



HAL
open science

Timed Alignments

Thomas Chatain, Neha Rino

► **To cite this version:**

Thomas Chatain, Neha Rino. Timed Alignments. 2022 4th International Conference on Process Mining (ICPM), Oct 2022, Bolzano, Italy. pp.112-119, 10.1109/ICPM57379.2022.9980687. hal-03703712

HAL Id: hal-03703712

<https://hal.science/hal-03703712>

Submitted on 4 Jul 2022

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.



Distributed under a Creative Commons Attribution 4.0 International License

Timed Alignments

Thomas Chatain

*LMF, ENS Paris-Saclay,
CNRS, Université Paris-Saclay, Inria
Gif-sur-Yvette, France*

thomas.chatain@ens-paris-saclay.fr
ORCID 0000-0002-1470-5074

Neha Rino

*LMF, ENS Paris-Saclay,
CNRS, Université Paris-Saclay, Inria
Gif-sur-Yvette, France*

neha.rino@ens-paris-saclay.fr

Abstract—The subject of this paper is to study conformance checking for timed models, that is, process models that consider both the sequence of events in a process as well as the timestamps at which each event is recorded. Time-aware process mining is a growing subfield of research, and as tools that seek to discover timing related properties in processes develop, so does the need for conformance checking techniques that can tackle time constraints and provide insightful quality measures for time-aware process models. In particular, one of the most useful conformance artefacts is the alignment, that is, finding the minimal changes necessary to correct a new observation to conform to a process model. In this paper, we set our problem of timed alignment and solve two cases each corresponding to a different metric over time processes. For the first, we have an algorithm whose time complexity is linear both in the size of the observed trace and the process model, while for the second we have a quadratic time algorithm for linear process models.

Index Terms—Conformance checking, Alignments, Timestamps, Time Petri nets

I. INTRODUCTION

A. Conformance Checking and Alignments

Process mining studies vast systems through their event logs, and seeks to extract meaningful ways to model the underlying patterns or processes that govern the behaviour of the system in order to better understand the system, or predict future behaviour [1]. Once such a process model is obtained, through machine learning or other related techniques, it is natural to ask how one is sure the obtained model is a reasonable approximation of the system’s behaviour at all, especially given the lack of explainability in the blackbox approach ML takes to producing solutions. This is where conformance checking comes into the picture, as it is the art of judging the performance of a process model by relating modelled and observed behaviour of a process to each other, without depending on the origin of the model [2]. Observed behaviour comes in the form of traces in an event log, as a sequence of events occurring during the functioning of the system, while process models are blueprints that describe what the underlying processes of any given system are supposed to look like. Measures of

how well a model reflects system behaviour include fitness, precision, generalisation, and simplicity. We often do not want the system to always precisely generate any and all possible future system behaviour, but neither should they simply regurgitate the event log and accept no new system behaviours. What is often much more useful is a process model that can, up to some small error factor, approximate any reasonable future system behaviour.

In Arya Adriansyah’s seminal thesis [3], we obtain the notion of an alignment, that is, the minimal series of corrections needed to transform an observed event trace into the execution of the process model that most closely mimics it. It is often seen as an execution of the synchronised product of the process model and a trace model i.e. a simple model that captures exactly one event log trace. This gives us a series of edits, usually insertions or deletions, that can transform a process trace into the observed trace. Alignments thereby help pinpoint exactly where inevitable deviations from expected behaviour occur, and the more distant the aligning word of a model is to its observed trace, the worse the model is at reflecting real system behaviour.

B. Time-aware Processes

Process models can be represented using a variety of formal objects, including Petri nets. They both offer a graphical means by which to represent concurrent systems, and a formal semantics for their execution, which allows one to mathematically analyse the same. Assuming the event logs are a list of words over a finite alphabet (the set of possible discrete events), the problem of calculating the alignment has been extensively studied [3] [4]. The notion of distance used on these words used is usually either Hamming distance or Levenshtein’s edit distance. It is natural to want to study explicitly timed systems as by considering events along with their timestamps when mining processes, we can study the minimum delay between two events, or the maximum duration the system takes to converge upon a state, or check deadlines, all of which are highly relevant in real world applications [5] [6] [7]. In addition, one may want to predict the timestamps of processes [8]. In the process mining community, there are ways to use existing process model notation in order to denote

Neha Rino was funded by the International Master’s Scholarships Program IDEX of Université Paris-Saclay.

some time constraints. BPMN 2.0 comes equipped with *timer events* and can record absolute, relative, and cyclical time constraints. For our purposes, we use an extension of Petri nets equipped with the ability to express the relative and cyclical time constraints that BPMN equipped with timers can. What we are referring to are time Petri nets, that is, Petri nets augmented with the ability to record and check the duration it takes to fire a transition once enabled, using which one can impose certain constraints on the relationships between the timestamps of different events.

Time-aware process mining seeks to study both what sort of underlying processes govern system behaviour, and what sort of time constraints they can impose on when certain events can occur [9] [10] [11]. The need for such a setting is clear, as specifying and checking concrete time constraints regarding the durations of events can ensure critical information about a system's behaviour. As time-aware process mining grows popular, new quality measures and conformance checking techniques must be developed that are sensitive to temporal constraints, but so far in the study of alignments as a conformance checking artefact, we notice that the process model used is never time-aware. For this, crucially, one needs to define distance functions that can meaningfully compare and separate different time processes. This paper seeks to provide a framework by which to do the very same, and set and solve the alignment problem for time-aware processes.

C. Three Distances, and Algorithms to Align them

In this paper, we propose three different distance functions over timed words, and study the alignment problem for the first two. The first is essentially Manhattan distance, and hence very standard, and has a quadratic time algorithm in the case of a structurally restricted class of models (linear causal processes). For all other types of models, we present an encoding in simplex that solves the problem, albeit with exponential worst-time complexity. The second distance function utilises the structure of the process model to study distances between time processes, and we present a straightforward algorithm for solving the alignment in this case, whose time complexity is linear in the size of the causal process and the transition set of the model. The second setting, and the new notion of delay edits that it represents, provide insight into time processes and their relevant constraints.

II. PRELIMINARIES

We represent events as pairs (a, t) where $a \in \Sigma$ is the name of the action, and t denotes the time at which said action was taken.

Definition 1. A *timed trace* is a sequence $\gamma \in (\Sigma \times \mathbb{R}^+)^*$ of *timed events*, seen as a *timed word*.

We will often ignore the untimed parts of timed words, i.e., the projection onto Σ^* , leaving just the timestamps, a sequence belonging to \mathbb{R}^{+*} .

The timed process model we use here is a labelled time Petri net.

Definition 2 (Labelled Time Petri Net). A *labelled time Petri net* (or *TPN*) is a tuple $N = (P, T, F, SI, \Sigma, \lambda, M_0, M_f)$, where P, T are disjoint sets of places and transitions respectively, $F \subseteq (P \times T) \cup (T \times P)$ is the flow relation, $SI : T \rightarrow \mathbb{I}$ is the static interval function, where $SI(t) = (Eft(t), Lft(t))$ such that Eft stands for earliest firing time, and Lft for latest firing time, $\lambda : T \rightarrow \Sigma$ is the labelling function, labelling transitions with actions from the action set Σ , and $M_0, M_f : P \rightarrow \mathbb{N}$ are the initial and final markings.

Given a transition $t \in T$ we define the pre-set of t as $\bullet t = \{p \in P \mid (p, t) \in F\}$ and its post-set similarly is defined as $t \bullet = \{p \in P \mid (t, p) \in F\}$, (the presets and post-sets of places are defined similarly). A transition t of a time Petri net is *enabled* at marking M iff $\forall p \in \bullet t : M(p) > 0$. The set of all enabled transitions at a marking M is denoted by $Enabled(M)$.

A *state* of a TPN $N = (P, T, F, SI, \Sigma, \lambda, M_0, M_f)$ is a pair $S = (M, I)$, where M is a marking of N and $I : Enabled(M) \rightarrow \mathbb{R}^+$ is called the *clock function*. The initial state is $(M_0, \mathbf{0})$, where $\mathbf{0}$ is the zero function.

A transition t is fireable from state $S = (M, I)$ after delay $\theta \in \mathbb{T}$ iff t is enabled at M , and updating the clock function to increment by delay θ will keep t 's new clock value $I(t) + \theta$ in the range $[Eft(t), Lft(t)]$ as determined by the static interval function.

Once a fireable transition t is fired, the marking and clock function are both updated to reflect the firing, as defined below :

Definition 3 (Firing Rule). When a transition t fires after time θ from state $S = (M, I)$, the new state $S' = (M', I')$ is given as follows :

$$M'' = M' \cup t \bullet \text{ where } M' = M \setminus \bullet t$$

$$I'(t) = \begin{cases} I(t) + \theta & \text{If } t \in Enabled(M') \\ 0 & \text{If } t \in Enabled(M'') \setminus Enabled(M') \\ Undefined & \text{Otherwise} \end{cases}$$

This is also denoted by $(M, I)[t](M'', I')$

An important feature of time Petri nets is the notion of urgency, that is, if t is enabled at marking M and has clock value $Lft(t)$, it must fire, or another transition must fire at the same instant disabling t .


A valid execution of the model begins at the initial marking, fires a sequence of transitions (representing a series of activities occurring at certain times) and at the end of the firing sequence reaches M_f , with any clock function I .

Definition 4 (Language of a time Petri net). A *word* $w = (a_0, a_1, \dots, a_n) \in \Sigma^*$ is in the language of the labelled time Petri net $\mathcal{L}(N)$ if there is a fireable sequence of transitions $(t_0, t_1 \dots t_n) \in T^*$ such that $\lambda((t_0, t_1, \dots, t_n)) = w$ and

they transform the initial marking into the final one, that is, for some clock function I on M_f ,

$$(M_0, \mathbf{0})[t_0, t_1, \dots, t_n)(M_f, I)$$

Example 1. Consider the following example of a time Petri net N :



One possible execution of N would be for the firing sequence

$$w = (a, 1)(b, 2)(d, 3)(e, 4)(f, 5)$$

The initial marking only has a enabled, and firing a at time 1 updates the marking by removing the token from a 's preplace and filling its two post-places, thereby enabling b, c, d, e . Now, note by urgency that c can never fire past time 2 as then b would have already fired, but in this execution this doesn't matter as b fires at 2 regardless, filling one preplace of f . Then d fires, having been enabled for 2 units of time, thereby resetting the enabling time of both d and e to zero, so that when e does fire it has only been most recently enabled for 1 unit of time. Now, finally, f is enabled, and fires after 1 unit of having been enabled.

Now we will want to build some vocabulary to help us talk about individual executions of timed words on time Petri nets.

Definition 5 (Causal Net). A causal net $CN = (B, E, G)$ is a finitary, acyclic net where

$$\forall b \in B : |b^\bullet| \leq 1 \wedge |\bullet b| \leq 1.$$

This can be viewed as the original Petri net itself, but every time a place is revisited, it is copied afresh to ensure that the execution only ever moves forward. Hence, multiple elements of B or E map to the same element in P or T respectively, a notion that is expressed via the map p defined below.

Definition 6 (Homomorphism). Let N be a time Petri net with place set P and transition set T , and $CN = (B, E, G)$ be a causal net. A mapping $p : B \cup E \rightarrow P \cup T$ is a homomorphism if $p(B) \subseteq P$, $p(E) \subseteq T$ and $\forall e \in E$ the restriction of p to $\bullet e$ is a bijection between $\bullet e$ and $\bullet p(e)$ and the restriction of p to e^\bullet is a bijection between e^\bullet and $p(e)^\bullet$ and the restriction of p to $\text{Min}(CN)$ is a bijection between $\text{Min}(CN)$ and M_0 .

Definition 7 (Causal Process). A causal process of a TPN is a pair (CN, p) where CN is a causal net and p is a homomorphism from CN to TPN.

Using p , elements of CN are identified with their corresponding *originals* in the time Petri net, and as a result, any causal process corresponds uniquely to an untimed run on the untimed version of a given time Petri net.

Definition 8 (Timing Function). A timing function $\tau : E \rightarrow \mathbb{R}$ is a function from events of a causal process into time values.

III. THE ALIGNMENT PROBLEM IN TIMED SETTINGS

The alignment problem, given a trace from an event log and a process model, involves finding a valid execution of the process model that is closest to the trace under some metric.

Definition 9 (The General Alignment Problem). Given a process model N denoted by a Time Petri Net and a timed trace σ we wish to find a timed word $\gamma \in \mathcal{L}(N)$ such that $d(\sigma, \gamma) = \min_{x \in \mathcal{L}(N)} d(\sigma, x)$ for some distance function d on timed words.

In the untimed setting, this is viewed as a problem of minimizing cost over a series of edit moves, either insertions (a model move) or deletions (a trace move). When aligning timed words, clearly there are two crucial aspects to the problem, which are, how one deals with deviations in the action labels themselves (the action labels), versus how one deals with deviations in just the timing properties of the word (the timestamps). The alignment problem has been extensively studied for the untimed case [3] [4], but the timed setting is more complex.

Example 2. Consider the process model of Example 1 and an observed trace $\sigma_1 = (a, 0)(a, 1)(b, 2)(d, 3)(e, 3)(f, 5)$. Clearly, there is an extra a event that cannot have occurred as no process trace has more than one a , and also, the e event fires too early, it should have waited at least one time unit after the firing of event d .

On the other hand, consider an observed trace of the form $\sigma_2 = (a, 1)(b, 1)(d, 3)(e, 4)(f, 5)$. Here, either we can assume the b fired too early, and correct to $(b, 2)$, or that was supposed to be a c event as those two are parallel choices.

Thirdly, consider an observed trace of the form $\sigma_3 = (a, 1)(d, 1)(d, 2)(e, 4)(f, 5)$. Clearly one of the d 's needs to be converted to a b or a c event, but which d ? The reasonable choices seem to be either correcting $(d, 1)$ to $(c, 1)$ or $(d, 2)$ to $(b, 2)$ or $(c, 2)$. Timestamps are attached to the letters they belong to, so in a sense, firing event b at time 2 as compared to firing event c at time 1 are not easily comparable, as the time constraints on different events are often different. In this manner, it is not clear how meaningful it would be to define edit moves that transform timestamps that are attached to different tasks.

This is why, we first restrict our attention to the case where the untimed part of σ does match the model, that is to say, if $\sigma = (w, \tau)$ where w is the action labels and τ is the timestamp sequence. then we assume that w

has a valid causal process (CN, p) over N . Clearly this degenerate case of the problem needs to be solved if the general timed alignment problem is to be solved, and we argue that this case is interesting and complex in its own right.

This ensures that we can compare the timing constraints on analogous parts of the process, and hence reason meaningfully about deviances from the model that are purely based on the timing of the process. This gives us the following problem :

Definition 10 (The Purely Timed Alignment Problem). *Given a process model N denoted by a Time Petri Net and a timed trace (w, σ) , and a valid causal process of the untimed part w , we wish to find a valid timing function γ such that $d(\sigma, \gamma) = \min_{x \in \mathcal{L}(N)} d(\sigma, x)$ for some distance function d on timing functions.*

Hence, for the first few sections of this paper, we will ignore the untimed parts of the words, i.e we assume the word σ we wish to align comes with a valid causal process (CN, p) of the model, and seek to find timing functions on said causal process that are both valid, and minimize distance to σ under the metric of choice.

We shall revisit the general timed alignment problem in section VI, and show how our methods can be adapted using existing techniques to provide approaches for solving the general problem.

A. Edit Moves on Timing Functions

Now we come to the problem of deciding how to compare two timing functions over the same causal process, and quantify how close they are to each other.

Much like Levenshtein's edit distance, popularly used in the untimed case of the alignment problem, we view the definition of these distances as an exercise in cost minimisation over the set of all transformations between two words. In order to formalise the same, we need to define what the valid moves of such a transformation could be. We define *moves* as functions that map one timing function over (CN, p) to another. What sort of functions are useful notions of transformation on a timed system?

Example 3. *Let us go back to example 1, and study the process model N presented there.*

Now, say we had the following words that did not fit the model, and we wished to analyse how best to modify them to fit them back into the model.

We start with $(u_1, \sigma_1) = (abdef)(1, 3, 4, 6, 6)$. One candidate for the closest valid execution would be if the timing function σ_1 were replaced by $\gamma = (1, 2, 4, 6, 6)$, and it feels reasonable to say that the cost for aligning σ_1 to this γ is 1. A way to arrive at this conclusion is by trying to execute σ 's firing sequence, and noticing that only the guard for b fails. If b were shifted to fire at 2 instead, the whole run would execute without a hitch. This sort of local, almost typographical error can often happen in systems, and it is the simplest kind to fix.

We start with $(u_1, \sigma_2) = (abdef)(1, 2, 5, 9, 12)$. The closest valid timing function would try to preserve the positions of a and b but the moment it tries to fire d , σ_2 runs late throughout. The closest we can get, intuitively, is to fire every transition in the lower branch as late as possible, giving $\gamma = (1, 2, 4, 8, 11)$. Now, when trying to compare these firing sequences, we can view it just like we did for σ_1 , as d , e and f all fire later than they should, each timestamp is moved back once, giving an aligning cost of 3. There is however, another way to see this deviation. This cascading chain of errors can be fixed if d is moved back to fire at 4, and all the relative relationships between d and its successors are preserved. This views the tasks e and f as only caring about when d ended, which makes sense, because they depend on d 's completion. This means, the switch from the timestamp series $(5, 9, 12)$ to $(4, 8, 11)$ can be viewed as only a cost 1 edit. This is a slightly more complex error to conceive of, but it reflects the fact that if a delay at the beginning caused the whole process to run late, the important thing to fix is just that initial offset, and the rest of the process will now conform to the model as needed.

Based on the above example, we naturally arrive at two types of moves.

We define a *stamp move* as a move that translates the timing function only at a point, i.e., that edits a particular element of the timestamp series τ .

Definition 11 (Stamp Move). *Given a timing function $\gamma : E \rightarrow \mathbb{R}$, formally, we define this as :*

$$\forall x \in \mathbb{R}, e \in E : \text{stamp}(x, e)(\gamma) = \gamma' \text{ where}$$

$$\forall e' \in E : \gamma'(e') = \begin{cases} \gamma(e') + x & e' = e \\ \gamma(e') & \text{otherwise} \end{cases}$$

The next type of move we describe is the more novel and interesting *delay move*. Here, we sought to leverage the structure of the process model itself, by reflecting the causal relationships the pre-order G causes. If an event is G -reachable from another, that means the first event is strictly in the causal history of the second so changing the timestamp of such a causal predecessor has consequences for all of its causal descendents, while leaving any causally unrelated (i.e, non- G -reachable) events undisturbed.

In other words, a delay move at e will preserve relative relationships between timestamps in the future, at the cost of shifting the timestamp of every causal descendent of e by the same amount.

Definition 12 (Delay Move). *Given a timing function $\gamma : E \rightarrow \mathbb{R}^n$, we define a delay move applied to it as follows :*

$$\forall x \in \mathbb{R}, e \in E : \text{delay}(x, e)(\gamma) = \gamma' \text{ where}$$

$$\forall e' \in E : \gamma'(e') = \begin{cases} \gamma(e') + x & e' \geq_G e \\ \gamma(e') & \text{otherwise} \end{cases}$$

Now, armed with these types of moves, three natural notions of distance can be constructed.

Definition 13 (Stamp Only Distance : d_t). Given any two timing functions τ_1, τ_2 over the same causal process (CN, p) , we define the stamp-only distance d_t as follows :

$$d_t(\tau_1, \tau_2) = \min\{\text{cost}(m) | m \in \text{Stamp}^*, m(\tau_1) = \tau_2\}$$


Definition 14 (Delay Only Distance : d_θ). Given any two timing functions τ_1, τ_2 over the same causal process (CN, p) , we define the delay-only distance d_θ as follows :

$$d_\theta(\tau_1, \tau_2) = \min\{\text{cost}(m) | m \in \text{Delay}^*, m(\tau_1) = \tau_2\}$$

Definition 15 (Mixed Moves Distance : d_N). And thirdly, given any two timing functions τ_1, τ_2 over the same causal process (CN, p) , we define the mixed move distance $d_N(\tau_1, \tau_2)$ as follows :

$$\min\{\text{cost}(m) | m \in (\text{Stamp} \cup \text{Delay})^*, m(\tau_1) = \tau_2\}$$

Example 4. Consider the following example :



Now for this N , let the observed trace $\sigma = (3, 4, 5) \notin \mathcal{L}(N)$.

The best d_t alignment for the example in the diagram below is $\gamma = (1, 3, 4)$ with minimum cost $d_t(\sigma, \gamma) = 4$.

The best $d_t = \theta$ alignment for the example in the diagram below is also $\gamma = (1, 3, 4)$, but this time with minimum cost $d_\theta(\sigma, \gamma) = 3$, evidenced by the move sequence $(\text{delay}(-2, 1)\text{delay}(+1, 2))$.

And lastly, the best d_N alignment for the example in the diagram below is also $\gamma = (1, 3, 4)$, the sequence of moves being one stamp and one delay move at the start, $m = \text{stamp}(-1, 1)\text{delay}(-1, 1)$, and now with minimum cost $d_N(\sigma, \gamma) = 2 < \min\{d_t(\sigma, \gamma), d_\theta(\sigma, \gamma)\}$.

IV. RESULTS AND ALGORITHM : STAMP ONLY SETTING

Lemma 1 (Stamp only distance : d_t). d_t as defined above is equivalent to the Manhattan distance or taxicab distance between two points of \mathbb{R}^n , where $|E| = n$, say $\mathbf{t} = (t_1, t_2, \dots, t_n)$ and $\mathbf{s} = (s_1, s_2, \dots, s_n)$, defined as

$$d_t(\mathbf{t}, \mathbf{s}) = \sum_{i=1}^n |t_i - s_i|$$

This lemma also allows us to restate the alignment problem for d_t , as we see below.

A. Casting Alignment as a Linear Programming Problem

The alignment problem in the stamp-only setting can be viewed as a problem of convex optimization, minimizing the stamp only cost from a fixed observed trace σ , which by Lemma 1 is $\text{cost}(\tau|_i) = \sum_{j=1}^i |\tau_j - \sigma_j|$ over the set of valid timestamp series τ for the model N . What remains is to show that the space of valid runs of the model is a convex set, and this set has been shown to be convex in [12].

Hence, we claim that using an efficient encoding of the target function using the firing domain formulation

described above, the problem can be solved using the simplex algorithm, which is used to solve convex optimization problems in good time in practice, and having worst case exponential time complexity.

Theorem 2. Given a bounded time Petri Net N , an observed trace σ , and its causal process $CN = (B, E, G)$, the alignment problem can be viewed as seeking the vector $(\gamma_1, \gamma_2, \dots, \gamma_{|E|})$ that minimizes the quantity $\sum_{e \in E} |\gamma_e - \sigma_e|$ over the set $\gamma \in \mathcal{L}(N)$.

This can be cast as a linear programming problem, i.e.,

$$\underset{\alpha - \beta \in \mathcal{C}, \alpha, \beta \geq 0}{\text{minimize}} \quad \alpha + \beta$$

Where \mathcal{C} is a convex set.

Remark 1. By the above theorem we see that the alignment problem for d_t is easily cast as a simplex instance, and can be solved by any linear programming solver. Hence it can be solved in polynomial average running time, but this approach does have theoretically exponential worst time complexity. This bound is not necessarily tight, but either improving it or proving its tightness are left for future work. While the general case seems to be hard to solve efficiently, there is a subclass of process models for which the problem is significantly more tractable, and moreover, has a quadratic time solution.

B. Summing Minimum Cost Graphs

We are given a process model N , an observed trace σ , and its underlying linear causal process $CN = (B, E, G)$ with homomorphism p for the untimed run of σ on N .

Suppose we tried to express the cost of aligning a prefix of σ to some timing function τ on the truncation of the linear causal process up to the i th transition, call it cost .

$$\text{cost}(\tau|_i) = \sum_{j=1}^i |\tau_j - \sigma_j|$$

Now, note the following equation holds for the $(i+1)$ th static interval constraint $[a, b]$:

$$\min_{\tau|_{i+1} \in \mathcal{L}^{i+1}(N)} \text{cost}(\tau|_{i+1}) =$$

$$\min_{\tau_{i+1} - \tau_i \in [a, b]} (|\tau_{i+1} - \sigma_{i+1}| + \min_{\tau|_i \in \mathcal{L}^i(N)} \text{cost}(\tau|_i))$$

This suggests that the minimal cost of aligning prefixes σ to prefixes of the model can be recursively calculated as a function of the value of the last timestamp in the alignment of a given prefix.

$$g_i(d) = \min_{\tau|_i \in \mathcal{L}^i(N) \wedge \tau_i = d} \text{cost}(\tau|_i)$$

Now, the previous equation can be expressed in terms of the g functions as follows :

$$g_{i+1}(d) = |d - \sigma_{i+1}| + \min_{d' \in [a, b]} g_i(d')$$

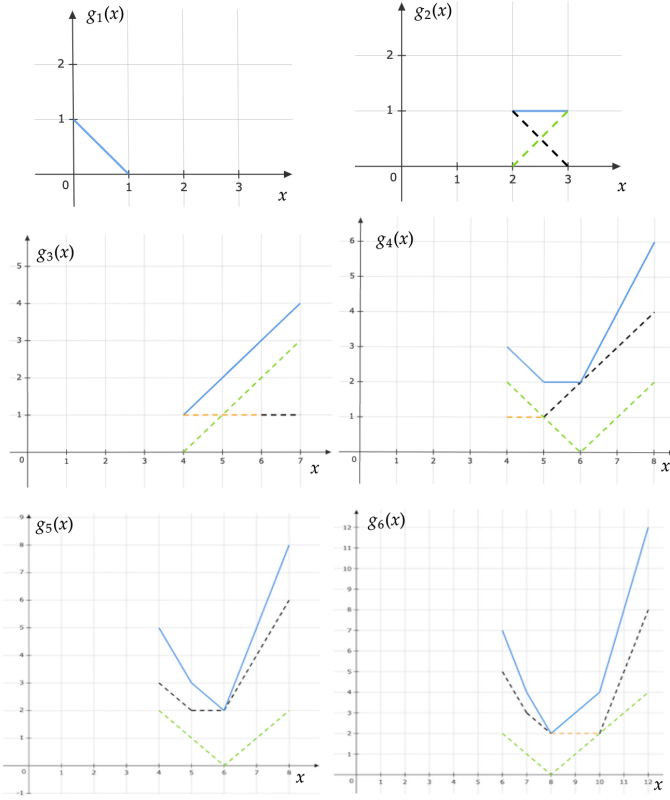



Figure 1. Visual representation of a run of the algorithm on Example 5

If we can calculate the family of functions $\forall i \leq n : g_i$ efficiently, then the minimal cost for aligning σ to N is simply $\min_d g_n(d)$

Example 5. We take a moment to look at an example of the computation of this family of functions we proposed by considering the following d_t -alignment problem :



The observed trace is $\sigma = (1, 2, 4, 6, 6, 8)$.

We begin to compute the family of functions g_i , starting with $g_1(x) = |1 - x|$ in the domain $[0, 1]$. Utilising the recursion, we proceed to compute g_2 . The new domain is $[2, 3]$, a translation of the previous one because the guard on the second transition is the point $[2, 2]$. Now, we want the sum of both $|x - 2|$ and $\min_{x' - x \in [2, 2]} g_1(x')$. The first of these functions, i.e., the contribution of the last timestamp to the error, is pictured in figure 1 in green dashed lines. The second component is pictured in black dashed lines, and is a translation of g_1 , as $\min_{x' - x \in [2, 2]} g_1(x') = g_1 - 2$.

We construct g_3 similarly, by summing the contribution

of the last timestamp, $|x - 4|$, to $\min_{x' - x \in [2, 4]} g_2(x')$. We now find an orange segment in the second component : This is used to demarcate any additional segments in $\min_{x' - x \in [a_i, b_i]} g_{i-1}(x')$ not present in g_{i-1} (if any). We see the effect of the min operation is to attach a segment with the same length as the i th static interval constraint, at the lowest point of g_{i-1} . Now we begin to see a simple way to construct these graphs visually : The next black graph is a translation of the previous blue one, with an orange line segment of the same length as the next static interval constraint inserted at the lowest point, and the next green graph is $|x - \sigma_i|$ in the new domain, and summing these two gives the next blue graph. Even summation is easier than it looks, as these are all piecewise linear graphs, so one can simply increment the slopes of all the segments of the black-orange curve by one if they're to the right of σ_i , and by minus one if they're to the left, and translate the height up as needed. Clearly the orange, black and green segments of the next graph are obtained like this, and g_4 is their sum.

The procedure for constructing g_5 and g_6 are identical, and now finally we see that the minimum cost of 2 is achieved in g_6 by setting the last timestamp to 8. Now, we know the contribution of the 8 to the error is zero, so this means we can deduce the corresponding value in g_5 must have been 2 as well, achieved when the fifth timestamp was set to 6. Now 6 gives a contribution of 0 again to the error, letting us look at the preimage of 2 once again on g_4 , and so on. We can backtrack in this manner to build an optimal alignment, such as in this case, $(0, 2, 5, 6, 6, 8)$.

With this context, we present the following algorithm that progressively calculates the graphs of g_i (represented as lists of segments, each characterised by their slopes and x-projections of endpoints) each round, and claim it is both correct and efficient. The algorithm precisely executes the above procedure, where the subfunction GRAPHMIN takes $g_i(d)$ and outputs

$$\min_{d-d' \in [a, b]} g(d') = \begin{cases} g(d-a) & l+a \leq d < m+a \\ g(m) & m+a \leq d \leq m'+b \\ g(d+b) & m'+b < d \leq h+b \end{cases}$$

Which is exactly the process of taking the last blue graph and creating the next orange-black graph, and subfunction GRAPHADDMOD takes $\min_{d-d' \in [a, b]} g(d')$ and adds $|d - \sigma_{i+1}|$ to it, which is precisely adding the green component. Lastly, the BACKTRACK algorithm backtracks through the list of graphs $\{g_i\}_n$ constructed to reverse engineer a word that aligns with the minimal cost calculated. A more detailed algorithm for backtracking is included in the appendix.

Theorem 3. Algorithm 1 is correct, i.e., given a linear causal process of a time Petri Net N , its result γ has the properties :

- 1) $\gamma \in \mathcal{L}(N)$

Algorithm 1 d_t Algorithm

Input : σ , a list of static interval constraints for the linear model SI

Output : $cost = \min_{x \in \mathcal{L}(N)} d_t(x, \sigma)$, $\gamma \in \mathcal{L}(N)$ such that $d_t(\gamma, \sigma) = cost$

Object : $graph = \{(left, slope, right) | graph[i].left = graph[i-1].right\}$

procedure STAMPOONLYALGO(σ, SI)

Initialise System

$[a, b] \leftarrow SI[i]$

$graph \leftarrow \{(a, 0, b)\}$

GRAPHADDMOD($graph, \sigma[0]$)

$graphlist \leftarrow \{graph\}$

$i \leftarrow 1$

$cost \leftarrow |a - \sigma[0]|$

while $i < |\sigma|$ **do**

$[a, b] \leftarrow SI[i]$

GRAPHMIN($graph, a, b$)

GRAPHADDMOD($graph, \sigma[i]$)

$graphlist.append(graph)$

$cost \leftarrow cost + |a - \sigma[i]|$

end while

$i \leftarrow 0$

$s \leftarrow graph[0].slope$

while $graph[i].slope < 0$ **do**

$s \leftarrow graph[i].slope$

$cost \leftarrow cost + s * (graph[i].right - graph[i].left)$

$i \leftarrow i + 1$

end while

$\gamma \leftarrow BACKTRACK(\sigma, graphlist, cost, SI)$

Return : $cost, \gamma$

end procedure

$$2) \forall x \in \mathcal{L}(N) : d_t(\gamma, \sigma) \leq d_t(x, \sigma)$$

Remark 2. Note that the subfunctions GRAPHMIN and GRAPHADDMOD both run through the list of segments of the graph once each, and hence are linear in the sizes of their inputs, and Algorithm 1 calls each of these subfunctions once for each letter of the observed trace, each time on a graph with size linear in the current prefix. So the cost calculation section of Algorithm 1 has quadratic time complexity in the size of the input.

The subfunction BACKTRACK runs a binary search on each stored graph in graphlist to find a trace in the language that does indeed give the minimum cost, letter by letter, and so overall takes $O(n \log n)$ time, keeping the overall time complexity $O(n^2)$ where $n = |\sigma|$.

V. RESULTS AND ALGORITHM : DELAYS ONLY SETTING


With delay edits, which constitute a new way of thinking transformations over time-series, comes a new perspective with which we can view timing functions over causal processes. There is of course the standard definition, $\tau : E \rightarrow \mathbb{R}^+$ that assigns to each event a timestamp that records exactly when the event occurs.

Instead, thinking along the lines of durations between events occurring, we will often benefit from considering the following representation when speaking about delay moves, defined as to view a timed word not in terms of its absolute timestamps, but by the delays between them.

Definition 16 (Flow Function). Given a causal process (CN, p) over N , where $CN = (B, E, G)$, and a (not necessarily valid) timing function $\tau : E \rightarrow \mathbb{R}^+$, we first define the flow function of τ , $f_\tau : E \rightarrow \mathbb{R}^+$ such that

$$f_\tau(e) = \begin{cases} \tau(e) & \bullet \bullet e = \emptyset \\ \tau(e) - \tau(e') & e' \in \bullet \bullet e, \\ & \tau(e') = \max_{e'' \in \bullet \bullet e} \{\tau(e'')\} \cup \{0\} \end{cases}$$

Example 6. Consider the following process model N_2 :



For this net and the word

$$w = (a, 1)(b, 2)(c, 5)(d, 6)(e, 6)(f, 8)(g, 8)$$

we can confuse transition labels with event labels, so we can define its flow function as

$$a \mapsto 1, b \mapsto 2, c \mapsto 4, d \mapsto 4, e \mapsto 4, f \mapsto 2, g \mapsto 2$$

Where the order of events is clear (such as here, where it is alphabetical, or in linear causal processes) we can write the flow function as a tuple, as follows :

$$f_w = (1, 2, 4, 4, 2, 2)$$

Note that e' as defined here is the latest causal predecessor of e . Hence, f_τ as defined produces exactly the time durations that the guards of each transition in the model checks, i.e. the clock function values during the run.

As defined, we see that if a word is in the language of the model, then its f function maps events to values that lie within the constraint that the event's corresponding transition demands, that is, it is perfectly aligned with the model with distance 0. This condition is unfortunately not sufficient (due to urgency), but it is quite close to the exact condition necessary for a word to be in the language of a time Petri net.

Also note that given the underlying causal process and the resulting f_τ , we can reconstruct τ quite straightforwardly as

$$\forall e \in E : \tau(e) = \sum_{e' \leq_G e} f_\tau(e')$$

Lemma 4 (Delay only distance : d_θ). d_θ as defined previously is equivalent to the following formulation directly using the flow function:

$$d_\theta(\tau_1, \tau_2) = \sum_{e \in E} |f_{\tau_1}(e) - f_{\tau_2}(e)|$$

The flow function is hence a dual representation of timing functions, much like line graphs for graphs, as if a timing function traditionally labels its transition nodes with timestamps, the flow function labels edges leading up to transitions with the duration since said transition was enabled. This gives us a new way to formulate the alignment problem in terms of minimising distance from the flow vector of σ .

Example 7. Let us study the process model N_2 from example 6 once again : Say we want to align

$$w = (a, 1)(b, 2)(c, 5)(d, 6)(e, 6)(f, 8)(g, 8)$$

to N_1 . We recall that we calculated its flow function to be

$$a \mapsto 1, b \mapsto 2, c \mapsto 4, d \mapsto 4, e \mapsto 4, f \mapsto 2, g \mapsto 2$$

As membership in the language only depends on the flow values, we try to pick flow components that are as close to w 's as possible while keeping the alignment in the language, i.e. minimize $\sum_{e \in E} |f_w(e) - f_u(e)|$. Choosing naively to optimise each summand gives us the flow function

$$a \mapsto 1, b \mapsto 1, c \mapsto 3, d \mapsto 6, e \mapsto 8, f \mapsto 1, g \mapsto 3$$

From which we can reconstruct the closest word

$$u = (a, 1)(b, 1)(c, 4)(d, 7)(e, 7)(f, 8)(g, 10)$$

Clearly if one tries to move any component of u 's flow function any closer to that of w , the word will no longer be in the language.

We could reason this way because d_θ allows local edits that do not affect global membership in the language of the time Petri net, as both the time Petri net and delay edits concern themselves only with the duration of time elapsed between a transition's enabling and its firing. This relationship is especially clear in extended free choice time Petri nets, where the only events that determine the fireability of an event are its causal history and events with the exact same preset which would naturally disable it by firing. Hence, we present Algorithm 2, which minimizes the delay cost at each transition locally, while still overall ensuring membership in the language of the model. The main step in the algorithm is line 12, the rest is just to ensure the time process we obtain is a valid one.

Theorem 5. Algorithm 2 is correct on extended free choice time Petri nets, that is, its result γ has the properties :

- 1) (CN, p, γ) is a valid time process of N .
- 2) $d_\theta(\gamma, \sigma) \leq d_\theta(x, \sigma)$ for all x such that (CN, p, x) is a valid time process of N .

Algorithm 2 Local d_θ Algorithm

- 1: **Input** : $N, CN = (B, E, G), p, \sigma : E \rightarrow \mathbb{R}^+$
- 2: **Output** : $\gamma : E \rightarrow \mathbb{R}^+$ such that γ is a valid timing function and $d_\theta(\gamma, \sigma) \leq d_\theta(x, \sigma)$ for all valid timing functions x
- 3: $Curr = Min(CN)$
- 4: $Enabled = \{t \mid t \subseteq p(Curr)\}$
- 5: $FD = \{(t, Eft(t), l, 0) \mid t \in Enabled, l = \min_{t' \bullet = t} Lft(t')\}$
- 6: $Soon = \{(t, eft, l, toe) \in FD \mid l + toe \text{ is minimal in } FD\}$
- 7: **while** $E \neq \emptyset$ **do**
- 8: Pick $(t, eft, l, toe) \in Soon, t \in p(E)$
- 9: $E \leftarrow E \setminus \{e \mid p(e) = t\}$
- 10: $e \leftarrow p^{-1}(t)$
- 11: $Curr \leftarrow \{Curr \setminus \{\bullet e\}\} \cup \{e \bullet\}$
- 12: $f_\gamma(e) = \operatorname{argmin}_{x \in [eft, l]} |x - f_\sigma(e)|$
- 13: **for all** $t' \in T \wedge \bullet t' = \bullet t$ **do**
- 14: $Enabled \leftarrow Enabled \setminus \{t'\}$
- 15: $FD \leftarrow FD \setminus \{(t', e', l', toe')\}$
- 16: **end for**
- 17: **for all** $t' \in T \wedge t' \bullet = t \bullet$ **do**
- 18: $Enabled \leftarrow Enabled \cup \{t'\}$
- 19: $eft' \leftarrow Eft(t')$
- 20: $l' \leftarrow \min_{t'' \bullet = t'} Lft(t'')$
- 21: $toe' \leftarrow toe + f_\gamma(e)$
- 22: $FD \leftarrow FD \cup \{(t', eft', l', toe')\}$
- 23: **end for**
- 24: $Soon = \{(t, eft, l, toe) \in FD \mid l + toe \text{ is minimal in } FD\}$
- 25: **end while**
- 26: **return** γ


Remark 3. Note that Algorithm 2 runs through the events of the causal process E exactly once, and each time runs through the set of currently enabled transitions, and the total set of transitions, and hence its time complexity is linear in the size of the input $|CN|$, and the size of the transition set, i.e. $O(|E||T|)$.

VI. THE GENERAL TIMED ALIGNMENT PROBLEM

Now that we have provided methods to tackle the purely timed alignment problem for two metrics, we return to the general timed alignment problem. One approach for solving the general timed alignment problem would be to begin by completely ignoring the timed aspect, and aligning only the untimed part of the word w and the process model N (now reduced to the underlying untimed Petri net) using the \mathbb{A}^* algorithm [3]. This will produce a word w' with a valid causal process. This new untimed word w' can then be given a timing function based on that of w , repeating or deleting timestamps for any insertion or deletion moves. This yields a valid causal process and a potentially invalid timing function over it, which is then cast as a purely timed alignment problem.

This approach is quite naive, as it may have a tendency to exaggerate the importance of action label errors. For example, consider the following scenario :

Example 8. Consider the following net N_3 where the final marking is presumed to be the sink place:



Now, let the word we seek to align be $w = (a, 100)(a, 100)(a, 100)$.

The alignment the approach we described above would provide is $w' = (a, 0)(a, 0)(a, 0)$.

But if we assign even a tenth of the cost of action edits to timestamp edits, clearly the closer word in the language was $w'' = (b, 100)(a, 100)(a, 100)$.

The issue here is the tradeoff between timestamp edits and action label edits. A better approach would be to assign a particular cost c_A to action edits and another c_T to time edits, and devise an algorithm to minimise their sum. A potential fix to our previous idea would be to essentially retain the above approach, but instead of just doing it for the best untimed alignment, take a selection of good untimed alignments (up to some allowable threshold), and try to align timestamps for each untimed candidate, and then minimise the total cost over this set. This alleviates some of the bias towards preserving action labels by allowing for action deviations in case the timestamp alignment proves particularly easy, but this is still just an initial approach to the wider problem, and we hope to find better methods in the future.

VII. IMPLEMENTATION

We have implemented the stamp only alignment algorithm in python, available at <https://github.com/NehaRino/TimedAlignments>. This algorithm is quadratic in time complexity, and runs well in practice, as seen below

Trace Length	10	100	1000
Running Time (seconds)	0.001	0.1	9.8

VIII. PERSPECTIVES AND CONCLUSION

In this paper, we posed the alignment problem for timed processes, devised three metrics with which to study alignments for timestamp sequences, and solved the purely timed alignment problem for the first two metrics proposed. As far as we know, this is the first step in conformance checking for time-aware process mining, and much further work can be inspired from this point. The alignment problem for the third metric d_N is a first future direction. Secondly, for both metrics studied here the class of models for which the alignment problem was solved efficiently are structurally restricted (being linear causal processes for d_t and extended free choice time Petri nets for d_θ) and it would be interesting to see how to broaden the scope of these methods to larger classes of process models. Thirdly, further investigation in the general timed

alignment problem is necessary, as our proposed approach here is rather rudimentary and can certainly be improved. Lastly, there are a number of other conformance artefacts that can be set and studied in the timed setting, such as anti-alignments [13], and one can better develop all such conformance checking methods to account for timed process models.

REFERENCES

- [1] W. M. P. van der Aalst, *Process Mining - Data Science in Action, Second Edition*. Springer, 2016.
- [2] J. Carmona, B. F. van Dongen, A. Solti, and M. Weidlich, *Conformance Checking - Relating Processes and Models*. Springer, 2018. [Online]. Available: <https://doi.org/10.1007/978-3-319-99414-7>
- [3] A. Adriansyah, “Aligning observed and modeled behavior,” Ph.D. dissertation, Technische Universiteit Eindhoven, 2014.
- [4] M. Boltenhagen, T. Chatain, and J. Carmona, “A discounted cost function for fast alignments of business processes,” in *BPM 2021, Proceedings*, ser. LNCS, vol. 12875. Springer, 2021, pp. 252–269. [Online]. Available: https://doi.org/10.1007/978-3-030-85469-0_17
- [5] S. Cheikhrouhou, S. Kallel, N. Guermouche, and M. Jmaiel, “The temporal perspective in business process modeling: a survey and research challenges,” *Service Oriented Computing and Applications*, vol. 9, pp. 75–85, 2014.
- [6] J. Eder, E. Panagos, and M. Rabinovich, “Time constraints in workflow systems,” in *CAiSE*, 1999.
- [7] A. Nguyen, S. Chatterjee, S. Weinzierl, L. Schwinn, M. Matzner, and B. Eskofier, *Time Matters: Time-Aware LSTMs for Predictive Business Process Monitoring*, 03 2021, pp. 112–123.
- [8] W. Aalst, H. Schonenberg, and M. Song, “Time prediction based on process mining,” *Inf. Syst.*, vol. 36, pp. 450–475, 04 2011.
- [9] A. Rogge-Solti, R. Mans, W. M. P. van der Aalst, and M. Weske, “Repairing event logs using timed process models,” in *On the Move to Meaningful Internet Systems: OTM 2013, Proceedings*, ser. LNCS, vol. 8186. Springer, 2013, pp. 705–708. [Online]. Available: https://doi.org/10.1007/978-3-642-41033-8_89
- [10] R. Conforti, M. L. Rosa, A. H. M. ter Hofstede, and A. Augusto, “Automatic repair of same-timestamp errors in business process event logs,” in *BPM 2020, Proceedings*, ser. LNCS, vol. 12168. Springer, 2020, pp. 327–345. [Online]. Available: https://doi.org/10.1007/978-3-030-58666-9_19
- [11] W. M. P. van der Aalst and L. F. R. Santos, “May I take your order? - on the interplay between time and order in process mining,” in *Business Process Management Workshops - BPM 2021 International Workshops*, ser. Lecture Notes in Business Information Processing, A. Marrella and B. Weber, Eds., vol. 436. Springer, 2021, pp. 99–110. [Online]. Available: https://doi.org/10.1007/978-3-030-94343-1_8
- [12] B. Barthomieu and M. Menasche, “A state enumeration approach for analyzing time petri nets,” 1982. [Online]. Available: <https://projects.laas.fr/tina/papers/atpn82.pdf>
- [13] T. Chatain, M. Boltenhagen, and J. Carmona, “Anti-alignments - measuring the precision of process models and event logs,” *Inf. Syst.*, vol. 98, p. 101708, 2021. [Online]. Available: <https://doi.org/10.1016/j.is.2020.101708>
- [14] S. P. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2014.
- [15] T. Aura and J. Lilius, “Time processes for time petri-nets,” in *Application and Theory of Petri Nets 1997, 18th International Conference, ICATPN '97, Proceedings*, ser. LNCS, P. Azéma and G. Balbo, Eds., vol. 1248. Springer, 1997, pp. 136–155. [Online]. Available: https://doi.org/10.1007/3-540-63139-9_34

A. Stamp Only Algorithm

We start off with a proof of Lemma 1

Proof. The way we see this is as follows. Given a causal process $CN = (B, E, G)$ and two timing functions (not necessarily constituting valid timings) τ_1, τ_2 , let

$$d_t(\tau_1, \tau_2) = \sum_{e \in E} |\tau_1(e) - \tau_2(e)|$$

This is both achievable by stamp only moves by doing the corresponding $\tau_1(e) - \tau_2(e)$ stamp move to the second word, and is the least cost any such stamp only path can take as Manhattan distance obeys the triangle inequality, and every stamp move results in a word at exactly the same Manhattan distance from its predecessor as the cost of said stamp move. \square

Now, we can proceed to the proof of Theorem 2 :

Proof. Lemma 1 (General Form Lemma) of [12] states that the firing domains of state classes for any bounded time Petri net may be expressed as solution sets of systems of inequalities of the following form :

$$\begin{cases} a_i \leq t(i) \leq b_i & \forall i \\ t(j) - t(k) \leq c_{jk} & \forall j \neq k \end{cases}$$

As $\gamma \in \mathcal{L}(N)$ this means it obeys the set of inequalities corresponding to the state class containing the final marking. Now, the above constraint is convex, so γ varies over a convex set. As σ is a constant vector, translation by it will keep the space convex, so $\gamma - \sigma \in \mathcal{C}$ for some convex set \mathcal{C} .

Now, we wish to minimize $|\gamma - \sigma|$, but it would be better to view this objective as a linear function. To this end we use a trick from [14] and use the following variable translation :

$$\alpha, \beta \geq 0 : \alpha - \beta = \gamma - \sigma$$

As defined above, α and β are the absolute values of the nonnegative and negative components of the vector $\gamma - \sigma$, respectively, so their difference is $\gamma - \sigma$ and their sum is $|\gamma - \sigma|$. Clearly now, we have completed the transformation, and our alignment problem is equivalently cast as

$$\underset{\alpha - \beta \in \mathcal{C}, \alpha, \beta \geq 0}{\text{minimize}} \quad \alpha + \beta$$

\square

Before we prove Theorem 3, we prove a small, relevant lemma.

Lemma 6. *For any convex piecewise linear function g , the following function g' is also convex and piecewise linear.*

$$g'(d) = \min_{d-d' \in [a,b]} g(d')$$

Proof. For any piecewise linear convex function g , the function

$$g'(d) = \min_{d-d' \in [a,b]} g(d')$$

is also piecewise linear. If the domain of g was $[l, h]$, the domain of g' is $[l+a, h+b]$.

Say the leftmost minimum of g (if g has multiple minima they are all on the same flat line segment) is at $m \in [l, h]$. For all $x \in [l+a, m+a]$ and $x-y \in [a, b]$ and $y \in [l, h]$ implies $y \in [x-b+a, x-a] \cap [l, h] \subseteq [l, h] \cap [l-b+am] = [l, m]$ meaning for each x , in the possible domain of y $g(y)$ is a strictly decreasing function, hence we have

$$\forall x \in [l+a, m+a] : \min_{x-y \in [a,b]} g(y) = g(x-a)$$

Now, let the rightmost minimum of g is m' . For all $x \in [m+a, m'+b]$ and $x-y \in [a, b]$ and $y \in [l, h]$ implies $y \in [x-b, x-a] \cap [l, h]$, and $[x-b, x-a] \cap [m, m'] \neq \emptyset$. Now, this means for each $x \in [m+a, m'+b]$ the domain of y has a minimum of g , so we have

$$\forall x \in [m+a, m'+b] : \min_{x-y \in [a,b]} g(y) = g(m)$$

Lastly, let $x \in [m'+b, h+b]$ and $x-y \in [a, b]$ and $y \in [l, h]$. This implies $y \in [x-b, x-a] \cap [l, h] = [m', h]$, so for each $x \in [m+a, m'+b]$ the domain of y is in a region where g is strictly increasing, so we have

$$\forall x \in [m+a, m'+b] : \min_{x-y \in [a,b]} g(y) = g(x+b)$$

So we see that for any convex piece-wise linear function g ,

$$\min_{d-d' \in [a,b]} g(d') = \begin{cases} g(d-a) & l+a \leq d < m+a \\ g(m) & m+a \leq d \leq m'+b \\ g(d+b) & m'+b < d \leq h+b \end{cases}$$

Hence g' is also convex and piecewise linear, and has at most one more linear segment than g . \square

We include here pseudocode for the BACKTRACK subfunction :

We may now prove Theorem 3 as follows :

Proof. We are given a process model N , an observed trace σ , and its underlying linear causal process $CN = (B, E, G)$ with homomorphism p for the untimed run of σ on N . As discussed earlier, for linear process models, we define

$$g_i(d) = \min_{\tau|_i \in \mathcal{L}^i(N) \wedge \tau_i = d} \text{cost}(\tau|_i)$$

This family of functions is seen to obey the following recursion:

$$g_{i+1}(d) = |d - \sigma_{i+1}| + \min_{d-d' \in [a,b]} g_i(d')$$

If we can calculate the family of functions $\forall i \leq n : g_i$ efficiently, then the minimal cost for aligning σ to N is simply

$$\min_d g_n(d)$$

Algorithm 3 Backtrack to find the best word

Input : σ , $graphlist$, $cost$, SI
Output : $\gamma \in \mathcal{L}(N)$ such that $d_t(\gamma, \sigma) = cost$
for $i = 0, i < |graphlist|$ **do**
 Find $ylist[i]$, the y values of $graphlist[i]$
end for
 $i \leftarrow |sigma| - 1$
while $i \geq 0$ **do**
 $[a, b] = SI[i + 1]$
 Truncate $ylist[i]$ and $graphlist[i]$ so $\gamma_{i+1} - x \in [a, b]$.
 $k \leftarrow \text{BINARYSEARCH}(ylist[i], cost)$
 $xsegment \leftarrow graphlist[i][k]$
 $\gamma_i \leftarrow xsegment^{-1}(cost)$
 $cost \leftarrow cost - |\gamma_i - \sigma_i|$
 $i \leftarrow i - 1$
end while

Hence to prove the correctness of Theorem 3 we only need show that the auxiliary subfunctions called do indeed do their job correctly, as the main function just implements the recursion.

To this end, we first notice that the family of functions $g_i(d)$ are convex and piecewise linear. This can be seen through induction :

Firstly $g_1(d) = |d - \sigma_1|$ so the base case is covered.

Now suppose g_{i-1} is known to be convex and piecewise linear for some $i \geq 2$. By Lemma 6 we see that $\min_{d-d' \in [a_i, b_i]} g_{i-1}(d')$ is also convex and piecewise linear, and so is $|d - \sigma_i|$ and both these properties are preserved by addition, so $g_i(d)$ must also be convex and piecewise linear.

Now that we have this result, our representation of the graph as a sequence of line segments, each represented as a triple (left x endpoint, slope, right x endpoint) is justified. The y value of any graph g_i can be back calculated because we know the value of the cost of the leftmost endpoint, it represents the word obtained by firing every transition at the earliest possible time.

Now subfunction GRAPHADDMOD is adding the point σ_i if it is inside the domain of the function, and changing the slopes of the segments before σ_i by subtracting one, and of those after σ_i by adding one, which has the effect of adding the modulus component $|d - \sigma_i|$ to the previous graph.

Subfunction GRAPHMIN is even simpler, as Lemma 6 suggests all it does is translate the whole graph forward by a , and then translate the strictly increasing portion from $[m', h]$ by an addition $b - a$, adding an extra flat segment of length $b - a$ in between to keep the two portions of the function connected, and so on input $g(d)$, $[a, b]$ it outputs

$$\min_{d-d' \in [a, b]} g(d').$$

Subfunction BACKTRACK takes the minimum cost obtained by the calculation of the g_n function, finds the value of d , i.e. γ_n for which the minimum is achieved, and subtracts the cost aligning only the last place incurs,

thereby finding the minimum cost for aligning the $n - 1$ length prefix. It then proceeds to do the same iteratively for each prefix, reverse-engineering a trace γ for which the total minimum cost of aligning to σ is achieved. \square

B. Preliminaries for Delay Only Algorithm

Now, in order to study *timed* executions, we want to be able add a timing function to the causal process we defined above, thereby allowing us to record when different transitions are taken. The following definitions, properties and theorems in this subsection cover the material developed in Aura and Lilius' article on Time Processes [15].

Given the definition of the timing function, we try to see how a time process unfolds. First the initial events, that is, those enabled at $Min(CN)$ happen, and as each event occurs, new events are enabled and disabled. In order to keep track of this, we first define the auxiliary Cut function that represents the effect of having all the events in a subset fire simultaneously (we only use Cut on sets of events that can concurrently fire), as defined below.

$$Cut(E') = (E' \bullet \cup Min(NS)) \setminus \bullet E'.$$

Now we can define the *time of enabling* for a transition t of a time Petri net in a set of conditions B' of its causal process as $TOE(B', t) =$

$$max(\{\tau(\bullet b) | b \in B' \setminus Min(CN) \wedge p(b) \in \bullet t\} \cup \{0\})$$

We must now verify that this timing function does indeed represent a valid execution that obeys all the static interval constraints, as below :

Definition 17 (Valid Timing). *A timing function τ is a valid timing of the causal process iff*

$$\forall e \in E : \tau(e) \geq TOE(\bullet e, p(e)) + Eft(p(e))$$

$$\forall e \in E : \forall t \in Enabled(p(C_e)) : \tau(e) \leq TOE(C_e, t) + Lft(t)$$

where $C_e = Cut(Earlier(e))$.

Here, checking every element of C_e might seem unnecessary but in certain nets, complex dependencies between transitions can cause causally unrelated transitions to force transitions to fire or be disabled, due to urgency. This phenomenon is known as confusion, and hence to guard against this, a timing function must check that it never leaves a section of the process behind, completing the firing or disabling of every event in $Earlier(e)$ before it can reason about the fireability of e .

Definition 18 (Time Process). *A time process of a time Petri net N is a triple (CN, p, τ) where τ is a valid timing of (CN, p) which is a causal process of N .*

A few last notions here will assist us when we consider a simpler class of time Petri nets, that allow for much easier validity checking.

Definition 19 (Extended Free Choice). *A time Petri net is extended free choice iff for all two transitions t and t' , $\bullet t \cap \bullet t' \neq \emptyset$ implies $\bullet t = \bullet t'$.*

This class of time Petri nets ensure that the net is *confusion-free*, that is, causally unrelated events cannot affect the fireability of other events. This means that in order to check for the validity of a timing function, checking all of C_e as before is no longer necessary.

How do we construct a valid timing function on the fly? Given a partial time process of a TPN, we want to be able to study the effect of firing particular transitions amongst the set of transitions currently enabled. In order to do so, we notice the following class of transitions.

A transition t is a *choice* at $B' \subseteq B$ iff B' is a co-set that maps injectively to places and $\bullet t = p(B')$, and a choice is an *extension* of the process iff $B' \subseteq \text{Cut}(E)$.

These transitions reflect exactly the transitions enabled at a particular moment in the evolution of the process. Either they must be fired or disabled, by the urgency condition.

Given the above, Aura and Lilius characterise a method by which one can build partial time processes inductively, building forward while maintaining the validity of the timing function, denoted as keeping the process complete with respect to the timing function :

Definition 20. *A causal process (CN, p) of a time Petri Net where $CN = (B, E, G)$ is said to be complete with respect to a timing function τ iff for every extension transition t of the process,*

$$\max\{\tau(e) | e \in E\} \leq \text{TOE}(\text{Cut}(E), t) + \text{Lft}(t).$$

Now given this, the following theorem due to Aura and Lilius [15] holds :

Theorem 7. *Let N be an extended free choice time Petri net and (CN, p) a causal process of N , where $CN = (B, E, G)$. A timing function τ is valid iff the following criteria hold :*

- 1) $\forall e \in E : \text{Eft}(p(e)) \leq \tau(e) - \text{TOE}(\bullet e, p(e))$
- 2) $\tau(e) - \text{TOE}(\bullet e, p(e)) \leq \min\{\text{Lft}(t) | \bullet t = \bullet p(e)\}$
- 3) (CN, p) is complete with respect to τ , i.e., for every extension transition t of the process,

$$\max\{\tau(e) | e \in E\} \leq \text{TOE}(\text{Cut}(E), t) + \text{Lft}(t)$$

C. Proof of correctness of Delay Only Algorithm

We start off with a quick proof for Lemma 4 :

Proof. By noting that delay moves on a trace τ translate exactly to stamp moves on f_τ , we can deduce by the same argument as Lemma 1 that the distance function d_θ can be seen to be identical to the Manhattan distance on flow functions, i.e., the above formulation. \square

Proof. By theorem 7 proved in Aura and Lilius' article [15], we know that in order to ensure that this is a valid

time process, we need only ensure three inequalities hold as the time process γ evolves.

- 1) $\forall e \in E : \text{Eft}(p(e)) \leq \gamma(e) - \text{TOE}(\bullet e, p(e))$
By definition $f_\gamma(e) = \gamma(e) - \text{TOE}(\bullet e, p(e))$, and by line 6 of the algorithm the variable eft is correctly initialised to store $\text{Eft}(p(e))$, and line 17 ensures the above inequality holds.
- 2) $\gamma(e) - \text{TOE}(\bullet e, p(e)) \leq \min\{\text{Lft}(t) | \bullet t = \bullet p(e)\}$
In similar vein, line 6 also ensures that l is initialised precisely to store $\min\{\text{Lft}(t) | \bullet t = \bullet p(e)\}$, and line 17 again enforces the left hand side to be less than l .
- 3) (CN, p) is complete with respect to γ , i.e., for every extension transition t of the process,

$$\max\{\gamma(e) | e \in E\} \leq \text{TOE}(\text{Cut}(E), t) + \text{Lft}(t)$$

At every point when assigning a timestamp to an event, it is ensured that it belongs to *Soon*, a set defined to contain only those enabled events that have the least value of $l + \text{toe}$, or, $\text{TOE}(\text{Cut}(E), p(e)) + \text{Lft}(p(e))$. This ensures that whenever an event e is assigned a timestamp $\gamma(e)$, for all extension transitions t (which must either be causal descendents of e or already enabled when e was picked, thereby being in *Enabled*, by extended free choice) we know that $\gamma(e) \leq \text{TOE}(\text{Cut}(E), p(e)) + \text{Lft}(p(e)) \leq \text{TOE}(\text{Cut}(E), t) + \text{Lft}(t)$.

Hence, (CN, p, γ) is a valid time process of N .

As for its optimality, we see that any other valid time process would have to obey the same inequalities, in particular for another valid timing function γ' it would have to obey

$$f_{\gamma'}(e) \in [\text{Eft}(p(e)), \min_{\bullet t' = \bullet p(e)} \text{Lft}(t')] = [a, b]$$

and hence at any event e where γ differs from γ' it would have lower than or equal cost for that event as

$$f_\gamma(e) = \operatorname{argmin}_{x \in [a, b]} |f_\sigma(e) - x|$$

Hence, the algorithm is correct. \square

D. A Note on Linear Causal Processes


We take a moment to discuss the structurally restricted class of processes that our stamp only-algorithm works for, that is, causal processes whose graphical structure is that of a straight line. These reflect the executions of time Petri nets that do not have any branching points ($e \in E : |e^\bullet| > 1$) and naturally as a consequence, no points of synchrony ($e \in E : |\bullet e| > 1$). This essentially reflects the executions of automata, allowing for exclusive branching ($p \in P : |p^\bullet| > 1$) and iteration (that is, cycles in the model) only. It loses the ability to capture properties of a concurrent nature, as the run never splits into more than one token. On the other hand, studying the language of this restricted class and the alignment problem over it

becomes substantially simpler, the G pre-order becomes total, and so to begin with, this is a significantly more tractable class of time processes.


Example 9. Going back to the net N in example 1, and the firing sequence we analysed then

$$w = (a, 1)(b, 2)(d, 3)(e, 4)(f, 5)$$

We now construct the causal process of this execution, and see that it is indeed branching, as the very first transition, a , itself has two post-places, which violates the no branching condition $e \in E : |e^\bullet| \leq 1$.



Example 10. On the other hand, consider the following time Petri net N_1 :



Now, as all the branching and joining happens at places rather than transitions, its executions all have linear causal processes, such as the above causal process for the untimed word $w = aabcea$.

This causal process is unrolled from the Petri net, exactly the way runs of finite state automata are simple paths over the graph of the automaton. We throughout assume that wherever needed, such a causal process can be obtained efficiently.