

# Task Assignment in Tree-Like Hierarchical Structures

Cem Evrendilek<sup>a</sup>, Ismail Hakki Toroslu<sup>\*,b</sup>, Sasan Hashemi<sup>b</sup>

<sup>a</sup>*İzmir University of Economics  
Computer Engineering Department  
35330 İzmir, Turkey*

<sup>b</sup>*Middle East Technical University  
Computer Engineering Department  
06531 Ankara, Turkey*

---

## Abstract

Most large organizations, such as corporations, are hierarchical organizations. In hierarchical organizations each entity in the organization, except the root entity, is a sub-part of another entity. In this paper we study the task assignment problem to the entities of tree-like hierarchical organizations. The inherent tree structure introduces an interesting and challenging constraint to the standard assignment problem. When a task is assigned to an entity in a hierarchical organization, the whole entity, including its sub-entities, is responsible from the execution of that particular task. In other words, if an entity has been assigned to a task, neither its descendants nor its ancestors can be assigned to a task. Sub-entities cannot be assigned as they have an ancestor already occupied. Ancestor entities cannot be assigned since one of their sub-entities has already been employed in an assignment. In the paper, we formally introduce this new version of the assignment problem called Maximum Weight Tree Matching (*MWTM*), and show its NP-hardness. We also propose an effective heuristic solution based on an iterative LP-relaxation to it.

*Key words:* task assignment, hierarchy constraints, NP-hardness, heuristic solution, integer linear programming, linear programming relaxation

---

## 1. Introduction

In the standard assignment problem (or as sometimes referred to linear assignment problem) [1], the number of tasks and the number of agents are equal, and a scalar value is used to represent the cost/performance of assigning a task to an agent. The objective of the assignment problem is to determine an assignment such that each task is assigned to a different agent and the summation

---

\*Corresponding author. Tel.: +90 312 210 5585

*Email addresses:* [cem.evrendilek@ieu.edu.tr](mailto:cem.evrendilek@ieu.edu.tr) (Cem Evrendilek),  
[toroslu@ceng.metu.edu.tr](mailto:toroslu@ceng.metu.edu.tr) (Ismail Hakki Toroslu), [sasan@ceng.metu.edu.tr](mailto:sasan@ceng.metu.edu.tr) (Sasan Hashemi)

of the costs/profits of the assignment is minimized/maximized. Many different variations of this problem have already been studied including *Generalized Assignment Problem* [2, 3, 4, 5]. In this work, we also investigate a new version of the standard assignment problem which appears in real-life applications.

In real-life, most of the large organizations such as corporations, governments, military etc., have hierarchical structures. Hierarchical organizations are nothing but trees where each node corresponds to an entity in the organization, and entity sub-entity relationships are represented as parent-child relationships.

In the standard assignment problem, agents are flat, and have got no structure imposed on them one task is assigned to one agent. However, in Maximum Weight Tree Matching (*MWTM*) problem, since agents are organized as a tree, and sub-entities in the tree represent sub-parts of the agents, an additional constraint, named hereafter as *hierarchy constraint*, is introduced to the assignment problem: When a task is assigned to an agent, no other assignment can be made to its sub-entities, as they are assumed to be a part of an agent already assigned. This constraint indirectly implies another constraint. Since an agent should be assigned to a task as a whole along with its sub-parts, when one of its sub-parts has already been assigned, then it cannot be assigned itself to any task. In other words, if an agent is assigned to a task, none of its ancestors in the tree can be assigned at all. In a more general term, on every path from the root to a leaf in a tree, there could be at most a single assignment. This, in turn, is easily seen to lead to the observation that the number of leaves in the tree should be at least equal to the number of tasks to be executed. Otherwise no feasible assignment exists.

A simpler version of *MWTM* problem where each node has the same assignment weight for all the tasks to be performed has been introduced in [6]. It is called “tree like weighted set packing” in [6] since the set-subset relationships form a tree, and the weight assigned to each set (or each node in the tree) can be interpreted as the weight of assigning a task to that node. The same hierarchy (independence) constraint has been enforced to prevent the selection of two sets having set-subset relationships (either directly or indirectly), and finally the number of sets to be packed (selected) is given to maximize the total weight. That problem effectively becomes a simpler version of the problem studied in this paper, and an effective dynamic programming solution to it has been developed in [6].

Although many different versions of assignment problems have been defined and explored, there are only a very few problems remotely related to *MWTM* problem reported in the literature, such as [7], and [8]. Similar to *MWTM*, both of these problems introduce different kinds of set constraints on the vertices of a bipartite graph, and they have both been shown to be NP-hard. Therefore, heuristic solutions have been proposed, namely a greedy heuristic for [7], and a genetic algorithm based solution for [8], and these solutions have been shown to be quite effective.

*MWTM* problem has already been introduced in [9], and a generic heuristic (genetic algorithm) has been used to solve it. In [9], it has been shown that GA works quite effectively in terms of solution quality for randomly generated

inputs. Although the number of iterations were not very large, due to the cost of each genetic operator among the chromosome populations, each iteration takes a considerably long time to complete, and therefore we have observed that the execution times are much higher to reach to the level of near-optimal results obtained with the approach proposed in this paper. Since GA approach uses a generic heuristic (slightly customized for the problem), it is actually not fair to compare it with our problem-specific heuristic, which is much more effective. Moreover, although GA approach has been applied to different sized inputs, significant input parameters have not been explored in its evaluation in [9] corresponding to the structure and the distributions of the weights. That is why we have compared the quality of the solutions of our heuristic proposed in this paper with that of ILP only which produces the optimal (whenever possible). This paper has the following additional contributions to [9]:

- The problem is shown to be NP-hard,
- An Integer Linear Programming (ILP) model of the problem is given,
- An iterative Linear Programming (LP) relaxation solution is developed,
- The effectiveness of the proposed iterative LP-relaxation solution is verified through extensive tests.

Iterative LP-relaxation or rounding algorithms have previously been used. A factor 2 approximation algorithm is presented in [10] for finding a minimum-cost subgraph having at least a specified number of edges in each cut. This class of problems defined in [10] includes the generalized Steiner network problem also known as the survivable network design problem. The algorithm in [10] first solves the linear relaxation of ILP formulation of the problem, and then iteratively rounds off the solution. The approach taken in [10] has been generalized and formalized in [11]. In order to exploit the full power of LP, a new technique called iterative rounding has been introduced in [11]. Iterative rounding is used in [11] to iteratively recompute the best fractional solution while maintaining the rounding of the previous phases. Although an iterative rounding based heuristic solution is developed in this paper for *MWTM*, the presence of the hierarchy constraint does not simply lend itself to the consideration of fractional values from the highest to the smallest.

The rest of the paper has been organized as follows. The next section formally introduces the problem, and proves its NP-hardness. Section 3, describes a mathematical (integer linear programming) formulation, and Section 4 presents how its relaxation to LP can be iteratively used as an effective heuristic. Section 5 describes the experiments and their results. Finally, the last section presents concluding remarks.

## 2. Problem Description and its NP-Hardness

We will now introduce *Maximum Weight Tree Matching (MWTM)* problem formally.

**Definition 1.** A tree  $T$  with  $n$  nodes rooted at a node  $r$ , and a separate set of  $m$  tasks are given. Associated with each node  $i$  in  $T$  is a real valued function  $w_{i,j}$  denoting the weight of assigning node  $i$  to task  $j$  for all  $i \in \{1..n\}$  and  $j \in \{1..m\}$ . The problem of finding an assignment of all tasks to nodes in  $T$  with the maximum total weight in such a way that the assignment between nodes and tasks forms a matching, and no node assigned to a task is allowed to have any ancestors (or descendants) which have also been assigned to a task is named *MWTM*.

It should be noted that the requirement for the weight function to be defined for all combinations of nodes and tasks in *MWTM* stems from a deliberate decision. *MWTM* is more restricted than its possible variants where some combinations of nodes and tasks can be forbidden. As *MWTM* can be reduced directly to these more general forms, NP-hardness of them would easily follow once *MWTM* is shown to be NP-hard.

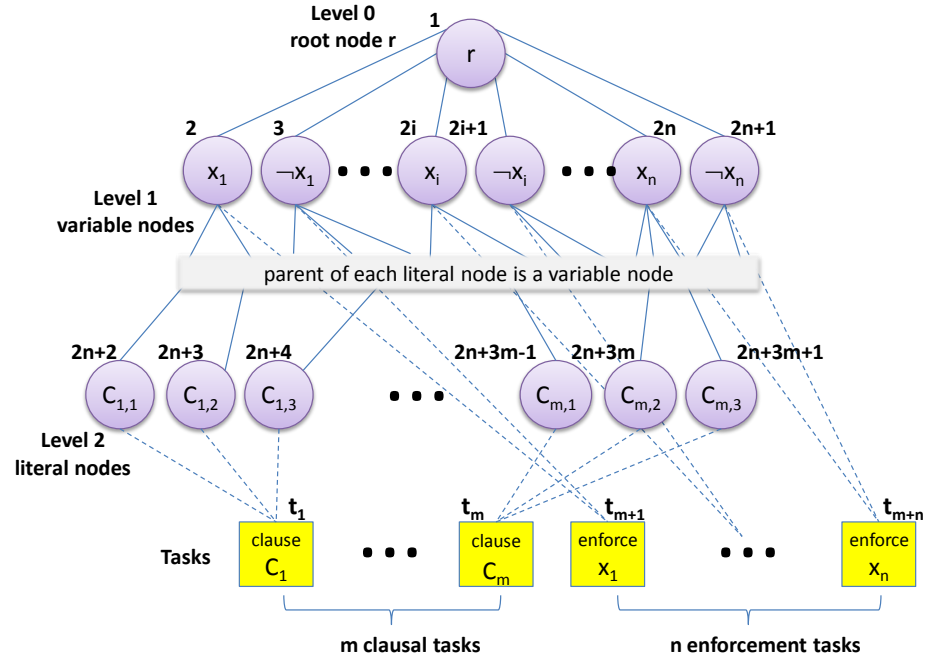


Figure 1: Transforming an E3-SAT instance to the corresponding instance of *MWTM*. The solid lines between the nodes are the tree edges while the dashed lines between the nodes and the tasks correspond to the weight function  $w_{i,j}$ .

The constraint associated with the hierarchical structure of the tree dictates that no two nodes on the same path from the root  $r$  to a leaf node in  $T$  can ever be simultaneously assigned in a solution to an instance of *MWTM*.

**Definition 2.** Two paths in a tree from the root to any two distinct nodes are said to be *independent paths* if and only if none of the two paths is a subset of the other.

In the light of this definition, the hierarchy constraint can simply be restated as the requirement that the paths from the assigned nodes to the root are all pairwise independent.

*MWTM* can be shown to be NP-hard by a polynomial time reduction from E3-SAT which is a variant of 3-satisfiability (3-SAT) problem. E3-SAT (resp. 3-SAT) is defined to be the problem of deciding whether a satisfying truth assignment is possible for the variables of a given Boolean formula in Conjunctive Normal Form (CNF) where each clause is a disjunction of exactly (resp. at most) three literals each of which is either a variable or its negation. 3-SAT is one of Karp's 21 NP-complete problems [12]. Any given instance of 3-SAT can be easily transformed to a corresponding instance of E3-SAT by introducing three new dummy variables,  $d_1$ ,  $d_2$ , and  $d_3$ . While only  $d_1$  is inserted into the clauses with one literal, both  $d_1$  and  $d_2$  are inserted into the clauses with two literals. In order to make sure in any satisfying assignment that the dummy variables can only be set to false, the conjunction of all maxterms of the dummy variables except  $d_1 + d_2 + d_3$  are finally appended to the clauses each with exactly three literals now. The NP-completeness of E3-SAT is hence confirmed.

A given instance of E3-SAT problem is transformed to a corresponding instance of *MWTM* in time polynomial in the size of the input Boolean expression. Let a given instance of E3-SAT have  $n$  variables denoted by  $x_i$  where  $i \in [1..n]$  and a 3-CNF formula  $C_1 \wedge C_2 \wedge \dots \wedge C_m$  where each  $C_i$  represented by  $C_{i,1} \vee C_{i,2} \vee C_{i,3}$  is a disjunction of three literals corresponding to either a variable or its negation. The transformation starts by introducing the root node designated by  $r$  to the initially empty tree  $T$  of the corresponding *MWTM* instance at level 0. The root node  $r$  is numbered as 1. For each variable  $x_i$ , two child nodes to root  $r$  are then created numbered  $2i$  for  $x_i$ , and  $2i + 1$  for  $\neg x_i$  corresponding to assigning *true* and *false* respectively to this variable. The parents of all such nodes are set to point to node  $r$ . As there are  $n$  distinct variables in the given E3-SAT instance, the root  $r$  of  $T$  in the corresponding *MWTM* instance becomes populated with a total of  $2n$  children at level 1 of  $T$  after this step. These are called *variable nodes* (see Figure 1). In the final step of the construction of  $T$ , for each literal  $C_{i,j}$  where  $i \in [1..m]$ , and  $j \in [1..3]$ , a node numbered  $1 + 2n + 3(i - 1) + j$  is created. The parent of such a node is set to  $2k$  if  $C_{i,j} = x_k$ , and to  $2k + 1$  otherwise if  $C_{i,j} = \neg x_k$  where  $k \in [1..n]$ . What this step achieves in effect for each node corresponding to assigning *true* to  $x_k$  or to its negation  $\neg x_k$  at level 1 is the creation of as many children at the next level 2 under the relevant variable node as there are occurrences of the corresponding variable in the clauses of the given E3-SAT instance. The tree  $T$  constructed is shown in Figure 1. While parent-child relationships are indicated by solid lines in this figure, dashed lines depict the weight function  $w_{i,j}$ . It should be noted that the variable nodes at level 1 will have as many children as there are occurrences of the corresponding literal at level 2 which

is implied by the existence of multiple edges emanating from a variable node while the nodes corresponding to literals in clauses at level 2 will have a single edge to their parent as shown in the figure. The nodes at level 2 are accordingly called *literal nodes*.

Once we obtain the tree  $T$  in  $MWTM$  instance corresponding to the given instance of E3-SAT, we also set the number of tasks to  $m + n$ . Each task  $t_i$  for  $i \in [1..m]$  corresponds to satisfying a clause  $C_i$ . We call these *clausal tasks*. Each task  $t_i$  for  $i \in [m + 1..m + n]$  among the rest of the tasks, however, are used to enforce that the corresponding variable  $x_{i-m}$  is set to either one of *true* or *false* consistently over all clauses. We call such tasks *enforcement tasks*.

Apparently, the total number of nodes in  $T$  in the corresponding instance of  $MWTM$  is given by  $1 + 2n + 3m$  where  $n$  and  $m$  are the number of variables and clauses respectively specified in the given E3-SAT instance. The number of tasks, on the other hand, is  $n + m$ . The concluding step of the transformation is to appropriately set the corresponding values  $w_{i,j}$  for all nodes  $i \in [1..1+2n+3m]$  in  $T$ , and all tasks  $j \in [1..m + n]$  as shown in Equation 1 below:

$$w_{i,j} = \begin{cases} 0, & \text{if } i = 1 \wedge j \in [1..m + n] \\ 1, & \text{if } i \in [2..2n + 1] \wedge j = m + \lfloor \frac{i}{2} \rfloor \\ 1, & \text{if } i \in [2n + 2..2n + 3m + 1] \wedge j = \lfloor \frac{i-2n-2}{3} \rfloor + 1 \\ 0, & \text{otherwise} \end{cases} \quad (1)$$

The weights of carrying out any one task by the root node are all initialized to zero. For a variable node  $i \in [2..2n + 1]$  at level 1 corresponding to  $x_{\lfloor \frac{i}{2} \rfloor}$  or  $\neg x_{\lfloor \frac{i}{2} \rfloor}$  depending on whether  $i$  is even or odd respectively, however, the weights of executing tasks are set in such a way that a consistent assignment of truth values to individual variables can be enforced. The only task whose execution by node  $i$  can have a positive contribution to the solution is therefore the corresponding enforcement task  $t_{m+\lfloor \frac{i}{2} \rfloor}$ . At level 2 are the literal nodes ranging from  $2n + 2$  to  $1 + 2n + 3m$  corresponding to the literals in the clauses of the given E3-SAT instance. Since each literal can accordingly be set to satisfy a clause in which it occurs, the weight  $w_{i,j}$  of assigning a level 2 node  $i$  representing a literal  $C_{p,q}$  to clausal task  $t_j$  corresponding to the clause  $C_p$  itself is appropriately set to 1 to reflect a feasible assignment. Therefore, the equalities  $i = 1 + 2n + 3(p - 1) + q$ , and  $j = p$  must hold. Noting that  $q$  can only take on the values 1 through 3 inclusive readily gives  $p = \lfloor \frac{i-2n-2}{3} \rfloor + 1$ , and  $q = (i - 2n - 2) \bmod 3 + 1$ . All other combinations of nodes and tasks have weight 0.

It should be pointed out that an  $MWTM$  instance so constructed would always lend itself to feasible solutions since the number of leaf nodes in  $T$  is greater than or equal to the number of tasks. This last inequality can be seen to hold by noting that  $m \geq (2n - t)/3$  where  $t \in [0..n]$  denotes the number variable nodes without any children in  $T$  based on the assumption that at least one of a variable or its negation is used in one of  $m$  clauses in the given E3-SAT instance. As  $m$ ,  $n$ , and  $t$  are all non-negative,  $m \geq (2n - t)/3 = \frac{2}{3}n - \frac{1}{3}t \geq \frac{1}{2}n - \frac{1}{3}t \geq \frac{1}{2}n - \frac{1}{2}t$  is easily obtained. Multiplying both sides of  $m \geq \frac{1}{2}n - \frac{1}{2}t$  by two, and then

adding  $m$  to both sides, we obtain  $3m \geq m + n - t$ , and then  $3m + t \geq m + n$  by rearranging. The feasibility of the corresponding *MWTM* instances obtained through the transformation described are hence confirmed.

Given the transformation described, we make the following straightforward observation to be used in a lemma to follow.

**Observation 1.** *In any solution with total weight  $n + m$  to the corresponding *MWTM* instance obtained from a given *E3-SAT* instance through the transformation described, a literal node at level 2 can be assigned to a related clausal task iff no other literal node corresponding to its negation at the same level has already been allocated.*

PROOF. In any solution with total weight  $n + m$  to the corresponding *MWTM* instance after the transformation depicted, all the enforcement tasks should have already been assigned to variable nodes at level 1 in  $T$ . Only the children at level 2 of the unassigned variable nodes at level 1 now, by the definition of *MWTM*, can be used to fulfill the tasks corresponding to the clauses which ensure a consistent assignment.  $\square$

The following lemma can now be proved easily.

**Lemma 1.** *A given *E3-SAT* instance with  $n$  variables, and  $m$  clauses is satisfiable iff the corresponding *MWTM* instance obtained through the transformation described above has a solution with total weight  $n + m$ .*

PROOF. Let us first prove the sufficiency part: If a given *E3-SAT* instance is satisfiable then there exists an assignment of truth values to all  $n$  variables which makes all  $m$  clauses evaluate to true. This, in turn, implies that at least one literal in every clause can be made true. The corresponding *MWTM* instance is then easily seen to have an optimal assignment with weight  $n + m$ : Each task corresponding to a clause in this scheme is assigned to one node at level 2 corresponding to one of the literals satisfying this clause while each enforcement task is assigned to the node at level 1 representing the negation of the literal evaluating to true in a satisfying truth assignment to the given *E3-SAT* instance. This is indeed a matching since each task is matched to a different node and no node which is a parent of an already assigned literal node is assigned to a task. The latter is guaranteed by the fact that if a literal node is assigned to a clausal task, then the node corresponding to the negation of this literal at a higher level can only be used to accomplish the respective enforcement task. The total weight is also the maximum possible as no weight value can be greater than 1.

In order to prove the necessity part, let us assume that there exists a solution with total weight  $n + m$  to the corresponding *MWTM* instance. A truth assignment for the given *E3-SAT* instance can be obtained by setting each variable  $x_i$  to true if the corresponding enforcement task  $m + i$  is assigned to node  $2i + 1$ , and to false if the assignment is to node  $2i$ . This truth assignment definitely satisfies 3-CNF expression of the given *E3-SAT* instance by Observation 1 above.  $\square$

To illustrate the idea in the reduction process, let us consider the following example.

**Example 1.** A 3-CNF formula  $(p \vee \neg q \vee \neg p) \wedge (p \vee r \vee \neg s) \wedge (q \vee r \vee s)$  with 4 variables, and 3 clauses is given. While the variables are named  $p$ ,  $q$ ,  $r$ , and  $s$ , the clauses are denoted by  $C_1 = (p \vee \neg q \vee \neg p)$ ,  $C_2 = (p \vee r \vee \neg s)$ , and  $C_3 = (q \vee r \vee s)$ . The corresponding tree structure obtained through the transformation just described is given in Figure 2, while the accompanying weights of assigning nodes to tasks are shown in Figure 3. While the nodes are numbered from 1 through 18, tasks are called  $t_{C_1}$ ,  $t_{C_2}$ , and  $t_{C_3}$  corresponding to the clausal tasks, and  $t_p$ ,  $t_q$ ,  $t_r$ , and  $t_s$  corresponding to the enforcement tasks.

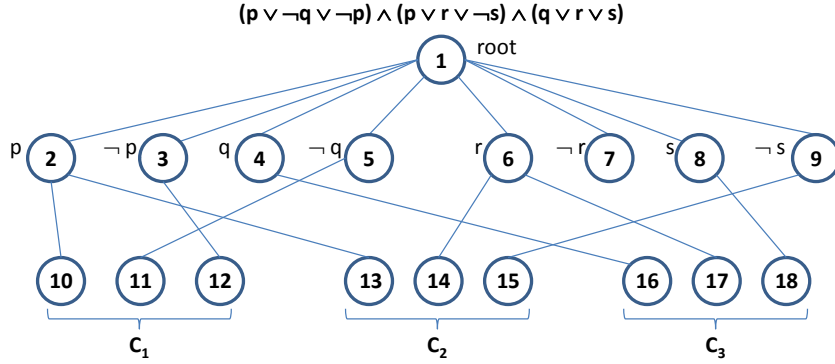


Figure 2: The tree for the corresponding *MWTM* instance obtained from the example E3-SAT instance.

The given 3-CNF Boolean expression is satisfiable if and only if *MWTM* instance identified with the corresponding tree structure given in Figure 2, and accompanying weight function depicted in Figure 3 has an assignment with a total weight of  $3 + 4 = 7$ .

It should be noted that the construction constrains the weight values for variable nodes 2 through 9 in the table to the left of Figure 3 through enforcement tasks in such a way that a node corresponding to a variable and another node corresponding to its negation can not at the same time contribute to a solution. An inspection of the weight values for literal nodes 10 through 18 in the table to the right in Figure 3 reveals similarly that only one of three such nodes can contribute to a solution through a corresponding clausal task.

If a given 3-CNF formula is satisfiable, any satisfying truth assignment induces a straightforward matching of the nodes to all the available tasks such that the variable nodes at level 1 evaluating to false with respect to the given truth assignment are all assigned to their corresponding enforcement tasks leaving only their negations for a consistent instantiation over the entire set of clauses. There are several ways to satisfy the above example formula. Let us assume that we pick an assignment as follows: Variables  $p$  and  $q$  are both assigned to *true* while  $r$  and  $s$  can be assigned randomly. If we reflect these choices on the



	$t_{C1}$	$t_{C2}$	$t_{C3}$	$t_p$	$t_q$	$t_r$	$t_s$		$t_{C1}$	$t_{C2}$	$t_{C3}$	$t_p$	$t_q$	$t_r$	$t_s$
<b>1</b>	0	0	0	0	0	0	0	<b>10</b>	1	0	0	0	0	0	0
<b>2</b>	0	0	0	1	0	0	0	<b>11</b>	1	0	0	0	0	0	0
<b>3</b>	0	0	0	1	0	0	0	<b>12</b>	1	0	0	0	0	0	0
<b>4</b>	0	0	0	0	1	0	0	<b>13</b>	0	1	0	0	0	0	0
<b>5</b>	0	0	0	0	1	0	0	<b>14</b>	0	1	0	0	0	0	0
<b>6</b>	0	0	0	0	0	1	0	<b>15</b>	0	1	0	0	0	0	0
<b>7</b>	0	0	0	0	0	1	0	<b>16</b>	0	0	1	0	0	0	0
<b>8</b>	0	0	0	0	0	0	1	<b>17</b>	0	0	1	0	0	0	0
<b>9</b>	0	0	0	0	0	0	1	<b>18</b>	0	0	1	0	0	0	0

Figure 3:  $w_{i,j}$  values for the corresponding *MWTM* instance obtained from the example E3-SAT instance.

instance of *MWTM* obtained, tasks  $t_{C1}$  and  $t_{C2}$  are assigned respectively to nodes 10 and 13 both corresponding to  $p$  while task  $t_{C3}$  is assigned to node 16 corresponding to  $q$ . Then, we can assign enforcement task  $t_p$  to node 3 corresponding to  $\neg p$ , task  $t_q$  to node 5 corresponding to  $\neg q$ . Finally, we assign task  $t_r$  to either one of the nodes 6 or 7, and task  $t_s$  to either one of the nodes 8 or 9. For the last two,  $r$  and  $s$ , the choice is not really important, since neither one of these variables has been used in satisfying the clauses.  $\square$

The transformation described in this section is certainly polynomial in the size of the given E3-SAT instance. This is easily observed by noting that the number of nodes created to form a tree in the corresponding instance of *MWTM* is  $2n + 3m + 1$ , and  $O(n + m)$  additional processing is needed in the worst case for every node as its weight to  $n + m$  tasks are all initialized to either 0 or 1 resulting in a total time proportional to  $O((n + m)^2)$ .

**Theorem 2.** *MWTM problem is NP-hard.*

PROOF. It follows easily from Lemma 1, and the fact that transformation is polynomial in the size of the given E3-SAT instance.  $\square$

MAX-E3-SAT is an optimization problem which generalizes E3-SAT in such a way that instead of a satisfying assignment, it finds an assignment satisfying the maximum number of clauses in a given 3-CNF formula. As shown in [13], it is NP-hard to approximate satisfiable MAX-E3-SAT instances to within a factor  $7/8 + \epsilon$  of the optimal for any  $\epsilon \in (0, 1]$ . It is accordingly noted at this point that we can slightly modify the illustrated transformation from E3-SAT to *MWTM* to obtain a transformation also from MAX-E3-SAT to *MWTM*. First, the weights of assigning the variable nodes to the corresponding enforcement tasks are set to  $m$  (the number of clauses). Then,  $m$  additional dummy nodes whose weights of executing any one of the tasks have all been initialized to zero are introduced as children directly to the root. It is now easily seen that any given

instance of MAX-E3-SAT denoted by  $\Pi_1$  has a solution with value  $k^*$  if and only if the corresponding instance of *MWTM* denoted by  $\Pi_2$  has a solution with a value of  $k^* + mn$ . This polynomial time reduction can also be used to establish that *MWTM* cannot have a polynomial-time approximation scheme (PTAS). Otherwise, we could use it to obtain a  $7/8 + \epsilon$  approximation algorithm for MAX-E3-SAT, and hence a contradiction. In order to see this, let us assume that *MWTM* has a  $1 - \delta$  approximation where  $\delta \in (0, 1]$ . For a given instance of MAX-E3-SAT, the corresponding instance of *MWTM* is first obtained in polynomial time using the transformation just depicted. Setting  $\delta = \frac{1}{mn}(1/8 - \epsilon)$ , the approximation algorithm for *MWTM* is run next on the transformed instance to return  $k + mn \geq (1 - \delta)(k^* + mn)$  where  $k$  and  $k^*$  are the number of clausal tasks in the approximate and optimal solutions respectively. We can then write the inequality  $(k^* + mn) - (k + mn) \leq (k^* + mn) - (1 - \delta)(k^* + mn)$ . Arranging the left and the right hand sides, we obtain  $k^* - k \leq \delta(k^* + mn)$ . For sufficiently large values of  $m$  and  $n$ , the inequality can be rewritten as  $k^* - k \leq \delta mn k^*$  which is, in turn, arranged to give  $(1 - \delta mn)k^* \leq k$ . Substituting the value for  $\delta$ ,  $(7/8 + \epsilon)k^* \leq k$  is readily obtained contradicting the fact that no such approximation is possible unless  $P = NP$ . A very trivial result can hence be stated as in the following corollary.

**Corollary 3.** *There exists no  $1 - \epsilon$  approximation algorithm for *MWTM* problem where  $\epsilon \in (0, 1]$  unless  $P = NP$ .*

### 3. ILP Formulation of *MWTM* Problem

In an instance of *MWTM*, the number of nodes organized as a tree,  $T$ , and the number of tasks are given by  $n$  and  $m$  respectively. The weight of executing each task  $j$  by a node  $i$  is also denoted by  $w_{i,j}$  where  $i \in [1..n]$  and  $j \in [1..m]$ . Let  $r$  designate the root of this tree,  $T$ . Let us denote by  $\lambda \subseteq \{1..n\}$  the leaf nodes of  $T$ . Each unique path from the root  $r$  to a leaf node  $k \in \lambda$  is represented by a set of nodes on this path which is denoted by  $\Pi_k$ . Integer Linear Programming (ILP) formulation of *MWTM* problem can thus be given as:

$$\text{maximize } \sum_{i=1}^n \sum_{j=1}^m w_{i,j} * x_{i,j} \quad (2)$$

subject to

$$\sum_{j=1}^m x_{i,j} \leq 1, \forall i \in \{1..n\} \quad (3)$$

$$\sum_{i=1}^n x_{i,j} = 1, \forall j \in \{1..m\} \quad (4)$$

$$\sum_{i \in \Pi_k} \sum_{j=1}^m x_{i,j} \leq 1, \forall k \in \lambda \quad (5)$$

$$x_{i,j} \in \{0, 1\}, \forall i \in \{1..n\} \text{ and } \forall j \in \{1..m\} \quad (6)$$

The inequality in (3) simply means a node can be assigned to at most one task. The constraint in (4) is used to enforce that every task is executed by a single node. In order to enforce that on any path leading to a leaf node, at most one node can be assigned to a task, (5) is used. Finally, (6) is there to make sure that decision variables  $x_{i,j}$  can take on the integer values 0 and 1 only. The given ILP formulation can readily be relaxed to an LP by removing the last constraint (6) which restricts  $x_{i,j}$  values to either 0 or 1.

#### 4. Bottom-Up Assignment Heuristic

In this section, a heuristic solution is developed in an effort to solve *MWTM* effectively. When ILP formulation is relaxed by removing the last constraint (6) to obtain an LP model,  $x_{i,j}$  can take on fractional values in the range  $[0, 1]$ . To cope with these fractional values in order to come up with a feasible integer solution, **BOTTOM-UP-ASSIGNMENT** (BOA) procedure given in Algorithm 1 is used.

Before giving a detailed explanation of BOA, a high level description of the heuristic can be presented as follows: First, a call is made to obtain a solution to LP relaxation of ILP formulation of *MWTM*. Then, this possibly fractional solution is converted to a feasible, partial 0-1 solution where leaf nodes with greater fractional assignments are favored. The remaining nodes and tasks that are still not allocated at this current episode, if any, form a smaller instance of *MWTM* which is simply handed over to a subsequent iteration. At this successive iteration, a new call to LP relaxation for the smaller instance is issued. This process is repeated as long as there are tasks not assigned yet. The entire heuristic hence works its way via making leaf-assignments between successive calls to LP.

BOA in Algorithm 1 assumes that the number of tasks and nodes are represented by  $m$  and  $n$  respectively. The number of tasks,  $m$ , is greater than 1 to address only the non-trivial instances of *MWTM*. It is also assumed that the root of the tree is dummy, i.e., it cannot be assigned to a task as the other nodes would be rendered useless otherwise. The input to this algorithm are a tree  $T$  with  $n$  nodes, and weights  $w_{i,j}$  for each node-task pair  $(i, j)$  of performing task  $j$  by node  $i$ .

BOA starts by initializing the set  $\alpha$  of assignments at line 1 to be empty. At line 1, both sets *tasksLeft* and *nodesLeft* used to keep track of the remaining tasks and the remaining nodes respectively are initialized. The call to LP, next at line 1, takes as parameters the original *MWTM* instance along with the assignments made so far to construct and also solve the LP formulation given by (2) through (5) of the given *MWTM* instance with respect to the set of already made assignments in  $\alpha$ . If a feasible solution exists, a 2-dimensional array  $x$  of possibly fractional values are returned by this call. The effect of the parameter  $\alpha$  is to set all  $x_{i,j}$  values to 1 in the corresponding LP formulation for all node-task pairs  $(i, j) \in \alpha$ . This definitely ensures that neither the ancestors nor the

---

**Algorithm 1:** BOTTOM-UP-ASSIGNMENT( $T, w, n, m$ )

---

**Input:**  $T$  is a tree modeling the parent-child relationships among  $n$  nodes rooted at node  $r$ ;  
 $w$  is a 2-dimensional array where  $w_{i,j}$  denotes the weight of assigning node  $i$  to task  
 $j$  for all  $i \in \{1..n\}$ , and  $j \in \{1..m\}$ . The number of tasks, given by  $m$ , satisfies  
 $m > 1$  as  $m = 1$  case is trivial to handle; The weight function is such that  $w_{r,j} = 0$   
for all  $j \in \{1..m\}$  since the root can only be assigned when there is only one task in  
the problem instance.

**Output:** A feasible assignment  $\alpha$  of  $m$  tasks to nodes in  $T$ .

```
1  $\alpha \leftarrow \emptyset$ ;  
2  $tasksLeft \leftarrow \{1..m\}$ ;  $nodesLeft \leftarrow \{1..n\}$ ;  
  // a call to LP with substitutions  $x_{i,j} \leftarrow 1 \forall (i,j) \in \alpha$   
3  $x_{i,j} \leftarrow LP(T, w, n, m, \alpha) \forall i \in \{1..n\}$  and  $j \in \{1..m\}$ ; // check for a feasible solution!  
4  $T' \leftarrow deleteNodes(T, \{(i,j) \in \alpha\})$ ; // delete all nodes assigned  
5  $\lambda \leftarrow leaves(T')$ ;  
  // leaves with a non-zero assignment are examined in decreasing order of  $x_{i,j}$  values  
6 while ( $\max_{x_{i,j} \neq 0} (i \in \lambda \cap nodesLeft)$  and ( $j \in tasksLeft$ ) can be found) do  
7    $\alpha \leftarrow \alpha \cup \{(i,j)\}$ ; // record this assignment  
8    $\lambda \leftarrow \lambda - \{i\}$ ;  
9    $tasksLeft \leftarrow tasksLeft - \{j\}$ ;  $nodesLeft \leftarrow nodesLeft - \{i\}$ ;  
10   $\Pi_i \leftarrow$  set of nodes on the path from  $i$  to  $r$  in  $T'$ ;  
  // remove all the nodes from  $i$  up to  $r$  in  $T'$  from consideration  
11  foreach  $k \in (\Pi_i - \{i\})$  do  $nodesLeft \leftarrow nodesLeft - \{k\}$ ;  
12   $T' \leftarrow deleteNodes(T', \{i\})$ ; // delete  $i$  in  $T'$   
13 end  
14 if ( $tasksLeft \neq \emptyset$ ) then  
15    $T' \leftarrow deleteNodes(T', \lambda)$ ; // to give ancestors a chance  
16    $leavesLeft \leftarrow nodesLeft \cap leaves(T')$ ;  $nodesLeftInT' \leftarrow nodesLeft \cap nodes(T')$ ;  
17   if ( $nodesLeftInT'$  has nodes with  $x_{i,j} \neq 0$ ) and ( $|leavesLeft| \geq |tasksLeft|$ ) then  
18     go to step 1;  
19   else  
20     go to step 1;  
21   end  
22 end  
23 return assignment  $\alpha$ ;
```

---

descendants of already assigned nodes in any feasible solution can have a non-zero assignment value  $x_{i,j}$  associated with them. Line 1 deletes all the nodes in  $T$  assigned so far by BOA to obtain a new tree  $T'$ . This tree  $T'$  along with  $nodesLeft$ ,  $tasksLeft$ , and the unmodified weight function  $w$  actually identify a residual *MWTM* instance obtained by reflecting the current assignments in  $\alpha$  made so far into the original instance. This is achieved simply by pruning the assigned nodes, and hence their descendants from  $T$  to obtain  $T'$  as well as keeping the set  $nodesLeft$  synchronized in the algorithm by removing the ancestors of these already assigned nodes from it to enforce the hierarchy constraint. It should be observed at this point that the most recent LP relaxation formulation at line 1 corresponds exactly to this residual *MWTM* instance as represented by the current values of the variables in  $(T', nodesLeft, tasksLeft, w)$  held at the time when line 1 gets executed. The first parameter to function  $deleteNodes()$  is immutable, and is not modified in the function. After the set  $\lambda$  is populated with a copy of the leaf nodes in  $T'$  at line 1, those leaves with a non-zero assignment in it are examined in the order of non-increasing  $x_{i,j}$  values in the while-loop between lines 1 through 1. The loop iterates as long as the maximum value assignment with a non-zero  $x_{i,j}$  between the leaf nodes and the tasks not assigned yet can be found. Among the leaves in  $\lambda$  which have not been assigned

to a task yet, only the ones not removed due to the hierarchy constraint are considered eligible as reflected by the expression  $(i \in \lambda \cap \text{nodesLeft})$  where *nodesLeft* keeps track of the remaining nodes in the original tree  $T$  which have yet been neither assigned nor left in a non-assignable state as a result of the hierarchy constraint. Existence of such an  $x_{i,j}$  value requires that an assignment between node  $i$  and task  $j$  gets recorded as illustrated at line 1. This line amounts effectively to setting  $x_{i,j}$  to 1. The following two lines 1 and 1 updates accordingly the set of leaves not considered yet and the sets of remaining tasks and nodes after the assignment just made. Since this recent assignment of node  $i$  also necessitates that the nodes on the path from node  $i$  up to the root  $r$  are removed from any further consideration for a possible assignment, such nodes computed at line 1 are accordingly deleted from *nodesLeft* at line 1, and get the right treatment. As the final statement at line 1 in the body of the while-loop, node  $i$  is pruned from  $T'$  by the respective call. The thread of control is transferred to line 1 to test whether any tasks have been left not assigned, as soon as it breaks out of the loop. If all tasks have already been assigned, the set of assignments constructed so far is returned at line 1 as the solution. Otherwise, the remaining leaf nodes are first deleted from  $T'$  at line 1 to give their ancestors a chance before a new call to LP is made. Then at line 1, a set of leaves in  $T'$  that are also in *nodesLeft* denoted by *leavesLeft*, and a set of all the nodes in  $T'$  that are also in *nodesLeft* represented by *nodesLeftInT'* are computed. Finally, at line 1, a conditional check consisting of the conjunction of two expressions is performed. The former expression evaluates to true if the nodes in  $T'$  that can still be used for further assignments have non-zero  $x_{i,j}$  values with  $j \in \text{tasksLeft}$ . The latter expression called the *feasibility invariant* is maintained throughout the entire execution of the algorithm. It basically ensures that the number of the leaf nodes still assignable are always greater than the number of the remaining tasks. If both expressions evaluate to true, execution continues by setting  $\lambda$  to the leaf nodes in the updated  $T'$  at line 1 to get ready for the subsequent execution of the while-loop once more, and otherwise a jump to line 1 occurs where a new invocation to LP occurs. All the deletions performed at line 1 in  $T'$  are effectively rolled back at line 1.

Both *deleteNodes()* and *leaves()* which are based on post-order traversal run in time proportional to the number of nodes in the tree they operate on. An implementation making an efficient evaluation at the start of every iteration of the while-loop possible employs max-heaps one for every task not assigned yet whose roots are also organized as a max-heap. Overall running time complexity of the heuristic is, however, dominated by the calls to LP at line 1. As after each call, if a feasible solution exists, BOA assigns at least one task before the next call to LP, the total number of LP calls made is equal to the number of tasks,  $m$ , in the worst case. Since LP lends itself to polynomial solutions [14, 15], BOA is easily demonstrated to be also polynomial in its worst case running time. The overhead originating from the repetitive nature of the heuristic is discussed also in the next section, and it is shown through experiments that the actual observed value for the number of times the call at line 1 to LP gets executed is almost constant on the average.

If there is a 0-1 assignment to ILP formulation of a given *MWTM* instance, its LP relaxation has certainly a fractional assignment with total weight at least that of ILP. In such a case, this fractional assignment can always be converted to a feasible 0-1 assignment by BOA in Algorithm 1. In an effort to prove this, a series of lemmas will be presented and some observations regarding the algorithm will be made.

A trivial observation could be made at this point by simply noting that a condition which ensures that the number of leaf nodes is greater than or equal to the number of tasks in a given instance of *MWTM* is both necessary and sufficient for the existence of a solution.

**Lemma 4.** *A given instance of MWTM represented by  $(T, w, m, n)$  where  $T$  is a tree, and  $w(i, j)$  is the weight of assigning node  $i$  in  $T$  to task  $j$  for all combinations of  $i \in \{1..n\}$  and  $j \in \{1..m\}$  has a solution if and only if  $|\lambda| \geq m$  where  $\lambda$  denotes the set of leaf nodes in  $T$ .*

PROOF. Let us prove the sufficiency part first. If a given *MWTM* instance  $(T, w, m, n)$  has a solution, then there exists an assignment of  $m$  nodes in  $T$  to  $m$  tasks. The hierarchy constraint in the definition of *MWTM* problem requires, in turn, that no two among these  $m$  nodes has a parent-child relationship, and they are, hence, on  $m$  mutually independent paths (see Definition 2). Therefore, the number of leaves in  $T$  denoted by  $|\lambda|$  cannot be less than the number of available independent paths from these  $m$  nodes to the root.

In order to prove the necessity part, we proceed as follows: As there are as many as  $|\lambda| \geq m$  leaf nodes, any subset of  $m$  leaves out of  $\lambda$  can be freely picked, and assigned to available tasks in a random order. Since each node that gets picked is on an independent path ensuring that the hierarchy constraint is not violated, a feasible solution is hence obtained.  $\square$

**Definition 3.** In a feasible solution to LP relaxation of a given *MWTM* instance, a node  $i$  in tree  $T$  associated with at least one non-zero  $x_{i,j}$ , and yet, not having any such descendants in  $T$  is defined to be an *effective leaf* with respect to the corresponding LP relaxation solution. The set of all such nodes is termed *effective leaves*.

In the light of this definition, the following lemma can now be stated regarding an LP relaxation formulation corresponding to a given *MWTM* instance.

**Lemma 5.** *If LP relaxation to a given MWTM instance has a solution, then the number of effective leaf nodes in the corresponding LP relaxation is greater than or equal to the number of tasks in the given problem instance.*

PROOF. If a given *MWTM* instance's LP relaxation has a solution, then the constraints (3) through (5) must hold. Therefore, we obtain by summing Equation (4) over all possible  $j$  values:

$$\sum_{j=1}^m \sum_{i=1}^n x_{i,j} = m \tag{7}$$

Let  $\lambda_e$  denote the set of effective leaves in  $T$  with respect to the particular LP relaxation solution. Since no nodes other than those in  $\lambda_e$  and their ancestors can have a non-zero  $x_{i,j}$  value associated with them, we next sum Inequality (5) over all the effective leaf nodes to obtain:

$$\sum_{k \in \lambda_e} \sum_{i \in \Pi_k} \sum_{j=1}^m x_{i,j} \leq |\lambda_e| \quad (8)$$

As the sum of individual  $x_{i,j}$  values in (7) is less than or equal to the sum in (8) over all paths leading to effective leaf nodes, we conclude:

$$m \leq |\lambda_e| \quad (9)$$

□

We can now establish the following lemma by noting that the number of effective leaves with respect to the corresponding LP relaxation solution of a given *MWTM* instance actually forms a lower bound for the number of leaf nodes in  $T$ .

**Lemma 6.** *The corresponding LP relaxation of a given MWTM instance,  $(T, w, m, n)$ , has a solution if and only if  $|\lambda| \geq m$  where  $\lambda$  denotes the set of leaf nodes in  $T$ .*

PROOF. As to the sufficiency; if LP relaxation has a solution, then, by Lemma 5,  $|\lambda_e| \geq m$  where  $\lambda_e$  is the set of effective leaves. As it is known that  $|\lambda| \geq |\lambda_e|$ ,  $|\lambda| \geq m$  follows easily.

In order to prove the necessity, on the other hand, we observe by Lemma 4 that if  $|\lambda| \geq m$ , then the given *MWTM* instance has a solution. This latter result definitely implies the existence of a solution to the corresponding LP relaxation formulation. □

**Definition 4.** An execution of BOA between successive calls to LP at line 1 is called an *iteration*.

**Theorem 7.** *BOA heuristic in Algorithm 1 returns a feasible solution whenever there exists one.*

PROOF. The algorithm will keep repeatedly performing iterations until all tasks in *tasksLeft* are exhausted, and finally an assignment is obtained. Every single iteration of BOA is launched at line 1 in an attempt to discover an assignment for a smaller residual *MWTM* instance which is identified by the values that the tree  $T'$ , *nodesLeft*, *tasksLeft*, and the unmodified weight function  $w$  have right after the statement at line 1 gets executed. The most recent LP relaxation formulation at line 1 corresponds exactly to this instance whose full recognition is achieved interestingly enough later at the next statement.

It is known by Lemma 6 that when the number of leaf nodes in *nodesLeft* (given by  $leaves(T') \cap nodesLeft$  as would be computed at line 1) is greater

than or equal to the number of remaining tasks in  $tasksLeft$  in the residual  $MWTM$  instance at the start of an iteration  $i$  before a call to LP, there must exist a feasible solution to the corresponding LP relaxation formulation at line 1. The existence of a feasible solution to the corresponding LP relaxation, in turn, implies by Lemma 5 that the number of effective leaf nodes in  $T'$  computed at line 1 is greater than or equal to the number of tasks in  $tasksLeft$ . Therefore, at least one assignment between an effective leaf and an available task will be performed in the while-loop between lines 1 through 1 in every iteration of the algorithm, and BOA will eventually terminate.

An additional observation can be made by noting that the feasibility variant cannot be violated so long as the while-loop iterates since it holds at the start, and the only type of modification allowed in the body of the loop is the assignment of an effective leaf to an available task. Such an assignment, however, removes exactly one leaf node and one task from consideration ensuring that the feasibility invariant is still maintained.

Once the control breaks out of the while-loop, either a feasible solution by BOA is returned if there are no more tasks left, or otherwise all the useless leaf nodes which survived the previous while-loop are deleted at line 1. These leaf node deletions are the only deletions that can possibly violate the feasibility invariant. Hence, once such a violation is detected at line 1, a jump at line 1 initiates the next iteration where all such deletions are effectively rolled back by reconstructing  $T'$  from scratch. In case there are no such violations, control goes once more to the while-loop. Consequently, the feasibility invariant is maintained from one iteration to the next throughout the entire execution of the algorithm.

BOA will then always find a feasible solution as long as the feasibility constraint holds at the start of the first iteration. But this is already guaranteed by Lemma 4, hence completing the proof.  $\square$

It is clearly not easy, if not impossible, to generate a feasible 0-1 assignment at once using only possibly fractional non-zero assignments obtained from the corresponding LP relaxation solution for all  $MWTM$  instances. The difficulty stems from the fact that the distribution of fractional assignment values returned by the corresponding LP relaxation solution may not easily lend itself to an integer valued assignment for all tasks without violating the hierarchy constraint. Therefore, as it is done in BOA, LP might need to be called iteratively in order to cover the tasks that have not been already assigned in the previous iterations. In an attempt to reduce the number of iterations, however, as many fractional assignment values as allowed by the feasibility invariant is checked at each iteration in BOA as to their eligibility to contribute to a feasible solution. Moreover, some zero valued assignments in previous iterations may come out non-zero in subsequent iterations leading to solutions with smaller total weights. Thus, we prefer to use the earliest LP results with non-zero assignments as much as possible.

An example is provided below for a better understanding of how BOA operates.



**Example 2.** An *MWTM* instance is depicted in Figure 4. The corresponding tree representing the hierarchical structure of an organization has 6 nodes numbered in level order as shown in the figure where the root node is denoted by 1. It is assumed in this particular example that the organization has 3 tasks to be executed not explicitly shown in the figure. The weights of executing these tasks, namely  $t_1$ ,  $t_2$ , and  $t_3$  are given in this order as a triple inside each node. In this example, the corresponding ILP formulation will produce the optimal solution with the following assignments highlighted with the corresponding weight values in red in Figure 4:

- Task  $t_1$  is assigned to node 4 with weight 6,
- Task  $t_2$  is assigned to node 5 with weight 4,
- Task  $t_3$  is assigned to node 3 with weight 8.

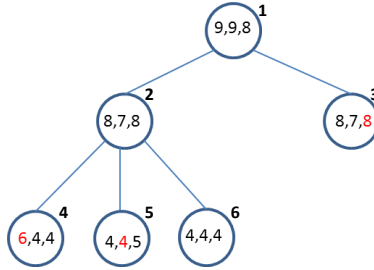


Figure 4: A sample tree structure with 6 nodes, and 3 tasks not explicitly shown. Weights of executing each task  $t_1$ ,  $t_2$ , and  $t_3$  are given in this order inside each node as triples. Red values correspond to node-task assignments obtained from ILP solution.

The next figure, Figure 5, presents a solution obtained by the corresponding LP relaxation on the same problem instance (assignments are shown in green). The solution is as follows:

- Task  $t_1$  is assigned to nodes 3 and 4 both with the same fractional value 0.5, contributing to the total weight by  $7(= 8/2 + 6/2)$ ,
- Task  $t_2$  is assigned to nodes 5 and 6 both with the same value 0.5 again, contributing to the total weight by  $4(= 4/2 + 4/2)$ ,
- Finally, Task  $t_3$  is assigned to nodes 2 and 3 both with the same value 0.5, causing this time an increase of  $8(= 8/2 + 8/2)$  in the total weight.

Since LP is allowed to make fractional assignments, the weight 19 of the solution achieved by LP is even higher than the optimal 18 found by ILP. The direct application of LP unfortunately cannot produce an integer assignment for the given *MWTM* instance. BOA in Algorithm 1, however, will work its way to a feasible solution as follows on this example:

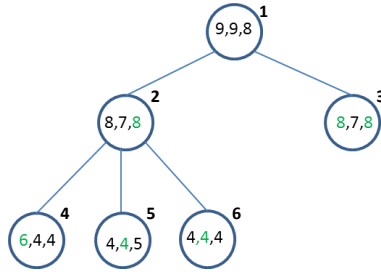


Figure 5: Green values correspond to fractional node-task assignments obtained from the corresponding LP relaxation solution where task  $t_1$  is assigned to both nodes 3 and 4, task  $t_2$  is assigned to nodes 5 and 6, and finally task  $t_3$  is assigned to nodes 2 and 3 all with the same value  $\frac{1}{2}$ .

- After a call to LP is made at line 1 in the first iteration, potentially fractional assignment values with non-zero  $x_{i,j}$  will be processed from the largest to the smallest for the leaf nodes of the tree. In this example, all the assignment values happen to be the same, namely 0.5. Such leaves may, therefore, be processed in any order. Although different heuristics may also be developed for breaking ties such as considering the depths of nodes or favoring nodes with higher  $w_{i,j}$  values, we assume for the sake of this example that the assignments with the same value are processed in increasing order of node identifiers and then in increasing order of task numbers. As a result, first, task  $t_1$  is assigned to node 3 with weight 8. Then, task  $t_2$  gets assigned to node 5 with weight 4. This assignments in the first iteration are shown in green as depicted in Figure 6. At this point, there is obviously no leaf node left with a non-zero assignment that can be used to make any further assignments, leaving task  $t_3$  hence unassigned. It should be noted that while these assignments are made, all the nodes violating the hierarchy constraint are also removed from consideration. This is evidently reflected by leaving only the nodes 4 and 6 in *nodesLeft*.
- Once it is realized that no more assignments are possible, the remaining leaves, namely 4 and 6, are deleted at line 1 from the tree  $T'$  leaving only the nodes 1 and 2 in it. As there are no nodes in the tree that are also in *nodesLeft*, a jump to line 1 initiates the second iteration of the algorithm.
- With *tasksLeft* =  $\{t_3\}$  and *nodesLeft* =  $\{4, 6\}$  at the start of the second iteration, the only remaining task is  $t_3$ , and the remaining nodes that are eligible for assignments are 4 and 6. Now a call is made to LP formulated with the assignments made in the first iteration in mind. This formulation corresponds exactly to an *MWTM* instance where the tree denoted by  $T'$  is obtained at line 1 by pruning nodes 3 and 5 from the original tree denoted by  $T$ , and the set of eligible nodes and target tasks to be matched are as dictated by the values of *nodesLeft* and *tasksLeft* at the moment. The algorithm, hence, terminates by assigning the only remaining task  $t_3$

to either node 4 or node 6 with weight 4. Figure 6 shows this assignment in the second iteration in orange. The total weight achieved by BOA is hence  $8 + 4 + 4 = 16$ , which is slightly less than the optimal ILP solution.

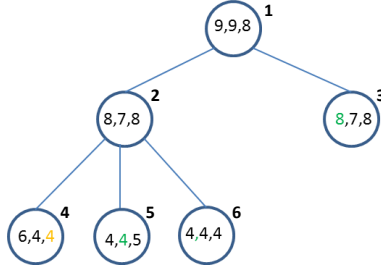


Figure 6: The assignments obtained by two LP calls. The first call generates assignments for the tasks  $t_1$  and  $t_2$  (green), and the second call generates the assignment for the task  $t_3$  (orange).

□

Another example is presented now to demonstrate the feasibility invariant at line 1 of BOA in Algorithm 1.

**Example 3.** Figure 7 is an example to an *MWTM* instance where deletions by BOA at line 1 of many leaves rooted at the same node renders this parent as an effective leaf, and the subsequent assignment of this parent to a new task runs the risk of an unanticipated decrease in the number of leaves leading to unfeasibility.

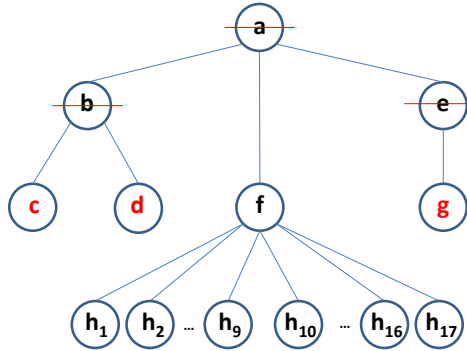


Figure 7: An instance illustrating the need for the invariant at line 1 of the algorithm.

Let us assume that there does exist a weight function  $w_{i,j}$  such that  $x_{i,j}$  values obtained by an LP call are as given in Figure 8. Then, the sum of

	$t_1$	$t_2$	$t_3$	$t_4$	$t_5$		$t_1$	$t_2$	$t_3$	$t_4$	$t_5$		$t_1$	$t_2$	$t_3$	$t_4$	$t_5$
<del>a</del>	<del>0</del>	<del>0</del>	<del>0</del>	<del>0</del>	<del>0</del>	$h_1$	0	0	0.1	0	0	$h_{10}$	0	0	0	0.1	0
<del>b</del>	<del>0.5</del>	<del>0</del>	<del>0</del>	<del>0</del>	<del>0</del>	$h_2$	0	0	0.1	0	0	$h_{11}$	0	0	0	0.1	0
c	0	0.4	0.1	0	0	$h_3$	0	0	0.1	0	0	$h_{12}$	0	0	0	0.1	0
d	0	0.4	0	0.1	0	$h_4$	0	0	0.1	0	0	$h_{13}$	0	0	0	0.1	0
<del>e</del>	<del>0.5</del>	<del>0.2</del>	<del>0</del>	<del>0</del>	<del>0.2</del>	$h_5$	0	0	0.1	0	0	$h_{14}$	0	0	0	0.1	0
f	0	0	0	0	0.8	$h_6$	0	0	0.1	0	0	$h_{15}$	0	0	0	0.1	0
g	0	0	0	0.1	0	$h_7$	0	0	0.1	0	0	$h_{16}$	0	0	0	0.1	0
						$h_8$	0	0	0.1	0	0	$h_{17}$	0	0	0	0.1	0
						$h_9$	0	0	0.1	0	0						

Figure 8: Assignments obtained by the first call to LP for the same instance in Figure 7.

assignments to each of the five tasks is 1, and each path from a leaf node to the root denoted by  $a$  has total weight less than or equal to 1. If the algorithm is run, after the first call to LP, leaf node  $d$  with the maximum assignment value 0.4 (breaking ties arbitrarily) gets assigned to task  $t_2$ . This assignment makes the next highest assignment value 0.4 for node  $c$  unusable leaving us with the only option of 0.1 as to the next largest assignment value. Assuming that node  $c$  is now picked to be assigned to task  $t_3$ , followed by the same valued assignment of node  $g$  to task  $t_4$ , all leaves  $h_1$  through  $h_{17}$  are also rendered useless. As a result, they are all deleted at line 1. This, in turn, would give way through a jump at line 1 to assigning node  $f$  to task  $t_5$  within the body of while-loop, if it were not for the invariant at line 1 in the algorithm. Such an assignment clearly would have left no nodes that can be assigned to task  $t_1$ .

The feasibility invariant at line 1 of the algorithm ensures that a sufficient number of leaves to a possible next iteration is always maintained.  $\square$

## 5. Experiments

In order to measure the performance of LP-relaxation based heuristic BOA in Algorithm 1, several experiments have been performed for varying problem parameters. The parameters employed, and their values are as follows:

1.  $\#Nodes$ : It represents the number of nodes in the tree in a given *MWTM* instance. In order to generate a variety of tree sizes, the following values are employed in the experiments: 16 (*small tree*), 32, 64, and 128 (*large tree*).
2. *Average Degree*: This parameter is defined to be the average degree of a node in the tree in a given instance of *MWTM*. It is tuned throughout the experiments to control the type of trees generated in a scale ranging from *deep* to *shallow* for fixed values of  $\#Nodes$  parameter. The values used in the experiments are 1.5 (deep tree), 2.0, and 2.5 (shallow tree).
3.  $\#Tasks / \#Nodes$ : It is defined to be the ratio of the number of the tasks to the number of the nodes in the tree associated with a given *MWTM*

instance. This parameter is used to generate a range of *MWTM* instances changing from those with a very few tasks called *sparse* to those with a large number of tasks called *dense* in proportion to the tree size. The values used are 0.125 (sparse), 0.25, and 0.5 (dense). As this ratio increases, the flexibility to use non-leaf nodes for assignments decreases.

4. *Weight Distribution*: The weight of assigning a node to a task has a value chosen from the range  $[1.. \frac{\#Nodes}{2}]$ . The following 3 weight distributions are used: i) the weights are increasing from the root to the leaves, ii) the weights are decreasing from the root to the leaves, and iii) the weights are assigned randomly without regard to the respective depths of the nodes.

For each combination of these four parameters, a total of  $4 * 3 * 3 * 3 = 108$  different test cases are formed. For each test case, 20 instances of the problem are then randomly generated, and their averages are taken in the experiments. We record the total number of LP calls made at line 1 in BOA for every instance. Corresponding to each instance, both the execution time and the solution obtained are also recorded once for the corresponding ILP formulation which gives the optimal solution, and once for BOA expected to return a suboptimal solution.

All the tests were run on a machine with a 4 GB of RAM and an Intel Core 2 Duo T9550 2.66 Ghz mobile processor. Microsoft Solver Foundation 3.0 was employed as LP/ILP solver library, and the code was developed in C#5.0.

The results of the experiments are presented through a series of seven tables in this section. These tables all share a common structure. As the topmost two rows are used to set the values for the parameters *Average Degree* and  $\#Tasks / \#Nodes$ , the leftmost two columns display the values for the parameters *Weight Distribution* and *#Nodes*. The last six tables, on the other hand, can be logically grouped into three each with two tables. While the first table in a group presents a comparison between the execution times of ILP and BOA, the second evaluates the quality of the solutions by BOA against the optimal. These three groups correspond to the three distinct values that the *Average Degree* parameter can take on, namely 2.5, 2.0, 1.5, and are presented in this order. Of the four parameters only one, namely the *Average Degree*, is fixed, and the average results are given for all combinations of the other three parameters in these groups of tables. Finally, an additional row labeled *Method* is inserted as the third from the top to allow us to specify either ILP or BOA in these tables. It should be noted that the cells at the same position in both tables in the same group correspond to the exact same combination of parameter values.

The colors yellow and green are used consistently to highlight the cells containing *NaN* and  $\infty$  respectively in all the tables. The cells in yellow marked with *NaN* in a table mean that there exists no feasible solution. For some combinations of parameters no feasible solution was possible. Especially when the instances get dense, and the trees associated with them become deep, as would be expected, it becomes more difficult to find a feasible solution satisfying the hierarchy constraint. Such configurations are characterized with high  $\#Tasks / \#Nodes$  values, and with the low values of the *Average Degree* param-

ter. The results in the tables to follow confirm this expectation. All such cases leading to infeasibility are shown in yellow. Moreover, when the weight distribution is such that it is decreasing from the root to the leaves, finding an optimal solution becomes even more difficult using ILP. Under these circumstances, the execution time for ILP grows very quickly after the number of nodes become larger than 16. We do not include these extremely large execution times in the tables, and indeed we have canceled those solutions without finding the optimal values. All such cells are displayed in green marked with an  $\infty$  symbol. The existence of feasible solutions by BOA in the corresponding cells, on the other hand, is an evidence for the existence of the optimal solutions for those cases as well. In order to verify, therefore, the quality of a solution by BOA in these situations, we make use of the corresponding possibly fractional LP relaxation solution as a potential upper bound. A quick inspection of the relevant cells reveals that the difference is very small even in these cases which definitely guarantees an even smaller distance to the actual optimal. It is hence suspected that BOA might even have achieved it.

Weight Distribution	Average Degree →	1.5			2			2.5		
	#Tasks/#Nodes →	1/8	1/4	1/2	1/8	1/4	1/2	1/8	1/4	1/2
	#Nodes ↓									
Weights are increasing from root to leaves	16	1	1	1	1	1	1	1	1	1
	32	1	1	1	1	1	1	1	1	1
	64	1	1	<i>NaN</i>	1	1	1	1	1	1
	128	1	1	<i>NaN</i>	1	1	1	1	1	1
Weights are decreasing from root to leaves	16	1	1	1	1	1	1.06	1	1	1.1
	32	1.05	1.05	<i>NaN</i>	1.1	1.15	1	1	1.2	1.05
	64	1.15	1.15	1	1.15	1.2	1.1	1.05	1.3	1.2
	128	1	1.45	<i>NaN</i>	1.2	1.3	1.35	1.2	1.35	1.45
Weights are random from root to leaves	16	1	1	1	1	1	1	1	1	1.05
	32	1	1	1	1	1	1.14	1	1.05	1.1
	64	1.1	1.2	<i>NaN</i>	1	1	1.11	1	1	1.1
	128	1.1	1.15	<i>NaN</i>	1	1.2	1.18	1	1.15	1.3

Figure 9: The average number of times LP is called at line 1 in BOA in Algorithm 1 for all test cases. The cells in yellow are marked with the symbol *NaN* to mean that there exists no feasible solution.

The table in Figure 9 displays the average number of LP invocations performed at line 1 in BOA in Algorithm 1 for each of 108 different test cases. As the table clearly reflects, the number of times the call to the corresponding LP relaxation gets executed is very close to 1. The cells marked with *NaN* all correspond to the test cases for which no feasible solutions exist as explained above.

The two tables in Figure 10 and Figure 11 display the execution times, and the solutions respectively when the parameter representing the average degree of a node in the tree is set to 2.5 which corresponds to shallow trees. There are only 3 out of 36 test cases where BOA is slightly slower in Figure 10. These

correspond to the test cases where: i)  $\#Tasks / \#Nodes = 1/8$ , *Weight Distribution* = random,  $\#Nodes = 128$ , ii)  $\#Tasks / \#Nodes = 1/4$ , *Weight Distribution* = increasing,  $\#Nodes = 64$ , and iii)  $\#Tasks / \#Nodes = 1/4$ , *Weight Distribution* = random,  $\#Nodes = 64$ . BOA, on the other hand, achieves optimal or almost optimal solutions as seen in Figure 11 for these test cases. Also an examination of the cells corresponding to these test cases in the table in Figure 9 reveals that they all have the value one.

Weight Distribution	Average Degree	2.5					
	#Tasks/#Nodes	1/8		1/4		1/2	
	Method	ILP	BOA	ILP	BOA	ILP	BOA
	#Nodes						
Weights are increasing from root to leaves	16	5.46875	0.78125	7.03125	0.78125	5.46875	2.34375
	32	8.59375	5.46875	12.5	7.03125	22.65625	17.96875
	64	27.34375	18.75	49.21875	50.78125	126.5625	111.71875
	128	135.9375	118.70118	275.78125	246.875	695.3125	624.4629
Weights are decreasing from root to leaves	16	$\infty$	2.34375	$\infty$	3.125	$\infty$	5.46875
	32	$\infty$	10.9375	$\infty$	14.84375	$\infty$	50
	64	$\infty$	81.25	$\infty$	194.53125	$\infty$	596.875
	128	$\infty$	565.625	$\infty$	2003.125	$\infty$	6573.0957
Weights are random from root to leaves	16	11.551465	4.55059	14.05167	5.0755	17.6022	7.3
	32	23.802615	8.225945	32.6041	19.927	59.732	44.605
	64	44.30491	39.1046	121.09	126.666	980.199	608.2272
	128	220.126	435.179	1699.978	1500.96	19688.367	3722.5981

Figure 10: The execution times when the average degree of a tree node parameter is set to 2.5 corresponding to shallow trees. The symbol  $\infty$  in a blue cell indicates a very large value.

Weight Distribution	Average Degree	2.5					
	#Tasks/#Nodes	1/8		1/4		1/2	
	Method	ILP	BOA	ILP	BOA	ILP	BOA
	#Nodes						
Weights are increasing from root to leaves	16	47.2	47.2	93.3	93.3	171.5	171.5
	32	243.55	243.55	465.8	465.8	881.3	881.3
	64	1175.6	1175.6	2238.15	2238.15	4409.8	4409.8
	128	5379.8	5379.8	10930.95	10930.95	21637.4	21637.4
Weights are decreasing from root to leaves	16	1792.3	1792.05	3571.65	3569.65	7100.8	7097.4
	32	3531.3	3528.35	7018.25	7009.7	13757.45	13749.65
	64	6850.6792	6824.85	13298.025	13264.9	25818.322	25791.2
	128	12295.824	12250.2	22864.054	22811.95	43858.147	43799.2
Weights are random from root to leaves	16	15.6	15.55	31.3	31.2	60	59.95
	32	63.45	63.45	125.35	125.3	246.9	246.5
	64	254.7	254.65	507	506.85	1009.05	1008.25
	128	1020.95	1020.9	2037.7	2037.2	4059.25	4058.4

Figure 11: The solutions obtained when the average degree of a tree node parameter is set to 2.5. The values in the blue cells are the estimated upper bounds obtained by the corresponding possibly fractional LP relaxation solutions.

These execution time anomalies observed to occur when BOA finds an almost optimal solution in only one iteration can therefore be explained by the overhead introduced by BOA. When BOA obtains an almost optimal solution with a single

LP call, it would be natural to also expect ILP itself to discover the optimal integer assignments quickly. As BOA has some additional computations, its running time for such cases would be slightly more than that of ILP.

Even when it takes forever to compute the optimal by ILP, the values in the corresponding blue cells in Figure 10 are all available for BOA as an indication of its running time performance. In terms of solution quality, BOA always achieves optimal solutions when *Weight Distribution* is such that it is increasing from the root to the leaves. Otherwise, the solutions obtained as shown in Figure 11 are so close to the corresponding optimal values that it is easily seen to perform within 1% of even the upper bounds obtained via the corresponding LP relaxation solution.

Weight Distribution	Average Degree	2.0					
	#Tasks/#Nodes	1/8		1/4		1/2	
	Method	ILP	BOA	ILP	BOA	ILP	BOA
#Nodes							
Weights are increasing from root to leaves	16	4.6875	1.5625	6.25	2.34375	4.3402778	5.2083333
	32	8.59375	3.90625	10.15625	8.59375	20.833333	15.625
	64	25	18.75	53.90625	42.96875	136.36364	113.63636
	128	144.38477	114.0625	290.625	255.46875	652.64423	621.39423
Weights are decreasing from root to leaves	16	$\infty$	0.78125	$\infty$	3.90625	$\infty$	5.5803571
	32	$\infty$	8.59375	$\infty$	17.96875	$\infty$	24.147727
	64	$\infty$	70.3125	$\infty$	186.71875	$\infty$	294.03409
	128	$\infty$	675	$\infty$	1981.25	$\infty$	4321.0227
Weights are random from root to leaves	16	10.50115	2.85	12.17652	4.22551	19.302387	7.3009267
	32	18.877	8.601	41.45527	21.577755	32.276818	26.730682
	64	58.682465	66.808455	142.84314	155.81968	1610.3582	540.5999
	128	272.5094	399.55046	2003.8215	1224.0517	24866.999	8480.4036

Figure 12: The execution times when the average degree of a tree node parameter is set to 2.0. The symbol  $\infty$  in a blue cell indicates a very large value.

The tables in Figure 12 and Figure 13 display the execution times, and the solutions respectively when the *Average Degree* parameter is set to 2.0. There are this time 4 out of 36 test cases where BOA turns out to be slower than the ILP solver library, and these correspond to the cells in Figure 12 characterized by: i)  $\#Tasks/\#Nodes = 1/8$ , *Weight Distribution* = random,  $\#Nodes = 64$ , ii)  $\#Tasks/\#Nodes = 1/8$ , *Weight Distribution* = random,  $\#Nodes = 128$ , iii)  $\#Tasks/\#Nodes = 1/4$ , *Weight Distribution* = random,  $\#Nodes = 64$ , and iv)  $\#Tasks/\#Nodes = 1/2$ , *Weight Distribution* = increasing,  $\#Nodes = 16$ . An inspection of the respective cells corresponding to these test cases in both Figure 9 and Figure 13 confirms once more that BOA finds solutions with optimal or almost optimal values in exactly one iteration making a single LP call. As a result, the previous analysis stating that ILP performs very fast for the instances whose LP formulations also return integer assignments still holds.

BOA always achieves optimal or very close to optimal solutions as shown in Figure 13. For example, when  $\#Tasks/\#Nodes = 1/2$  for a 128-node tree, and the weights are randomly distributed among all nodes, the ILP produces the optimal goal value as 4047.9167 and BOA heuristic generates 4036.5. This is



Weight Distribution	Average Degree	2.0					
	#Tasks/#Nodes	1/8		1/4		1/2	
	Method	ILP	BOA	ILP	BOA	ILP	BOA
	#Nodes						
Weights are increasing from root to leaves	16	55.15	55.15	119.15	119.15	209.83333	209.83333
	32	292.8	292.8	575.65	575.65	1053.75	1053.75
	64	1439.5	1439.5	2796.5	2796.5	5004.1818	5004.1818
	128	6768.4	6768.4	13728	13728	23750.462	23750.462
Weights are decreasing from root to leaves	16	1786.65	1785.6	3554.5	3551.9	7082.7857	7079.0714
	32	3510.05	3504.85	6895.35	6885.9	13679.364	13678.909
	64	6594.1538	6572.65	12790.208	12748.6	24993.333	24975.909
	128	11110.642	11051.7	20462.263	20355.1	36935.409	36903.909
Weights are random from root to leaves	16	15.55	15.55	31.35	31.3	59.533333	59.533333
	32	62.85	62.85	125.2	125.05	246.09091	246.09091
	64	254	253.9	506.05	505.75	998.0625	997.125
	128	1020.2	1020.15	2037.3	2036.85	4047.9167	4036.5

Figure 13: The solutions obtained when the average degree of a tree node parameter is set to 2.0. The values in the blue cells are the estimated upper bounds obtained by the corresponding possibly fractional LP relaxation solutions.

one of the cases with the largest difference between the optimal solution and our heuristic solution. Even in this case, the difference between the two solutions is much less than 1%. For some cases where we have used LP relaxation solutions as upper bounds instead of ILP, the differences are slightly higher. For example, when  $\#Tasks/\#Nodes = 1/4$  for a 128-node tree, and the weights are decreasing from the root to the leaves, the upper bound to the optimal is 20462.263, and BOA achieves 20355.1. Even for this upper bound the difference is very small. Potentially, BOA might even have the same solution as the actual optimal, or else would have definitely achieved a closer value to the actual optimal.

Weight Distribution	Average Degree	1.5					
	#Tasks/#Nodes	1/8		1/4		1/2	
	Method	ILP	BOA	ILP	BOA	ILP	BOA
	#Nodes						
Weights are increasing from root to leaves	16	3.90625	0.78125	3.125	4.6875	21.484375	5.859375
	32	4.6875	3.90625	11.71875	8.59375	46.875	15.625
	64	24.21875	16.40625	43.75	45.3125	NaN	NaN
	128	150.78125	120.3125	298.4375	245.3125	NaN	NaN
Weights are decreasing from root to leaves	16	∞	3.125	∞	5.7565789	∞	3.125
	32	∞	7.03125	∞	17.96875	NaN	NaN
	64	∞	39.0625	∞	108.59375	∞	109.375
	128	∞	560.15625	∞	1403.125	NaN	NaN
Weights are random from root to leaves	16	11.0514	3.4	15.2019	5.7506	13.12665	8.501
	32	17.8772	8.8261	50.5064	38.02982	51.006	23.003
	64	87.93	89.761395	619.223	152.217	NaN	NaN
	128	467.584	527.416	2981.5173	1508.0479	NaN	NaN

Figure 14: The execution times when the average degree of a tree node parameter is set to 1.5 corresponding to deep trees. While the cells in yellow marked with the symbol *NaN* represent the parameter combinations for which there are no feasible solutions, the symbol  $\infty$  in a blue cell indicates a very large value.

Figure 14 and Figure 15 display the execution times, and the solutions respectively when the parameter representing the average degree of a tree node is set to 1.5 which corresponds to deep trees. In 4 out of the 36 test cases presented in Figure 14, BOA executes longer in figuring out a solution. The first two of these correspond to the cases where the parameter  $\#Nodes$  is set to either 64 or 128 when  $\#Tasks/\#Nodes = 1/8$  and *Weight Distribution* is random. The cells corresponding to these two test cases in Figure 9 have both the value 1.1. Furthermore, it is seen from the corresponding cells in Figure 15 that BOA finds solutions very close to optimal. The last two test cases correspond, however, to the combinations of parameters when  $\#Nodes$  is set to either 16 or 64 when  $\#Tasks/\#Nodes = 1/4$  and *Weight Distribution* is such that it is increasing from the root to the leaves. A quick inspection of the corresponding cells for the last two test cases in the corresponding tables reveals that BOA found the optimal solutions after a single LP invocation. So the prior justification is still valid.

Weight Distribution	Average Degree	1.5					
	#Tasks/#Nodes	1/8		1/4		1/2	
	Method	ILP	BOA	ILP	BOA	ILP	BOA
#Nodes							
Weights are increasing from root to leaves	16	71.5	71.5	154.2	154.2	212.875	212.875
	32	415.3	415.3	698.95	698.95	510	510
	64	2035.35	2035.35	3861.45	3861.45	NaN	NaN
	128	9762.8	9762.8	19354.3	19354.3	NaN	NaN
Weights are decreasing from root to leaves	16	1786.3	1784.4	3545.4211	3542.6316	7073.2	7071.8
	32	3444.8	3431.6	6847.4	6833.4	NaN	NaN
	64	6297.7083	6273.2	11957.885	11934.85	25421.6	25403
	128	9810.1167	9734.15	15903.15	15859.6	NaN	NaN
Weights are random from root to leaves	16	15.55	15.55	30.3	30.2	58.5	58.5
	32	63.25	63.2	125.1	124.6	242	242
	64	253.7	253.65	503.45	502.95	NaN	NaN
	128	1019.6	1019.35	2032.4	2030.85	NaN	NaN

Figure 15: The solutions obtained when the average degree of a tree node parameter is set to 1.5. While the cells in yellow marked with the symbol *NaN* represent the parameter combinations for which there are no feasible solutions, the values in the blue cells are the estimated upper bounds obtained by the corresponding possibly fractional LP relaxation solutions.

The results of the experiments show that for all cases BOA generates goal values very close to the optimal obtained by ILP. The results are either exactly the same, or there is a very small difference. Besides, in the latter case, the distance to the optimal is always much less than 1%. Moreover with *Weight Distribution* increasing from the root to the leaves, BOA always finds optimal solutions.

When the parameter *Weight Distribution* is such that it decreasing from the root to the leaves, it takes forever to compute the optimal by ILP as shown by the corresponding cells marked  $\infty$  throughout the tables. Under the same setting, BOA, on the other hand, returns in polynomial time almost optimal solutions that are within 1% of even the upper bounds obtained via the corresponding LP relaxation solution.

In only 11 out of a total of 108 different test cases, ILP runs faster than ILP. All 11 of these execution time anomalies are seen occur when BOA discovers an almost optimal solution after at most 1 or 1.1 LP calls on the average. These test cases are therefore thought to correspond most probably to the instances that can be solved efficiently by ILP. In such a case ILP can essentially find a solution by making only a very few LP relaxation calls via a branch and bound algorithm. It is then easily anticipated that the additional overhead posed by BOA leaves it behind ILP.

## 6. Conclusion

In this paper we have introduced a new version of the assignment problem, called as *MWTM* problem. In *MWTM*, as is the case with the standard assignment problem, a one-to-one assignment is sought between a set of tasks and a set of agents (nodes) to maximize the total profit (weight) value. Moreover, there is an additional constraint in *MWTM* preventing some combinations of the assignments. Since agents are organized in a tree structure representing hierarchical (agent - sub-agent) relationships, when an agent is assigned to a task, none of its sub-agents or super-agents can be assigned to any other task. This problem is shown to be NP-hard. Therefore, we proposed an iterative LP-relaxation solution to it. Through experiments we have shown that our heuristic solution is very effective, and produces either the optimal solution, or a solution very close to the optimal in a very reasonable time performing only a few iterations. In most cases the solution is achieved within a single iteration.

## References

- [1] R. E. Burkard, M. Dell'Amico, S. Martello, Assignment Problems, SIAM, 2009.
- [2] R. Cohen, L. Katzir, D. Raz, An efficient approximation for the generalized assignment problem, Information Processing Letters 100 (4) (2006) 162–166.
- [3] L. Fleischer, M. X. Goemans, V. S. Mirrokni, M. Sviridenko, Tight approximation algorithms for maximum general assignment problems, in: SODA'06, ACM, 2006, pp. 611–620.
- [4] S. Geetha, K. P. K. Nair, A variation of the assignment problem, European Journal of Operational Research 68 (3) (1993) 422–426.
- [5] R. D. Armstrong, Z. Jin, On solving a variation of the assignment problem, European Journal of Operational Research 87 (1) (1995) 142–147.
- [6] M. Gulek, I. H. Toroslu, A dynamic programming algorithm for tree-like weighted set packing problem, Information Sciences 180 (20) (2010) 3974–3979.

- [7] F. Wang, S. Zhou, N. Shi, Group-to-group reviewer assignment problem, *Computers & Operations Research* 40 (5) (2013) 1351–1362.
- [8] T. Shima, S. J. Rasmussen, A. G. Sparks, K. M. Passino, Multiple task assignments for cooperating uninhabited aerial vehicles using genetic algorithms, *Computers & Operations Research* 33 (11) (2006) 3252–3269.
- [9] M. Gulek, I. H. Toroslu, A genetic algorithm for maximum-weighted tree matching problem, *Applied Soft Computing* 10 (4) (2010) 1127–1131.
- [10] K. Jain, A factor 2 approximation algorithm for the generalized steiner network problem, *Combinatorica* 21 (1) (2001) 39–60.
- [11] K. Jain, Enhancing techniques in LP based approximation algorithms, Ph.D. thesis, Georgia Institute of Technology (August 2000).
- [12] R. M. Karp, Reducibility among combinatorial problems, in: R. E. Miller, J. W. Thatcher (Eds.), *Complexity of Computer Computations*, The IBM Research Symposia Series, Plenum Press, New York, 1972, pp. 85–103.
- [13] J. Håstad, Some optimal inapproximability results, *Journal of ACM* 48 (4) (2001) 798–859.
- [14] L. Khachiyan, A polynomial algorithm in linear programming, *Soviet Math. Dokl.* 20 (1) (1979) 191–194.
- [15] N. Karmarkar, A new polynomial-time algorithm for linear programming, *Combinatorica* 4 (4) (1984) 373–395.