



HAL
open science

A Discounted Cost Function for Fast Alignments of Business Processes

Mathilde Boltenhagen, Thomas Chatain, Josep Carmona

► **To cite this version:**

Mathilde Boltenhagen, Thomas Chatain, Josep Carmona. A Discounted Cost Function for Fast Alignments of Business Processes. BPM 2021 - 19th International Conference on Business Process Management, Sep 2021, Rome, France. 10.1007/978-3-030-85469-0_17 . hal-03362277

HAL Id: hal-03362277

<https://inria.hal.science/hal-03362277v1>

Submitted on 1 Oct 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Discounted Cost Function for Fast Alignments of Business Processes

Mathilde Boltenhagen¹, Thomas Chatain¹, and Josep Carmona²

¹ Université Paris-Saclay, ENS Paris-Saclay, CNRS, Inria, LSV, Gif-sur-Yvette, (France)
{boltenhagen, chatain}@lsv.fr

² Universitat Politècnica de Catalunya, Barcelona (Spain)
jcarmona@cs.upc.edu

Abstract. Alignments are a central notion in conformance checking. They establish the best possible connection between an observed trace and a process model, exhibiting the closest model run to the trace. Computing these alignments for huge amounts of traces, coming from big logs, is a computational bottleneck. We show that, for a slightly modified version of the distance function between traces and model runs, we significantly improve the execution time of an A*-based search algorithm. We show experimentally that the alignments found with our modified distance approximate very nicely the optimal alignments for the classical distance.

1 Introduction

Conformance checking techniques enable establishing a relation between modeled and observed behavior [5]. The techniques on this field are grounded on solving a very particular problem, known as *alignment* [1]: given a process model that describes a certain process, and a trace which contains a potential observation of this process, to decide if the trace is in the language of the model, and if not, to pinpoint where it deviates. Computing alignments is not necessarily the ultimate goal of an analysis, but instead can be used to further enhance a process model with the evidences found in the data, e.g. depicting in the model explicitly the bottlenecks of the underlying process [20].

As happened a decade ago with process discovery, process mining software vendors are seriously considering incorporating conformance checking capabilities in their solutions, since conformance checking unleashes the full potential of model-based reasoning on top of real process data. However, whilst process discovery techniques may rely on light techniques, based on the *directly-follows* relations computed in event log [10,3] without significantly sacrificing the quality of the obtained solution, conformance checking techniques instead need to fight the state-space explosion problem, due to the need to explore, when computing alignments, the behavior underlying the process model. Hence, industry adoption for conformance checking is still not widespread.

The current process mining field is living an interesting paradox: whilst it is widely accepted that the capabilities of discovering huge process models exist, when it comes to analysing these discovered models through conformance checking techniques, only approximate techniques for deriving alignments are used in practice. In the next section we provide a complete overview of current alternatives for alignment computation.

In this paper, however, we somehow go back to the roots, and adopt the seminal work from Arya Adriansyah’s PhD thesis [1] as main alignment algorithm. In the aforementioned work, an A^* graph search algorithm is proposed over the graph corresponding to the state space of *synchronous product net*. This net is made out of the initial process model and a *trace net* corresponding to the input trace. The cost function that governs the A^* search is typically a *standard* cost function which assigns unitary costs to all the possible types of deviations.

We consider a rather simple, yet powerful idea that is motivated from the following use case: for certain processes, the costs associated to deviations at early stages of the process’ execution are more important than the ones at the end. For instance, consider a loan application process that has two decisions: one at the beginning, assessing the type of customer (gold, silver, normal), and one at the end, determining whereas the loan was received in a labour day or not. It is normal that the stage in which these decisions are made in any possible execution of the process reveal their importance. For instance, if for the company it is very important to know the type of the customer so that other further information needs to be gathered depending on the customer’s type, then it is likely that the corresponding process has the type of customer decision close to the start of any possible execution. On the contrary, if the day when the loan was received is not so important, then it is likely that the corresponding events will be pushed to the end of the traces.

Having this in mind, one can instantiate the A^* algorithm to make the cost function exponentially biased to this use case: giving more importance (higher cost) to the deviations that occur in early stages of the alignments, and exponentially reducing the cost as the search algorithm progresses. Importantly, this *discounted* cost function has a huge impact on the size of the search space required for the A^* search, since the cost asymmetry makes the search space to rapidly shrink after the first alignment steps are made.

For processes which follow the aforementioned use case, this cost function puts the search focus in the right place, deriving alignments that aim at synchronizing modeled and observed behavior in the important decisions that are made at the beginning. Interestingly, this idea can also be used for processes that are not following this trend, since although putting the focus at the beginning may not be the most likely explanation, the computational alleviation can make the problem tractable, where other techniques fail.

In this paper we formalize this simple idea, and show the great impact in performance with respect to several variations of the A^* search proposed in the last years. Surprisingly, this improvement is with a very minor loss in quality: as we will see in the experiments, for well-known and accepted benchmarks, the proposed techniques is often able to produce alignments very close to the optimal ones.

This paper is organized as follows: in the next section we provide a detailed overview of the different techniques to compute alignments. Then in Section 3 we provide the necessary definitions to understand the technique of this paper. In Sections 4 and 5 we provide the formal definition and corresponding algorithmic adaptations for the discounted cost function presented in this paper. Then in Section 6 an evaluation of the proposed technique is reported, and Section 7 concludes the paper.

2 Related Work

The seminal work in [1] proposed the notion of alignment and developed a technique based on A* to compute optimal alignments for a particular class of process models. The approach represents the state-of-the-art technique for computing alignments, and can be adapted (at the expense of increasing significantly the memory footprint) to provide all optimal alignments. Alternatives to A* have appeared in recent years: in the approach presented in [6], the alignment problem is mapped as an *automated planning* instance. Automata-based techniques have also appeared [13,9]. The techniques in [13] (recently extended in [14]) rely on state-space exploration and determination of the automata corresponding to both the event log and the process model, whilst the technique in [9] is based on computing several subsets of activities and projecting the alignment instances accordingly.

The work in [16] presented the notion of *approximate* alignment to alleviate the computational demands by proposing a recursive paradigm on the basis of the structural theory of Petri nets. In spite of resource efficiency, the solution is not guaranteed to be executable. Alternatively, the technique in [17] presents a framework to reduce a process model and the event log accordingly, with the goal of alleviating the computation of alignments. The obtained alignment, called *macro-alignment* since some of the positions are high-level elements, is expanded based on the information gathered during the initial reduction. Techniques using local search have recently been also proposed very recently [18].

Decompositional techniques have been presented [19,11,23] that, instead of computing optimal alignments, they focus on the *crucial problem* of whether a given trace fits or not a process model. These techniques vertically decompose the process model into pieces satisfying certain conditions (so only *valid* decompositions [19], which satisfy restrictive conditions on the labels and connections forming a decomposition, guarantee the derivation of a real alignment). Later on, the notion of *recomposition* has been proposed on top of decompositional techniques, in order to obtain optimal alignments whenever possible by iterating the decompositional methods when the required conditions do not hold [8]. In contrast to the aforementioned vertical decomposition techniques, our methodology does not require this last consolidation step of partial solutions, and therefore can be a fast alternative to these methods at the expense of losing the guarantee of optimality.

We believe our work has similarities and synergies with two recent works. In [4], a symbolic algorithm to maximize the number of synchronous moves in the alignment is proposed. Recently, in [21], an online alignment technique with a window-based backwards exploration is proposed. In both cases, the technique of this paper can be used to enhance the aforementioned techniques.

3 Preliminaries

We represent event data as log traces and process models as labeled Petri nets.

Definition 1 (Log Traces). Let Σ be a set of activities. We define a log L as a finite multiset of sequences $\sigma \in \Sigma^*$, which we refer to as log traces.

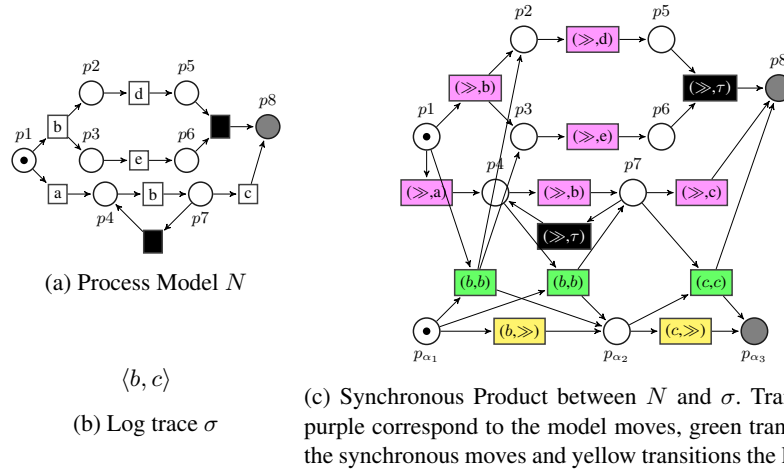


Fig. 1: Synchronous Product for Alignments between N and σ that has the Marking Reachability Problem.

Definition 2 (Process Model (Labeled Petri Net) [12]). A Process Model defined by a labeled Petri net system (or simply Petri net) is a tuple $N = \langle P, T, F, m_0, m_f, \Sigma, \lambda \rangle$, where P is the set of places, T is the set of transitions (with $P \cap T = \emptyset$), $F \subseteq (P \times T) \cup (T \times P)$ is the flow relation, m_0 is the initial marking, m_f is the final marking, Σ is an alphabet of actions and $\lambda : T \rightarrow \Sigma \cup \{\tau\}$ labels every transition by an activity or as silent.

Semantics. The semantics of Petri nets is given in term of firing sequences. Given a node $x \in P \cup T$, we define its pre-set $\bullet x \stackrel{\text{def}}{=} \{y \in P \cup T \mid (y, x) \in F\}$ and its post-set $x^\bullet \stackrel{\text{def}}{=} \{y \in P \cup T \mid (x, y) \in F\}$. A marking is an assignment of a non-negative integer to each place. A transition t is *enabled* in a marking m when all places in $\bullet t$ are marked. When a transition t is enabled, it can *fire* by removing a token from each place in $\bullet t$ and putting a token to each place in t^\bullet . A marking m' is *reachable* from m if there is a sequence of firings $\langle t_1 \dots t_n \rangle$ that transforms m into m' , denoted by $m[t_1 \dots t_n]m'$. The set of reachable markings from m_0 is denoted by $[m_0]$. A Petri net is *k-bounded* if no marking in $[m_0]$ assigns more than k tokens to any place. A Petri net is *safe* if it is 1-bounded. A *full run* of a Petri net N is a firing sequence $m_0[t_1 \dots t_n]m_f$ from the initial marking m_0 to the final marking m_f . A Petri net is *easy sound* [2] if it has at least one full run, i.e. m_f is reachable from m_0 .

In this paper we assume safe and easy sound Petri nets.

Definition 3 (Alignments). Given a log trace $\sigma = \langle \sigma_1, \dots, \sigma_m \rangle \in L$ of alphabet Σ , and a process model $N = \langle P, T, F, m_0, m_f, \Sigma, \lambda \rangle$, an alignment of σ with N is a finite sequence $\varphi = \langle (\sigma'_1, u'_1), \dots, (\sigma'_p, u'_p) \rangle$ of moves such that:

- each move is either: a synchronous move $(a, t) \in \Sigma \times T$ with $a = \lambda(t)$, a log move (a, \gg) (where \gg is a special ‘skip’ symbol), or a model move (\gg, t) ,

- dropping the \gg symbols from the left projection $\langle \sigma'_1, \dots, \sigma'_p \rangle$ of φ , yields σ ;
- dropping the \gg symbols from the right projection $\langle u'_1, \dots, u'_p \rangle$ of φ , yields a full run u of N .

Example 1. Fig 1 shows a process model (Fig. 1a) and a log trace (Fig. 1b). An alignment of N and σ is $\varphi = \langle (\gg, a), (b, b), (c, c) \rangle$.

In order to compare the quality of alignment quality, we define a cost function which penalizes log moves and model moves.

Definition 4 (Classical Alignment Cost Function, Optimal Alignments). *For every alignment φ between a model N and a log trace σ , the Classical Alignment Cost Function \mathcal{C} assigns a cost 0 to synchronous moves and a cost 1 to log moves or model moves. The cost of an alignment is the sum of the costs of its moves. An alignment between a model N and a log trace σ is optimal if it minimizes the cost.*

Example 2. The alignment $\varphi = \langle (\gg, a), (b, b), (c, c) \rangle$ between N of Fig. 1a and σ of Fig. 1b costs 1 and is optimal.

The optimal alignment cost given with the classical alignment cost function \mathcal{C} gives the Levenshtein edit distance between a run of M and the trace σ .

Definition 5 (Levenshtein Edit distance). *The Levenshtein Edit Distance $dist(u, v)$ between two words u and $v \in \Sigma^*$ is the minimal number of edits needed to transform u to v . In our case, edits can be deletions or insertions of a letter in words.*

The main methods of the literature to compute optimal alignments are Dijkstra-based algorithms which often implies the construction of the *Synchronous Product* between the given process model and a sequential Petri net representing the log trace [1].

Definition 6 (Synchronous Product for Alignments). *For a log trace $\sigma = \langle \sigma_1, \dots, \sigma_m \rangle$ defined in $\Sigma \setminus \gg$ and a process model $N = \langle P, T, F, m_0, m_f, \Sigma, \lambda \rangle$, the Synchronous Product $SN = \langle P_{SN}, T_{SN}, F_{SN}, m_{SN_0}, m_{SN_f}, \Sigma, \lambda_{SN} \rangle$ of alignments between N and σ is given by a product of Petri nets such as:*

- $N_\sigma = \langle P_\sigma, T_\sigma, F_\sigma, m_{\sigma_0}, m_{\sigma_f}, \Sigma, \lambda_\sigma \rangle$ is a translation of σ to a sequential Petri net with: $P_\sigma = \{P_{\sigma_0}, \dots, P_{\sigma_m}\}$, $T_\sigma = \{t_{\sigma_i} = \lambda_\sigma(\sigma_i) \mid i \in \{1, \dots, m\}\}$, $F_\sigma = \{(P_{\sigma_{i-1}}, t_{\sigma_i}), (t_{\sigma_i}, P_{\sigma_i}) \mid i \in \{1, \dots, m\}\}$, $m_{\sigma_0} = \{P_{\sigma_0} : 1\}$, $m_{\sigma_f} = \{P_{\sigma_m} : 1\}$,
- $P_{SN} = P \cup P_\sigma$
- $T_{SN} = T \gg \cup T_\sigma^\gg \cup T_S$, where $T \gg = \{(\gg, t) \mid t \in T\}$ represents the model moves, $T_\sigma^\gg = \{(t, \gg) \mid t \in T_\sigma\}$ represents the log moves, $T_S = \{(t_1, t_2) \mid t_1 \in T, t_2 \in T_\sigma \text{ and } \lambda(t_1) = \lambda_\sigma(t_2)\}$ represents the synchronous moves,
- $F_{SN} = F \cup F_\sigma \cup \{(P_i, t_i) \mid t_i = (t_1, t_2) \in T_{SN}, t_1 \neq \gg, t_2 \neq \gg, P_i \in \bullet t_1 \cap \bullet t_2\} \cup \{(t_i, P_i) \mid t_i = (t_1, t_2) \in T_{SN}, t_1 \neq \gg, t_2 \neq \gg, P_i \in t_1 \bullet \cap t_2 \bullet\}$
- $m_{SN_0} = m_0 \cup m_{\sigma_0}$,
- $m_{SN_f} = m_f \cup m_{\sigma_f}$,
- λ_{SN} maps every $t \in T_{SN}$ to its move.

Example 3. Fig 1 shows the synchronous product for alignments of the process model N given in Fig. 1a and the log trace σ of Fig. 1b.

The Dijkstra-based algorithm for finding optimal alignments, explores the reachability graph of the synchronous product of Def. 6. Weights are given by the transitions fired to reach the markings, according to the type of move that they represent. The best firing sequences found for reaching a marking is the less costly one. The algorithm that we present in Section 5 is an adaptation of this classical Dijkstra-based algorithm for alignments. As we are using easy-sound Petri nets as process models, the Synchronous Products for Alignments are easy-sound which implies termination of the Dijkstra algorithm with the condition to reach the final marking m_{SN_f} [24].

4 Discounted Cost Function and Properties

The classical alignment cost function corresponds to Levenshtein edit distance between a run of a process model and a log trace, where additions and deletions represent model and log moves. In this section, we introduce the *Discounted Edit Distance* and its impact when using it as alignment cost function.

The idea of this Discounted Edit Distance is to penalize more insertions and deletions when they occur at the beginning of the strings, and less when they occur later.

Definition 7 (Discounted Edit Distance). We define the Discounted Edit Distance between two strings u and v (of length $|u|$ and $|v|$ respectively) with discount parameter $E \geq 1$ by $\mathcal{D}_E^{u,v}(|u|, |v|, 0)$ where:

$$\mathcal{D}_E^{u,v}(i, j, k) = \begin{cases} 0 & \text{if } (i = 0 \text{ and } j = 0) \\ \mathcal{D}_E^{u,v}(i-1, j-1, k+1) & \text{if } (u_i = v_j) \\ \mathcal{D}_E^{u,v}(i, j-1, k+1) + E^{-k} & \text{if } i = 0 \\ \mathcal{D}_E^{u,v}(i-1, j, k+1) + E^{-k} & \text{if } j = 0 \\ \min \begin{cases} \mathcal{D}_E^{u,v}(i-1, j, k+1) + E^{-k} \\ \mathcal{D}_E^{u,v}(i, j-1, k+1) + E^{-k} \end{cases} & \text{otherwise.} \end{cases}$$

thus allowing equation $((u_i = v_j))$ for free and insertions and deletions at cost E^{-k} where k refers to the position where the edit occurs.

Lemma 1. For $E = 1$, the Discounted Edit Distance corresponds to the Levenshtein distance.

In practice, relevant values for the discount parameter E are slightly larger than 1. For $E = 2$, the discount is already very severe since an edit at position k costs more than the sum of all the following edits.

Lemma 2. With the Discounted Edit Distance, for $E = 2$, an edit at position k costs more than the sum of all the following edits.

Proof. For u and v , two words, let k be the position of a non-free cost in $\mathcal{D}_E^{u,v}$. We note its cost $c(k)$ and for $E = 2$, $c(k) = 2^{-k}$.

The next edits can occur at positions $j \in \{k + 1, \dots, n\}$ where, in the worst case, $n = |u| + |v|$. We write $S(j, n)$ the sum of costs. The maximal value of this sum appears when only non-free edits are used by the discounted edit distance:

$$S(k, n) = \sum_{j=k+1}^n c(j) = 2^{-(k+1)} + 2^{-(k+2)} + \dots + 2^{-n} = 2^{-k} - 2^{-n}$$

Hence, $c(k) = 2^{-k} > S(k, n)$. Otherwise, in the best case, there is no edit after position k and the cost of the edit at position k is higher than a null sum. \square

Example 4. Let $u = \langle x, a, b \rangle$ and $v = \langle a, y, b \rangle$. The discounted edit distance between u and v is $\mathcal{D}_E^{u,v} = E^{-1} + E^{-3}$. If $E = 1$, the distance equals to 2 and is the Levenshtein edit distance where deleting x costs 1 and adding y costs 1.

Similarly to the Levenshtein edit distance, the Discounted Edit Distance can be applied to alignments.

Definition 8 (Discounted Cost Function for Alignments). For an alignment $\varphi = \langle (\sigma'_1, u'_1), \dots, (\sigma'_p, u'_p) \rangle$ between a process model N and a log trace σ , the Discounted Cost Function for Alignments assigns a cost 0 to every synchronous move and E^{-k} to every pair (σ'_k, u'_k) that is either a log move either a model move, where $k \in \{1, \dots, p\}$, $p \in \mathbb{N}$ is the length of the alignment, and $E \geq 1$ is the discount parameter.

For $E = 1$, the Discounted Cost Function for Alignment is equivalent to the Classical Alignment Cost Function. However when $E > 1$, the costs of moves are dynamic and depend on the number of previous moves of the alignment.

Example 5. For the alignment $\varphi = \langle (\gg, a), (b, b), (c, c) \rangle$, presented in the first example, the cost of (\gg, a) is E^{-1} because it is the first move of the alignment. For $\varphi' = \langle (b, b), (\gg, d), (\gg, e), (\gg, \tau) \rangle$ which is certainly not an optimal alignment but still an alignment of $\sigma = \langle b, c \rangle$ and the Petri net N of Fig. 1a, the cost of (\gg, e) is E^{-3} .

Lemma 3. For $E > 1$, a non-free move ω of position j , any move of position $k > j$ costs less than ω .

As a consequence, algorithms for computing optimal discounted alignments will tend to align in priority the prefixes of the log traces. Suffixes are less costly. From Lemma 2, when the discount parameter is $E = 2$, a non-free move of position j is more costly than the sum of all the next costs.

Example 6. In Example 1, we saw that the optimal alignment by using the classical alignment cost function between $\sigma = \langle b, c \rangle$ and the model N of Fig. 1a is $\varphi = \langle (\gg, a), (b, b), (c, c) \rangle$ of cost 1. However, by using the discounted cost function with $E = 2$, optimal alignments are of $\varphi' = \langle (b, b), (\gg, d), (\gg, e), (\gg, \tau) \rangle$ and $\varphi'' = \langle (b, b), (\gg, e), (\gg, d), (\gg, \tau) \rangle$ of cost $2^{-2} + 2^{-3}$. This is due to the position of the move (\gg, a) in φ which is extremely costly.

5 Using the Discounted Cost Function in an A*-based Algorithm for Discounted Alignments

To compute alignments by using the discounted cost function, we present an A*-based algorithm which assigns weights to the explored states according to the discounted cost function for alignment. Let E be the discount parameter. Then, to a state reached by a move ω occurring in position i , will be assigned the weight of its predecessor, increased by the cost

$$h(\omega, i, E) \stackrel{\text{def}}{=} (0 \text{ if } \omega \text{ is a synchronous move, } E^{-i} \text{ otherwise}).$$

As a result of Lemma 3, this heuristic aims at aligning prefix first more than suffixes.

The function h , based on the discounted cost function, is easily incorporated in the state-of-the-art algorithms for computing alignments. We present two versions of the incorporation, one by using the synchronous product of the process model and the log trace, and another one that avoids the computation of the product by exploring the process model along with the trace.

5.1 Algorithm for Computing Optimal Discounted Alignments

Our algorithm Alg. 1, noted $A^*SP_{D=E}$, is adapted from [5]. It proceeds by exploring the state space of the synchronous product of the process model with a sequential Petri net representing the log trace as defined in Def. 6. An optimal alignment corresponds to the shortest path between the initial marking to the final marking of the synchronous product. For the discounted cost function, a state contains a marking and the length of the current sequence to reach the marking. For instance line 2 contains $\langle m_0, 0 \rangle$ which corresponds to the initial state.

Each state is associated to a cost which is stored in the function d . Cost of states are given by the optimal firing sequences of transitions to reach the states. Some transitions of the synchronous product correspond to the log and model moves and are costly. The other transitions are the synchronous moves and are free. In Fig .1c, the green transition are free and all the other aren't.

Our discounted cost function h appears on line 19. It determines if the move is costly and attributes the cost given the current length of the alignment, i.e, the position of the move j , and the parameter E .

The algorithm returns the best sequences of moves to reach the final marking.

Role of Length in States In the classical version of alignment computation, the state contains the markings only. However, the length of the current alignment plays an important role in the discounted cost function. Indeed, the first visit of a marking might not be the optimal one, as it is the case in the classical version of alignments. A same marking ω can be reached with different firing sequences of moves of different lengths. The first path that gives the first visit of the marking ω is the shortest one. Let's call this short path γ_{short} , and γ_{long} a longer path from the initial marking to this marking ω . We have $|\gamma_{long}| > |\gamma_{short}|$. However, γ_{short} might contain future costly moves to reach the final marking. If those moves are of position lower than $|\gamma_{long}|$, it questions the optimality of γ_{short} . We give an example of this situation below.

Example 7. Let's suppose that silent transition labelled by τ costs for this example. To reach marking $\omega = \{p_4 : 1, p_{\alpha_2} : 1\}$, the algorithm can play $\gamma = \langle (\gg, a), (\gg, b), (\gg, \tau a), (b, b) \rangle$ whose cost is $E^{-1} + E^{-2} + E^{-3}$. This firing sequence costs as much as $\gamma' = \langle (\gg, a), (\gg, b), (b, \gg) \rangle$ which reaches the same marking ω . However we notice that γ has a synchronisation at position 4 but we don't know yet what appear at position 4 for γ' . Then both paths should be kept.

Notify that we tackled the problem of optimality of the alignment with the discounted cost function. For $E > 1$, this optimality does not correspond to the optimal classical alignment.

Algorithm 1: Computation of Optimal Alignments for the Discounted Cost Function ($A^*SP_{\mathcal{D}=E}$)

Input : $SP = ((P, T, F, \lambda, m_0, m_f))$: synchronous product,
 E : discount parameter

```

1  $A \leftarrow \emptyset$  ▷ Initialise closed set
2  $X \leftarrow \{\langle m_0, 0 \rangle\}$  ▷ Initialise open set
3  $p(\langle m_0, 0 \rangle) = \text{null}$  ▷ Initialise predecessor function
4  $\forall m \in RS(SN) \forall i \in \{0, \dots, \max_{|u|, u \in R_{\text{uns}}(SP)} d(\langle m, i \rangle) = \infty$  ▷ Initialise function d
5  $d(\langle m_0, 0 \rangle) = 0$ 
6 while  $X \neq \emptyset$  ▷ While not all states visited
7 do
8    $\langle m, i \rangle \leftarrow \langle m, i \rangle \in X$  minimizing  $d(m)$  ▷ Get the closest state
9   if  $m = m_f$  then
10     break while
11    $A \leftarrow A \cup \{\langle m, i \rangle\}$  ▷ Add state to the closed set
12    $X \leftarrow X \setminus \{\langle m, i \rangle\}$  ▷ Remove state to the open set
13   for  $t \in T$  with  $m[t]m'$  do
14     if  $\langle m', i + 1 \rangle \notin A$  ▷ Reaching a state not yet visited
15     then
16        $X \leftarrow X \cup \{\langle m', i + 1 \rangle\}$  ▷ Add state to the open set
17        $c \leftarrow d(\langle m, i \rangle) + h(t, i + 1, E)$  ▷ Get the current cost of reaching state
18
19       if  $c < d(\langle m', i + 1 \rangle)$  ▷ If current cost is better than known cost
20       then
21          $d(\langle m', i + 1 \rangle) \leftarrow c$  ▷ Update distance function
22          $p(\langle m', i + 1 \rangle) \leftarrow (t, \langle m, i \rangle)$  ▷ Update predecessor function
23
24  $\gamma \leftarrow \langle t_0, \dots, t_n \rangle$  such that
     $t_n = \#_1(p(\langle m_f, |\gamma| \rangle)), t_{n-1} = \#_1(p(\#_2(p(\langle m_f, |\gamma| \rangle)))$  etc. until the
    initial state  $\langle m_0, 0 \rangle$  is reached recursively.
Output:  $d(\langle m_f, |\gamma| \rangle), \gamma$ 

25 Function  $h(t, i, E)$  :
26   if  $t$  is a synchronous move then
27     return 0
28   else
29     return  $E^{-i}$  ▷ Discounted cost

```

Comparison to Classical Alignments Due to the discount parameter E in the discounted cost function, our heuristic prioritizes the alignment of the beginning of the log

trace. In the algorithm, this difference with the classical alignment algorithm appears line 9 of Alg. 1 where the markings that minimize the cost are much more different with the discounted cost function than by using the classical cost function for alignments. Indeed, when costs are all equivalent, many paths compete in the search for the optimal alignment. However, with very different costs, the number of paths with similar costs is low, thus reducing the search space. This characteristic is well observed in practice with a reduction of runtime.

Example 8. For the example of Fig. 1, there is a first choice between a and b . For large E , the decision is quickly given thus disabling testing the depth of the other paths. For instance, with $E = 2$, the log sequences of type $\langle a, b, c \rangle$, $\langle a, b, \tau, b, c \rangle$, $\langle a, b, \tau, b, \tau, b, c \rangle$ won't be explored because they cannot have a better discounted cost. Then the search space is very reduced.

5.2 A Heuristic for Reducing the Search Space of the Algorithm

The search space of $A^*SP_{\mathcal{D}=E}$ is large and even larger than the Dijkstra-based algorithm for alignments due to the incorporation of the lengths in the states. To reduce it, we come back to the classical states that contain the markings only. Every $\langle m, i \rangle$ of Alg. 1 is reduced to m (like in [5]).

With this reduction, only the first path that reaches the marking is used. When several concurrent firing sequences of equal cost exist, line 8 picks one as the optimal path and line 11 classifies the marking in the closed set. Then the other firing sequences of equal cost (but different length) for this marking are not considered anymore.

Example 9. For the marking $\omega = \{p4 : 1, p_{\alpha_2} : 1\}$ of the synchronous product given in Fig. 1c, two firing sequences compete for the minimization of the cost (in case that silent transition costs). Indeed both $\gamma = \langle (\gg, a), (\gg, b), (\gg, \tau), (b, b) \rangle$ and $\gamma' = \langle (\gg, a), (\gg, b), (b, \gg) \rangle$ have a cost of $E^{-1} + E^{-2} + E^{-3}$ and reach the marking ω . Hence, when using markings only in the states, the algorithm picks one of the sequence as the optimal one and adds ω in the closed set. Later, it does not consider the other firing sequence. However, we saw in Example 7 that γ' is better than γ by using the discounted cost function but it can go to the hatch in the reduced version.

With this reduction of the search space, the modified algorithm is not guaranteed anymore to return the optimal discounted alignments, but the gain in runtime is extremely significant. Moreover, in practice, the loss of quality is very limited: we observed that the alignments found by the modified algorithm have very similar discounted cost than the optimal discounted alignments.

Process Model Exploration Along with Trace Exploration (noted $A^*PT_{\mathcal{D}=E}$) Finally, in order to speed up the exploration, the alignment algorithm can simulate the synchronous product without explicitly constructing it. The idea is to replace the markings of the synchronous product by the markings of the process model along with the position of the trace. For instance, marking $\{p4 : 1, p_{\alpha_2} : 1\}$ of the synchronous product given in Fig. 1c is translated to $(\{p4 : 1\}, 1)$ where $\{p4\}$ is the marking in N and

1 the position in σ . Then the final state is $(mf, |\sigma|)$. The predecessor function p of this algorithm is slightly different to obtain γ the sequence of moves. To insert a move in the alignment, the algorithm requires to determine its type by simply comparing the labels of the fired transition and the activity of the current position in the trace.

6 Experiments: Discounted Alignments as a Heuristic for Approximating Classical Alignments

This section aims at showing a comparison of quality and runtimes between the proposed version of alignments and existing ones. We first present general comparisons. Then we stress on the impacts of the discounted parameter. Finally, we zoom on particular cases and added other methods similar to our.

6.1 Comparison with respect to baselines

Inputs We experimented the algorithms for both artificial and real-life logs and the corresponding models. Artificial set is taken from [15] and contains large models. For real-life logs, we used data given in the Business Process Intelligence Challenges from 2012 to 2020. We mined the process models of those logs with methods of the literature³. First, we applied a preprocessing method introduced by [7] to extract good prototypes of the logs⁴. This preprocessing step allows us to obtain not perfectly fitting models when using miners, interesting for alignments comparison. In fact, the method aims a finding more precise models. Then, we launched two different discovery algorithms on the found prototypes: the inductive miner [10] and the split miner 2.0 [3]. As the latter tool gives BPMN models, we use ProM plugins to transform them into Petri nets.

We computed the alignments on *variants* only, i.e., unique sequences of activities. This choice of using variants only allows to correctly compare the method’s runtimes in case that one reduce the log to variants and not another one. Table 2 gives an overview of the inputs.

Implementation The two algorithms presented in the previous section have been implemented in a branch of pm4py⁵.

Comparison We compare our alignment results to the four current methods of the state-of-the-art implemented in pm4py which are: -the Dijkstra search on the Synchronous Product without heuristic (*DSP*) [1], -a Dijkstra that consumes less memory by using a similar idea of our second algorithm (*DLM*), -an A*-based algorithm on the state-space of the synchronous product that incorporates an heuristic on reaching the final marking (*A*SP_{mf}*) [22] and -its less-memory version (*A*LM_{mf}*). To compare the

³ available at <https://github.com/BoltMaud/A-Discounted-Cost-Function-for-Fast-Alignments-of-Business-Processes-Sources>

⁴ Prototype Selection plugin of Prom software with default settings

⁵ currently available at <https://github.com/BoltMaud/pm4py-core>

Log	#variants	$ \Sigma $	$\max_{\sigma \in Log} len(\sigma)$	Model	$ T $	$ P $	$ F $
L1	453	36	37	M1	39	40	92
L2	500	32	52	M2	34	34	80
L3	462	109	217	M3	123	108	276
L4	496	44	176	M4	52	36	106
L5	500	32	71	M5	33	35	78

(a) Artificial Logs and Models

Log	#variants	$ \Sigma $	$\max_{\sigma \in Log} len(\sigma)$	Model Miner	$ T $	$ P $	$ F $
BPI2012	4366	24	175	IM	34	24	68
				SM	30	23	60
BPI2018 _{pa}	3832	24	100	IM	22	24	60
				SM	20	15	40
BPI2019	11973	42	990	IM	18	13	38
				SM	13	10	26
BPI2020 _{dd}	99	17	24	IM	15	11	32
				SM	14	9	28
BPI2020 _{rp}	89	19	20	IM	31	26	74
				SM	23	12	46

(b) Real-life Logs and Models

Fig. 2: Input Description

runtimes, we exclude log-based implementations and used the trace-based version to avoid the use of the parallelism between variants that can be added to any version at the case level. We recall notation of our methods $A^*SP_{\mathcal{D}=E}$, for the version that uses the synchronous product, and $A^*PT_{\mathcal{D}=E}$, for the second one that explores only the process model and the trace.

Results The aim of this section is to compare the *runtimes* to compute the alignments and the *quality* of the obtained alignments. The *quality* of an alignment found by a heuristic method, is defined as the ratio (in %) between the classical cost (number of model or log moves) of the optimal alignment and the classical cost of the alignment found by the method. In Fig.3a we give the quality of each method.

Similarly, in Fig.3b, each line shows the sum of the runtimes of alignment computations by a method, expressed in percentage of the runtime of the *DSP* method. We use *DSP* method, which is the Dijkstra-based algorithm, as a reference because it represents the complete alignment algorithm.


For instance, we see in Fig. 3b that the runtime of the *DLM* algorithm is 20% of the runtime of the *DSP* method. For our heuristic $A^*PTD = 2$, the average runtime is around 10% of the reference method *DSP* (which corresponds to a gain of 90% of runtime), for an average quality between 90 and 85% of the optimal alignments.

We did not represent in the charts the runtimes for methods A^*SP_{mf} and A^*LM_{mf} (implemented in pm4py) since they are much higher than the others: A^*SP_{mf} ran up to 30 times longer than the *DSP* and A^*LM_{mf} up to 7 times longer. We invite the reader to find the results and scripts on github ⁶.


6.2 Influence of the Discount Parameter E on the Quality and Runtime

Fig.3a shows that the quality decreases when the parameter E of the discounted cost function raises. However the gain in term of runtimes is high when $E > 1$. This small lost of quality is correlated to the gain of runtime that we observe in Fig.3b.

⁶ available at <https://github.com/BoltMaud/A-Discounted-Cost-Function-for-Fast-Alignments-of-Business-Processes-Sources>



(a) Quality of the alignments obtained by different methods. Quality is defined as the ratio (in %) between the classical cost (number of model or log moves) of the optimal alignment and the classical cost of the alignment found by the method. The first line is the baseline and present the optimal approach which has no loss in quality, i.e., we observe 100% of quality. In general, one can see that the loss of quality with our heuristics is rather limited.



(b) Runtime (in % of the runtime of the *DSP* Algorithm). A percentage lower than 100 corresponds to a gain of runtime compared to the *DSP* method.

Fig. 3: Comparison of quality and runtimes of different methods.

The output of those experiments is the correlation between the parameter E and the compromise between quality and runtime of alignments. For $E = 1$, one gets exact alignments but runtimes are slow. For higher E one can extract very fast alignments but the quality is reduced.

In another hand, we want to raise awareness on the method $A^*PTD = 1$ which corresponds to the optimal alignment without the construction of the synchronous product. The method gives exact alignments for reduced runtimes. Method *DLM* also doesn't construct the synchronous product but we see in Fig.3a that there is a lost of quality.

6.3 Comparison with similar approaches

Additive Comparisons The omission of ProM and other tool results in the previous section is due to the differences between the output formats which made difficult the comparison of quality and runtimes. However, in this section we zoom in particular cases, i.e., by running a log sequence only, thus making human interpretation possible. We add *PNR* the results given by the PNetReplayer package of [2] in ProM and *REC_{ILP}* the results given by [16]. We present only our versions for $E = 2$ which provide fast results.




Fig. 4: Particular alignments that draws advantages and disadvantages of our methods. Like above, DSP is the baseline.

Specific Inputs We choose 3 models and traces that have particular characteristics. First, we run the alignment between the first trace of $BPI2020_{rp}$ log and model IM because this couple ($BPI2020_{rp}$, IM) gives the least differences between the methods (Fig. 4a). The model has only 2 parallelism patterns and no loop. Then, we run the alignment of the first trace of $L5$ and model $M5$ where our method specifically works well (Fig. 4b). This model contains a concurrent pattern including 28 transitions and one loop. Finally we present an alignment of the fifth trace of $L3$ which is very long (215 activities) and its model which has many loops and many parallelism patterns (Fig. 4c). The latter aims at showing the weakness of our method.

Results From the three charts of Fig. 4, we observe that all the methods usually find perfect alignments. This conclusion is done with the y-axis. The x-axis shows the runtimes and has a logarithmic scale due to the large differences between the methods. Then, the charts give an overview of the gain/loss of runtime compared to the baseline DSP versus the loss of quality. Best methods are on the top left of the charts. For $BPI2020_{rp}$ and $L5$, we see that our implementations noted $A*SPD = 2$ and $A*TD = 2$ neighbor method of ProM (noted PNR). The versions using less memory seem to be much less efficient sometimes even for less quality (see DLM for $L3$). The method REC_{ILP} worked only for the second model. $M3$ is too large for the Gurobi open source version and format of model 2020_{rp} is not accepted by the tool. Alignments of Fig. 4c highlights our weakness. Despite very fast results, the quality is only of 20% of the optimal one. This is explained by the parameter E . Due to the size of the trace (215) and thus the run, the high base of logarithm E and the loops in the model, the algorithms face a situation where all the markings have the same cost (E^n where n is very large borders zero). We launched the same input for $E = 1.1$ and our method finds the optimal alignment.

7 Conclusion

In this paper, we present a novel cost function for alignments. By using the position of the moves, our discounted cost function penalizes deviations of business processes that

appear at early stage of the process execution. While the first aim is to align prefixes first, we nicely see that the proposed discounted cost function gives a heuristic for classical alignments. We implemented two versions of an A*-based algorithm that incorporate this heuristic and we experimented with artificial and real-life logs. The outputs of the experiments clearly show that the loss of quality, in terms of log and model moves, is correlated to the gain of runtimes. This is due to the parameter E of our discounted cost function that forces prefix-first alignments. The compromise between quality and runtime can easily be set by using this parameter.

As future work, we suggest to combine the discounted cost function with other heuristics used for alignments. Also, the idea of using a discounted cost for alignments may be more or less relevant depending on the application that one is targeting. Among the multiple applications of alignments in conformance checking, some may be more or less resilient to the little loss of quality that we accept when using heuristics. In some settings, even, the alignments obtained for our discounted cost function may be more relevant than classical alignments, typically when the application justifies to penalize early deviations more than late ones.

References

1. Arya Adriansyah. *Aligning observed and modeled behavior*. PhD thesis, Technische Universiteit Eindhoven, 2014.
2. Arya Adriansyah, Boudewijn F van Dongen, and Wil MP van der Aalst. Memory-efficient alignment of observed and modeled behavior. *BPM Center Report*, 3:1–44, 2013.
3. Adriano Augusto, Raffaele Conforti, Marlon Dumas, Marcello La Rosa, and Artem Polyvyanyy. Split miner: automated discovery of accurate and simple business process models from event logs. *Knowledge and Information Systems*, 59(2):251–284, 2019.
4. Vincent Bloemen, Sebastiaan J. van Zelst, Wil M. P. van der Aalst, Boudewijn F. van Dongen, and Jaco van de Pol. Maximizing synchronization for aligning observed and modelled behaviour. In *Business Process Management - 16th International Conference, BPM 2018, Sydney, NSW, Australia, September 9-14, 2018, Proceedings*, pages 233–249, 2018.
5. Josep Carmona, Boudewijn F. van Dongen, Andreas Solti, and Matthias Weidlich. *Conformance Checking - Relating Processes and Models*. Springer, 2018.
6. Massimiliano de Leoni and Andrea Marrella. Aligning real process executions and prescriptive process models through automated planning. *Expert Syst. Appl.*, 82:162–183, 2017.
7. Mohammadreza Fani Sani, Mathilde Boltenhagen, and Wil van der Aalst. Prototype selection based on clustering and conformance metrics for model discovery. *arXiv*, 2019.
8. Wai Lam Jonathan Lee, H. M. W. Verbeek, Jorge Munoz-Gama, Wil M. P. van der Aalst, and Marcos Sepúlveda. Recomposing conformance: Closing the circle on decomposed alignment-based conformance checking in process mining. *Inf. Sci.*, 466:55–91, 2018.
9. Sander J. J. Leemans, Dirk Fahland, and Wil M. P. van der Aalst. Scalable process discovery and conformance checking. *Softw. Syst. Model.*, 17(2):599–631, 2018.
10. Sander JJ Leemans, Dirk Fahland, and Wil MP van der Aalst. Discovering block-structured process models from event logs containing infrequent behaviour. In *International conference on business process management*, pages 66–78. Springer, 2013.
11. Jorge Munoz-Gama, Josep Carmona, and Wil M. P. van der Aalst. Single-entry single-exit decomposed conformance checking. *Inf. Syst.*, 46:102–122, 2014.
12. T. Murata. Petri nets: Properties, analysis and applications. *Proceedings of the IEEE*, 77(4):541–574, April 1989.

13. Daniel Reißner, Raffaele Conforti, Marlon Dumas, Marcello La Rosa, and Abel Armas-Cervantes. Scalable conformance checking of business processes. In *OTM CoopIS*, , Rhodes, Greece, pages 607–627, 2017.
14. Daniel Reißner, Abel Armas-Cervantes, Raffaele Conforti, Marlon Dumas, Dirk Fahland, and Marcello La Rosa. Scalable alignment of process models and event logs: An approach based on automata and s-components. *Information Systems*, 94:101561, 2020.
15. Farbod Taymouri and Josep Carmona. Model and event log reductions to boost the computation of alignments. In *International Symposium on Data-Driven Process Discovery and Analysis*, pages 1–21. Springer, 2016.
16. Farbod Taymouri and Josep Carmona. A recursive paradigm for aligning observed behavior of large structured process models. In *14th International Conference of Business Process Management (BPM), Rio de Janeiro, Brazil, September 18 - 22*, pages 197–214, 2016.
17. Farbod Taymouri and Josep Carmona. Model and event log reductions to boost the computation of alignments. In Paolo Ceravolo, Christian Guetl, and Stefanie Rinderle-Ma, editors, *Data-Driven Process Discovery and Analysis*, pages 1–21. Springer International Publishing, 2018.
18. Farbod Taymouri and Josep Carmona. Computing alignments of well-formed process models using local search. *ACM Trans. Softw. Eng. Methodol.*, 29(3):15:1–15:41, 2020.
19. Wil M. P. van der Aalst. Decomposing Petri nets for process mining: A generic approach. *Distributed and Parallel Databases*, 31(4):471–507, 2013.
20. Wil M. P. van der Aalst. *Process Mining - Data Science in Action, Second Edition*. Springer, 2016.
21. Sebastiaan J. van Zelst, Alfredo Bolt, Marwan Hassani, Boudewijn F. van Dongen, and Wil M. P. van der Aalst. Online conformance checking: relating event streams to process models using prefix-alignments. *Int. J. Data Sci. Anal.*, 8(3):269–284, 2019.
22. Sebastiaan J van Zelst, Alfredo Bolt, and Boudewijn F van Dongen. Tuning alignment computation: An experimental evaluation. In *ATAED@ Petri Nets/ACSD*, pages 6–20, 2017.
23. H. M. W. Verbeek and W. M. P. van der Aalst. Merging alignments for decomposed replay. In Fabrice Kordon and Daniel Moldt, editors, *Application and Theory of Petri Nets and Concurrency: 37th International Conference, PETRI NETS 2016, Toruń, Poland, June 19-24, 2016. Proceedings*, pages 219–239. Springer International Publishing, Cham, 2016.
24. Glynn Winskel. Petri nets, algebras, morphisms, and compositionality. *Information and Computation*, 72(3):197–238, 1987.