

1 A Stream Reasoning System for Maritime 2 Monitoring

3 **Georgios M. Santipantakis, Akrivi Vlachou, Christos Doulkeridis**

4 University of Piraeus, Karaoli and Dimitriou 80, 18534 Piraeus, Greece

5 {gsant|avlachou|cdouk}@unipi.gr

6 **Alexander Artikis**

7 Department of Maritime Studies, University of Piraeus, Greece

8 Institute of Informatics & Telecommunications, NCSR “Demokritos”, Ag. Paraskevi, Greece

9 a.artikis@unipi.gr

10 **Ioannis Kontopoulos**

11 Institute of Informatics & Telecommunications, NCSR “Demokritos”, Ag. Paraskevi, Greece

12 ikon@iit.demokritos.gr

13 **George A. Vouros**

14 University of Piraeus, Karaoli and Dimitriou 80, 18534 Piraeus, Greece

15 georgev@unipi.gr

16 — Abstract —

17 We present a stream reasoning system for monitoring vessel activity in large geographical areas.
18 The system ingests a compressed vessel position stream, and performs online spatio-temporal link
19 discovery to calculate proximity relations between vessels, and topological relations between vessel
20 and static areas. Capitalizing on the discovered relations, a complex activity recognition engine,
21 based on the Event Calculus, performs continuous pattern matching to detect various types of
22 dangerous, suspicious and potentially illegal vessel activity. We evaluate the performance of the
23 system by means of real datasets including kinematic messages from vessels, and demonstrate
24 the effects of the highly efficient spatio-temporal link discovery on performance.

25 **2012 ACM Subject Classification** activity recognition and understanding

26 **Keywords and phrases** event pattern matching, Event Calculus

27 **Digital Object Identifier** 10.4230/LIPIcs.TIME.2018.20

28 **Acknowledgements** This work is supported by the datAcron project, which has received fund-
29 ing from the European Union’s Horizon 2020 research and innovation programme under grant
30 agreement No 687591 (<http://datacron-project.eu>).

31 **1 Introduction**

32 Nowadays, maritime surveillance systems that keep track of the daily operation of vessel fleets
33 constitute an essential tool for large shipping companies, coast guards, as well as governmental
34 agencies, due to their immense effect on economy and environment [10]. Advances in
35 navigation technology enable the real-time provision of vessel positional information for the
36 benefit of surveillance systems. The Automatic Identification System (AIS)¹, for example, is
37 a tracking system for identifying and locating vessels at sea through data exchange. The
38 acquisition of positional data is achieved either by AIS base stations along coastlines, or

¹ <http://www.imo.org/en/OurWork/Safety/Navigation/Pages/AIS.aspx>



© Santipantakis et al;

licensed under Creative Commons License CC-BY

25th International Symposium on Temporal Representation and Reasoning (TIME 2018).

Editors: Natasha Alechina, Kjetil Nørkvåg, and Wojciech Penczek; Article No. 20; pp. 20:1–20:16

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

39 even by satellites when out of range of terrestrial networks. As this data is streamed in a
 40 maritime surveillance system, several operations need to be performed in real-time, including
 41 data integration and complex maritime activity recognition.

42 We present a monitoring system that exploits spatio-temporal relations between vessels,
 43 or vessels and areas of interest (e.g., protected areas), which are calculated online by a
 44 dedicated component for spatio-temporal link discovery (stLD). These relations are provided
 45 as input to a complex activity recognition component, which is based on the ‘Event Calculus
 46 for Run-Time reasoning’ (RTEC) [3]. This is an Event Calculus [12] implementation with
 47 various optimization techniques for continuous narrative assimilation on data streams.

48 We address the problem of online spatio-temporal link discovery over streaming and
 49 archival data. This link discovery (LD) problem is challenging because of (a) the streaming
 50 nature of data, and (b) the complexity of evaluating spatio-temporal similarity functions
 51 to determine the links between spatio-temporal entities. Typically, LD tasks are solved by
 52 *filter-and-refine* algorithms that identify a set of candidate entities (for linking) in the *filtering*
 53 *step*, which need to be verified in the *refinement step*. The refinement step is the principal
 54 cost factor that determines the overall performance of LD, since it requires the evaluation
 55 of distance/similarity functions on complex geometrical entities. For the case of stream to
 56 static LD, we propose a filtering technique that improves the efficiency of the filtering step
 57 by eliminating entities that cannot be linked, thereby reducing the number of entities that
 58 have to be considered during refinement. For the case of streaming data only, we provide an
 59 efficient method for streaming LD, identifying proximity relations between moving vessels.

60 In earlier work, we presented a maritime monitoring system which employed RTEC for
 61 complex activity recognition [20]. In that work, RTEC performed spatial calculations to
 62 determine whether a vessel is close to a port, or within an area of interest. Furthermore, it
 63 approximated crudely the *nearby* relation between vessels by checking whether a pair of vessels
 64 is located within the same cell of a grid. In this work, we placed emphasis on the detection
 65 of spatial relations and developed a separate component for highly efficient spatio-temporal
 66 link discovery. Moreover, RTEC has at its disposal additional spatial relations—whether a
 67 vessel is *nearby* some area—and a much more accurate account of proximity between vessels.

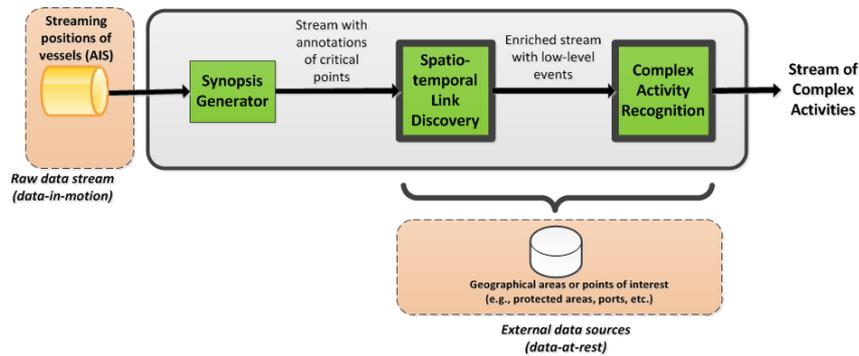
68 To summarise, this paper makes the following contributions:

- 69 ■ We present a stream reasoning system integrating a component for spatio-temporal
 70 link discovery (stLD), and a component for recognizing complex activities by means of
 71 temporal pattern matching.
- 72 ■ We propose a technique for stream-based LD, which improves the efficiency of filtering,
 73 by eliminating entities that cannot be linked, thereby reducing the number of entities
 74 that have to be considered during refinement.
- 75 ■ We evaluate the performance of the system by means of real datasets including AIS
 76 kinematic messages from vessels.

77 **2 System Architecture**

78 Our work is performed in the context of the datAcron research project², which aims at
 79 advancing the management of voluminous and heterogeneous data-at-rest (archival data)
 80 and data-in-motion (streaming data) sources, so as to promote the safety and effectiveness
 81 of critical operations of moving objects in large geographical areas. The architecture of the
 82 datAcron prototype for complex activity recognition is depicted in Figure 1.

² <http://datacron-project.eu/>



■ **Figure 1** The datAcron prototype architecture.

83 The main input is streaming positions of vessels, in the form of AIS messages. As this
 84 streaming data flows in the system, trajectory compression takes place, by annotating a
 85 subset of the original vessel positions as *critical points/events*. The *Synopses Generator*
 86 component (see Figure 1) provides algorithms for trajectory reconstruction and compression,
 87 by cleaning erroneous data and eliminating vessel positions that do not significantly affect
 88 the quality of trajectory representation. Then, critical points are linked with archival data
 89 online, namely marine areas or points of interest, such as protected areas (Natura2000) or
 90 ports (World Port Index). Link discovery is performed at the spatio-temporal level, thus
 91 identifying those areas (or points) that are related to a given position of a vessel. Relations
 92 may be topological (e.g., a vessel is located *within* an area) or proximity-based (e.g., a vessel
 93 is *nearby* a port). In addition, vessel positions are linked to each other, by processing the
 94 streaming data only; this results in discovering proximity relations between moving vessels
 95 (e.g., a vessel is *nearby* another vessel) in an online fashion. Subsequently, a complex activity
 96 recognition module consumes the stream of spatial relations and critical points to recognize
 97 various types of vessel activity. The *Complex Activity Recognition* component is based on
 98 the ‘Event Calculus for Run-Time reasoning’ (RTEC) [3].

99 The innovative feature of the proposed architecture is the integration of an optimized
 100 component for online discovery of spatial relations with an activity recognition engine. In
 101 the following, we delve into the technical details of spatio-temporal link discovery (Section 3)
 102 and complex activity recognition (Section 4) for maritime surveillance.

103 3 Spatio-temporal Link Discovery

104 We begin by providing the definitions and problem setting for spatio-temporal link discovery.
 105 Then, we focus on: (a) topological and proximity relations between the positions of moving
 106 entities and static geographical areas of interest (Section 3.1), and (b) proximity relations
 107 between the positions of moving entities (Section 3.2). The former is a case of streaming to
 108 static link discovery, while the latter is a case of streaming to streaming link discovery.

109 Let $p = (p.x, p.y, p.t)$ denote a spatio-temporal point corresponding to a vessel’s V position
 110 $(p.x, p.y)$ at a given time $p.t$, and \mathcal{A} a dataset that consists of geographical areas represented
 111 as polygons. A polygon $A \in \mathcal{A}$ is represented as a set of points a_i , i.e., $A = \{a_1, a_2, \dots, a_n\}$,
 112 and we write $a_i \in A$ to denote that a_i is included in the representation of A .

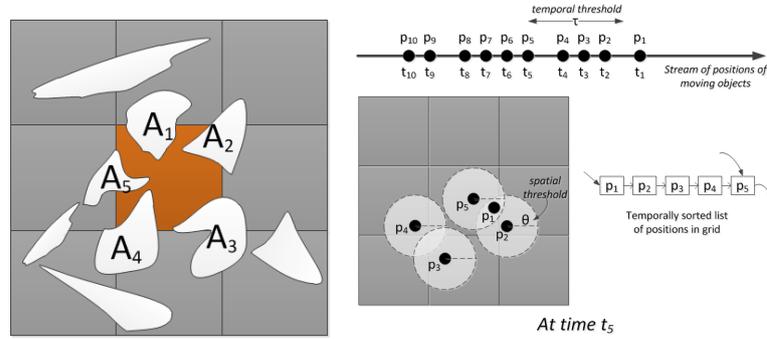
113 Further, let $d(p, p')$ denote the distance between the spatial positions $(p.x, p.y)$ and
 114 $(p'.x, p'.y)$ of two vessels, whereas $t(p, p')$ stands for their temporal difference. Without
 115 loss of generality, in this paper, we employ the Haversine distance function to quantify

116 the spatial distance of two points, whereas $t(p, p') = |p.t - p'.t|$. However, our approach is
 117 readily applicable to other domains, where other distance functions (e.g., Euclidean) are
 118 more appropriate. We abuse notation slightly by denoting $d(p, A)$ the distance between a
 119 point p and an area A (polygon). This is also known as MINDIST, and is defined as the
 120 distance of p to the closest point of A (one point of the perimeter), if p is outside of A .
 121 Otherwise, $d(p, A) = 0$. Below are the formal descriptions of the spatial relations discovered
 122 by the stLD component.

123 ► **Definition 1.** $withinArea(V, A)$: Given a spatio-temporal position p of a vessel V and an
 124 area $A \in \mathcal{A}$, $withinArea(V, A)$ is true, if p is enclosed in A .

125 ► **Definition 2.** $nearbyArea(V, A, \theta)$: Given a spatio-temporal point p of a vessel V , an area
 126 $A \in \mathcal{A}$, and a distance threshold θ , $nearbyArea(V, A, \theta)$ is true, if $d(p, A) \leq \theta$.

127 ► **Definition 3.** $nearby(V_1, V_2, \theta, \tau)$: Given two spatio-temporal points p and p' of vessels V_1
 128 and V_2 respectively, a distance threshold θ , and a temporal threshold τ , $nearby(V_1, V_2, \theta, \tau)$
 129 is true, if $d(p, p') \leq \theta$ and $t(p, p') \leq \tau$.



■ **Figure 2** Link discovery examples.

130 3.1 Stream to Static LD

131 **Discovery of Topological Relations: Within.** Given a *target* dataset T , a *source*
 132 dataset S , and a relation r , the goal of link discovery is to detect the pairs $(\sigma, \tau) \subseteq S \times T$,
 133 where $\sigma \in S$ and $\tau \in T$, s.t. (σ, τ) satisfies r . Consider the case of relation ‘within’ between
 134 a moving entity p , whose current spatio-temporal position is streamed into our system, and
 135 a set of static geographical areas of interest $\mathcal{A} = \{A_1, \dots, A_n\}$. The aim of link discovery is
 136 to identify all areas $A_i \in \mathcal{A}$ which enclose the spatial position $(p.x, p.y)$ of p .

137 A brute force link discovery algorithm would have to perform the geometrical test between
 138 p and all areas in \mathcal{A} , thereby performing $O(n)$ comparisons, where $n = |\mathcal{A}|$. To avoid this
 139 prohibitively expensive cost, the state-of-the-art LD methods (e.g., [16, 22]) employ a *space*
 140 *tiling* approach, which essentially partitions the space in cells and assigns each area to its
 141 overlapping cells. Then, to compute the relation “within”, the position of a vessel is compared
 142 only against those (say c) areas overlapping the cell, thus resulting in $O(c)$ comparisons, and
 143 typically $c \ll n$. Practically, this method belongs to the filter-and-refine paradigm, where

Algorithm 1 Spatio-temporal LD algorithm for relation “within” using mask.

```

1: Input: Grid cells  $C = \{c_1, \dots, c_m\}$ , Areas  $\mathcal{A} = \{A_1, \dots, A_n\}$ , position  $p (p.x, p.y)$ 
2: Output: Subset of areas  $\mathcal{A}^w \subseteq \mathcal{A}$  that enclose the position  $(p.x, p.y)$  of  $p$ 
3: Requires: Grid has been constructed and areas have been assigned to overlapping cells
4:  $\mathcal{A}^w \leftarrow \emptyset$ 
5: locate cell  $c_i$  that encloses  $p$ 
6: if within( $p$ ,  $mask(c_i)$ ) then
7:   return  $\mathcal{A}^w$ 
8: else
9:   for each  $A_j \in c_i$  do
10:    if within( $p$ ,  $A_j$ ) then
11:       $\mathcal{A}^w \leftarrow \mathcal{A}^w \cup A_j$ 
12: return  $\mathcal{A}^w$ 

```

144 in the filtering step only a small set of c (out of n) candidate areas are identified, and in
 145 the refinement step the candidates are examined one-by-one to discover the subset of areas
 146 actually enclosing the vessel. Since the refinement step is the most costly processing part,
 147 any efficient link discovery algorithm should minimize the number of candidates.

148 In several applications of spatial link discovery, such as the maritime domain, grid cells
 149 contain a significant amount of “empty space”, namely the space of the cell that overlaps
 150 with no areas. Our observation is that positions of moving vessels located in the empty space
 151 of a cell induce high processing cost, as they must be compared to all areas in the cell in
 152 vain, since no “within” links can be produced. Motivated by this observation, we propose a
 153 technique to explicitly represent the empty space within cells as yet another area. Thus, for
 154 each grid cell, we construct an artificial area called *mask*, which is defined as the difference
 155 between the cell and the union of areas overlapping with the cell, i.e. $mask = c - (c \cap \bigcup(A)_i)$.
 156 Notice that the intersection with c before computing the set difference is only used to cover
 157 the case where areas are larger the cells. Figure 2a shows an example of the mask of a
 158 cell; the middle cell overlaps with areas $\{A_1, \dots, A_5\}$, and the mask of the cell is the area
 159 represented in orange color.

160 Having the mask of a cell as yet another area, we can devise an efficient algorithm for
 161 link discovery that eagerly avoids comparisons to areas for positions located in the empty
 162 space. In practice, after we identify the enclosing cell of a position of a vessel, we first
 163 compare it to the mask of the cell, to check if it is enclosed in the empty space. If this single
 164 comparison returns true, we stop processing this position, thereby saving c comparisons (c
 165 denotes the average number of areas in a cell). For the typical case where a cell contains
 166 several areas, this technique can save significant computational cost, as will be demonstrated
 167 in the empirical analysis presented in Section 5.

168 Algorithm 1 presents the pseudo-code for discovering a ‘within’ link between the position
 169 p of a vessel and the areas that enclose it. As a prerequisite, the grid has already been
 170 constructed and the static areas in the dataset R have been assigned to cells. This is a
 171 pre-processing step. In the first step, the cell c_i that encloses p is determined (line 5). This
 172 operation is performed in constant time $O(1)$ in the case of equi-grids. Then, we check if p is
 173 contained in the mask $mask(c_i)$ of cell c_i (line 6). If it is contained, no further processing is
 174 required, and the algorithm terminates returning the empty set. If it is not contained, then
 175 we check for containment against all areas A_j in cell c_i (line 9). For those areas A_j that
 176 contain p , we add them to the result set \mathcal{A}^w (line 11) and eventually return them as result.

177 The lines 9–11 of Algorithm 1 are processed in parallel, i.e., each iteration in the for-loop
 178 is carried out by a different thread/‘worker’. The number of concurrent workers is usually a
 179 predefined constant to allow uninterrupted system operation (in our experiments we employ
 180 8 workers). We have enabled multi-thread processing using a *pool of tasks*, populated with
 181 the refinement tasks of $within(p, A_j)$. As soon as a worker is available and the pool contains
 182 tasks, the next task is selected and assigned to the worker for processing. Also, the algorithm
 183 is amenable to parallelization beyond the scope of a single machine, by simply partitioning
 184 the stream of positions of vessels to the available processing nodes, each of which runs an
 185 instance of Algorithm 1 on an off-line constructed grid.

186 **Discovery of Proximity Relations: Nearby.** The above technique is applicable also
 187 for link discovery of a proximity relation, such as the ‘nearby’ relation, between vessel position
 188 p and a set of static areas \mathcal{A} . The ‘nearby’ relation is defined using a spatial threshold θ ,
 189 and retrieves the subset of areas in \mathcal{A} that are located at most at distance θ from p .

190 The technique of using the mask needs to be slightly adjusted to work in the case of
 191 relation ‘nearby’. The main adjustment concerns the way the mask of each cell is computed.
 192 We expand each area A_i by θ and then the cell’s mask is computed as previously, only using
 193 the expanded areas (A_i^θ) instead of the actual areas A_i . To differentiate this mask of a cell
 194 c_i from the one used in the previous algorithm, we denote it by $mask^\theta(c_i)$.

195 The LD algorithm for relation ‘nearby’ operates as follows. The grid is constructed
 196 exactly as before, using the original areas $\{A_i\}$. First, the cell c_i that encloses p is located.
 197 If $mask^\theta(c_i)$ contains p , then we can safely stop processing, since no area is nearby p . This
 198 is because $mask^\theta(c_i)$ has been constructed based on expanded areas. If this pruning is not
 199 successful, we need to examine all cells $c_i \in C'$ that overlap with a circle centered at p with
 200 radius θ , since they may contain results. We examine each area A_j in c_i , and if the distance
 201 of p to A_j is lower or equal to the threshold θ , then A_j is added to the result \mathcal{A}^n . To avoid
 202 computing the distance of p to an area A_j multiple times (due to A_j assignment to multiple
 203 cells), we maintain the already examined areas in a set (\mathcal{P}). In summary, the algorithm
 204 avoids processing points that would safely produce no results, due to the use of the mask
 205 technique.

206 3.2 Stream to Stream LD

207 Streaming link discovery of ‘nearby’ relations between positions of moving vessels requires
 208 meticulous use of the available memory, since it is not feasible to store the complete data
 209 stream in memory. Our solution relies on the use of a grid data structure on the spatial
 210 domain, similarly to the case of Section 3.1. The grid is used to maintain the positions
 211 arriving in the stream. When a new vessel position p arrives in the stream, in the filtering
 212 step, we examine the cells that intersect with a circle centered at $(p.x, p.y)$ with radius θ ,
 213 and retrieve all vessel positions p'_i that satisfy the spatial constraint, i.e., $d(p, p'_i) \leq \theta$. Then,
 214 in the refinement step, we identify the subset of vessel positions $\{p'_i\}$ that in addition satisfy
 215 the temporal constraint, i.e., $t(p, p'_i) \leq \tau$. This solution avoids the exhaustive comparison of
 216 p to all other positions that have arrived before p . Algorithm 2 is invoked for each incoming
 217 p and describes this solution.

218 However, in order to manage the available memory effectively, we need to find a suitable
 219 way to clean up the grid, since vessel positions that have arrived before more than τ time
 220 units will never produce a ‘nearby’ link to a new vessel position. A second reason to perform
 221 this cleaning operation is efficiency. If no cleaning were performed, then too many (old) vessel
 222 positions would be retrieved that satisfy the spatial constraint, but would be eliminated due
 223 to the temporal constraint, leading to wasteful processing.

Algorithm 2 Spatio-temporal LD algorithm for relation ‘nearby’ between vessels.

```

1: Input: Grid cells  $C = \{c_1, \dots, c_m\}$ , vessel position  $p (p.x, p.y)$ , thresholds  $\theta, \tau$ 
2: Output: Set  $R$  of pairs of vessel positions  $(p, p')$  satisfying Definition 3
3:  $R \leftarrow \emptyset$ 
4: locate cells  $C$  that overlap with circle at  $p$  with radius  $\theta$ 
5: for each  $c_j \in C$  do
6:   for each  $p'_i \in c_j$  do
7:     if  $d(p, p'_i) \leq \theta$  then
8:       if  $t(p, p'_i) \leq \tau$  then
9:          $R \leftarrow R \cup (p, p'_i)$ 
10: add  $p$  to grid  $C$ 
11: return  $R$ 

```

224 A naive approach for cleaning would be to scan all grid cells (set at time t_{now}) and delete
225 vessel positions p whose timestamp $p.t$ is more than τ units ago, i.e., $t_{now} - p.t > \tau$. However,
226 this operation has linear complexity wrt. the number of positions in the grid, and incurs
227 non-negligible overhead, as it must be performed during stream processing. To address this
228 problem, we introduce an auxiliary, bookkeeping data structure that efficiently detects the
229 vessel positions that need to be deleted. For this purpose, we maintain a list of pointers
230 to the vessel positions in the grid. The list is in temporal order, since vessel positions are
231 inserted in the list in the order in which they arrive in the stream. Then, cleaning can be
232 performed efficiently, by traversing the list and deleting vessel positions (from the grid and
233 list) until a position p is found with timestamp $t_{now} - p.t \leq \tau$. The maintenance cost of this
234 list is small; insertions have $O(1)$ cost, whereas deletions have linear cost to the number of
235 vessel positions that need to be deleted.

236 Figure 2b shows an example of the bookkeeping structure. Assuming that $t_{now} = t_5$,
237 then the grid and the list contain the five vessel positions p_1, \dots, p_5 , as depicted. In case
238 no cleaning is performed, then p_1 would be identified as candidate for ‘nearby’ to p_5 , which
239 would then be rejected due to the temporal threshold, since $p_5.t - p_1.t > \tau$. In case cleaning
240 has already been performed, p_1 would have already been deleted from the grid, thus we
241 would have avoided the cost of examining p_1 .

242 An interesting issue that deserves further discussion is the frequency of performing the
243 cleaning operation. An *eager* strategy could invoke the cleaning operation prior to processing
244 every new vessel position that arrives in the stream. This would guarantee that all vessel
245 positions satisfying the spatial distance threshold would be results, without the need to check
246 the temporal threshold. However, the eager strategy would result in paying the overhead of
247 grid cleanup for every new position. Another approach is to follow a *lazy* strategy, where
248 the cleaning operation is invoked periodically, thus limiting the induced overhead at the
249 expense of retrieving some candidate vessel positions that may be rejected by the temporal
250 threshold. We evaluate the frequency of performing the grid cleanup in Section 5 presenting
251 the empirical evaluation.

252 4 Complex Activity Recognition

253 The complex activity recognition (CAR) component of the datAcron prototype consumes
254 the stream of spatial relations detected by stLD, as well as the critical points expressing
255 compressed trajectories, to detect various types of vessel activity (see Section 2). The upper
256 part (above the two horizontal lines) of Table 1 presents the input to the CAR component.

■ **Table 1** Complex Activity Recognition: Input events are presented above the two horizontal lines, while the output stream is presented below these lines. The input events above the single horizontal line are detected by stLD, while the remaining input events are emitted by the trajectory compression component. All input events, except $nearBy(V_1, V_2)$, are instantaneous, while all output activities are durative.

Event/Activity	Description
$withinArea(V, A)$	A vessel V is within some area A of interest
$nearbyArea(V, A)$	A vessel V is close to some area A of interest
$nearPorts(V)$	A vessel V is close to one or more ports
$nearby(V_1, V_2)$	Two vessels V_1 and V_2 are close to each other
$gapStart(V)$	A vessel V stopped sending position signals
$gapEnd(V)$	A vessel V resumed sending position signals
$slowMotionStart(V)$	A vessel started moving at a low speed
$slowMotionEnd(V)$	A vessel stopped moving at a low speed
$stopStart(V)$	A vessel started being idle
$stopEnd(V)$	A vessel stopped being idle
$changeInSpeedStart(V)$	A vessel started changing its speed
$changeInSpeedEnd(V)$	A vessel stopped changing its speed
$changeInHeading(V)$	A vessel changed its heading
$gap(V)$	A vessel V has a communication gap in the open sea
$stopped(V)$	A vessel V is stopped in the open sea
$slowMotion(V)$	A vessel V moves slowly in the open sea
$illegalFishing(V)$	A vessel V is potentially engaged in illegal fishing
$loitering(V)$	A vessel V is suspiciously idle
$rendezVous(V_1, V_2)$	Two vessels V_1 and V_2 are having a rendez-vous
$highSpeedIn(V, AreaType)$	A vessel V exceeds the speed limit of an area of $AreaType$

257 The output of this component, i.e. the list of recognized activities, specified in collaboration
 258 with the domain experts of datAcron, is presented in the lower part of Table 1.

259 4.1 Run-Time Event Calculus

260 The CAR component of datAcron is based on the ‘Event Calculus for Run-Time reasoning’
 261 (RTEC) [3]. RTEC is a Prolog implementation of the Event Calculus [12], designed to compute
 262 continuous narrative assimilation queries for pattern matching on data streams. RTEC has a
 263 formal, declarative semantics—complex (vessel) activity formalisations are (locally) stratified
 264 logic programs [21]. Moreover, RTEC includes various optimisation techniques for efficient
 265 pattern matching, such as ‘windowing’, whereby all input events that took place prior to the
 266 current window are discarded/‘forgotten’. Details about the reasoning algorithms of RTEC,
 267 including a complexity analysis, may be found in [3]. In what follows, we illustrate the use
 268 of RTEC for specifying maritime activities, focusing on the effects of stLD on CAR.

269 The time model in RTEC is linear and includes integer time-points. Variables start with
 270 an upper-case letter, while predicates and constants start with a lower-case letter. Where F
 271 is a *fluent*—a property that is allowed to have different values at different points in time—the
 272 term $F = V$ denotes that fluent F has value V . An *event description* in RTEC includes
 273 rules that define the event instances with the use of the `happensAt` predicate, the effects of
 274 events on fluents with the use of the `initiatedAt` and `terminatedAt` predicates, and the values

■ **Table 2** Main predicates of RTEC.

Predicate	Meaning
$\text{happensAt}(E, T)$	Event E occurs at time T
$\text{holdsAt}(F = V, T)$	The value of fluent F is V at time T
$\text{holdsFor}(F = V, I)$	I is the list of the maximal intervals for which $F = V$ holds continuously
$\text{initiatedAt}(F = V, T)$	At time T a period of time for which $F = V$ is initiated
$\text{terminatedAt}(F = V, T)$	At time T a period of time for which $F = V$ is terminated
$\text{union_all}(L, I)$	I is the list of maximal intervals produced by the union of the lists of maximal intervals of list L
$\text{intersect_all}(L, I)$	I is the list of maximal intervals produced by the intersection of the lists of maximal intervals of list L
$\text{relative_complement_all}(I', L, I)$	I is the list of maximal intervals produced by the relative complement of the list of maximal intervals I' with respect to every list of maximal intervals of list L

275 of the fluents with the use of the `holdsAt` and `holdsFor` predicates. Table 2 summarizes the
 276 main predicates of RTEC. Complex activities are typically durative (see the lower part of
 277 Table 2), thus in CAR the task generally is to compute the maximal intervals for which a
 278 fluent expressing a complex activity has a particular value continuously. Below, we discuss
 279 the representation of fluents/complex maritime activities, and briefly present the way we
 280 compute their maximal intervals.

281 4.2 Maritime Pattern Representation

282 For a fluent F , $F = V$ holds at a particular time-point T if $F = V$ has been initiated by
 283 an event at some time-point earlier than T , and has not been terminated at some other
 284 time-point in the meantime. This is an implementation of the law of *inertia*. The time-points
 285 at which $F = V$ is initiated (terminated) are computed with the use of `initiatedAt` (`terminatedAt`)
 286 rules. Consider the following example:

$$\begin{aligned}
 & \text{initiatedAt}(\text{gap}(V) = \text{true}, T) \leftarrow \\
 & \quad \text{happensAt}(\text{gapStart}(V), T), \text{ not happensAt}(\text{nearPorts}(V), T) \\
 & \text{terminatedAt}(\text{gap}(V) = \text{true}, T) \leftarrow \\
 & \quad \text{happensAt}(\text{gapEnd}(V), T)
 \end{aligned} \tag{1}$$

289 $\text{gap}(V)$ is a Boolean fluent denoting a communication gap for some vessel V , i.e. V stops
 290 transmitting AIS messages in the open sea. In some cases, the absence of AIS messages is
 291 suspicious and thus we need to record it. $\text{gapStart}(V)$ and $\text{gapEnd}(V)$ are instantaneous
 292 critical events indicating, respectively, the time-points in which a vessel V stops and resumes
 293 sending AIS messages (see Table 1). ‘not’ is negation by failure. $\text{nearPorts}(V)$ is an event
 294 emitted by stLD when vessel V is close to a port. Thus, rule-set (1) states that $\text{gap}(V) = \text{true}$
 295 is initiated if the trajectory compression component reports a gapStart event for V , and
 296 stLD does not report that V is close to a port. Furthermore, $\text{gap}(V) = \text{true}$ is terminated
 297 when V resumes communications. Given rule-set (1), RTEC computes the maximal intervals
 298 I for which $\text{gap}(V) = \text{true}$ holds continuously, i.e. $\text{holdsFor}(\text{gap}(V) = \text{true}, I)$, by finding all
 299 time-points T_s at which $\text{gap}(V) = \text{true}$ is initiated, and then, for each T_s , computing the first
 300 time-point T_e after T_s at which $\text{gap}(V) = \text{true}$ is terminated.

301 Most of the maritime patterns (defined in collaboration with domain experts) concern
 302 vessel activity close to, or within, some area of interest, such a Natura area. To simplify the
 303 structure of such patterns, we defined the auxiliary fluent *inArea* as follows:

$$\begin{aligned}
 & \text{initiatedAt}(\text{inArea}(V, \text{protected}) = \text{true}, T) \leftarrow \\
 & \quad \text{happensAt}(\text{withinArea}(V, A), T), \text{protected}(A) \\
 & \text{terminatedAt}(\text{inArea}(V, \text{protected}) = \text{true}, T) \leftarrow \\
 304 & \quad \text{happensAt}(\text{nearbyArea}(V, A), T), \text{protected}(A) \tag{2} \\
 & \text{terminatedAt}(\text{inArea}(V, \text{protected}) = \text{true}, T) \leftarrow \\
 305 & \quad \text{happensAt}(\text{start}(\text{gap}(V) = \text{true}), T)
 \end{aligned}$$

306 *inArea*(*V*, *AreaType*) records the maximal intervals in which a vessel *V* is in an area of
 307 some type. *withinArea*(*V*, *A*) and *nearbyArea*(*V*, *A*) are spatial events emitted by stLD,
 308 stating, respectively, that *V* is within, or close to area *A*. *protected*(*A*) is an atemporal
 309 predicate that stores protected areas. *start*(*F* = *V*) (respectively *end*(*F* = *V*)) is a built-in
 310 RTEC event taking place at each starting (ending) point of each maximal interval for which
 311 *F* = *V* holds continuously. According to rule-set (2), *inArea*(*V*, *protected*) = true is initiated
 312 when stLD detects a *withinArea*(*V*, *A*) event for vessel *V* and protected area *A*. Furthermore,
 313 *inArea*(*V*, *protected*) = true is terminated when there is information that *V* has exited the
 314 area, i.e. when a *nearbyArea*(*V*, *A*) event is emitted for the protected area *A*, or when *V*
 315 stops transmitting position signals (in this case we do not know the location of *V*).

316 Similar rule-sets define *inArea* for other types of area. Moreover, *inArea* is represented as
 317 a Boolean fluent since areas of different type may overlap, and thus a vessel may be within
 318 various types of area at the same time.

319 In previous work, RTEC performed both temporal and atemporal reasoning for CAR
 320 [20]. For example, RTEC performed spatial calculations to determine whether a vessel is
 321 within some area of interest or close to a port. In this work, all spatial relations necessary
 322 for CAR are computed by stLD. Thus, the evaluation of the *initiatedAt* rule of rule-set (2),
 323 for example, is reduced to checking for (*withinArea*) facts (in the input stream). This way,
 324 RTEC is used only for temporal reasoning for which it is optimized.

325 The absence of the *nearbyArea* relation in earlier work [20] did not allow us to compute
 326 the *intervals* during which a vessel is in some area. These intervals allow us to develop more
 327 accurate patterns of maritime activity—consider e.g. illegal fishing:

$$\begin{aligned}
 & \text{holdsFor}(\text{illegalFishing}(V) = \text{true}, I) \leftarrow \\
 328 & \quad \text{fishing}(V), \text{holdsFor}(\text{slowMotion}(V) = \text{true}, I_1), \\
 & \quad \text{holdsFor}(\text{inArea}(V, \text{protected}) = \text{true}, I_2), \text{intersect_all}([I_1, I_2], I)
 \end{aligned} \tag{3}$$

330 In addition to the domain-independent definition of *holdsFor*, an event description may
 331 include domain-specific *holdsFor* rules, such as rule (3), used to define the values of a fluent
 332 *F*, e.g. *illegalFishing*, in terms of a Boolean function on the values of other fluents. Domain-
 333 specific *holdsFor* rules make use of interval manipulation constructs: *union_all*, *intersect_all*,
 334 and *relative_complement_all*. These are presented in Table 2. *fishing* is an atemporal predicate
 335 recording fishing vessels. According to rule (3), the list *I* of maximal intervals during which
 336 a fishing vessel *V* is said to be engaged in illegal fishing is computed by determining the
 337 list *I*₁ of maximal intervals during which *V* is moving in slow motion in the open sea (these
 338 intervals are computed given the instantaneous *slowMotionStart* and *slowMotionEnd* critical
 339 events), the list *I*₂ of maximal intervals during which *V* is in a protected area, and then
 340 calculating the list *I* representing the intersections of the maximal intervals in *I*₁ and *I*₂.

341 The interval manipulation constructs of RTEC support the following type of definition: for
 342 all time-points *T*, *F* = *V* holds at *T* if and only if some Boolean combination of fluent-value

343 pairs holds at T . For a wide range of fluents, this is a much more concise definition than the
 344 traditional style of Event Calculus representation, i.e. identifying the various conditions under
 345 which the fluent is initiated and terminated so that maximal intervals can then be computed
 346 using the domain-independent `holdsFor`. For instance, the representation of *illegalFishing* by
 347 means of `initiatedAt` and `terminatedAt` predicates requires four rules.

348 In addition to patterns for individual vessels, the knowledge base of the CAR component
 349 of `datAcron` includes formalizations of activities for pairs of vessels. Consider *rendezVous*:

$$\begin{aligned}
 & \text{holdsFor}(\text{rendezVous}(V_1, V_2) = \text{true}, I) \leftarrow \\
 & \quad \text{holdsFor}(\text{nearby}(V_1, V_2) = \text{true}, I_1), \text{ holdsFor}(\text{slowMotion}(V_1) = \text{true}, I_2), \\
 350 & \quad \text{holdsFor}(\text{stopped}(V_1) = \text{true}, I_3), \text{ union_all}([I_2, I_3], I_4), \\
 & \quad \text{holdsFor}(\text{slowMotion}(V_2) = \text{true}, I_5), \text{ holdsFor}(\text{stopped}(V_2) = \text{true}, I_6), \\
 351 & \quad \text{union_all}([I_5, I_6], I_7), \text{ intersect_all}([I_1, I_4, I_7], I)
 \end{aligned} \tag{4}$$

352 *nearby*(V_1, V_2) = true is a fluent denoting whether two vessels V_1 and V_2 are close to each
 353 other. This relation is computed by stLD. *stopped*(V) = true expresses that vessel V has
 354 stopped in the open sea. The intervals of this fluent-value pair are computed with the use
 355 of the instantaneous *stopStart* and *stopEnd* critical events (see Table 1). According to rule
 356 (4), vessels V_1 and V_2 are said to have a rendez-vous when they are close to each other, and
 357 each of them is stopped or moving in slow motion in the open sea. In previous work, we
 358 approximated very crudely the *nearby*(V_1, V_2) relation by checking whether V_1 and V_2 are
 359 located within the same cell of a grid. This approximation produced several false negatives
 360 (vessels located close to each other but in different cells) as well as false positives (vessels
 361 located within the same cell but not close to each other). In this work, we can avoid these
 362 issues, due to the efficient spatial relation detection of stLD.

363 5 Experimental Evaluation

364 We present the results of our experimental study using real maritime datasets.

365 5.1 Experimental Setup

366 **Platform.** All experiments presented below were performed on a computer with 8 cores
 367 (Intel(R) Core(TM) i7-7700 CPU @ 3.6GHz) and 16 GB of RAM, running Ubuntu 16.04
 368 LTS 64-bit with Linux Kernel 4.8.0-53-generic, Java 8, and SWI Prolog 7.2.3 for RTEC.

369 **Datasets.** Our main dataset contains AIS kinematic messages from vessels sailing in the
 370 Atlantic Ocean around Brest, France, spanning between 1 October 2015 to 31 March 2016
 371 (<https://zenodo.org/record/1167595#.Ww1Wjp99LJ9>). AIS messages from 5,050 vessels
 372 are compressed by the trajectory synopsis module (see Figure 1), keeping only critical points
 373 (see Table 1). Each critical point is accompanied by mobility information such as speed and
 374 heading. The dataset contains 4,765,647 critical points. These are consumed by stLD which
 375 produces 3,204,206 spatial events that are additionally provided to CAR as input. Spatial
 376 events determine whether a vessel is in or near an area of interest, and if two vessels are close
 377 to each other in space and time.

378 We also employ a set of spatial datasets describing areas and points of interest, as follows:
 379 ■ Natura2000 regions: Natura 2000 is an index of protected regions for biodiversity in the
 380 European Union. It is set up to ensure the survival of Europe’s most valuable species
 381 and habitats, indexing 3,522 polygons.

Granularity	Average Number of Areas in Cell	Cells Constructed
0.5°	5.74	2,852
1°	13.54	858
2°	36.09	272
3°	64.14	147

■ **Table 3** Distribution of areas from target dataset for each grid configuration.

- 382 ■ Fishing areas: This dataset includes 5,076 polygons generated from raster images depicting
383 the fishing intensity in European waters (as reported by European Union).
384 ■ The World Port Index (WPI) including a listing of 3,685 ports throughout the world,
385 describing their location, characteristics, known facilities, and available services. Ports in
386 this dataset are represented as points, and WPI is used for discovering proximity relations
387 between vessel positions and ports.

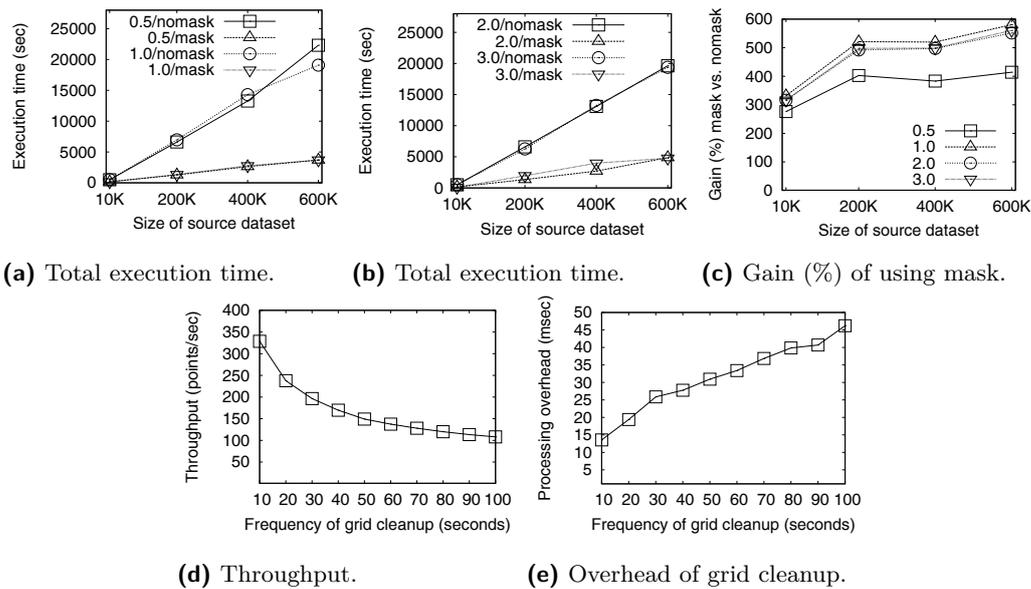
388 **Metrics.** Our main metric is the achieved throughput of the stream reasoning system,
389 assessed by delving into the individual performance of its two constituent components,
390 stLD and CER. Other auxiliary metrics are additionally presented, in order to explain
391 performance, such as processing time and recognition time for stLD and complex activity
392 recognition respectively. We also measure the number of *unrewarding comparisons* of each
393 LD configuration, i.e., the number of comparisons that did not result in a relation.

394 **Parameters.** In the experiments evaluating the performance of link discovery, we vary
395 the input size and granularity of the grid (i.e., cell size). In the case of stream to static LD (i.e.,
396 ‘within’ and ‘nearby’ relations between moving objects and regions), we setup multiple equi-
397 grid configurations varying in granularity $\{0.5^\circ, 1^\circ, 2^\circ, 3^\circ\}$ (degrees in latitude/longitude),
398 using the datasets of Natura2000 and fishing regions, totaling 8,599 regions. An equi-grid of
399 0.5° granularity means that the size of a cell is $0.5^\circ \times 0.5^\circ$. Notice that we construct the
400 grid over the complete static datasets, as typically done in link discovery tasks. We evaluate
401 the performance and scalability of stLD, using subsets of critical points datasets, i.e., $\{10K,$
402 $200K, 400K, 600K\}$ entries. Also, we evaluate the gain obtained by using the mask technique
403 (for the stream to static link discovery method of Section 3.1), and using the bookkeeping
404 technique (for the stream to stream link discovery method of Section 3.2). All topological
405 relations (such as within, but also those required for computing the mask) are provided by
406 the Java Topology Suite (JTS) v.1.13. Finally, we evaluate the effect of varying the number
407 of cores on CAR, as well as the effect of increasing the window size.

408 5.2 Link Discovery

409 **Stream to Static LD.** We refer to the areas organized in a grid as *target* dataset, whereas
410 *source* dataset refers to the dataset of streaming vessel positions. We use as *target* dataset the
411 regions of Natura2000 and fishing areas, totaling 8,599 areas. This static data is organized
412 off-line in main-memory using grid configurations of varying granularity for the complete
413 geographical space covered by the areas. For the streaming *source* dataset, we employ subsets
414 of the critical points dataset of different sizes, in order to evaluate its effect. The distribution
415 of areas for each grid configuration is shown in Table 3. The second column shows the
416 average number of areas in a cell, while the third column indicates the number of non-empty
417 cells that were constructed (i.e., hold at least one of the given areas).

418 Figures 3a–3c show the benefits of using the proposed technique with mask. Figure 3a
419 depicts the total execution time required for processing a subset of critical points, whose size



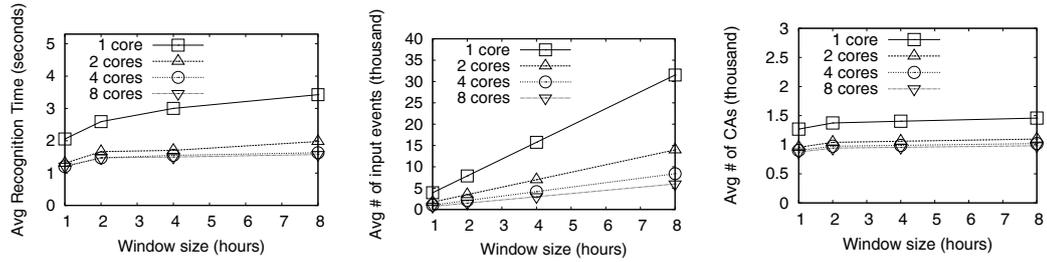
■ **Figure 3** Spatio-temporal LD. Stream to static: (a)–(c); Stream to stream: (d) and (e).

420 is indicated in the x-axis, using two grids with granularity 0.5° and 1.0° respectively. We
 421 observe that the mask technique achieves better execution time for all grid configurations
 422 and input sizes. Figure 3b shows both grid configurations of granularity 2.0° and 3.0° .
 423 In terms of throughput, the use of the mask achieves to increase the throughput up to a
 424 factor of 5, compared to not using the mask. Figure 3c quantifies this gain, as the ratio
 425 of unrewarding comparisons, i.e., the number of unrewarding comparisons without mask,
 426 divided by the number of unrewarding comparisons with mask. For instance, the ratio of
 427 unrewarding comparisons for input size 200K and granularity 0.5° is 402.7%. This result
 428 clearly demonstrates the advantage of using the mask technique.

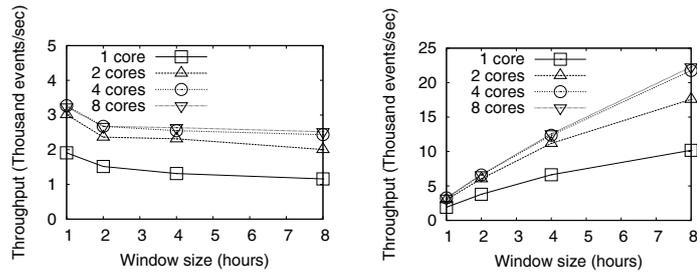
429 **Stream to Stream LD.** We evaluate the effect of the grid cleaning technique on
 430 throughput, using a temporal threshold of 30 seconds. Figure 3d shows how throughput
 431 is affected when varying the frequency of grid cleanup from 10–100 sec. The chart shows
 432 that the more frequent cleanup is performed, the higher the achieved throughput. When the
 433 cleanup is delayed, the processing time spent on comparisons is increased, affecting the total
 434 execution time, and decreasing the average throughput. The next question to be answered
 435 concerns the overhead imposed by frequent grid cleanup. Figure 3e shows the average time
 436 spent for grid cleanup in milliseconds. We observe that the overhead is very small, and
 437 increases when the cleaning is not performed regularly. Recall that the bookkeeping structure
 438 is a list of spatio-temporal positions of vessels ordered by their temporal values, and the
 439 disposal of part of the list on each call allows us to avoid searching every cell of the grid for
 440 ‘expired’ entries. In summary, the cleanup operation has minimal overhead, which makes it
 441 applicable with high frequency, thereby increasing the achieved throughput.

442 5.3 Complex Activity Recognition

443 RTEC performs continuous query processing, computing the maximal intervals of fluents
 444 representing complex maritime activities. At each query time Q_i , the input events that
 445 fall within a specified sliding window ω are taken into consideration. All input events that



(a) Average recognition time. (b) Average number of input events. (c) Average number of recognised activities.



(d) Throughput for slide step $\beta = 1$ hour. (e) Throughput for slide step $\beta = \omega$.

■ **Figure 4** Complex activity recognition.

446 took place before or at $Q_i - \omega$ are discarded/‘forgotten’. This way, the cost of reasoning
 447 is independent of the data stream size. At Q_i , the complex activity intervals computed
 448 by RTEC are those that can be derived from input events that occurred in the interval
 449 $(Q_i - \omega, Q_i]$, as recorded at time Q_i . When the range ω is longer than the slide step β , it
 450 is possible that an input event occurs in the interval $(Q_i - \omega, Q_{i-1}]$ but arrives at RTEC only
 451 after Q_{i-1} ; its effects are taken into account at query time Q_i . Window parameters for ω
 452 and slide step β are defined by the domain experts along with the maritime patterns, since
 453 they reflect the time horizon over which interesting phenomena may be detected. Usually, ω
 454 spans many minutes or even hours for capturing meaningful events across a vessel’s route.

455 Figure 4 presents the experimental results for increasing window sizes ω : 1 hour to 8
 456 hours. Figure 4a presents the average recognition time per window ω , while Figures 4b and 4c
 457 show, respectively, the average number of input events and recognized activities per window.
 458 Recall that Table 1 lists the types of input event and recognized activity. The slide step β
 459 is set to 1 hour. The empirical analysis shows that RTEC is capable of real-time maritime
 460 activity recognition—e.g. RTEC recognizes ($\approx 1,500$) maritime activities given a window of 8
 461 hours ($\approx 31,000$ input events) in less than 3.5 sec, even when operating on a single processing
 462 core of the desktop computer. (The use of multiple cores will be discussed shortly.) Figure
 463 4d presents the throughput. As the window size ω increases, throughput decreases since
 464 the average recognition time per window increases (see Figure 4a). Figure 4e presents the
 465 throughput when the slide step is the same as the window size ($\beta = \omega$), i.e. windows are
 466 non-overlapping. In this case, the average recognition time per window increases, but only
 467 very slightly (due to the higher cost of ‘forgetting’ more past events), and thus not presented
 468 here. (The average numbers of input events and recognized activities per window are also
 469 similar to the case of $\beta = 1$ hour). Figure 4e shows that, as the window size ω increases,
 470 throughput increases. When ω increases, the number of windows/query-times decreases (we

471 have significantly less windows than the case of $\beta = 1$ hour), while the average recognition
472 time per window increases only slightly.

473 We also run RTEC in parallel, by launching different instances of the engine, each
474 operating on a different processing core. Each RTEC instance was responsible for activity
475 recognition in a separate sub-area of the dataset, receiving input events only from vessels in
476 that sub-area. To avoid false negatives on the borders of the sub-areas, we allowed for some
477 overlap. Figure 4 presents the results when 2, 4 and 8 processing cores are used. Figure 4a
478 presents the average recognition of the worst-performing RTEC instance, while Figures 4b
479 and 4c show the average input events and recognised activities for that instance. Figures
480 4d and 4e present the throughput. As shown in Figure 4, the benefits of parallelization can
481 be significant. A more refined segmentation of the dataset into sub-areas, optimizing load
482 allocation, would have had more profound effects on performance.

483 6 Discussion

484 We presented a stream reasoning system for maritime monitoring, which supports complex
485 activity recognition assisted by online spatio-temporal discovery of relations between vessels
486 and areas of interest. Compared to earlier work, the proposed system optimizes the com-
487 putation of spatial relations, leading to improved system performance. Our experimental
488 evaluation on real datasets demonstrated the efficiency of the proposed prototype.

489 Even though the topic of link discovery has attracted much interest lately (see [15] for a
490 recent survey), there is not much work on spatio-temporal link discovery nor on link discovery
491 over streaming data. Our work tackles explicitly these topics. Generic LD frameworks include
492 LIMES [17] and SILK [11]. LIMES [17] is an LD framework for metric spaces that uses the
493 triangular inequality in order to avoid processing all pairs of objects. LIMES employs the
494 concept of exemplars, which are used to represent areas in the multidimensional space, and
495 prunes entire areas (and the respective enclosed entities) during link discovery. SILK [11]
496 is an LD framework proposing a novel blocking method called MultiBlock, which uses a
497 multidimensional index. The spatial LD methods [16, 22] apply grid partitioning (a.k.a.
498 space tiling) on sources A and B, in order to perform efficiently the *filtering step*. Then, in
499 the *refinement step*, different optimizations are employed in order to minimize the number of
500 computations necessary to produce the correct result set. However, all these works focus on
501 topological relations, and it is not straightforward to extend them for proximity relations.

502 Various languages and systems have been proposed in the literature for complex even-
503 t/activity recognition [7, 1]. RTEC has a formal, declarative semantics—activity patterns in
504 RTEC are (locally) stratified logic programs. In contrast, most complex event processing
505 languages, event query languages, data stream processing languages and commercial pro-
506 duction rule systems, rely on an informal and/or procedural semantics [8]. Furthermore,
507 RTEC explicitly represents complex activity intervals (unlike e.g. [9, 8, 4]) and thus avoids
508 the related logical problems [18]. Maritime activities form hierarchies, in the sense that the
509 formulation of one activity is used to define other, higher-level activities. We defined e.g. po-
510 tentially illegal fishing in terms of slow motion (recall rule (3)). In contrast to state-of-the-art
511 recognition systems, such as the Esper engine (<http://www.espertech.com/esper/>) and
512 SASE (<http://sase.cs.umass.edu/>), RTEC can naturally express hierarchical knowledge
513 by means of well-structured specifications, and consequently employ caching techniques to
514 avoid unnecessary re-computations [3]. Concerning the Event Calculus literature, a key
515 feature of RTEC is that it includes a windowing technique. No other Event Calculus system
516 (including [6, 5, 13, 19, 2, 14]) ‘forgets’ or represents concisely the data stream history.

517 — **References** —

- 518 **1** E. Alevizos, A. Skarlatidis, A. Artikis, and G. Paliouras. Probabilistic complex event
519 recognition: A survey. *ACM Comput. Surv.*, 50(5):71:1–71:31, 2017.
- 520 **2** Alexander Artikis and Marek J. Sergot. Executable specification of open multi-agent sys-
521 tems. *Logic Journal of the IGPL*, 18(1):31–65, 2010.
- 522 **3** Alexander Artikis, Marek J. Sergot, and Georgios Paliouras. An event calculus for event
523 recognition. *IEEE Trans. Knowl. Data Eng.*, 27(4):895–908, 2015.
- 524 **4** H. Beck, M. Dao-Tran, and T. Eiter. LARS: A Logic-Based Framework for Analytic
525 Reasoning over Streams. Technical Report INFSYS RR-1843-17-03, Institute of Information
526 Systems, TU Vienna, 2017.
- 527 **5** Iliano Cervesato and Angelo Montanari. A calculus of macro-events: Progress report. In
528 *Proceedings of TIME*, pages 47–58, 2000.
- 529 **6** L. Chittaro and A. Montanari. Efficient temporal reasoning in the cached event calculus.
530 *Computational Intelligence*, 12(3):359–382, 1996.
- 531 **7** G. Cugola and A. Margara. Processing flows of information: From data stream to complex
532 event processing. *ACM Computing Surveys*, 44(3):15, 2012.
- 533 **8** Gianpaolo Cugola and Alessandro Margara. TESLA: a formally defined event specification
534 language. In *Proceedings of DEBS*, pages 50–61, 2010.
- 535 **9** C. Dousson and P. Le Maigat. Chronicle recognition improvement using temporal focusing
536 and hierarchisation. In *Proceedings of IJCAI*, pages 324–329, 2007.
- 537 **10** Bilal Idiri and Aldo Napoli. The automatic identification system of maritime accident risk
538 using rule-based reasoning. In *Proceedings of SoSE*, pages 125–130, 2012.
- 539 **11** Robert Isele, Anja Jentzsch, and Christian Bizer. Efficient multidimensional blocking for
540 link discovery without losing recall. In *Proceedings of WebDB*, 2011.
- 541 **12** Robert A. Kowalski and Marek J. Sergot. A logic-based calculus of events. *New Generation*
542 *Comput.*, 4(1):67–95, 1986.
- 543 **13** R. Miller and M. Shanahan. Some alternative formulations of the event calculus. In
544 *Computational Logic: Logic Programming and Beyond*, LNAI 2408, pages 452–490. 2002.
- 545 **14** M. Montali, F. M. Maggi, F. Chesani, P. Mello, and W. M. P. van der Aalst. Monitoring
546 business constraints with the event calculus. *ACM TIST*, 5(1):17:1–17:30, 2013.
- 547 **15** Markus Nentwig, Michael Hartung, Axel-Cyrille Ngonga Ngomo, and Erhard Rahm. A
548 survey of current link discovery frameworks. *Semantic Web*, 8(3):419–436, 2017.
- 549 **16** Axel-Cyrille Ngonga Ngomo. ORCHID - reduction-ratio-optimal computation of geo-spatial
550 distances for link discovery. In *Proceedings of ISWC*, pages 395–410, 2013.
- 551 **17** Axel-Cyrille Ngonga Ngomo and Sören Auer. LIMES - A time-efficient approach for large-
552 scale link discovery on the web of data. In *Proceedings of IJCAI*, pages 2312–2317, 2011.
- 553 **18** A. Paschke. ECA-RuleML: An approach combining ECA rules with temporal interval-based
554 KR event/action logics and transactional update logics. Technical Report 11, Technische
555 Universität München, 2005.
- 556 **19** Adrian Paschke and Martin Bichler. Knowledge representation concepts for automated
557 SLA management. *Decision Support Systems*, 46(1), 2008.
- 558 **20** K. Patroumpas, E. Alevizos, A. Artikis, M. Vodas, N. Pelekis, and Y. Theodoridis. Online
559 event recognition from moving vessel trajectories. *GeoInformatica*, 21(2):389–427, 2017.
- 560 **21** T. Przymusiński. On the declarative semantics of stratified deductive databases and logic
561 programs. In *Foundations of Deductive Databases and Logic Programming*. Morgan, 1987.
- 562 **22** Mohamed Ahmed Sherif, Kevin Dreßler, Panayiotis Smeros, and Axel-Cyrille Ngonga
563 Ngomo. Radon - rapid discovery of topological relations. In *Proceedings of AAAI*, pages
564 175–181, 2017.