1 (Inserting a new point or segment π^* into a BoT at node π)

```

1: function INSERT( $\pi^*, \pi$ )
2:   if  $\pi^* = \emptyset$  then return
3:   if  $\pi = \emptyset$  then REPLACE( $\pi, \pi^*$ )
4:   else
5:     Define  $S' := \pi.S \setminus \text{cl}(R_{\text{up}}(\pi^*))$ 
6:     if  $S' = \emptyset$  then

```

```

7:       REMOVE_NODE( $\pi$ )
8:       INSERT( $\pi^*, \pi$ )
9:   else
10:    if  $\exists S'', S'''$  s.t.  $S' = S'' \cup S''' \wedge \text{cl}(S'') \cap$ 
         $\text{cl}(S''') = \emptyset \wedge S'' \leq S'''$  then
11:      Create new node  $\pi' = (S''', \emptyset, \pi.r)$ 
12:       $\pi.r \leftarrow \pi'$ 
13:       $\pi.S \leftarrow S''$ 
14:    else  $\pi.S \leftarrow S'$ 
15:    INSERT( $\pi^* \cap R_{\text{left}}(\pi), \pi.l$ )
16:    INSERT( $\pi^* \cap R_{\text{right}}(\pi), \pi.r$ ).

```

We skip the trivial cases and focus on line 5 onward in Algorithm 1. The set S' is obtained by removing from $\pi.S$ the set of points dominated by π^* . If $S' = \emptyset$, π is entirely dominated by π^* and can be removed; the procedure is then called recursively on the new subtree. If $S' \neq \emptyset$, then it might be the union of at most two subsegments S'' and S''' . If so, S''' becomes the segment of a new node π' , which is assigned the right subtree of π , and π has its segment restricted to S'' . Otherwise, $\pi.S$ is changed to S' . Insertion is then called recursively on the left and right subtrees of π .

The following property is equivalent to the total order mentioned in Section 2.

Property 1. For any $\pi \in \Pi$, all nodes in the subtree of $\pi.l$ (resp. $\pi.r$) are located completely within $R_{\text{left}}(\pi)$ (resp. $R_{\text{right}}(\pi)$).

Algorithm 2 (Remove a node that has been shown to be dominated)

```

1: function REMOVE_NODE( $\pi$ )
2:   if  $\text{size}(\pi) = 1$  then  $\pi \leftarrow \emptyset$ 
3:   else
4:     if  $\text{size}(\pi.l) > \text{size}(\pi.r)$  then
5:        $\tilde{\pi} \leftarrow \text{FIND\_RIGHTMOST\_NODE}(\pi.l)$ 
6:     else  $\tilde{\pi} \leftarrow \text{FIND\_LEFTMOST\_NODE}(\pi.r)$ 
7:     REPLACE( $\pi, \tilde{\pi}$ )
8:     REMOVE_NODE( $\tilde{\pi}$ ).

```

For a tree with t nodes at the time of an insertion, the worst-case complexity of INSERT is $O(t)$ as the entire tree might have to be visited. For example, consider the tree containing all points $\{(i, M - i), i = 1, 2, \dots, M - 1\}$

for any $M > 2$. Inserting $\pi^* = (0, 0)$, which results in π^* replacing the whole tree, requires visiting all nodes of the tree.

3.1. Correctness

Proposition 1. INSERT removes any portion of a currently stored node π , which is dominated by an inserted node π^* .

Proof. Obvious from steps 7 (if $S' = \emptyset$), 13 (if $S' = S'' \cup S'''$), and 14 (S' is a segment) of Algorithm 1.

Proposition 2. INSERT adds node π^* , or a portion, to the tree if and only if it is not dominated by any node currently stored in the tree.

Proof. (If) Let S' be a portion of $\pi^*.S$ not dominated by any $\pi \in \Pi$; then $(S', \emptyset, \emptyset)$ is inserted as a leaf node after a finite number of recursive calls—see lines 15 and 16.

(Only if) We show the contrapositive. Suppose there is a portion $\tilde{S}$ of $\pi^*.S$ that is dominated by $\pi \in \Pi$. If $\pi = \pi_0$, then $\text{INSERT}(\pi^*, \pi)$ has the effect (line 5) of removing $\tilde{S}$ from the set of points to be inserted, while the two recursive calls at lines 15 and 16 are made for $\pi^* \cap R_{\text{left}}(\pi)$ and $\pi^* \cap R_{\text{right}}(\pi)$, both excluding $\tilde{S}$. If $\pi \neq \pi_0$ and $\pi \in \text{subtree}(\pi_0.l)$ (resp. $\pi \in \text{subtree}(\pi_0.r)$), the recursive call $\text{INSERT}(\pi^* \cap R_{\text{left}}(\pi_0), \pi_0.l)$ (resp. $\text{INSERT}(\pi^* \cap R_{\text{right}}(\pi_0), \pi_0.r)$) guarantees that a call to $\text{INSERT}(\pi^*, \pi)$ will be placed, i.e., $\tilde{S}$ will be eliminated by subsequent recursive calls. $\square$

According to Propositions 1 and 2, the end state of the BoT contains all and only nondominated solutions.

Proposition 3. INSERT retains Property 1.

Proof. The result is trivial for the cases in lines 2 and 3 of Algorithm 1. The same holds if $S' = \emptyset$ (line 6), as it is easy to prove that REMOVEDNODE retains Property 1. If, for the node π^* to be inserted, $\pi^*.S \subset \bigcup_{\pi \in \Pi} R_{\text{up}}(\pi)$, then π^* is dominated and by Proposition 2 it will not be inserted, thereby not modifying the BoT. By the same

proposition, any segment $\pi^*.S$ or portion that neither dominates nor is dominated by any node will be added to the BoT so as to satisfy Property 1: INSERT recursively runs lines 15 and 16 until said portion is added as a leaf.

Assume now that $\pi^*.S \cap R_{\text{dn}}(\pi) \neq \emptyset$ for at least one $\pi \in \Pi$. Since π^* is not dominated, two cases arise: (i) $\pi.S \setminus \text{cl}(R_{\text{up}}(\pi^*))$ is a single segment S' ; (ii) $\pi.S \setminus \text{cl}(R_{\text{up}}(\pi^*))$ is the union of two disjoint segments S'' and S''' with $S'' \leq S'''$.

In case (i), the property holds after running line 14, as $\pi.S$ is replaced by its subset S' . Case (ii) is dealt with on lines 10–13: since $S'' \leq S'''$, $\tilde{\pi}.S \leq S''$ for all $\tilde{\pi} \in \text{subtree}(\pi.l)$ and $S''' \leq \hat{\pi}.S$ for all $\hat{\pi} \in \text{subtree}(\pi.r)$. Thus, replacing $\pi.S$ with S'' ensures that Property 1 is maintained for $\pi.l$, while placing $\pi' := (S''', \emptyset, \pi.r)$ as the right child of π ensures that the same holds for $\pi.r$. By construction $\pi \leq \pi'$, and this concludes the proof. $\square$

3.2. Illustrative Example

We use the points and segments in Figure 1(a) as input to a BoT and show a few of the nontrivial steps of developing it. Assume that these solutions are obtained from five separate slice problems and that the Pareto sets of these slice problems are: (i) the singleton (1, 19), (ii) the piecewise linear curve connecting (1, 17) and (9, 13), (iii) the one connecting (6, 16) and (11, 4), (iv) the singleton (5, 11), and (v) the piecewise linear curve connecting (8, 7) and (17, 2). The points and segments that define these Pareto sets are inserted into the BoT in the order of (iii), (iv), (ii), (v), (i). Piecewise linear curves are inserted as individual line segments from left to right.

The first call is on an empty tree, hence $\pi^* = [(6, 16), (7, 10)]$ becomes the tree. The next point is $\pi^* = [(7, 10), (10, 5)]$; note that $\pi^* \subset R_{\text{right}}(\pi_0)$ and should be inserted at $\pi_0.r$. Since $\pi_0.r = \emptyset$, π^* replaces $\pi_0.r$. The insertion of $[(10, 5), (11, 4)]$ is analogous.

Next consider Pareto set (iv). Let $\pi^* \leftarrow (5, 11)$ and call $\text{INSERT}(\pi^*, \pi_0)$. Observe Figure 3(a). Because π^* partially

Figure 3. (Color online) An Example That Shows the Effect of the Insertion of a Segment on a Set of Nondominated Segments

Figure 4. (Color online) Rebalancing a Tree Upon Insertion of Point (5, 11)

dominates π_0 , we remove the dominated portion of π_0 by letting $S' = \pi_0 \cdot S \setminus \text{cl}(R_{\text{up}}(\pi^*))$. As a consequence, $\pi^* \subset R_{\text{left}}(S')$, and, since $\pi_0 \cdot l = \emptyset$, π^* becomes the left child of π_0 . Figure 4(a) shows the BoT after π^* has been inserted. We leave it to the reader to consider Pareto set (ii). Rebalancing the subtree rooted at $\pi_0 \cdot l$ after processing this set yields the BoT shown in Figure 4(b).

Consider now the insertion of Pareto set (v). Let $\pi^* \leftarrow [(8, 7), (14, 3)]$ and call $\text{INSERT}(\pi^*, \pi_0)$. Clearly $\pi^* \subset R_{\text{right}}(\pi_0)$, and is hence inserted to $\pi_0 \cdot r$. Observe from Figure 3(b) that π^* partially dominates $\pi_0 \cdot r$. This time, though, the portion of $\pi_0 \cdot r$, which is dominated, is the center section of the segment. This means that $\pi_0 \cdot r$ must be split into two nodes π_1 and π_2 . Node π_1 takes the place in the tree where $\pi_0 \cdot r$ originally was, and the left subtree of $\pi_0 \cdot r$ becomes the left subtree of π_1 . Node π_2 becomes the right child of π_1 and the right subtree of $\pi_0 \cdot r$ becomes the right subtree of π_2 . Subsequently, $\pi^* \subset R_{\text{right}}(\pi_1)$ and thus π^* is inserted to π_2 (which is now $\pi_0 \cdot r \cdot r$). Notice that $\pi_0 \cdot r \cdot r$ partially dominates π^* and that it is the center portion of π^* that is dominated. Thus the calls to $\text{INSERT}(\pi^* \cap R_{\text{left}}(\pi_0 \cdot r \cdot r), \pi_0 \cdot r \cdot r \cdot l)$ and $\text{INSERT}(\pi^* \cap R_{\text{right}}(\pi_0 \cdot r \cdot r), \pi_0 \cdot r \cdot r \cdot r)$ each causes a portion of π^* to be inserted at $\pi_0 \cdot r \cdot r \cdot l$ and $\pi_0 \cdot r \cdot r \cdot r$ respectively. Since $\pi_0 \cdot r \cdot r \cdot l = \emptyset$, $\pi^* \cap R_{\text{right}}(\pi_0 \cdot r \cdot r)$ becomes $\pi \cdot r \cdot r \cdot l$. Since $\pi_0 \cdot r \cdot r \cdot r$ is the segment (10, 5) to (11, 4), it is clear that another portion of π^* needs to be removed, and then the remainder of π^* becomes $\pi_0 \cdot r \cdot r \cdot r \cdot r$.

The remaining insertions are analogous to those that we have described. After yet another rebalance, the final BoT is as in Figure 4(c).

4. Computational Experiments

We implemented the BoT in the C programming language and performed two tests. The first test addresses the efficiency with which a large number of randomly generated solutions can be stored in a BoT, using different rebalancing techniques. The second test addresses the utility of a BoT when used in the two BB algorithms for BOMILP by Belotti et al. (2012) and Adelgren

and Gupte (2016). All tests were run on Clemson University's Palmetto Cluster. Specifically, an HP SL250s server node with a single Intel E5-2665 CPU core with 16 GB RAM running Scientific Linux 6.4 was used.

In all of these experiments, we compare the performance of the BoT with that of a dynamic linked list (L). Like a BoT, the linked list takes points and segments in $\mathbb{R}^2$ as input and stores only the nondominated subset of all input. All segments $S = (x_1, x_2, y_1, y_2)$ of the linked list are stored in nondecreasing order of x_1 , so that parsing all elements of the list produces the same output as an *in-order* visit of a BoT. Inserting a point or segment S' consists in comparing it with every stored point or segment in the list, until a segment S is encountered such that $S' \subset R_{\text{left}}(S)$. During each comparison, dominated solutions are discarded. Although only a few, if any, elements of the list might be changed with the insertion of S' (because we do not know where these elements are located and do not have more sophisticated search mechanisms on such a list), insertion has an *average* complexity of $O(t)$. Such lists have been used for storing nondominated solutions in both the pure integer (Sun and Steuer 1996) and mixed-integer cases (Mavrotas and Diakoulaki 2005, Vincent et al. 2013).

Maintaining a balanced tree is one of the most costly operations, as shown in the tests below. As its only purpose is efficiency, it does not need to be applied at every step. Hence, we decided to further consider the rebalancing operations and use an alternative strategy that is less computationally costly, but still keeps the BoT fairly balanced.

We use the strategy of Overmars and Van Leeuwen (1982): for each nonleaf node π , the subtrees of $\pi \cdot l$ and $\pi \cdot r$ must contain no more than $(1/(2 - \delta))\text{size}(\pi)$ nodes, where δ is a preselected value in the open interval (0, 1). This causes the depth of the tree to be, at most, $\log_{2-\delta} t$, where t is the number of nodes in the tree.

Overmars and Van Leeuwen (1982) also suggest rebalancing by traversing the path travelled by an inserted solution in the reverse order and checking whether or not the balance criterion is satisfied at each

of these nodes. This saves one from having to check the balance criterion at every node in the tree since the only places where it could have been altered are at nodes along this path. In our case, though, when a line segment is inserted into a BoT, it often does not remain intact, but may be separated into many smaller segments, each traversing its own path through the tree before finally being added. For this reason, rearranging the tree after insertion is troublesome, and we experimented with the following alternative approaches:

A0. No rebalancing is used.

A1. Before inserting a point or segment at the root node, check the balance criterion at every node in the tree and rebalance where necessary. This approach guarantees that the balance of the tree is maintained, though at a high computational cost.

A2. Check the balance criterion after the k -th insertion (we used $k = 100$ in our tests), then check the balance criterion when the tree size increased by 101% w.r.t. the size at the previous check. This approach significantly decreases the complexity of rebalancing, but eliminates the balance guarantee.

A3. Check the balance criterion at any node where insertion is recursively called. This approach has a much lower complexity than A1, and would cause balance to be maintained at the root node, and along any frequently travelled paths in the tree. However, again the guarantee of balance is lost.

A4. Combine approaches A2 and A3: check the balance criterion of the entire tree after the k -th insertion ($k = 100$), then check the balance criterion again when the tree size increases by 800% w.r.t. the size at the previous check. In between these checks of the entire tree, check the balance criterion for any node where insertion is recursively called.

Approach A4 allows for maintaining a fairly well balanced tree by applying approach A2 much more infrequently than if using approach A2 alone. Clearly this has a higher complexity than approach A3, but it may be less than that of approach A2 and allow for a more balanced tree.

We implemented each of these approaches in our first experiment, described in Section 4.1. We utilize approach A2 when performing our second experiment, which is described in Section 4.2 because for most of our tests, A2 performed comparably to A0 in terms of CPU time, but always maintained a more balanced tree.

4.1. Insertion of Large Number of Random Points

This first test has two main purposes: (i) to compare the efficiency of a BoT with that of a dynamic list when storing nondominated solutions, and (ii) to determine the best rebalancing approach w.r.t. tree depth and time.

The test consists of repeating the following procedure until N insertions have been made into a BoT

or the dynamic list. First, generate a random integer $i \in [1, 6]$ and a random number $r_1 \in (0, 10)$. Then, if $i > 1$, for each $j \in \{2, \dots, i\}$ a random number $c_j \in (0, 1)$ is generated and we define $r_j = r_1 + \sum_{\ell=2}^j c_\ell$. Next, for each $j \in \{1, \dots, i\}$ the following are computed: (i) $y_j = (10.5 - r_j)^2/5 - k$, and (ii) $x_j = r_j + (5 - k)$. Here k is a dynamic value which is defined as 1 at the start of the test and increases by μ/N each time the above process is repeated, and $\mu \in \mathbb{R}$ determines how much the solutions should “improve” over the course of the test. If $i = 1$, the singleton (x_1, y_1) is inserted into the structure; otherwise, the points $(x_1, y_1), \dots, (x_i, y_i)$ are arranged in order of increasing x values and then the line segments connecting each adjacent pair of points are inserted into the structure. All instances used for these tests and the code used to create them are available online at <https://github.com/Nadelgren/IJOC-Efficient>.

We performed this test 100 times for each combination of the values $N \in \{10^4, 10^5, 10^6, 10^7\}$ and $\mu \in \{0, 0.001, 0.01, 0.1, 1, 10\}$. We used various values for δ and found that the results were quite similar, but we decided to use a value of $\delta = 0.3$. For each test we recorded the total insertion time for both algorithms, the final depth of a BoT, and the final number of nodes stored in the BoT and in the list.

With μ close to zero, there is little or no separation between early generated solutions and later generated ones, and all are likely to be Pareto. With large values of μ , there is significant separation between early generated solutions and later generated ones, the latter being much more likely to be Pareto. Figure 5 shows an example of solutions generated during this experiment for $\mu = 0.1, 1$ and 10 and for $N = 100$; of these, only nondominated solutions are in the BoT at the end of the test.

The minimum, maximum, and geometric means of the CPU times and final depths of the tree can be found in Tables 1–3. The symbols A0–A4 and L indicate runs in which the various rebalancing approaches of a BoT and the dynamic list were used for storing solutions. Also, entries in Tables 1–3 that contain dashes are those for which no results are available because individual runs took over 12 hours to complete and were therefore terminated. The symbols $\otimes$ and $\dagger$ indicate results for which, due to the large amount of time taken for each individual run, we were unable to perform the test 100 times. For these results, each test was instead run five times (for $\otimes$) and three times (for $\dagger$).

We use performance profiles (Dolan and Moré 2002) to show the relative effectiveness of the various rebalancing approaches, in terms of CPU time and tree depth, in Figure 6. We omit the results of the list implementation and rebalancing approach A1 as they performed poorly in terms of CPU time when compared to the other approaches, and only show profiles for

Figure 5. (Color online) Example of Randomly Generated Solutions; $N = 100$

$N = 10^7$ and for a few values of μ for reasons of space. For $\mu = 0$ (Figure 6(a)), A2, shown by the curve with + symbols, has the best time performance, although it does not dominate A0 (the solid line), which in turn is at least 30% worse than A2 in half of the instances but never worse than 40%. Algorithms A3 and A4 perform similarly and are dominated by A0 and A2. For $\mu = 0.001$, instead (see Figure 6(b)), A0 dominates A2, which is, nevertheless, never more than 15% worse than the best performance. Algorithms A3 and A4 (lines with * and $\circ$, respectively) perform worse still, and this is confirmed for $\mu = 1$ (see Figure 6(c)), where A4 fares slightly better than A3. This is confirmed for $\mu = 10$ (not shown).

While A2 and A0 perform well in general, A0 carries the risk of an unbalanced tree, as reported in Figure 6(d), where tree depth is shown to be up to 3.5 times worse than the best. Algorithms A3 and A4 have the best performance in terms of maximum tree depth, which, however, reflects in poor time performance as shown in Figure 6(b).

As shown in Tables 1–3, with approaches A0, A2, A3, and A4, the BoT is able to process inserted solutions much more quickly than the dynamic list; A1 is more efficient than the list, but far slower than the other approaches. The performance difference grows with N .

Also, for most values of N and μ , A0 typically performs the best in terms of running time, followed

Table 1. Time and Depth of the Tree for the Random Points Test ($\mu = 0, 0.001$)

N	Rebal type	Time (s)			Depth			Time (s)			Depth		
		Min	Avg	Max	Min	Avg	Max	Min	Avg	Max	Min	Avg	Max
10 ⁴	A0	0.08	0.09	0.10	34	42.6	53	0.03	0.03	0.04	25	31.6	41
	A1	2.22	2.31	2.55	16	16.2	18	0.28	0.30	0.36	12	12.0	13
	A2	0.16	0.20	0.35	16	17.1	18	0.05	0.05	0.08	12	13.6	16
	A3	0.11	0.16	0.38	16	17.0	18	0.04	0.05	0.1	12	13.1	14
	A4	0.11	0.16	0.39	16	17.0	18	0.04	0.05	0.09	12	13.1	14
	L	5.65	5.94	6.27	—	—	—	0.35	0.37	0.41	—	—	—
10 ⁵	A0	4.52	4.85	4.94	42	51.3	62	0.64	0.70	0.79	27	36.7	52
	A1	448.54	476.75	520.94	19	19.2	20	9.09	9.57	11.02	13	14.0	15
	A2	6.31	9.91	32.54	19	20.2	22	0.80	0.91	1.12	14	17.0	19
	A3	5.83	9.81	41.62	19	19.8	21	0.84	0.99	1.30	14	14.9	16
	A4	5.69	9.77	43.83	19	19.7	21	0.81	0.99	1.31	14	14.9	16
	L	668.19	694.63	747.21	—	—	—	19.30	19.65	20.73	—	—	—
10 ⁶	A0	126.01	128.91	135.84	45	54.6	70	17.32	18.39	18.79	34	42.2	54
	A1	—	—	—	—	—	—	427.00	434.40	465.09	17	17	17
	A2	124.02	143.02	254.55	21	22.4	24	18.93	19.82	21.71	17	21.0	24
	A3	137.08	168.69	328.82	21	21.0	22	20.12	21.35	24.55	16	16.3	17
	A4	136.23	165.19	306.11	20	21.0	22	19.85	21.07	24.2	16	16.44	17
	L	—	—	—	—	—	—	—	—	—	—	—	—
10 ⁷	A0	1,684.59	2,122.15	2,809.04	50	57.9	68	431.54	459.64	485.31	36	46.3	63
	A1	—	—	—	—	—	—	39,654.42	39,922.79	40,136.91	17	17.0	17
	A2	1,287.45	1,915.42	4,908.04	22	24.4	26	423.87	466.23	504.79	20	26.2	32
	A3	1,599.44	2,422.43	6,690.78	21	21.3	23	468.92	491.24	536.21	18	18.1	19
	A4	1,611.95	2,413.61	6,880.00	21	21.3	23	456.67	482.73	522.69	17	18.0	19
	L	—	—	—	—	—	—	†63,936.51	67,937.70	70,214.80	—	—	—

Table 2. Time and Depth of the Tree for the Random Points Test ($\mu = 0.01, 0.1$)

N	Rebal type	Time (s)			Depth			Time (s)			Depth		
		Min	Avg	Max	Min	Avg	Max	Min	Avg	Max	Min	Avg	Max
10 ⁴	A0	0.02	0.02	0.02	14	21.6	36	0.01	0.01	0.02	11	12.8	16
	A1	0.11	0.11	0.13	10	10.2	11	0.05	0.05	0.05	8	8.7	10
	A2	0.03	0.03	0.03	10	12.2	16	0.02	0.02	0.02	8	10.5	14
	A3	0.03	0.03	0.03	10	11.1	12	0.02	0.02	0.02	9	9.4	11
	A4	0.03	0.03	0.03	10	11.1	12	0.02	0.02	0.02	9	9.4	11
	L	0.11	0.12	0.14	—	—	—	0.04	0.04	0.06	—	—	—
10 ⁵	A0	0.34	0.36	0.39	18	25.4	47	0.23	0.23	0.24	13	15.1	21
	A1	3.14	3.22	3.35	12	12.1	13	1.18	1.20	1.30	10	10.1	12
	A2	0.41	0.43	0.47	13	15.9	20	0.26	0.26	0.27	10	14.2	18
	A3	0.47	0.50	0.55	12	13.0	14	0.33	0.34	0.36	10	11.1	12
	A4	0.48	0.50	0.56	12	13.0	14	0.34	0.34	0.37	10	11.1	12
	L	4.12	4.27	4.47	—	—	—	1.25	1.31	1.37	—	—	—
10 ⁶	A0	6.61	7.07	7.88	21	29.6	43	3.50	3.58	3.76	16	18.0	21
	A1	91.56	96.47	105.81	13	13.9	15	32.01	32.41	32.85	12	12.1	13
	A2	6.92	7.37	7.84	16	19.8	25	3.69	3.78	3.92	14	17.7	20
	A3	8.31	8.83	9.56	14	14.9	16	4.89	5.03	5.25	12	13.0	14
	A4	8.24	8.72	9.14	14	14.9	16	4.90	5.04	5.21	12	13	14
	L	211.81	215.19	224.98	—	—	—	—	—	—	—	—	—
10 ⁷	A0	94.29	184.30	188.04	23	33.2	48	68.06	70.71	74.91	19	20.9	24
	A1	1,874.98	4,396.94	6,018.77	15	15.1	16	948.30	1,027.48	1,132.85	13	13.9	15
	A2	93.15	187.70	191.03	18	23.5	29	69.08	71.76	74.85	16	20.7	23
	A3	108.64	206.37	210.70	16	16.3	17	85.63	88.87	93.43	14	14.8	15
	A4	109.23	200.11	207.89	16	16.3	17	84.15	87.15	92.13	14	14.8	15
	L	2,786.57	14,687.80	15,234.72	—	—	—	2,136.03	2,210.22	2,452.43	—	—	—

Table 3. Time and Depth of the Tree for the Random Points Test ($\mu = 1, 10$)

N	Rebal type	Time (s)			Depth			Time (s)			Depth		
		Min	Avg	Max	Min	Avg	Max	Min	Avg	Max	Min	Avg	Max
10 ⁴	A0	0.01	0.01	0.01	9	11.2	16	0.01	0.01	0.01	9	10.9	17
	A1	0.03	0.04	0.04	7	8.0	9	0.03	0.03	0.03	7	7.4	9
	A2	0.01	0.01	0.02	8	9.8	13	0.01	0.01	0.01	7	8.9	12
	A3	0.01	0.02	0.02	8	8.6	10	0.01	0.01	0.01	7	8.0	9
	A4	0.01	0.02	0.02	8	8.6	10	0.01	0.01	0.01	7	8.0	9
	L	0.03	0.03	0.04	—	—	—	0.03	0.03	0.04	—	—	—
10 ⁵	A0	0.19	0.19	0.19	11	12.6	17	0.16	0.16	0.17	12	16.7	24
	A1	0.55	0.56	0.61	8	8.8	10	1.31	1.40	1.49	8	8.53	10
	A2	0.20	0.20	0.21	9	12.3	17	0.23	0.23	0.25	9	11.1	15
	A3	0.25	0.25	0.26	9	9.3	10	0.28	0.29	0.31	8	9.0	10
	A4	0.25	0.25	0.26	9	9.3	10	0.28	0.29	0.31	8	9.0	10
	L	0.50	0.52	0.57	—	—	—	2.51	2.67	2.91	—	—	—
10 ⁶	A0	2.39	2.42	2.47	14	15.3	19	1.92	1.94	1.96	11	12.6	20
	A1	12.15	12.44	12.85	10	10.2	11	5.62	5.72	5.89	8	8.7	10
	A2	2.47	2.50	2.57	13	15.2	19	1.96	1.98	2.01	10	12.5	20
	A3	3.48	3.54	3.64	10	11.1	12	2.53	2.56	2.61	9	9.41	11
	A4	3.48	3.55	3.63	10	11.1	12	2.54	2.57	2.62	9	9.41	11
	L	13.18	13.44	13.87	—	—	—	5.20	5.38	5.54	—	—	—
10 ⁷	A0	34.49	34.97	35.78	16	18.0	21	23.28	23.97	24.25	13	15.23	19
	A1	317.90	321.25	329.82	12	12.1	13	121.22	124.55	126.14	10	10.1	12
	A2	35.15	36.10	37.20	16	18.0	21	24.24	24.45	24.87	13	15.2	19
	A3	48.55	49.79	52.08	12	13.0	14	34.57	35.44	36.21	10	11.2	13
	A4	48.54	49.84	51.46	12	13.0	14	35.22	35.56	36.78	10	11.2	13
	L	419.88	435.56	458.88	—	—	—	131.85	134.998	137.2	—	—	—

Figure 6. Performance Profiles of the Rebalancing Algorithms A0 (Solid Line), A2 (+), A3 (*), and A4 (o)

Notes. All profiles are w.r.t. results for $N = 10^7$. For each algorithm A , a point (x, y) on A 's curve indicates that for $y\%$ of the instances A 's performance is at most x times worse than the best algorithm.

closely by A2. However, A3 and A4 perform the best in terms of tree depth. Because an increasing μ implies an increasing number of eliminated solutions at each insertion (and, in general, fewer solutions stored in the data structure at any time), the time taken to insert solutions decreases as the value of μ increases. Furthermore, the larger the value of μ , the smaller the gap between the CPU time spent by the linked list and the BoT. We conclude that the BoT is more scalable and performs better than a linked list for storing randomly generated solutions. Of all rebalancing algorithms tested, only A1 appears to be inefficient, while no rebalancing at all (algorithm A0) seems to have a limited impact on the performance. The results suggest that A2 seems to combine good performance in terms of time and maximum tree depth, and is therefore the rebalancing algorithm of choice for our next tests.

4.2. Using a BoT in Branch-and-Bound Algorithms

To test the utility of a BoT more in depth when utilized in a BB, we performed another set of tests in which a variety of instances of BOMILP were solved

using branch and bound. To ensure the results were not too dependent on one solver, we ran our tests on two implementations, described in Belotti et al. (2012) and Adelgren and Gupte (2016). The latter utilizes this data structure not only for storing found solutions, but also to check for domination of bound sets, and hence, for fathoming. In addition, it uses *warm starting*, i.e., it generates an initial set of points before the BB is started. This BB was able to solve larger instances in under eight hours and also did not run into numerical difficulties when solving instances of type II from Boland et al. (2014). For both BB solvers, we have used approach A2 for rebalancing.

We used three classes of BOMILP instances: two sets of instances described and used in Belotti et al. (2012) and Boland et al. (2014), and one that we have created. The latter set contains instances, which we call *geometrical*, that, while of rather simple structure, contain by design a large number of Pareto points. Testing these instances forces the BB algorithms to create numerous BB nodes and solutions, thus assessing a BoT in a more realistic, large-scale setting where the number of

solutions inserted are closer to those of our first experiment. These are very simple BOMILPs with a large Pareto set for which a single BB node is solved very quickly. The model is as follows:

$$\begin{aligned}
 & \max x \\
 & \max y \\
 & \text{s.t. } y - s_i x \leq b_i - a_i s_i \quad \forall i \in \{1, \dots, k\}, \\
 & \quad x - 0.1z \leq 0.025, \\
 & \quad x - 0.1z \geq -0.025, \\
 & \quad x, y \geq 0, z \in \mathbb{Z},
 \end{aligned} \tag{4}$$

where $a_1=0$, $a_{i+1}=a_i+N/k$ for $i \in \{1, \dots, k\}$, $b_i = \sqrt{N^2 - a_i^2}$ for $i \in \{1, \dots, k+1\}$, $s_i = (b_{i+1} - b_i)/(a_{i+1} - a_i)$ for $i \in \{1, \dots, k\}$, and N and k are parameters. All geometrical instances and those of Belotti et al. (2012) are available online at the aforementioned repository, <https://github.com/Nadelgren/IJOC-Efficient>. For the instances by Boland et al. (2014), we refer to the same article.

The results from the branch-and-bound tests are given in Tables 4–7. We use T and L to represent the implementations using a BoT and the dynamic list, respectively. From Belotti et al. (2012), 30 instances

Table 4. BB Tests: Results for Algorithm of Belotti et al. (2012)

Size	M	Time		Time (insert)		#BBnodes (×10 ²)	#inserts (×10 ³)	#nodes	Tree depth	
		T	L	T	L					
Instances from Belotti et al. (2012) (averaged)										
60	10	15.0	15.1	0.007	0.007	8.11	1.790	62.6	6.3	
	25	11.6	11.6	0.005	0.004	6.32	1.155	60.0	6.6	
	50	16.4	16.1	0.012	0.007	8.22	1.838	64.1	6.4	
80	10	38.0	38.2	0.011	0.009	13.55	2.729	82.9	6.9	
	25	32.6	32.9	0.005	0.006	11.42	2.046	82.1	7.0	
	50	42.6	42.9	0.011	0.008	14.11	2.988	80.6	6.9	
Instances from Boland et al. (2014), type 1										
80	50	101	102	0.93	1.70	38	68	931	11	
		76	78	0.21	1.05	34	120	565	10	
		88	90	0.25	1.06	47	65	857	11	
		61	61	0.15	0.62	28	59	900	11	
		66	66	0.11	0.37	44	52	706	10	
		77	78	0.24	0.85	37	70	779	11	
		200	44	45	0.10	0.27	22	38	926	11
	32	33	0.06	0.27	15	38	571	10		
	49	50	0.13	0.29	26	27	860	11		
	38	38	0.17	0.32	18	38	903	11		
	44	45	0.07	0.30	29	31	717	11		
	41	42	0.10	0.29	21	34	783	11		
	300	57	57	0.49	0.75	21	35	936	11	
	33	34	0.05	0.32	15	37	568	10		
	48	49	0.05	0.19	24	24	862	11		
	37	38	0.08	0.30	16	35	906	11		
	43	43	0.08	0.22	27	28	724	11		
	43	43	0.10	0.31	20	31	786	11		
	160	500	13,584	13,862	49.30	177.68	1,502	3,274	2,272	20
			15,992	15,790	14.80	69.07	2,469	2,807	2,102	22
			11,721	11,857	8.12	89.31	2,030	1,858	2,176	19
5,354			5,442	8.91	90.58	819	1,319	4,600	14	
2,384			2,384	2.89	17.53	364	572	2,366	22	
7,987			8,043	10.88	70.49	1,176	1,667	2,575	19	
2,000			10,933	11,088	34.37	165.50	1,219	2,484	2,272	14
12,505		12,489	13.33	59.33	1,954	2,156	2,111	12		
10,623		10,606	6.35	74.90	1,847	1,448	2,193	14		
4,846		5,005	8.66	84.26	709	958	4,597	14		
2,028		2,062	2.28	14.83	304	461	2,379	13		
6,775		6,857	8.95	62.04	989	1,280	2,583	13		
3,000		7,775	7,852	12.58	79.05	1,193	2,418	2,274	13	
11,940		12,051	13.67	62.07	1,900	2,090	2,109	13		
9,872		10,298	7.32	74.13	1,817	1,406	2,195	13		
4,451		4,532	6.56	68.13	696	917	4,604	14		
1,994		2,012	2.41	14.59	297	449	2,382	13		
6,054		6,162	7.24	51.48	968	1,240	2,585	13		

Table 5. BB Tests: Algorithm of Adelgren and Gupte (2016)

Size	Time		Time (insert)		Time (other)		#BBnodes	#inserts		#nodes		Tree depth	
	T	L	T	L	T	L		T	L	T	L		
Instances from Belotti et al. (2012) (averaged)													
60	5	5	0.00	0.00	0.00	0.01	76	189		65	64	6	
80	13	13	0.00	0.00	0.01	0.01	109	247		83		7	
Instances from Boland et al. (2014), type 1													
80	10	11	0.01	0.09	0.01	0.15	363	6,429		943	930	11	
	5	5	0.02	0.03	0.00	0.13	220	5,066		590	583	12	
	18	19	0.01	0.16	0.02	0.45	580	9,491		885	883	10	
	6	7	0.02	0.05	0.01	0.15	235	4,919		928	920	11	
	7	7	0.02	0.07	0.01	0.09	269	4,795		737	732	10	
	8	8	0.02	0.07	0.01	0.16	311	5,923		804	797	11	
	205	211	0.12	1.47	0.09	3.44	1,418	45,501		2,293	2,285	17	
	129	139	0.09	1.91	0.12	2.80	997	41,534		2,160	2,136	13	
160	104	107	0.07	0.64	0.08	1.59	924	29,895		2,250	2,203	13	
	400	414	0.30	4.86	0.13	9.76	2,315	79,421		4,785	4,675	17	
	108	111	0.03	0.82	0.07	1.94	743	22,397		2,489	2,408	13	
	164	170	0.09	1.48	0.10	3.11	1,176	39,850		2,658	2,610	14	
	6,266	6,713	1.70	205.34	1.16	189.91	6,436	486,497	487,183	11,184	10,970	48	
	4,892	5,457	1.84	243.26	0.74	242.09	5,278	354,556	354,343	12,070	11,689	18	
	2,742	2,924	1.41	105.24	0.30	65.32	3,367	246,820		12,488	12,138	15	
	5,291	5,629	1.56	176.53	0.57	132.42	5,925	364,204	363,980	12,863	12,531	16	
320	2,404	2,522	1.38	49.21	0.43	60.80	3,773	223,789		9,955	9,703	18	
	4,035	4,329	1.57	135.49	0.58	119.31	4,803	322,156	322,168	11,664	11,361	21	
	Instances from Boland et al. (2014), type 2												
	800	1	1	0.00	0.00	0.00	0.00	39	254		54		6
		2	2	0.00	0.00	0.01	0.00	46	279		64		7
		4	4	0.00	0.00	0.00	0.00	98	589		90		7
		9	9	0.00	0.01	0.01	0.01	180	900		129	128	7
		3	3	0.00	0.00	0.00	0.00	75	440		80	79	7
1,250	9	9	0.00	0.00	0.01	0.01	147	867		140		8	
	19	19	0.00	0.00	0.00	0.02	310	1,922	1,935	200	201	8	
	19	19	0.02	0.00	0.02	0.06	302	1,962		245		9	
	46	46	0.01	0.01	0.00	0.06	542	3,328		281		9	
	20	20	0.01	0.00	0.01	0.03	294	1,816	1,819	210	210	8	
2,500	514	515	0.04	0.13	0.29	1.42	2,893	7,650		306		9	
	944	945	0.01	0.18	0.31	2.28	3,831	9,409	9,406	370		9	
	1,243	1,244	0.02	0.21	0.45	2.73	4,007	9,849		451	449	10	
	3,185	3,199	0.02	0.45	0.86	6.24	8,468	18,665		386	385	9	
	1,177	1,180	0.02	0.22	0.43	2.73	4,404	10,725	10,724	375	374	9	

are available for each problem size; the set of points encountered by the BB by the same authors in one of these instances is depicted in Figure 7. There are five (resp. four) instances from Boland et al. (2014) for each problem size of Type I (resp. II). The BB by Belotti et al. (2012) works by creating M initial Pareto points via solving M single-objective MIPs; hence, we tested using different values of M . In Tables 4 and 5, we report geometric means of the results on instances from Belotti et al. (2012), and provide results for every instance from Boland et al. (2014) (with geometric means reported in bold). We report the total solve time, the time spent inserting solutions (Time (insert)), the time spent doing other data structure operations such as fathoming (Time (other), only reported for the BB

by Adelgren and Gupte (2016) as it is negligible otherwise), the number of BB nodes, the number of insertions, the number of nodes in the structure at termination (#nodes), and the depth of the BoT. Again, problem sizes are reported as the number of variables, which in most cases also equals the number of constraints (for instances from Boland et al. 2014, Type II, the number of constraints is slightly greater than the number of variables).

In principle, a BB algorithm is expected to follow the same path regardless of the data structure, with the same number of BB nodes explored and of solutions stored. While this is true for most tests, there were discrepancies when using the BB by Adelgren and Gupte (2016) for solving instances from Belotti

Table 6. BB Tests: Algorithm of Belotti et al. (2012) on Geometrical Instances

N	k	Time		Time (insert)		#BBnodes ($\times 10^2$)	#inserts ($\times 10^3$)	#nodes ($\times 10^3$)	Tree depth
		T	L	T	L				
500	1	15	25	0.8	6.8	10	35	5	24
	5	15	16	0.6	2.3	10	35	5	23
	25	21	35	1.0	6.2	10	34	5	25
	125	48	47	0.5	2.2	10	33	5	27
		22	28	0.7	3.8	10	34	5	25
1,000	1	36	39	2.5	10.3	20	70	10	25
	5	39	41	2.5	11.2	20	70	10	26
	25	47	58	2.6	13.8	20	69	10	28
	125	103	115	3.1	16.6	20	67	10	37
		51	57	2.6	12.7	20	69	10	29
2,000	1	111	185	9.0	95.7	40	140	20	23
	5	100	186	8.3	96.2	40	140	20	30
	25	133	212	9.2	96.9	40	139	20	32
	125	224	345	9.7	113.7	40	137	20	40
		135	224	9.0	100.4	40	139	20	31
4,000	1	370	966	41.9	681.9	80	280	40	32
	5	350	1,070	53.0	774.3	80	280	40	31
	25	409	1,025	42.6	691.1	80	279	40	30
	125	650	1,308	48.4	751.8	80	277	40	27
		431	1,085	46.2	723.7	80	279	40	30
8,000	1	1,952	6,275	339.2	5,185.5	160	560	80	39
	5	1,861	4,786	289.5	3,893.4	160	560	80	38
	25	2,113	5,363	329.6	4,339.8	160	559	80	34
	125	3,793	6,189	329.2	4,679.2	160	556	80	89
		2,323	5,619	321.3	4,499.8	160	559	80	46
16,000	1	8,973	22,803	1,627.5	19,345.8	320	1,120	160	41
	5	8,059	26,171	1,638.7	22,958.1	320	1,120	160	41
	25	12,250	23,252	1,970.6	20,019.5	320	1,119	160	41
	125	13,464	29,552	1,674.1	24,950.9	320	1,116	160	118
		10,450	25,305	1,722.3	21,702.8	320	1,119	160	53

Table 7. BB Tests: Algorithm of Adelgren and Gupte (2016) on Geometrical Instances

N	k	Time		Time (insert)		Time (insert)		#BBnodes ($\times 10^2$)	#inserts ($\times 10^3$)	#nodes ($\times 10^3$)	Tree depth
		T	L	T	L	T	L				
500	1	12	56	1.93	16.13	0.09	23.85	10	20	5	16
	5	119	158	1.08	11.37	0.16	21.86	10	20	5	15
	25	131	160	0.85	9.59	0.12	18.19	10	20	5	15
	125	163	200	0.93	13.19	0.14	20.78	10	20	5	14
		75	130	1.13	12.34	0.12	21.07	10	20	5	15
1,000	1	39	273	5.31	90.61	0.16	104.85	20	40	10	16
	5	458	707	3.13	60.20	0.45	133.60	20	40	10	15
	25	511	762	3.09	66.47	0.38	135.67	20	40	10	15
	125	619	895	3.08	69.18	0.38	149.48	20	40	10	16
		273	603	3.55	70.77	0.32	129.83	20	40	10	15
2,000	1	133	1,346	14.91	389.46	0.50	546.32	40	80	20	17
	5	1,837	3,314	16.28	349.33	1.45	756.48	40	80	20	18
	25	2,032	3,251	13.38	305.68	1.54	607.06	40	81	21	20
	125	2,298	3,725	13.47	316.52	1.68	760.57	40	81	20	18
		1,034	2,711	14.46	338.72	1.17	660.93	40	81	20	18
4,000	1	996	6,695	80.32	1,596.64	3.07	2,645.33	80	160	40	18
	5	8,530	15,352	106.16	1,518.61	8.59	3,974.04	80	160	40	18
	25	8,273	14,081	99.59	1,423.79	5.36	3,189.86	80	162	42	18
	125	10,466	16,394	91.13	1,386.03	7.60	3,239.40	80	160	40	18
		5,208	12,411	93.79	1,479.00	5.73	3,228.40	80	160	40	18

Figure 7. (Color online) Solutions Found by the BB Method in Belotti et al. (2012) with 80 Variables and 80 Constraints

et al. (2012) and Boland et al. (2014). We show the number of BB nodes, of insertions, and of nodes in Table 5 when these differ, in general by a small amount but up to 3% in a few cases. Extra tests suggest that these discrepancies are due to small differences in the solutions contained in either data structure: the insertion of a solution π into a BoT and a list that have, initially, the same set of solutions may result in different sets if, for instance, the extremes of segment $\pi.S$ are very close to previous solutions, say within a tolerance of 10^{-4} , since the operations of Algorithm 1 allow for multiple changes in a subtree of the BoT. This effect can be amplified as it may result in one BB node to be fathomed or explored, the latter resulting in an entire new BB subtree being explored.

Tables 4 and 5 show that the time spent by the BoT is much less than that spent using the linked list. Because the total BB time is much greater than insertion time, the performance advantage of the BoT over the list is less influential yet clear. However, these results confirm those of experiment 1: even within a BB, insertion into a linked list often takes orders of magnitude longer than for a BoT. For all instances of the two considered classes, the number of insertions amounts to up to a few million, yet the final structure only has a few thousand solutions (see e.g., Table 4, instances by Boland et al. 2014, for size 160); this suggests a pattern similar to that with larger values of μ (see Tables 1–3).

Both solvers that we used are rather rudimentary implementations—the alternative method by Boland et al. (2014) easily outperforms the BB solver by Belotti et al. (2012). This explains why the BB times reported can be hours even if the BB algorithms visit up to only a few thousand BB nodes, well below the current state-of-the-art BB algorithms for single-objective MILP, where millions, or tens of millions, of nodes can be explored in the same time for much larger instances.

The advantages of BoT over list are more apparent on geometrical instances, which, albeit simple, allow for a thorough stress test of both the BB solvers and the data structures. For these instances, the time spent by the data structures is a significant portion of the total CPU time. All BB nodes are solved very quickly, and as a result many solutions are found and many BB nodes are explored in a short time. Inserting all of these solutions, while efficient for a BoT, requires much more computational effort for a linked list, to the point that the amount of time spent inserting solutions in a large linked list becomes a substantial percentage of the total CPU time.

These results suggest that an efficient data structure will be essential when faster and more stable BB solvers (for instance, using quicker fathoming rules as in Belotti et al. 2016) become available: as the time for each BB node decreases and the number of inserted solutions increases, the data structure must be efficient enough that the insertion function only takes a small percentage of the total CPU time.

5. Concluding Remarks

We have introduced the *biobjective tree*, a variant of a binary tree to efficiently store nondominated solutions of BOMILPs. The BoT is equipped with an insertion procedure for adding points and possibly eliminating several other points that are dominated by the newly inserted one. The BoT also has the desirable property that an *in-order* pass produces a sorted list of nondominated points. We tested the practical value of the BoT with two experiments. The results show that a BoT provides a more efficient method for storing solutions to BOMILP than a list of points. They also show that a BoT is a very useful tool when used in BB methods for solving BOMILPs.

It is worth noting here that we have focused on a simple BST data structure, instead of one with better balancing properties, because we wanted to avoid the travails of studying and implementing complicated rebalancing procedures, mixed with the already taxing duty of maintaining other desirable properties of a BoT, i.e., the aforementioned ability to hold nondominated points at any step.

Data structures like the BoT are preferable over lists when inserting large numbers of points in an undefined order, such as in a BB algorithm. Because current state-of-the-art BB algorithms for BOMILP can tackle relatively small problems, one can observe the impact of the data structure in those classes of problems where BB node solution is fast; i.e., those problems that admit a simple and compact structure yet contain a very large number of nondominated points. We speculate that future implementations of a BB algorithm for BOMILP will expose even more clearly the advantages of a BST data structure. However, algorithms that find solutions

in a nonrandom order might benefit from more specific, and perhaps less sophisticated, data structures.

Acknowledgments

Part of the work was carried out at the Department of Mathematical Sciences, Clemson University (Nathan Adelgren).

References

- Adelgren N, Gupte A (2016) Branch-and-bound for biobjective mixed-integer programming. Accessed January 9, 2018, http://www.optimization-online.org/DB_HTML/2016/10/5676.html.
- Belotti P, Soylu B, Wiecek MM (2012) A branch-and-bound algorithm for biobjective mixed-integer programs. Technical report, Clemson University, Clemson, SC. http://www.clemson.edu/ces/math/technical_reports/belotti.bb-bicriteria.pdf.
- Belotti P, Soylu B, Wiecek MM (2016) Fathoming rules for biobjective mixed integer linear programs: Review and extensions. *Discrete Optim.* 22(B):341–363.
- Boland N, Charkhgard H, Savelsbergh M (2014) The triangle splitting method for biobjective mixed integer programming. Lee J, Vygen J, eds. *17th Internat. Conf. Integer Programming Combinatorial Optim. (IPCO)*, Lecture Notes in Computer Science, Vol. 8494 (Springer, Boston), 162–173.
- De Loera JA, Hemmecke R, Köppe M (2009) Pareto optima of multicriteria integer linear programs. *INFORMS J. Comput.* 21(1): 39–48.
- Dolan ED, Moré JJ (2002) Benchmarking optimization software with performance profiles. *Math. Programming* 91(2):201–213.
- Ehrgott M (2005) *Multicriteria Optimization* (Springer, Berlin).
- Jozefowiez N, Laporte G, Semet F (2012) A generic branch-and-cut algorithm for multiobjective optimization problems: Application to the multilabel traveling salesman problem. *INFORMS J. Comput.* 24(4):554–564.
- Kiziltan G, Yucaoglu E (1983) An algorithm for multiobjective zero-one linear programming. *Management Sci.* 29(12):1444–1453.
- Knuth DE (1998) *The Art of Computer Programming: Sorting and Searching*, Vol. 3 (Pearson Education, London).
- Lenstra HW Jr (1983) Integer programming with a fixed number of variables. *Math. Oper. Res.* 8(4):538–548.
- Mavrotas G, Diakoulaki D (2005) Multi-criteria branch and bound: A vector maximization algorithm for mixed 0-1 multiple objective linear programming. *Appl. Math. Comput.* 171(1):53–71.
- Overmars MH, Van Leeuwen J (1982) Dynamic multi-dimensional data structures based on quad- and k -d trees. *Acta Informatica* 17(3):267–285.
- Özpeynirci Ö, Köksalan M (2010) An exact algorithm for finding extreme supported nondominated points of multiobjective mixed integer programs. *Management Sci.* 56(12):2302–2315.
- Przybylski A, Gandibleux X, Ehrgott M (2010) A two phase method for multi-objective integer programming and its application to the assignment problem with three objectives. *Discrete Optim.* 7(3):149–165.
- Ralphs TK, Saltzman MJ, Wiecek MM (2006) An improved algorithm for solving biobjective integer programs. *Ann. Oper. Res.* 147(1):43–70.
- Sourd F, Spanjaard O (2008) A multiobjective branch-and-bound framework: Application to the biobjective spanning tree problem. *INFORMS J. Comput.* 20(3):472–484.
- Soylu B, Yıldız GB (2016) An exact algorithm for biobjective mixed integer linear programming problems. *Comput. Oper. Res.* 72(Suppl C):204–213.
- Stidsen T, Andersen KA, Dammann B (2014) A branch and bound algorithm for a class of biobjective mixed integer programs. *Management Sci.* 60(4):1009–1032.
- Sun M (2006) A primogenitary linked quad tree data structure and its application to discrete multiple criteria optimization. *Ann. Oper. Res.* 147(1):87–107.
- Sun M, Steuer RE (1996) Quad-trees and linear lists for identifying nondominated criterion vectors. *INFORMS J. Comput.* 8(4): 367–375.
- Vincent T, Seipp F, Ruzika S, Przybylski A, Gandibleux X (2013) Multiple objective branch and bound for mixed 0-1 linear programming: Corrections and improvements for the biobjective case. *Comput. Oper. Res.* 40(1):498–509.