

Edge Finding Filtering Algorithm for Discrete Cumulative Resources in $\mathcal{O}(kn \log n)$

Petr Vilím

ILOG S.A. an IBM Company, 9, rue de Verdun, BP 85
F-94253 Gentilly Cedex, France
petr.vilim@cz.ibm.com

Abstract. This paper presents new Edge Finding algorithm for discrete cumulative resources, i.e. resources which can process several activities simultaneously up to some maximal capacity C . The algorithm has better time complexity than the current version of this algorithm: $\mathcal{O}(kn \log n)$ versus $\mathcal{O}(kn^2)$ where n is number of activities on the resource and k is number of distinct capacity demands. Moreover the new algorithm is slightly stronger and it is able to handle optional activities. The algorithm is based on the Θ -tree – a binary tree data structure which already proved to be very useful in filtering algorithms for unary resource constraints.

Note: This version of the paper differs from the original: it fixes several mistakes (mostly typos). All fixes are marked by footnotes and they include author of the fix.

Version: October 28th 2010

1 Introduction

Nowadays, constraint based scheduling engines like IBM ILOG CP-Optimizer [1] allow to describe and solve very complex scheduling problems involving a variety of different constraints. This paper describes a filtering algorithm called Edge Finding for one of them – for discrete cumulative resource constraint.

Let us demonstrate the problem on a simple example on Figure 1. Note that this example may be just a small part of much more complex problem.




Fig. 1. An example: est_D can be updated from 0 to 4.

There are three equivalent workers (a resource with capacity $C = 3$) who must perform four different activities $T = \{A, B, C, D\}$. Activity A requires all three workers ($c_A = 3$) for one hour ($p_A = 1$), activity B requires only one worker ($c_B = 1$) but for 3 hours ($p_B = 3$) and remaining two activities C and D require two workers each ($c_C = c_D = 2$), activity C for 2 hours ($p_C = 2$) and activity D for 3 hours ($p_D = 3$). Moreover the earliest possible starting time of activities A and D is zero ($est_A = est_D = 0$), for activities B and C it is 2 ($est_B = est_C = 2$). Latest possible completion time (deadline) for activities A , B and C is 5 ($lct_A = lct_B = lct_C = 5$), activity D does not have a deadline ($lct_D = \infty$).

Looking closely to the problem we can see that there is no way for D to start before 4. Therefore we can update est_D from 0 to 4 and this way limit the search space of the problem. The rest of this paper describes the algorithm (called Edge Finding) which performs such an update.

Note that there are also filtering algorithms for discrete cumulative resource other than Edge Finding. For example Timetable propagation [2], Not-First/Not-Last [3], Energetic Reasoning [2], Max Energy propagation [4] or Extended Edge Finding [5]. However Timetable and Edge Finding are the ones which are used most of the time.

Let us quickly review existing Edge Finding algorithms for discrete cumulative resources. To the author’s knowledge the current state-of-the-art algorithm can be found in [5]. In this paper Luc Mercier and Pascal Van Hentenryck proved that the original cumulative Edge Finding algorithm with time complexity $\mathcal{O}(n^2)$ in [2] is incomplete, and therefore they designed new (complete) algorithm with time complexity $\mathcal{O}(kn^2)$.

This paper further improves the algorithm [5] in several aspects:

- Θ -tree data structure improves time complexity from $\mathcal{O}(kn^2)$ to $\mathcal{O}(kn \log n)$.
- Better usage of relation “ends before end” makes the filtering a little bit stronger. There are cases when the new algorithm propagates while the old one does not (Section 6.2).
- The algorithm can be easily adapted to handle optional activities. Although propagation of optional activities can be further improved [6], it is already pretty strong. We will show how to handle optional activities at the end of the paper (Section 9).
- The algorithm is based on modified Edge Finding rules which are more suitable for propagation. We provide a proof that the new rules are not weaker than the original ones (Section 8).

Like algorithm [5] the new algorithm has two phases:

Detection phase tries to discover necessary relative positions of activities on the resource. The result of this phase is a partial knowledge of a relation “ends before end” (see later), which will be used in the next phase. For the example on Figure 1 the algorithm in this phase detects that activity D must end after the end of $\{A, B, C\}$.

Adjustment phase utilizes results of the previous phase to improve temporal bounds of activities – earliest start times and latest completion times.

In comparison with algorithm [5] the time complexities of both phases are improved. For detection phase from $\mathcal{O}(n^2)$ to $\mathcal{O}(n \log n)$, for adjustment phase from $\mathcal{O}(kn^2)$ to $\mathcal{O}(kn \log n)$. For simplicity we will describe each phase separately.

We present only the algorithm to update earliest start times (not latest completion times) because the algorithm for update of latest completion times is symmetrical.

Note also that Edge Finding algorithm is not idempotent and therefore it is usually repeated until a fixpoint is reached.

2 Basic Notation

Let us formalize the notation already used in the introduction. The input of the algorithm is a discrete cumulative resource with capacity $C \in \mathbb{N}^+$ and a set of activities T ($|T| = n$) which must be processed by the resource. Each activity $i \in T$ is characterized by the following attributes:

- earliest possible start time $est_i \in \mathbb{N}$,
- latest possible completion time $lct_i \in \mathbb{N}$,
- processing time (duration) $p_i \in \mathbb{N}$ – a constant,
- required capacity $c_i \in \mathbb{N}$ – a constant $c_i \leq C$.

Activities are not preemptive, that is, if an activity i starts at time t it must continue without interruption until time $t + p_i$ where it ends. During the whole processing from t to $t + p_i$ it requires capacity c_i from the resource. At any time, the total capacity required from the resource cannot exceed the maximum capacity C . We define k to be number of distinct capacity demands $k = |\{c_l, l \in T\}|$.

Values est_i and lct_i can change – they can be updated by other filtering algorithms or by the Edge Finding algorithm itself. Therefore the input of the Edge Finding algorithm are current bounds est_i and lct_i , the output are new (updated) bounds.

Note that at any time the following inequality must hold for every activity i :

$$est_i + p_i \leq lct_i$$

If it does not hold then the problem does not have any solution and the propagation ends (a *fail*).

From the processing time and required capacity we can compute an energy of an activity i :

$$e_i = c_i p_i$$

The energy corresponds to the area occupied by the activity on the resource when depicted like on Figure 1.

The notation for earliest start time, latest completion time and energy can be easily extended for a set of activities $\Omega \subseteq T$:

$$\begin{aligned} \text{est}_\Omega &= \min \{\text{est}_i, i \in \Omega\} \\ \text{lct}_\Omega &= \max \{\text{lct}_i, i \in \Omega\} \\ e_\Omega &= \sum_{i \in \Omega} e_i \end{aligned}$$

3 Earliest Completion Time, Energy Envelope

A critical role in the algorithm plays a computation of possible earliest completion time of a set of intervals $\Theta \subseteq T$. This computation was already described in detail in [4], therefore here we only quickly repeat the main idea. We defined a lower bound $\text{Ect}(\Theta)$ of earliest completion time of a set of activities $\Theta \subseteq T$ as:

$$\text{Ect}(\Theta) = \left\lceil \frac{\text{Env}(\Theta)}{C} \right\rceil$$

where $\text{Env}(\Theta)$ is so-called *energy envelope* of set Θ :

$$\text{Env}(\Theta) = \max_{\Omega \subseteq \Theta} \{C \text{est}_\Omega + e_\Omega\} \quad (1)$$

3.1 Cumulative Θ -tree

The paper [4] also describes how to compute $\text{Env}(\Theta)$. The idea is to organize set Θ in a balanced binary tree which we call cumulative Θ -tree. Activities are represented by leaf nodes and sorted by est_i from left to right. Each node v of the tree holds the following values:

$$e_v = e_{\text{Leaves}(v)} \quad (2)$$

$$\text{Env}_v = \text{Env}(\text{Leaves}(v)) \quad (3)$$

Where $\text{Leaves}(v)$ is a set of all activities represented by leaves of the subtree rooted in v .

Figure 2 shows a Θ -tree for $\Theta = \{A, B, C, D\}$ from Figure 1. Notice that the energy envelope of the represented set Θ is equivalent to the value Env of the root node. We can conclude that $\text{Ect}(\{A, B, C, D\}) = \lceil 16/3 \rceil = 6$.

For a leaf node v representing an activity $i \in T$ the values in the tree are set to:

$$e_v = e_i \quad \text{Env}_v = \text{Env}(\{i\})$$

For internal nodes v these values can be computed recursively from their children nodes $\text{left}(v)$ and $\text{right}(v)$ as shown in the following proposition.

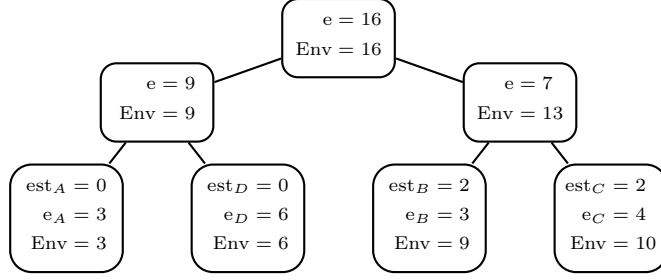


Fig. 2. An example of a Θ -tree for $\Theta = \{A, B, C, D\}$ from Figure 1.

Proposition 1. For an internal node v , values e_v and Env_v can be computed by the following recursive formulas:

$$e_v = e_{\text{left}(v)} + e_{\text{right}(v)} \quad (4)$$

$$\text{Env}_v = \max \{ \text{Env}_{\text{left}(v)} + e_{\text{right}(v)}, \text{Env}_{\text{right}(v)} \} \quad (5)$$

Proof. See [4]. □

Thanks to formulas (4) and (5), computation of values e_v and Env_v can be integrated within standard operations with balanced binary trees without changing their usual time complexities.

4 Relation “Ends before End”

Before going into details about Edge Finding, let us introduce a notion of *ends before end*. We say that an activity j ends before the end of an activity i (denoted by $j \prec i$) if in all solutions $\text{end}(j) \leq \text{end}(i)$. The goal of the Edge Finding algorithm is to discover as much of the relation \prec as possible and use it to update temporal bounds. Until a solution is found the relation “ends before end” is known only partially. Therefore in the following we will use the notation $j \prec i$ in the sense “it is known that in all solutions j ends before the end of i ”.

The notation for “ends before end” can be extended also to sets of activities:

$$\forall \Omega \subset T, \forall i \in T \setminus \Omega : \quad \Omega \prec i \Leftrightarrow (\forall j \in \Omega : j \prec i)$$

5 Edge Finding: Detection Rule

Let us start by definition of a *left cut of T by activity j* :

$$\text{LCut}(T, j) = \{l, l \in T \ \& \ \text{lct}_l \leq \text{lct}_j\}$$

To detect the relation \prec we will use the following rule:

$$\forall j \in T, i \in T \setminus \text{LCut}(T, j) :$$

$$(\text{Ect}(\text{LCut}(T, j) \cup \{i\}) > \text{lct}_j \Rightarrow \text{LCut}(T, j) \prec i)$$

The idea of this rule follows. The set $\text{LCut}(T, j)$ must be processed before lct_j . So if there is not enough time to process i together with $\text{LCut}(T, j)$ then the activity i must end after $\text{LCut}(T, j)$. Note that this rule is different from the original Edge Finding rule. We will show that this new rule is not weaker later in Section 8.

The rule above can be rewritten using energy envelope into a form more suitable for the algorithm:

$$\forall j \in T, i \in T \setminus \text{LCut}(T, j) : \\ (\text{Env}(\text{LCut}(T, j) \cup \{i\}) > C \text{lct}_j \Rightarrow \text{LCut}(T, j) \prec i) \quad (\text{EF1})$$

Our goal is to discover as much of the relation \prec as possible. Therefore for each activity $i \in T$ we are looking for an activity $j \in T$ with maximal¹ lct_j such that $\text{LCut}(T, j) \prec i$ can be detected by the rule (EF1). This is the task for the first part of the algorithm.

6 Detection Algorithm

Notice that once we prove by (EF1) that $\text{LCut}(T, j) \prec i$ then it is pointless to evaluate the rule (EF1) for the same activity i and any $j' \in T$ such that $\text{lct}_{j'} \leq \text{lct}_j$ because it cannot bring any new information ($\text{LCut}(T, j') \subseteq \text{LCut}(T, j)$).

The algorithm is similar to the Edge Finding algorithm for unary resource [7]. We iterate over activities j in non-increasing order by lct_j and we maintain a set $\Lambda \subseteq T \setminus \text{LCut}(T, j)$ of all activities i for which we still did not find a set which must end before end of i . In each step of the algorithm we check whether there is some activity $i \in \Lambda$ such that the rule (EF1) proves that $\text{LCut}(T, j) \prec i$. In other words, we test whether the following inequality holds:

$$\max_{i \in \Lambda} \{\text{Env}(\text{LCut}(T, j) \cup \{i\})\} > C \text{lct}_j$$

- If it holds then we find the responsible activity $i \in \Lambda$ and conclude that $\text{LCut}(T, j) \prec i$. Therefore we can remove i from Λ .
- If it does not hold then we move activity j into Λ and continue by next activity j (because there is no activity i such that $\text{LCut}(j) \prec i$ can be proved by (EF1)).

To formalize the algorithm let us define:

$$\text{Env}(\Theta, \Lambda) = \max_{i \in \Lambda} \{\text{Env}(\Theta \cup \{i\})\}$$

Although we did not show yet how to compute $\text{Env}(\Theta, \Lambda)$ we can already present the resulting Algorithm 1.1. The result of the computation is the array `prec` which has the following meaning:

$$\forall i \in T : \{l, l \in T \ \& \ \text{lct}_l \leq \text{prec}[i]\} \prec i$$

¹ Maximality of lct_j assures that for any other $j' \in T$ satisfying $\text{LCut}(T, j') \prec i$ by (EF1) we have $\text{LCut}(T, j') \subseteq \text{LCut}(T, j)$.

Algorithm 1.1. Edge Finding: Detection

```

1  for  $i \in T$  do
2     $\text{prec}[i] := -\infty$ ;
3   $\Theta := T$ ;
4   $\Lambda := \emptyset$ ;
5  for  $j \in T$  in non-increasing order of  $\text{lct}_j$  do begin
6    while  $\text{Env}(\Theta, \Lambda) > C \text{lct}_j$  do begin
7       $i :=$  activity from  $\Lambda$  responsible for  $\text{Env}(\Theta, \Lambda)$ ;
8       $\text{prec}[i] := \text{lct}_j$ ; // means:  $\text{LCut}(T, j) \leq i$ 
9       $\Lambda := \Lambda \setminus \{i\}$ ;
10   end;
11    $\Theta := \Theta \setminus \{j\}$ ;
12    $\Lambda := \Lambda \cup \{j\}$ ;
13 end;

```

In the algorithm, $\Theta = \text{LCut}(T, j)$ unless there are duplicities in lct_j (the algorithm is correct even with such duplicities). In the following we will concentrate on the computation of $\text{Env}(\Theta, \Lambda)$ and the proof that the algorithm 1.1 has time complexity $\mathcal{O}(n \log n)$.

6.1 Computation of $\text{Env}(\Theta, \Lambda)$

The idea is to extend cumulative Θ -tree into Θ - Λ -tree in a similar way it was done for unary resource in [7]. The cumulative Θ - Λ -tree is a balanced binary tree where each leaf represents one activity from the set Θ or Λ . Leaves are sorted from left to right according to est_i . Each node of the tree holds the following values:

$$\begin{aligned}
 e_v &= e_{\text{Leaves}(v) \cap \Theta} \\
 e_v^{\Lambda} &= e_{\text{Leaves}(v) \cap \Theta} + \max_{i \in \text{Leaves}(v) \cap \Lambda} \{e_i\} \\
 \text{Env}_v &= \text{Env}(\text{Leaves}(v) \cap \Theta) \\
 \text{Env}_v^{\Lambda} &= \text{Env}(\text{Leaves}(v) \cap \Theta, \text{Leaves}(v) \cap \Lambda)
 \end{aligned}$$

Notice that $\text{Env}(\Theta, \Lambda)$ is equivalent to Env_v^{Λ} in the root node. For an example of the cumulative Θ - Λ -tree see Figure 3.

For a leaf node v an activity $i \in \Theta \cup \Lambda$ these values are set to:

$$\begin{aligned}
 e_v &= \begin{cases} e_i & \text{if } i \in \Theta \\ 0 & \text{if } i \in \Lambda \end{cases} & e_v^{\Lambda} &= \begin{cases} -\infty & \text{if } i \in \Theta \\ e_i & \text{if } i \in \Lambda \end{cases} \\
 \text{Env}_v &= \begin{cases} C \text{est}_i + e_i & \text{if } i \in \Theta \\ -\infty & \text{if } i \in \Lambda \end{cases} & \text{Env}_v^{\Lambda} &= \begin{cases} -\infty & \text{if } i \in \Theta \\ C \text{est}_i + e_i & \text{if } i \in \Lambda \end{cases}
 \end{aligned}$$

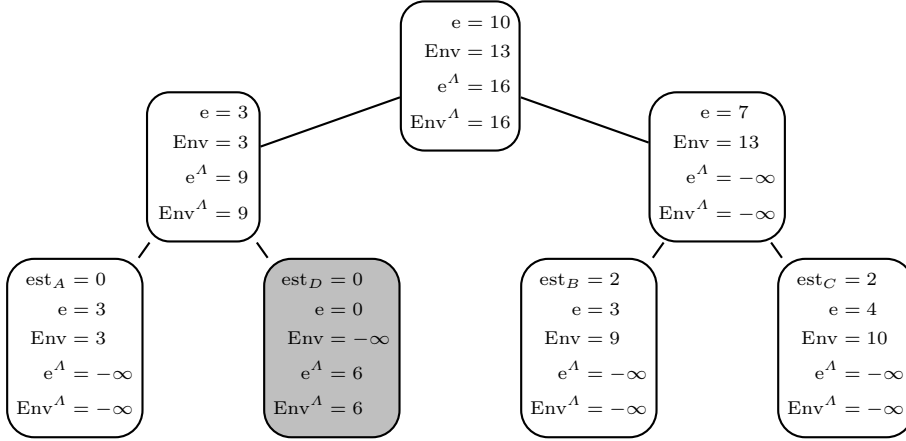


Fig. 3. An example of a Θ - A -tree for $\Theta = \text{LCut}(T, A) = \{A, B, C\}$ and $A = \{D\}$ from Figure 1. We see that $\text{Env}(\Theta, A) = 16$ which is more than $\text{C}lct_A = 15$ and therefore $\{A, B, C\} \prec D$.

For internal nodes v these values are computed recursively from their children nodes $\text{left}(v)$ and $\text{right}(v)$:

Proposition 2. For an internal node v values e_v , e_v^A , Env_v and Env_v^A can be computed by the following formulas:

$$e_v = e_{\text{left}(v)} + e_{\text{right}(v)} \quad (6)$$

$$e_v^A = \max \left\{ e_{\text{left}(v)}^A + e_{\text{right}(v)}, e_{\text{left}(v)} + e_{\text{right}(v)}^A \right\} \quad (7)$$

$$\text{Env}_v = \max \left\{ \text{Env}_{\text{left}(v)} + e_{\text{right}(v)}, \text{Env}_{\text{right}(v)} \right\} \quad (8)$$

$$\text{Env}_v^A = \max \left\{ \text{Env}_{\text{left}(v)}^A + e_{\text{right}(v)}, \text{Env}_{\text{left}(v)} + e_{\text{right}(v)}^A, \text{Env}_{\text{right}(v)}^A \right\} \quad (9)$$

Proof. First notice that formulas (6) and (8) are the same as formulas (4) and (5) in Proposition 1. Addition of new leaves representing A into the tree cannot invalidate these formulas because for these leaves v we have $e_v = 0$ and $\text{Env}_v = -\infty$. Therefore formulas (6) and (8) hold by Proposition 1.

Formula (7) is simple to prove. It is enough to realize that the difference between computation of e_v by (6) and computation of e_v^A is that it is allowed to use one of the activities $i \in A$. This activity i can be either in the left subtree of v (and in this case we can use $e_{\text{left}(v)}^A$ instead of $e_{\text{left}(v)}$) or in the right subtree of v (and we can use $e_{\text{right}(v)}^A$ instead of $e_{\text{right}(v)}$). Putting this together we transform formula (6) into (7).

It remains to prove formula (9). Again the difference between computation of Env_v and Env_v^A is that it is allowed to use one of the activities $i \in A$. This activity can be either in the left subtree of v (and in this case we can use $\text{Env}_{\text{left}(v)}^A$ instead of $\text{Env}_{\text{left}(v)}$) or in the right subtree of v (and we can use $\text{Env}_{\text{right}(v)}^A$ instead of




Fig. 4. An example: $\{M, N\} \prec O$ but the rule (EF1) is not able to detect it.

$\text{Env}_{\text{right}(v)}^A$ or $e_{\text{right}(v)}^A$ instead of $e_{\text{right}(v)}$ but not both). This way we transform formula (8) into (9). \square

Thanks to these recursive formulas it is possible to recompute internal values within standard operations with balanced binary trees without changing their time complexity. Therefore lines 9, 11 and 12 of Algorithm 1.1 has time complexity $\mathcal{O}(\log n)$ and line 6 has time complexity $\mathcal{O}(1)$. To prove that time complexity of the whole Algorithm 1.1 is $\mathcal{O}(n \log n)$ it remains to show that time complexity of line 7 is also $\mathcal{O}(\log n)$.

The activity $i \in \Lambda$ responsible for $\text{Env}(\Theta, \Lambda)$ can be found by following a path from the root of the tree to the responsible leaf. In each internal node we can recognize in $\mathcal{O}(1)$ whether the responsible activity is in the left or right subtree by analyzing which part of the formulas (9) or (7) was used in the given node:

$$\text{responsible}_{e^A}(v) = \begin{cases} \text{responsible}_{e^A}(\text{left}(v)) & \text{if } e^A(v) = e_{\text{left}(v)}^A + e_{\text{right}(v)} \\ \text{responsible}_{e^A}(\text{right}(v)) & \text{if } e^A(v) = e_{\text{left}(v)} + e_{\text{right}(v)}^A \end{cases}$$

$$\text{responsible}_{\text{Env}^A}(v) = \begin{cases} \text{responsible}_{\text{Env}^A}(\text{right}(v)) & \text{if } \text{Env}^A(v) = \text{Env}_{\text{right}(v)}^A \\ \text{responsible}_{e^A}(\text{right}(v)) & \text{if } \text{Env}^A(v) = \text{Env}_{\text{left}(v)} + e_{\text{right}(v)}^A \\ \text{responsible}_{\text{Env}^A}(\text{left}(v)) & \text{if } \text{Env}^A(v) = \text{Env}_{\text{left}(v)}^A + e_{\text{right}(v)} \end{cases}$$

We start the search in the root node r looking for $\text{responsible}_{\text{Env}^A}(r)$ and continue down the tree using the formulas above (and possibly switching from $\text{responsible}_{\text{Env}^A}(v)$ to $\text{responsible}_{e^A}(v)$ on the path) until we reach a leaf.

6.2 Improving Detection

Consider the example on Figure 4. In this example we can see that in every solution $\text{end}(M) \leq \text{end}(O)$ because the maximum possible value for $\text{end}(M)$ is $\text{lct}_M = 5$ and the minimum possible value for $\text{end}(O)$ is $\text{est}_O + p_O = 5$. Therefore $M \prec O$. Similarly $N \prec O$. However Edge Finding rule (EF1) is not able to detect that $\{M, N\} \prec O$ and we miss the update of est_O from 0 to 5. It is a similar situation to Detectable Precedences for unary resource described in [7].

The idea is to improve the propagation by improving the knowledge of the relation \prec stored in the array `prec`:

$$\text{prec}[i] := \max \{ \text{prec}[i], \text{est}_i + p_i \}$$

It takes time $\mathcal{O}(n)$ to update all `prec[i]` according to the formula above.




Fig. 5. An example: est_Z can be updated from 0 to 2.

7 Time Bound Adjustment

Let us return again to the example on Figure 1. The algorithm presented in the previous chapter detected that $\{A, B, C\} \prec D$. We will try to use this knowledge to update est_D . Notice that activity A is actually not important for the update (but it was important in the previous phase to realize that $\{A, B, C\} \prec D$), it is a set $\Omega = \{B, C\} \subset \Theta$ which determines new est_D . With this set Ω we can compute new value of est_D denoted as est'_D :

$$est'_D = est_\Omega + \left\lceil \frac{e_\Omega - (C - c_D)(lct_\Omega - est_\Omega)}{c_D} \right\rceil = 2 + \left\lceil \frac{7 - (3 - 2)(5 - 2)}{2} \right\rceil = 4$$

However when $LCut(T, j) \prec i$ we cannot use just any subset $\Omega \subseteq LCut(T, j)$ to compute update of est_i as we did for est_D above. Consider the example on Figure 5. Here $\{W, X, Y\} \prec Z$ but we cannot use $\Omega = \{Y\}$ because the result would be invalid:

$$est_\Omega + \left\lceil \frac{e_\Omega - (C - c_Z)(lct_\Omega - est_\Omega)}{c_Z} \right\rceil = 6 + \left\lceil \frac{1 - (2 - 1)(7 - 6)}{1} \right\rceil = 6$$

The valid update would be to set est_Z to 2, not to 6. The reason that we cannot use $\Omega = \{Y\}$ for update of est_Z is that there is not enough energy in $\Omega = \{Y\}$ to be in potential conflict with Z .

Let us generalize the idea demonstrated on these examples. When $LCut(T, j) \prec i$ then we want to update est_i the following way:

$$LCut(T, j) \prec i \quad \Rightarrow \quad est'_i := \max \{ \text{update}(j, c_i), est_i \} \quad (\text{EF2})$$

where:

$$\text{update}(j, c) = \max_{\substack{\Omega \subseteq LCut(T, j) \\ e_\Omega > (C - c)(lct_\Omega - est_\Omega)}} \left\{ est_\Omega + \left\lceil \frac{e_\Omega - (C - c)(lct_\Omega - est_\Omega)}{c} \right\rceil \right\}$$

The condition $e_\Omega > (C - c)(lct_\Omega - est_\Omega)$ makes sure that we do not make any invalid update as described above.

In the following we will describe how to compute values $\text{update}(j, c)$. When all values $\text{update}(j, c)$ are computed then update of est_i using array `prec` and formula (EF2) is trivial.

Let's assume for simplicity that there are no duplicates in the set $\{\text{lct}_j, j \in T\}$. Therefore if we sort activities T by increasing lct_j in a sequence j_1, j_2, \dots, j_n we get:

$$\text{LCut}(T, j_1) \subsetneq \text{LCut}(T, j_2) \subsetneq \dots \subsetneq \text{LCut}(T, j_n)$$

So when we compute value $\text{update}(j_l, c)$ we do not have to iterate again on all possible subsets $\Omega \subseteq \text{LCut}(T, j_l)$, we can use the fact that we already considered part of them in the computation of $\text{update}(j_{l-1}, c)$. I.e. in the outermost cycle of the algorithm we iterate over all $c \in \{c_m, m \in T\}$ and in the inner cycle we iterate over all $j_l \in T$ and compute:

$$\text{update}(j_l, c) = \begin{cases} \text{diff}(j_1, c) & \text{when } l = 1 \\ \max \{ \text{update}(j_{l-1}, c), \text{diff}(j_l, c) \} & \text{when } l > 1 \end{cases} \quad (10)$$

where:

$$\text{diff}(j, c) = \max_{\substack{\Omega \subseteq \text{LCut}(T, j) \\ e_\Omega > (C-c)(\text{lct}_j - \text{est}_\Omega)}} \left\{ \text{est}_\Omega + \left\lceil \frac{e_\Omega - (C-c)(\text{lct}_j - \text{est}_\Omega)}{c} \right\rceil \right\} \quad (11)$$

So in the computation of $\text{diff}(j, c)$ we “pretend” that all sets $\Omega \subseteq \text{LCut}(T, j)$ has $\text{lct}_\Omega = \text{lct}_j$. This is not true, there may be sets $\Omega \subsetneq \text{LCut}(T, j)$ such that $\text{lct}_\Omega < \text{lct}_j$. However these sets are correctly considered during computation of $\text{diff}(j', c)$ such that $\text{lct}_{j'} = \text{lct}_\Omega$.

Let's define² function $\text{maxest}(j, c)$ as:

$$\text{maxest}(j, c) = \max \{ \text{est}_\Omega, \Omega \subseteq \text{LCut}(T, j) \ \& \ e_\Omega > (C-c)(\text{lct}_j - \text{est}_\Omega) \}$$

Notice that for a particular set Ω_m which defines $\text{maxest}(j, c)$, i.e. $\text{est}_{\Omega_m} = \text{maxest}(j, c)$, we have:

$$\text{est}_{\Omega_m} + \left\lceil \frac{e_{\Omega_m} - (C-c)(\text{lct}_j - \text{est}_{\Omega_m})}{c} \right\rceil > \text{est}_{\Omega_m} = \text{maxest}(j, c)$$

and therefore $\text{diff}(j, c) > \text{maxest}(j, c)$. Now we will show that:

$$\text{diff}(j, c) = \max_{\substack{\Omega \subseteq \text{LCut}(T, j) \\ \text{est}_\Omega \leq \text{maxest}(j, c)}} \left\{ \text{est}_\Omega + \left\lceil \frac{e_\Omega - (C-c)(\text{lct}_j - \text{est}_\Omega)}{c} \right\rceil \right\} \quad (12)$$

The reason follows. The original condition was more restrictive than the new one: in (12) we iterate over more sets Ω than in (11). However for every additional set Ω we have $\text{est}_\Omega \leq \text{maxest}(j, c)$ and $e_\Omega \leq (C-c)(\text{lct}_j - \text{est}_\Omega)$ therefore:

$$\text{est}_\Omega + \left\lceil \frac{e_\Omega - (C-c)(\text{lct}_j - \text{est}_\Omega)}{c} \right\rceil \leq \text{est}_\Omega \leq \text{maxest}(j, c)$$

And we already know that $\text{diff}(j, c) > \text{maxest}(j, c)$. Therefore newly added sets cannot influence the resulting maximum value in formula (12).

² Thanks to Luc Mercier and Joseph Scott for correcting min/max error in the paper.

Formula (12) is algebraically equivalent to:

$$\text{diff}(j, c) = \left\lceil \frac{\text{Env}(j, c) - (C - c) \text{lct}_j}{c} \right\rceil \quad (13)$$

where:

$$\text{Env}(j, c) = \max_{\substack{\Omega \subseteq \text{LCut}(T, j) \\ e_\Omega > (C - c)(\text{lct}_j - \text{est}_\Omega)}} \{C \text{est}_\Omega + e_\Omega\} \quad (14)$$

We can split each set Ω by $\text{maxest}(j, c)$ into two parts:

$$\begin{aligned} \Omega_1 &= \{l, l \in \Omega \ \& \ \text{est}_l \leq \text{maxest}(j, c)\} \\ \Omega_2 &= \{l, l \in \Omega \ \& \ \text{est}_l > \text{maxest}(j, c)\} \end{aligned}$$

And then $C \text{est}_\Omega + e_\Omega = C \text{est}_{\Omega_1} + e_{\Omega_1} + e_{\Omega_2}$. Let's apply this idea on formula (14). We define:

$$\begin{aligned} \alpha(j, c) &= \{l, l \in \text{LCut}(T, j) \ \& \ \text{est}_l \leq \text{maxest}(j, c)\} \\ \beta(j, c) &= \{l, l \in \text{LCut}(T, j) \ \& \ \text{est}_l > \text{maxest}(j, c)\} \end{aligned}$$

And (14) is equivalent to:

$$\begin{aligned} \text{Env}(j, c) &= \max_{\substack{\Omega_1 \subseteq \alpha(j, c) \\ \Omega_2 \subseteq \beta(j, c)}} \{C \text{est}_{\Omega_1} + e_{\Omega_1} + e_{\Omega_2}\} = \\ &= e_{\beta(j, c)} + \text{Env}(\alpha(j, c)) \end{aligned} \quad (15)$$

We can compute $\text{Env}(\alpha(j, c))$ by building Θ -tree for the set $\alpha(j, c)$ as shown in Proposition 1. However it is more suitable for the algorithm to build Θ -tree for the whole set $\text{LCut}(T, j)$ and cut it into two parts just before the computation of $\text{Env}(\alpha(j, c))$. The *cut* operation splits the tree into two trees, it is done in such a way that all activities $l \in \text{LCut}(T, j)$ such that $\text{est}_l \leq \text{maxest}(j, c)$ go to the left part while the others go into the right part. See Figure 6 for an example. The cut operation has time complexity $\mathcal{O}(\log n)$ and it splits the set $\text{LCut}(T, j)$ into sets $\alpha(j, c)$ and $\beta(j, c)$. The value $\text{Env}(\alpha(j, c))$ can be found in the root node of the Θ -tree for $\alpha(j, c)$, and $e_{\beta(j, c)}$ can be found in the root node of the Θ -tree for $\beta(j, c)$.

It remains to show how to compute $\text{maxest}(j, c)$. The value $\text{maxest}(j, c)$ was defined as:

$$\text{maxest}(j, c) = \min \{\text{est}_\Omega, \Omega \subseteq \text{LCut}(T, j) \ \& \ e_\Omega > (C - c)(\text{lct}_j - \text{est}_\Omega)\}$$

The condition $e_\Omega > (C - c)(\text{lct}_j - \text{est}_\Omega)$ is algebraically equivalent to:

$$(C - c) \text{est}_\Omega + e_\Omega > (C - c) \text{lct}_j \quad (16)$$

Notice that the left part of this inequality is very similar to the computation of energy envelope, just C is replaced by $(C - c)$. Let us define a new variant of energy envelope Env^c :

$$\text{Env}^c(\Theta) = \max_{\Omega \subseteq \Theta} \{(C - c) \text{est}_\Omega + e_\Omega\}$$

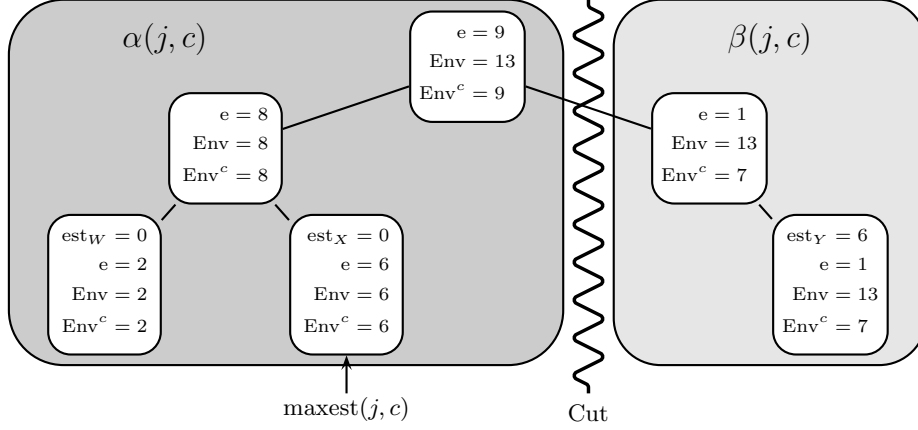


Fig. 6. Example: computation of $\text{Env}(j, c)$ for $c = 1$ and $j = Y$ from example on Figure 5. Therefore $\text{LCut}(T, j) = \{W, X, Y\}$. Situation just before the cut.

The computation of Env^c can be done again using Θ -tree by Proposition 1. We can compute Env and Env^c in the same Θ -tree as shown on Figure 6. Now we can see that because of condition (16) it must hold³:

$$\text{Env}^c(\beta(j, c)) \leq (C - c) \text{lct}_j$$

but if we would include activities l with $\text{est}_l = \text{maxest}(j, c)$ into the right tree (in other words if we would do the cut more on the left) then this condition would not hold. That allows to find a leaf l with $\text{est}_l = \text{maxest}(j, c)$ by following a path from the root the leaf as shown in Algorithm 1.2. Using this procedure we can compute all values $\text{update}(j, c)$ by Algorithm⁴ 1.3 with time complexity $\mathcal{O}(kn \log n)$. Note that once $\text{update}(j, c)$ is computed we can trivially update values est_i using (EF2).

8 Relation with standard Edge Finding

We will show that the algorithm described in the paper does not miss any update done by Edge Finding algorithm described in [5]. It is enough to prove that the original Edge Finding propagation rules are subsumed by the new rules (EF1) and (EF2).

The traditional Edge Finding rule is:

$$\forall i \in T, \forall \Theta \subseteq T \setminus \{i\} : C(\text{lct}_\Theta - \text{est}_{\Theta \cup \{i\}}) < e_{\Theta \cup \{i\}} \Rightarrow \text{est}_i := \max(\text{est}_i, \text{newest}_i)$$

where:

$$\text{newest}_i = \max_{\substack{\Omega \subseteq \Theta \\ e_\Omega > (C-c)(\text{lct}_\Omega - \text{est}_\Omega)}} \left\{ \text{est}_\Omega + \left\lceil \frac{e_\Omega - (C-c)(\text{lct}_\Omega - \text{est}_\Omega)}{c_i} \right\rceil \right\} \quad (17)$$

³ Thanks to Albert Oliveras Llunell for correcting a typo in the original paper.

⁴ Thanks to Joseph Scott for correcting a typo on line 9.

Algorithm 1.2. Computation of $\text{maxest}(j, c)$ using Θ -tree for $\text{LCut}(T, j)$

```
1  $v := \text{root};$ 
2  $E := 0;$ 
3 while  $v$  is not a leaf node do begin
4   if  $\text{Env}^c(\text{right}(v)) + E > (C - c) \text{lct}_j$  then
5      $v := \text{right}(v);$ 
6   else begin
7      $E := E + e_{\text{right}(v)};$ 
8      $v := \text{left}(v);$ 
9   end;
10 end;
11  $l :=$  activity represented by leaf  $v$ ;
12 return  $\text{est}_l$ ;
```

Algorithm 1.3. Computation of all $\text{update}(j, c)$

```
1 for  $c \in \{c_m, m \in T\}$  do begin
2    $\Theta := \emptyset;$ 
3    $\text{upd} := -\infty;$ 
4   for  $j \in T$  in non-decreasing order by  $\text{lct}_j$  do begin
5      $\Theta := \Theta \cup \{j\};$ 
6      $\text{maxest} := \text{maxest}(j, c);$  // see Algorithm 1.2
7      $(\alpha, \beta) := \text{Cut}(\Theta, \text{maxest});$ 
8      $\text{Env}(j, c) := e(\beta) + \text{Env}(\alpha);$  // see (15)
9      $\text{diff} := \left\lceil \frac{\text{Env}(j, c) - (C - c) \text{lct}_j}{c} \right\rceil;$  // see (13)
10     $\text{upd} := \max(\text{upd}, \text{diff});$  // see (10)
11     $\text{update}(j, c) := \text{upd};$ 
12     $\Theta := \text{join}(\alpha, \beta);$ 
13  end;
14 end;
```

Let's consider an activity i and sets Θ and Ω which achieves the best update by the rule above. Then we can define j to be an activity from Θ such that $\text{lct}_j = \text{lct}_\Theta$. And because $\Theta \subseteq \text{LCut}(T, j)$ we can see that the rule (EF1) holds for i and j . And because $\Omega \subseteq \Theta \subseteq \text{LCut}(T, j)$ the update by the rule (EF2) must be at least the same as by the original rule (17).

9 Optional Activities

Optional activity is an activity which may or may not be present in the resulting schedule [1]. Optional activities makes modeling of certain types of problems much easier (for example dealing with alternatives) and it also allows the CP engine to propagate better. Therefore it is very important that Edge Finding algorithm can handle optional activities.

To handle optional activities we can use the same idea as suggested in [4]: instead of changing the algorithm we can just change its input data. If an activity j is optional, we set for the algorithm $\text{lct}_j = \infty$ regardless the real value of lct_j . This way the algorithm can never conclude that $j < i$ for any activity i because from the point of view of the algorithm the activity j can be always scheduled later than i . Therefore optional activities will be influenced by non-optional ones, but non-optional activities will not be influenced by optional ones.

Note that propagation for optional activities could be further improved as suggested for unary resource in [6]. However it would probably result in increase of time complexity of the algorithm.

10 Experimental Results

Speed of the presented algorithm was tested against incomplete algorithm [2] by measuring time needed for initial propagation. These tests was done on cumulative job-shop instances with resources of capacity 2 (note that in this case $k = 1$). For $n = 20$ activities on resource the presented algorithm is on average faster by factor 1.34, for $n = 30$ it is faster by factor 1.60, for $n = 40$ by 1.99, for $n = 60$ by 2.68, for $n = 100$ by 4.15 and for $n = 200$ by factor 7.35.

11 Conclusions

This paper presents a new Edge Finding algorithm for discrete capacity resources. The new algorithm is stronger than the state-of-the-art algorithm [5], it is faster (in term of time complexity) and it can handle optional activities. The algorithm is successfully used by CP-Optimizer [1] starting from version 2.0.

References

1. : IBM ILOG CP Optimizer <http://www.ilog.com/products/cpoptimizer/>.
2. Philippe Baptiste, C.L.P., Nuijten, W.: Constraint-Based Scheduling: Applying Constraint Programming to Scheduling Problems. Kluwer Academic Publishers (2001)
3. Schutt, A., Wolf, A., Schrader, G.: Not-first and not-last detection for cumulative scheduling in $O(n^3 \log n)$. In: 16th International Conference on Applications of Declarative Programming and Knowledge Management, INAP 2005, Springer (2005) 66–80
4. Vilím, P.: Max energy filtering algorithm for discrete cumulative resources. In Willem-Jan van Hoeve, J.N.H., ed.: Proceedings of CP-AI-OR 2009. Volume 5547 of Lecture Notes in Computer Science., Springer-Verlag (2009) 294–308
5. Mercier, L., Hentenryck, P.V.: Edge finding for cumulative scheduling. *Inform. Journal of Computing* **20**(1) (2008) 143–153
6. Kuhnert, S.: Efficient edge-finding on unary resources with optional activities. In: Proceedings of INAP 2007 and WLP 2007. (2007)
7. Vilím, P.: Global Constraints in Scheduling. PhD thesis, Charles University in Prague, Faculty of Mathematics and Physics, Department of Theoretical Computer Science and Mathematical Logic (2007)