

A Generic Geometrical Constraint Kernel in Space and Time for Handling Polymorphic k -Dimensional Objects

N. Beldiceanu¹, M. Carlsson², E. Poder¹, R. Sadek¹, and C. Truchet³

¹ École des Mines de Nantes, LINA FRE CNRS 2729, FR-44307 Nantes, France
{Nicolas.Beldiceanu, Emmanuel.Poder, Rida.Sadek}@emn.fr

² SICS, P.O. Box 1263, SE-164 29 Kista, Sweden

Mats.Carlsson@sics.se

³ Université de Nantes, LINA FRE CNRS 2729, FR-44322 Nantes, France
Charlotte.Truchet@univ-nantes.fr

Abstract. This paper introduces a geometrical constraint kernel for handling the location in space and time of polymorphic k -dimensional objects subject to various geometrical and time constraints. The constraint kernel is generic in the sense that one of its parameters is a set of constraints on subsets of the objects. These constraints are handled globally by the kernel.

We first illustrate how to model several placement problems with the constraint kernel. We then explain how new constraints can be introduced and plugged into the kernel. Based on these interfaces, we develop a generic k -dimensional lexicographic sweep algorithm for filtering the attributes of an object (i.e., its shape and the coordinates of its origin as well as its start, duration and end in time) according to all constraints where the object occurs. Experiments involving up to hundreds of thousands of objects and 1 million integer variables are provided in 2, 3 and 4 dimensions, both for simple shapes (i.e., rectangles, parallelepipeds) and for more complex shapes.

1 Introduction and Presentation of the Kernel

This paper introduces a constraint kernel $geost(k, \mathcal{O}, \mathcal{S}, \mathcal{C})$ for handling in a generic way a variety of geometrical constraints \mathcal{C} in space and time between polymorphic k -dimensional objects \mathcal{O} ($k \in \mathbb{N}^+$), where each object takes a shape among a set of shapes described by \mathcal{S} during a given time interval and at a given position in space. This line of research can be seen as a continuation and generalisation of previous work on non-overlapping parallelepipeds [1–4].

Each shape is defined as a finite set of shifted boxes, where each shifted box is described by a box in a k -dimensional space at a given offset (from the origin of the shape) with given sizes. More precisely, a *shifted box* $s = sbox(sid, t[], l[]) \in \mathcal{S}$ is an entity defined by its shape id $s.sid$, shift offset $s.t[d]$, $0 \leq d < k$, and sizes $s.l[d]$ ($s.l[d] > 0$, $0 \leq d < k$). All attributes of a shifted box are integer values. Then, a *shape* is defined as the union of shifted boxes sharing the same shape id. Each *object* $o = object(id, sid, x[], start, duration, end) \in \mathcal{O}$ is an entity defined by its unique object id $o.id$, shape id $o.sid$, origin $o.x[d]$, $0 \leq d < k$, start in time $o.start$, duration in time

$o.duration$ ($o.duration \geq 0$) and end in time $o.end$.¹ All these attributes correspond to domain variables.² Typical constraints from the list of constraints \mathcal{C} can express, for instance, the fact that a given subset of objects from \mathcal{O} do not pairwise overlap or that they are all included within a given bounding box. Constraints always have two first arguments \mathcal{A}_i and \mathcal{O}_i (followed by possibly some additional arguments) which resp. specify:

- A list of distinct dimensions (integers in $[0, k - 1]$) that the constraint considers.
- A list of identifiers of the objects to which the constraint applies.

Example 1. Assume we have a 3D placement problem (i.e., $k = 3$) involving a set of parallelepipeds \mathcal{P} and one subset \mathcal{P}' of \mathcal{P} , where we want to express the fact that (1) no parallelepipeds of \mathcal{P} should overlap, and (2) no parallelepipeds of \mathcal{P}' should be piled. Constraints (1) and (2) resp. correspond to $non-overlapping([0, 1, 2], \mathcal{P})$ and to $non-overlapping([0, 1], \mathcal{P}')$. Within the first $non-overlapping$ constraint, the argument $[0, 1, 2]$ expresses the fact that we consider a non-overlapping constraint according to dimensions 0, 1 and 2 (i.e., given any pair of parallelepipeds p' and p'' of \mathcal{P} there should exist at least one dimension d ($d \in \{0, 1, 2\}$) where the projections of p' and p'' on d do not overlap). Similarly, the argument $[0, 1]$ of the second non-overlapping constraint expresses the fact that, given any pair of parallelepipeds p' and p'' of \mathcal{P}' , there should exist at least one dimension d ($d \in \{0, 1\}$) where p' and p'' do not overlap).

$geost(k, \mathcal{O}, \mathcal{S}, \mathcal{C})$ is defined in the following way: given a constraint $ctr_i(\mathcal{A}_i, \mathcal{O}_i)$ from the list of constraints \mathcal{C} between a subset of objects $\mathcal{O}_i \subseteq \mathcal{O}$ according to the attributes \mathcal{A}_i , let \mathcal{MC}_i denote the sets of cliques stemming from the objects of \mathcal{O}_i that all overlap in time.³ The constraints of $geost(k, \mathcal{O}, \mathcal{S}, \mathcal{C})$ hold if and only if $\forall ctr_i \in \mathcal{C}, \forall \mathcal{O}_{\mathcal{MC}_i} \in \mathcal{MC}_i : ctr_i(\mathcal{A}_i, \mathcal{O}_{\mathcal{MC}_i})$ holds.

Example 2. Fig. 1 presents a typical example of a dynamic 2D placement problem where one has to place four objects, in time and within a given box, so that objects that overlap in time do not overlap. Parts (A), (B), (C) and (D) resp. represent the potential shapes associated with the four objects to place, where the origin of each object is represented by a black square ■. Part (E) shows the position of the four objects of the example as the time varies, where the first, second, third and fourth objects were resp. assigned shapes S_1, S_5, S_8 and S_9 :

- During the first time interval $[2, 9]$ we have only object O_1 at position $(1, 2)$.
- Then, at instant 10 objects O_2 and O_3 both appear. Their origins are resp. placed at positions $(2, 1)$ and $(4, 1)$.
- At instant 14 object O_1 disappears and is replaced by object O_4 . The origin of O_4 is fixed at position $(1, 1)$. Finally, at instant 22 all three objects O_2, O_3 and O_4 disappear.

The corresponding arguments are:

¹ The time dimension is treated specially since the $duration$ attribute may not be fixed, which is not the case for the sizes of a shifted box. Also, the geometrical constraints only apply on objects that intersect in time.

² A *domain variable* v is a variable ranging over a finite set of integers denoted by $\text{dom}(v)$; let \underline{v} and \bar{v} resp. denote the minimum and maximum possible values for v .

³ In fact, these cliques (of an interval graph) are only used for defining the declarative semantics of $geost$'s constraints.

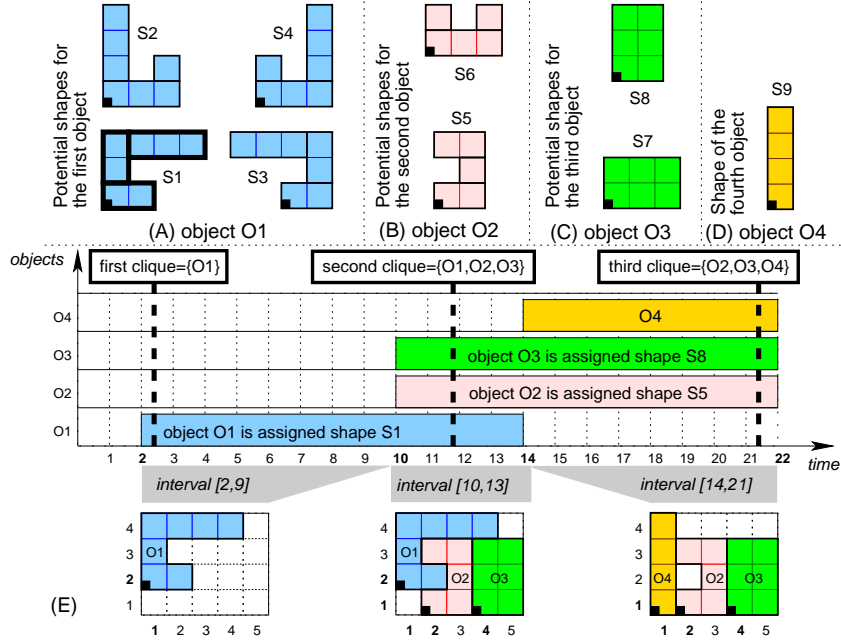


Fig. 1. Example with 4 objects, 9 shapes, one *non-overlapping* and one *included* constraints

```

01 geost(2,
02   [object(1,1,[1,2], 2,12,14), object(2,5,[2,1],10,12,22),
03     object(3,8,[4,1],10,12,22), object(4,9,[1,1],14, 8,22)],
04   [sbox(1,[0,0],[2,1]), sbox(1,[0,1],[1,2]), sbox(1,[1,2],[3,1]),
05     sbox(2,[0,0],[3,1]), sbox(2,[0,1],[1,3]), sbox(2,[2,1],[1,1]),
06     sbox(3,[0,0],[2,1]), sbox(3,[1,1],[1,2]), sbox(3,[2,2],[3,1]),
07     sbox(4,[0,0],[3,1]), sbox(4,[0,1],[1,1]), sbox(4,[2,1],[1,3]),
08     sbox(5,[0,0],[2,1]), sbox(5,[1,1],[1,1]), sbox(5,[0,2],[2,1]),
09     sbox(6,[0,0],[3,1]), sbox(6,[0,1],[1,1]), sbox(6,[2,1],[1,1]),
10     sbox(7,[0,0],[3,2]), sbox(8,[0,0],[2,3]), sbox(9,[0,0],[1,4])],
11   [non-overlapping([0,1],[1,2,3,4]), included([0,1],[1,2,3,4],[1,1],[5,4])])

```

Its first argument 2 is the number of dimensions of the placement space we consider. Its second and third arguments resp. describe the four objects and the shifted boxes of the nine shapes we have. For instance, the 3 boxes of shape S_1 (depicted by 3 thick rectangles in Part (A) of Fig. 1) respectively correspond to the 3 boxes declared at line 04 of the example. Finally, its last argument gives the list of geometrical constraints imposed by *geost*: the first constraint expresses a non-overlapping constraint between the four objects, while the second constraint imposes the four objects to be located within the box containing all points (x, y) such that $1 \leq x \leq 1 + 5 - 1$ and $1 \leq y \leq 1 + 4 - 1$. The constraints of *geost* hold since the four objects do not simultaneously overlap in time and in space and since they are completely included within the previous box (i.e., see Part (E) of Fig. 1).

Within the scope of *geost*($k, \mathcal{O}, \mathcal{S}, \mathcal{C}$), this paper presents a filtering algorithm that prunes the domain of each attribute of every object $o = \text{object}(id, sid, x[], start, duration, end) \in \mathcal{O}$. All values found infeasible are deleted from the shape attribute *sid*; for the other attributes (i.e., the origin $x[]$, the start, the duration and the end), the minimum and maximum are adjusted.

The paper is organised as follows. Section 2 provides an overview of placement problems that can be modelled with the constraints currently available in *geost*. Section 3 presents the overall architecture of the geometrical kernel. It explains how to define geometrical constraints in terms of a programming interface by the geometrical kernel. Section 4 focusses on the main contribution of this paper: a multi-dimensional lexicographic sweep algorithm used for filtering the attributes of an object of *geost*. Section 5 evaluates the scalability of the *geost* kernel as well as its ability to deal with a variety of specific placement problems. Before we conclude, Section 6 compares *geost* with related work and suggests future directions.

2 Modelling Problems with *geost*

As illustrated by Fig. 2 in the context of *non-overlapping*, *geost* allows to model directly a large number of placement problems:

- Case (A) corresponds to a non-overlapping constraint among three segments.
- The second and third cases (B,C) correspond to a non-overlapping constraint between rectangles where (B) is a special case where the sizes of all rectangles in the second dimension are equal to 1; this can be interpreted as a *machine assignment problem*.
- Case (D) corresponds to a non-overlapping constraint between rectangles where each rectangle can have two orientations. This is achieved by associating with each rectangle two shapes of respective sizes $l \times h$ and $h \times l$. Since their orientation is not initially fixed, the *included* constraint enforces the three rectangles to be included within the bounding box defined by the origin's coordinates 1, 1 and sizes 8, 3.
- Case (E) corresponds to a non-overlapping constraint between more complex objects where each object is described by a given set of rectangles.
- Case (F) describes a placement problem where one has to first assign each rectangle to a strip so that all rectangles that are assigned to the same strip do not overlap.
- Case (G) corresponds to a non-overlapping constraint between parallelepipeds.
- Case (H) can be interpreted as a non-overlapping constraint between parallelepipeds that are assigned to the same container. The first dimension corresponds to the identifier of the container, while the next three dimensions are associated with the position of a parallelepiped inside a container.
- Case (I) describes a rectangle placement problem over three consecutive time-slots: rectangles assigned to the same time-slot should not overlap in time. We initially start with the three rectangles 1, 2 and 3. Rectangle 3 is no longer present at instant 2 (the triangle ▼ within rectangle 3 at time 1 indicates that rectangle 3 will disappear at the next time-point), while rectangle 4 appears at instant 2 (the triangle ▲ within rectangle 4 at time 2 denotes the fact that the rectangle 4 appears at instant 2). Finally, rectangle 2 disappears at instant 3 and is replaced by rectangle 5.

3 Standard Representation of Geometrical Constraints

The key idea for handling multiple geometrical constraints in a common kernel is the following. For each type of geometrical constraint found in \mathcal{C} (also called *exter-*




Fig. 2. Nine typical examples of use of *geost*

nal constraints), one has to provide a service that computes necessary conditions (also called *internal constraints*) for a given object and shape. Given an external geometrical constraint $ectr_i(\mathcal{A}_i, \mathcal{O}_i)$ ($\mathcal{A}_i \subseteq \{0, 1, \dots, k-1\}$, $\mathcal{O}_i \subseteq \mathcal{O}$), one of its object $o \in \mathcal{O}_i$ and one potential shape s of o , such a necessary condition generated by $ectr_i$, o and s is a unary⁴ constraint $ictr(o.x)$ such that: $o.sid = s \wedge ectr_i(\mathcal{A}_i, \mathcal{O}_i) \Rightarrow ictr(o.x)$. Now, the key to being able to globally treat such necessary conditions in the kernel is to give them a uniform representation. We have chosen the following one:

- A constraint $outbox(t, l)$ on $o.x$ holds iff $o.x$ is located outside the shifted box defined by its origins point $t[d]$, $0 \leq d < k$, and sizes $l[d]$, $0 \leq d < k$ (i.e., $\exists d \in [0, k-1] \mid o.x[d] < t[d] \vee o.x[d] > t[d] + l[d] - 1$).

Thus, an outbox corresponds to a box-shaped set of points that are infeasible for $o.x$. The purpose of the introduction of outboxes is to have a common representation for the kernel, suitable for the k -dimensional lexicographic sweep algorithm presented in the next section, which considers all the outboxes, for a selected object and shape, in one run.

Consequently, for each type of external geometrical constraint, found in \mathcal{C} a service $GenOutboxes(ectr_i, o, s) : (ictrs)$, responsible for generating outboxes, must be provided. This service is assumed to generate outboxes that intersect the domains of the origin coordinates of o . Also, if all attributes mentioned by $ectr_i$ belonging to objects other than o are fixed, those outboxes are assumed to be necessary and sufficient conditions, lest the kernel accept false solutions.

Example of External Geometrical Constraints. We now illustrate some external geometrical constraints that are currently available within the constraint kernel. As we saw in the introduction, an external constraint always has at least two arguments that resp.

⁴ Unary, since it involves the k coordinates of a *single* object.

correspond to a list of distinct dimensions and to a list of object identifiers to which the constraint apply.

The *included* and *non-overlapping* external constraints. The *included*($\mathcal{A}_i, \mathcal{O}_i, t, l$) and the *non-overlapping*($\mathcal{A}_i, \mathcal{O}_i$) external constraints take as input a list of distinct dimensions \mathcal{A}_i in $\{0, 1, \dots, k-1\}$ and a list \mathcal{O}_i of distinct object identifiers of *geost*. In addition, the *included* constraint considers a shifted box defined by its origin point $t[d]$, $0 \leq d < k$, and size $l[d]$, $0 \leq d < k$.

The *included* constraint enforces for each object o (with $o.id \in \mathcal{O}_i$) and for any corresponding shifted box s (with $o.sid = s.sid$) the condition $\forall d \in \mathcal{A}_i \mid t[d] \leq o.x[d] + s.t[d] \wedge o.x[d] + s.t[d] + s.l[d] - 1 \leq t[d] + l[d] - 1$ (i.e., s is included within the shifted box attribute defined by the parameters t and l of the *included* constraint). Depending on which shape of an object we actually consider, the *included* constraint can be translated to $2k$ outbox constraints.

The *non-overlapping* constraint enforces the following condition: given two distinct objects o and o' (with $o.id, o'.id \in \mathcal{O}_i$) that overlap in time, no shifted box s (with $o.sid = s.sid$) should overlap any shifted box s' (with $o'.sid = s'.sid$); i.e. it should hold that $\exists d \in \mathcal{A}_i \mid o.x[d] + s.t[d] + s.l[d] \leq o'.x[d] + s'.t[d] \vee o'.x[d] + s'.t[d] + s'.l[d] \leq o.x[d] + s.t[d]$ (i.e., there exists a dimension where they do not intersect). While focussing on an object o we can easily generate an outbox constraint for each object o' that should not overlap o by reusing the results of [2].

4 The Geometrical Kernel: a Generic k -dimensional Lexicographic Sweep Algorithm

In this section, we first present the sweep algorithm used for filtering the coordinates of the origin of an object o of *geost* when each object has one single shape. We initially assume that time is treated exactly like the space dimensions, i.e. that the $o.x$ array is extended by one element. Toward the end of this section, we explain in detail how to treat the time attributes of an object. We also assume for now that the shape attribute is fixed, and explain later how to handle multiple potential shapes for an object (i.e., polymorphism). We now introduce some notation used throughout this section.

Notation. Assume v and w are vectors of scalars of k components. Then $v \leftarrow w$ denotes the element-wise assignment of w to v , $w + d$ (resp. $w - d$) denotes the element-wise addition of d (resp. $-d$) to w . Given a scalar d , $0 \leq d \leq k-1$, $\text{rot}(v, d, k)$ denotes the vector $(v[d], v[(d+1) \bmod k], \dots, v[(d-1) \bmod k])$. That is, in the rotated vector, $v[d]$ is the most significant element, which is what we need when running the sweep algorithm on dimension d .

The Sweep Algorithm. This algorithm first considers all outboxes \mathcal{IC}_o derived from \mathcal{C} where object o actually appears, and then performs a recursive traversal of the placement space for each coordinate and direction (i.e., min or max). Without loss of generality, assume we want to adjust the minimum value of the d^{th} coordinate $o.x[d]$,

$0 \leq d < k$, of the origin of o . The algorithm starts its recursive traversal of the placement space at point $c = \text{rot}(\underline{o.x}, d, k)$ and could in principle explore all points of the domains of $o.x$, one by one, in increasing lexicographic order, until a point is found that is not inside any outbox, in which case $c[0]$ is the computed new minimum value. To make the search efficient, instead of moving each time to the successor point, we arrange the search so that it skips points that are known to be inside some outbox.⁵

Thus, we compute the lexicographically smallest point c' such that:

1. c' is lexicographically greater than or equal to c ,
2. every element of c' is in the domain of the corresponding element of $o.x$,
3. c' is not inside any outbox of \mathcal{IC}_o .

If no such c' exists, the constraint fails. Otherwise, the minimum value of $o.x[d]$ is adjusted to $c'[0]$. As we saw, the sweep algorithm moves in increasing lexicographic order a point c from its lexicographically smallest potential feasible position to its lexicographically largest potential feasible position through all potential points. The algorithm uses the following data structures:

- The current position c of the sweep.
- A vector $n[0..k-1]$ that records knowledge about already encountered sets of infeasible points while moving c from its first potential feasible position. The vector n is always element-wise greater than c and maintained as follows. Let inf, sup denote the vectors $\text{inf} = \text{rot}(\underline{o.x}, d, k)$ and $\text{sup} = \text{rot}(\overline{o.x} + 1, d, k)$:
 - Initially, $n = \text{sup}$.
 - Whenever an outbox f containing c is found, n is updated by taking the element-wise minimal value of n and the upper boundary of $\text{rot}(f, d, k)$, indicating the fact that new candidate points can be found beyond that value.
 - Whenever we skip to the next candidate point, we reset the elements of n that were used to the corresponding values of sup .

The following invariant holds for the vector n , and is used when advancing c to the next candidate point. Let i be the smallest j such that $n[j+1] = \text{sup}[j+1] \wedge \dots \wedge n[k-1] = \text{sup}[k-1]$ and suppose c is known to be in some outbox. Then, the next point, lexicographically greater than c and not yet known to be in any outbox, is $(c[0], \dots, c[i-1], n[i], \text{inf}[i+1], \dots, \text{inf}[k-1])$.

Algorithm 1 implement this idea. The algorithm prunes the bounds of each coordinate of every object wrt. its relevant outboxes, iterating to fix-point.

Efficiency. The main inefficiency in this sweep algorithm lies in searching the set of outboxes (line 4 of PruneMin). In order to make this search more efficient, we can make the sweep algorithm more sophisticated by the following modifications to PruneMin:

- We extend the state of the algorithm by an *event point series*, ordered in lexicographically increasing order. These events correspond to the lexicographically smallest (insert events) and largest (delete events) relevant infeasible point associated with each outbox $ictr_o \in \mathcal{IC}_o$. They are sorted in lexicographically increasing order, and we maintain a pointer into the series in sync with point c .

⁵ Potential holes in the domains are reflected in outboxes.

VARIABLES

x1 in 1..4, y1 in 2..4
 x2 in 4..4, y2 in 6..6
 x3 in 2..4, y3 in 8..9
 x4 in 7..7, y4 in 1..1
 x5 in 1..8, y5 in 1..8, y5 <> 7

EXTERNAL CONSTRAINT (non-overlapping)

```
geost([object(1,1,[x1,y1],0,1,1),object(2,2,[x2,y2],0,1,1),
      object(3,3,[x3,y3],0,1,1),object(4,4,[x4,y4],0,1,1),
      object(5,5,[x5,y5],0,1,1)],
      [shape(1,[0,2],[0,1]),shape(2,[0,3],[0,1]),shape(3,[0,1],[0,1]),
      shape(4,[0,1],[0,3]),shape(5,[0,5],[0,4])],
      [non-overlapping([0,1],[1,2,3,4,5])])
```

INTERNAL CONSTRAINTS GENERATED FOR FILTERING THE ORIGIN OF THE FIFTH OBJECT, i.e. (x5,y5) (ICTRS)

```
ctr1: outbox([1,1],[2,2]) ctr3: outbox([1,8],[2,1])
ctr2: outbox([1,3],[6,4]) ctr4: outbox([3,1],[5,3])
ctr5: outbox([1,7],[8,1])
```




Fig. 3. Illustration of the lexicographic sweep algorithm for adjusting the minimum value of the abscissa of rectangle r_5

```

PROCEDURE FilterCtrs( $k, \mathcal{O}, \mathcal{S}, \mathcal{C}$ ) : bool
1:  $nonfix \leftarrow \mathbf{true}$  // fixpoint not yet reached
2: while  $nonfix$  do
3:    $nonfix \leftarrow \mathbf{false}$  // assumes no filtering will be done
4:   for all  $o \in \mathcal{O}$  do
5:      $I \leftarrow \bigcup_{e \in \mathcal{C}} \text{GenOutboxes}(e, o, o.sid)$  // build the set of outboxes on  $o$ 
6:      $I \leftarrow I \cup \bigcup_{0 \leq d < k} \text{possible outboxes corresponding to holes in } o.x[d]$ 
7:     for  $d \leftarrow 0$  to  $k - 1$  do
8:       if  $\neg \text{PruneMin}(o, d, k, I) \vee \neg \text{PruneMax}(o, d, k, I)$  then
9:         return false // no feasible origin
10:      else if  $o.x$  was pruned then
11:         $nonfix \leftarrow \mathbf{true}$  // fixpoint not yet reached
12:      end if
13:    end for
14:  end for
15: end while
16: return true // feasible origin

PROCEDURE PruneMin( $o, d, k, I$ ) : bool
1:  $b \leftarrow \mathbf{true}$  //  $b = \mathbf{true}$  while we have not failed
2:  $c \leftarrow \underline{o.x}$  // initial position of the point
3:  $n \leftarrow \overline{o.x} + 1$  // upper limits+1 in the different dimensions
4: while  $b \wedge \exists f \in I \mid c \in f$  do
5:    $n \leftarrow \min(n, f.t + f.l)$  // update vector  $n$  according to an outbox  $f$  containing  $c$ 
6:    $b \leftarrow \mathbf{false}$  // no new point to jump to yet
7:   for  $j \leftarrow k - 1$  downto  $0$  do
8:      $j' \leftarrow (j + d) \bmod k$  // rotation wrt.  $d, k$ 
9:      $c[j'] \leftarrow n[j']$  // use vector  $n$  to jump
10:     $n[j'] \leftarrow \overline{o.x}[j'] + 1$  // reset component of  $n$  to maximum value
11:    if  $c[j'] \leq \overline{o.x}[j']$  then
12:       $b \leftarrow \mathbf{true}$  // jump target found
13:       $j \leftarrow 0$  // exit for loop
14:    else
15:       $c[j'] \leftarrow \underline{o.x}[j']$  // reset component of  $c$ , for exhausted a dimension
16:    end if
17:  end for
18: end while
19: if  $b$  then
20:    $\underline{o.x}[d] \leftarrow \max(\underline{o.x}[d], c[d])$ 
21: end if
22: return  $b$ 

```

Algorithm 1: FilterCtrs is the main filtering algorithm associated with $geost(k, \mathcal{O}, \mathcal{S}, \mathcal{C})$, where k , \mathcal{O} , \mathcal{S} and \mathcal{C} resp. correspond to the number of dimensions, to the objects, to the shapes and to the external geometrical constraints. PruneMin adjusts the lower bound of the d^{th} coordinate of the origin of object o where I is the set of outboxes associated with object o (since PruneMax is similar to PruneMin it is omitted). The given fixpoint loop is an over-simplification. The implementation maintains a set of objects that need filtering. Whenever an object o is pruned, all non-fixed objects connected to o by an external constraint are added to this set. When the set becomes empty, the fixpoint is reached.

- We maintain the set of *active outboxes*, corresponding to all outboxes $ictr_o \in \mathcal{IC}_o$ such that c is between its lexicographically smallest and largest infeasible points. This set is initially empty.
- When c is initialized in line 2 as well as when c is incremented in lines 7-17, the relevant events up to point c from the event point series are processed, and the corresponding outboxes are added to or deleted from the set of active outboxes.
- In line 4, only the active outboxes are considered.

Example 3. Fig. 3 illustrates the k -dimensional lexicographic sweep algorithm in the context of $k = 2$. Parts (A) and (B) provide the variables of the problem (i.e., the abscissa and ordinate of each rectangle r_1, r_2, r_3, r_4 and r_5) as well as the non-overlapping constraint between the five previous rectangles. On Part (D) we have represented the extreme possible feasible positions of each rectangle (i.e., rectangles r_1 to r_4): for instance the leftmost lower corner of rectangle r_1 can only be fixed at positions $(1, 2), (1, 3), (1, 4), (2, 2), (2, 3), (2, 4), (3, 2), (3, 3), (3, 4), (4, 2), (4, 3)$ and $(4, 4)$. Parts (C) to (L) of Fig. 3 detail the different steps of the algorithm for adjusting the minimum value of the abscissa of rectangle r_5 . Part (C) provides the outboxes associated with the fact that we want to prune the coordinates of r_5 : constraints ctr_1, ctr_2, ctr_3 and ctr_4 resp. correspond to the fact that rectangle r_5 should not overlap rectangles r_1, r_2, r_3 and r_4 , while constraint ctr_5 represents the fact that the ordinate of r_5 should be different from 7. Part (D) represents the initialisation phase of the algorithm where we have all five outboxes with their respective lexicographically smallest infeasible point (i.e., $(1, 1)$ for $ctr_1, (1, 3)$ for $ctr_2, (1, 7)$ for $ctr_5, (1, 8)$ for ctr_3 and $(3, 1)$ for ctr_4). Part (E) represents the first step of the sweep algorithm where we start the traversal of the placement space at point $c = (1, 1)$. We first transfer to the list of active outboxes all outboxes for which the first lexicographically smallest infeasible point is lexicographically greater than or equal to the current position of the sweep $c = (1, 1)$ (i.e., constraint $ctr_1 = \text{outbox}([1, 1], [2, 2])$). We then search through the list of active constraints (represented on the figure by a box with the legend ACTRS on top of it) the first constraint for which $c = (1, 1)$ is infeasible. In fact, since ctr_1 is infeasible (represented on the figure by a box with the legend CONFLICT on top of it) we compute the vector $f = (3, 3)$ that tells how to get the next potentially feasible point in the different dimensions. Consequently the sweep moves to the next position $(1, 3)$ (see Part (F)) and the process is repeated until we finally find a feasible point for all outboxes (i.e., point $(3, 8)$ in Part (L)). Note that, when the lexicographically largest infeasible point associated with an active outbox is lexicographically less than the current position of the sweep, we remove that constraint from the list of active outboxes. This is for instance the case in Part (I), where we remove constraint ctr_3 from the list of active outboxes (since its lexicographically largest infeasible point $(2, 8)$ is lexicographically less than the position of the sweep $c = (3, 1)$).

Complexity. Rather than analysing the complexity of the *geost* kernel for a fixed k , which depends both on the type of each external constraint (i.e., the complexity of a given external constraint for generating all its corresponding outboxes as well as their number), we rather focus on PruneMin for adjusting the minimum value of the d^{th} coordinate of the origin of an object. Assuming that the maximum number of outboxes is equal to n we give an upper bound on the maximum number of jumps of PruneMin (i.e., the maximum number of times the sweep is moved).

First note that we always jump to an upper border (+1) of an outbox (i.e., see line 5 of PruneMin) or that we reset some coordinates of the sweep to its minimum value (i.e., see line 15 of PruneMin). Consequently, all coordinates of the sweep are always equal to an upper border (+1) of some outboxes or to a minimum possible value. Since we

want to evaluate the maximum number of jumps, let us assume that for every dimension d ($0 \leq d < k$) the upper limits of all the n outboxes are distinct. Having this in mind we can construct a maximum of $(n + 1)^k$ points. Even if we found a systematic construction where this number of jumps is reached, the performance evaluation of Section 5 indicates that we can handle a reasonable number of objects for $k = 2, 3, 4$.

From a memory consumption point of view, the algorithm only records the coordinates of the sweep from one invocation to the next, in order not to restart the search from scratch (i.e., $2k$ points for each object).

Handling Time. Given an object $o \in \mathcal{O}$ of *geost*, the sweep algorithm that we have introduced in the previous section can be easily adapted to handle the start in time $o.start$, duration in time $o.duration$ and end in time $o.end$. Beside maintaining bound consistency for the constraint $o.end = o.start + o.duration$, we add an extra *time* dimension to the geometric coordinates of object o . Roughly, this new time coordinate corresponds to $o.start$ resp. $o.end$ depending on whether we are adjusting the minimum or maximum.

Handling Polymorphism. In order to handle the fact that objects can have several potential shapes we modify the previous algorithm in the following way. For adjusting the minimum value of the coordinate of the origin of an object that has more than one shape we call the sweep algorithm for each potential shape of the object (i.e., for each value of its shape variable). Then we take the smallest minimum value obtained (i.e., we use constructive disjunction) and prune the shape variable of an object if we did not find any feasible point for a given potential shape of that object.

Other Internal Constraints. The standard representation of geometrical constraints given in Section 3 is an over-simplification. For some constraints, e.g. distance constraints, outboxes are not a suitable representation, as the set of forbidden coordinates cannot be covered by a small number of boxes. Therefore, the constraint kernel internally handles other representations of necessary conditions, with an appropriate internal API. For details, see the technical report [5].

5 Performance Evaluation

We evaluate the implementation⁶ of the *geost* kernel from three perspectives:

Wanting to measure the speed and the scalability of the sweep algorithm for finding a first solution on loosely constrained placement problems (i.e., 20% spare space), we generated one set of random problem instances of m k -dimensional boxes for $k \in \{2, 3, 4\}$ involving $t \in \{1, 16, 256, 1024\}$ distinct types of boxes, and for $m \in \{1024, 2048, \dots, 262144\}$. The results for $k = 2$ are shown in Fig. 4 (top left) and indicates that the approach is sensible to the number of distinct types of boxes. It can typically pack 1024 2D, 3D and 4D distinct boxes in at most 200 msec. The longest

⁶ The experiments were run in SICStus Prolog 4 compiled with gcc -02 version 4.0.2 on a 3GHz Pentium IV with 1MB of cache.

time, 13694 seconds (close to 4 hours), was obtained for packing 262144 4D parallelepipeds (over 1 million domain variables) with a memory consumption of 351MB.

Wanting to get an idea of the performance of the *geost* kernel on very tight placement problems (i.e., 0% spare space), we considered the *perfect squared squares problem* [1, 6] as well as the *3D pentominoes problem* [7]:

- A *perfect squared square of order n* is a square that can be tiled with n smaller squares where each of the smaller squares has a different integer size. We used the data available (i.e., the size of the small squares to pack) from the catalogue [8] and tested the corresponding 207 instances. The labelling strategy is roughly to repeat the following, first for the x dimension, then for the y dimension:
 1. Find the smallest position where some square can be placed.
 2. Find a square to place in that position.
- *Pentominoes* are pieces made of 5 connected unit cubes laid on a plane surface. Their shapes look like the 12 letters $F, I, L, P, N, T, U, V, W, X, Y$ and Z . We considered the problem of finding the different ways of putting 12 distinct shapes that can be reflected and rotated in a box having a volume of 60 unit cubes. Our labelling strategy is roughly to repeat the following:
 1. Find a slot in the space that has not yet been filled by some piece.
 2. Find a piece that can fill that slot.

Fig. 4 (top right) and Table 1 respectively report, for the squared squares and the pentominoes problems, the time and number of backtracks for exploring all the search space⁷ without breaking any symmetry. For the squared squares problems the maximum time of 1585 seconds was spent on problem 48; on the other hand, 148 problems were completely solved within 60 seconds. For the 3D pentomino packing instances, performance results for comparison can be found in [7]. However, they stop the search when the first 100 solutions have been found, so the results are only partly comparable.

Finally, wanting to compare the *geost* kernel with a recent exact state of the art method for the 2D orthogonal packing problem [4], we reused the benchmarks proposed by Clautiaux et al. [9]. This is a feasibility problem which consists in determining whether a set of rectangles that cannot be rotated, can be packed or not into a rectangle of fixed size. In these instances the discrepancy between the sum of the areas of the rectangles to pack and the area of the big rectangle vary from 0% to 20%. We have 41 instances involving between 10 and 23 rectangles. Moreover, from these 41 instances, 26 instances are not feasible. In order to break symmetries between multiple rectangles of the same shape we added lexicographic ordering constraints. All x coordinates were labelled followed by all y coordinates, by decreasing rectangle size. Values were tried by increasing value. Fig. 4 (bottom) compares our results with the ones reported in [4]. Note that the sequence order for the curves differs, since the instances of each curve are ordered by increasing y value. We solved all instances and are comparable with [4], although 8 instances are much easier for [4] and 10 instances are much easier for us.

Note that for the last three problems (i.e., Squared Squares, Pentominoes and 2D orthogonal packing) extra filtering algorithms mostly based on cumulative relaxation were integrated within our kernel. Since this paper focusses on the constraint kernel and because of space limitations these methods were not detailed.

⁷ Finding all solutions and proving that there is no other solution.

6 Related Work and Future Directions

The rectangles packing problem has been studied by Clautiaux et al. [9, 4] where scheduling-based reasoning is used [1]. The use of sweep algorithms in constraint filtering algorithms was introduced in [3] and applied to the non-overlapping 2D rectangles constraints. This paper generalizes and extends that work in several ways.

- The 2D sweep is generalized to a lexicographic sweep, independent of the number of dimensions.
- The notion of forbidden regions for non-overlapping rectangles is generalized to necessary conditions for general geometric constraints.

The idea of generating necessary conditions is reminiscent of indexicals [10], a.k.a. projection constraints [11]. An indexical for a constraint $c(x_1, \dots, x_n)$ computes a unary constraint on a single variable x_i , i.e. a set S of values such that $c \Rightarrow x_i \in S$, in reaction to domain changes in x_1, \dots, x_n . The constraint kernel then immediately enforces $x_i \in S$. Our kernel generalizes this in two ways:

- We compute necessary conditions in the form of k -dimensional forbidden regions.
- We treat all such forbidden regions, for a selected object and shape, in one run of the sweep algorithm. Projecting a single forbidden region on one coordinate often does not yield any pruning, whereas considering the union of forbidden regions is much more effective.

Dal Palù et al. in [12] proposed a constraint solver specialized for 3D discrete domains. Their solver was targeted to the study of problems in molecular, chemical and crystal structures. Our work, however, remains in the setting of mainstream finite domain constraint systems, whereas our kernel internally handles k -dimensional objects.

Even though the *geost* kernel has been designed over discrete domains, it could rather easily be extended to continuous domains with the coordinates of the objects approximated by the floating-point numbers \mathbb{F} . Since switching from \mathbb{N} to \mathbb{F} may cause rounding errors at this level, the sweep algorithm needs to handle these rounding errors when moving the sweep out of an outbox constraint. If the projections of the forbidden regions on all dimensions are intervals of real bounds we can proceed as follows. On continuous domains, an outbox will have a very thin strip at the border where the feasibility of the corresponding internal constraint is unknown. The region inside this strip is strictly forbidden, and outside, the constraints is certainly satisfied. The outbox must be computed including this strip, by taking lower and upper approximations of the region's coordinates. In that case, the solutions are guaranteed to be valid, but the solver may not be complete, because it may (rarely) happen that the real forbidden region allows positions that are forbidden by its approximation.

This research was conducted under the European Union project “Net-WMS”, a major task of which is to study packing problems in warehouse management. In this context, our constraint kernel is a step towards being able to capture a large set of packing rules in a constraint programming setting. Future work involves extending our set of external geometric constraints to include such packing rules.

7 Conclusion

The main contribution of this paper is a geometrical constraint kernel for handling the location in space and time of polymorphic k -dimensional objects subject to various geometrical and time constraints. The constraint kernel is generic in the sense that one of its parameters is a set of constraints on subsets of the objects. These constraints are handled globally by the kernel.

We have presented a sweep algorithm for filtering the attributes of the objects. Thank to its architecture, new geometric constraints can be plugged into this sweep algorithm without modifying it. The strong point of this sweep algorithm is that it considers all the geometrical constraints for a selected object and shape in one run. As a first result, more deduction can be performed by combining sets of forbidden points coming from multiple geometrical constraints. Secondly, it can handle within one single constraint problems involving up several tens of thousands of objects without memory consumption problems, which is often a weak point for constraint programming environment. We have also shown that we could handle tight 2D or 3D placement problems, which were traditionally solved by specific approaches.

Acknowledgements

This research was conducted under European Union Sixth Framework Programme Contract FP6-034691 “Net-WMS”.

configuration	backtracks (1st)	time (1st)	backtracks (all)	time (all)	solutions
$20 \times 3 \times 1$	1434	1840	47381	49740	8
$15 \times 4 \times 1$	290	560	888060	939060	1472
$12 \times 5 \times 1$	1594	1850	3994455	4112870	4040
$10 \times 6 \times 1$	111	260	9688985	10726810	9356
$10 \times 3 \times 2$	1267	2370	1203511	1778980	96
$6 \times 5 \times 2$	157	730	n/a	n/a	n/a
$5 \times 4 \times 3$	3567	14930	n/a	n/a	n/a

Table 1. Performance evaluation. 3D pentomino packing instances. Time in milliseconds. “n/a” corresponds to a quantity that was not available with a time-out of several hours.

References

1. A. Aggoun and N. Beldiceanu. Extending CHIP in order to solve complex scheduling and placement problems. *Mathl. Comput. Modelling*, 17(7):57–73, 1993.
2. N. Beldiceanu, Q. Guo, and S. Thiel. Non-overlapping constraints between convex polytopes. In T. Walsh, editor, *Proc. CP’2001*, volume 2239 of *LNCS*, pages 392–407. Springer-Verlag, 2001.
3. N. Beldiceanu and M. Carlsson. Sweep as a generic pruning technique applied to the non-overlapping rectangles constraints. In T. Walsh, editor, *Proc. CP’2001*, volume 2239 of *LNCS*, pages 377–391. Springer-Verlag, 2001.




Fig. 4. Performance evaluation. Top left: scalability, $t \in \{1, 16, 256, 1024\}$. Top right: Perfect Squared Squares, runtime and backtracks. Bottom: 2D Orthogonal Packing, runtime (left) and backtracks (right). Time in milliseconds. In each curve, the instances are ordered by increasing y value.

4. F. Clautiaux, A. Joulet, J. Carlier, and A. Moukrim. A new constraint programming approach for the orthogonal packing problem. *Computers and Operation Research*, to appear.
5. N. Beldiceanu, M. Carlsson, E. Poder, R. Sadek, and C. Truchet. A generic geometrical constraint kernel in space and time for handling polymorphic k -dimensional objects. SICS Technical Report T2007:08, Swedish Institute of Computer Science, 2007.
6. P. Van Hentenryck. Scheduling and packing in the constraint language cc(FD). In M. Zweben and M. Fox, editors, *Intelligent Scheduling*. Morgan Kaufmann Publishers, 1994.
7. A. Colmerauer and B. Gilletta. Solving the three-dimensional pentamino puzzle. Technical report, Laboratoire d'Informatique de Marseille, 1999. <http://www.lim.univ-mrs.fr/~colmer/ArchivesPublications/Gilletta/misc99.pdf>.
8. C. J. Bouwkamp and A. J. W. Duijvestijn. Catalogue of simple perfect squared squares of orders 21 through 25. Technical Report EUT Report 92-WSK-03, Eindhoven University of Technology, The Netherlands, November 1992.
9. F. Clautiaux, J. Carlier, and A. Moukrim. A new exact method for the two-dimensional orthogonal packing problem. *European Journal of Operational Research*, to appear.
10. Pascal Van Hentenryck, Vijay Saraswat, and Yves Deville. Constraint processing in cc(FD). Manuscript, 1991.
11. Gregory Sidebottom. *A Language for Optimizing Constraint Propagation*. PhD thesis, Simon Fraser University, 1993.
12. Alessandro Dal Pal'u, Agostino Dovier, and Enrico Pontelli. A new constraint solver for 3D lattices and its application to the protein folding problem. In Geoff Sutcliffe and Andrei Voronkov, editors, *Logic for Programming, Artificial Intelligence, and Reasoning, 12th International Conference, LPAR 2005*, volume 3835 of *LNCS*, pages 48–63. Springer-Verlag, 2005.