

Handling of Past and Future with Phenesthe+

Manolis Pitsikalis Alexei Lisitsa
Patrick Totzke

Department of Computer Science, University of Liverpool, United Kingdom

{e.pitsikalis, a.lisitsa, totzke}@liverpool.ac.uk

Writing temporal logic formulae for properties that combine instantaneous events with overlapping temporal phenomena of some duration is difficult in classical temporal logics. To address this issue, in previous work we introduced a new temporal logic with intuitive temporal modalities specifically tailored for the representation of both instantaneous and durative phenomena. We also provided an implementation of a complex event processing system, Phenesthe, based on this logic, that has been applied and tested on a real maritime surveillance scenario.

In this work, we extend our temporal logic with two extra modalities to increase its expressive power for handling future formulae. We compare the expressive power of different fragments of our logic with Linear Temporal Logic and dyadic first-order logic. Furthermore, we define correctness criteria for stream processors that use our language. Last but not least, we evaluate empirically the performance of Phenesthe+, our extended implementation, and show that the increased expressive power does not affect efficiency significantly.

1 Introduction

Temporal logics are widely used in many domains as they allow the formalisation of time dependent properties. For example in philosophy temporal logics can be used to reason about issues involving the temporal domain [28], in computer science and specifically, in monitoring, temporal logics are used to specify and monitor specific properties of a system [21, 6], in complex event processing or recognition [12, 3, 7]—which is the focus of this work—temporal logics are used for specifying temporal phenomena and detecting them in streams of information. Naturally, each temporal logic comes with its own focus and limitations.

The starting point for many monitoring systems is Linear Temporal Logic (LTL) [24], where formulae are interpreted over single event sequences. This makes it difficult to incorporate concurrent activities such as, for instance, those carried out by a chef following some recipe to create a meal (see Figure 1). They may prepare multiples dishes of the same course in parallel, however the preparation of each dish happens on different overlapping intervals. This is difficult to formalise with logics that talk about single traces of events. Temporal logics that allow the representation of concurrent activities directly are those of Halpern and Shoham (HS) [16] and Allen’s Algebra [2]. Both logics are interpreted over sets of (possibly overlapping, discrete) time intervals, and model instantaneous events via point intervals ($[t,t]$).

In previous work [22], we introduced a temporal logic similar to those of Halpern/Shoham and Allen, which is interpreted over separate (sets of) time intervals and instantaneous events, which was specifically designed for a maritime application domain. It allows to easily specify concurrent activities and related start and endpoints. Deliberately absent in our logics (and related prior ones) are explicit negations/complements of formulae that hold on overlapping intervals. We implemented an event processing



Figure 1: Example of instantaneous and durative temporal phenomena. *turnOn/Off* is instantaneous and true when the stove is turned on/off, *stoveIsOn* is durative and holds when the stove is on, and finally *preparingDish* is also durative and holds when a dish is being prepared.

system¹ and evaluated it on a real maritime monitoring scenario [23]—whereby maritime experts authored maritime phenomena of interest—proving that our system is capable of producing phenomena detections in real time.

In this work, we extend our language with two extra temporal modalities: minimal range (\leadsto), and interval filtering (*filter*). The minimal range operator allows to capture durative temporal phenomena that start at the latest occurrence of a starting condition before the stopping condition, for example, the last working period of machine since someone operated it until it broke down. Filtering is very important for specifying the duration constraints of a temporal phenomenon, e.g., a steak is cooked rare if its on a hot pan for approximately 90 seconds on each side. Both examples are impossible to formalise in the original version of our temporal logic. It is evident that \leadsto has similar semantics to the ‘until’ operator of LTL. However, formulae that utilise ‘until’ are true on instants of time, while formulae that utilise ‘ \leadsto ’ are true on intervals. While in terms of expressive power the two operators are similar, \leadsto , in practice, allows efficient computations and more concise formulae. Similarly, in the case of filtering, writing an LTL formula for filtering periods based on some fixed threshold is possible, however the formula is not trivial and its length depends the threshold. Comparing LTL and our temporal logic, we show that the fragment of our original language that is comparable with LTL was at most as expressive as the pure past fragment of LTL, while the same fragment but with the addition of the ‘ \leadsto ’ and ‘*filter*’ operators has equal expressive power to LTL. Concerning our full temporal logic, a comparison with LTL is impossible since the structures on which their respective semantics are based, are incomparable, however we show that our language is expressible in Dyadic First Order logic (DFO) and is strictly less expressive. As expected, including temporal modalities that involve the future requires additional steps for complex event processing. In order to guarantee that our extended implementation, Phenesthe+, is correct, inspired from runtime monitoring and verification [1, 5], we define criteria for *proper* stream processors of our language and discuss how Phenesthe+ conforms to them. Finally, we illustrate through experimental evaluation that the efficiency of Phenesthe+, is not significantly compromised. Therefore the contributions of this paper are:

- We extend the expressiveness of our language for representing “look ahead happenings”,
- We formally study the expressive power of the temporal logic introduced in [22] and its extension,
- We define criteria for *proper* stream processors utilising our temporal logic,
- We showcase that our stream processing engine is capable of performing real-time complex event processing by adopting a maritime surveillance use-case.

The paper is organised as follows. First in Section 2 we describe our temporal logic. Next, in Section 3 we illustrate through examples inspired by the maritime domain the usage of the new temporal modalities.

¹<https://manospits.github.io/Phenesthe/>

In Section 4 we study the expressive power of our temporal logic. Then, in Section 5 we describe the requirements a stream processor should satisfy for processing formulae of our language, while in Section 6 we empirically evaluate Phenesthe+. Finally, in Section 7 we present work related to ours, summarise, and discuss further directions.

2 The Language of Phenesthe

The key components of our language are instantaneous events, durative disjoint states and durative, possibly non disjoint, dynamic temporal phenomena. In what follows, ‘temporal phenomena’ includes all of the three aforementioned categories.

Syntax. Formally, our Temporal Phenomena Definition Language (*TPhL*) is described by the triplet $\langle \mathcal{P}^{esd}, L, \Phi \rangle$, where \mathcal{P}^{esd} is a predicate set defined by the union of the event, state or dynamic temporal phenomenon predicates sets (in symbols $\mathcal{P}^{e/s/d}$ resp.); L is a set defined by the union of the set of the logical connectives $\{\wedge, \vee, \neg, \in\}$, the set of temporal operators, $\{\dashv, \curlywedge, \sqcup, \sqcap, \setminus, \text{filter}_{\square}\}$ where the ‘ \square ’ symbol may be one of the following symbols $\{<, \geq, =\}$, the set of temporal relations $\{\text{before, meets, overlaps, finishes, starts, equals, contains}\}$ and finally the set of the $\{\text{start, end}\}$ operators; Φ is the set of formulae defined by the union of the formulae sets Φ^* , Φ^- and $\Phi^=$. We assume that the set of predicate symbols includes those with atemporal and fixed semantics, such as arithmetic comparison operators etc., however for simplification reasons in what follows we omit their presentation. Formulae of Φ^* describe instantaneous temporal phenomena, and formulae of Φ^- describe durative temporal phenomena that hold (are true) in disjoint maximal intervals, finally formulae of $\Phi^=$ describe durative temporal phenomena that may hold in non-disjoint intervals. Figure 1 shows an example of an event (*turnOn/Off*), a state (*stoveIsOn*) and a dynamic temporal phenomenon (*preparingDish*).

Therefore given a set of event, state and dynamic phenomena predicates \mathcal{P}^{esd} the formulae of *TPhL* are defined as follows:

$$\begin{aligned}
\phi &:= \phi^* \mid \phi^- \mid \phi^= \\
\phi^* &:= \mathcal{P}^e(a_1, \dots, a_k) \mid \neg \phi^* \mid \phi^* [\wedge, \vee] \phi^* \mid \text{start}(\phi^-) \mid \text{end}(\phi^-) \mid \phi^* \in \phi^- \\
\phi^- &:= \mathcal{P}^s(a_1, \dots, a_k) \mid \phi^* [\dashv, \curlywedge] \phi^* \mid \phi^- [\sqcup, \sqcap, \setminus] \phi^- \mid \phi^- \text{filter } n \text{ (where } n \in \mathbb{N} \cup \{\infty\}) \\
\phi^= &:= \mathcal{P}^d(a_1, \dots, a_k) \mid [\phi^-, \phi^=] [\text{meets, overlaps, equals}] [\phi^-, \phi^=] \\
&\quad \mid [\phi^*, \phi^-, \phi^=] [\text{starts, finishes}] [\phi^-, \phi^=] \\
&\quad \mid [\phi^-, \phi^=] \text{contains} [\phi^*, \phi^-, \phi^=] \\
&\quad \mid [\phi^*, \phi^-, \phi^=] \text{before} [\phi^*, \phi^-, \phi^=]
\end{aligned}$$

where a_1, \dots, a_k correspond to terms denoting atemporal properties.

Semantics. We assume time is discrete and represented by the natural numbers $T = \mathbb{N}$ ordered via the ‘ $<$ ’ relation. In what follows we assume that for all models discussed in this paper time is represented by \mathbb{N} . For the formulae sets Φ^* , Φ^- and $\Phi^=$ we define the model $\mathcal{M} = \langle T, I, <, V^*, V^-, V^= \rangle$ where $V^* : \mathcal{P}^e \rightarrow 2^T$, $V^- : \mathcal{P}^s \rightarrow 2^I$, $V^= : \mathcal{P}^d \rightarrow 2^I$ are valuations, and $I = \{[ts, te] : ts < te \text{ and } ts, te \in T\} \cup \{[ts, \infty) : ts \in T\}$ is the set of time intervals of T . Intervals of the form $[ts, \infty)$ denote that a phenomenon started being true at ts , and continues being true forever. Intervals of the form $[ts, te]$ denote that a phenomenon started being true at ts and stopped being true at te . In what follows, we will use the abbreviated version for bounded quantifiers, i.e., $\forall_{z_1}^{<z_2} x(\dots)$ denotes $\forall x (x > z_1 \wedge x < z_2) \rightarrow (\dots)$, and $\exists_{z_1}^{<z_2} x(\dots)$ denotes $\exists x (x > z_1 \wedge x < z_2) \wedge (\dots)$.

Given a model \mathcal{M} , the validity of a formula $\phi \in \Phi^+$ at a timepoint $t \in T$ (in symbols $\mathcal{M}, t \models \phi$) is determined by the rules below, starting with the boolean connectives.

- $\mathcal{M}, t \models P^e(a_1, \dots, a_n)$ iff $t \in V^*(P^e(a_1, \dots, a_n))$.
- $\mathcal{M}, t \models \neg\phi$ iff $\mathcal{M}, t \not\models \phi$.
- $\mathcal{M}, t \models \phi[\wedge, \vee]\psi$ iff $\mathcal{M}, t \models \phi$ [and, or] $\mathcal{M}, t \models \psi$.

Next we define the semantics for start, end and \in which allow interaction between formulae of Φ^+ and Φ^- via the starting, ending, and intermediate points of intervals at which Φ^- formulae hold.

- $\mathcal{M}, t \models \text{start}(\phi)$ iff $\exists te. \mathcal{M}, [t, te] \models \phi$ or $\mathcal{M}, [t, \infty) \models \phi$, where $\mathcal{M}, [t, te] \models \phi$ denotes the validity of a formula $\phi \in \Phi^-$ at an interval $[t, te]$ as defined below.
- $\mathcal{M}, t \models \text{end}(\phi)$ iff $\exists ts. \mathcal{M}, [ts, t] \models \phi$.
- $\mathcal{M}, t \models \phi \in \psi$ iff $\mathcal{M}, t \models \phi$ and $\exists^{<t} ts. \exists_{\geq t} te. \mathcal{M}, [ts, te] \models \psi$.

Given a model \mathcal{M} , the validity of a formula $\phi \in \Phi^-$ at a time interval $i = [ts, te] \in I$ (in symbols $\mathcal{M}, [ts, te] \models \phi$) is defined as follows. We start with the \curvearrowright and \curvearrowleft operators which allow specifying minimal or maximal intervals between instants where formulae of Φ^+ are true.

- $\mathcal{M}, i \models P^s(a_1, \dots, a_n)$ iff $i \in V^-(P^s(a_1, \dots, a_n))$.
- $\mathcal{M}, [ts, te] \models \phi \curvearrowright \psi$ iff $\mathcal{M}, ts \models \phi$ and $\mathcal{M}, te \models \psi \wedge \neg\phi$ and $\forall_{>ts}^{<te} t. [\mathcal{M}, t \not\models \phi \text{ and } \mathcal{M}, t \not\models \psi \wedge \neg\phi]$. Therefore, $\phi \curvearrowright \psi$ holds for the intervals that start at the latest instant ts at which ϕ is true and end at first instant te after ts where $\psi \wedge \neg\phi$ is true.
- $\mathcal{M}, [ts, te] \models \phi \curvearrowleft \psi$ iff $\mathcal{M}, ts \models \phi$ and $\mathcal{M}, te \models \psi \wedge \neg\phi$ and $\forall_{>ts}^{<te} t. \mathcal{M}, t \not\models \psi \wedge \neg\phi$ and $\forall^{<ts} ts'. \mathcal{M}, ts' \models \phi \rightarrow \exists_{>ts'}^{<ts} te'. \mathcal{M}, te' \models \psi \wedge \neg\phi$.

Essentially, $\phi \curvearrowleft \psi$ holds for the disjoint maximal intervals that start at the earliest instant ts where ϕ is true and end at the earliest instant te where ψ is true and ϕ is false.

- $\mathcal{M}, [ts, \infty) \models \phi \curvearrowleft \psi$ iff $\mathcal{M}, ts \models \phi$ and $\forall_{>ts} t. \mathcal{M}, t \not\models \psi \wedge \neg\phi$ and $\forall^{<ts} ts'. \mathcal{M}, ts' \models \phi \rightarrow \exists_{>ts'}^{<ts} te'. \mathcal{M}, te' \models \psi \wedge \neg\phi$.

Therefore a formula $\phi \curvearrowleft \psi$ may hold indefinitely if there does not exist an instant after ts at which $\psi \wedge \neg\phi$ is satisfied. For simplification reasons in the semantics below we omit intervals open at the right to infinity since they can be treated in a similar manner.

We continue with the definition of semantics for \sqcup, \sqcap and \setminus which correspond to the usual set operations but for time intervals.

- $\mathcal{M}, [ts, te] \models \phi \sqcup \psi$ iff²
 - exists a sequence of length $k > 1$ of intervals $i_1, \dots, i_k \in I$ where $i_k = [ts_k, te_k]$, $ts = ts_1$ and $te = te_k$ such that:
 1. $\forall \alpha \in [1, k-1]: te_\alpha \in i_{\alpha+1}, ts_\alpha < ts_{\alpha+1}$ and $te_\alpha < te_{\alpha+1}$,
 2. $\forall \beta \in [1, k]: \mathcal{M}, [ts_\beta, te_\beta] \models \phi$ or $\mathcal{M}, [ts_\beta, te_\beta] \models \psi$, and
 3. $\nexists i_\gamma = [ts_\gamma, te_\gamma] \in I - \{i_1, \dots, i_k\}$ where $\mathcal{M}, [ts_\gamma, te_\gamma] \models \phi$ or $\mathcal{M}, [ts_\gamma, te_\gamma] \models \psi$ and $ts_1 \in i_\gamma$ or $te_k \in i_\gamma$
 - or, $\mathcal{M}, [ts, te] \models \phi$ or $\mathcal{M}, [ts, te] \models \psi$ and $\nexists i_\gamma = [ts_\gamma, te_\gamma] \in I - \{[ts, te]\}$ where $\mathcal{M}, [ts_\gamma, te_\gamma] \models \phi$ or $\mathcal{M}, [ts_\gamma, te_\gamma] \models \psi$ and $ts \in i_\gamma$ or $te \in i_\gamma$.

²A first order definition of the semantics of \sqcup is possible but more lengthy.

For a sequence of intervals, conditions (1-2) ensure that intervals, at which ϕ or ψ are valid, overlap or touch will coalesce, while condition (3) ensures that the resulting interval is maximal. In the case of a single interval, the conditions ensure that at the interval $[ts, te]$ ϕ or ψ is valid, and that $[ts, te]$ is maximal. In simple terms, the temporal union $\phi \sqcup \psi$ holds for the intervals where at least one of ϕ or ψ hold. The above definition of temporal union follows the definitions of temporal coalescing presented in [8, 14].

- $\mathcal{M}, [ts, te] \models \phi \setminus \psi$ iff $\exists [ts', te'] \in I$ where $\mathcal{M}, [ts', te'] \models \phi$, $[ts, te] \subseteq [ts', te']$ (i.e., $[ts, te]$ subinterval of $[ts', te']$), $\forall [ts_\psi, te_\psi] \in I$ where $\mathcal{M}, [ts_\psi, te_\psi] \models \psi$, $[ts, te] \cap [ts_\psi, te_\psi] = \emptyset$ and finally $[ts, te]$ is maximal. In plain language, the temporal difference of formulae ϕ, ψ holds for the maximal subintervals of the intervals at which ϕ holds but ψ doesn't hold.
- $\mathcal{M}, [ts, te] \models \phi \sqcap \psi$ iff $\exists [ts_\phi, te_\phi], [ts_\psi, te_\psi] \in I$ where $\mathcal{M}, [ts_\phi, te_\phi] \models \phi$, $\mathcal{M}, [ts_\psi, te_\psi] \models \psi$ and $\exists [ts, te] \in I$ where $[ts, te] \subseteq [ts_\phi, te_\phi]$, $[ts, te] \subseteq [ts_\psi, te_\psi]$ and $[ts, te]$ is maximal. In other words, the temporal intersection of two formulae of Φ^- holds for the intervals at which both formulae hold.

We finish the semantics for formulae of Φ^- , with the semantics of the filter operator, which allows specifying constraints on the length of intervals at which formulae of Φ^- hold.

- $\mathcal{M}, [ts, te] \models \phi \text{ filter}_{\{<, \geq, =\}} n$ iff $\mathcal{M}, [ts, te] \models \phi$ and $te - ts \{<, \geq, =\} n$.

Due to space limitations and for this part only we will adopt point intervals to refer to instants. This will allow us to define the semantics for formulae of $\Phi^=$ without specifying different rules for involved sub-formulae of $\Phi^=$. In other words given a $\phi \in \Phi^=$ we will denote the satisfaction relation $\mathcal{M}, t \models \phi$ as $\mathcal{M}, [t, t] \models \phi$. Given a model \mathcal{M} , the validity of a formula $\phi \in \Phi^=$ at a time interval $[ts, te] \in I$ (in symbols $\mathcal{M}, [ts, te] \models \phi$) is defined as follows:

- $\mathcal{M}, [ts, te] \models P^d(a_1, \dots, a_n)$ iff $[ts, te] \in V^=(P^d(a_1, \dots, a_n))$.
- $\mathcal{M}, [ts, te] \models \phi \text{ before } \psi$ iff $\exists te'. \exists_{>te'} ts'. [\mathcal{M}, [ts, te'] \models \phi \text{ and } \mathcal{M}, [ts', te] \models \psi \text{ and } \forall ts''. \forall_{>te'}^{<ts'} te''.$
 $\mathcal{M}, [ts'', te''] \not\models \phi \text{ and } \forall_{>te'}^{<ts'} ts''. \forall te''. \mathcal{M}, [ts'', te''] \not\models \psi]$. In our approach the 'before' relation holds only for intervals where the pair of instants or intervals at which the participating formulae are true or hold, are contiguous. For example, for the intervals $[1, 2]$, $[1, 3]$ and $[5, 6]$ only $[1, 3]$ is before $[5, 6]$. We chose to limit the intervals satisfying the before relation, as in practice it is usually the case that the interval directly before another one is required for specifying a dynamic phenomenon.
- $\mathcal{M}, [ts, te] \models \phi \text{ meets } \psi$ iff $\exists t. \mathcal{M}, [ts, t] \models \phi \text{ and } \mathcal{M}, [t, te] \models \psi$.
- $\mathcal{M}, [ts, te] \models \phi \text{ overlaps } \psi$ iff $\exists_{>ts}^{<te} ts'. \exists_{>ts'}^{<te} te'. [\mathcal{M}, [ts, te'] \models \phi \text{ and } \mathcal{M}, [ts', te] \models \psi]$.
- $\mathcal{M}, [ts, te] \models \phi \text{ finishes } \psi$ iff $\exists_{>ts}^{<te} ts'. [\mathcal{M}, [ts', te] \models \phi \text{ and } \mathcal{M}, [ts, te] \models \psi]$.
- $\mathcal{M}, [ts, te] \models \phi \text{ starts } \psi$ iff $\exists_{>ts}^{<te} te'. [\mathcal{M}, [ts, te'] \models \phi \text{ and } \mathcal{M}, [ts, te] \models \psi]$.
- $\mathcal{M}, [ts, te] \models \phi \text{ equals } \psi$ iff $\mathcal{M}, [ts, te] \models \phi \text{ and } \mathcal{M}, [ts, te] \models \psi$.
- $\mathcal{M}, [ts, te] \models \phi \text{ contains } \psi$ iff $\mathcal{M}, [ts, te] \models \phi \text{ and } \exists_{>ts} ts'. \exists_{<te} te'. \mathcal{M}, [ts', te'] \models \psi$.

3 Examples of maritime properties expressed in *TPhL*

We demonstrate the usability of *TPhL* and the new temporal modalities by adopting a maritime monitoring scenario. When it comes to maritime surveillance there are several resources available; for example the Automatic Identification System (AIS) allows the transmission of timestamped positional and ancillary data from vessels, maritime areas in the form of polygons can be used for producing vessel-area

relations and so on. Similar to [22], we assume the input consists of AIS messages along with spatial events relating vessels to areas of interest e.g., port areas, fishing areas and so on. Therefore our task here involves detecting maritime phenomena of interest i.e., the instants, time periods at which they are true over a maritime input stream. Below we formalise some maritime temporal phenomena³ that utilise the new temporal modalities ($\text{q}\rightarrow$ and filter).

Fishing warning. Illegal fishing is a very important issue. Vessels engaged in illegal fishing typically declare fake ship-types. Consider the formalisation below for detecting suspicious stops in fishing areas.

$$\begin{aligned} &\text{state_phenomenon } \textit{fishing_warning}(V, F) : \\ &((\textit{in_fishing_area}(V, F) \wedge \neg \textit{vessel_type}(V, \textit{fishing})) \sqcap \textit{stopped}(V)) \text{ filter}_{\geq 600}. \end{aligned}$$

state_phenomenon is a keyword for declaring the phenomenon type, $\textit{in_fishing_area}$ is a user defined state that holds for the time periods a vessel V is within a fishing area F , while $\textit{stopped}$ is a state that holds for the time periods a vessel is stopped. Finally, $\textit{vessel_type}(V, T)$ is an atemporal predicate that is true when vessel V has type T . Therefore, a vessel performs a $\textit{fishing_warning}$, if it is not a fishing vessel, and it is stopped within a fishing area for a period longer than 10 min (600 sec). Here, filtering is used for minimising false detections occurring from AIS errors (e.g., zero speed) or normal activities.

Port waiting time. Monitoring the waiting time of vessels since they entered a port and until they get moored is highly useful for various operational and logistical reasons (e.g., efficient planning of resources). However, some vessels may enter and leave a port without mooring—due to weather conditions for example. We formalise port waiting time below.

$$\begin{aligned} &\text{state_phenomenon } \textit{waiting_time}(V, P) : \\ &\text{start}(\textit{in_port}(V, P)) \text{ q}\rightarrow \text{start}(\textit{moored}(V, P)). \end{aligned}$$

$\textit{in_port}$ is a state that holds when a vessel is in a port, while \textit{moored} is a state that holds when a vessel is moored at a port. Note that the left and right arguments of $\text{q}\rightarrow$ are formulae of Φ^* , therefore if these formulae were used in other definitions we could have defined corresponding events. Consequently, the $\textit{waiting_time}$ state holds for the minimal periods between the time a vessel enters a port and the time the vessel starts being moored. Here we are interested in the minimal period, as we want to detect only the cases where a vessel entered a port and got moored.

4 Expressiveness

In this section we study the expressive power of our language. We consider three syntactic fragments of $TPhL$. The first one, denoted as $TPhL_o^-$, corresponds to the original version of the language (w/o $\text{q}\rightarrow$, filter) and excluding formulae of $\Phi^=$ (recall that $\Phi^=$ formulae hold on possibly non-disjoint intervals). The second is $TPhL^-$, which is the same as $TPhL_o^-$ but includes $\text{q}\rightarrow$ and filter, while the third, $TPhL$ corresponds to the complete language. Figure 2 (left) illustrates the syntactic relation between $TPhL$, $TPhL^-$ and $TPhL_o^-$. In more detail, we will show that $TPhL_o^-$ is equally expressive as pure past LTL , $TPhL^-$ has equal expressive power to LTL , and finally $TPhL$ is strictly less expressive than DFO . The relations in terms of expressive power between the different language fragments are illustrated in Figure 2 (right).

³The complete set of definitions is available in our online repository <https://github.com/manospits/Phenesthe/tree/future>.




Figure 2: Syntactic relation between the fragments of $TPhL$ (left). Expressive relations between different fragments of $TPhL$, LTL and DFO (right). A fragment A is strictly more expressive from a fragment B if they are connected via $A \rightarrow B$. Double lined edges denote equal expressive power.

4.1 Preliminaries

Before we continue with our analysis, as a reminder we present the syntax of LTL with past, and First Order Monadic Logic of Order (*FOMLO*).

LTL. The formulae of $LTL[XU, YS]$, given a set of propositions \mathcal{P} are defined as follows:

$$\phi ::= \perp \mid p \mid \neg\phi_1 \mid \phi_1 \wedge \phi_2 \mid X\phi_1 \mid \phi_1 U \phi_2 \mid Y\phi_1 \mid \phi_1 S \phi_2;$$

where X , and U stand for the next and until modalities, while Y , and S stand for previous and since. The formulae of $LTL[XU, YS]$ are interpreted over a discrete, linear model of time, formally represented as $\mathcal{M}_{TL} = \langle T, <, V^{TL} \rangle$, where T is equal to \mathbb{N} , $<$ is the linear order and $V^{TL} : \mathcal{P} \rightarrow 2^T$ is the interpretation function, mapping each proposition to a set of time instants. The satisfaction relation, i.e., that a formula ϕ is true at t , is defined as $\mathcal{M}_{TL}, t \models \phi$. The semantics of LTL are defined as usual; more specifically in what follows we assume the reflexive⁴ semantics of S and U . We denote the pure past fragment of LTL i.e., LTL without X and U as $LTL[YS]$.

FOMLO. Given a countable set of variables x, y, z, \dots , the formulae of *FOMLO* over a set of unary predicate symbols Σ are defined a follows:

$$\begin{aligned} \text{atomic} &::= x < y \mid x = y \mid P(x) \text{ (where } P \in \Sigma) \\ \phi &::= \text{atomic} \mid \neg\phi_1 \mid \phi_1 \vee \phi_2 \mid \phi_1 \wedge \phi_2 \mid \exists x. \phi_1 \mid \forall x. \phi_1 \end{aligned}$$

We interpret *FOMLO* formulae over structures of the form $\mathcal{M}_{FO} \langle T, <, V^{FO} \rangle$, where T is equal to \mathbb{N} , ' $<$ ' is the linear order while $V^{FO} : \Sigma \rightarrow 2^T$ is the interpretation of Σ . $\mathcal{M}_{FO}, t_1, t_2, \dots, t_n \models \phi(x_1, x_2, \dots, x_n)$ denotes the satisfaction of a formula ϕ with free variables x_1, x_2, \dots, x_n when they are interpreted as elements t_i of \mathcal{M}_{FO} . The semantics of the formulae are defined as usual (see for example [25]). We also define the $FOMLO^-$ fragment of *FOMLO*. Syntactically a formula with one free variable $\phi(x)$, is a formula of the fragment if any bounded variable in the negated normal form of $\phi(x)$ is bounded to be $\leq x$. Semantically, this means that for all models a formula $\phi(x)$ of $FOMLO^-$ satisfies:

$$\forall t \mathcal{M}_{FO}, t \models \phi(x) \leftrightarrow \mathcal{M}_{FO}[0, t], t \models \phi(x) \quad (1)$$

where $\mathcal{M}_{FO}[0, t]$ is a finite model starting from 0 and ending up to position t inclusive. Intuitively, formulae of $FOMLO^-$ can talk only about the past and the present. In what follows, given a set of propositions $\mathcal{P} = \{p_1, \dots, p_k\}$ and a set of predicate symbols $\Sigma = \{p_1(x), \dots, p_k(x)\}$ a *FOMLO* model \mathcal{M}_{FO} is *faithful* to \mathcal{M}_{TL} iff $\forall_{\geq 1}^k i. V^{FO}(p_i(x)) = V^{TL}(p_i)$.

DFO. Finally, on our expressiveness analysis we will also consider *DFO*, which in contrast to *FOMLO*, uses dyadic predicate symbols e.g., $p(x, y)$.

⁴As we work with discrete linear orders, this choice makes no difference.

4.2 $TPhL_o^-$ and Pure Past LTL

In this section we will show that $TPhL_o^-$ is expressively equal to the pure past fragment of LTL , i.e., $LTL[YS]$. The $TPhL_o^-$ fragment is described by the triplet $\langle \mathcal{P}^{es}, L_o^-, \Phi \rangle$, where \mathcal{P}^{es} is a set defined by the union of event and state predicate sets ($\mathcal{P}^{e/s}$); L_o^- is a set defined by the union of the set of the logical connectives $\{\wedge, \vee, \neg, \in\}$ and the set of temporal operators $\{\rightarrow, \sqcup, \sqcap, \backslash\}$. Formulae $TPhL_o^-$ are evaluated over $\mathcal{M}^- = \langle T, I, <, V^*, V^- \rangle$ models which are defined in a similar manner to the models presented in Section 2.

Given a finite set of event propositions⁵ $\mathcal{P}^e = \{e_1, \dots, e_k\}$, and a finite set of state propositions $\mathcal{P}^s = \{s_1, \dots, s_k\}$ we say the $FOMLO^-$ model $\mathcal{M}_{FO} = \langle T, <, V^{FO} \rangle$ is *faithful* to the model $\mathcal{M}^- = \langle T, I, <, V^*, V^- \rangle$ of $TPhL_o^-$ if it has the following properties:

- for any proposition e in \mathcal{P}^e , $V^*(e) = V^{FO}(e_t)$, and
- for any proposition s in \mathcal{P}^s , $V^-(s) = \rho(V^{FO}(s_+), V^{FO}(s_\infty), V^{FO}(s_-))$

where e_t corresponds to the monadic predicate $e_t(x)$ and the triplet (s_+, s_∞, s_-) corresponds to the monadic predicates $s_+(x)$, $s_\infty(x)$, $s_-(x)$ which are true on instants corresponding to the start, intermediate, and end of an interval respectively at which s is true. $\rho : 2^I \times 2^I \times 2^I \rightarrow 2^I$ is a partial mapping from three set of points to a set of intervals of I . Given three sets S, B and E , corresponding to starting, intermediate and ending points resp., ρ is defined as follows:

$$\rho(S, B, E) = \begin{aligned} & \{ [ts, te] : ts < te \wedge ts \in S \wedge te \in E \wedge \forall_{ts < t < te} t. (t \in B \wedge t \notin S \wedge t \notin E) \} \\ & \cup \{ [ts, \infty) : ts \in S \wedge \nexists_{t > ts} te. te \in E \wedge \forall_{t > ts} t. t \in B \} \end{aligned}$$

For all $i \in I$ where a formula $\phi \in \Phi^-$ is true ρ is bijective (recall that $\phi \in \Phi^-$ formulae always hold on disjoint intervals). For our expressiveness study, we will use the following theorem.

Theorem 1 (Gabbay et al. 1980 [15]). *For every formula of $FOMLO^-$ $\phi(x)$ we can find an $LTL[YS]$ formula θ such that $\forall t. \mathcal{M}_{TL}, t \models \phi \leftrightarrow \mathcal{M}_{FO}, t \models \theta$ for all \mathcal{M}_{FO} and their faithful models \mathcal{M}_{TL} .*

Proof. Dual proof of Theorem 2.2 in [15]. □

Theorem 2. *For every formula ϕ of $TPhL_o^-$, for all $TPhL_o^-$ models \mathcal{M}^- and their faithful $FOMLO^-$ models \mathcal{M}_{FO} :*

1. *if $\phi \in \Phi^*$, there exists a formula $\phi(t)$ with one free variable of $FOMLO^-$ such that $\mathcal{M}^-, t \models \phi$ iff $\mathcal{M}_{FO}, t \models \phi(t)$.*
2. *if $\phi \in \Phi^-$, there exist formulae $\phi^{+/\infty/-}(t)$ with one free variable such that $\mathcal{M}^-, [ts, te] \models \phi$ iff*

$$\mathcal{M}_{FO}, ts \models \phi^+(ts) \wedge \forall_{ts < t < te} t. \mathcal{M}_{FO}, t \models \phi^\infty(t) \wedge \mathcal{M}_{FO}, te \models \phi^-(te)$$

$$\text{and, } \mathcal{M}^-, [ts, \infty) \models \phi \text{ iff } \mathcal{M}_{FO}, ts \models \phi^+(ts) \wedge \forall_{t > ts} t. \mathcal{M}_{FO}, t \models \phi^\infty(t)$$

Proof. The proof is straightforward by direct translations. We define the translation τ^* from formulae of Φ^* to $FOMLO^-$ formulae as follows:

- $\tau^*(e, x) = e(x)$
- $\tau^*(\neg\phi, x) = \neg\tau^*(\phi, x)$

⁵In what follows, for simplicity, we will refer to atomic predicates of $TPhL$ as propositions.

- $\tau^*(\phi[\wedge, \vee]\psi, x) = \tau^*(\phi, x)[\wedge, \vee]\tau^*(\psi, x)$
- $\tau^*(\phi \in \psi, x) = \tau^*(\phi, x) \wedge (\tau_+^-(\psi, x) \vee \tau_{\in}^-(\psi, x) \vee \tau_{\in}^-(\psi, x))$ (We define $\tau_{+/\in/-}^-$ below.)

Considering that disjoint intervals can be recreated by their starting points, intermediate and endpoints, we define the τ_+^- , τ_{\in}^- , τ_-^- translation functions respectively, from formulae of Φ^- to $FOMLO^-$ as follows.

- $\tau_{\{+, \in, -\}}^-(s, x) = s_{\{+, \in, -\}}(x)$
- $\tau_+^-(\phi \rightarrow \psi, x) = \tau^*(\phi, x) \wedge \forall^{<x}z. [\tau^*(\phi, z) \rightarrow \exists_{>z}^{<t}z'. \tau^*(\psi \wedge \neg\phi, z')]$
- $\tau_{\in}^-(\phi \rightarrow \psi, x) = \exists^{<x}z. [\tau_+^-(\phi \rightarrow \psi, z) \wedge \forall_{>z}^{<x}z'. \neg\tau^*(\psi \wedge \neg\phi, z')]$
- $\tau_-^-(\phi \rightarrow \psi, x) = \tau^*(\psi \wedge \neg\phi, x) \wedge \exists^{<x}z. [\tau_+^-(\phi \rightarrow \psi, z) \wedge \forall_{>z}^{<x}z'. \neg\tau^*(\psi \wedge \neg\phi, z')]$
- $\tau_+^-(\phi \sqcup \psi, x) = [\tau_+^-(\phi, x) \wedge \neg\tau_+^-(\psi, x) \wedge \neg\tau_{\in}^-(\psi, x) \wedge \neg\tau_-^-(\psi, x)] \vee [\tau_+^-(\phi, x) \wedge \tau_+^-(\psi, x)] \vee [\tau_+^-(\psi, x) \wedge \neg\tau_+^-(\phi, x) \wedge \neg\tau_{\in}^-(\phi, x) \wedge \neg\tau_-^-(\phi, x)]$
- $\tau_{\in}^-(\phi \sqcup \psi, x) = [\tau_+^-(\phi, x) \vee \tau_{\in}^-(\psi, x)] \vee [\tau_+^-(\psi, x) \vee \tau_{\in}^-(\phi, x)] \vee [\tau_+^-(\phi, x) \wedge \tau_-^-(\psi, x)] \vee [\tau_-^-(\phi, x) \wedge \tau_+^-(\psi, x)]$
- $\tau_-^-(\phi \sqcup \psi, x) = [\tau_-^-(\phi, x) \wedge \neg\tau_+^-(\psi, x) \wedge \neg\tau_{\in}^-(\psi, x) \wedge \neg\tau_-^-(\psi, x)] \vee [\tau_-^-(\phi, x) \wedge \tau_-^-(\psi, x)] \vee [\tau_-^-(\psi, x) \wedge \neg\tau_+^-(\phi, x) \wedge \neg\tau_{\in}^-(\phi, x) \wedge \neg\tau_-^-(\phi, x)]$

The remaining translations are similar to the ones already presented and therefore omitted. It is easy to see that the conditions for an instant to be the starting, intermediate or endpoint of a formula of $TPhL^-$ is described by $FOMLO^-$ formulae. Consequently, given a formula ϕ of Φ^+ , $\mathcal{M}^-, t \models \phi \leftrightarrow \mathcal{M}_{FO}, t \models \tau^*(\phi, t)$. Given a formula ϕ of Φ^- it holds:

$$\mathcal{M}^-, [ts, te] \models \phi \text{ iff } \mathcal{M}_{FO}, ts \models \tau_+^-(\phi, ts) \wedge \forall_{>ts}^{<te}t. \mathcal{M}_{FO}, t \models \tau_{\in}^-(\phi, t) \wedge \mathcal{M}_{FO}, te \models \tau_-^-(\phi, te)$$

Finally, given a formula ϕ of Φ^- it holds:

$$\mathcal{M}^-, [ts, \infty) \models \phi \text{ iff } \mathcal{M}_{FO}, ts \models \tau_+^-(\phi, ts) \wedge \forall_{>ts}t. \mathcal{M}_{FO}, t \models \tau_{\in}^-(\phi, t) \quad \square$$

From, Theorems 1 and 2 we deduct that:

Theorem 3. For every formula ϕ of $TPhL_o^-$, for all models \mathcal{M}^- and their faithful models \mathcal{M}_{TL}^6 :

- if $\phi \in \Phi^+$ then there exists a formula ϕ_t of $LTL[YS]$ such that $\mathcal{M}^-, t \models \phi$ iff $\mathcal{M}_{TL}, t \models \phi_t$.
- if $\phi \in \Phi^-$ then there exist formulae $\phi_t^+, \phi_t^{\in}, \phi_t^-$ of $LTL[YS]$ such that $\mathcal{M}^-, [ts, te] \models \phi$ iff:

$$\mathcal{M}_{TL}, ts \models \phi_t^+ \wedge \forall_{>ts}^{<te}t. \mathcal{M}_{TL}, t \models \phi_t^{\in} \wedge \mathcal{M}_{TL}, te \models \phi_t^-$$

$$\text{and, } \mathcal{M}^-, [ts, \infty) \models \phi \text{ iff } \mathcal{M}_{TL}, ts \models \phi_t^+ \wedge \forall_{>ts}t. \mathcal{M}_{TL}, t \models \phi_t^{\in}$$

Now we will show that $LTL[YS]$ is expressible in $TPhL_o^-$. We define the translation $\tau^r : \Phi_t \rightarrow \Phi^+$ where Φ_t is the set of formulae of $LTL[YS]$ and Φ^+ is a subset of $TPhL_o^-$ formulae, as follows:

- $\tau^r(p) = p$
- $\tau^r(\neg\phi) = \neg\tau^r(\phi)$
- $\tau^r(\phi[\wedge, \vee]\psi) = \tau^r(\phi)[\wedge, \vee]\tau^r(\psi)$
- $\tau^r(Y\phi) = (\tau^r(\phi) \vee \neg\tau^r(\phi)) \in (\tau^r(\phi) \rightarrow \neg\tau^r(\phi)) \wedge \neg\text{start}(\tau^r(\phi) \rightarrow \neg\tau^r(\phi))$
- $\tau^r(\phi S \psi) = (\tau^r(\phi) \vee \neg\tau^r(\phi)) \in (\tau^r(\psi) \rightarrow \neg\tau^r(\phi))$

Note that in the case of Y and S, $\tau^r(\phi) \vee \neg\tau^r(\phi)$ is true everywhere but is restricted via the \in modality. It is clear that, given a finite set of propositions \mathcal{P} of $LTL[YS]$, for all models \mathcal{M}_{TL} and their corresponding faithful $TPhL_o^-$ models (for all $p \in \mathcal{P}$ it holds $V^{TL}(p) = V^*(p_e)$ where p_e is an event proposition), and for all $LTL[YS]$ formulae ϕ it holds $\mathcal{M}_{TL}, t \models \phi$ iff $\mathcal{M}^-, t \models \tau^r(\phi)$.

⁶We omit the definition of faithful models of $TPhL$ and LTL as they are defined in a similar manner to $TPhL$ and $FOMLO$.

4.3 $TPhL^-$ and $LTL[XU, YS]$

In this section we will show that $TPhL^-$, the extension of $TPhL_o^-$ (i.e., with filter and \heartsuit), has equal expressive power with $LTL[XU, YS]$. Our approach is similar to the previous section, however this time we will translate formulae of $TPhL^-$ to $FOMLO$ (instead of $FOMLO^-$). We will use Kamp's theorem:

Theorem 4 (Kamp [17]). *Given any FOMLO formula $\phi(x)$ with one free variable, there is an LTL formula θ , such that $\theta \equiv \phi(x)$ for all models \mathcal{M}_{FO} and \mathcal{M}_{TL} .*

Consequently, the only thing that remains to prove that $TPhL^-$ is expressible in $LTL[XU, YS]$ is to show that formulae involving the minimal range operator (\heartsuit) and filtering (filter) are expressible in $FOMLO$. Similar to the proof of Theorem 2, this is straightforward by extending the translation functions $\tau_+, \tau_-, \tau_{\infty}$ for supporting ' \heartsuit ' and 'filter'. We begin with the translation of formulae that involve the minimal range operator (\heartsuit):

$$\begin{aligned}\tau_+(\phi \heartsuit \psi, x) &= \tau^*(\phi, x) \wedge \exists_{>x} te. \left[\tau^*(\psi \wedge \neg\phi, te) \wedge \forall_{>x}^{<te} t. \left[\neg\tau^*(\phi, x) \wedge \neg\tau^*(\psi \wedge \neg\phi, t) \right] \right] \\ \tau_{\infty}^-(\phi \heartsuit \psi, x) &= \exists_{<x} ts. \exists_{>x} te. \left[\tau^*(\phi, ts) \wedge \tau^*(\psi \wedge \neg\phi, te) \wedge \forall_{>ts}^{<te} t. \left[\neg\tau^*(\phi, t) \wedge \neg\tau^*(\psi \wedge \neg\phi, t) \right] \right] \\ \tau_-(\phi \heartsuit \psi, x) &= \exists_{<x} ts. \left[\tau^*(\phi, ts) \wedge \tau^*(\psi \wedge \neg\phi, x) \wedge \forall_{>ts}^{<x} t. \left[\neg\tau^*(\phi, t) \wedge \neg\tau^*(\psi \wedge \neg\phi, t) \right] \right]\end{aligned}$$

Essentially, the translation of $\phi \heartsuit \psi$ is similar to the translation of $\phi \rightsquigarrow \psi$, however in this case it is clear that there is a need for future FOMLO formulae. Concerning the translation of formulae involving filtering (filter), first we define formulae $C^{k+}(x_0, \phi, \psi)$ as follows:

$$\begin{aligned}C^{k+}(x_0, \phi, \psi) &= \exists_{>x_0} x_1 \cdots \exists_{>x_{i-1}} x_i \cdots \exists_{>x_{k-1}} x_k \cdot \#_{>x_0}^{<x_1} x_{0,1} \cdot \#_{>x_i}^{<x_{i+1}} x_{i,i+1} \cdots \#_{>x_{k-1}}^{<x_k} x_{k-1,k} \\ &\quad \left[\phi(x_1) \wedge \cdots \wedge \phi(x_{k-1}) \wedge \psi(x_k) \right]\end{aligned}$$

denoting that ϕ is true from x_1 to x_{k-1} , ψ is true at x_k , and all x_i are contiguous and right of x_0 . Similar to C^{k+} define the C^{k-} for the left direction from x_0 . Here, we will only define the translations for the filter $_{<}$ case as the remaining cases can be easily defined in a similar manner.

$$\begin{aligned}\tau_+(\phi \text{ filter}_{<} n, x) &= \tau_+(\phi, x) \wedge \left[C^{1+}(x, \tau_{\infty}^-(\phi), \tau_-(\phi)) \vee \cdots \vee C^{n-1+}(x, \tau_{\infty}^-(\phi), \tau_-(\phi)) \right] \\ \tau_{\infty}^-(\phi \text{ filter}_{<} n, x) &= \tau_{\infty}^-(\phi, x) \wedge \left[C^{1-}(x, \tau_{\infty}^-(\phi), \tau_+(\phi)) \wedge C^{n-2+}(x, \tau_{\infty}^-(\phi), \tau_-(\phi)) \right. \\ &\quad \vee \left[C^{2-}(x, \tau_{\infty}^-(\phi), \tau_+(\phi)) \wedge C^{n-3+}(x, \tau_{\infty}^-(\phi), \tau_-(\phi)) \right] \vee \cdots \\ &\quad \left. \vee \left[C^{n-2-}(x, \tau_{\infty}^-(\phi), \tau_+(\phi)) \wedge C^{1+}(x, \tau_{\infty}^-(\phi), \tau_-(\phi)) \right] \right] \\ \tau_-(\phi \text{ filter}_{<} n, x) &= \tau_-(\phi, x) \wedge \left[C^{1-}(x, \tau_{\infty}^-(\phi), \tau_+(\phi)) \vee \cdots \vee C^{n-1-}(x, \tau_{\infty}^-(\phi), \tau_+(\phi)) \right]\end{aligned}$$

It can be seen that although a translation of $\phi \text{ filter}_{<} n$ exists, the size of the translated formula is linear to n . Note that a translation with smaller size might be possible, however for our expressiveness study it is not required to find the optimal translation.

Given all of the above, from Theorem 4, it is clear that the analog of Theorem 3 also holds for $TPhL^-$ and $LTL[XU, YS]$. For the opposite direction, it suffices to show that there are τ^r translations from formulae involving the remaining LTL modalities, i.e., X and U, to $TPhL^-$ formulae. For convenience we first define $c(\phi) = \phi \rightsquigarrow \neg\phi$ where $\phi \in \Phi$. Essentially, $c(\phi)$ holds for the maximal intervals $[ts, te]$ or $[ts, \infty)$ for which $\forall_{\geq ts}^{<te} t. \mathcal{M}, t \models \phi$ or $\forall_{\geq ts} t. \mathcal{M}, t \models \phi$ respectively. Therefore we define the corresponding

τ^r translations as follows.

$$\begin{aligned}\tau^r(\mathbb{X}\phi) &= (\tau^r(\phi) \in c(\tau^r(\phi))) \wedge \neg \text{end}(c(\tau^r(\phi))) \vee \text{start}(\neg\tau^r(\phi) \wp \text{end}(c(\neg\tau^r(\phi)))) \\ \tau^r(\phi \cup \psi) &= ((\tau^r(\phi) \vee \neg\tau^r(\phi)) \in (\text{start}(c(\tau^r(\phi))) \sqcap c(\neg\tau^r(\psi))) \wp \\ &\quad (\text{end}(c(\tau^r(\phi))) \sqcap c(\neg\tau^r(\psi))) \wedge \tau^r(\psi)) \vee \tau^r(\psi)\end{aligned}$$

Concerning the translation of $\mathbb{X}\phi$, the left part of the disjunction holds true for all instants included in an interval $[ts, te]$ where $c(\tau^r(\phi))$ is true, while the right part is true at the start of an interval $[t, t+1]$ where $\neg\tau^r(\phi) \wp \text{end}(c(\neg\tau^r(\phi)))$ is true. In the case of $\tau^r(\phi \cup \psi)$, the translation can be divided into two parts: the first part uses the inclusion operator between the tautology $\tau^r(\phi) \vee \neg\tau^r(\phi)$ (true everywhere) and the minimal range formula between (a) the start of a period at which both ϕ and $\neg\psi$ are true for all points (excluding the end) and (b) the end of a period $[ts, te]$ at which both ϕ and $\neg\psi$ are true and ψ holds at te , thus capturing the cases where ϕ is true before ψ becomes true; the second part of the translation is the case of $\tau^r(\psi)$ which captures single instances of ψ . Considering all of the above, we can now say that the $TPhL^-$ fragment of $TPhL$ has equal expressive power with $LTL[XU, YS]$.

4.4 Expressiveness of $TPhL$

Concerning the complete language $TPhL$, a comparison with LTL is not possible as the structures on which semantics is based are incomparable, even for atomic entities. This is because dynamic temporal phenomena may hold on non-disjoint intervals which by default require the half plane of a 2-dimensional temporal space for their representation. When compared to DFO , it can be easily seen that for all formulae of $TPhL$ there are equivalent formulae with two free variables of DFO . For example, a dynamic temporal phenomenon proposition p can be represented by a dyadic predicate $p(x, y)$ such that $\mathcal{M}, [ts, te] \models p \leftrightarrow \mathcal{M}_{DFO}, ts, te \models p(x, y)$. In a similar manner to the previous sections, we can define translations from $TPhL$ to DFO for the remaining formulae—this time however with two free variables corresponding to starting and ending instants. However, the reverse direction does not hold. This can be shown with the following example. Consider the DFO formula $\phi(x, y) = \neg p(x, y)$, since negation is not included for formulae⁷ of $\Phi^=$ there is no formula ϕ_t of $TPhL$ such that $\mathcal{M}, [ts, te] \models \phi_t \leftrightarrow \mathcal{M}_{DFO}, ts, te \models \neg p(x, y)$ for all models. Therefore $TPhL$ is strictly less expressive than DFO .

5 Stream processing

In this section, we formally present the correctness criteria for stream processing with the $TPhL$ language. Given a stream, i.e., an arbitrary long sequence of time associated atomic formulae of Φ , the evaluation at a given instant t , of formulae that refer only to the past (ϕ of $TPhL_o^-$) is an easy task, as their truth value can be determined for all $t' \leq t$ (see Equation (1)). However, this is not the case for formulae such as $\tau^r(\mathbb{X}p)$ and $\tau^r(p \cup q)$ that refer to the future, as their truth value at an instant t may depend on future information ($> t$). Consequently, in order to guarantee *correctness*, *monotonicity* and *punctuality*—we will define these shortly—, the two valued semantics of $TPhL$, are not sufficient for the evaluation of formulae on constantly evolving streams. In order to treat the issue of evaluations with unknown status—i.e., when all required information is not available at current time—, we follow an approach similar to [6]. We extend the semantics of $TPhL$ for stream processing, to utilise three values: true (\top), false (\perp) and unknown ($?$). Due to space limitations, we will omit the presentation of the three valued

⁷We chose to omit negation from formulae of $\Phi^=$ as it would affect significantly the performance of our implementation.

semantics⁸ in this paper, instead we will focus on formalising the notions of stream, stream processor, and define the properties of correctness, punctuality and monotonicity for stream processors of *TPhL*.

A *stream*, at any instant t can be represented by the finite model $\mathcal{M}_t = \langle T_t, I_t, <, V^+, V^-, V^= \rangle$ where $T_t = \{0, 1, \dots, t\}$, $I_t = T_t \times T_t \cup \{[ts, \infty) : ts \in T_t\}$, and $V^+, V^-, V^=$ are valuation functions defined in similar manner to Section 2.

A *stream processor*, in symbols \mathcal{SP}_t , is defined by the triplet $\langle \Lambda_t^+, \Lambda_t^-, \Lambda_t^= \rangle$ where $t \in T$, $\Lambda_t^+ : \Phi \times T_t \rightarrow \{\top, \perp, ?\}$, $\Lambda_t^- : \Phi \times (I_t^c \cup I_t^+) \rightarrow \{\top, \perp, ?\}$, and $\Lambda_t^= : \Phi \times (I_t^c \cup I_t^+) \rightarrow \{\top, \perp, ?\}$, are formulae valuation functions assigning truth values on formulae-instants/intervals pairs and $I_t^c = T_t \times T_t$ and $I_t^+ = \{[ts, t+) : ts, t \in T_t\}$. Intervals of I_t^c , (e.g., $[ts, te)$) denote that a phenomenon started at ts and ended at te , while intervals of I_t^+ , (e.g., $[ts, t+)$) denote that a phenomenon started at ts , and continues to be true/unknown at t but does not end at t . Intervals of I_t^+ are useful for capturing the truth value of valuations that are true but are still ongoing—see for example the semantics of $\phi \rightsquigarrow \psi$ for intervals open to ∞ . We assume that the input phenomena are ordered and their truth value is never unknown. Now, we define the *correctness*, *punctuality* and *monotonicity* properties for \mathcal{SP} .

Correctness. A stream processor has the *correctness* property iff given any stream, for all t and for any $\phi \in \Phi$ evaluation by \mathcal{SP}_t (i.e., via $\Lambda^{+/-/=}$) that is true (false) at an instant t_i or interval i , ϕ is also true (false) (i.e., via the semantics of Section 2) at t_i or i in \mathcal{M}_t .

Monotonicity. A stream processor has the *monotonicity* property iff given any stream, for all t and for $\phi \in \Phi$ evaluation by \mathcal{SP}_t that is true (false) at an instant t_i or interval i , will also be evaluated to be true (false) at t_i or i by all $\mathcal{SP}_{t'}$ with $t' > t$.

Punctuality. A stream processor has the *punctuality* property iff given any stream, for any ϕ , and for all instants t_i or intervals i if there exists minimum $t \geq t_i$ such that ϕ is true (false) for all $t' \geq t$ in all \mathcal{M}'_t at t_i or t then \mathcal{SP}_t evaluates ϕ to be true (false) at an instant t_i or interval i .

We say that a stream processor for *TPhL* is *proper* iff it has all the three aforementioned properties. It is easy to see that some formulae, given certain streams, can never be true but always stay unknown. For example, consider the formula $\tau^r(Gp)$, where G is the ‘globally’ *LTL* operator, and a stream where p is true at all instants; at any given point t in time, the stream processor is agnostic to the future, therefore in order to maintain the *monotonicity* property, $\tau^r(Gp)$ will be evaluated by \mathcal{SP}_t for all $t' \leq t$ to be unknown.

Phenesthe+, is a *proper* stream processor of *TPhL*. While we will not present a formal proof in this paper, we will briefly discuss its processing and its implementation. Phenesthe+ is a complex event processing engine that given an input stream and a set of temporal phenomena definitions, will produce an output stream of temporal phenomena detections i.e., phenomena associated with a set of instants or intervals at which they are true. Compared to automata based methods, the phenomena are compiled via rewriting into an internal Prolog representation which is then later used for processing. This procedure is linear with respect to the size of the formulae involved. The phenomena definitions can be hierarchical, and their processing, if possible, can happen in parallel. Phenesthe+ detects phenomena by performing temporal queries over tumbling temporal windows of size equal to a user defined step (ST) size. A temporal window contains all the new information that has arrived since the last temporal query, as well as information from previous windows that has a possible future use. For example, given the formula $\phi \rightsquigarrow \psi$, if ϕ is true in the current window but ψ is not, then ϕ must be retained. Note that information from previous windows was also retained in the previous version of Phenesthe for valuations that are true but ongoing, or involved dynamic temporal phenomena. All information outside the temporal window that does not have ‘future use’ is discarded. From a practical perspective it is not viable to keep everything

⁸The complete three valued semantics are available in https://manospits.github.io/files/Three_valued_semantics.pdf.

from the past that can contribute to a future detection. Therefore, we allow setting a maximum limit for past information (retaining threshold RT). When this threshold is active Phenesthe+ is no longer *proper* with respect to the full stream, but remains *proper* for the part of the stream that is bounded by RT . Similarly, the *punctuality* property depends on ST , if $ST = 1$, then Phenesthe+ is punctual, however if $ST > 1$, detections will be produced at the latest $ST - 1$ time units after their punctual time.

In terms of complexity, evaluation of formulae in Phenesthe+ happens via single-scan or in the worst case, i.e., when overlapping intervals are involved, polynomial algorithms with respect to the size of the structure (current temporal window). It has to be noted that while $TPhL^-$ has equal expressive power with LTL, it can accomplish efficient processing of phenomena definitions by utilising intervals to represent set of points. For example an interval $[ts, te]$ produced by the evaluation of the formula $\phi \mapsto \psi$ requires only two points for the representation of all the instants included in $[ts, te]$, therefore contributing significantly to space and processing time economy.

6 Experimental Evaluation

We presented the theoretical basis of $TPhL$. Now, we will evaluate the efficiency of our extended stream processing engine on a reproducible⁹ maritime monitoring scenario.

Experimental setup For our experimental evaluation we use a public dataset containing AIS vessel data, transmitted over a period of 6 months, from October 1st, 2015 to March 31st, 2016, in the area of Brest, France [26] along with spatio-temporal events relating vessels with areas (in total $\approx 16M$ input events). We run our experiments on machine with an Intel i7-3770 CPU running Ubuntu 20.04.6 LTS. The set of maritime phenomena we detect as well as the input events are summarised in Table 1. We compare stream processing efficiency when the set of maritime phenomena definitions includes and does not include phenomena marked with \ddagger , i.e., phenomena that utilise \wp and, or filter or depend on phenomena that utilise them.

Experimental results The results of our evaluation are illustrated in Figure 3. We perform complex event processing with $ST = 3h$, and $RT = \{2, 4, 8, 16\}$ days. Figure 3 (left) shows the average processing time for each experiment. The results show that the addition of \ddagger phenomena does not affect processing efficiency significantly, but also that Phenesthe+ is capable of producing detections in less than 2 seconds (multithreaded) when the data retaining threshold is set to 16 days. Note that the performance gain by running the multithreaded version of Phenesthe+ depends on the dependencies between phenomena and the computation of the processing order. For example, *fishing_warning* and *waiting_time* can be processed in parallel while *fishing_warning* and *suspicious_trip* cannot. In [22] we describe the computation of the processing order. With the addition of the new temporal phenomena, we limited the number of phenomena that can be processed in parallel. The results of Figure 3 (left) confirm this. We also perform complex event processing with $ST = 24h$, and set $RT = \infty$ (i.e., keep non-redundant information forever)—recall that our dataset involves a 6 month period. Similar to the previous experiments we compare performance when Phenesthe+ is executed in parallel or serial manner, and with or without temporal phenomena marked with \ddagger (see Table 1). Figure 3 (middle) shows the average processing time while Figure 3 (right) shows the average number of input entities plus retained instants/ intervals for each case. The results show, that in terms of processing time, performance is not significantly affected when including \ddagger phenomena in both serial and parallel processing even when Phenesthe+ retains all information. In more detail, apart from the input events (on average 90K per temporal query) when we include \ddagger phenomena the number of retained instants/intervals increases on average by 20K, therefore bringing the

⁹<https://github.com/manospits/Phenesthe/tree/future>

Table 1: Input and output phenomena description. ‘IE’, ‘UE’, ‘US’ and ‘UD’ stand for Input/User Event/State/Dynamic temporal phenomenon. Phenomena with † have future dependencies while phenomena with ‡ utilise the new temporal modalities or depend on phenomena that utilise them. The last column lists approximately the number of input or output instants/intervals.

Type	Phenomenon	Description	Number
IE	$ais(V, S, C, H)$	AIS transmitted information (vessel ID, speed, course, heading)	15.8M
	$enters/leaves\{Port, Fishing\}(V, A)$	Vessel enters/leaves port/fishing area.	160K
UE	$stop_start/end(V)$	Start/end of a stop.	800K
	$in_port, fishing_area(V, A)$	In fishing/port area.	70K
	$stopped(V)$	Stopped vessel.	300K
US	$underway(V)$	Vessel underway.	132K
	$moored(V)$	Moored vessel.	323K
	‡ $fishing_warning(V, F)$	Warning: Non fishing vessel possibly engaged in fishing.	7K
	‡ $waiting_time(V, P)$	Port waiting time.	42K
	‡ $long_waiting_time(V, P)$	Warning: waiting time longer than a threshold.	28K
	$unusual_stop(V)$	Warning: vessel performs a stop in an unexpected area.	27K
	‡ $possible_malfunction(V)$	Warning: Vessel might have a malfunction.	3K
	† $trip(V, PA, PB)$	Vessel trip from PA to PB.	39K
UD	‡ $suspicious_trip(V, PA, PB)$	Trip from PA to PB contained warnings.	3K
	† $fishing_trip(V, PA, FA, PB)$	Fishing trip from PA to PB contained fishing in FA.	6K

total number of input+retained entities up to 120K. Even in this setting Phenesthe+, produces detections in approximately 2 and 4 seconds (serial and parallel respectively).

7 Related work & Discussion

There are several very expressive temporal logics. The HS logic [16] is a very powerful logic for representing both instantaneous and durative temporal phenomena. When time is linear and the intervals homogeneous (therefore non overlapping) the HS logic is equally expressive with *LTL* but is exponentially more succinct [9]. In its original version, the HS logic does not make any assumptions on the nature of intervals. In this paper, we showed the *TPhL*⁻ has equal expressive power to *LTL*, therefore concerning the linear HS variant studied in [9], *TPhL*⁻ is equally expressive. It is well known, that the chop operator of Venema’s CDT logic [27] is inexpressible in the HS logic [20, 10]. *TPhL* supports the chop operator in the form of meets. While a formal expressiveness comparison of *TPhL* with the HS or CDT would be desirable, the omission of negation from $\Phi^=$ formulae makes this a challenging task.

Concerning our criteria for proper stream processors of *TPhL*, as mentioned earlier, their concepts are not entirely new. For example “correctness” has a similar notion with “soundness” in run-time verification [5] (i.e., the output should be correct with respect to the specification). Likewise, the “monotonicity” property as we have defined it, in run-time monitoring appears as the irrevocability property respectively for monitors [1]—a monitor that has the “irrevocability” property is unable to revoke the acceptance or the rejection of a trace. Finally, the “punctuality” property can be related to the “tightness” property of monitors [1], under which monitors are restricted to make a choice as soon as there is sufficient information available. In this work, we utilise the similar notions from run-time monitoring and verification for the task of complex event processing.




Figure 3: Experimental results. Average processing time per query with $ST = 3h$ and $RT = \{2, 4, 8, 16\}$ days or $ST = 24h$ and $RT = \infty$ (left) and (middle) respectively. Average number of input plus retained instants/ intervals per query when $ST = 24h$ and $RT = \infty$ (right).

From a complex event processing perspective, LARS is a logic based framework for reasoning over streams [7]. While the language of LARS is expressible in LTL , the reverse direction does not hold, since LARS does not support ‘until’. A well known runtime monitoring system with point-based semantics is LOLA [13]. While LOLA does not allow durative phenomena, we saw in section 4 that formulae that hold on disjoint intervals can be expressed using point based modalities. It is not possible, however, to model formulae that hold on overlapping intervals. Furthermore, in the worst case LOLA requires memory equal to the size of the trace so far, which is not practical for large industrial applications such as maritime monitoring. In Phenesthe+ we allow the user to choose the retaining threshold. A complex event recognition framework is RTEC [4]. RTEC is a logic based formalism whereby events and fluents are expressed with a variant of the Event Calculus [18]. While there isn’t a formal study of the expressive power of the language of RTEC, its semantics suggest that it has at most equal expressive power with pure past LTL . Bauer et. al. [6], propose three valued semantics for monitoring LTL formulae. In our work, we also use three valued semantics, however for a more general case, as our language allows the representation of temporal phenomena that hold on overlapping intervals, which cannot be modeled in LTL . Team semantics for LTL [19] or Hyper LTL [11] offer a promising direction towards the representation of concurrent temporal phenomena, however they are limited to a finite number of concurrent traces. In $TPhL$ a dynamic temporal phenomenon may hold on possibly infinite overlapping intervals.

Closing, in this paper we presented $TPhL$ and studied the expressive power of its different fragments. Specifically, we showed that $TPhL_o^-$ has equal expressive power with pure past LTL while its extension, $TPhL^-$, has equal expressive power with LTL . Concerning the complete logic $TPhL$, we showed that it is strictly less expressive than dyadic first-order logic. Moreover, we defined criteria for *proper* stream processors that use our language, and evaluated Phenesthe+, our stream processing implementation on real maritime data. Our results, show that Phenesthe+ is suitable for the task of maritime monitoring as it produces results in real-time. While the application of our experiment involved the maritime domain, Phenesthe+ is generic, and can be applied in other areas.

Regarding future work, we aim to study the expressive power of a theoretical variant of $TPhL$ that includes negation on formulae of Φ^- in comparison with two-dimensional modal logics. Furthermore, as one of the main motivations for the creation of $TPhL$ was facilitating writing temporal formulae, we plan to compare succinctness of $TPhL^-$ formulae with LTL formulae. Finally, we aim to apply Phenesthe+ for human activity monitoring in smart homes.

References

- [1] Luca Aceto, Antonis Achilleos, Adrian Francalanza, Anna Ingólfssdóttir & Karoliina Lehtinen (2019): *Adventures in Monitorability: From Branching to Linear Time and Back Again*. *Proc. ACM Program. Lang.* 3(POPL), doi:10.1145/3290365.
- [2] James F. Allen (1983): *Maintaining knowledge about temporal intervals*. *Communications of the ACM* 26(11), p. 832–843, doi:10.1145/182.358434.
- [3] Darko Anicic, Paul Fodor, Sebastian Rudolph, Roland Stühmer, Nenad Stojanovic & Rudi Studer (2010): *A Rule-Based Language for Complex Event Processing and Reasoning*, p. 42–57. *Lecture Notes in Computer Science* 6333, Springer Berlin Heidelberg, doi:10.1007/978-3-642-15918-3_5.
- [4] Alexander Artikis, Marek Sergot & Georgios Paliouras (2015): *An Event Calculus for Event Recognition*. *IEEE Transactions on Knowledge and Data Engineering* 27(4), pp. 895–908, doi:10.1109/TKDE.2014.2356476.
- [5] Ezio Bartocci & Yliès Falcone, editors (2018): *Lectures on Runtime Verification - Introductory and Advanced Topics*. *Lecture Notes in Computer Science* 10457, Springer, doi:10.1007/978-3-319-75632-5.
- [6] Andreas Bauer, Martin Leucker & Christian Schallhart (2011): *Runtime Verification for LTL and TLTL*. *ACM Trans. Softw. Eng. Methodol.* 20(4), doi:10.1145/2000799.2000800.
- [7] Harald Beck, Minh Dao-Tran & Thomas Eiter (2018): *LARS: A Logic-based framework for Analytic Reasoning over Streams*. *Artificial Intelligence* 261, pp. 16–70, doi:10.1016/j.artint.2018.04.003.
- [8] Michael H Bohlen, Renato Busatto & Christian S. Jensen (1998): *Point-versus interval-based temporal data models*. In: *Proceedings 14th International Conference on Data Engineering*, pp. 192–200, doi:10.1109/ICDE.1998.655777.
- [9] Laura Bozzelli, Alberto Molinari, Angelo Montanari, Adriano Peron & Pietro Sala (2018): *Interval vs. Point Temporal Logic Model Checking: An Expressiveness Comparison*. *ACM Trans. Comput. Logic* 20(1), doi:10.1145/3281028.
- [10] Davide Bresolin, Dario Della Monica, Valentin Goranko, Angelo Montanari & Guido Sciavicco (2008): *Decidable and Undecidable Fragments of Halpern and Shoham’s Interval Temporal Logic: Towards a Complete Classification*. In Iliano Cervesato, Helmut Veith & Andrei Voronkov, editors: *Logic for Programming, Artificial Intelligence, and Reasoning*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 590–604, doi:10.1007/978-3-540-89439-1_41.
- [11] Michael R. Clarkson, Bernd Finkbeiner, Masoud Koleini, Kristopher K. Micinski, Markus N. Rabe & César Sánchez (2014): *Temporal Logics for Hyperproperties*. In Martín Abadi & Steve Kremer, editors: *Principles of Security and Trust - Third International Conference, POST 2014, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2014, Grenoble, France, April 5-13, 2014, Proceedings*, *Lecture Notes in Computer Science* 8414, Springer, pp. 265–284, doi:10.1007/978-3-642-54792-8_15.
- [12] Gianpaolo Cugola & Alessandro Margara (2010): *TESLA: a formally defined event specification language*. In: *DEBS ’10*, ACM Press, p. 50, doi:10.1145/1827418.1827427.
- [13] Ben D’Angelo, Sriram Sankaranarayanan, Cesar Sanchez, Will Robinson, Bernd Finkbeiner, Henny B. Sipma, Sandeep Mehrotra & Zohar Manna (2005): *LOLA: Runtime Monitoring of Synchronous Systems*. In: *Proceedings of the 12th International Symposium on Temporal Representation and Reasoning, TIME ’05*, IEEE Computer Society, USA, p. 166–174, doi:10.1109/TIME.2005.26.
- [14] Andreas Dohr, Christiane Engels & Andreas Behrend (2018): *Algebraic Operators for Processing Sets of Temporal Intervals in Relational Databases*. In Natasha Alechina, Kjetil Nørsvåg & Wojciech Penczek, editors: *25th International Symposium on Temporal Representation and Reasoning (TIME 2018)*, *Leibniz International Proceedings in Informatics (LIPIcs)* 120, Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, pp. 11:1–11:16, doi:10.4230/LIPIcs.TIME.2018.11.
- [15] Dov Gabbay, Amir Pnueli, Saharon Shelah & Jonathan Stavi (1980): *On the Temporal Analysis of Fairness*. In: *Proceedings of the 7th ACM SIGPLAN-SIGACT Symposium on Principles of Program-*

- ming Languages*, POPL '80, Association for Computing Machinery, New York, NY, USA, p. 163–173, doi:10.1145/567446.567462.
- [16] Joseph Y. Halpern & Yoav Shoham (1991): *A propositional modal logic of time intervals*. *Journal of the ACM* 38(4), p. 935–962, doi:10.1145/115234.115351.
- [17] Johan Anthony Wilem Kamp (1968): *Tense Logic and the Theory of Linear Order*. Ph.D. thesis, University of California, Los Angeles.
- [18] Robert Kowalski & Marek Sergot (1986): *A logic-based calculus of events*. *New Generation Computing* 4(1), pp. 67–95, doi:10.1007/BF03037383.
- [19] Andreas Krebs, Arne Meier, Jonni Virtema & Martin Zimmermann (2018): *Team Semantics for the Specification and Verification of Hyperproperties*. In Igor Potapov, Paul G. Spirakis & James Worrell, editors: *43rd International Symposium on Mathematical Foundations of Computer Science, MFCS 2018, August 27–31, 2018, Liverpool, UK, LIPIcs 117*, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 10:1–10:16, doi:10.4230/LIPIcs.MFCS.2018.10.
- [20] Kamal Lodaya (2000): *Sharpening the Undecidability of Interval Temporal Logic*. In Jifeng He & Masahiko Sato, editors: *Advances in Computing Science - ASIAN 2000, 6th Asian Computing Science Conference, Penang, Malaysia, November 25–27, 2000, Proceedings, Lecture Notes in Computer Science 1961*, Springer, pp. 290–298, doi:10.1007/3-540-44464-5_21.
- [21] Corto Mascle, Daniel Neider, Maximilian Schwenger, Paulo Tabuada, Alexander Weinert & Martin Zimmermann (2020): *From LTL to RLTL Monitoring: Improved Monitorability through Robust Semantics*. In: *Proceedings of the 23rd International Conference on Hybrid Systems: Computation and Control, HSCC '20*, Association for Computing Machinery, New York, NY, USA, pp. 170–204, doi:10.1145/3365365.3382197.
- [22] Manolis Pitsikalis, Alexei Lisitsa & Shan Luo (2021): *Representation and Processing of Instantaneous and Durative Temporal Phenomena*. In Emanuele De Angelis & Wim Vanhoof, editors: *Logic-Based Program Synthesis and Transformation - 31st International Symposium, LOPSTR 2021, Tallinn, Estonia, September 7–8, 2021, Proceedings, Lecture Notes in Computer Science 13290*, Springer, pp. 135–156, doi:10.1007/978-3-030-98869-2_8.
- [23] Manolis Pitsikalis, Alexei Lisitsa, Patrick Totzke & Simon Lee (2022): *Making Sense of Heterogeneous Maritime Data*. In: *2022 23rd IEEE International Conference on Mobile Data Management (MDM)*, pp. 401–406, doi:10.1109/MDM55031.2022.00089.
- [24] Amir Pnueli (1977): *The temporal logic of programs*. In: *18th Annual Symposium on Foundations of Computer Science (sfcs 1977)*, IEEE, Providence, RI, USA, p. 46–57, doi:10.1109/SFCS.1977.32.
- [25] Alexander Rabinovich (2014): *A Proof of Kamp's theorem*. *Logical Methods in Computer Science* Volume 10, Issue 1, doi:10.2168/LMCS-10(1:14)2014.
- [26] Cyril Ray, Richard Dréo, Elena Camossi, Anne-Laure Joussemme & Clément Iphar (2019): *Heterogeneous integrated dataset for Maritime Intelligence, surveillance, and reconnaissance*. *Data in Brief* 25, p. 104141, doi:10.1016/j.dib.2019.104141.
- [27] Yde Venema (1990): *Expressiveness and completeness of an interval tense logic*. *Notre Dame Journal of Formal Logic* 31(4), doi:10.1305/ndjfl/1093635589.
- [28] Peter Øhrstrøm & Per Hasle (2006): *Modern temporal logic: The philosophical background*. In Dov M. Gabbay & John Woods, editors: *Logic and the Modalities in the Twentieth Century, Handbook of the History of Logic* 7, North-Holland, pp. 447–498, doi:10.1016/S1874-5857(06)80032-4.