To appear in: *Proceedings of the 6th International Conference on Principles of Knowledge Representation and Reasoning (KR'98)*, Trento, Italy, June 1998.

A Canonical Model of the Region Connection Calculus

Jochen Renz
Institut für Informatik
Albert-Ludwigs-Universität
D-79110 Freiburg, Germany

Abstract

Although the computational properties of the Region Connection Calculus RCC-8 are well studied, reasoning with RCC-8 entails several representational problems. This includes the problem of representing arbitrary spatial regions in a computational framework, leading to the problem of generating a realization of a consistent set of RCC-8 formulas. A further problem is that RCC-8 performs reasoning about topological space, which does not have a particular dimension. Most applications of spatial reasoning, however, deal with two- or three-dimensional space. Therefore, a consistent set of RCC-8 formulas might not be realizable within the desired dimension. In this paper we address these problems and develop a canonical model of RCC-8 which allows a simple representation of regions with respect to a set of RCC-8 formulas, and, further, enables us to generate realizations in any dimension $d \geq 1$, even when regions are constrained to be (sets of) polytopes. For three- and higher-dimensional space this can also be done for internally connected regions.

1 INTRODUCTION

The Region Connection Calculus (RCC) (Randell *et al.* 1992) is a topological approach to qualitative spatial representation and reasoning (Cohn 1997) where spatial regions are regular subsets of a topological space. Of particular interest for application purposes is RCC-8, a sub-calculus of RCC that uses eight mutually exhaustive and pairwise disjoint base relations. The computational properties of RCC-8 have been studied thoroughly (Nebel 1995; Renz and Nebel 1997) and efficient reasoning mechanisms were identified.

Despite this, there are still several problems with representing spatial regions within RCC-8. As the calculus is based on topology, spatial regions might be arbitrary subsets of a topological space which are not necessarily analytically describable; therefore, it appears to be difficult to represent spatial regions in a computational framework.

Another representational drawback of using RCC-8 is that a topological space does not have a particular dimension, whereas most applications of qualitative spatial reasoning deal only with two- or three-dimensional space. It might be possible that a set of RCC-8 formulas is consistent but not realizable within a particular dimension. Lemon (1996) gave an example of a set of spatial formulas which is realizable in three dimensional space but not in two dimensional space if regions are internally connected. Lemon used this result to argue that spatial logics like RCC are not an adequate formalism for representing space.

A further problem, which also depends on the ability to represent spatial regions, is finding a realization of a consistent and realizable set of spatial formulas in a particular dimension, instead of just knowing whether the set is realizable or not.

In this paper, we will refer to these representational topics. In order to represent arbitrary spatial regions, it is necessary to have a *canonical model* of RCC-8, i.e., a structure that allows to model any consistent sentence of the calculus. Topological space is of course a canonical model, but, as described above, this does not seem to be very useful for representing regions. Therefore, we will present a new canonical model of RCC-8 that permits a simple representation of spatial regions by reducing them to their necessary topological features with respect to their spatial relations. Based on this model, we will prove that for any consistent set of spatial formulas there are realizations in any dimension $d \geq 1$ when regions are allowed to be internally disconnected. This is still true even when regions are constrained to be sets of polytopes. Actually, internal connectedness of regions is not at all forced in the RCC-theory, so RCC can still be seen as an adequate representation formalism of space. We will also argue that forcing internal connectedness of all regions is too restrictive when dealing with spatial regions. Nevertheless, we will prove that in three- and higher dimensional space every consistent set of spatial formulas can always be realized with internally connected regions. Using the new canonical model for representing spatial regions, it becomes possible to determine realizations of consistent sets of spatial formulas. We will give algorithms for generating realizations of both internally connected and disconnected regions.

The remainder of the paper is structured as follows: In Section 2 we introduce RCC-8 and some basic topological notions. Section 3 sketches the modal encoding of RCC-8 and presents the new canonical model of RCC-8. In Section 4 we give a topological interpretation of this model which is used in Section 5 to prove the results about realizations in particular dimensions. Section 6 describes how models of sets of spatial relations can be determined and how realizations can be generated. In Section 7 we will discuss our results.

2 QUALITATIVE SPATIAL REPRESENTATION WITH RCC

RCC is a topological approach to qualitative spatial representation and reasoning where *spatial regions* are regular subsets of a topological space (Randell *et al.* 1992). Relationships between spatial regions are defined in terms of the relation $C(r, s)$ which is true if and only if the closure of region r is connected to the closure of region s, i.e., if they share a common point. Regions themselves do not have to be internally connected, i.e., a region may consist of different disconnected pieces. The domain of *spatial variables* (denoted as X, Y, Z) is the whole topological space.

RCC-8 (Randell *et al.* 1992) uses a set of eight pairwise disjoint and mutually exhaustive relations, called *base relations*, denoted as DC, EC, PO, EQ, TPP, NTPP, TPP^{-1}, and NTPP^{-1}, with the meaning of *DisConnected, Externally Connected, Partial Overlap, EQual, Tangential Proper Part, Non-Tangential Proper Part*, and their converses. Examples for these relations are shown in Figure 1.

Sometimes it is not known which of the eight base relations holds between two regions, but it is possible to exclude some of them. In order to represent this, unions of base relations can be used.

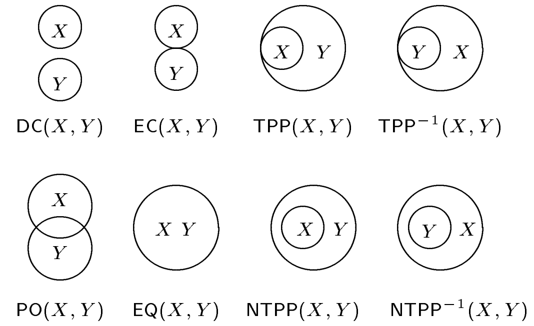

Figure 1: Two-dimensional examples for the eight base relations of RCC-8.

Since base relations are pairwise disjoint, this results in 2^8 different relations. A *spatial formula* $S(X, Y)$ is a relation between two spatial variables, a *spatial configuration* is a set Θ of spatial formulas. Θ is *consistent* if it is possible to find a *realization* of Θ, i.e., a model where every spatial variable is instantiated by a spatial region such that all relations hold between the regions. A consistent instantiation of the spatial variables X, Y, Z will be denoted as $\mathtt{X}, \mathtt{Y}, \mathtt{Z}$, respectively. Computational properties of reasoning with RCC-8 were studied in (Nebel 1995; Renz and Nebel 1997).

As we will go further into topology, we will define some common topological terms:

Definition 2.1 *Let* $\mathcal{U}$ *be a set, the* universe. *A* topology *on* $\mathcal{U}$ *is a family* T *of subsets of* $\mathcal{U}$, *with*

1. *if* $O_1, O_2 \in T$, *then* $O_1 \cap O_2 \in T$,
2. *if* $O_i \in T$ *for* $i \in I$, *then* $\bigcup O_i \in T$,
3. $\emptyset, \mathcal{U} \in T$.

A topological space *is a pair* $\langle \mathcal{U}, T \rangle$. *Every subset* $O \subset \mathcal{U}$ *with* $O \in T$ *is* open.

If the particular topology T on a set $\mathcal{U}$ is not important, we say that $\mathcal{U}$ is a topological space.

Definition 2.2 *Let* $\mathcal{U}$ *be a topological space,* $M \subset \mathcal{U}$ *be a subset of* $\mathcal{U}$ *and* $p \in \mathcal{U}$ *be a point in* $\mathcal{U}$.

- M *is* closed *if* $\mathcal{U} \setminus M$ *is open.*
- $N \subset \mathcal{U}$ *is said to be a* neighborhood *of* p *if there is an open subset* $O \subset \mathcal{U}$ *such that* $p \in O \subset N$.
- p *is said to be an* interior point *of* M *if there is a neighborhood* N *of* p *contained in* M. *The set of all interior points of* M *is called the* interior *of* M, *denoted* $i(M)$.

- *p is said to be an* exterior point *of M if there is a neighborhood N of p that contains no point of M. The set of all exterior points of M is called the* exterior *of M, denoted $e(M)$.*

- *p is said to be a* boundary point *of M if every neighborhood N of p contains at least one point in M and one point not in M. The set of all boundary points of M is called the* boundary *of M, denoted $b(M)$.*

- *The* closure *of M is the smallest closed set which contains M, i.e., $M \cup b(M)$.*

- *For any arbitrarily given point $p \in \mathcal{U}$, the family of all neighborhoods of p in $\mathcal{U}$ is called the* neighborhood system *of p in $\mathcal{U}$.*

A neighborhood system $\mathcal{N}$ has the property that every finite intersection of members of $\mathcal{N}$ belongs to $\mathcal{N}$.

It is possible to use any topological space which is a model for the RCC axioms as specified in (Randell *et al.* 1992). Gotts (1996) has shown that every regular connected topological space is a model for the RCC axioms (see also Section 7). So, whenever we refer to a topological space in the remainder of the paper, we mean a regular connected topological space.

3 MODAL ENCODING & CANONICAL MODELS

After making a brief introduction to modal logic, we will introduce the modal encoding of RCC-8 and a canonical model for this encoding.

3.1 MODAL LOGIC & KRIPKE SEMANTICS

Propositional modal logic (Fitting 1993; Chellas 1980) extends classical propositional logic by additional unary *modal operators* $\Box_i$. A common semantic interpretation of modal formulas is the *Kripke semantics* which is based on a *Kripke frame* $\mathcal{F} = \langle W, \mathcal{R} \rangle$ consisting of a set of *worlds* W and a set $\mathcal{R}$ of *accessibility relations* between the worlds, where $R \subseteq W \times W$ for every accessibility relation $R \in \mathcal{R}$. A different accessibility relation $R_{\Box_i} \in \mathcal{R}$ is assigned to every modal operator $\Box_i$. If $u, v \in W$, $R \in \mathcal{R}$, and uRv holds, we say that v is *R-accessible* from u or v is an *R-successor* of u.

A *Kripke model* $\mathcal{M} = \langle W, \mathcal{R}, \pi \rangle$ uses an additional valuation π that assigns each world and each propositional atom a truth value $\{true, false\}$. Using a Kripke model, a modal formula can be interpreted with

$S(X,Y)$	*Model Constraints*	*Entailment Constraints*
DC	$\neg(\mathsf{X} \wedge \mathsf{Y})$	$\neg\mathsf{X}, \neg\mathsf{Y}$
EC	$\neg(\mathbf{I}\mathsf{X} \wedge \mathbf{I}\mathsf{Y})$	$\neg(\mathsf{X} \wedge \mathsf{Y}), \neg\mathsf{X}, \neg\mathsf{Y}$
PO	—	$\neg(\mathbf{I}\mathsf{X} \wedge \mathbf{I}\mathsf{Y}), \mathsf{X} \rightarrow \mathsf{Y}, \mathsf{Y} \rightarrow \mathsf{X}, \neg\mathsf{X}, \neg\mathsf{Y}$
TPP	$\mathsf{X} \rightarrow \mathsf{Y}$	$\mathsf{X} \rightarrow \mathbf{I}\mathsf{Y}, \mathsf{Y} \rightarrow \mathsf{X}, \neg\mathsf{X}, \neg\mathsf{Y}$
TPP^{-1}	$\mathsf{Y} \rightarrow \mathsf{X}$	$\mathsf{Y} \rightarrow \mathbf{I}\mathsf{X}, \mathsf{X} \rightarrow \mathsf{Y}, \neg\mathsf{X}, \neg\mathsf{Y}$
NTPP	$\mathsf{X} \rightarrow \mathbf{I}\mathsf{Y}$	$\mathsf{Y} \rightarrow \mathsf{X}, \neg\mathsf{X}, \neg\mathsf{Y}$
NTPP^{-1}	$\mathsf{Y} \rightarrow \mathbf{I}\mathsf{X}$	$\mathsf{X} \rightarrow \mathsf{Y}, \neg\mathsf{X}, \neg\mathsf{Y}$
EQ	$\mathsf{X} \rightarrow \mathsf{Y}, \mathsf{Y} \rightarrow \mathsf{X}$	$\neg\mathsf{X}, \neg\mathsf{Y}$

Table 1: Encoding of the base relations in modal logic

respect to the set of worlds, the accessibility relations, and the valuation. For example, a propositional atom a holds in a world w of the Kripke model $\mathcal{M}$ (written as $\mathcal{M}, w \Vdash a$) if and only if $\pi(w, a) = true$. An arbitrary modal formula is interpreted according to its inductive structure. A modal formula $\Box_i\varphi$, e.g., holds in a world w of the Kripke model $\mathcal{M}$, i.e., $\mathcal{M}, w \Vdash \Box_i\varphi$, if and only if φ holds in all $R_{\Box_i}$-successors of w. $\mathcal{M}, w \Vdash \neg\Box_i\varphi$ if and only if there is an $R_{\Box_i}$-successor of w where φ does not hold.

Different modal operators can be distinguished according to their different accessibility relations. In this paper we are using a so-called S4-operator and an S5-operator. The accessibility relation of an S4-operator is reflexive and transitive, the accessibility relation of an S5-operator is reflexive, transitive, and euclidean.

3.2 MODAL ENCODING OF RCC8

The encoding of RCC-8 in propositional modal logic was introduced by Bennett (1995) and extended in (Renz and Nebel 1997). In both cases the encoding is restricted to regular closed regions. The encoding is based on a set of *model* and *entailment constraints* for each base relation, where model constraints must be true and entailment constraints must not be true. Bennett encoded these constraints in modal logic by transforming every spatial variable to a propositional atom and introducing an S4-operator $\mathbf{I}$ which he interpreted as an interior operator (Bennett 1995). In order to distinguish between spatial variables and the corresponding propositional atoms we will write propositional atoms as X, Y. Table 1 displays the constraints for the eight base relations. In order to combine them to a single modal formula, Bennett introduced an S5-operator[1] $\Box$, where $\Box\varphi$ is written for every model constraint φ and $\neg\Box\psi$ for every entailment constraint ψ (Bennett 1995). All constraints of a single base relation are then combined conjunctively to a single modal

[1]Bennett called this a *strong* S5-operator as all worlds are $R_\Box$-accessible from each other, i.e., $R_\Box = W \times W$.

formula. In order to represent unions of base relations, the modal formulas of the corresponding base relations are combined disjunctively. In this way every spatial formula $S(X,Y)$ can be transformed to a modal formula $m_1(S(X,Y))$. Additional constraints $m_2(X)$ are necessary to guarantee that only regular closed regions are used (Renz and Nebel 1997): every region must be equivalent to the closure of its interior, and the complement of a region must be equivalent to its interior.

$$m_2(X) = \Box(\mathsf{X} \leftrightarrow \neg\mathbf{I}\neg\mathbf{I}\mathsf{X}) \wedge \Box(\neg\mathsf{X} \leftrightarrow \mathbf{I}\neg\mathsf{X}).$$

So, any set of spatial formulas Θ can be written as a single modal formula $m(\Theta)$

$$m(\Theta) = \bigwedge_{S(X,Y)\in\Theta} m_1(S(X,Y)) \quad \wedge \bigwedge_{X\in Reg(\Theta)} m_2(X),$$

where $Reg(\Theta)$ is the set of spatial variables of Θ.

3.3 A CANONICAL MODEL OF RCC8

The modal encoding of RCC-8 can be interpreted by Kripke models. As the modal encoding of RCC-8 is equivalent to a set of RCC-8 formulas (Bennett 1995), a canonical model of RCC-8 is a structure that allows a Kripke model for the modal encoding of any consistent set of spatial formulas Θ. In order to obtain a canonical model, we distinguish different levels of worlds of W (Renz and Nebel 1997). A world w is of *level 0* if there is no world $v \neq w$ with $vR_{\mathbf{I}}w$. A world w is of *level l* if there is a world v of level $l-1$ with $vR_{\mathbf{I}}w$ and there is no world $u \neq w$ of a level higher than $l-1$ with $uR_{\mathbf{I}}w$. Based on this hierarchy of worlds, we will define the canonical model of RCC-8.

Definition 3.1 *An* RCC-8-structure $\mathcal{S}_{RCC8} = \langle W_{\mathcal{S}}, \{R_\Box, R_{\mathbf{I}}\}, \pi_{\mathcal{S}}\rangle$ *has the following properties:*

1. $W_{\mathcal{S}}$ *contains only worlds of level* 0 *and* 1.
2. *For every world u of level* 0 *there are exactly* $2n$ *worlds v of level* 1 *with* $uR_{\mathbf{I}}v$. *These* $2n+1$ *worlds form an* RCC-8-cluster *(see Figure 2).*
3. *For every world v of level* 1 *there is exactly one world u of level* 0 *with* $uR_{\mathbf{I}}v$.
4. *For all worlds* $w, v \in W_{\mathcal{S}}$: $wR_{\mathbf{I}}w$ *and* $wR_\Box v$.

$\mathcal{S}_{RCC8}$ *contains* RCC-8-*clusters with all possible valuations*[2] *with respect to* $R_{\mathbf{I}}$. *A set of* RCC-8-*clusters* $\mathcal{M} = \langle W, \{R_\Box, R_{\mathbf{I}}\}, \pi\rangle \subset \mathcal{S}_{RCC8}$ *is an* RCC-8-model *of* $m(\Theta)$ *if* $\mathcal{M}, w \Vdash m(\Theta)$ *for a world* $w \in W$. *In a* polynomial RCC-8-model *the number of worlds is polynomially bounded by the number of regions n.*

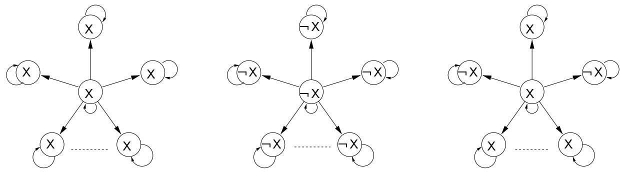

Figure 2: Three possible RCC-8-clusters of $\mathcal{S}_{RCC8}$. Worlds are drawn as circles, the arrows indicate $R_{\mathbf{I}}$.

In (Renz and Nebel 1997) it was proven that if $m(\Theta)$ is satisfiable, there is a polynomial RCC-8-model $\mathcal{M}$ with $\mathcal{M}, w \Vdash m(\Theta)$ that uses $O(n^2)$ different worlds of level 0 – one world of level 0 for every entailment constraint. So the RCC-8-structure $\mathcal{S}_{RCC8}$ is a canonical model[3] of the modal encoding of RCC-8. In order to obtain a "topological" canonical model for the topological calculus RCC-8, we give in the next section a topological interpretation of RCC-8-models.

4 TOPOLOGICAL INTERPRETATION OF THE CANONICAL MODEL

The modal encoding of RCC-8 was obtained by introducing a modal operator $\mathbf{I}$ corresponding to the topological interior operator and transferring the topological properties and axioms to modal logic. Using the intended interpretation of $\mathbf{I}$ as an interior operator, it is unclear how the Kripke models we consider, especially the accessibility relations $R_\Box$ and $R_{\mathbf{I}}$, can be interpreted topologically. In this section we present a way of topologically interpreting RCC-8-models such that all parts of the models can be interpreted consistently on a topological level. The $\mathbf{I}$-operator will not be interpreted as an interior operator, but we will prove that it suffices the intended interpretation.

Because $\mathbf{I}$ is an S4-operator and because of the additional constraints $m_2(X)$, exactly one of the following formulas is true for every world w of $\mathcal{M}$ and every propositional atom X (see Figure 2).

1. $\mathcal{M}, w \Vdash \mathbf{I}\mathsf{X}$
2. $\mathcal{M}, w \Vdash \mathbf{I}\neg\mathsf{X}$
3. $\mathcal{M}, w \Vdash \mathsf{X} \wedge \neg\mathbf{I}\mathsf{X}$

[2] As the number of spatial variables is countable, the number of RCC-8-clusters is also countable.

[3] The RCC-8-structure does not cover all possible Kripke models of $m(\Theta)$. The goal of a canonical model is just to provide a model for any consistent sentence of a calculus, not to cover all possible models.

Consider a particular world w. Then the set of all spatial variables can be divided into three disjoint sets of spatial variables according to which of the three possible formulas is true in w. Let $\mathcal{I}_w$, $\mathcal{E}_w$, and $\mathcal{B}_w$ be the sets where the first, the second, and the third formula is true in w, respectively, i.e., $\mathcal{M}, w \Vdash \mathbf{I}\mathsf{X} \wedge \mathbf{I}\neg\mathsf{Y} \wedge (\mathsf{Z} \wedge \neg\mathbf{I}\mathsf{Z})$ for all $X \in \mathcal{I}_w$, $Y \in \mathcal{E}_w$, and $Z \in \mathcal{B}_w$.

When looking at points in a topological space, for every region there are three different kinds of points: interior points, exterior points, and boundary points of a region. If a point is interior or exterior point of a region, there is a neighborhood of the point such that all points of the neighborhood are also inside or outside the region, respectively. If a point is boundary point of a region, every neighborhood contains points inside and points outside the region (see Definition 2.2).

There seems to be a correspondence between worlds and points of a topological space, and between the accessibility relation $R_{\mathbf{I}}$ and topological neighborhoods. In the following lemma we further investigate this correspondence by deriving topological constraints from the modal formulas.

Lemma 4.1 *Let X and Y be two spatial variables of Θ. Depending on which sets $\mathcal{I}_w, \mathcal{E}_w$, or $\mathcal{B}_w$ they are contained in for a world w, the following relations between X and Y are impossible. This has some topological consequences on possible instantiations $\mathtt{X}, \mathtt{Y}$:*

X	Y	*Impossible Relations*	*Consequences*
$\mathcal{I}_w$	$\mathcal{I}_w$	DC, EC	$i(\mathtt{X}) \cap i(\mathtt{Y}) \neq \emptyset$
$\mathcal{I}_w$	$\mathcal{E}_w$	TPP, NTPP, EQ	$i(\mathtt{X}) \cap e(\mathtt{Y}) \neq \emptyset$
$\mathcal{I}_w$	$\mathcal{B}_w$	DC, EC, TPP, NTPP, EQ	$i(\mathtt{X}) \cap b(\mathtt{Y}) \neq \emptyset$
$\mathcal{E}_w$	$\mathcal{E}_w$	–	–
$\mathcal{E}_w$	$\mathcal{B}_w$	TPP^{-1}, NTPP^{-1}, EQ	$e(\mathtt{X}) \cap b(\mathtt{Y}) \neq \emptyset$
$\mathcal{B}_w$	$\mathcal{B}_w$	DC, NTPP, NTPP^{-1}	$b(\mathtt{X}) \cap b(\mathtt{Y}) \neq \emptyset$[4]

Proof: Most entries in the table follow immediately from the encoding of the relations in modal logic. The only more difficult entry is the relation $\mathsf{EC}(X, Y)$ in the third line of the table. This relation is not possible because of the property $\Box(\mathsf{Y} \rightarrow \neg\mathbf{I}\neg\mathbf{I}\mathsf{Y})$ which states that for any world w that satisfies Y there is a world v with $wR_{\mathbf{I}}v$ that satisfies $\mathbf{I}\mathsf{Y}$. As v also satisfies $\mathbf{I}\mathsf{X}$, the model constraint of $\mathsf{EC}(X, Y)$ is violated, so this relation is not possible. The topological consequences result from distinguishing the impossible from the possible relations. □

[4]If $\mathsf{PO}(X, Y)$ holds, X and Y do not necessarily have a common boundary point if one of them is not internally connected. However, assuming $b(\mathtt{X}) \cap b(\mathtt{Y}) \neq \emptyset$ in this case does not contradict any spatial formula since RCC-8 is not expressive enough to distinguish different kinds of partial overlap.

It can be seen that when, e.g., $\mathbf{I}\mathsf{X}$ and $\mathbf{I}\mathsf{Y}$ hold in a world w, then $\mathtt{X}$ and $\mathtt{Y}$ must have a common interior. So, there is a common interior point of $\mathtt{X}$ and $\mathtt{Y}$ where w can be mapped to.

Theorem 4.2 *Let Θ be a consistent set of spatial formulas, $m(\Theta)$ be the modal encoding of Θ, and $\mathcal{M} = \langle W, \{R_\Box, R_{\mathbf{I}}\}, \pi\rangle$ be an RCC-8-model of $m(\Theta)$. Then there is a function $p : W \mapsto \mathcal{U}$ that maps each world $w \in W$ to a point $p(w) \in \mathcal{U}$ and a function $N : W \mapsto 2^{\mathcal{U}}$ that assigns each world $w \in W$ a neighborhood $N(w)$ of $p(w)$ such that $p(w)$ is in the interior of $\mathtt{X}$ if $\mathcal{M}, w \Vdash \mathbf{I}\mathsf{X}$ holds, $p(w)$ is in the exterior of $\mathtt{X}$ if $\mathcal{M}, w \Vdash \mathbf{I}\neg\mathsf{X}$ holds, $p(w)$ is on the boundary of $\mathtt{X}$ if $\mathcal{M}, w \Vdash \mathsf{X} \wedge \neg\mathbf{I}\mathsf{X}$ holds, and $p(u) \in N(w)$ if and only if $wR_{\mathbf{I}}u$ holds.*[5]

Proof: Let w be a world of W and $\mathcal{I}_w, \mathcal{E}_w$, and $\mathcal{B}_w$ be the corresponding sets of spatial variables. We assume that there is a realization of Θ such that there is at least one point in the topological space that is in the interior of every $\mathtt{X}$, in the exterior of every $\mathtt{Y}$, and on the boundary of every $\mathtt{Z}$ simultaneously ($X \in \mathcal{I}_w$, $Y \in \mathcal{E}_w$, $Z \in \mathcal{B}_w$). It follows from Lemma 4.1 that this is true for every pair of regions. As RCC-8 permits only binary constraints between spatial variables and regions are allowed to be internally disconnected, this assumption holds. We further assume that p maps w to one of these points.

Because of Definition 2.2, there must be neighborhoods $N_X(w)$ and $N_Y(w)$ of $p(w)$ for every $X \in \mathcal{I}_w$ and every $Y \in \mathcal{E}_w$ such that $N_X(w)$ is in the interior of $\mathtt{X}$ and $N_Y(w)$ is disjoint with $\mathtt{Y}$. Also, for every $Z \in \mathcal{B}_w$, every neighborhood $N_Z(w)$ of $p(w)$ contains points inside and outside $\mathtt{Z}$. All these neighborhoods are members of the neighborhood system of $p(w)$, so their intersection $N(w)$ is also a neighborhood of $p(w)$ where all $R_{\mathbf{I}}$-successors of w can be mapped to. □

Using the above defined functions p and N, $\mathcal{M}, w \Vdash \mathbf{I}\mathsf{X}$ can be interpreted as "there is a neighborhood $N(w)$ of $p(w)$ such that all points of $N(w)$ are in $\mathtt{X}$". This obeys the intended interpretation of $\mathbf{I}$ as an interior operator, as $\mathcal{M}, w \Vdash \mathsf{X}$ means that $p(w)$ is in $\mathtt{X}$ and $\mathcal{M}, w \Vdash \mathbf{I}\mathsf{X}$ means that $p(w)$ is in the interior of $\mathtt{X}$.

The function N, as defined in Theorem 4.2, can be replaced by any function $N' : W \mapsto 2^{\mathcal{U}}$, with $N'(w) \subseteq N(w)$ for all $w \in W$, if $N'(w)$ is a member of the neighborhood system of $p(w)$. p has to be changed accordingly. In particular, we will regard in the following

[5]The properties for $R_\Box$ ($p(u) \in \mathcal{U}$ if $wR_\Box u$ holds and $p(w) \in \mathcal{U}$) can be omitted as we already defined N and p such that only points of $\mathcal{U}$ are used.

all neighborhoods as d-dimensional spheres.

In order to make the following argumentation easier, a world mapped to an interior point of X is denoted *interior world* of X, a world mapped to an exterior point of X *exterior world* of X, and a world mapped to a boundary point of X *boundary world* of X. Accordingly, a region is called *interior, exterior* or *boundary region* of a world. In particular, a world w with $\mathcal{M}, w \Vdash \mathbf{I}\mathsf{X}$ is an interior world of X, a world w with $\mathcal{M}, w \Vdash \mathbf{I}\neg\mathsf{X}$ is an exterior world of X, and a world w with $\mathcal{M}, w \Vdash \mathsf{X} \wedge \neg\mathbf{I}\mathsf{X}$ is a boundary world of X.

5 RCC8-MODELS AND THE DIMENSION OF SPACE

In the previous section we have shown how the RCC-8-models introduced in Section 3 can be mapped to topological space, but we still have no information about the dimension of the topological space. In this section we investigate the influence of dimension on the possibility to map the RCC-8-models to the topological space, i.e., which dimension is required in order to find a realization of a consistent set of spatial formulas Θ. We will start with proving that for any RCC-8-model there is a realization in two-dimensional space. It is sufficient to prove this only for sets of base relations as every realization of Θ uses only base relations.[6] For this proof it is important to keep in mind that regions do not have to be internally connected, i.e., they might consist of different disconnected pieces. It will turn out that our proof leads to realizations in any dimension $d \geq 1$. Finally, for three- and higher-dimensional space we will prove that every consistent set Θ can also be realized with internally connected regions.

For the following examination we restrict regions to be sets of d-dimensional polytopes. Sets are required since regions might consist of several disconnected pieces where each piece is a single polytope. This restriction will be lifted later and the results can be generalized to arbitrary regular regions.

Let Θ be a consistent set of spatial formulas and $\mathcal{M}$ be an RCC-8-model of $m(\Theta)$, the modal encoding of Θ. Suppose that only two-dimensional regions are permitted, i.e., the topological space is a two-dimensional plane $\mathcal{U}$. All worlds of $\mathcal{M}$ are mapped to points of $\mathcal{U}$ as specified in Theorem 4.2. The general intuition of the proof is that every RCC-8-cluster, i.e., every world of level 0 together with its $R_{\mathbf{I}}$-successors is mapped to an independent neighborhood such that each neighborhood can be placed on an arbitrary but distinct position on the plane. Each neighborhood will then be extended to different closed sets that form the pieces of the spatial regions. In the following we will study the requirements neighborhoods have to meet in order to guarantee two-dimensional realizations.

[6] The relation EQ can be omitted as any pair of spatial variables X and Y with $\mathsf{EQ}(X, Y)$ can be combined to a single spatial variable.

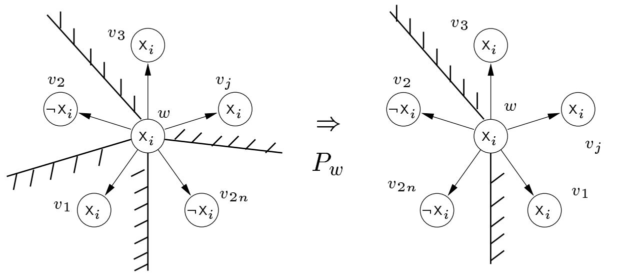

Figure 3: Permutation P_w of the $R_{\mathbf{I}}$-successors of a world w. The solid line indicates the boundary of X, the hashed region the interior of X.

For every spatial variable X_i $(1 \leq i \leq n)$ and every world w of level 0, we define a *region vector* $r_i^w = (r_{i,1}^w, \ldots, r_{i,2n}^w)$ that represents the affiliation of the $2n$ $R_{\mathbf{I}}$-successors of w to X_i, i.e., $r_{i,j}^w = 1$ if $\mathcal{M}, v_j \Vdash \mathsf{X}_i$ and $r_{i,j}^w = 0$ if $\mathcal{M}, v_j \not\Vdash \mathsf{X}_i$ where v_j is the jth $R_{\mathbf{I}}$-successor of w. Since in the two-dimensional case the neighborhood $N(w)$ is a disc, we suppose that the points $p(v_j)$ corresponding to the $R_{\mathbf{I}}$-successors v_j of w are clock-wisely ordered within the disc according to j. If $p(w)$ is a boundary point of $\mathtt{X}_i$, some values of r_i^w are 1 and some are 0. Otherwise all values of r_i^w are either 1 (if $p(w)$ is contained in $\mathtt{X}_i$) or 0 (if $p(w)$ is not contained in $\mathtt{X}_i$).

Lemma 5.1 *If for every world w of level 0 there is a permutation P_w of the values of r_i^w such that $(r_{i,P_w(1)}^w, \ldots, r_{i,P_w(2n)}^w)$ is a bitonic sequence for all $1 \leq i \leq n$, then the neighborhoods $N(w)$ can be placed in a two-dimensional plane such that all spatial relations are satisfied within the neighborhoods.*

Proof: If r_i^w is a bitonic sequence, i.e., the values of r_i^w are in a form $0^e 1^f 0^g$ or $1^e 0^f 1^g$ for $e, f, g \geq 0$, and $p(w)$ is a boundary point of $\mathtt{X}_i$, then the mappings of the worlds of level 1 corresponding to the values of r_i^w can be separated into points inside $\mathtt{X}_i$ and points outside $\mathtt{X}_i$ by at most two line segments meeting at $p(w)$ (see Figure 3). These line segments can be regarded as the part of the boundary of $\mathtt{X}_i$ which is inside $N(w)$. So, neighborhoods can be separated in an interior and an exterior part of a region by a one-dimensional boundary. Therefore all neighborhoods can be placed in a two-dimensional plane. As the permutation of the $R_{\mathbf{I}}$-successors has no influence on the relations between the regions, all spatial relations between the regions hold within the neighborhoods. □

Actually, a permutation as described in the previous lemma is not necessary to guarantee two-dimensional realizations. A region might look as shown on the left of Figure 3, but in this case we restrict the shape and the internal connection of the regions by the neighborhoods we are using which is not at all desirable. However, a permutation as described in Lemma 5.1 is necessary for one-dimensional realizations and realizations with internally connected regions.

Since a permutation P_w is only necessary for boundary worlds, we will in the following regard only those particular $\mathsf{RCC\text{-}8}$-models $\mathcal{M}$ for which only those worlds are boundary worlds of regions which are explicitly forced to be boundary worlds of these regions by the constraints. Therefore, we have to take a closer look at which worlds are introduced as boundary worlds of some regions by the entailment constraints, and which worlds are forced to be boundary worlds of regions by the constraints. As a world w of level 0 is forced to be a boundary world of X if $\mathcal{M}, w \Vdash \mathsf{X}$ and $\mathcal{M}, v \not\Vdash \mathsf{X}$ hold for a world v with $wR_{\mathbf{I}}v$, we have to find out which of the model and entailment constraints force $\mathcal{M}, w \Vdash \mathsf{X}$ if $\mathcal{M}, v \not\Vdash \mathsf{X}$ holds or force $\mathcal{M}, v \not\Vdash \mathsf{X}$ if $\mathcal{M}, w \Vdash \mathsf{X}$ holds.

Proposition 5.2 *Boundary worlds are introduced only by the following relations (see Table 1):*

1. $\mathsf{EC}(X,Y)$: $\neg\Box(\neg(\mathsf{X} \wedge \mathsf{Y}))$ *introduces a boundary world of X and Y because of* $\Box(\neg(\mathbf{I}\mathsf{X} \wedge \mathbf{I}\mathsf{Y}))$.
2. $\mathsf{TPP}(X,Y)$: $\neg\Box(\mathsf{X} \rightarrow \mathbf{I}\mathsf{Y})$ *introduces a boundary world of X and Y because of* $\Box(\mathsf{X} \rightarrow \mathsf{Y})$.
3. $\mathsf{TPP}^{-1}(X,Y)$: $\neg\Box(\mathsf{Y} \rightarrow \mathbf{I}\mathsf{X})$ *introduces a boundary world of X and Y because of* $\Box(\mathsf{Y} \rightarrow \mathsf{X})$.

Apart from the above worlds that are introduced as boundary worlds of particular regions, worlds can also be forced to be boundary worlds of other regions.

Proposition 5.3 *A world w is forced to be a boundary world of X only with the following constraints:*

1. $\Box(\mathsf{X} \rightarrow \mathsf{Y})$: *If w is a boundary world of Y and X is true in w, then w must also be a boundary world of X.*
2. $\Box(\mathsf{Y} \rightarrow \mathsf{X})$: *If w is a boundary world of Y and $\neg\mathsf{X}$ is true in an $R_{\mathbf{I}}$-successor of w, then w must also be a boundary world of X.*
3. $\Box(\neg(\mathbf{I}\mathsf{X} \wedge \mathbf{I}\mathsf{Y}))$: *If X and Y are true in w, then w must be a boundary world of X and Y.*

For the constraints $\Box(\mathsf{X} \rightarrow \mathsf{Y})$ and $\Box(\mathsf{Y} \rightarrow \mathsf{X})$, w must already be a boundary world of some other region, so w must be introduced by one of the relations $\mathsf{EC}(X,Y)$, $\mathsf{TPP}(X,Y)$, or $\mathsf{TPP}^{-1}(X,Y)$. If w is forced to be a boundary world of X and Y with the constraint $\Box(\neg(\mathbf{I}\mathsf{X} \wedge \mathbf{I}\mathsf{Y}))$, then X and Y must both be true in w. This can only be forced when there is a $Z_1 \in Reg(\Theta)$ with $\mathsf{TPP}(Z_1, X)$ and Z_1 is true in w, a $Z_2 \in Reg(\Theta)$ with $\mathsf{TPP}(Z_2, Y)$ and Z_2 is true in w, and w is a boundary world of Z_1 and Z_2 introduced by $\mathsf{EC}(Z_1, Z_2)$. So, in any case when a world is forced to be a boundary world of some region it must already be a boundary world of other regions introduced as described in Proposition 5.2.

We will now have a look at how regions must be related in order to force a world to be a boundary world of these regions using the constraints of Proposition 5.3. Suppose that w is a boundary world of X and Y introduced by either $\mathsf{EC}(X,Y)$ or $\mathsf{TPP}(X,Y)$.[7] We will write $X|Y$ in order to express that we can either use X or Y but always the same. With one of the following constraints it can be forced that w is also a boundary world of $Z \neq X, Y$ (v is an $R_{\mathbf{I}}$-successor of w):

$$\begin{array}{ll}
\Box(\mathsf{Z} \rightarrow (\mathsf{X}|\mathsf{Y})) \text{ and } \mathcal{M}, w \Vdash \mathsf{Z} & (\rightsquigarrow \mathsf{TPP}(Z, X|Y)) \\
\Box(\neg(\mathbf{I}\mathsf{Z} \wedge \mathbf{I}(\mathsf{X}|\mathsf{Y}))) \text{ and } \mathcal{M}, w \Vdash \mathsf{Z} & (\rightsquigarrow \mathsf{EC}(Z, X|Y)) \\
\Box((\mathsf{X}|\mathsf{Y}) \rightarrow \mathsf{Z}) \text{ and } \mathcal{M}, v \Vdash \neg\mathsf{Z} & (\rightsquigarrow \mathsf{TPP}^{-1}(Z, X|Y))
\end{array}$$

$\mathcal{M}, w \Vdash \mathsf{Z}$ is forced with the following constraint:

$$\Box(\mathsf{U} \rightarrow \mathsf{Z}) \text{ and } \mathcal{M}, w \Vdash \mathsf{U} \qquad (\rightsquigarrow \mathsf{TPP}(U, Z))$$

$\mathcal{M}, v \Vdash \neg\mathsf{Z}$ is forced with the following constraints:

$$\begin{array}{ll}
\Box(\mathsf{Z} \rightarrow \mathsf{U}) \text{ and } \mathcal{M}, v \Vdash \neg\mathsf{U} & (\rightsquigarrow \mathsf{TPP}^{-1}(U, Z)) \\
\Box(\neg(\mathbf{I}\mathsf{Z} \wedge \mathbf{I}\mathsf{U})) \text{ and } \mathcal{M}, v \Vdash \mathsf{U} & (\rightsquigarrow \mathsf{EC}(U, Z))
\end{array}$$

When we compose these relations (written as $\circ$) , we obtain the possible relations between U and $X|Y$.

$R(U,Z)$	$S(Z, X\|Y)$	$(R \circ S)(U, X\|Y)$
TPP	TPP	$\mathsf{TPP}, \mathsf{NTPP}$
TPP	EC	DC, EC
TPP^{-1}	TPP^{-1}	$\mathsf{TPP}^{-1}, \mathsf{NTPP}^{-1}$
EC	TPP^{-1}	DC, EC

As w is a boundary world of X and Y, $\mathsf{DC}(U, X|Y)$ and $\mathsf{NTPP}(U, X|Y)$ are not possible together with $\mathcal{M}, w \Vdash \mathsf{U}$, and $\mathsf{NTPP}^{-1}(U, X|Y)$ is not possible together with $\mathcal{M}, v \Vdash \neg\mathsf{U}$. In order to force $\mathcal{M}, w \Vdash \mathsf{U}$,

[7] We omit $\mathsf{TPP}^{-1}(X,Y)$ as $\mathsf{TPP}^{-1}(X,Y) = \mathsf{TPP}(Y,X)$ and the order is not important.

there must be a sequence of spatial variables U_i with $\mathsf{TPP}(U_1, U)$, $\mathsf{TPP}(U_{i+1}, U_i)$, until there is a U_m that is equal to X or Y, so $\mathsf{TPP}(X, Z)$ or $\mathsf{TPP}(Y, Z)$ must hold. In order to force $\mathcal{M}, v \Vdash \neg\mathsf{U}$, there must be a sequence of spatial variables U_i with $\mathsf{TPP}^{-1}(U_1, U)$, $\mathsf{TPP}^{-1}(U_{i+1}, U_i)$, and $\mathsf{EC}(U_j, U_{j-1})$ and $\mathcal{M}, v \Vdash \mathsf{U}_j$ must hold. In order to force $\mathcal{M}, v \Vdash \mathsf{U}_j$, there must be a sequence of TPP-related spatial variables, as described above, until one of them is equal to X or Y, so $\mathsf{EC}(Z, X)$ or $\mathsf{EC}(Z, Y)$ must hold. This results in only three different possibilities of how w is forced to be a boundary world of Z if w was introduced as a boundary world of X and Y.

a. $\mathsf{TPP}(X, Y), \mathsf{TPP}(X, Z)$, and $\mathsf{TPP}(Z, Y)$ hold.

b. $\mathsf{EC}(X, Y), \mathsf{TPP}(X, Z)$, and $\mathsf{EC}(Z, Y)$ hold.

c. $\mathsf{EC}(X, Y), \mathsf{TPP}(Y, Z)$ and $\mathsf{EC}(Z, X)$ hold.

As different spatial variables Z_i, Z_j, for which w is forced to be a boundary world of, all have the boundary world w in common, only the relations $\mathsf{EC}, \mathsf{PO}, \mathsf{TPP}$, or TPP^{-1} can hold between them.

We have shown that only those worlds are boundary worlds which are introduced as boundary worlds of some regions by the entailment constraints, and, further, that other regions are only forced to be boundary regions of these worlds when they are related in a particular way. This will be used in the following lemma.

Lemma 5.4 *Let $\mathcal{M}$ be an* RCC-8*-model. Then two different types of $R_{\mathbf{I}}$-successors are sufficient for every world w of level 0.*

Proof: If w is not a boundary world of some region, all $R_{\mathbf{I}}$-successors of w satisfy exactly the same formulas as w. Otherwise, w is introduced as a boundary world by either $\mathsf{EC}(X, Y)$ or $\mathsf{TPP}(X, Y)$ (see Proposition 5.2). Let w be forced to be a boundary world of the spatial variables Z_i. For $\mathsf{EC}(X, Y)$, some of the $R_{\mathbf{I}}$-successors of w satisfy X but not Y, and some satisfy Y but not X, the others neither satisfy X nor Y. For all Z_i with $\mathsf{TPP}(X, Z_i)$ and $\mathsf{EC}(Z_i, Y)$ and all Z_j with $\mathsf{TPP}(Y, Z_j)$ and $\mathsf{EC}(Z_j, X)$, all $R_{\mathbf{I}}$-successors of w satisfy Z_i if they satisfy X and satisfy Z_j if they satisfy Y. So all $R_{\mathbf{I}}$-successors of w that satisfy X satisfy the same formulas, and all $R_{\mathbf{I}}$-successors of w that satisfy Y satisfy the same formulas. For the $R_{\mathbf{I}}$-successors v of w which do not satisfy X or Y, there are only two requirements: if $\mathsf{TPP}(Z_{k'}, Z_k)$ holds then Z_k must be true in v whenever $\mathsf{Z}_{k'}$ is true in v; if $\mathsf{EC}(Z_k, Z_{k'})$ holds then Z_k and $\mathsf{Z}_{k'}$ must not both be true in v. However, there is no constraint that forces the existence of these worlds, so it can be assumed that all $R_{\mathbf{I}}$-successors

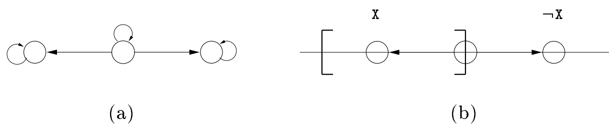

Figure 4: (a) shows a reduced RCC-8-cluster of the reduced RCC-8-structure. (b) shows how a neighborhood can be placed in one-dimensional space. The two brackets indicate a one-dimensional region `X` where the neighborhood represents a boundary point of `X`.

of w satisfy either X or Y. As the respective worlds all satisfy the same formulas, two different kinds of $R_{\mathbf{I}}$-successors of the boundary world w introduced by $\mathsf{EC}(X, Y)$ are sufficient.

For $\mathsf{TPP}(X, Y)$, all $R_{\mathbf{I}}$-successors of w that satisfy X also satisfy Y, all $R_{\mathbf{I}}$-successors of w that do not satisfy Y also do not satisfy X, and some $R_{\mathbf{I}}$-successors of w satisfy Y but not X. For all Z_i with $\mathsf{TPP}(X, Z_i)$ and $\mathsf{TPP}(Z_i, Y)$, all $R_{\mathbf{I}}$-successors of w satisfy Z_i if they satisfy X. For the $R_{\mathbf{I}}$-successors of w that satisfy Y but not X, there is only one requirement, namely, that Z_k must be true whenever $\mathsf{Z}_{k'}$ is true in these worlds for any two spatial variables $Z_k, Z_{k'}$ with $\mathsf{TPP}(X, Z_k), \mathsf{TPP}(Z_{k'}, Y)$ and $\mathsf{TPP}(Z_k, Z_{k'})$. However, there is again no constraint that forces the existence of these worlds, so it can be assumed that all $R_{\mathbf{I}}$-successors of w satisfy X if they satisfy Y. □

Whether a boundary world w is introduced by $\mathsf{EC}(X, Y)$ or by $\mathsf{TPP}(X, Y)$, in both cases two different kinds of $R_{\mathbf{I}}$-successors are sufficient. Thus, by grouping together the respective $R_{\mathbf{I}}$-successors for every world w of level 0 of $\mathcal{M}$, we can always find a permutation of the worlds of level 1 such that r^w is a bitonic sequence for all regions.

Instead of having $2n$ $R_{\mathbf{I}}$-successors for every world of level 0 from which we know that they belong to only two different types, it is sufficient two use only two $R_{\mathbf{I}}$-successors for every world of level 0. This leads to a very simple canonical model shown in Figure 4a. We call this a *reduced* RCC-8*-structure* and the corresponding models *reduced* RCC-8*-models*. They are defined in the same way as in Definition 3.1 except that we have exactly two worlds of level 1 instead of $2n$ worlds.

We can now apply Lemma 5.1 and place all neighborhoods independently on the plane while all relations between spatial regions hold within the neighborhoods. Thereby, neighborhoods corresponding to non-boundary worlds are homogeneous in the sense that all points within one of these neighborhoods have the same topological properties. Neighborhoods corresponding to a boundary world w consist of two homogeneous parts corresponding to the two $R_{\mathtt{I}}$-successors of w. These two parts are divided by the common boundary of the boundary regions of w (see Figure 5a).

In order to obtain a realization, we have to find regions such that the relations between them hold in the whole plane and not just within the neighborhoods. Since regions do not have to be internally connected, it is possible to compose every region out of pieces resulting from the corresponding neighborhoods, i.e., for every neighborhood a region is affiliated with, we generate a piece of that region. As the neighborhoods are open sets and regions as well as their pieces must be regular closed sets, we have to *close* every neighborhood, i.e., find a closed set $\mathtt{X}^w$ for every region $\mathtt{X}$ and every neighborhood $N(w)$ with $\mathcal{M}, w \Vdash \mathsf{X}$ such that all relations hold between the regions composed of the pieces. As all neighborhoods are independent of each other, we only have to make sure that the relations of the different pieces corresponding to a single neighborhood do not violate the relations of the compound regions. This can be done independently for every neighborhood.

Consider a particular neighborhood $N(w)$. If w is not a boundary world, then only the relations PO, TPP, NTPP, and their converse are possible between the regions affiliated with $N(w)$, since they share $N(w)$ as their common interior. For closing the neighborhood $N(w)$, all pieces must fulfill the "part of" relations whereas the PO relations cannot be violated as long as the corresponding pieces have a common interior.

One possibility to fulfill the "part of" relations is using a hierarchy H_Θ of the regions, where a region $\mathtt{X}$ is of level $H_\Theta(\mathtt{X}) = 1$ if there is no region $\mathtt{Y}$ which is part of $\mathtt{X}$. A region $\mathtt{X}$ is of level $H_\Theta(\mathtt{X}) = k$ if there is a region $\mathtt{Y}$ of level $H_\Theta(\mathtt{Y}) = k - 1$ which is part of $\mathtt{X}$ and if there is no region $\mathtt{Z}$ which is part of $\mathtt{X}$ and has a higher level than $H_\Theta(\mathtt{Z}) = k - 1$.[8] The pieces of all regions affiliated with $N(w)$ must then be chosen according to H_Θ, i.e., pieces of regions of the same level are equal for this particular neighborhood and are non-tangential proper part of all pieces of regions of a higher level. We choose the single pieces to be rectangles.

If w is a boundary world, the boundary regions of w are only affiliated with one part of $N(w)$ and their pieces must share the common boundary. Therefore, both parts of $N(w)$ must be closed separately according to H_Θ (see Figure 5c). In the same way as for two-dimensional space, neighborhoods can be placed in any higher dimensional space and closed therein according to H_Θ. As the three points corresponding to a world w of level 0 and its two $R_{\mathtt{I}}$-successors can always be aligned, $N(w)$ can also be placed on a line. Thus, all neighborhoods can be placed independently in a one-dimensional space and closed as intervals according to H_Θ (see Figure 4b).

[8]This corresponds to the finish time of depth-first search for each vertex of a graph G_Θ where regions are vertices V_Θ and "part of" relations are directed edges E_Θ, computable in time $O(V_\Theta + E_\Theta)$ (Cormen *et al.* 1990, p.477ff)

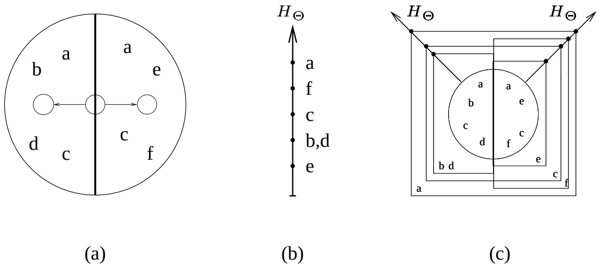

Figure 5: (a) shows the two-dimensional neighborhood of a boundary world which is divided into two parts by the common boundary of the boundary regions b, d, e, and f. (b) shows a possible hierarchy H_Θ of regions. In (c) the neighborhood is closed with respect to H_Θ.

Theorem 5.5 *Every consistent set of spatial formulas can be realized in any dimension $d \geq 1$ where regions are (sets of) d-dimensional polytopes.*

So far all regions consist of as many pieces as there are neighborhoods affiliated with them, i.e., $O(n^2)$ many pieces for every region. We can further show that for three- and higher-dimensional space all regions can also be realized as internally connected. For this we construct a $d + 1$-dimensional realization of internally connected regions by connecting all pieces of the same regions of a d-dimensional realization of internally disconnected regions.

Theorem 5.6 *Every consistent set Θ of spatial formulas can be realized with internally connected regions in any dimension $d \geq 3$ where regions are polytopes.*

Proof: Suppose that Θ is consistent. With the following construction we obtain a three-dimensional realization of internally connected regions starting from a two-dimensional realization. (1a) Place all neighborhoods on a circle in the plane determined by the x- and y-axes such that the common boundary of every neighborhood corresponding to a boundary world points to the center of the circle, the z-axis (see Figure 6a). (1b) Close all neighborhoods according to the hierarchy H_Θ such that all pieces of regions are rectangles. (2a) Pro-

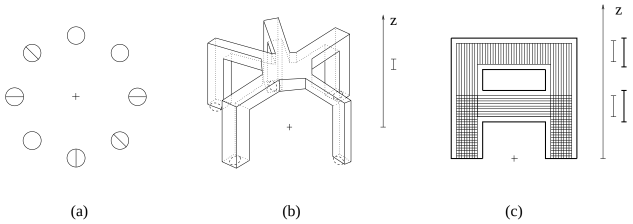

Figure 6: Construction of the three-dimensional realization. (a) placing the two-dimensional neighborhoods on a circle. (b) connecting the pieces of a region on a particular level. (c) connecting the pipes of a region (bold line) that contains the vertically and the horizontally hashed regions.

ceed from this two-dimensional realization according to H_Θ by first choosing pairwise distinct intervals on the positive z-axis for every region with $H_\Theta = 1$, i.e., for the regions that do not contain any other region. (2b) Build a pipe parallel to the z-axis for every piece of these regions starting from the plane (z=0) up to the endpoint of the corresponding interval. (2c) Connect the pipes of the same region within the range of the corresponding interval using pipes pointing to the z-axis (see Figure 6b). (3) Next the regions with $H_\Theta = 2$ are connected, i.e., those regions that only contain already connected regions. To do this, (3a) choose intervals on the z-axis for these regions such that the intervals contain all intervals of the contained regions but do not overlap with any other interval. (3b) Build a pipe for every piece up to the endpoint of the corresponding interval with the largest z-value, and (3c) connect the pipes of every region within the range of all corresponding intervals (see Figure 6c). (4) Repeat step 3. successively for every level of H_Θ until all regions are connected. (5) Finally, close all neighborhoods on the negative z-axis according to H_Θ.

Obviously, with this construction all regions are internally connected. Furthermore all internally connected three-dimensional regions hold the same base relations as the two-dimensional realizations from which we started the construction. This is because all intervals on the z-axis are either contained in each other or are distinct, they have no common boundary points. All intervals corresponding to region X are contained in the intervals of region Y if and only if $\mathsf{NTPP}(X,Y)$ or $\mathsf{TPP}(X,Y)$. When two regions are disconnected they remain disconnected as they are not affiliated with the same neighborhoods. Two externally connected regions remain externally connected because every neighborhood was placed on the circle such that the common boundary points to its center. Therefore, if two of these regions are both affiliated with the same neighborhood, their pipes are externally connected and the horizontal connection of the single pipes is distinct. All other requirements of relations as, e.g., a common boundary point are already met by the pipes.

With the same construction, a $d+1$-dimensional realization of internally connected regions can be obtained from a d-dimensional realization of internally disconnected regions. All constructions kept the polytopic shape of the regions, so every region can be realized as a (d-dimensional) polytope. □

The restriction of regions to be polytopes can immediately be generalized to an arbitrary shape of regions.

6 APPLICABILITY OF THE CANONICAL MODEL

In the previous sections we reported about the existence of (reduced) RCC-8-models and how they can be mapped to topological spaces of different dimensions. In this section we study how RCC-8-models can be determined and how a realization can be generated from them. As there is a (reduced) RCC-8-model $\mathcal{M}$ for every consistent set of spatial relations Θ, and as it is always possible to generate a realization of $\mathcal{M}$, RCC-8 models are suitable for representing spatial regions with respect to their relations. RCC-8-models represent the characteristic points and information about their neighborhoods of a possible realization.

6.1 DETERMINATION OF RCC-8-MODELS

Given a set of spatial formulas Θ, we have to find a reduced RCC-8-model $\mathcal{M}$ for the modal encoding of RCC-8 such that only those worlds are boundary worlds of regions which are forced to be by the constraints. The Kripke frame of $\mathcal{M}$, i.e., the number of worlds and their accessibility relations are already known from the entailment constraints, but we have to find a valuation for every world and every region. For some worlds and some regions the valuation is already given from the constraints, for some it can be inferred using the constraints, for others it can be chosen. In order to make the inference step as easy as possible, we use the propositional encoding of RCC-8 with respect to a Kripke frame where every world w and every spatial variable X is transformed to a propositional atom X_w which is true if and only if X holds in w (Renz and Nebel 1997). The valuation of $\mathcal{M}$ can then be obtained from the satisfying assignment of the propositional formula. Even if the encoding of the reduced

RCC-8-models is not a Horn formula,[9] unit-resolution plus additional choices is sufficient for finding a satisfying assignment. As all clauses of the propositional encoding use worlds of the same RCC-8-cluster, the inference step is independent for every cluster. From Proposition 5.2 it is known which RCC-8-clusters contain a boundary world. Suppose that an RCC-8-cluster contains a boundary world, then the valuation of the two regions which introduced the boundary world can be chosen in all worlds of the RCC-8-cluster according to the relation of the two regions. The valuations of the other regions are either determined by unit-resolution or can be chosen according to their other valuations: If the valuation of a particular region in some world of the RCC-8-cluster is true, then the other valuations are also chosen as true, otherwise all valuations are chosen as false. If an RCC-8-cluster does not contain a boundary world, all worlds of the RCC-8-cluster have the same valuation. If the valuation of a region is not determined by unit-resolution it is chosen as false. With these choices a satisfying assignment is always found, even though the propositional formula is not Horn. As there are $O(n^2)$ worlds and n regions, there are $O(n^4)$ clauses (Renz and Nebel 1997), so a reduced RCC-8-model can be determined in time $O(n^4)$.

6.2 GENERATING A REALIZATION

Suppose we have given a reduced RCC-8-model of a consistent set of RCC-8 formulas Θ. We have to distinguish the tasks of generating a realization of internally connected and disconnected regions. A realization of disconnected regions in d-dimensional space can be obtained by placing the $O(n^2)$ different neighborhoods in the d-dimensional space and close each neighborhood as specified in Section 5. For this, the hierarchy H_Θ of regions must be known, which can be computed in time $O(n + P_\Theta)$ where $P_\Theta \in O(n^2)$ is the number of "part of" relations in Θ (see Footnote 8). Let $A_\Theta \in O(n)$ be the maximal number of regions affiliated with a neighborhood, then the closure of a neighborhood can be computed in time $O(A_\Theta)$.

Theorem 6.1 *Given a reduced* RCC-8 *model of a set of* RCC-8 *relations* Θ*, a realization in* d*-dimensional space* ($d \geq 1$) *can be generated in time* $O(n^2 A_\Theta)$ *when regions are allowed to be disconnected.*

In order to generate a realization of internally connected regions we can use the construction of the proof to Theorem 5.6. For every region we have to find the corresponding intervals on the z-axis. The number of intervals of a particular region X is equal to the number of regions with $H_\Theta = 1$ that are contained in X. Let $I_\Theta \in O(n)$ be the maximal number of regions with $H_\Theta = 1$ that are contained in a region.

Theorem 6.2 *Given a reduced* RCC-8 *model of a set of* RCC-8 *relations* Θ*, a realization of internally connected regions in* d*-dimensional space* ($d \geq 3$) *can be generated in time* $O(n^2 A_\Theta I_\Theta)$.

If $P_\Theta \in O(n^2)$ is the maximal number of neighborhoods affiliated with a region, every region can be realized as a polytope with $O(P_\Theta I_\Theta)$ vertices.

7 DISCUSSION & RELATED WORK

There is some work on identifying canonical models for the RCC axioms, i.e., determining what mathematical structures fulfill all the RCC axioms, as, e.g., every region has a non-tangential proper part (Randell *et al.* 1992). Gotts (1996) found that every connected and regular topological space is a model for the RCC axioms. Stell and Worboys (1997) identified a whole class of models base on Heyting structures. Both approaches only describe models for the RCC axioms, i.e., what kind of regions can be used at all. When additional constraints expressing relationships between regions are added, these results do not say anything about models anymore. [10] They are also by no means constructive, as they do not provide a way of effectively representing regions or generating realizations.

Previous approaches on dealing with dimension and internal connectedness of regions tried to specify predicates and suitable axioms in order to restrict dimension and connectedness of regions (Bennett 1996; Gotts 1994). As all regions must have the same dimension anyway, using our results it is not necessary to specify the dimension of regions explicitly if internally disconnected regions are permitted. If internally connected regions are required, these predicates only have an influence on the consistency of a set of spatial relations in one- or two-dimensional applications. In three- and higher-dimensional space all regions may be either internally connected or disconnected. Forcing internal connectedness of regions in two-dimensional space leads to difficult computational problems as there are no complete algorithms for dealing with this task. As Grigni et al. (1995) pointed out,

[9] This is because of the constraint $\Box(\mathsf{X} \to \neg\mathbf{I}\neg\mathbf{I}\mathsf{X})$ which is transformed to $\bigwedge_{w\in W_0}(\neg\mathsf{X}_w \vee \mathsf{X}^1_w \vee \mathsf{X}^2_w)$ in the notation of (Renz and Nebel 1997).

[10] Consider the hypothetical case of a set of spatial formulas which is realizable in two- but not in one-dimensional space. Then, a one-dimensional space is still a model of the RCC axioms, but not of RCC-8.

a well-known open problem in graph theory which is NP-hard but not known to be in NP (Kratochvíl 1991; Kratochvíl and Matoušek 1991) can be reduced to the consistency problem for two-dimensional internally connected regions.

It is certainly the better approach to have an additional connectedness predicate than forcing all regions to be internally connected which is done, e.g., by the similar calculus of Egenhofer (1991), as there are many applications where regions are in fact disconnected. Within the area of geographical information systems, e.g., which offer a great variety of possible applications, many countries or other geographical entities are not internally connected regions. In areas like computer vision it is often dealt with two-dimensional projections of the three-dimensional space where many connected objects are perceived as disconnected objects due to occlusion. In robot navigation, maps are often two-dimensional cuttings of a three-dimensional space.

With the result on realizations in one-dimensional space it becomes possible now to use RCC-8 for temporal reasoning tasks, in particular when non-convex intervals are allowed and the direction of the time is not important, as, e.g., in some scheduling problems. This is in contrast to previous approaches that used temporal calculi for spatial reasoning (Guesgen 1989).

8 SUMMARY

We identified a canonical model of RCC-8 based on Kripke semantics. In order to obtain a "topological" canonical model, we gave a topological interpretation of the Kripke models such that regions can be represented by points in the topological space and information about the neighborhood of these points with respect to the spatial relations holding between the regions. Using this canonical model, we proved that every consistent set of spatial formulas has a realization in any dimension when regions are not forced to be internally connected, which is the case for regions as used by RCC-8. Furthermore, we proved that for three- and higher dimensional space there is always a realization with internally connected regions. Further, we give for the first time algorithms for generating realizations of either internally connected or disconnected regions. Future work includes analyzing the usage of the canonical model for dealing with the special case of two-dimensional internally connected regions as well as analyzing the cognitive meaning of the canonical model (Knauff *et al.* 1997).

Acknowledgments

I would like to thank Bernhard Nebel and Tony Cohn for the fruitful discussions on the topic of the paper. This research has been supported by DFG as part of the project FAST-QUAL-SPACE, which is part of the DFG special research effort on "Spatial Cognition".

References

B. Bennett (1995). Modal logics for qualitative spatial reasoning. *Bulletin of the IGPL*, 4(1).

B. Bennett (1996). Carving up space: steps towards construction of an absolutely complete theory of spatial regions. In *Proc. JELIA'96*, LNCS 1126.

B.F. Chellas (1980). *Modal Logic: An Introduction.* Cambridge University Press, Cambridge, UK.

A.G. Cohn (1997). Qualitative spatial representation and reasoning techniques. In *Proc. KI-97*, LNCS 1303.

T.H. Cormen, C.E. Leiserson, and R.L. Rivest (1990). *Introduction to Algorithms.* MIT Press.

M.J. Egenhofer (1991). Reasoning about binary topological relations. In *Proc. SSD'91*, LNCS 525.

M.C. Fitting (1993). Basic modal logic. In *Handbook of Logic in Artificial Intelligence and Logic Programming – Vol. 1*, Oxford, Clarendon Press.

N.M. Gotts (1994). How far can we C? defining a 'doughnut' using connection alone. In *Proc. KR94*.

N.M. Gotts (1996). An axiomatic approach to topology for spatial information systems. Technical Report 96-25, Univ. of Leeds, School of Computer Studies.

M. Grigni, D. Papadias, and C. Papadimitriou (1995). Topological inference. In *Proc. IJCAI'95*.

H. Guesgen (1989). Spatial reasoning based on Allen's temporal logic. Technical Report TR-89-049, ICSI, Berkeley.

M. Knauff, R. Rauh, and J. Renz (1997). A cognitive assessment of topological spatial relations: Results from an empirical investigation. In *Proc. COSIT'97*, LNCS 1329.

J. Kratochvíl and J. Matoušek (1991). String graphs requiring exponential representations. *Journal of Combinatorial Theory, Series B*, 53:1–4.

J. Kratochvíl (1991). String graphs II: recognizing string graphs is NP-hard. *Journal of Combinatorial Theory, Series B*, 52:67–78.

O. Lemon (1996). Semantical foundations of spatial logics. In *Proc. KR'96*.

B. Nebel (1995). Computational properties of qualitative spatial reasoning: First results. In *Proc. KI-95*, LNCS 981.

D.A. Randell, Z. Cui, and A.G. Cohn (1992). A spatial logic based on regions and connection. In *Proc. KR'92*.

J. Renz and B. Nebel (1997). On the complexity of qualitative spatial reasoning: A maximal tractable fragment of the Region Connection Calculus. In *Proc. IJCAI'97*. Extended version under review for journal publication.

J. Stell and M. Worboys (1997). The algebraic structure of sets of regions. In *Proc. COSIT'97*, LNCS 1329.