



HAL
open science

1-Safe Petri nets and special cube complexes: equivalence and applications

Jérémie Chalopin, Victor Chepoi

► **To cite this version:**

Jérémie Chalopin, Victor Chepoi. 1-Safe Petri nets and special cube complexes: equivalence and applications. 2018. hal-01863455v1

HAL Id: hal-01863455

<https://hal.science/hal-01863455v1>

Preprint submitted on 28 Aug 2018 (v1), last revised 25 Apr 2019 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

1-SAFE PETRI NETS AND SPECIAL CUBE COMPLEXES: EQUIVALENCE AND APPLICATIONS

JÉRÉMIE CHALOPIN AND VICTOR CHEPOI

ABSTRACT. Nielsen, Plotkin, and Winskel (1981) proved that every 1-safe Petri net N unfolds into an event structure \mathcal{E}_N . By a result of Thiagarajan (1996 and 2002), these unfoldings are exactly the trace regular event structures. Thiagarajan (1996 and 2002) conjectured that regular event structures correspond exactly to trace regular event structures. In a recent paper (Chalopin and Chepoi, 2017, 2018), we disproved this conjecture, based on the striking bijection between domains of event structures, median graphs, and CAT(0) cube complexes. On the other hand, in Chalopin and Chepoi (2018) we proved that Thiagarajan’s conjecture is true for regular event structures whose domains are principal filters of universal covers of finite special cube complexes.

In the current paper, we prove the converse: to any finite 1-safe Petri net N one can associate a finite special cube complex \mathbf{X}_N such that the domain of the event structure \mathcal{E}_N (obtained as the unfolding of N) is a principal filter of the universal cover $\tilde{\mathbf{X}}_N$ of \mathbf{X}_N . This establishes a bijection between 1-safe Petri nets and finite special cube complexes and provides a combinatorial characterization of trace regular event structures.

Using this bijection and techniques from graph theory and geometry (MSO theory of graphs, bounded treewidth, and bounded hyperbolicity) we disprove yet another conjecture by Thiagarajan (from the paper with S. Yang from 2014) that the monadic second order logic of a 1-safe Petri net (i.e., of its event structure unfolding) is decidable if and only if its unfolding is grid-free. It was proven by Thiagarajan and Yang, 2014 that the MSO logic is undecidable if the unfolding is not grid-free.

Our counterexample is the trace regular event structure which arises from a virtually special square complex \mathbf{Z} with one vertex, four edges, and three squares. The domain of this event structure \mathcal{E}_Z is the principal filter of the universal cover $\tilde{\mathbf{Z}}$ of \mathbf{Z} in which to each vertex we added a pendant edge. The graph of the domain of \mathcal{E}_Z has bounded hyperbolicity (and thus the associated event structure \mathcal{E}_Z is grid-free) but has infinite treewidth. The MSO theory of the event structure \mathcal{E}_Z is undecidable because adding pendant edges and using results of Seese and Courcelle, we reduce the decidability of the MSO theory of the graph of the domain of \mathcal{E}_Z to the MSO theory of the event structure \mathcal{E}_Z . However, the MSO theory of the graph of a domain is decidable if and only if this graph has bounded treewidth.

CONTENTS

1. Introduction	2
2. On Thiagarajan’s conjectures	3
2.1. The nice labeling conjecture	3
2.2. The conjecture on regular event structures	4
2.3. The conjecture on decidability of MSO logic of trace regular event structures	5
3. Event structures and net systems	6
3.1. Event structures and their domains	6
3.2. Mazurkiewicz traces	6
3.3. Regular trace event structures	7
3.4. Net systems and their event structure unfoldings	7
3.5. The MSO theory of trace event structures	9
4. Domains, median graphs, and CAT(0) cube complexes	10
4.1. Median graphs	10
4.2. Nonpositively curved cube complexes	11
4.3. Domains versus median graphs/CAT(0) cube complexes	13
4.4. Special cube complexes	14

5.	Geodesic traces and prime traces	15
5.1.	Geodesic traces	15
5.2.	Prime geodesic traces	17
6.	Directed NPC complexes	18
6.1.	Directed median graphs	18
6.2.	Directed NPC cube complexes	19
7.	Directed special cube complexes	19
7.1.	The results	19
7.2.	Proof of Theorem 7.5	21
7.3.	Proof of Proposition 7.6	21
8.	1-Safe Petri nets and special cube complexes	22
8.1.	The results	22
8.2.	Proof of Theorem 8.2	24
8.3.	Proof of Theorem 8.4 and Proposition 8.5	25
9.	Decidability of the MSO theory of net systems and of their domains	27
9.1.	The results	27
9.2.	Treewidth	28
9.3.	Hyperbolicity	29
9.4.	Context-free graphs	29
9.5.	Some results from MSO theory	30
9.6.	Grids	31
9.7.	Proof of Theorem 9.1	33
9.8.	Proof of Proposition 9.2	36
9.9.	The MSO theory of hairings of event structures	36
10.	Counterexamples to Conjecture 3.4	39
10.1.	Proof of Theorem 10.1	40
10.2.	Another counterexample to Conjecture 3.4	46
11.	Conclusion	48
	Note	49
	Acknowledgements	49
	References	49

1. INTRODUCTION

Finite 1-safe Petri nets, also called net systems, are natural models of asynchronous concurrency. Nielsen, Plotkin, and Winskel [36] proved that every net system $N = (S, \Sigma, F, m_0)$ unfolds into an event structure $\mathcal{E}_N = (E, \leq, \#, \lambda)$ describing all possible executions of N : the events of \mathcal{E}_N are all prime Mazurkiewicz traces on the set of transitions of N , equipped with the causal dependency and conflict relations. Later results of Nielsen, Rozenberg, and Thiagarajan [37] show in fact that 1-safe Petri nets and event structures represent each other in a strong sense. An event structure [36, 51, 52] is a partially ordered set of the occurrences of actions, called events, together with a conflict relation. The partial order captures the causal dependency of events. The conflict relation models incompatibility of events so that two events that are in conflict cannot simultaneously occur in any state of the computation. Consequently, two events that are neither ordered nor in conflict may occur concurrently. The domain of an event structure consists of all computation states, called configurations. Each computation state is a subset of events subject to the constraints that no two conflicting events can occur together in the same computation and if an event occurred in a computation then all events on which it causally depends have occurred too. Therefore, the domain of an event structure is the set of all finite configurations ordered by inclusion. The future (or the principal filter) of a configuration is the set of all finite configurations containing it.

In a series of papers [42, 46–48], Thiagarajan formulated (alone or with co-authors) three important conjectures (1) about the local-to-global behavior of event structures (the *nice labeling conjecture*): *any event structure of finite degree admits a finite nice labeling*, (2) on the relationship between event structures and net systems: *regular event structures are exactly the unfoldings of net systems* and (3) about the decidability of the Monadic Second Order theory (MSO theory) of net systems: *grid-free net systems are exactly the net systems with decidable MSO theory*. The last two conjectures were motivated by the fact that in each case, one of the two implications holds and by evidences and important particular cases for which the converse implication also holds. For example, it was proven in [46, 47] that unfoldings of net systems are exactly the trace regular event structures, and thus the second conjecture asks whether a regular event structure is trace regular.

In the previous papers [17] and [13, 14] we provided counterexamples to the first two conjectures. In the current paper, we will provide a counterexample to the third conjecture about the decidability of the MSO theory of grid-free net systems. The three counterexamples are based on different ideas and techniques, however, they all use the bijections between domains of event structures, median graphs, and $\text{CAT}(0)$ cube complexes. Median graphs is the most important class of graphs in metric graph theory and $\text{CAT}(0)$ cube complexes play an essential role in geometric group theory and the topology of 3-manifolds. Even if the three conjectures turned out to be false, the work on them raised many important open questions and the current paper establishes a surprising bijection between 1-safe Petri nets (trace regular event structures) and finite special cube complexes. Notice that special cube complexes, introduced by Haglund and Wise [27, 28], played an essential role in the recent solution of the famous virtual Haken conjecture for hyperbolic 3-manifolds by Agol [1, 2].

2. ON THIAGARAJAN’S CONJECTURES

We continue with an informal description of Thiagarajan’s conjectures, of some related work on them, and of the results of this paper.

2.1. The nice labeling conjecture. The *nice labeling conjecture* was formulated by Rozoy and Thiagarajan in [42] and asserts that

Conjecture 1. *Every event structure with finite degree admits a nice labeling with a finite number of labels.*

A nice labeling is a labeling of events with the letters from some finite alphabet such that any two co-initial events (i.e., any two events which are concurrent or in minimal conflict) have different labels. The nice labelings of event structures arise when studying the equivalence of three different models of distributed computation: labeled event structures, net systems, and distributed monoids. The nice labeling conjecture can be viewed as a question about a local-to-global finite behavior of such models.

A counterexample to nice labeling conjecture was constructed in [17]. It is based on the bijection between domains of event structures, pointed median graphs, and $\text{CAT}(0)$ cube complexes and on the Burling’s construction [12] of 3-dimensional box hypergraphs with clique number 2 and arbitrarily large chromatic numbers. Assous, Bouchitté, Charretton, and Rozoy [4] proved that the event structures of degree 2 admit nice labelings with 2 labels and noticed that Dilworth’s theorem implies that the conflict-free event structures of degree n have nice labelings with n labels. They also showed that finding the least number of labels in a nice labeling of a finite event structure is NP-hard. Santocanale [44] proved that all event structures of degree 3 with tree-like partial orders have nice labelings with 3 labels. Chepoi and Hagen [19] proved that the nice labeling conjecture holds for event structures with 2-dimensional domains, i.e., for event structures not containing three pairwise concurrent events.

For $\text{CAT}(0)$ cube complexes a question related to the nice labeling conjecture was independently formulated by F. Haglund, G. Niblo, M. Sageev, and the second author of this paper: *is it true that the 1-skeleton of any $\text{CAT}(0)$ cube complex of finite degree can be isometrically embedded into the Cartesian product of a finite number of trees?* A negative answer to this question

was obtained in [19], based on a modification of the counterexample from [17]. However, in [19] it was shown that the answer is positive for 2-dimensional CAT(0) cube complexes. Haglund [26] proved that this embedding question has a positive answer for hyperbolic CAT(0) cube complexes. Modifying the argument of [26], we can also show that the nice labeling conjecture is true for event structures with hyperbolic domains.

2.2. The conjecture on regular event structures. To deal with net systems, Thiagarajan [46,47] introduced the notions of regular event structure and trace regular event structure. The main difference is that the regularity of event structures is defined for unlabeled event structures while trace regularity is defined under the much stronger assumption of a given trace regular labeling. These definitions were motivated by the fact that the event structures \mathcal{E}_N arising from *finite* 1-safe Petri nets N are regular: Thiagarajan [46] in fact proved that event structures of *finite* 1-safe Petri nets correspond to regular trace event structures. This lead Thiagarajan to conjecture (see also Conjecture 3.3 below) in [46,47] that

Conjecture 2. *Regular event structures and regular trace event structures are the same.*

Equivalently, this can be reformulated in the following way: *an event structure \mathcal{E} is isomorphic to the event structure unfolding of a net system if and only if \mathcal{E} is regular.*

Nielsen and Thiagarajan [38] established this conjecture for all regular conflict-free event structures and Badouel, Darondeau, and Raoult [6] proved it for context-free event domains, i.e., for domains whose underlying graph is a context-free graph sensu Müller and Schupp [35]. Morin [33] showed that any event structure admitting a regular nice labeling is trace regular.

In a recent paper [13], we presented a counterexample to Thiagarajan's Conjecture 2 based on a geometric and combinatorial view on event structures. To deal with regular event structures, we showed in [13] how to construct regular event domains from CAT(0) cube complexes. By a result of Gromov [24], CAT(0) cube complexes are exactly the universal covers of nonpositively curved cube (NPC) complexes. Of particular importance for us are the CAT(0) cube complexes arising as universal covers of *finite* NPC complexes. We adapted the universal cover construction to directed NPC complexes (Y, o) and showed that every principal filter of the directed universal cover (\tilde{Y}, \tilde{o}) is the domain of an event structure. Furthermore, if the NPC complex Y is finite, then this event structure is regular. Motivated by this result, we called an event structure *strongly regular* if its domain is the principal filter of the directed universal cover $\tilde{\mathbf{Y}} = (\tilde{Y}, \tilde{o})$ of a finite directed NPC complex $\mathbf{Y} = (Y, o)$. Our counterexample to this Thiagarajan's conjecture is a strongly regular event structure not admitting a finite regular nice labeling. It is derived from Wise's [53,54] nonpositively curved square complex \mathbf{X} obtained from a tile set with six tiles.

In view of this counterexample, one can ask the following two important questions:

Question 2.1. Are the event structures arising as unfoldings of finite 1-safe Petri nets strongly regular?

Question 2.2. Under which conditions a regular event structure is trace regular?

Haglund and Wise [27,28] detected five types of pathologies which may occur in NPC complexes. They called the NPC complexes without such pathologies *special*. The main motivation for introducing and studying special cube complexes was the profound idea of Wise that the famous virtual Haken conjecture for hyperbolic 3-manifolds can be reduced to solving problems about special cube complexes. In a breakthrough result, Agol [1,2] completed this program and solved the virtual Haken conjecture using the deep theory of special cube complexes developed by Haglund and Wise [27,28]. The main ingredient in this proof is Agol's theorem that finite NPC complexes whose universal covers are hyperbolic are virtually special (i.e., they admit finite covers which are special cube complexes).

In [13] we proved that Thiagarajan's conjecture is true for event structures whose domains arise as principal filters of universal covers of finite special cube complexes. Using the result of Agol, we specified this result and showed that Thiagarajan's conjecture is true for strongly regular event structures whose domains occur as principal filters of hyperbolic CAT(0) cube

complexes that are universal covers of finite directed NPC complexes. Since context-free domains are hyperbolic, this result can be viewed in some sense as a partial generalization of the result of Badouel et al. [6].

In the current paper, we establish the converse to the previous result of [13]: we prove that to any 1-safe Petri net N one can associate a finite directed labeled special cube complex \mathbf{X}_N such that the domain of the event structure \mathcal{E}_N (obtained as the unfolding of N) is a principal filter of the universal cover $\tilde{\mathbf{X}}_N$ of \mathbf{X}_N . This proves that *the trace regular event structures are exactly the special strongly regular event structures* and that *the trace labeling is obtained via the covering map*. This shows that all event structures arising as unfoldings of finite 1-safe Petri nets are strongly regular, answering in the positive Question 2.1. This also shows that specialness must be added to strong regularity to ensure a positive answer to Thiagarajan's Conjecture 2. Therefore, the trace regular event structures can be characterized as the event structures whose domains arise from *finite special* cube complexes. This establishes a surprising bijection between 1-safe Petri nets (fundamental objects in concurrency) and special cube complexes (fundamental objects in geometric group theory).

2.3. The conjecture on decidability of MSO logic of trace regular event structures.

Thiagarajan and Yang [48] defined the monadic second order (MSO) theory $\text{MSO}(\mathcal{E}_N)$ of an event structure unfolding $\mathcal{E}_N = (E, \leq, \#, \lambda)$ of a net system $N = (S, \Sigma, F, m_0)$ as the MSO theory of the relational structure $(E, (R_a)_{a \in \Sigma}, \leq)$ (see Subsection 3.5 for definitions). This immediately leads to the following fundamental question:

Question 2.3. When $\text{MSO}(\mathcal{E}_N)$ is decidable?

It turns out that the MSO theory of trace event structures is not always decidable: [48] presented such an example suggested by I. Walukiewicz. To circumvent this example, Thiagarajan and Yang formulated the notion of a grid event structure and they showed that the MSO theory of event structures containing grids is undecidable. This leads Thiagarajan to conjecture in [48] that (see also Conjecture 3.4 below):

Conjecture 3. *The MSO theory of a trace regular event structure \mathcal{E}_N is decidable if and only if \mathcal{E}_N is grid-free.*

Notice also that preceding [48], Madhusudan [32] proved that the MSO theory of a trace event structure is decidable provided quantifications over sets are restricted to conflict-free subsets of events. In particular, this shows that the MSO theory of conflict-free trace regular event structures is decidable.

With the event structure \mathcal{E}_N one can associate two other MSO logics: the MSO logic $\text{MSO}(\vec{G}(\mathcal{E}_N))$ of the directed graph $\vec{G}(\mathcal{E}_N)$ of the domain $\mathcal{D}(\mathcal{E}_N)$ of \mathcal{E}_N and the MSO logic $\text{MSO}(G(\mathcal{E}_N))$ of the undirected graph $G(\mathcal{E}_N)$ of the domain. This leads to the next question:

Question 2.4. When $\text{MSO}(\vec{G}(\mathcal{E}_N))$ (respectively, $\text{MSO}(G(\mathcal{E}_N))$) is decidable?

As we will prove in this paper, the decidability of each of $\text{MSO}(G(\mathcal{E}_N))$ and $\text{MSO}(\vec{G}(\mathcal{E}_N))$ is equivalent to the fact that $G(\mathcal{E}_N)$ has finite treewidth and to the fact that $\vec{G}(\mathcal{E}_N)$ is a context-free graph. This completely answers Question 2.4. We also prove that if $\text{MSO}(\vec{G}(\mathcal{E}_N))$ is decidable, then $\text{MSO}(\mathcal{E}_N)$ is decidable (the converse is not true). We introduce the notion of hairing of an event structure \mathcal{E}_N , which is an event structure $\dot{\mathcal{E}}_N$ obtained from \mathcal{E}_N by adding an event e_c for each configuration c of the domain in a such a way that e_c is in conflict with all events except those from c (those events precede e_c). We prove that $\text{MSO}(\dot{\mathcal{E}}_N)$ is decidable if and only if $\text{MSO}(\vec{G}(\mathcal{E}_N))$ is decidable, i.e., if and only if $G(\mathcal{E})$ has finite treewidth. All these results provide partial answers to Question 2.3.

Using these results, we construct a counterexample to Thiagarajan's Conjecture 3. Namely, we construct an NPC square complex \mathbf{Z} with one vertex, four edges, and three squares. We show that \mathbf{Z} is virtually special and thus any principal filter of the universal cover of \mathbf{Z} is the domain of a trace regular event structure \mathcal{E}_Z . The hairing $\dot{\mathcal{E}}_Z$ of \mathcal{E}_Z is still trace regular. We show that

the graphs $G(\mathcal{E}_Z)$ and $G(\dot{\mathcal{E}}_Z)$ have infinite treewidth and bounded hyperbolicity. The first result implies that $\text{MSO}(\dot{\mathcal{E}}_Z)$ is undecidable while the second result shows that $\dot{\mathcal{E}}_Z$ is grid-free.

3. EVENT STRUCTURES AND NET SYSTEMS

3.1. Event structures and their domains. An *event structure* is a triple $\mathcal{E} = (E, \leq, \#)$, where

- E is a set of *events*,
- $\leq \subseteq E \times E$ is a partial order of *causal dependency*,
- $\# \subseteq E \times E$ is a binary, irreflexive, symmetric relation of *conflict*,
- $\downarrow e := \{e' \in E : e' \leq e\}$ is finite for any $e \in E$,
- $e \# e'$ and $e' \leq e''$ imply $e \# e''$.

Two events e', e'' are *concurrent* (notation $e' \parallel e''$) if they are order-incomparable and they are not in conflict. The conflict $e' \# e''$ between two elements e' and e'' is said to be *minimal* (notation $e' \#_{\mu} e''$) if there is no event $e \neq e', e''$ such that either $e \leq e'$ and $e \# e''$ or $e \leq e''$ and $e \# e'$. We say that e is an *immediate predecessor* of e' (notation $e < e'$) if and only if $e \leq e', e \neq e'$, and for every e'' if $e \leq e'' \leq e'$, then $e'' = e$ or $e'' = e'$.

Given two event structures $\mathcal{E}_1 = (E_1, \leq_1, \#_1)$ and $\mathcal{E}_2 = (E_2, \leq_2, \#_2)$, a map $f : E_1 \rightarrow E_2$ is an isomorphism if f is a bijection such that $e \leq_1 e'$ iff $f(e) \leq_2 f(e')$ and $e \#_1 e'$ iff $f(e) \#_2 f(e')$ for every $e, e' \in E_1$. If such an isomorphism exists, then \mathcal{E}_1 and \mathcal{E}_2 are said to be isomorphic; notation $\mathcal{E}_1 \equiv \mathcal{E}_2$.

A *configuration* of an event structure $\mathcal{E} = (E, \leq, \#)$ is any finite subset $c \subseteq E$ of events which is *conflict-free* ($e, e' \in c$ implies that e, e' are not in conflict) and *downward-closed* ($e \in c$ and $e' \leq e$ implies that $e' \in c$) [52]. Notice that \emptyset is always a configuration and that $\downarrow e$ and $\downarrow e \setminus \{e\}$ are configurations for any $e \in E$. The *domain* of an event structure is the set $\mathcal{D} := \mathcal{D}(\mathcal{E})$ of all configurations of \mathcal{E} ordered by inclusion; (c', c) is a (directed) edge of the Hasse diagram of the poset $(\mathcal{D}(\mathcal{E}), \subseteq)$ if and only if $c = c' \cup \{e\}$ for an event $e \in E \setminus c$. An event e is said to be *enabled* by a configuration c if $e \notin c$ and $c \cup \{e\}$ is a configuration. Denote by $en(c)$ the set of all events enabled at the configuration c . Two events are called *co-initial* if they are both enabled at some configuration c . Note that if e and e' are co-initial, then either $e \#_{\mu} e'$ or $e \parallel e'$. It is easy to see that two events e and e' are in minimal conflict $e \#_{\mu} e'$ if and only if $e \# e'$ and e and e' are co-initial. The *degree* $\text{deg}(\mathcal{E})$ of an event structure \mathcal{E} is the least positive integer d such that $|en(c)| \leq d$ for any configuration c of \mathcal{E} . We say that \mathcal{E} has *finite degree* if $\text{deg}(\mathcal{E})$ is finite. The *future* (or the *principal filter*) $\mathcal{F}(c)$ of a configuration c is the set of all configurations c' containing c : $\mathcal{F}(c) = \uparrow c := \{c' \in \mathcal{D}(\mathcal{E}) : c \subseteq c'\}$, i.e., $\mathcal{F}(c)$ is the principal filter of c in the ordered set $(\mathcal{D}(\mathcal{E}), \subseteq)$.

A *labeled event structure* $\mathcal{E}^{\lambda} = (\mathcal{E}, \lambda)$ is defined by an *underlying event structure* $\mathcal{E} = (E, \leq, \#)$ and a *labeling* λ that is a map from E to some alphabet Σ . Two labeled event structures $\mathcal{E}_1^{\lambda_1} = (\mathcal{E}_1, \lambda_1)$ and $\mathcal{E}_2^{\lambda_2} = (\mathcal{E}_2, \lambda_2)$ are isomorphic (notation $\mathcal{E}_1^{\lambda_1} \equiv \mathcal{E}_2^{\lambda_2}$) if there exists an isomorphism f between the underlying event structures \mathcal{E}_1 and \mathcal{E}_2 such that $\lambda_2(f(e_1)) = \lambda_1(e_1)$ for every $e_1 \in E_1$.

A labeling $\lambda : E \rightarrow \Sigma$ of an event structure \mathcal{E} defines naturally a labeling of the directed edges of the Hasse diagram of its domain $\mathcal{D}(\mathcal{E})$ that we also denote by λ . A labeling $\lambda : E \rightarrow \Sigma$ of an event structure \mathcal{E} is called a *nice labeling* if any two events that are co-initial have different labels [42]. A nice labeling of \mathcal{E} can be reformulated as a labeling of the directed edges of the Hasse diagram of its domain $\mathcal{D}(\mathcal{E})$ subject to the following local conditions:

Determinism: the edges outgoing from the same vertex of $\mathcal{D}(\mathcal{E})$ have different labels;

Concurrency: the opposite edges of each square of $\mathcal{D}(\mathcal{E})$ are labeled with the same labels.

In the following, we use interchangeably the labeling of an event structure and the labeling of the edges of its domain.

3.2. Mazurkiewicz traces. A (*Mazurkiewicz*) *trace alphabet* is a pair $M = (\Sigma, I)$, where Σ is a finite non-empty alphabet set and $I \subset \Sigma \times \Sigma$ is an irreflexive and symmetric relation called the *independence relation*. The relation $D := (\Sigma \times \Sigma) \setminus I$ is called the *dependence relation*. As

usual, Σ^* is the set of finite words with letters in Σ . For $\sigma \in \Sigma^*$, $last(\sigma)$ denotes the last letter of σ . The independence relation I induces the equivalence relation \sim_I , which is the reflexive and transitive closure of the binary relation \leftrightarrow_I : if $\sigma, \sigma' \in \Sigma^*$ and $(a, b) \in I$, then $\sigma ab\sigma' \leftrightarrow_I \sigma ba\sigma'$. The \sim_I -equivalence class containing $\sigma \in \Sigma^*$ is called a (*Mazurkiewicz*) *trace* and will be denoted by $\langle \sigma \rangle$. The trace $\langle \sigma \rangle$ is *prime* if σ is non-null and for every $\sigma' \in \langle \sigma \rangle$, $last(\sigma) = last(\sigma')$. The partial ordering relation \sqsubseteq between traces is defined by $\langle \sigma \rangle \sqsubseteq \langle \tau \rangle$ (and $\langle \sigma \rangle$ is said to be a *prefix* of $\langle \tau \rangle$) if there exist $\sigma' \in \langle \sigma \rangle$ and $\tau' \in \langle \tau \rangle$ such that σ' is a prefix of τ' .

3.3. Regular trace event structures. In this subsection, we recall the definitions of regular event structures, regular trace event structures, and regular nice labelings of event structures. We closely follow the definitions and notations of [38, 46, 47]. Let $\mathcal{E} = (E, \leq, \#)$ be an event structure. Let c be a configuration of \mathcal{E} . Set $\#(c) = \{e' : \exists e \in c, e\#e'\}$. The *event structure rooted at c* is defined to be the triple $\mathcal{E} \setminus c = (E', \leq', \#')$, where $E' = E \setminus (c \cup \#(c))$, \leq' is \leq restricted to $E' \times E'$, and $\#'$ is $\#$ restricted to $E' \times E'$. It can be easily seen that the domain $\mathcal{D}(\mathcal{E} \setminus c)$ of the event structure $\mathcal{E} \setminus c$ is isomorphic to the principal filter $\mathcal{F}(c)$ of c in $\mathcal{D}(\mathcal{E})$ such that any configuration c' of $\mathcal{D}(\mathcal{E})$ corresponds to the configuration $c' \setminus c$ of $\mathcal{D}(\mathcal{E} \setminus c)$.

For an event structure $\mathcal{E} = (E, \leq, \#)$, define the equivalence relation $R_{\mathcal{E}}$ on its configurations in the following way: for two configurations c and c' set $cR_{\mathcal{E}}c'$ if and only if $\mathcal{E} \setminus c \equiv \mathcal{E} \setminus c'$. The *index* of an event structure \mathcal{E} is the number of equivalence classes of $R_{\mathcal{E}}$, i.e., the number of isomorphism types of futures of configurations of \mathcal{E} . The event structure \mathcal{E} is *regular* [38, 46, 47] if \mathcal{E} has finite index and finite degree.

Now, let $\mathcal{E}^\lambda = (\mathcal{E}, \lambda)$ be a labeled event structure. For any configuration c of \mathcal{E} , if we restrict λ to $\mathcal{E} \setminus c$, then we obtain a labeled event structure $(\mathcal{E} \setminus c, \lambda)$ denoted by $\mathcal{E}^\lambda \setminus c$. Analogously, define the equivalence relation $R_{\mathcal{E}^\lambda}$ on its configurations by setting $cR_{\mathcal{E}^\lambda}c'$ if and only if $\mathcal{E}^\lambda \setminus c \equiv \mathcal{E}^\lambda \setminus c'$. The index of \mathcal{E}^λ is the number of equivalence classes of $R_{\mathcal{E}^\lambda}$. We say that an event structure \mathcal{E} admits a *regular nice labeling* if there exists a nice labeling λ of \mathcal{E} with a finite alphabet Σ such that \mathcal{E}^λ has finite index.

We continue by recalling the definition of regular trace event structures from [46, 47]. For a trace alphabet $M = (\Sigma, I)$, an *M -labeled event structure* is a labeled event structure $\mathcal{E}^\phi = (\mathcal{E}, \lambda)$, where $\mathcal{E} = (E, \leq, \#)$ is an event structure and $\lambda : E \rightarrow \Sigma$ is a labeling function which satisfies the following conditions:

- (LES1) $e\#_\mu e'$ implies $\lambda(e) \neq \lambda(e')$,
- (LES2) if $e < e'$ or $e\#_\mu e'$, then $(\lambda(e), \lambda(e')) \in D$,
- (LES3) if $(\lambda(e), \lambda(e')) \in D$, then $e \leq e'$ or $e' \leq e$ or $e\#e'$.

We call λ a *trace labeling* of \mathcal{E} with the trace alphabet $M = (\Sigma, I)$. The conditions (LES2) and (LES3) on the labeling function ensures that the concurrency relation \parallel of \mathcal{E} respects the independence relation I of M . In particular, since I is irreflexive, from (LES3) it follows that any two concurrent events are labeled differently. Since by (LES1) two events in minimal conflict are also labeled differently, this implies that λ is a finite nice labeling of \mathcal{E} .

An M -labeled event structure $\mathcal{E}^\lambda = (\mathcal{E}, \lambda)$ is *regular* if \mathcal{E}^λ has finite index. Finally, an event structure \mathcal{E} is called a *regular trace event structure* [46, 47] if there exists a trace alphabet $M = (\Sigma, I)$ and a regular M -labeled event structure \mathcal{E}^λ such that \mathcal{E} is isomorphic to the underlying event structure of \mathcal{E}^λ . From the definition immediately follows that every regular trace event structure is also a regular event structure.

3.4. Net systems and their event structure unfoldings. In the following presentation of finite 1-safe Petri nets and their unfoldings to event structures, we closely follow the paper by Thiagarajan and Yang [48]. A *net system* (or, equivalently, a *finite 1-safe Petri net*) is a quadruplet $N = (S, \Sigma, F, m_0)$ where S and Σ are disjoint finite sets of *places* and *transitions* (called also *actions* or *events*), $F \subseteq (S \times \Sigma) \cup (\Sigma \times S)$ is the *flow relation*, and $m_0 \subseteq S$ is the *initial marking*. For $v \in S \cup \Sigma$, set $\bullet v = \{u : (u, v) \in F\}$ and $v \bullet = \{u : (v, u) \in F\}$. A *marking* of N is a subset of S . The *transition relation* (or the *firing rule*) $\longrightarrow \subseteq 2^S \times \Sigma \times 2^S$ is defined by $m \xrightarrow{a} m'$ if $\bullet a \subseteq m$, $(a \bullet - \bullet a) \cap m = \emptyset$, and $m' = (m - \bullet a) \cup a \bullet$. The transition relation \longrightarrow is extended to sequences of transitions as follows (this new relation is also denoted by \longrightarrow):




FIGURE 1. The net system N^* has 12 transitions (represented by rectangles) and 10 places (represented by circles).

(1) $m \xrightarrow{\varepsilon} m$ for any marking m and (2) if $m \xrightarrow{\sigma} m'$ for $\sigma \in \Sigma^*$ and $m' \xrightarrow{a} m''$ for $a \in \Sigma$, then $m \xrightarrow{\sigma a} m''$. $\sigma \in \Sigma^*$ is called a *firing sequence* at m if there exists a marking m' such that $m \xrightarrow{\sigma} m'$. Denote by FS the set of firing sequences at m_0 . A marking m is *reachable* if there exists a firing sequence σ such that $m_0 \xrightarrow{\sigma} m$.

Given a net system $N = (S, \Sigma, F, m_0)$, there is a canonical way to associate a Σ -labeled event structure \mathcal{E}_N with N . The trace alphabet associated with N is the pair (Σ, I) , where $(a, b) \in I$ iff $(a^\bullet \cup \bullet a) \cap (b^\bullet \cup \bullet b) = \emptyset$. Observe that the trace alphabet (Σ, I) is independent of the initial marking of N . Given the trace alphabet (Σ, I) , we call the traces of the form $\langle \sigma \rangle$, $\sigma \in FS$ *firing traces* of N . Denote by $\mathcal{FT}(N)$ the set of all firing traces of N and by $\mathcal{PFT}(N)$ the subset of $\mathcal{FT}(N)$ consisting of prime firing traces.

Example 3.1. In Figure 1, we present a net system N^* with 12 transitions $h_1, h'_1, h_2, h'_2, h_3, h'_3, h_4, h'_4, v_1, v'_1, v_2, v'_2$ and 10 places $H_1, H_2, H_3, H_4, V_1, V_2, C_1, C_2, C_3, C_4$. The initial marking is given by the places containing tokens in the figure.

The trace alphabet (Σ, I) associated with the net system N^* has 12 letters $h_1, h'_1, h_2, h'_2, h_3, h'_3, h_4, h'_4, v_1, v'_1, v_2, v'_2$. The letter v_1 is dependent from the letters v'_1, v_2, v'_2 (because of the place V_1 and/or V_2), h'_2 , and h_4 (because of C_1). The letter h_1 is dependent from the letters h'_1, h_4, h'_4 (because of the place H_1), h'_2, h_2 (because of H_2), h'_3 , and v'_2 (because of C_2). For the remaining letters, the dependency relation is defined in a similar way.

Observe that the letters h_1 and h_3 are independent, but there is no firing trace containing h_1 and h_3 as consecutive letters.

Following [36], the *event structure unfolding* of N is the event structure $\mathcal{E}_N = (E, \leq, \#, \lambda)$, where

- E is the set of prime firing traces $\mathcal{PFT}(N)$,
- \leq is \sqsubseteq , restricted to $E \times E$,
- Let $e, e' \in E$. Then $e \# e'$ iff there does not exist a firing trace $\langle \sigma \rangle$ such that $e \sqsubseteq \langle \sigma \rangle$ and $e' \sqsubseteq \langle \sigma \rangle$,
- $\lambda : E \rightarrow \Sigma$ is given by $\lambda(\langle \sigma \rangle) = \text{last}(\sigma)$.

The following result establishes the equivalence between unfoldings of net systems and regular trace event structures:

Theorem 3.2 ([47, Theorem 1]). *An event structure \mathcal{E} is a regular trace event structure if and only if there exists a net system N such that \mathcal{E} and \mathcal{E}_N are isomorphic.*

This lead Thiagarajan to conjecture in [46, 47] that

Conjecture 3.3. An event structure \mathcal{E} is isomorphic to the event structure \mathcal{E}_N arising from a finite 1-safe Petri net N if and only if \mathcal{E} is regular.

3.5. The MSO theory of trace event structures. We start with the definition of monadic second-order logic (MSO-logic). Let A be a universe and $\mathcal{A} = (A, (R_i)_{i \in I})$, where $R_i \subseteq A^{n_i}$ for $i \in I$ be a *relational structure*. The MSO logic of \mathcal{A} has two types of variables: individual (or first-order) variables and set (or second-order) variables. The individual variables range over the elements of A and are denoted by x, y, z , etc. The set variables range over subsets of A and are denoted X, Y, Z , etc. *MSO-formulas* over the signature of \mathcal{A} are constructed from the atomic formulas $R_i(x_1, \dots, x_{n_i})$, $x = y$, and $x \in X$ (where $i \in I$, $x_1, \dots, x_{n_i}, x, y$ are individual variables and X is a set variable) using the Boolean connectives \neg, \vee, \wedge , and quantifications over first order and second order variables. The notions of free variables and bound variables are defined as usual. A formula without free occurrences of variables is called an *MSO-sentence*. If $\varphi(x_1, \dots, x_n, X_1, \dots, X_m)$ is an MSO-formula such that at most the individual variables among x_1, \dots, x_n and the set variables among X_1, \dots, X_m occur freely in φ , and $a_1, \dots, a_n \in A$ and $A_1, \dots, A_m \subseteq A$, then $\mathcal{A} \models \varphi(a_1, \dots, a_n, A_1, \dots, A_m)$ means that φ evaluates to true in \mathcal{A} when x_i evaluates to a_i and X_j evaluates to A_j . The *MSO theory* of \mathcal{A} , denoted by $\text{MSO}(\mathcal{A})$, is the set of all MSO-sentences φ such that $\mathcal{A} \models \varphi$. The MSO theory of \mathcal{A} is said to be *decidable* if there exists an algorithm deciding for each MSO-sentence φ in $\text{MSO}(\mathcal{A})$, whether $\mathcal{A} \models \varphi$.

Let $\mathcal{E}_N = (E, \leq, \#, \lambda)$ be a regular trace event structure, which is the event structure unfolding of a net system $N = (S, \Sigma, F, m_0)$ (by Theorem 3.2, any regular trace event structure admits such a representation). Thiagarajan and Yang [48] defined the MSO theory $\text{MSO}(\mathcal{E}_N)$ of \mathcal{E}_N as the MSO theory of the relational structure $(E, (R_a)_{a \in \Sigma}, \leq)$, where E is the set of events, $R_a \subseteq E$ is the set of a -labeled events for $a \in \Sigma$, and $\leq \subseteq E \times E$ is the precedence relation. The MSO theory of a net system N is the MSO theory of its event structure unfolding [48].

As shown in [48], the conflict relation $\#$, the concurrency relation \parallel , and the notion of a configuration of \mathcal{E} , as well as other connectives of propositional logic such as \wedge, \Rightarrow (implies) and \equiv (if and only if), universal quantification over individual and set variables ($\forall x(\varphi), \forall X(\varphi)$), the set inclusion relation \subseteq ($X \subseteq Y$), can be defined as well. The conflict and concurrency relations $\#$ and \parallel of \mathcal{E} are defined in [48] in the following way:

- $x \hat{\#} y := \neg(x \leq y) \wedge \neg(y \leq x) \wedge \bigvee_{(a,b) \in D} (R_a(x) \wedge R_b(y))$.
- $x \# y := \exists x' \exists y' (x' \leq x \wedge y' \leq y \wedge x' \hat{\#} y')$.
- $x \parallel y := \neg(x \leq y) \wedge \neg(y \leq x) \wedge \neg(x \# y)$.

An interpretation \mathcal{I} assigns to every individual variable an event in E and every set variable, a subset of E . Then \mathcal{E}_N satisfies a formula φ under an interpretation \mathcal{I} , denoted by $\mathcal{E}_N \models_{\mathcal{I}} \varphi$, if the following holds [48]:

- $\mathcal{E}_N \models_{\mathcal{I}} R_a(x)$ iff $\lambda(\mathcal{I}(x)) = a$.
- $\mathcal{E}_N \models_{\mathcal{I}} x \leq y$ iff $\mathcal{I}(x) \leq \mathcal{I}(y)$.
- $\mathcal{E}_N \models_{\mathcal{I}} x \in X$ iff $\mathcal{I}(x) \in \mathcal{I}(X)$.

- $\mathcal{E}_N \models_{\mathcal{I}} \exists x(\varphi)$ iff there exists $e \in E$ and an interpretation \mathcal{I}' such that $\mathcal{E} \models_{\mathcal{I}'} \varphi$ where \mathcal{I}' satisfies the conditions: $\mathcal{I}'(x) = e$, $\mathcal{I}'(y) = \mathcal{I}(y)$ for every individual variable y other than x , and $\mathcal{I}'(X) = \mathcal{I}(X)$ for every set variable X .
- $\mathcal{E}_N \models_{\mathcal{I}} \exists X(\varphi)$ iff there exists $E' \subseteq E$ and an interpretation \mathcal{I}' such that $\mathcal{E} \models_{\mathcal{I}'} \varphi$ where \mathcal{I}' satisfies: $\mathcal{I}'(X) = E'$, $\mathcal{I}'(x) = \mathcal{I}(x)$ for every individual variable x , and $\mathcal{I}'(Y) = \mathcal{I}(Y)$ for every set variable Y other than X .
- $\mathcal{E}_N \models_{\mathcal{I}} \neg\varphi$ and $\mathcal{E} \models_{\mathcal{I}} \varphi_1 \vee \varphi_2$ are defined in the standard way.

$\mathcal{E} \models \varphi$ will denote that \mathcal{E} is a model of the sentence φ .

It turns out that the MSO theory of trace event structures is not always decidable: Fig. 1 of [48] presented an example of such an event structure suggested by Igor Walukiewicz. To circumvent this example, Thiagarajan and Yang formulated the following notion.

The event structure $\mathcal{E} = (E, \leq, \#)$ is *grid-free* [48] if there does not exist three pairwise disjoint sets X, Y, Z of E satisfying the following conditions:

- $X = \{x_0, x_1, x_2, \dots\}$ is an infinite set of events with $x_0 < x_1 < x_2 < \dots$.
- $Y = \{y_0, y_1, y_2, \dots\}$ is an infinite set of events with $y_0 < y_1 < y_2 < \dots$.
- $X \times Y \subseteq \parallel$.
- There exists an injective mapping $g : X \times Y \rightarrow Z$ satisfying: if $g(x_i, y_j) = z$ then $x_i < z$ and $y_j < z$. Furthermore, if $i' > i$ then $x_{i'} \not< z$ and of $j' > j$ then $y_{j'} \not< z$.

The Σ -labelled event structure $(E, \leq, \#, \lambda)$ is said to be *grid-free* if the unlabeled event structure $(E, \leq, \#)$ is grid-free. The net system N is *grid-free* if the event structure \mathcal{E}_N is grid-free. As noticed in [48], Walukiewicz's net system is not grid-free. Thiagarajan and Yang [48] proved that if a net system N is not grid-free, then the MSO theory $\text{MSO}(\mathcal{E}_N)$ is not decidable. Thiagarajan conjectured that the converse holds:

Conjecture 3.4. The MSO theory of a net system N is decidable iff N is grid-free.

4. DOMAINS, MEDIAN GRAPHS, AND $\text{CAT}(0)$ CUBE COMPLEXES

In this section, we recall the bijections between domains of event structures and median graphs/ $\text{CAT}(0)$ cube complexes established in [3] and [9], and between median graphs and 1-skeleta of $\text{CAT}(0)$ cube complexes established in [16] and [41].

4.1. Median graphs. Let $G = (V, E)$ be a simple, connected, not necessarily finite graph. The *distance* $d_G(u, v)$ between two vertices u and v is the length of a shortest (u, v) -path, and the *interval* $I(u, v)$ between u and v consists of all vertices on shortest (u, v) -paths, that is, of all vertices (metrically) *between* u and v :

$$I(u, v) := \{x \in V : d_G(u, x) + d_G(x, v) = d_G(u, v)\}.$$

An induced subgraph H of G (or the corresponding vertex set) is called *convex* if it includes the interval of G between any of its vertices. An induced subgraph H of G is called *gated* if for any vertex $v \in V(G) \setminus V(H)$ there exists a unique vertex $v' \in V(H)$ such that $v' \in I(v, u)$ for any $u \in V(H)$ (the vertex v' is called the *gate* of v in H). Any gated subgraph is convex, but the converse is not true for general graphs. A graph $G = (V, E)$ is *isometrically embeddable* into a graph $H = (W, F)$ if there exists a mapping $\varphi : V \rightarrow W$ such that $d_H(\varphi(u), \varphi(v)) = d_G(u, v)$ for all vertices $u, v \in V$.

A graph G is called *median* if the interval intersection $I(x, y) \cap I(y, z) \cap I(z, x)$ is a singleton for each triplet x, y, z of vertices. Median graphs are bipartite. Basic examples of median graphs are trees, hypercubes, rectangular grids, and Hasse diagrams of distributive lattices and of median semilattices [7]. With any vertex v of a median graph $G = (V, E)$ is associated a *canonical partial order* \leq_v defined by setting $x \leq_v y$ if and only if $x \in I(v, y)$; v is called the *basepoint* of \leq_v . Since G is bipartite, the Hasse diagram G_v of the partial order (V, \leq_v) is the graph G in which any edge xy is directed from x to y if and only if the inequality $d_G(x, v) < d_G(y, v)$ holds. We call G_v a *pointed median graph*. There is a close relationship between pointed median graphs and median semilattices. A *median semilattice* is a meet semilattice (P, \leq) such that (i)

for every x , the *principal ideal* $\downarrow x = \{p \in P : p \leq x\}$ is a distributive lattice, and (ii) any three elements have a least upper bound in P whenever each pair of them does.

Theorem 4.1 ([5]). *The Hasse diagram of any median semilattice is a median graph. Conversely, every median graph defines a median semilattice with respect to any canonical order \leq_v .*

Median graphs can be obtained from hypercubes by amalgams and median graphs are themselves isometric subgraphs of hypercubes [8, 34]. The canonical isometric embedding of a median graph G into a (smallest) hypercube can be determined by the so called *Djoković-Winkler* (“parallelism”) relation Θ on the edges of G [23, 50]. For median graphs, the equivalence relation Θ can be defined as follows. First say that two edges uv and xy are in relation Θ' if they are opposite edges of a 4-cycle $uvxy$ in G . Then let Θ be the reflexive and transitive closure of Θ' . Any equivalence class of Θ constitutes a cutset of the median graph G , which determines one factor of the canonical hypercube [34]. The cutset (equivalence class) $\Theta(xy)$ containing an edge xy defines a convex split $\{W(x, y), W(y, x)\}$ of G [34], where $W(x, y) = \{z \in V : d_G(z, x) < d_G(z, y)\}$ and $W(y, x) = V \setminus W(x, y)$ (we call the complementary convex sets $W(x, y)$ and $W(y, x)$ *half-spaces*). Conversely, for every convex split of a median graph G there exists at least one edge xy such that $\{W(x, y), W(y, x)\}$ is the given split. We denote by $\{\Theta_i : i \in I\}$ the equivalence classes of the relation Θ (in [9], they were called parallelism classes). For an equivalence class $\Theta_i, i \in I$, we denote by $\{A_i, B_i\}$ the associated convex split. We say that Θ_i *separates* the vertices x and y if $x \in A_i, y \in B_i$ or $x \in B_i, y \in A_i$. The isometric embedding φ of G into a hypercube is obtained by taking a basepoint v , setting $\varphi(v) = \emptyset$ and for any other vertex u , letting $\varphi(u)$ be all parallelism classes of Θ which separate u from v .

From the definition it follows that any median graph G satisfies the following *quadrangle condition*: for any four vertices u, v, w, z with $d(v, z) = d(w, z) = 1$ and $2 = d(v, w) \leq d(u, v) = d(u, w) = d(u, z) - 1$, there exists a common neighbor x of v and w such that $d(u, x) = d(u, v) - 1$. In fact, x is the median of the triplet u, v, w . Since this median is unique, the vertex x in quadrangle condition is also unique.

We conclude this subsection with the following simple but useful local characterization of convex sets of median graphs (which holds for much more general classes of graphs):

Lemma 4.2 ([15]). *Let S be a connected subgraph of a median graph G . Then S is a convex subgraph if and only if S is locally-convex, i.e., $I(x, y) \subseteq S$ for any two vertices x, y of S having a common neighbor in S .*

We also recall that convex subgraphs and gated subgraphs of median graphs are the same:

Lemma 4.3 ([15]). *A subgraph H of a median graph G is convex if and only if H is gated.*

4.2. Nonpositively curved cube complexes. A 0-cube is a single point. A 1-cube is an isometric copy of the segment $[-1, 1]$ and has a cell structure consisting of 0-cells $\{\pm 1\}$ and a single 1-cell. An n -cube is an isometric copy of $[-1, 1]^n$, and has the product structure, so that each closed cell of $[-1, 1]^n$ is obtained by restricting some of the coordinates to $+1$ and some to -1 . A *cube complex* is obtained from a collection of cubes of various dimensions by isometrically identifying certain subcubes. The *dimension* $\dim(X)$ of a cube complex X is the largest value of d for which X contains a d -cube. A *square complex* is a cube complex of dimension 2. The 0-cubes and the 1-cubes of a cube complex X are called *vertices* and *edges* of X and define the graph $X^{(1)}$, the 1-skeleton of X . We denote the vertices of $X^{(1)}$ by $V(X)$ and the edges of $X^{(1)}$ by $E(X)$. For $i \in \mathbb{N}$, we denote by $X^{(i)}$ the i -skeleton of X , i.e., the cube complex consisting of all j -dimensional cubes of X , where $j \leq i$. A *square complex* X is a combinatorial 2-complex whose 2-cells are attached by closed combinatorial paths of length 4. Thus, one can consider each 2-cell as a square attached to the 1-skeleton $X^{(1)}$ of X . All cube complexes occurring in this paper are *simple* [27] in the sense that two distinct squares cannot meet along two consecutive edges.

The *star* $\text{St}(v, X)$ of a vertex v of X is the subcomplex spanned by all cubes containing v . The *link* of a vertex $x \in X$ is the simplicial complex $\text{Link}(x, X)$ with a $(d - 1)$ -simplex for each

d -cube containing x , with simplices attached according to the attachments of the corresponding cubes. More generally, the *link* of a k -dimensional cube Q of X is the simplicial complex $\text{Link}(Q, X)$ with a $(d - k - 1)$ -simplex for each d -cube containing Q , with simplices attached according to the attachments of the corresponding cubes. Note that in the definition of the link, the simplices are added with multiplicity: if x (or Q) belongs to a cube Q' in multiple ways, then Q' contributes to the link with multiple (disjoint) simplices. For example, if X is a 1-dimensional complex with only one 0-cube x and only one 1-cube e (a loop around x), then $\text{Link}(x, X)$ consists of two disjoint 0-simplices.

The link $\text{Link}(x, X)$ is said to be a *flag (simplicial) complex* if each $(d+1)$ -clique in $\text{Link}(x, X)$ spans an d -simplex. A cube complex X is *flag* if $\text{Link}(x, X)$ is a flag simplicial complex for every vertex $x \in X$. This flagness condition of $\text{Link}(x, X)$ can be restated as follows: whenever three $(k + 2)$ -cubes of X share a common k -cube containing x and pairwise share common $(k + 1)$ -cubes, then they are contained in a $(k + 3)$ -cube of X . A cube complex X is called *simply connected* if it is connected and if every continuous mapping of the 1-dimensional sphere S^1 into X can be extended to a continuous mapping of the disk D^2 with boundary S^1 into X . Note that X is connected iff $G(X) = X^{(1)}$ is connected, and X is simply connected iff $X^{(2)}$ is simply connected. Equivalently, a cube complex X is simply connected if X is connected and every cycle C of its 1-skeleton is null-homotopic, i.e., it can be contracted to a single point by elementary homotopies.

Given two cube complexes X and Y , a *covering (map)* is a surjection $\varphi: Y \rightarrow X$ mapping cubes to cubes and such that φ induces an isomorphism between $\text{Link}(v, Y)$ and $\text{Link}(\varphi(v), X)$. When the 1-skeleton $X^{(1)}$ of X does not contain loops or multiple edges, the condition on the links is equivalent to the following condition on the stars: $\varphi|_{\text{St}(v, Y)}: \text{St}(v, Y) \rightarrow \text{St}(\varphi(v), X)$ is an isomorphism for every vertex v in Y . The space Y is then called a *covering space* of X . For any vertex v of X , any vertex \tilde{v} of Y such that $\varphi(\tilde{v}) = v$ is called a *lift* of v . It is well-known that if X and Y are flag cube complexes, Y is a covering space of X if and only if the 2-skeleton $Y^{(2)}$ of Y is a covering space of $X^{(2)}$. A *universal cover* of X is a simply connected covering space; it always exists and it is unique up to isomorphism [29, Sections 1.3 and 4.1]. The universal cover of a complex X will be denoted by \tilde{X} . In particular, if X is simply connected, then its universal cover \tilde{X} is X itself.

An important class of cube complexes studied in geometric group theory and combinatorics is the class of nonpositively curved and CAT(0) cube complexes. We continue by recalling the definition of CAT(0) spaces. A *geodesic triangle* $\Delta = \Delta(x_1, x_2, x_3)$ in a geodesic metric space (X, d) consists of three points in X (the vertices of Δ) and a geodesic between each pair of vertices (the sides of Δ). A *comparison triangle* for $\Delta(x_1, x_2, x_3)$ is a triangle $\Delta(x'_1, x'_2, x'_3)$ in the Euclidean plane \mathbb{E}^2 such that $d_{\mathbb{E}^2}(x'_i, x'_j) = d(x_i, x_j)$ for $i, j \in \{1, 2, 3\}$. A geodesic metric space (X, d) is defined to be a *CAT(0) space* [24] if all geodesic triangles $\Delta(x_1, x_2, x_3)$ of X satisfy the comparison axiom of Cartan–Alexandrov–Toponogov: *If y is a point on the side of $\Delta(x_1, x_2, x_3)$ with vertices x_1 and x_2 and y' is the unique point on the line segment $[x'_1, x'_2]$ of the comparison triangle $\Delta(x'_1, x'_2, x'_3)$ such that $d_{\mathbb{E}^2}(x'_i, y') = d(x_i, y)$ for $i = 1, 2$, then $d(x_3, y) \leq d_{\mathbb{E}^2}(x'_3, y')$.* A geodesic metric space (X, d) is *nonpositively curved* if it is locally CAT(0), i.e., any point has a neighborhood inside which the CAT(0) inequality holds. CAT(0) spaces can be characterized in several different natural ways and have many strong properties, see for example [11]. In particular, a geodesic metric space (X, d) is CAT(0) if and only if (X, d) is simply connected and is nonpositively curved. Gromov [24] gave a beautiful combinatorial characterization of CAT(0) cube complexes, which can be also taken as their definition:

Theorem 4.4 ([24]). *A cube complex X endowed with the ℓ_2 -metric is CAT(0) if and only if X is simply connected and the links of all vertices of X are flag complexes. If Y is a cube complex in which the links of all vertices are flag complexes, then the universal cover \tilde{Y} of Y is a CAT(0) cube complex.*

In view of the second assertion of Theorem 4.4, the cube complexes in which the links of vertices are flag complexes are called *nonpositively curved cube complexes* or shortly *NPC complexes*. As a corollary of Gromov's result, *for any NPC complex X , its universal cover \tilde{X} is CAT(0)*.

A square complex X is a *VH-complex* (*vertical-horizontal complex*) if the 1-cells (edges) of X are partitioned into two sets V and H called *vertical* and *horizontal* edges respectively, and the edges in each square alternate between edges in V and H . Notice that if X is a VH-complex, then X satisfies the Gromov's nonpositive curvature condition since no three squares may pairwise intersect on three edges with a common vertex, thus VH-complexes are particular NPC square complexes. A VH-complex X is a *complete square complex* (CSC) [54] if any vertical edge and any horizontal edge incident to a common vertex belong to a common square of X . By [54, Theorem 3.8], if X is a complete square complex, then the universal cover \tilde{X} of X is isomorphic to the Cartesian product of two trees. By a *plane* Π in \tilde{X} we will mean a convex subcomplex of \tilde{X} isometric to \mathbb{R}^2 tiled by the grid \mathbb{Z}^2 into unit squares.

We continue with the bijection between CAT(0) cube complexes and median graphs:

Theorem 4.5 ([16, 41]). *Median graphs are exactly the 1-skeleta of CAT(0) cube complexes.*

The proof of Theorem 4.5 presented in [16] is based on the following local-to-global characterization of median graphs:

Theorem 4.6 ([16]). *A graph G is a median graph if and only if its cube complex is simply connected and G satisfies the 3-cube condition: if three squares of G pairwise intersect in an edge and all three intersect in a vertex, then they belong to a 3-cube.*

A *midcube* of the d -cube c , with $d \geq 1$, is the isometric subspace obtained by restricting exactly one of the coordinates of d to 0. Note that a midcube is a $(d - 1)$ -cube. The midcubes a and b of X are *adjacent* if they have a common face, and a *hyperplane* H of X is a subspace that is a maximal connected union of midcubes such that, if $a, b \subset H$ are midcubes, either a and b are disjoint or they are adjacent. Equivalently, a hyperplane H is a maximal connected union of midcubes such that, for each cube c , either $H \cap c = \emptyset$ or $H \cap c$ is a single midcube of c . The *carrier* $N(X)$ of a hyperplane H of X is the union of all cubes intersecting H .

Theorem 4.7 ([43]). *Each hyperplane H of a CAT(0) cube complex X is a CAT(0) cube complex of dimension at most $\dim(X) - 1$ and $X \setminus H$ consists of exactly two components, called halfspaces.*

A 1-cube e (an edge) is *dual* to the hyperplane H if the 0-cubes of e lie in distinct halfspaces of $X \setminus H$, i.e., if the midpoint of e is in a midcube contained in H . The relation "dual to the same hyperplane" is an equivalence relation on the set of edges of X ; denote this relation by Θ and denote by $\Theta(H)$ the equivalence class consisting of 1-cubes dual to the hyperplane H (Θ is precisely the parallelism relation on the edges of the median graph $X^{(1)}$).

The following results summarize the well known and many times rediscovered convexity properties of halfspaces and carriers of CAT(0) cube complexes.

Theorem 4.8 ([34, 49]). *If H is a hyperplane of a CAT(0) cube complex X , then the carrier $N(H)$ of H and the two halfspaces defined by H restricted to the vertices of X induce convex and thus gated subgraphs of the 1-skeleton $G(X)$ of X . Any convex subgraph H of $G(X)$ is the intersection of the halfspaces of $G(X)$ containing H .*

Proposition 4.9 ([49]). *For any set \mathcal{H} of d pairwise intersecting hyperplanes of a CAT(0) cube complex X , the carriers of the hyperplanes of \mathcal{H} intersect in a d -cube of X .*

4.3. Domains versus median graphs/CAT(0) cube complexes. Theorems 2.2 and 2.3 of Barthélemy and Constantin [9] (this result was independently rediscovered by Ardilla et al. [3] in the language of CAT(0) cube complexes) establish the following bijection between event structures and pointed median graphs (in [9], event structures are called sites):




FIGURE 2. A self-intersecting hyperplane (a), a one-sided hyperplane (b), a directly self-oscultating hyperplane (c), an indirectly self-oscultating hyperplane (d), and a pair of hyperplanes that inter-oscultate (e).

Theorem 4.10 ([9]). *The (undirected) Hasse diagram of the domain $(\mathcal{D}(\mathcal{E}), \subseteq)$ of any event structure $\mathcal{E} = (E, \leq, \#)$ is a median graph. Conversely, for any median graph G and any basepoint v of G , the pointed median graph G_v is isomorphic to the Hasse diagram of the domain of an event structure.*

In [13] we presented a new proof of Theorem 4.10. In the current paper we will only need the canonical construction of an event structure from a pointed median graph (or pointed CAT(0) cube complex), presented in [9] and briefly recalled here. Suppose that v is an arbitrary but fixed basepoint of a median graph G . The events of the event structure \mathcal{E}_v are the hyperplanes of X . Two hyperplanes H and H' define concurrent events if and only if they cross. The hyperplanes H and H' are in precedence relation $H \leq H'$ if and only if $H = H'$ or H separates H' from v . Finally, the events defined by H and H' are in conflict if and only if H and H' do not cross and neither separates the other from v .

4.4. Special cube complexes. Consider a cube complex Y . Analogously to CAT(0) cube complexes, one can define the parallelism relation Θ' on the set of edges $E(Y)$ of Y by setting that two edges of Y are in relation Θ' if they are opposite edges of a common 2-cube of Y . Let Θ be the reflexive and transitive closure of Θ' and let $\{\Theta_i : i \in I\}$ denote the equivalence classes of Θ . For an equivalence class Θ_i , the hyperplane H_i associated to Θ_i is the cube complex consisting of the midcubes of all cubes of Y containing one edge of Θ_i . The edges of Θ_i are *dual* to the hyperplane H_i . Let $\mathcal{H}(Y)$ be the set of hyperplanes of Y . The *carrier* $N(H)$ of a hyperplane H of Y is the union of all cubes intersecting H . The *open carrier* $\mathring{N}(H)$ is the union of all open cubes intersecting H .

The hyperplanes of a cube complex Y do not longer satisfy the nice properties of the hyperplanes of CAT(0) cube complexes: they do not partition the complex in exactly two parts, they may self-intersect, self-oscultate, two hyperplanes may at the same time cross and oscultate, etc. Haglund and Wise [27] detected five types of pathologies which may occur in a cube complex (see Figure 2):

- (a) self-intersecting hyperplane;
- (b) one-sided hyperplane;
- (c) directly self-oscultating hyperplane;
- (d) indirectly self-oscultating hyperplane;
- (e) a pair of hyperplanes, which both intersect and oscultate.

A hyperplane H is *two-sided* if $\mathring{N}(H)$ is homeomorphic to the product $H \times (-1, 1)$, and there is a combinatorial map $H \times [-1, 1] \rightarrow X$ mapping $H \times \{0\}$ identically to H . As noticed in [27, p.1562], requiring that the hyperplanes of Y are two-sided is equivalent to defining an orientation on the dual edges of H such that all sources of such edges belong to one of the sets $H \times \{-1\}, H \times \{1\}$ and all sinks belong to the other one. This orientation is obtained by taking the equivalence relation generated by elementary parallelism relation: declare two

oriented edges e_1 and e_2 of Y *elementary parallel* if there is a square of Y containing e_1 and e_2 as opposite sides and oriented in the same direction. Such an orientation o of the edges of Y is called an *admissible* orientation of Y . Observe that Y admits an admissible orientation if and only if every hyperplane H of Y is two-sided (one can choose an admissible orientation for each hyperplane independently). Given a cube complex Y and an admissible orientation o of Y , (Y, o) is called a *directed* cube complex.

We continue with the definition of each of the pathologies (in which we closely follow [27, Section 3]). The hyperplane is *one-sided* if it is not two-sided (see Figure 2(b)).

Two hyperplanes H_1 and H_2 *intersect* if there exists a cube Q and two distinct midcubes Q_1 and Q_2 of Q such that $Q_1 \subseteq H_1$ and $Q_2 \subseteq H_2$, i.e., there exists a square with two consecutive edges e_1, e_2 such that e_1 is dual to H_1 and e_2 is dual to H_2 .

A hyperplane H of Y *self-intersects* if it contains more than one midcube from the same cube, i.e., there exist two edges e_1, e_2 dual to H that are consecutive in some square of Y (see Figure 2(a)).

Let v be a vertex of Y and let e_1, e_2 be two distinct edges incident to v but such that e_1 and e_2 are not consecutive edges in some square containing v . The hyperplanes H_1 and H_2 *osculate* at (v, e_1, e_2) if e_1 is dual to H_1 and e_2 is dual to H_2 . The hyperplane H *self-osculate* at (v, e_1, e_2) if e_1 and e_2 are dual to H . Consider a two-sided hyperplane H and an admissible orientation o of its dual edges. Suppose that H self-osculate at (v, e_1, e_2) . If v is the source of both e_1 and e_2 or the sink of both e_1 and e_2 , then we say that H *directly self-osculate* at (v, e_1, e_2) (see Figure 2(c)). If v is the source of one of e_1, e_2 , and the sink of the other, then we say that H *indirectly self-osculate* at (v, e_1, e_2) (see Figure 2(d)). Note that a self-osculation of a hyperplane H is either direct or indirect, and this is independent of the orientation of the edges dual to H .

Two hyperplanes H_1 and H_2 *inter-osculate* if they both intersect and osculate (see Figure 2(e)).

Haglund and Wise [27, Definition 3.2] called a cube complex Y *special* if its hyperplanes are two-sided, do not self-intersect, do not directly self-osculate, and no two hyperplanes inter-osculate. The definition of a special cube complex Y depends only of the 2-skeleton $Y^{(2)}$ [27, Remark 3.4]. Since no two hyperplanes of Y self-osculate, any special cube complex and its 2-skeleton satisfy the 3-cube condition. In fact, Haglund and Wise proved that special cube complexes can be seen as nonpositively curved complexes:

Lemma 4.11 ([27], Lemma 3.13). *If X is a special cube complex, then X is contained in a unique smallest nonpositively curved cube complex with the same 2-skeleton as X .*

In view of this lemma, in the following, we will consider only 2-dimensional special cube complexes, since they can always be canonically completed to NPC complexes that are also special.

5. GEODESIC TRACES AND PRIME TRACES

Let $M = (\Sigma, I)$ be a trace alphabet and let $\mathcal{E} = (E, \leq, \#)$ be an M -labeled event structure. Then the concurrency relation \parallel of \mathcal{E} coincides with the independence relation I . Let $\mathcal{D}(\mathcal{E})$ denote the domain of \mathcal{E} . Let $G(\mathcal{E})$ denote the covering graph of $\mathcal{D}(\mathcal{E})$ and $X(\mathcal{E})$ denote the CAT(0) cube complex of $G(\mathcal{E})$ pointed at vertex v_0 . Any vertex v of the median graph $G(\mathcal{E})$ corresponds to a configuration $c(v)$ of $\mathcal{D}(\mathcal{E})$; in particular, $c(v_0) = \emptyset$. In this auxiliary section we present characterizations of traces arising from geodesics of the domain (i.e., from the median graph $G(\mathcal{E})$).

5.1. Geodesic traces. Any shortest path $\pi = (v_0, v_1, \dots, v_{k-1}, v_k = v)$ from v_0 to a vertex v of $G(\mathcal{E})$ gives rise to an word $\sigma(\pi)$ of Σ^* : the i th letter of $\sigma(\pi)$ is the label $\lambda(v_{i-1}v_i)$ of the edge $v_{i-1}v_i$. We will say that a word $\sigma \in \Sigma^*$ is *geodesic* if $\sigma = \sigma(\pi)$ for a shortest path π between v_0 and a vertex v of $G(\mathcal{E})$. The trace $\langle \sigma \rangle$ of a geodesic word σ is called a *geodesic trace*.

Two shortest (v_0, v) -paths π and π' of $G(\mathcal{E})$ are called *homotopic* if they can be transformed one into another by a sequence of elementary homotopies, i.e., there exists a finite sequence

$\pi =: \pi_1, \pi_2, \dots, \pi_{k-1}, \pi_k := \pi'$ of shortest (v_0, v) -paths such that for any $i = 1, \dots, k-1$ the paths π_i and π_{i+1} differ only in a square $Q_i = (v_{j-1}, v_j, v_{j+1}, v'_j)$ of $X(\mathcal{E})$.

The following result is well-known; we provide a simple proof using median graphs.

Lemma 5.1. *Any two shortest (v_0, v) -paths π_1 and π_2 of $G(\mathcal{E})$ are homotopic.*

Proof. We proceed by induction on the distance $k = d(v_0, v)$. If $k = 2$, then the result is obvious because the paths π_1 and π_2 bound a square of $X(\mathcal{E})$. Let w_1 be the neighbor of v in π_1 and w_2 be the neighbor of v in π_2 . Observe that $d(v_0, w_1) = d(v_0, w_2) = d(v_0, v) - 1 = k - 1$. If $w_1 = w_2$, then the result holds by induction hypothesis. Otherwise, by the quadrangle condition, there exists a common neighbor x of w_1 and w_2 such that $d(v_0, x) = k - 2$. Let π'_1 be the subpath of π_1 from v_0 to w_1 , let π'_2 be the subpath of π_2 from v_0 to w_2 , and let π_3 be a shortest path from v_0 to x . By induction hypothesis, the path π'_1 is homotopic to the path $\pi''_1 = \pi_3 \cdot (x, w_1)$ and the path π'_2 is homotopic to the path $\pi''_2 = \pi_3 \cdot (x, w_2)$. Since vw_1xw_2 is a square of $X(\mathcal{E})$, the path $\pi_3 \cdot (x, w_1, v)$ is homotopic to the path $\pi_3 \cdot (x, w_2, v)$. Consequently, the paths π_1 and π_2 are homotopic. \square

Lemma 5.2. *If π and π' are two shortest (v_0, v) -paths of $G(\mathcal{E})$, then $\sigma(\pi)$ and $\sigma(\pi')$ belong to the same trace.*

Proof. By Lemma 5.1, the paths π and π' are homotopic. Thus there exists a finite sequence $\pi =: \pi_1, \pi_2, \dots, \pi_{k-1}, \pi_k := \pi'$ of shortest (v_0, v) -paths such that for any $i = 1, \dots, k-1$ the paths π_i and π_{i+1} differ only in a square $Q_i = (v_{j-1}, v_j, v_{j+1}, v'_j)$ of $X(\mathcal{E})$. Denote by σ_1 the sequence of labels of the edges of the path (v_0, \dots, v_j) . Analogously, denote by σ_2 the sequence of labels of the edges of the path $(v_{j+1}, \dots, v_k = v)$. The edges $v_{j-1}v_j$ and v'_jv_{j+1} are dual to the same hyperplane H_a , thus $\lambda(v_{j-1}v_j) = \lambda(v'_jv_{j+1}) = \lambda(H_a) = a$. Analogously, the edges $v_{j-1}v'_j$ and v_jv_{j+1} are dual to the same hyperplane H_b , whence $\lambda(v_{j-1}v'_j) = \lambda(v_jv_{j+1}) = \lambda(H_b) = b$. Since H_a and H_b intersect, the events corresponding to those hyperplanes are concurrent, therefore $(a, b) \in I$. Consequently, $\sigma(\pi_i) = \sigma_1ab\sigma_2 \leftrightarrow_I \sigma_1ba\sigma_2 = \sigma(\pi_{i+1})$, establishing that for any two consecutive paths π_i, π_{i+1} the words $\sigma(\pi_i)$ and $\sigma(\pi_{i+1})$ belong to the same trace, proving that $\sigma(\pi)$ and $\sigma(\pi')$ belong to the same Mazurkiewicz trace. \square

Lemma 5.3. *For any shortest (v_0, v) -path π of $G(\mathcal{E})$, the geodesic trace $\langle \sigma(\pi) \rangle$ consists exactly of all $\sigma(\pi')$ such that π' is a shortest (v_0, v) -path.*

Proof. From Lemma 5.2 it follows that $\{\sigma(\pi') : \pi' \text{ is a shortest } (v_0, v) \text{-path}\} \subseteq \langle \sigma(\pi) \rangle$. To prove the converse inclusion it suffices to show that if $\sigma_1, \sigma_2 \in \Sigma^*$, $(a, b) \in I$ and $\sigma_1ab\sigma_2 = \sigma(\pi')$ for a shortest (v_0, v) -path π' , then $\sigma_1ba\sigma_2 = \sigma(\pi'')$ for a shortest (v_0, v) -path π'' . Indeed, since $(a, b) \in I$, the hyperplanes H_a and H_b dual to the incident a - and b -edges of π' intersect in an ab -square Q . Moreover, the carriers of H_a and H_b intersect in Q . Since those carriers also contain the incident a - and b -edges of π' , it can be easily deduced that Q contains the a - and b -edges of π' . Let π'' be obtained from π' by replacing the ab -path by the ba -path of Q . Then obviously π'' is a shortest (v_0, v) -path and that $\sigma(\pi'') = \sigma_1ba\sigma_2$. This concludes the proof. \square

For a vertex v of $G(\mathcal{E})$, we will denote by $\langle \sigma_v \rangle$ the Mazurkiewicz geodesic trace of all shortest (v_0, v) -paths, i.e., the trace of the interval $I(v_0, v)$. Denote by $\mathcal{GT}(\mathcal{E})$ the set of all geodesic traces of \mathcal{E} . From Lemma 5.3, we immediately obtain the following corollary:

Corollary 5.4. *There exists a natural bijection $v \mapsto \langle \sigma_v \rangle$ between the set of vertices of $G(\mathcal{E})$ (i.e., configurations of \mathcal{E}) and the set $\mathcal{GT}(\mathcal{E})$ of geodesic traces of \mathcal{E} .*

Now we describe the precedence and the conflict relations between geodesic traces.

Lemma 5.5. *For two geodesic traces $\langle \sigma_u \rangle$ and $\langle \sigma_v \rangle$, we have $\langle \sigma_u \rangle \sqsubseteq \langle \sigma_v \rangle$ iff $u \in I(v_0, v)$.*

Proof. By definition of \sqsubseteq and Lemma 5.3, $\langle \sigma_u \rangle \sqsubseteq \langle \sigma_v \rangle$ iff there exists a shortest (v_0, u) -path π' and a shortest (v_0, v) -path π such that $\sigma(\pi')$ is a prefix of $\sigma(\pi)$. But this is equivalent to the fact that π' is a subpath of π which is equivalent to the fact that u belongs to the interval $I(v_0, v)$. \square

We will say that two geodesic traces $\langle \sigma_u \rangle$ and $\langle \sigma_v \rangle$ are in *conflict* if there does not exist a geodesic trace $\langle \sigma_w \rangle$ such that $\langle \sigma_u \rangle \sqsubseteq \langle \sigma_w \rangle$ and $\langle \sigma_v \rangle \sqsubseteq \langle \sigma_w \rangle$. In view of Lemma 5.5, this definition can be rephrased in the following way:

Lemma 5.6. *Two geodesic traces $\langle \sigma_u \rangle$ and $\langle \sigma_v \rangle$ are in conflict iff there does not exist a vertex w such that $u, v \in I(v_0, w)$.*

5.2. Prime geodesic traces. Recall that a trace $\langle \sigma \rangle$ is *prime* if σ is non-null and for every $\sigma' \in \langle \sigma \rangle$, $\text{last}(\sigma) = \text{last}(\sigma')$. We will characterize now prime geodesic traces of \mathcal{E} , in particular we will prove that they are in bijection with the hyperplanes (events) of \mathcal{E} .

We call an interval $I(v_0, v)$ *prime* if the vertex v has degree 1 in the subgraph induced by $I(v_0, v)$.

Lemma 5.7. *A geodesic trace $\langle \sigma_v \rangle$ is prime iff the interval $I(v_0, v)$ is prime.*

Proof. If $I(v_0, v)$ is prime and v' is the unique neighbor of v in $I(v_0, v)$, then for any shortest (v_0, v) -path π , we will have $\text{last}(\sigma(\pi)) = \lambda(v'v)$. Thus the geodesic trace $\langle \sigma_v \rangle$ is prime. Conversely, if the geodesic trace $\langle \sigma_v \rangle$ is prime, then $\text{last}(\sigma(\pi)) = \text{last}(\sigma(\pi'))$ for any two shortest (v_0, v) -paths. Since any two edges of $I(v_0, v)$ incident to v are labeled differently, the paths π and π' have the same last edge, i.e., v has degree 1 in $I(v_0, v)$. \square

Lemma 5.8. *Each hyperplane H of $X(\mathcal{E})$ (i.e., each event of \mathcal{E}) gives a unique prime geodesic trace $\langle \sigma_H \rangle := \langle \sigma_v \rangle$ defined by the prime interval $I(v_0, v)$, where v' is the gate of v_0 in the carrier $N(H)$ of the hyperplane H and v is the neighbor of v' such that the edge $v'v$ is dual to H . Conversely, for each prime geodesic trace $\langle \sigma_u \rangle$ there exists a unique hyperplane H such that $\langle \sigma_u \rangle = \langle \sigma_H \rangle$.*

Proof. For a hyperplane H , let v' and v be defined as in the formulation of the lemma. Let A and B be the two complementary halfspaces defined by H and suppose that $v' \in A$ and $v \in B$. We assert that the interval $I(v_0, v)$ is prime, i.e., v' is the unique neighbor of v in $I(v_0, v)$. Suppose by way of contradiction that v is adjacent in $I(v_0, v)$ to another vertex v'' . Since $v \in I(v', v'')$ and the halfspace A is convex, v'' cannot belong to this halfspace. Thus v'' belongs to B . By quadrangle condition, there exists a vertex w adjacent to v', v'' and one step closer to v_0 than v' and v'' . From the convexity of B we conclude that w belongs to A . Since $w \in A$ is adjacent to $v'' \in B$, w belongs to the carrier $N(H)$ of H . Since $d(v_0, w) < d(v_0, v')$ we obtain a contradiction with the assumption that v' is the gate of v_0 in $N(X)$. This contradiction establishes that the interval $I(v_0, v)$ is prime.

Conversely, let $\langle \sigma_u \rangle$ be a prime geodesic trace and let u' be the unique neighbor of u in $I(v_0, u)$. Let H be the hyperplane dual to the edge $u'u$ and A and B be the halfspaces defined by H with $u' \in A, u \in B$. We assert that u' is the gate of v_0 in the carrier $N(H)$ of H . Suppose that this is not true and let v' be the gate of v_0 in $N(H)$. Let $v'v$ be the edge incident to v' and dual to H . Then $v' \in I(v_0, u)$ and $v \in I(v', u) \subseteq I(v_0, u)$. Since B is convex, $I(v, u) \subseteq B$, thus u has a second neighbor in $I(v_0, u)$, contrary to the assumption that the interval $I(v_0, u)$ is prime. Hence u' is the gate of v_0 in $N(H)$, establishing that $\langle \sigma_u \rangle = \langle \sigma_H \rangle$. \square

Notice that the vertices v such that the interval $I(v_0, v)$ is prime are exactly the join irreducible elements of the poset $(\mathcal{D}(\mathcal{E}), \subseteq)$ (i.e., the nonminimal elements which cannot be written as the supremum of finitely many other elements). The bijection between the set $\mathbb{J}(X(\mathcal{E}))$ of join irreducibles and the set \mathcal{H} of hyperplanes was also established in [3].

Denote by $\mathcal{PGT}(\mathcal{E})$ the set of geodesic prime traces of \mathcal{E} . From Lemma 5.8, we immediately obtain the following corollary:

Corollary 5.9. *There exist natural bijections $H \mapsto \langle \sigma_v \rangle$ between the set of hyperplanes of $X(\mathcal{E})$ (i.e., events of \mathcal{E}), the set $\mathcal{PGT}(\mathcal{E})$ of prime geodesic traces of \mathcal{E} , and the set $\mathbb{J}(X(\mathcal{E}))$ of join irreducible elements of (\mathcal{D}, \subseteq) .*

Lemma 5.10. *For two hyperplanes H', H of $X(\mathcal{E})$ with prime geodesic traces $\langle \sigma_u \rangle$ and $\langle \sigma_v \rangle$, respectively, $H' \leq H$ holds iff $\langle \sigma_u \rangle \sqsubseteq \langle \sigma_v \rangle$.*

Proof. By Lemma 5.5 it suffices to show that $H' \leq H$ iff $u \in I(v_0, v)$. Let A', B' and A, B be the complementary halfspaces defined by H' and H , respectively, and suppose that $v_0 \in A' \cap A$. Let u' be the gate of v_0 in $N(H')$ and v' be the gate of v_0 in $N(H)$. Then $u' \in A'$ and u is the neighbor of u' in B' . Analogously, $v' \in A$ and v is the neighbor of v' in B . Notice also that u is the gate of v_0 in B' and that v is the gate of v_0 in B .

First suppose that $H' \leq H$, i.e., H' separates v_0 from H . This is equivalent to the inclusion $B \subseteq B'$. Since u is the gate of v_0 in B' and $v \in B$, this implies that $u \in I(v_0, v)$. Conversely, let $u \in I(v_0, v)$. Suppose by way of contradiction that $B \not\subseteq B'$, i.e., there exists a vertex $x \in B \cap A'$. Since $u' \in I(v_0, u) \subset I(v_0, v)$, the vertex v belongs to the halfspace B' . On the other hand, since v is the gate of v_0 in B and $x \in B$, we conclude that $v \in I(v_0, x)$. Since $v_0, x \in A'$ and $v \in B'$, we obtain a contradiction with the convexity of A' . This establishes that $u \in I(v_0, v)$ implies $H' \leq H$. \square

6. DIRECTED NPC COMPLEXES

Since we can define event structures from their domains, universal covers of NPC complexes represent a rich source of event structures. To obtain regular event structures, it is natural to consider universal covers of finite NPC complexes. Moreover, since domains of event structures are directed, it is natural to consider universal covers of NPC complexes whose edges are directed. However, the resulting directed universal covers are not in general domains of event structures. In particular, the domains corresponding to pointed median graphs given by Theorem 4.10 cannot be obtained in this way. In order to overcome this difficulty, in [13] we introduced directed median graphs and directed NPC complexes. Using these notions, we defined regular event structures starting from finite directed NPC complexes. In this section, we recall and extend these definitions and constructions.

6.1. Directed median graphs. A *directed median graph* is a pair $\vec{G} = (G, o)$, where G is a median graph and o is an orientation of the edges of G in a such a way that opposite edges of squares of G have the same direction. By transitivity of Θ , all edges from the same parallelism class Θ_i of G have the same direction. Since each Θ_i partitions G into two parts, o defines a partial order \prec_o on the vertex-set of G . For a vertex v of G , let $\mathcal{F}_o(v, G) = \{x \in V : v \prec_o x\}$ be the principal filter of v in the partial order $(V(G), \prec_o)$. For any canonical basepoint order \leq_v of G , (G, \leq_v) is a directed median graph. The converse is obviously not true: the 4-regular tree F_4 directed so that each vertex has two incoming and two outgoing arcs is a directed median graph which is not induced by a basepoint order.

Lemma 6.1 ([13]). *For any vertex v of a directed median graph $\vec{G} = (G, o)$, the following holds:*

- (1) $\mathcal{F}_o(v, G)$ induces a convex subgraph of G ;
- (2) the restriction of the partial order \prec_o on $\mathcal{F}_o(v, G)$ coincides with the restriction of the canonical basepoint order \leq_v on $\mathcal{F}_o(v, G)$;
- (3) $\mathcal{F}_o(v, G)$ together with \prec_o is the domain of an event structure;
- (4) for any vertex $u \in \mathcal{F}_o(v, G)$, the principal filter $\mathcal{F}_o(u, G)$ is included in $\mathcal{F}_o(v, G)$ and $\mathcal{F}_o(u, G)$ coincides with the principal filter of u with respect to the canonical basepoint order \leq_v on $\mathcal{F}_o(v, G)$.

A *directed (x, y) -path* of a directed median graph $\vec{G} = (G, o)$ is a (x, y) -path $\pi(x, y) = (x = x_1, x_2, \dots, x_{k-1}, x_k = y)$ of G in which any edge $x_i x_{i+1}$ is directed in \vec{G} from x_i to x_{i+1} .

Lemma 6.2. *Any directed path of a directed median graph \vec{G} is a shortest path of the median graph G .*

Proof. Since halfspaces of G are convex, a path $\pi(x, y)$ of G is a shortest path if and only if any hyperplane H of G intersects $\pi(x, y)$ in at most one edge. Since all edges of G dual to the same hyperplane H are directed in \vec{G} in the same way, H intersects a directed path $\pi(x, y)$ of \vec{G} in at most one edge. Hence the support of $\pi(x, y)$ is a shortest (x, y) -path in G . \square

6.2. Directed NPC cube complexes. A *directed NPC complex* is a directed cube complex (Y, o) , where Y is a nonpositively curved cube complex and o is an admissible orientation of Y . Recall that this means that o is an orientation of the edges of Y in a such a way that the opposite edges of the same square of Y have the same direction. For an edge xy , we will denote $o(xy)$ by \vec{xy} if x is the source and y is the sink of $o(xy)$ and by \vec{yx} otherwise. Note that there exist NPC complexes that do not admit any admissible orientation: consider a Möbius band of squares, for example. An admissible orientation o of Y induces in a natural way an orientation \tilde{o} of the edges of its universal cover \tilde{Y} , so that (\tilde{Y}, \tilde{o}) is a directed CAT(0) cube complex and $(\tilde{Y}^{(1)}, \tilde{o})$ is a directed median graph.

In the following, we need to consider directed colored NPC complexes and directed colored median graphs. A coloring ν of a directed NPC complex (Y, o) is an arbitrary map $\nu : E(Y) \rightarrow \Sigma$ where Σ is a set of colors. Note that a labeling is a coloring, but not the converse: *labelings* are precisely the colorings in which opposite edges of any square have the same color. In the following, we will denote a directed colored NPC complexes by bold letters like $\mathbf{Y} = (Y, o, \nu)$. Sometimes, we need to forget the colors and the orientations of the edges of these complexes. For a complex \mathbf{Y} , we denote by Y the complex obtained by forgetting the colors and the orientations of the edges of \mathbf{Y} (Y is called the *support* of \mathbf{Y}), and we denote by (Y, o) the directed complex obtained by forgetting the colors of \mathbf{Y} . We also consider directed colored median graphs that will be the 1-skeletons of directed colored CAT(0) cube complexes. Again we will denote such directed colored median graphs by bold letters like $\mathbf{G} = (G, o, \nu)$. Note that (uncolored) directed NPC complexes can be viewed as directed colored NPC complexes where all edges have the same color.

When dealing with directed colored NPC complexes, we consider only homomorphisms that preserve the colors and the directions of edges. More precisely, $\mathbf{Y}' = (Y', o', \nu')$ is a covering of $\mathbf{Y} = (Y, o, \nu)$ via a covering map φ if Y' is a covering of Y via φ and for any edge $e \in E(Y')$ directed from s to t , $\nu(\varphi(e)) = \nu'(e)$ and $\varphi(e)$ is directed from $\varphi(s)$ to $\varphi(t)$. Since any coloring ν of a directed colored NPC complex Y leads to a coloring of its universal cover \tilde{Y} , one can consider the colored universal cover $\tilde{\mathbf{Y}} = (\tilde{Y}, \tilde{o}, \tilde{\nu})$ of \mathbf{Y} .

When we consider principal filters in directed colored median graphs $\mathbf{G} = (G, o, \nu)$ (in particular, when G is the 1-skeleton of the universal cover $\tilde{\mathbf{Y}}$ of a directed colored NPC complex \mathbf{Y}), we say that two filters are isomorphic if there is an isomorphism between them that preserves the directions and the colors of the edges.

We now formulate the crucial regularity property of directed colored median graphs $(\tilde{Y}^{(1)}, \tilde{o}, \tilde{\nu})$ when (Y, o, ν) is finite.

Lemma 6.3 ([13]). *If $\mathbf{Y} = (Y, o, \nu)$ is a finite directed colored NPC complex, then $\tilde{\mathbf{Y}}^{(1)} = (\tilde{Y}^{(1)}, \tilde{o}, \tilde{\nu})$ is a directed median graph with at most $|V(Y)|$ isomorphism types of colored principal filters. In particular, if (Y, o) is a finite directed NPC complex, then $(\tilde{Y}^{(1)}, \tilde{o})$ is a directed median graph with at most $|V(Y)|$ isomorphism types of principal filters.*

Proposition 6.4 ([13]). *Consider a finite (uncolored) directed NPC complex (Y, o) . Then for any vertex \tilde{v} of the universal cover \tilde{Y} of Y , the principal filter $\mathcal{F}_{\tilde{o}}(\tilde{v}, \tilde{Y}^{(1)})$ with the partial order $<_{\tilde{o}}$ is the domain of a regular event structure with at most $|V(Y)|$ different isomorphism types of principal filters.*

We will call an event structure $\mathcal{E} = (E, \leq, \#)$ and its domain $\mathcal{D}(\mathcal{E})$ *strongly regular* if $\mathcal{D}(\mathcal{E})$ is isomorphic to a principal filter of the universal cover of some finite directed NPC complex. In view of Proposition 6.4, any strongly regular event structure is regular.

7. DIRECTED SPECIAL CUBE COMPLEXES

7.1. The results. Consider a finite NPC complex Y and let $\mathcal{H} = \mathcal{H}(Y)$ be the set of hyperplanes of Y . We define a canonical labeling $\lambda_{\mathcal{H}} : E(Y) \rightarrow \mathcal{H}$ by setting $\lambda_{\mathcal{H}}(e) = H$ if the edge e is dual to H . For any covering map $\varphi : \tilde{Y} \rightarrow Y$, $\lambda_{\mathcal{H}}$ is naturally extended to a labeling $\tilde{\lambda}_{\mathcal{H}}$ of $E(\tilde{Y})$ by setting $\tilde{\lambda}_{\mathcal{H}}(e) = \lambda_{\mathcal{H}}(\varphi(e))$. In [13] we proved that the strongly regular event structures

obtained from finite special cube complexes are trace regular event structures and that this characterizes special cube complexes:

Proposition 7.1 ([13]). *A finite NPC complex Y with two-sided hyperplanes is special if and only if there exists an independence relation I on $\mathcal{H} = \mathcal{H}(Y)$ such that for any admissible orientation o of Y , for any covering map $\varphi : \tilde{Y} \rightarrow Y$, and for any principal filter $\mathcal{D} = (\mathcal{F}_{\tilde{o}}(\tilde{v}, \tilde{Y}^{(1)}), \prec_{\tilde{o}})$ of (\tilde{Y}, \tilde{o}) , the canonical labeling $\tilde{\lambda}_{\mathcal{H}}$ is a regular trace labeling of \mathcal{D} with the trace alphabet (\mathcal{H}, I) .*

A finite NPC complex X is called *virtually special* [27, 28] if X admits a finite special cover, i.e., there exists a finite special NPC complex Y and a covering map $\varphi : Y \rightarrow X$. We will call a strongly regular event structure $\mathcal{E} = (E, \leq, \#)$ and its domain $\mathcal{D}(\mathcal{E})$ *cover-special* if $\mathcal{D}(\mathcal{E})$ is isomorphic to a principal filter of the universal cover of some virtually special complex with an admissible orientation.

Theorem 7.2 ([13]). *Any cover-special event structure \mathcal{E} admits a regular trace labeling, i.e., Thiagarajan's Conjecture 3.3 is true for cover-special event structures.*

In the following, we will need an extension of the second part of Proposition 7.1. Let Y be a finite cube complex with two-sided hyperplanes and let o be an admissible orientation of Y . Since the hyperplanes of Y are two-sided, there exists a bijection between the labelings of the edges of Y (i.e., colorings in which opposite edges of each square have equal colors) and the labelings of the hyperplanes of Y . Let $M = (\Sigma, I)$ be a trace alphabet. Extending the definition of trace labelings of domains of event structures (pointed CAT(0) cube complexes), we call a labeling $\lambda : E(Y) \rightarrow \Sigma$ of (Y, o) a *trace labeling* if the following conditions hold:

- (TL1) if there exists a square of Y in which two opposite edges are labeled a and two other opposite edges are labeled b , then $(a, b) \in I$;
- (TL2) for any vertex v of Y , any two distinct outgoing edges \vec{vx}, \vec{vy} have different labels and $(\lambda(\vec{vx}), \lambda(\vec{vy})) \in I$ iff \vec{vx} and \vec{vy} belong to a common square of Y ;
- (TL3) $(\lambda(\vec{xv}), \lambda(\vec{yv})) \in I$ iff \vec{xv} and \vec{yv} belong to a common square of Y ;
- (TL4) for any vertex v of Y , any two distinct incoming edges \vec{xv}, \vec{yv} have different labels and $(\lambda(\vec{xv}), \lambda(\vec{yv})) \in I$ iff \vec{xv} and \vec{yv} belong to a common square of Y .

Since for a trace labeling λ all edges dual to a hyperplane of Y have the same label, λ defines in a canonical way a labeling $\lambda : \mathcal{H} \rightarrow \Sigma$ of the hyperplanes \mathcal{H} of Y : for a hyperplane H , $\lambda(H) = \lambda(e)$ for any edge e dual to H . Notationally, for an edge xy of Y directed from x to y and its dual hyperplane H , we will write $\lambda(xy) = \lambda(\vec{xy}) = \lambda(H)$ to denote the (same) label of xy , \vec{xy} , and H .

Remark 7.3. Notice that (TL1) is a consequence of the three other axioms (TL2)-(TL4). Observe that (TL2)-(TL4) are equivalent to the condition that for any two incident edges $e_1, e_2 \in Y$, $(\lambda(e_1), \lambda(e_2)) \in I$ iff e_1 and e_2 belong to a common square of Y . Consequently, for any two letters $a, b \in \Sigma$ such that there are no hyperplanes $H_a, H_b \in \mathcal{H}$ labeled respectively a and b and intersecting or osculating, the axioms (TL1)-(TL4) hold for a and b , no matter whether (a, b) is in I or in D .

Remark 7.4. Even if formulated differently, the first three axioms (TL1)-(TL3) of a trace labeling of a directed cube complex can be viewed as a ‘‘local’’ reformulation of the axioms (LES1)-(LES3) of a trace labeling of the domain of an event structure. On the other hand, for domains of event structures the axiom (TL4) is implied by the axioms (LES1)-(LES3) because such domains are pointed median graphs and therefore, by the quadrangle condition, any two directed edges with the same sink belong to a square.

The existence of trace labelings characterizes the special cube complexes among finite cube complexes:

Theorem 7.5. *For a finite cube complex Y with two-sided hyperplanes the following conditions are equivalent:*

- (1) Y is special;

- (2) for any admissible orientation o of Y there exists a trace labeling λ of (Y, o) ;
(3) there exists an admissible orientation o of Y such that (Y, o) admits a trace labeling.

In view of Theorem 7.5, if a finite directed cube complex (Y, o) is given with a trace labeling, Y is supposed to be special. As before, given a directed special cube complex (Y, o) with a trace labeling $\lambda : E(Y) \rightarrow \Sigma$ for a trace alphabet $M = (\Sigma, I)$, let $\tilde{\mathbf{Y}} = (\tilde{Y}, \tilde{o}, \tilde{\lambda})$ denotes the directed labeled universal cover of Y and let $\varphi : \tilde{Y} \rightarrow Y$ denotes the covering map.

Proposition 7.6. *Let (Y, o) be a directed special cube complex with two-sided hyperplanes, $M = (\Sigma, I)$ be a trace alphabet, and $\lambda : E(Y) \rightarrow \Sigma$ be a trace labeling of Y . Then for any principal filter $\mathcal{D} = (\mathcal{F}_{\tilde{o}}(\tilde{v}, \tilde{Y}^{(1)}), \prec_{\tilde{o}})$ of $\tilde{\mathbf{Y}} = (\tilde{Y}, \tilde{o})$, the labeling $\tilde{\lambda}$ is a regular trace labeling of \mathcal{D} with the trace alphabet (Σ, I) .*

Remark 7.7. Since a trace labeling of a directed special cube complex (Y, o) also satisfies (TL4), the labeling $\tilde{\lambda}$ is a regular trace labeling not only of any principal filter $(\mathcal{F}_{\tilde{o}}(\tilde{v}, \tilde{Y}^{(1)}), \prec_{\tilde{o}})$ of $\tilde{\mathbf{Y}} = (\tilde{Y}, \tilde{o})$ but also of the whole complex $\tilde{\mathbf{Y}} = (\tilde{Y}, \tilde{o})$.

Remark 7.8. The canonical labeling $\lambda_{\mathcal{H}}$ in Proposition 7.1 is a trace labeling because all hyperplanes are labeled differently. Thus, Theorem 7.5 and Proposition 7.6 can be viewed as extensions of Proposition 7.1. However, formally Proposition 7.1 cannot be directly deduced from Theorem 7.5 and Proposition 7.6 because in those two results we assume that a trace labeling of a complex is satisfying (TL4).

7.2. Proof of Theorem 7.5. The implication (1) \implies (2) follows from Proposition 7.1 while (2) \implies (3) is trivial. To prove (3) \implies (1) suppose that o is an admissible orientation of Y such that (Y, o) admits a trace labeling $\lambda : E(Y) \rightarrow \Sigma$ with the trace alphabet $M = (\Sigma, I)$. We assert that Y is special.

First, if Y contains a self-intersecting hyperplane H , then there exists a square Q such that the four edges of Q are dual to H . Consequently, the four directed edges of Q are labeled $\lambda(H)$. Since o is an admissible orientation, there exists a vertex $v \in Q$ that has two outgoing edges with the same label, contradicting (TL2).

Now suppose that Y contains a hyperplane H that directly self-osculte at (v, e_1, e_2) . Let $e_1 = xv$ and $e_2 = yv$, and observe that with respect to the orientation o , either $e_1 = \vec{v\bar{x}}$ and $e_2 = \vec{v\bar{y}}$, or $e_1 = \vec{x\bar{v}}$ and $e_2 = \vec{y\bar{v}}$. This contradicts (TL2) in the first case and (TL4) in the second case, since $\lambda(e_1) = \lambda(e_2)$.

Finally, if Y contains two hyperplanes H_1 and H_2 that inter-osculte, then they osculte at (v, e'_1, e'_2) and they intersect on a square Q . Let $Q = (e_1, e_2, e_3, e_4)$. Suppose that the edges e'_1, e_1, e_3 are dual to H_1 and e'_2, e_2, e_4 are dual to H_2 . Hence $\lambda(e'_1) = \lambda(e_1) = \lambda(e_3) = \lambda(H_1)$ and $\lambda(e'_2) = \lambda(e_2) = \lambda(e_4) = \lambda(H_2)$. Since o is an admissible orientation, Q has a source s . By (TL2) applied at s , this implies that $(\lambda(e_1), \lambda(e_2)) \in I$. But if we consider the edges e'_1, e'_2 incident to v , by Remark 7.3, we have that $(\lambda(e_1), \lambda(e_2)) \notin I$, a contradiction. This completes the proof of Theorem 7.5.

7.3. Proof of Proposition 7.6. By Proposition 6.4, $\mathcal{D} = (\mathcal{F}_{\tilde{o}}(\tilde{v}, \tilde{Y}^{(1)}), \prec_{\tilde{o}})$ is the domain of a regular event structure \mathcal{E} . As explained in Subsection 4.3, the events of \mathcal{E} are the hyperplanes of \mathcal{D} . Hyperplanes \tilde{H} and \tilde{H}' are concurrent if and only if they cross, and $\tilde{H} \leq \tilde{H}'$ if and only if $\tilde{H} = \tilde{H}'$ or \tilde{H} separates \tilde{H}' from \tilde{v} . The events \tilde{H} and \tilde{H}' are in conflict iff \tilde{H} and \tilde{H}' do not cross and neither separates the other from v . Note that this implies that $\tilde{H} \prec \tilde{H}'$ iff \tilde{H} separates \tilde{H}' from v and \tilde{H} and \tilde{H}' osculte, and $\tilde{H} \#_{\mu} \tilde{H}'$ iff \tilde{H} and \tilde{H}' osculte and neither of \tilde{H} and \tilde{H}' separates the other from v . Notice also that each hyperplane \tilde{H}' of \mathcal{D} is the intersection of a hyperplane \tilde{H} of \tilde{Y} with \mathcal{D} .

Claim 7.9. $\tilde{\lambda}$ is a regular trace labeling of \mathcal{D} with the trace alphabet (Σ, I) .

First note that if \tilde{e}_1, \tilde{e}_2 are opposite edges of a square of \mathcal{D} , then $e_1 = \varphi(\tilde{e}_1)$ and $e_2 = \varphi(\tilde{e}_2)$ are opposite edges of a square of Y and thus $\tilde{\lambda}(\tilde{e}_1) = \lambda(e_1) = \lambda(e_2) = \tilde{\lambda}(\tilde{e}_2)$. Consequently, $\tilde{\lambda}$ is a labeling of the edges of \mathcal{D} . Since each labeling is a coloring, from Lemma 6.3, \mathcal{D} has at most

$|V(Y)|$ isomorphism types of labeled principal filters. Therefore, in order to show that $\tilde{\lambda}$ is a regular trace labeling of \mathcal{D} , we just need to show that $\tilde{\lambda}$ satisfies the conditions (LES1), (LES2), and (LES3).

For any two hyperplanes \tilde{H}_1, \tilde{H}_2 in minimal conflict in \mathcal{D} , there exist an edge \tilde{e}_1 dual to \tilde{H}_1 and an edge \tilde{e}_2 dual to \tilde{H}_2 such that \tilde{e}_1 and \tilde{e}_2 have the same source \tilde{u} . Note that since \tilde{H}_1 and \tilde{H}_2 are in conflict, \tilde{e}_1 and \tilde{e}_2 do not belong to a common square of \mathcal{D} . Moreover, if \tilde{e}_1 and \tilde{e}_2 are in a square \tilde{Q} in \tilde{Y} , then since there is a directed path from \tilde{v} to \tilde{u} , and since \tilde{u} is the source of \tilde{Q} , all vertices of \tilde{Q} are in $(\mathcal{F}_{\tilde{o}}(\tilde{v}, \tilde{Y}^{(1)}), \prec_{\tilde{o}}) = \mathcal{D}$. Consequently, the hyperplanes \tilde{H}_1 and \tilde{H}_2 osculate at $(\tilde{u}, \tilde{e}_1, \tilde{e}_2)$ in \tilde{Y} . Let $u = \varphi(\tilde{u})$, $e_1 = \varphi(\tilde{e}_1)$, and $e_2 = \varphi(\tilde{e}_2)$, and note that u is the source of e_1 and e_2 . Let H_1 and H_2 be the hyperplanes of Y that are respectively dual to e_1 and e_2 . Since φ is a covering map, e_1 and e_2 do not belong to a common square. Therefore $\lambda(e_1) \neq \lambda(e_2)$ and $(\lambda(e_1), \lambda(e_2)) \notin I$. Since $\tilde{\lambda}(\tilde{H}_1) = \lambda(e_1)$ and $\tilde{\lambda}(\tilde{H}_2) = \lambda(e_2)$, this establishes (LES1). This also establishes (LES2) when $\tilde{H}_1 \#_{\mu} \tilde{H}_2$.

Suppose now that $\tilde{H}_1 < \tilde{H}_2$ in \mathcal{D} . There exist an edge \tilde{e}_1 dual to \tilde{H}_1 and an edge \tilde{e}_2 dual to \tilde{H}_2 such that the sink \tilde{u} of \tilde{e}_1 is the source of \tilde{e}_2 . Since \tilde{H}_1 separates \tilde{H}_2 from \tilde{v} in \mathcal{D} , \tilde{H}_1 also separates \tilde{H}_2 from \tilde{v} in \tilde{Y} . Consequently, \tilde{e}_1 and \tilde{e}_2 do not belong to a common square of \tilde{Y} and the hyperplanes \tilde{H}_1 and \tilde{H}_2 osculate at $(\tilde{u}, \tilde{e}_1, \tilde{e}_2)$. Let $u = \varphi(\tilde{u})$, $e_1 = \varphi(\tilde{e}_1)$, and $e_2 = \varphi(\tilde{e}_2)$, and note that u is the sink of e_1 and the source of e_2 . Since φ is a covering map, e_1 and e_2 do not belong to a common square and therefore $(\lambda(e_1), \lambda(e_2)) \notin I$. Since $\tilde{\lambda}(\tilde{H}_1) = \lambda(e_1)$ and $\tilde{\lambda}(\tilde{H}_2) = \lambda(e_2)$, this establishes (LES2) when $\tilde{H}_1 < \tilde{H}_2$.


We prove (LES3) by contraposition. Consider two hyperplanes \tilde{H}_1, \tilde{H}_2 that are concurrent, i.e., they intersect in \mathcal{D} . Since \tilde{H}_1 and \tilde{H}_2 intersect in \tilde{Y} , there exists a square \tilde{Q} containing two consecutive edges \tilde{e}_1, \tilde{e}_2 that are respectively dual to \tilde{H}_1, \tilde{H}_2 . Let H_1 and H_2 be the hyperplanes of Y that are respectively dual to $e_1 = \varphi(\tilde{e}_1)$ and $e_2 = \varphi(\tilde{e}_2)$. Note that $\tilde{\lambda}(\tilde{e}_1) = \lambda(e_1)$ and $\tilde{\lambda}(\tilde{e}_2) = \lambda(e_2)$. Since φ is a covering map, e_1 and e_2 belong to a square in Y . Therefore $(\tilde{\lambda}(\tilde{H}_1), \tilde{\lambda}(\tilde{H}_2)) = (\lambda(e_1), \lambda(e_2)) \in I$, establishing (LES3).

8. 1-SAFE PETRI NETS AND SPECIAL CUBE COMPLEXES

8.1. The results. In this section we present the first main result of the paper, namely we show that to any net system $N = (S, \Sigma, F, m_0)$ one can associate a finite directed special cube complex $\mathbf{X}_N = (X_N, o)$ with a trace labeling $\lambda_N : E(X_N) \rightarrow \Sigma$ such that the domain $\mathcal{D}(\mathcal{E}_N)$ of the event structure unfolding \mathcal{E}_N of N is a principal filter of the universal cover $\tilde{\mathbf{X}}_N$ of \mathbf{X}_N .

Let $N = (S, \Sigma, F, m_0)$ be a net system. The transition relation $\longrightarrow \subseteq 2^S \times \Sigma \times 2^S$ defines a directed graph whose vertices are all markings of N and there is an arc from a marking m to a marking m' iff there exists a transition $a \in \Sigma$ such that $m \xrightarrow{a} m'$ (i.e., $\bullet a \subseteq m$, $(a \bullet - \bullet a) \cap m = \emptyset$, and $m' = (m - \bullet a) \cup a \bullet$). Denote by G_N the connected component of the support of this graph that contains the initial marking m_0 and call the undirected graph G_N the *marking graph* of N . Let $\vec{G}_N = (G_N, o)$ denote G_N whose edges are oriented according to \longrightarrow (for notational conveniences we use o instead of \longrightarrow) and call \vec{G}_N the *directed marking graph*. The marking graph G_N contains all markings reachable from m_0 but also it may contain other markings. Notice also that the directed marking graph \vec{G}_N is deterministic and codeterministic, i.e., for any vertex m and any transition $a \in \Sigma$ there exists at most one arc $m \xrightarrow{a} m'$ and at most one arc $m'' \xrightarrow{a} m$. We will say that two distinct transitions $a, b \in \Sigma$ are *independent* if $(\bullet a \cup a \bullet) \cap (\bullet b \cup b \bullet) = \emptyset$. Consider the trace alphabet (Σ, I) where $(a, b) \in I$ if and only if the transitions a and b are independent.

Definition 8.1. The 2-dimensional *cube complex* X_N of N is defined in the following way. The 0-cubes and the 1-cubes of X_N are the vertices and the edges of the marking graph G_N . A 4-cycle (m, m_1, m', m_2) of G_N defines a square of X_N iff there exist two (necessarily distinct) independent transitions $a, b \in \Sigma$ such that $m \xrightarrow{a} m_1, m \xrightarrow{b} m_2, m_1 \xrightarrow{b} m',$ and $m_2 \xrightarrow{a} m'$.

FIGURE 3. The special cube complex of the net system N^* .

The cube complex X_N can be transformed into a directed and colored cube complex $\mathbf{X}_N = (X_N, o, \lambda_N)$: an edge mm' of G_N is oriented from m to m' and $\lambda_N(mm') = a$ iff $m \xrightarrow{a} m'$ holds (clearly, Σ is the set of colors).

Theorem 8.2. *(X_N, o) is a finite directed special cube complex with two-sided hyperplanes and λ_N is a trace labeling of X_N with the trace alphabet (Σ, I) .*

By Lemma 4.11, X_N can be completed in a canonical way to a NPC complex that is also special. In the following, we will denote this completion also by X_N .

Example 8.3. The special cube complex X_{N^*} of the net system N^* from Example 3.1 is representend in Figure 3. In the figure, the leftmost vertices should be identified with the rightmost vertices that are on the same line. The leftmost edges are identified with the righmost edges that have the same label. Similarly, the lower vertices and edges should be identified with the upper vertices and edges. The dotted vertex (that appears on each corner of the figure) correspond to the initial marking of N^* described in Figure 1.

The complex X_{N^*} has 8 vertices, 32 edges, and 24 squares. A 4-cycle in the figure is a square of X_{N^*} if opposite edges have the same label (and direction) and if the labels appearing on the edges of the square correspond to independent transitions of N^* . For example, on the right bottom corner of the figure, the directed 4-cycle labeled by h_4 and v_1 is not a square of X_{N^*} because the transitions h_4 and v_1 are not independent (as explained in Example 3.1).

In the figure, the number (2 or 4) in the middle of each 4-cycle represent the number of squares of X_{N^*} on the vertices of this 4-cycle.

Let \tilde{X}_N denotes the universal cover of the special cube complex X_N and let $\varphi : \tilde{X}_N \rightarrow X_N$ denotes the covering map. Let $\tilde{\mathbf{X}}_N = (\tilde{X}_N, \tilde{o}, \tilde{\lambda}_N)$ be the directed colored CAT(0) cube complex, in which the orientation and the coloring are defined as in Section 7. For any lift \tilde{m}_0 of m_0 , denote by $\mathcal{E}_{X_N} = (E', \leq', \#, \tilde{\lambda}_N)$ the Σ -labeled event structure whose domain is the principal filter $\mathcal{F}_{\tilde{o}}(\tilde{m}_0, \tilde{X}_N^{(1)})$. Finally, let $\mathcal{E}_N = (E, \leq, \#, \lambda)$ be the event structure unfolding of N as defined in Subsection 3.4 and denote by $\mathcal{D}(\mathcal{E}_N)$ the domain of \mathcal{E}_N . The main result of this section is the following theorem:

Theorem 8.4. *The event structures $\mathcal{E}_N = (E, \leq, \#, \lambda)$ and $\mathcal{E}_{X_N} = (E', \leq', \#, \tilde{\lambda}_N)$ are isomorphic.*

Using Thiagarajan's characterization of trace regular event structures (Theorem 3.2), we establish the converse of Theorem 8.2.

Proposition 8.5. *For any finite (virtually) special cube complex X , any admissible orientation o of X , and any vertex \tilde{v} in the universal cover \tilde{X} of X , there exists a finite net system N such that the domain of the event structure \mathcal{E}_N is isomorphic to the principal filter $(\mathcal{F}_{\tilde{o}}(\tilde{v}_0, \tilde{X}^{(1)}), \leq_{\tilde{o}})$.*

By Theorem 8.2 and Proposition 8.5, we obtain a correspondence between trace regular event structures and special cube complexes, leading to the following corollary:

Corollary 8.6. *Any trace regular event structure is cover-special, and thus strongly regular.*

Remark 8.7. In [13], the following question was formulated: Is it true that any regular event structure is strongly regular?

In view of Corollary 8.6, if the answer to this question is negative, this would provide automatically other counterexamples to Thiagarajan's Conjecture 3.3, as the counterexamples provided in [13] are strongly regular event structures that are not trace regular.

8.2. Proof of Theorem 8.2. We will say that a square Q of X_N is an $\{a, b\}$ -square if two opposite edges of Q are labeled a and two other opposite edges are labeled b . Observe that by the definition of squares of X_N , if Q is an $\{a, b\}$ -square, then necessarily $(a, b) \in I$.

Each square Q of X_N has a unique source (a vertex m of Q whose two incident edges are directed from m) and a unique sink (a vertex m' of Q whose two incident edges are directed to m'). We restate the definition of squares of X_N in the following way:

Claim 8.8. *A vertex m of X_N is the source of an $\{a, b\}$ -square $Q = (m, m_1, m', m_2)$ of X_N iff $\bullet a \cup \bullet b \subseteq m$, $((a^\bullet - \bullet a) \cup (b^\bullet - \bullet b)) \cap m = \emptyset$, and $(\bullet a \cup a^\bullet) \cap (\bullet b \cup b^\bullet) = \emptyset$. In this case, $m_1 = m - \bullet a + a^\bullet$, $m_2 = m - \bullet b + b^\bullet$, $m' = m_1 - \bullet b + b^\bullet = m_2 - \bullet a + a^\bullet$ and $m \xrightarrow{a} m_1$, $m \xrightarrow{b} m_2$, $m_1 \xrightarrow{b} m'$, $m_2 \xrightarrow{a} m'$.*

From Claim 8.8 it follows that $\lambda_N(mm_1) = \lambda_N(m_2m')$ and $\lambda_N(mm_2) = \lambda_N(m_1m')$, therefore the coloring λ_N is a labeling of the directed complex (X_N, o) .

A transition $a \in \Sigma$ is called *degenerated* if $\bullet a = a^\bullet$.

Claim 8.9. *$a \in \Sigma$ is degenerated if for any arc $m \xrightarrow{a} m'$ of \vec{G}_N , we have $m = m'$, i.e., $m \xrightarrow{a} m'$ is a loop. Moreover, if a is degenerated, then $m \xrightarrow{a} m$ iff $\bullet a \subseteq m$.*

Proof. If a is degenerated and $m \xrightarrow{a} m'$, since $a^\bullet = \bullet a \subseteq m$ we conclude that $m' = (m - \bullet a) \cup a^\bullet = m$. Conversely, if $m \xrightarrow{a} m'$ is a loop, then the conditions $\bullet a \subseteq m$, $(a^\bullet - \bullet a) \cap m = \emptyset$, and $m = (m - \bullet a) \cup a^\bullet$ imply that $\bullet a = a^\bullet$. The second assertion trivially follows from the definition of $m \xrightarrow{a} m'$. \square

Notice that an $\{a, b\}$ -square (m, m_1, m', m_2) of X_N is either non-degenerated (its four vertices are pairwise distinct) or one of the transitions a or b is degenerated and another one not (if a is degenerated, then $m = m_1$ and $m_2 = m'$) or both transitions a and b are degenerated (in this case, $m = m_1 = m' = m_2$ and $\bullet a = a^\bullet = \bullet b = b^\bullet$).

For $a \in \Sigma$, a hyperplane H of X_N is called an *a -hyperplane* (or a hyperplane of type a) if all edges mm' dual to H are labeled a . An a -hyperplane H such that a is a degenerated transition is called a *degenerated hyperplane*. By Claim 8.9, a degenerated a -hyperplane can be viewed as a connected component of the subgraph of the marking graph G_N induced by all markings m having a loop labeled a .

Claim 8.10. *The hyperplanes of X_N are two-sided.*

Proof. Any degenerated hyperplane can always be viewed as a two-sided hyperplane. Now suppose that H is a non-degenerated a -hyperplane. If H is not two-sided, then there exists an edge mm' dual to H such that $m \xrightarrow{a} m'$ and $m' \xrightarrow{a} m$. Consequently, $\bullet a \subseteq m$, $(a^\bullet - \bullet a) \cap m = \emptyset$, $m' = (m - \bullet a) \cup a^\bullet$ and $\bullet a \subseteq m'$, $(a^\bullet - \bullet a) \cap m' = \emptyset$, $m = (m' - \bullet a) \cup a^\bullet$. Since $m = (m' - \bullet a) \cup a^\bullet$, $a^\bullet \subseteq m$. Since $(a^\bullet - \bullet a) \cap m = \emptyset$, we conclude that $a^\bullet \subseteq \bullet a$. If there exists $e \in \bullet a - a^\bullet$, then $e \notin m' = (m - \bullet a) \cup a^\bullet$, contrary to the assumption that $\bullet a \subseteq m'$. As a result, we deduce that $\bullet a = a^\bullet$, i.e., a is a degenerated transition, contrary to the choice of the hyperplane H . \square

Claim 8.11. *λ_N is a trace labeling of (X_N, o) .*

Proof. From the definition of the squares of X_N it follows that λ_N satisfies (TL1).

To prove (TL2), let m be a vertex of X_N with two outgoing edges $\overrightarrow{mm_1}, \overrightarrow{mm_2}$. Since \overrightarrow{G}_N is deterministic, $\lambda_N(mm_1) \neq \lambda_N(mm_2)$, say $\lambda_N(mm_1) = a$ and $\lambda_N(mm_2) = b$. We assert that $(a, b) \in I$ iff $\overrightarrow{mm_1}$ and $\overrightarrow{mm_2}$ belong to a common square of X_N . One direction is immediate: if $\overrightarrow{mm_1}$ and $\overrightarrow{mm_2}$ belong to a square of X_N , by the definition of the squares of X_N , we have $(a, b) = (\lambda_N(mm_1), \lambda_N(mm_2)) \in I$. Conversely, suppose that $(a, b) \in I$. Then $(\bullet a \cup a \bullet) \cap (\bullet b \cup b \bullet) = \emptyset$, i.e., the transitions a and b are independent. Since $m \xrightarrow{a} m_1$ and $m \xrightarrow{b} m_2$, we conclude that also $\bullet a \cup \bullet b \subseteq m$ and $((a \bullet - \bullet a) \cup (b \bullet - \bullet b)) \cap m = \emptyset$ hold. From Claim 8.8 we deduce that m is the source of an $\{a, b\}$ -square in which two neighbors of m are m_1 and m_2 . This establishes (TL2).

To prove (TL3), let $m \xrightarrow{a} m_1$ and $m_1 \xrightarrow{b} m'$ be two distinct arcs. Then we have to prove that $(a, b) \in I$ iff $\overrightarrow{mm_1}$ and $\overrightarrow{m_1m'}$ belong to a square of X_N . If $\overrightarrow{mm_1}$ and $\overrightarrow{m_1m'}$ belong to a square of X_N , then $\overrightarrow{mm_1}$ and $\overrightarrow{m_1m'}$ are consecutive edges of that square, and by the definition of the squares of X_N , this implies that $(a, b) \in I$. Conversely, let $(a, b) \in I$. By the definition of I , the transitions a and b are independent. Set $m_2 = m - \bullet b + b \bullet$. Suppose by way of contradiction that (m, m_1, m', m_2) is not a square of X_N . By Claim 8.8 and since $(\bullet a \cup a \bullet) \cap (\bullet b \cup b \bullet) = \emptyset$, either $\bullet a \cup \bullet b \not\subseteq m$ or $((a \bullet - \bullet a) \cup (b \bullet - \bullet b)) \cap m \neq \emptyset$. Since $\bullet a \subseteq m$ and $(a \bullet - \bullet a) \cap m = \emptyset$ because $m \xrightarrow{a} m_1$, we deduce that either $\bullet b \not\subseteq m$ or $(b \bullet - \bullet b) \cap m \neq \emptyset$. If $\bullet b \not\subseteq m$, since $(\bullet a \cup a \bullet) \cap \bullet b = \emptyset$ and $m_1 = m - \bullet a + a \bullet$, we deduce that $\bullet b \not\subseteq m_1$, contrary to the assumption that $m_1 \xrightarrow{b} m'$. On the other hand, if there exists $e \in (b \bullet - \bullet b) \cap m$, since $(b \bullet - \bullet b) \cap m_1 = \emptyset$ and $m_1 = m - \bullet a + a \bullet$, we conclude that $e \in \bullet a - a \bullet$. Hence $e \in \bullet a \cap b \bullet$, contrary to the assumption that a and b are independent. This proves that m_2 is an admissible marking and that (m, m_1, m', m_2) is a square of X_N , establishing (TL3).

Finally, we will establish (TL4). Let m' be a vertex of X_N and let $m_1 \xrightarrow{a} m'$ and $m_2 \xrightarrow{b} m'$ be two distinct arcs. Since \overrightarrow{G}_N is codeterministic, $a \neq b$. To prove (TL4) we have to show that $(a, b) \in I$ iff $\overrightarrow{m_1m'}$ and $\overrightarrow{m_2m'}$ belong to a common square of X_N . Again, one direction directly follows from the definition of the squares of X_N . Conversely, suppose that $(a, b) \in I$, i.e., $(\bullet a \cup a \bullet) \cap (\bullet b \cup b \bullet) = \emptyset$. Hence $(a \bullet - \bullet a) \cap (b \bullet - \bullet b) = \emptyset$. Set $m := m' - (a \bullet \cup b \bullet) + (\bullet a \cup \bullet b)$. We assert that (m, m_1, m', m_2) is a square of X_N with source m . Since the transitions a and b are independent, by Claim 8.8 it suffices to show that $\bullet a \cup \bullet b \subseteq m$ and $((a \bullet - \bullet a) \cup (b \bullet - \bullet b)) \cap m = \emptyset$. Both these properties directly follow from the definition of m . This establishes (TL4). \square

Theorem 8.2 now follows from Claims 8.10 and 8.11 and Theorem 7.5.

8.3. Proof of Theorem 8.4 and Proposition 8.5. Let $N = (S, \Sigma, F, m_0)$ be a net system. As above, FN denotes the set of all firing sequences at m_0 , i.e., all words $\sigma \in \Sigma^*$ for which there exists a marking m such that $m_0 \xrightarrow{\sigma} m$. The trace alphabet associated to N is the pair $M = (\Sigma, I)$ where $(a, b) \in I$ iff $(a \bullet \cup \bullet a) \cap (b \bullet \cup \bullet b) = \emptyset$. We called above the traces of the form $\langle \sigma \rangle$ for $\sigma \in FS$ firing traces and denoted them by $\mathcal{FT}(N)$. We also denoted by $\mathcal{PFT}(N)$ the firing traces which are prime. Let $\mathcal{E}_N = (E, \leq, \#, \lambda)$ be the M -labeled event structure unfolding of a net system N . Recall that the events of \mathcal{E}_N are the prime firing traces from $\mathcal{PFT}(N)$ and the label of an event $\langle \sigma \rangle$ is $\lambda(\langle \sigma \rangle) = \text{last}(\sigma)$. The precedence and the conflict relations in \mathcal{E}_N have been defined above.

Let also $\mathcal{E}_{X_N} = (E', \leq', \#', \tilde{\lambda}_N)$ be the Σ -labeled event structure whose domain is the principal filter $\mathcal{F}_{\tilde{o}}(\tilde{m}_0, \tilde{X}_N^{(1)})$ of the universal cover $\tilde{\mathbf{X}}_N = (\tilde{X}_N, \tilde{o}, \tilde{\lambda}_N)$ of the special cube complex (X_N, o, λ_N) of N . Let $\varphi : \tilde{X}_N \mapsto X_N$ denote the covering map. Let $G(\mathcal{E}_{X_N})$ denotes the median graph of \mathcal{E}_{X_N} . From Theorem 8.2 and Proposition 7.6 it follows that $\tilde{\lambda}_N$ is a trace labeling of the event structure \mathcal{E}_{X_N} . By Corollary 5.4 there exists a bijection between the set of configurations of \mathcal{E}_{X_N} and the set $\mathcal{GT}(\mathcal{E}_{X_N})$ of geodesic traces of \mathcal{E}_{X_N} . By Corollary 5.9 there exists a bijection $H \mapsto \langle \sigma_v \rangle$ between the set of hyperplanes (events) of \mathcal{E}_{X_N} and the set $\mathcal{PGT}(\mathcal{E}_{X_N})$ of prime geodesic traces of \mathcal{E}_{X_N} .

The following claim establishes a bijection between geodesic traces of \mathcal{E}_{X_N} and firing traces of N .

Claim 8.12. *Any geodesic trace $\langle \sigma_{\tilde{m}} \rangle$ of \mathcal{E}_{X_N} is a firing trace of N . Conversely, for any firing trace $\langle \sigma \rangle$ there exists a geodesic trace $\langle \sigma_{\tilde{m}} \rangle$ such that $\langle \sigma \rangle = \langle \sigma_{\tilde{m}} \rangle$.*

In particular, there is a bijection between prime geodesic traces of \mathcal{E}_{X_N} and the prime firing traces of N .

Proof. Each firing sequence σ of N corresponds to a directed path in the directed marking graph \vec{G}_N : if $\sigma = a_1 \dots a_k \in FS$ is a firing sequence, then there exists reachable markings m_1, \dots, m_{k+1} such that $\pi(\sigma) := m_0 \xrightarrow{a_1} m_1 \xrightarrow{a_2} \dots \xrightarrow{a_{k-1}} m_k \xrightarrow{a_k} m_{k+1}$ is a directed (m_0, m_{k+1}) -path of \vec{G}_N . Since G_N is the 1-skeleton of the special cube complex X_N , the directed universal cover $\tilde{\mathbf{X}}_N$ of X_N contains a directed path $\tilde{\pi}(\sigma)$ from \tilde{m}_0 to \tilde{m}_{k+1} whose image under the covering map φ is $\pi(\sigma)$. By Lemma 6.2, $\tilde{\pi}(\sigma)$ is a shortest $(\tilde{m}_0, \tilde{m}_{k+1})$ -path in the 1-skeleton of $\tilde{\mathbf{X}}_N$. Let σ' be another firing sequence such that there exists $(a_i, a_{i+1}) \in I$ such that $\sigma' = a_1 \dots a_{i-1} a_{i+1} a_i a_{i+2} \dots a_k$. Since the transitions a_i and a_{i+1} are independent, there exists m'_{i+1} such that $m_i \xrightarrow{a_{i+1}} m'_{i+1} \xrightarrow{a_i} m_{i+2}$ and $(m_i, m_{i+1}, m_{i+2}, m'_{i+1})$ is a square of X_N . Then σ' corresponds to the directed path $\pi(\sigma') = m_0 \xrightarrow{a_1} m_1 \xrightarrow{a_2} \dots \xrightarrow{a_{i-1}} m_i \xrightarrow{a_{i+1}} m'_{i+1} \xrightarrow{a_i} m_{i+2} \dots \xrightarrow{a_{k-1}} m_k \xrightarrow{a_k} m_{k+1}$ of \vec{G}_N . Analogously to $\pi(\sigma)$, $\tilde{\mathbf{X}}_N$ also contains a directed path $\tilde{\pi}(\sigma')$ with origin \tilde{m}_0 whose φ -image is $\pi(\sigma')$. Moreover, $(\tilde{m}_i, \tilde{m}_{i+1}, \tilde{m}_{i+2}, \tilde{m}'_{i+1})$ is a square of $\tilde{\mathbf{X}}_N$, thus $\tilde{\pi}(\sigma')$ can be obtained from $\tilde{\pi}(\sigma)$ by an elementary homotopy with respect to this square. As a conclusion of those two properties we obtain that the firing trace $\langle \sigma \rangle$ of N is contained in the geodesic trace $\langle \sigma_{\tilde{m}_{k+1}} \rangle$ of \mathcal{E}_{X_N} .

It remains to prove the converse inclusion, i.e., that any geodesic trace $\langle \sigma_{\tilde{m}} \rangle$ of \mathcal{E}_{X_N} is contained in a firing trace. Let $\tilde{\pi}$ be a shortest (\tilde{m}_0, \tilde{m}) -path in the graph $G(\mathcal{E}_{X_N})$. Then the edges of $\tilde{\pi}$ are directed from \tilde{m}_0 to \tilde{m} . The image of $\tilde{\pi}$ by the covering map φ is a directed (m_0, m) -path π in the graph \vec{G}_N . This path is not necessarily shortest or simple, however the words defined by the labels of edges of $\tilde{\pi}$ and π coincide: $\sigma(\tilde{\pi}) = \sigma(\pi)$. Since π is a directed (m_0, m) -path in the marking graph, necessarily $\sigma(\pi)$ is a firing sequence, yielding that $\sigma(\tilde{\pi}) \in FS$. By Lemma 5.3, the geodesic trace $\langle \sigma(\tilde{\pi}) \rangle$ consists exactly of all $\sigma(\tilde{\pi}')$ such that $\tilde{\pi}'$ is a shortest (\tilde{m}_0, \tilde{m}) -path. By Lemma 5.1, the paths $\tilde{\pi}$ and $\tilde{\pi}'$ are homotopic, i.e., there exists a finite sequence $\tilde{\pi} =: \tilde{\pi}_1, \tilde{\pi}_2, \dots, \tilde{\pi}_{k-1}, \tilde{\pi}_k := \tilde{\pi}'$ of shortest (\tilde{m}_0, \tilde{m}) -paths such that for any $i = 1, \dots, k-1$ the paths $\tilde{\pi}_i$ and $\tilde{\pi}_{i+1}$ differ only in a square $\tilde{Q}_i = (\tilde{m}_{j-1}, \tilde{m}_j, \tilde{m}_{j+1}, \tilde{m}'_j)$ of $\tilde{\mathbf{X}}_N$. Let π_i denote the image of the path $\tilde{\pi}_i$ under the covering map φ . Let also $Q_i := (m_{j-1}, m_j, m_{j+1}, m'_j) = \varphi(\tilde{Q}_i)$. Each π_i is a directed (m_0, m) -path of \vec{G}_N and each Q_i is a square of X_N . Moreover, $\sigma(\tilde{\pi}_i) = \sigma(\pi_i)$ for each i and the edges of the squares \tilde{Q}_i and Q_i are labeled in the same way. Each π_{i+1} is obtained from π_i by an elementary homotopy with respect to the square Q_i . From the definition of the squares of X_N it follows that there exists $(a_j, a_{j+1}) \in I$ such that $m_{j-1} \xrightarrow{a_j} m_j \xrightarrow{a_{j+1}} m_{j+1}$ and $m_{j-1} \xrightarrow{a_{j+1}} m'_j \xrightarrow{a_j} m_{j+1}$. But this implies that $\sigma(\pi_{i+1})$ is obtained from $\sigma(\pi_i)$ by exchanging a_j with a_{j+1} , yielding that $\sigma(\pi_{i+1})$ belongs to the trace of $\sigma(\pi_i)$. Since all $\sigma(\pi_i)$ are firing sequences of FN , this implies that they all belong to the firing trace of $\sigma(\pi)$. Since $\sigma(\tilde{\pi}) = \sigma(\pi)$, we conclude that the geodesic trace of $\sigma(\tilde{\pi})$ is included in the firing trace of $\sigma(\tilde{\pi})$. This concludes the proof of the equality between geodesic traces and firing traces.

Observe that if $\langle \sigma_{\tilde{m}} \rangle$ is a prime geodesic trace, then the interval $I(\tilde{m}_0, \tilde{m})$ is prime by Lemma 5.7. Since each $\sigma'_{\tilde{m}} \in \langle \sigma_{\tilde{m}} \rangle$ corresponds to a shortest $(\tilde{m}_0, \tilde{m}_k)$ -path, all such paths share the same last edge. Consequently, all the words in $\langle \sigma_{\tilde{m}} \rangle$ have the same last letter, and thus the corresponding firing trace is prime. Conversely, for any prime firing trace $\langle \sigma \rangle$, let \tilde{m} be the vertex of $\tilde{\mathbf{X}}_N$ such that $\langle \sigma \rangle = \langle \sigma_{\tilde{m}} \rangle$. Since $\langle \sigma \rangle$ is prime, all words in $\langle \sigma \rangle = \langle \sigma_{\tilde{m}} \rangle$ have the same last letter. Since two incoming arcs of \tilde{m} have different labels, this implies that \tilde{m} has only one incoming arc, i.e., the interval $I(\tilde{m}_0, \tilde{m})$ is prime. By Lemma 5.7, the geodesic trace $\langle \sigma_{\tilde{m}} \rangle$ is prime. \square

Claim 8.12 establishes a bijection between prime geodesic traces of \mathcal{E}_{X_N} and prime firing traces of N . Consequently, there is a bijection between the hyperplanes (events) of \mathcal{E}_{X_N} and

the hyperplanes (events) of \mathcal{E}_N and this bijection preserves labels. Therefore, to establish that the event structures \mathcal{E}_N and \mathcal{E}_{X_N} are isomorphic it remains to show that this bijection preserves the precedence and the conflict relations. By Lemma 5.10, for two hyperplanes H', H of \mathcal{E}_{X_N} with prime geodesic traces $\langle \sigma_{\tilde{u}} \rangle$ and $\langle \sigma_{\tilde{v}} \rangle$, respectively, we have $H' \leq H$ iff $\langle \sigma_{\tilde{u}} \rangle \sqsubseteq \langle \sigma_{\tilde{v}} \rangle$. On the other hand, for \mathcal{E}_N the precedence relation \leq is the prefix relation \sqsubseteq . Therefore the bijection between the events of \mathcal{E}_{X_N} and the events of \mathcal{E}_N preserves the precedence relation.

Finally, we will show that this bijection also preserves the conflict relation. Taking into account the bijection between firing traces and geodesic traces from Claim 8.12, the definition of the conflict relation in \mathcal{E}_N can be rephrased in the following way: two prime firing traces $\langle \sigma_{\tilde{u}} \rangle$ and $\langle \sigma_{\tilde{v}} \rangle$ are in conflict if and only there does not exist a firing trace $\langle \sigma_{\tilde{w}} \rangle$ such that $\langle \sigma_{\tilde{u}} \rangle$ and $\langle \sigma_{\tilde{v}} \rangle$ are prefixes of $\langle \sigma_{\tilde{w}} \rangle$. By Lemma 5.5, $\langle \sigma_{\tilde{u}} \rangle \sqsubseteq \langle \sigma_{\tilde{w}} \rangle$ iff $\tilde{u} \in I(\tilde{m}_0, \tilde{w})$. Consequently, two prime firing traces $\langle \sigma_{\tilde{u}} \rangle$ and $\langle \sigma_{\tilde{v}} \rangle$ are in conflict if and only if there does not exist a vertex \tilde{w} such that $\tilde{u}, \tilde{v} \in I(\tilde{m}_0, \tilde{w})$. By Lemma 5.6, there does not exist a vertex \tilde{w} such that $\tilde{u}, \tilde{v} \in I(\tilde{m}_0, \tilde{w})$ iff the prime geodesic traces $\langle \sigma_{\tilde{u}} \rangle$ and $\langle \sigma_{\tilde{v}} \rangle$ are in conflict in \mathcal{E}_{X_N} . This proves that the bijection between geodesic traces and firing traces preserves the conflict relations in \mathcal{E}_{X_N} and \mathcal{E}_N and finishes the proof of Theorem 8.4.

We conclude with the proof of Proposition 8.5. Consider a finite virtually special cube complex X with an admissible orientation o' and let Y be a finite special cover of X and o be the orientation of Y lifted from o' . Since the directed universal cover (\tilde{X}, \tilde{o}') of (X, o') is isomorphic to the directed universal cover (\tilde{Y}, \tilde{o}) of (Y, o) , it is enough to prove Proposition 8.5 for the finite special complex Y and the orientation o . By Theorem 7.5, there exists a trace labeling λ of (Y, o) . By Proposition 7.6, for any vertex \tilde{v}_0 of \tilde{Y} , the lift $\tilde{\lambda}$ of λ is a regular trace labeling of the principal filter $(\mathcal{F}_{\tilde{o}}(\tilde{v}, \tilde{Y}^{(1)}), \prec_{\tilde{o}})$ of (\tilde{Y}, \tilde{o}) . Hence, $(\mathcal{F}_{\tilde{o}}(\tilde{v}, \tilde{Y}^{(1)}), \prec_{\tilde{o}})$ is the domain of a trace regular event structure \mathcal{E} . By Thiagarajan's Theorem 3.2, there exists a net system N such that \mathcal{E} is isomorphic to \mathcal{E}_N . This ends the proof of Proposition 8.5.

9. DECIDABILITY OF THE MSO THEORY OF NET SYSTEMS AND OF THEIR DOMAINS

9.1. The results. Let $\mathcal{E} = (E, \leq, \#, \lambda)$ be a trace regular event structure and let $\mathcal{D}(\mathcal{E})$ denote the domain of \mathcal{E} . Let $G(\mathcal{E})$ denote the undirected covering median graph of $\mathcal{D}(\mathcal{E})$ and $\vec{G}(\mathcal{E}) = (G(\mathcal{E}), o)$ denote the directed graph of $\mathcal{D}(\mathcal{E})$. First we characterize the trace event structures for which the MSO theories of graphs $G(\mathcal{E})$ and $\vec{G}(\mathcal{E})$ are decidable.

Theorem 9.1. *For a trace regular event structure $\mathcal{E} = (E, \leq, \#, \lambda)$, the following conditions are equivalent:*

- (1) $\text{MSO}(\vec{G}(\mathcal{E}))$ is decidable;
- (2) $\text{MSO}_1(G(\mathcal{E}))$ is decidable;
- (3) $\text{MSO}_2(G(\mathcal{E}))$ is decidable;
- (4) $G(\mathcal{E})$ has finite treewidth;
- (5) the clusters of $G(\mathcal{E})$ have bounded diameter;
- (6) $\vec{G}(\mathcal{E})$ is context-free.

Similarly to a question about the decidability of the MSO theory of graphs of domains of trace regular event structures (i.e., of domains of event structure unfoldings of net systems), one can ask a similar question about the decidability of the MSO theory for the graphs (1-skeletons) of the universal covers of the special cube complexes X_N of net systems N . In this case, the following result holds:

Proposition 9.2. *Let $N = (S, \Sigma, F, m_0)$ be a net system, X_N be the special cube complex of N , and let $\vec{G}(\tilde{\mathbf{X}}_N)$ be the 1-skeleton of the directed labeled universal cover of X_N . Then the following conditions are equivalent:*

- (1) $\text{MSO}(\vec{G}(\tilde{\mathbf{X}}_N))$ is decidable;
- (2) $\text{MSO}_2(G(\tilde{\mathbf{X}}_N))$ is decidable;
- (3) $G(\tilde{\mathbf{X}}_N)$ has finite treewidth;

(4) $\vec{G}(\tilde{\mathbf{X}}_N)$ is context-free.

The proof of this result essentially follows from the result by Kuske and Lohrey [31] (see Theorem 9.13 below) that the decidability of the MSO theory of a directed graph \vec{G} of bounded degree whose automorphism group $\text{Aut}(\vec{G})$ has only finitely many orbits on \vec{G} is equivalent to the fact that \vec{G} is context-free and to the fact that its undirected support has finite treewidth. This result cannot be applied to prove Theorem 9.1 because $\text{Aut}(\vec{G}(\mathcal{E}))$ may have an infinite number of orbits (however this is true for $\vec{G}(\tilde{\mathbf{X}}_N)$).

To relate the MSO theory of the graph of the domain of a trace event structure with the MSO theory of the event structure, we introduce the notion of the *hairing* $\dot{\mathcal{E}} = (\dot{E}, \dot{\leq}, \dot{\#})$ of an event structure $\mathcal{E} = (E, \leq, \#)$. To obtain $\dot{\mathcal{E}}$, we add a *hair event* e_c for each configuration c of \mathcal{E} , i.e., $\dot{E} = E \cup E_C$ where $E_C = \{e_c : c \in \mathcal{D}(\mathcal{E})\}$. For any hair event e_c and any event $e \in \dot{E}$, we set $e \dot{\leq} e_c$ if $e \in c$ and $e \not\# e_c$ otherwise.

By the definition of $\dot{\mathcal{E}}$, its domain $\mathcal{D}(\dot{\mathcal{E}})$ is obtained from the domain $\mathcal{D}(\mathcal{E})$ of \mathcal{E} by adding an outgoing pendant edge $\overrightarrow{vw_v}$ to each vertex $v \in V(\mathcal{D}(\mathcal{E}))$. We call $\mathcal{D}(\dot{\mathcal{E}})$ the *hairing* of $\mathcal{D}(\mathcal{E})$. In a similar way we can define the *hairing* of any graph, in particular of the 1-skeleton of any special cube complex.

Proposition 9.3. *For a trace regular event structure $\mathcal{E} = (E, \leq, \#, \lambda)$, the hairing $\dot{\mathcal{E}}$ is also a trace regular event structure.*

Notice that the hair events of $\dot{\mathcal{E}}$ introduce a lot of conflicting events in $\dot{\mathcal{E}}$, and we use them to encode vertex variables as event variables in order to prove the following result:

Theorem 9.4. *For a trace regular event structure $\mathcal{E} = (E, \leq, \#, \lambda)$, $\text{MSO}(\dot{\mathcal{E}})$ is decidable if and only if $\text{MSO}(\vec{G}(\mathcal{E}))$ is decidable. In particular, $\text{MSO}(\dot{\mathcal{E}})$ is decidable if and only if $G(\mathcal{E})$ has finite treewidth.*

Remark 9.5. The condition on the treewidth of $G(\mathcal{E})$ in the previous theorem is independent of the choice of a particular trace labeling of \mathcal{E} . Therefore, one can rephrase the statement of the theorem in the following way: For any trace regular event structure $\mathcal{E} = (E, \leq, \#)$ such that $G(\mathcal{E})$ has bounded (respectively, unbounded) treewidth, $\text{MSO}(\dot{\mathcal{E}})$ is decidable (respectively, undecidable) for any regular trace labeling λ of $\dot{\mathcal{E}}$. Consequently, if there exists a regular trace labeling of $\dot{\mathcal{E}}$ such that $\text{MSO}(\dot{\mathcal{E}})$ is decidable (respectively, undecidable), then $\text{MSO}(\dot{\mathcal{E}})$ is decidable (respectively, undecidable) for all regular trace labelings of $\dot{\mathcal{E}}$.

Since $\text{MSO}(\mathcal{E})$ is a fragment of $\text{MSO}(\dot{\mathcal{E}})$, we obtain the following corollary of Theorem 9.4:

Corollary 9.6. *For any trace regular event structure $\mathcal{E} = (E, \leq, \#, \lambda)$, if $G(\mathcal{E})$ has finite treewidth, then $\text{MSO}(\mathcal{E})$ is decidable.*

9.2. Treewidth. Let $G = (V, E)$ be a simple graph, not necessarily finite. A *tree decomposition* [39] of G is a pair (T, f) , where $T = (V(T), E(T))$ is a tree and $f : V(T) \rightarrow 2^V \setminus \{\emptyset\}$ is a function such that the following holds:

- (i) $\bigcup_{t \in V(T)} f(t) = V$,
- (ii) for every edge $uv \in E$ of G there exists $t \in V(T)$ such that $u, v \in f(t)$,
- (iii) if $t', t'' \in V(T)$ and t lies on the unique path of T from t' to t'' , then $f(t') \cap f(t'') \subseteq f(t)$.

The supremum in $\mathbb{N} \cup \{\infty\}$ of the cardinalities $|f(t)|$, $t \in V(T)$, is called the *width* of the tree decomposition (T, f) . The graph G has *treewidth* $\leq b$ if there exists a tree decomposition of G of width $\leq b$. A graph G has *bounded* (or *finite*) *treewidth* if it has treewidth $\leq b$ for some $b \in \mathbb{N}$. The treewidth represents how close a graph is to a tree from a *combinatorial* point of view.

A graph H is a *minor* of a graph G if a graph isomorphic to H can be obtained from a subgraph G' of G by contracting edges. Equivalently, H is a minor of a connected graph G if G contains a subgraph G' such that there exists a partition of vertices of G' into connected subgraphs $\mathcal{P} = \{P_1, \dots, P_t\}$ and a bijection $f : V(H) \rightarrow \mathcal{P}$ such that if $uv \in E(H)$ then there exists an edge of G' running between the subgraphs $f(u)$ and $f(v)$ of \mathcal{P} (i.e., after contracting

each subgraph $P_i \in \mathcal{P}$ into a single vertex we will obtain a graph containing H as a spanning subgraph). Treewidth does not increase when taking a minor.

Since the treewidth of an $n \times n$ square grid is n , the treewidth of a graph G is always greater than or equal to the size of the largest square grid minor of G . In the other direction, the grid minor theorem by Robertson and Seymour [40] shows that there exists a function f such that the treewidth is at most $f(r)$ where r is the size of the largest square grid minor of G :

Theorem 9.7 ([40]). *A graph G has bounded treewidth if and only if the square grid minors of G have bounded size.*

9.3. Hyperbolicity. Similarly to nonpositive curvature, Gromov hyperbolicity is defined in metric terms. However, as for the CAT(0) property, the hyperbolicity of a CAT(0) cube complex can be expressed in a purely combinatorial way. A metric space (X, d) is δ -hyperbolic [11, 24] if for any four points v, w, x, y of X , $d(v, w) + d(x, y) \leq \max\{d(v, x) + d(w, y), d(v, y) + d(w, x)\} + 2\delta$. A graph $G = (X, E)$ endowed with its standard graph-distance d_G is δ -hyperbolic if the metric space (X, d_G) is δ -hyperbolic. At the difference of treewidth, Gromov hyperbolicity represents how close *metrically* a graph is to a tree.

In case of geodesic metric spaces and graphs, δ -hyperbolicity can be defined in other equivalent ways, e.g., via thin or slim geodesic triangles. For example, a geodesic metric space (X, d) is 2δ -hyperbolic, if all geodesic triangles $\Delta(x, y, z)$ of (X, d) are δ -slim, i.e., for any point u on the side $[x, y]$ the distance from u to $[x, z] \cup [z, y]$ is at most δ . This definition expresses the negative curvature of a geodesic metric space. A metric space (X, d) is *hyperbolic* if there exists $\delta < \infty$ such that (X, d) is δ -hyperbolic. In case of median graphs, i.e., of 1-skeletons of CAT(0) cube complexes, the hyperbolicity can be characterized in the following way:

Lemma 9.8 ([18, 25]). *Let X be a CAT(0) cube complex. Then its 1-skeleton $X^{(1)}$ is hyperbolic if and only if all isometrically embedded square grids are uniformly bounded.*

In Hagen's paper [25, Theorem 7.6], previous lemma is a consequence of another combinatorial characterization of hyperbolicity of CAT(0) cube complexes of bounded degrees. The *crossing graph* $\Gamma(X)$ of a CAT(0) cube complex X has the hyperplanes of X as vertices and pairs of intersecting hyperplanes as edges. We will say that a graph Γ has *thin bicliques* if there exists a natural number n such that any complete bipartite subgraph $K_{p,q}$ of Γ satisfies $p \leq n$ or $q \leq n$.

Theorem 9.9 ([25]). *A CAT(0) cube complex X with bounded degrees is hyperbolic if and only if its crossing graph $\Gamma(X)$ has thin bicliques.*

We call an event structure $\mathcal{E} = (E, \leq, \#)$ and its domain $\mathcal{D}(\mathcal{E})$ *hyperbolic* if $\mathcal{D}(\mathcal{E})$ is isomorphic to a principal filter of a directed CAT(0) cube complex, whose 1-skeleton is hyperbolic. We call an event structure $\mathcal{E} = (E, \leq, \#)$ and its domain $\mathcal{D}(\mathcal{E})$ *strongly hyperbolic regular* if there exists a finite directed NPC complex (X, o) such that \tilde{X} is hyperbolic and \mathcal{D} is a principal filter of $(\tilde{X}^{(1)}, \tilde{o})$. Note that an event structure can be strongly regular and hyperbolic without being strongly regular hyperbolic (see Remark 10.8).

9.4. Context-free graphs. Let G be an edge-labeled graph of uniformly bounded degree and v_0 be an arbitrary root (basepoint) of G . Let $S(v_0, k) = \{x \in V : d_G(v_0, x) = k\}$ denote the sphere of radius k centered at v_0 . A connected component Υ of the subgraph of G induced by $V \setminus S(v_0, k)$ is called an *end* of G . The vertices of $\Upsilon \cap S(v_0, k+1)$ are called *frontier points* and this set is denoted by $C(\Upsilon)$ [35] and called a *cluster*. There exists a bijection between the ends and the clusters: each end contains a unique cluster and conversely, for a cluster C , the unique end $\Upsilon(C)$ containing C consists of the union of all principal filters of the vertices $v \in C$ (with respect to the basepoint order).

Let $\Phi(G)$ and $\mathcal{C}(G)$ denote the set of all ends and all clusters of G , respectively. An *end-isomorphism* between two ends Υ and Υ' of G is a label-preserving mapping f between Υ and Υ' such that f is a graph isomorphism and f maps $C(\Upsilon)$ to $C(\Upsilon')$. Then G is called a *context-free graph* [35] if $\Phi(G)$ has only finitely many isomorphism classes under end-isomorphisms. Since G has uniformly bounded degree, each cluster $C(\Upsilon)$ is finite. Moreover, from the definition

of context-free graphs follows that a context-free graph G has only finitely many isomorphism classes of clusters, thus there exists a constant $\delta < \infty$ such that the diameter of any cluster of G is bounded by δ . By [18, Proposition 12] any graph G whose diameters of clusters is uniformly bounded by δ is δ -hyperbolic (in fact, G is quasi-isometric to a tree). Note that the converse is not true (see the 1-skeleton of the square complex \tilde{Z} described in Section 10.1).

9.5. Some results from MSO theory. In this subsection, we recall some results from MSO theory of undirected graphs, directed labeled graphs, lattices and posets, and event structures. These results either will be used below or are related to our work.

Among the MSO theories of various discrete structures, the MSO theory of undirected graphs is probably the most complete, with various and deep applications. For a comprehensive account of this theory, see the book by Courcelle and Engelfriet [22]. Let $G = (V, E)$ be an undirected and unlabeled graph. The MSO logic as introduced in Subsection 3.5 only allow quantifications over subsets of vertices of G . This theory is usually denoted by $\text{MSO}_1(G)$. In order to allow also quantifications over subsets of edges, an extended representation of a graph is used. This is the relational structure $G^e = (V \cup E, \text{inc})$, where

$$\text{inc} = \{(e, v) \in E \times V : \exists u \in V \text{ such that } e \in \{uv, vu\}\}.$$

The MSO theory of this relational structure G^e is usually denoted by $\text{MSO}_2(G)$. Seese [45] proved the following fundamental result about MSO_2 decidability:

Theorem 9.10 ([45]). *If $\text{MSO}_2(G)$ is decidable, then G has finite treewidth.*

The converse of Seese's theorem is not true: one can construct trees with undecidable MSO_2 theory. On the other, Courcelle [20] proved that for any natural integer k the class of all graphs of treewidth at most k has a decidable MSO_2 theory.

If $\text{MSO}_2(G)$ is decidable, then $\text{MSO}_1(G)$ is also decidable. Again, the reverse implication is not true. However, Courcelle [21] proved that the converse holds for graphs with bounded degrees:

Theorem 9.11 ([21]). *If G is a graph with uniformly bounded degrees and $\text{MSO}_1(G)$ is decidable, then $\text{MSO}_2(G)$ is also decidable.*

Now, consider labeled directed graphs. Let Σ be a finite alphabet. A Σ -labeled directed graph is a relational structure $\vec{G} = (V, (E_a)_{a \in \Sigma})$, where V is the set of vertices and $E_a \subseteq V \times V$ is the set of a -labeled directed edges. Denote by $\text{MSO}(\vec{G})$ the MSO theory of this relational structure. We associate to \vec{G} the unlabeled graph $G = (V, \bigcup_{a \in \Sigma} \{uv : u \neq v, (u, v) \in E_a \text{ or } (v, u) \in E_a\})$.

Müller and Schupp [35] proved the following fundamental theorem about Σ -labeled pointed context-free graphs of bounded degree (and directed according to the basepoint order):

Theorem 9.12 ([35]). *If \vec{G} is a context-free graph, then $\text{MSO}_1(\vec{G})$ is decidable.*

For a directed graph \vec{G} , denote by $\text{Aut}(\vec{G})$ its group of automorphisms. On the vertex set of \vec{G} we define the equivalence relation \sim by $u \sim v$ if there exists $f \in \text{Aut}(\vec{G})$ with $f(u) = v$. Kuske and Lohrey [31] established a kind of converse to Theorem 9.12:

Theorem 9.13 ([31]). *Let \vec{G} be a Σ -labeled connected graph of bounded degree such that $\text{Aut}(\vec{G})$ has only finitely many orbits on \vec{G} . Then the following properties are equivalent:*

- (1) $\text{MSO}(\vec{G})$ is decidable;
- (2) G has finite treewidth;
- (3) \vec{G} is context-free.

The formulation of our Theorem 9.1 is inspired by Theorem 9.13, however the proofs of two results are different.

Kuske [30] characterized the decidability of the MSO logic of distributive lattices. Let $\mathcal{L} = (L, \leq)$ be a distributive lattice. Let $\mathbb{J}(\mathcal{L})$ denote the set of join irreducible elements of \mathcal{L} . $\mathbb{J}(\mathcal{L})$ can be viewed as a poset, endowed with the partial order \leq of \mathcal{L} . An *antichain* is a set of pairwise

incomparable elements. Denote by $w(\mathcal{L})$ the *width* of \mathcal{L} , i.e., the supremum of the cardinalities of its antichains. For a poset (L, \leq) , $\text{MSO}(L, \leq)$ is the MSO theory of the relational structure (L, \leq) .

Theorem 9.14 ([30]). *Let \mathcal{L} be a distributive lattice. Then $\text{MSO}(\mathcal{L})$ is decidable if and only if $\text{MSO}(\mathbb{J}(\mathcal{L}))$ is decidable and the width $w(\mathcal{L})$ is bounded.*

Since distributive lattices are exactly the domains of conflict-free event structures and there exists a bijection between join irreducibles and the events of that event structure (Corollary 5.9), Theorem 9.14 can be viewed as a result about decidability of MSO theory of conflict-free event structures (graphs and event structures). That the MSO theory of trace conflict-free event structures is decidable follows from a more general result of Madhusudan [32]:

Theorem 9.15 ([32]). *The MSO theory of a trace event structure \mathcal{E} is decidable provided quantifications over sets are restricted to conflict-free subsets of events. In particular, if \mathcal{E} is conflict-free, then $\text{MSO}(\mathcal{E})$ is decidable.*

9.6. Grids. In this section we need to consider several types of square grids, which characterize different properties of event structures and their graphs. In this subsection, we will introduce some notational order between these notions and relate some of them. Recall that the infinite *square grid* Γ is the graph whose vertices correspond to the points in the plane with nonnegative integer coordinates and two vertices are connected by an edge whenever the corresponding points are at distance 1. The $n \times n$ square grid Γ_n is the subgraph of Γ whose vertices are all vertices of Γ with x - and y -coordinates in the range $0, \dots, n$. Γ and Γ_n can be viewed as directed graphs with respect to the basepoint order with respect to the corner $(0, 0)$. As we noticed above, Γ is the domain of the event structure consisting of two pairwise disjoint sets $X = \{x_0, x_1, x_2, \dots\}, Y = \{y_0, y_1, y_2, \dots\}$ of events, such that $x_0 < x_1 < x_2 < \dots$ and $y_0 < y_1 < y_2 < \dots$, and all events of X are concurrent with all events of Y . This event structure is conflict-free and trace regular. Below, if not specified, by Λ we denote either of the grids Γ or Γ_n . A *directed grid* $\vec{\Lambda}$ is a grid Λ with basepoint orientation with respect to the origin $(0, 0)$.

By Theorem 9.7, the treewidth of a graph is characterized by square grid minors. We will say that a square grid Λ is a *grid minor* of a graph G if Λ is a minor of G .

By Lemma 9.8, the hyperbolicity of a median graph (event domain or 1-skeleton of a CAT(0) cube complex) is characterized by isometrically embedded square grids. We will say that a square grid Λ is an *isometric grid* of a median graph $G = (V, E)$ if there exists an isometric embedding of Λ in G , i.e., a map $f : V(\Lambda) \rightarrow V$ such that $d_\Lambda(x, y) = d_G(x, y)$ for any two vertices $x, y \in V(\Lambda)$. An event structure characterization of isometric grids is provided below.

A stronger version of isometric grid is the notion of a flat grid. We will say that an isometric grid Λ is a *flat grid* of a median graph G if for any two vertices x, y of Λ at distance 2, any common neighbor z of x and y in G belongs to the grid Λ . Since any locally-convex connected subgraph of G is convex (Lemma 4.2), any flat grid is a convex (and thus gated) subgraph of G . If G is the 1-skeleton of a 2-dimensional cube complex, then any isometric grid is flat. If Λ is a flat grid of the graph $G(\mathcal{E})$ of an event domain $\mathcal{D}(\mathcal{E})$, then there are two disjoint subsets $X = \{x_0, x_1, x_2, \dots\}, Y = \{y_0, y_1, y_2, \dots\}$ of events of \mathcal{E} such that $x_0 < x_1 < x_2 < \dots$ and $y_0 < y_1 < y_2 < \dots$, and all events of X are concurrent with all events of Y .

The minor of a graph is defined by contracting edges. Minors are also implicitly used in the theory of event structures, namely, when the event structure $\mathcal{E} \setminus c$ rooted at a configuration c was defined (this notion was essential in the definition of regularity). The domain of $\mathcal{E} \setminus c$ is the principal filter $\mathcal{F}(c)$ of c . $\mathcal{F}(c)$ is a convex subgraph of $G(\mathcal{E})$, thus $\mathcal{F}(c)$ is the intersection of all halfspaces containing $\mathcal{F}(c)$. Therefore, $\mathcal{F}(c)$ can be obtained from the median graph $G(\mathcal{E})$ of the event structure \mathcal{E} by contracting all hyperplanes which do not intersect $\mathcal{F}(c)$.

Given a median graph G and a hyperplane H of its CAT(0) cube complex, the graph G' is obtained by *hyperplane-contraction* of G with respect to H if G' is obtained from G by simultaneously contracting all edges of G dual to H . Clearly, G' is also a median graph. We

will say that a median graph G' is a *strong-minor* of a median graph G if G' can be obtained from G by hyperplane-contraction of a set of hyperplanes of G .

Finally recall the event structure $\mathcal{E}_{TY} = (E, \leq, \#)$ occurring in the definition of grid-free event structures. Recall that E consists of three pairwise disjoint sets X, Y, Z satisfying the following conditions:

- $X = \{x_0, x_1, x_2, \dots\}$ is an infinite set of events with $x_0 < x_1 < x_2 < \dots$.
- $Y = \{y_0, y_1, y_2, \dots\}$ is an infinite set of events with $y_0 < y_1 < y_2 < \dots$.
- $X \times Y \subseteq \parallel$.
- There exists an injective mapping $g : X \times Y \rightarrow Z$ satisfying: if $g(x_i, y_j) = z$ then $x_i < z$ and $y_j < z$. Furthermore, if $i' > i$ then $x_{i'} \not< z$ and of $j' > j$ then $y_{j'} \not< z$.

The domain of \mathcal{E}_{TY} contains the infinite square grid Λ as a strong-minor. This grid corresponds to the events defined by the sets X and Y and is obtained by contracting all hyperplanes corresponding to the events in $E \setminus (X \cup Y \cup Z)$. On the other hand, the events from Z correspond to the hairs attached to the grid Λ in the definition of the hairing of an event structure. However, the relationship between the events of Z or the events of Z and a part of events of $X \cup Y$ is not specified, thus one cannot say more about the structure of the domain of \mathcal{E}_{TY} .

We continue with relationships between different types of grids. We start with isometric grids and hyperbolicity.

Lemma 9.16. *Let $\mathcal{E} = (E, \leq, \#)$ be an event structure of bounded degree. If the directed median graph $\vec{G}(\mathcal{E})$ contains an isometric $n \times n$ directed square grid, then E contains two disjoint conflict-free sets of events $A = \{x_0, x_1, \dots, x_{n-1}\}, B = \{y_0, y_1, \dots, y_{n-1}\}$ such that $x_i \parallel y_j$ for any two events $x_i \in A, y_j \in B$. Conversely, if for arbitrary m , E contains two disjoint sets of events $A = \{x_0, x_1, \dots, x_{m-1}\}, B = \{y_0, y_1, \dots, y_{m-1}\}$ such that $x_i \parallel y_j$ for any two events $x_i \in A, y_j \in B$, then the median graph $G(\mathcal{E})$ is not hyperbolic, and thus contains arbitrarily large isometric square grids.*

Proof. If $G(\mathcal{E})$ contains an isometric $n \times n$ directed grid $\vec{\Lambda}$, then let $X = \{x_0, \dots, x_{n-1}\}$ denote the events defining the edges of one side of $\vec{\Lambda}$ and let $Y = \{y_0, \dots, y_{n-1}\}$ denote the events defining the edges of another incident side of $\vec{\Lambda}$. Since each hyperplane H_{x_i} intersects each hyperplane H_{y_j} we conclude that the events of X are concurrent to the events of Y . It remains to show that two events of X or two events of Y cannot be in conflict. Pick for example any $x_i, x_j \in X$ with $i < j$ and suppose that X define horizontal edges of Λ and that the edges of X are directed from left to right. Then in $\vec{\Lambda}$ the hyperplane H_{x_i} separates the origin of the grid from the carrier of H_{x_j} . This implies that x_i and x_j cannot be in conflict.

To prove the converse, we will use Theorem 9.9 of [25]. By this theorem it suffices to show that for any n , the crossing graph $\Gamma(X(\mathcal{E}))$ contains a complete bipartite subgraph $K_{n,n}$. Suppose that the maximum degree of \mathcal{E} is d . Recall that the *Ramsey theorem* asserts that for any two nonnegative integers r and s there exists a least positive integer $R(r, s)$ such that any graph with at least $R(r, s)$ vertices either contains a stable set of size r or a clique of size s . Let $m \geq R(n, d + 1)$. Then E contains two disjoint sets of events $A = \{x_0, x_1, \dots, x_{m-1}\}, B = \{y_0, y_1, \dots, y_{m-1}\}$ such that $x_i \parallel y_j$ for any two events $x_i \in A, y_j \in B$. Recall that two events e and e' of E are concurrent if and only if their hyperplanes H_e and $H_{e'}$ intersect, i.e., H_e and $H_{e'}$ are adjacent in $\Gamma(X(\mathcal{E}))$. Consequently, H_{x_i} and H_{y_j} are adjacent in $\Gamma(X(\mathcal{E}))$ for any $x_i \in A$ and $y_j \in B$. Let Γ' (respectively, Γ'') be the subgraph of $\Gamma(X(\mathcal{E}))$ induced by the hyperplanes defined by the events of A (respectively, of B). Since Γ' contain $m \geq R(n, d + 1)$ vertices, by Ramsey's theorem, Γ' either contains a stable set A' of size n or a clique C' of size $d + 1$. In the second case we conclude that $X(\mathcal{E})$ contains $d + 1$ pairwise intersecting hyperplanes. By Proposition 4.9, this implies that $X(\mathcal{E})$ contains a $d + 1$ -cube Q . Since the orientation of the edges of $X(\mathcal{E})$ is admissible, Q contains a source of degree $d + 1$, contrary to the assumption that the maximum degree of \mathcal{E} is d . Consequently, Γ' contains a stable set A' of size n . Analogously, one can show that Γ'' contains a stable set B' of size n . But then $A' \cup B'$ induce the complete bipartite graph $K_{n,n}$ in the crossing graph $\Gamma(X(\mathcal{E}))$. \square

Proposition 9.17. *If the graph $G(\mathcal{E})$ of an event structure \mathcal{E} of bounded degree is hyperbolic, then \mathcal{E} is grid-free.*

Proof. Suppose by way of contradiction that the event structure \mathcal{E} is not grid-free, i.e., \mathcal{E} contains the three disjoint infinite sets of events X, Y, Z defining the event structure \mathcal{E}_{TY} . Since every event of X is concurrent with every event of Y , applying Lemma 9.16 with $A = X$ and $B = Y$, we deduce that $G(\mathcal{E})$ is not hyperbolic, a contradiction. \square

9.7. Proof of Theorem 9.1. Since for a Σ -labeled directed graph \vec{G} , the decidability of $\text{MSO}(\vec{G})$ implies the decidability of $\text{MSO}_1(G)$, (1) \Rightarrow (2). Since the degrees of vertices of $G(\mathcal{E})$ are uniformly bounded, the implication (2) \Rightarrow (3) follows from Courcelle's Theorem 9.11 [21]. The implication (3) \Rightarrow (4) is a particular case of Seese's Theorem 9.10 [45]. Finally, the implication (6) \Rightarrow (1) follows from the Müller and Schupp Theorem 9.12 [35] that the MSO theory of context-free graphs is decidable. It remains to establish the implications (4) \Rightarrow (5) and (5) \Rightarrow (6).

(4) \Rightarrow (5). Suppose by way of contradiction that $G(\mathcal{E})$ has clusters of arbitrarily large diameters. In this case, we will show that for any n one can construct in $G(\mathcal{E})$ a half of the square $n \times n$ grid as a minor (denote this half-grid by $\frac{1}{2}\Gamma_n$). Since $\frac{1}{2}\Gamma_n$ contains the $\frac{n}{2} \times \frac{n}{2}$ square grid, we deduce that $G(\mathcal{E})$ will contain arbitrarily large square grids as minors, contrary to the assumption that the treewidth of $G(\mathcal{E})$ is finite. We suppose that $\frac{1}{2}\Gamma_n$ has the set of vertices $\{z_{i,j} : 0 \leq i \leq n, 0 \leq j \leq n, \text{ and } i + j \leq n\}$.

Let v_0 be the basepoint. Recall that $S(v_0, k)$ is the sphere of radius k centered at v_0 . We will need the following properties of clusters of $G(\mathcal{E})$ (which hold for all median graphs):

Claim 9.18. *Let u, v be two vertices in a common cluster C of $G(\mathcal{E})$ located at distance k from v_0 . Then there exists a (u, v) -path $P'(u, v) = (u, p_1, q_1, p_2, q_2, \dots, p_{m-1}, q_{m-1}, p_m, v)$ such that $Q_1(u, v) = \{p_1, \dots, p_m\} \subseteq S(v_0, k + 1)$ and $Q_2(u, v) = \{q_1, \dots, q_{m-1}\} \subseteq C \subseteq S(v_0, k)$.*

Proof. Since u, v belong to a common cluster $C \subseteq S(v_0, k)$, there exists a (u, v) -path $P(u, v)$ all vertices of which have distance at least k from v_0 . Among all such paths, let $P'(u, v)$ be a path minimizing the sum $\sum_{w \in P'(u, v)} d(v_0, w)$. We assert that all vertices of $P'(u, v)$ have distance k or $k + 1$ to v_0 . Suppose by way of contradiction, that x is a furthest from v_0 vertex of $P'(u, v)$ and that $k' := d(v_0, x) \geq k + 2$. Let y and z are the neighbors of x in $P'(u, v)$. From the choice of x and since $G(\mathcal{E})$ is bipartite it follows that $d(v_0, y) = d(v_0, z) = k' - 1$. By quadrangle condition, there exists a vertex x' adjacent to y and z and having distance $k' - 2 \geq k$ from v_0 . Replacing in $P'(u, v)$ the vertex x by x' we will obtain a path $P'_0(u, v)$ in which all vertices have distance at least k from v_0 and having a smaller distance sum $\sum_{w \in P'_0(u, v)} d(v_0, w)$ than the path $P'(u, v)$. This contradicts the minimality choice of $P'(u, v)$. Therefore all vertices of $P'(u, v)$ have distance k or $k + 1$ from v_0 . Since the ends u, v of $P'(u, v)$ have distance k to v_0 and $G(\mathcal{E})$ is bipartite, the path $P'(u, v)$ is zigzagging, i.e., $P'(u, v) = (u, p_1, q_1, p_2, q_2, \dots, p_{m-1}, q_{m-1}, p_m, v)$ such that $Q_1(u, v) = \{p_1, \dots, p_m\} \subseteq S(v_0, k + 1)$ and $Q_2(u, v) = \{q_1, \dots, q_{m-1}\} \subseteq C \subseteq S(v_0, k)$. \square

Claim 9.19. *Let u, v be two vertices in a common cluster C of $G(\mathcal{E})$ located at distance k from v_0 . Then for any (u, v) -path $P_1(u, v) = (q_0 = u, p_1, q_1, p_2, q_2, \dots, p_{m-1}, q_{m-1}, p_m, v = q_m)$ such that $Q_1(u, v) = \{p_1, \dots, p_m\} \subseteq S(v_0, k + 1)$ and $Q_2(u, v) = \{q_1, \dots, q_{m-1}\} \subseteq C \subseteq S(v_0, k)$ there exists a sequence of vertices $Q_3(u, v) = \{r_1, \dots, r_{m'}\} \subseteq S(v_0, k - 1)$, such that $P_2(u, v) = (u, r_1, q_{i_1}, r_2, q_{i_2}, \dots, q_{i_{m'-1}}, r_{m'}, v)$ is a (u, v) -path of $G(\mathcal{E})$ and r_1 is adjacent to $q_0 = u, q_1, \dots, q_{i_1}$, r_2 is adjacent to $q_{i_1}, q_{i_1+1}, \dots, q_{i_2}$, etc, and $r_{m'}$ is adjacent to $q_{i_{m'-1}}, \dots, q_{m-1}, u = q_m$.*

Proof. Since $d(v_0, q_0) = d(v_0, q_1) = k$ and $d(v_0, p_1) = k + 1$, by quadrangle condition there exists a vertex r_1 adjacent to q_0 and q_1 and having distance $k - 1$ to v_0 . Let q_{i_1} be the last vertex of $Q_2(u, v)$ such that r_1 is adjacent to all vertices q_0, q_1, \dots, q_{i_1} . Again, since $d(v_0, q_{i_1}) = d(v_0, q_{i_1+1}) = k$ and $d(v_0, p_{i_1+1}) = k + 1$, by quadrangle condition there exists a vertex r_2 adjacent to q_{i_1} and q_{i_1+1} and having distance $k - 1$ to v_0 . Since r_1 is not adjacent to q_{i_1+1} , we have $r_1 \neq r_2$. Let q_{i_2} be the last vertex of $Q_2(u, v)$ such that r_2 is adjacent to all vertices

$q_{i_1}, q_{i_1+1}, \dots, q_{i_2}$. Continuing this way, we will define all vertices of $Q_3(u, v) = \{r_1, \dots, r_{m'}\}$. Then clearly $P_2(u, v) = (u, r_1, q_{i_1}, r_2, q_{i_2}, \dots, q_{i_{m'-1}}, r_{m'}, v)$ is a (u, v) -path. \square

We will call the union of paths $P_1(u, v)$ and $P_2(u, v)$ a *fence* and denote it by $F(u, v)$. Call $P_1(u, v)$ the *upper path* and $P_2(u, v)$ the *lower path* of $F(u, v)$. Notice that $P_1(u, v) \cap P_2(u, v) \subseteq C$. From the definition of clusters, all vertices of $P_2(u, v) \setminus P_1(u, v)$ also belong to a common cluster C' . If $d(u, v) = n'$, then both paths $P_1(u, v)$ and $P_2(u, v)$ have length at least n' . Thus, setting $u' := r_1, v' := r_{m'}$ and denoting by $P_1(u', v')$ the subpath of $P_2(u, v)$ comprised between u' and v' , we conclude that its length is at least $n' - 2$. On the other hand, the length of $P_1(u', v')$ is at most $n'' - 2$, where n'' is the length of $P_1(u, v)$. Applying Claim 9.19 to $P_1(u', v')$ we will define the path $P_2(u', v')$ and the fence $F(u', v')$. Notice that $P_1(u, v) = F(u, v) \cap F(u', v')$. Continuing this way, after $\frac{n'}{2} \leq n \leq n''$ steps, we will find two sequences of vertices $Q_u = (u = u_n, u_{n-1} = u', u_{n-2}, \dots, u_1, u_0 = w)$ and $Q_v = (v = v_n, v_{n-1} = v', v_{n-2}, \dots, v_1, v_0 = w)$ (constituting shortest (u, w) - and (v, w) -paths) and for each pair u_i, v_i a fence $F(u_i, v_i) = P_1(u_i, v_i) \cup P_2(u_i, v_i)$ such that any two consecutive fences $F(u_{i+1}, v_{i+1})$ and $F(u_i, v_i)$ intersect in the path $P_1(u_i, v_i)$. We will denote the union of all fences $F(u_i, v_i)$, $i = n, \dots, 0$, by F^* . We will also denote by C_i the cluster containing the vertices u_i and v_i (in particular, $C_n = C$).

We assert that F^* contains the half-grid $\frac{1}{2}\Gamma_n$ as a minor. Since $\frac{n'}{2} \leq n \leq n''$ and $n'' \geq n'$, we will be done. For this, for each vertex $z_{i,j}$ of $\frac{1}{2}\Gamma_n$ we will define a connected subgraph $Z_{i,j}$ of F^* satisfying the following properties:

- (1) $Z_{0,i} = \{u_i\}$ and $Z_{i,0} = \{v_i\}$ for each $i = 0, \dots, n$;
- (2) for each $k = 0, \dots, 2n$, if $i + j = k$, then $Z_{i,j}$ is a subpath of the lower path $P_2(u_k, v_k) \subseteq C_k \cup C_{k-1}$ of the fence $F(u_k, v_k)$ and $Z_{i,j}$ starts and ends at cluster C_i ;
- (3) for each $k = 0, \dots, 2n$ the paths $Z_{i,j}$ with $i + j = k$ are pairwise disjoint and are lexicographically ordered along $P_2(u_k, v_k)$ from u_k to v_k (i.e., for two pairs (i, j) and (i', j') with $i + j = i' + j' = k$, the path $Z_{i,j}$ appears before the path $Z_{i',j'}$ in $P_2(u_k, v_k)$ iff $i < i'$);
- (4) for each pair (i, j) , the first vertex of the path $Z_{i,j}$ is adjacent to the last vertex of the path $Z_{i-1,j}$ and the last vertex of $Z_{i,j}$ is adjacent to the first vertex of the path $Z_{i,j-1}$.

From last two conditions (3) and (4) we deduce that the paths $Z_{i,j}$ are pairwise disjoint and that contracting all such paths we will obtain $\frac{1}{2}\Gamma_n$ as a minor.

We will construct the paths $Z_{i,j}$ recursively. Suppose that the paths $Z_{i,j}$ satisfying the previous conditions have been defined for all pairs (i, j) such that $i + j \leq k$ and we have to define the paths $Z_{i,j}$ with $i + j = k + 1$. We proceed lexicographically on all such pairs. Consider a current pair (i, j) with $i + j = k + 1$. By induction assumption, the paths $Z_{i-1,j}$ and $Z_{i,j-1}$ have been defined. Let x denote the last vertex of the path $Z_{i-1,j}$ and y denote the first vertex of the path $Z_{i,j-1}$. By definition and induction hypothesis, the paths $Z_{i-1,j}$ and $Z_{i,j-1}$ are contained in the clusters $C_{k-1} \cup C_{k-2}$, are disjoint, and start and end at vertices of C_{k-1} . Consequently, x and y are vertices of C_{k-1} and x appears before y in the path $P_2(u_{k-1}, v_{k-1})$. In particular, x and y belong to the path $P_2(u_k, v_k)$. Traverse $P_2(u_k, v_k)$ from u_k to v_k . Denote by x' the vertex appearing after x in $P_2(u_k, v_k)$ and by y' the vertex of $P_2(u_k, v_k)$ appearing before y . Denote by $Z_{i,j}$ the subpath of $P_2(u_k, v_k)$ comprised between x' and y' . Then x' and y' are respectively the first and last vertices of $Z_{i,j}$.

We will show that $Z_{i,j}$ satisfies the conditions (2)-(4). Condition (4) follows from the definition of the vertices x, y, x', y' and of the path $Z_{i,j}$. Since $x, y \in C_{k-1}$, from the definition of $P_2(u_k, v_k)$ it follows that $x', y' \in C_k$. Hence $Z_{i,j}$ satisfies (2). Since the paths $Z_{i',j'}$ with $i' + j' = k - 1$ are lexicographically ordered, from the definition of the paths $Z_{i,j}$ with $i + j = k$ easily follows that such paths are also lexicographically ordered and pairwise disjoint. This concludes the proof of the implication.

(5) \Rightarrow (6). The implication follows from [6, Proposition 4.4] and the fact that trace event structures are recognizable by trace automata. Here we present a different (and hopefully simpler) proof. Let $\Phi(G(\mathcal{E}))$ be the set of ends of $G(\mathcal{E})$. We have to prove that $\Phi(G(\mathcal{E}))$ has only finitely many isomorphism classes under end-isomorphisms. Recall that there exists a bijection between

the ends of $\Phi(G(\mathcal{E}))$ and the clusters of $\mathcal{C}(G(\mathcal{E}))$ and that for a cluster C we denote by $\Upsilon(C)$ the end containing C . Let M be the size of the alphabet Σ . Since \mathcal{E} is a trace regular event structure, the degrees of vertices of $G(\mathcal{E})$ are uniformly bounded, say by some constant Δ . Suppose that the diameters of clusters are uniformly bounded by D . We will say that the sets of a set family \mathcal{S} have *constant size* if the sizes of all sets of \mathcal{S} depend only of M, Δ , and D . First notice that all clusters C have constant size. Indeed, pick a vertex v of C . Then C is included in the ball $B(v, D)$ of radius D centered at v . This ball has at most $K := \sum_{i=0}^D \Delta^i = O(\Delta^D)$ vertices.

Pick any cluster C . Since $(\mathcal{D}(\mathcal{E}), \subseteq)$ is a median meet semilattice, there exists the smallest median meet sub-semilattice $M(C)$ containing the set C , the *median closure* of C . Let $m_C \in M(C)$ denote the meet of C in $(\mathcal{D}(\mathcal{E}), \subseteq)$. Denote by $H(C)$ the subgraph of $G(\mathcal{E})$ induced by the set $\{v \in M(C) : m_C \leq v \text{ and } \exists x \in C, v \leq x\}$ of all vertices of $M(C)$ larger or equal than m_C and each smaller than a vertex of C .

Claim 9.20. *The graphs $H(C), C \in \mathcal{C}(G(\mathcal{E}))$, have constant size.*

Proof. Any $H(C)$ is included in the median closure $M(C)$ and any $M(C)$ is included in the convex hull $\text{conv}(C)$ of C . Therefore it suffices to prove that $\text{conv}(C)$ has constant size, namely that it has constant diameter. Pick $x, y \in \text{conv}(C)$. The distance $d(x, y)$ in $G(\mathcal{E})$ is the number of hyperplanes separating x and y . Since $\text{conv}(C)$ is the intersection of all halfspaces of $G(\mathcal{E})$ including C , the hyperplanes defining such halfspaces do not separate x and y . Therefore x and y can be separated only by hyperplanes separating vertices of C . There are at most D hyperplanes separating two given vertices of C , thus there are at most DK^2 hyperplanes separating vertices of C . Consequently, $d(x, y) \leq DK^2$, establishing that the diameter of $\text{conv}(C)$ is constant. \square

Suppose that the edges of each $H(C)$ are directed and labeled by λ as in $G(\mathcal{E})$. For each vertex $v \in H(C)$, let $r(v) = i$ if the principal filter $\mathcal{F}(v)$ of v belongs to the isomorphism class i . Call the edge- and vertex-labeled graph $(H(C), \lambda, r)$ the *recent past* of the cluster C . Since by Claim 9.20 all $H(C)$ have constant size, λ is finite, and there exists only a finite number of types of principal filters, we conclude that there exists only a finite number of types of recent pasts $\mathcal{P}_1, \dots, \mathcal{P}_n$.

Pick any isomorphism class \mathcal{P}_i and pick any two graphs $H(C)$ and $H(C')$ belonging to \mathcal{P}_i . Notice that since any isomorphism g between $H(C)$ and $H(C')$, preserves the orientation of edges, g maps the unique source m_C of $H(C)$ to the unique source $m_{C'}$ of $H(C')$. The set C of sinks of $H(C)$ is mapped to the set C' of sinks of $H(C')$. Since $r(m_C) = r(m_{C'})$, there exists an isomorphism f between the labeled principal filters $\mathcal{F}(m_C)$ and $\mathcal{F}(m_{C'})$.

Claim 9.21. *Any isomorphism g between $(H(C), \lambda, r)$ and $(H(C'), \lambda, r)$ coincides with f , i.e., for any vertex v of $H(C)$, $f(v) = g(v)$.*

Proof. Let $m_C \in P(C)$ and $m_{C'} \in P(C')$ be the meets of C and C' . Let g be an isomorphism between $(H(C), \lambda, r)$ and $(H(C'), \lambda, r)$. By induction on $k := d(v, m_C)$ we will prove that $g(v) = f(v)$. Since any isomorphism maps m_C to $m_{C'}$, this is true for $k = 0$. Since the labeling λ is deterministic, any isomorphism will map the edge $m_C v$ to the unique edge $m_{C'} v'$ labeled $\lambda(m_C v)$. Therefore $g(v) = v' = f(v)$ for any neighbor v of m_C in $H(C)$. Therefore, our assertion is also true for $k = 1$. Suppose that it is true for all vertices v of $H(C)$ such that $d(v, m_C) \leq k$ and pick a vertex w with $d(w, m_C) = k + 1$. Let v be a neighbor of w at distance k from m_C . By induction assumption, $g(v) = f(v)$, denote this vertex by v' . Set $w' := g(w)$ and $w'' := f(w)$. Since vw is an edge outgoing from v , $v'w'$ and $v'w''$ are edges outgoing from v' , both labeled by $\lambda(vw)$ ($v'w'$ is an edge of $H(C')$ but the vertex w'' and the edge $v'w''$ are not necessarily in $H(C')$). Since the labeling λ is deterministic, this is possible only if $w' = w''$, whence $g(w) = f(w)$. \square

The following claim concludes the proof of the implication (v) \Rightarrow (vi):

Claim 9.22. *f is an end-isomorphism between $\Upsilon(C)$ and $\Upsilon(C')$.*

Proof. Since $\Upsilon(C)$ is the union of all principal filters $\mathcal{F}(v), v \in C$, we have $\Upsilon(C) \subseteq \mathcal{F}(m_C)$, thus f is well-defined on $\Upsilon(C)$. Since f is a bijective map between $\mathcal{F}(m_C)$ and $\mathcal{F}(m_{C'})$, f is an injective map from $\Upsilon(C)$ to $\mathcal{F}(m_{C'})$. By Claim 9.21, f maps C to C' , thus the f -image of any principal filter $\mathcal{F}(v)$ with $v \in C$ is a principal filter $\mathcal{F}(f(v))$ with $f(v) \in C'$, thus $f(\Upsilon(C)) \subseteq \Upsilon(C')$. Since any vertex of $\Upsilon(C')$ belongs to at least one principal filter $\mathcal{F}(v')$ with $v' \in C'$ and f bijectively maps C to C' , f is a surjective map from $\Upsilon(C)$ to $\Upsilon(C')$. Since f is also injective on $\Upsilon(C)$, f is a bijection between $\Upsilon(C)$ and $\Upsilon(C')$. Since any edge xy of $\Upsilon(C)$ belongs to at least one principal filter $\mathcal{F}(v)$ with $v \in C$, f maps xy to an edge $x'y'$ of $\Upsilon(C')$ and $\lambda(x'y') = \lambda(xy)$ holds. Since the same property holds for edges of $\Upsilon(C')$, this establishes that f is an end-isomorphism between $\Upsilon(C)$ and $\Upsilon(C')$. \square

9.8. Proof of Proposition 9.2. For a covering map $\varphi : Y \rightarrow X$, an automorphism $\alpha : Y \rightarrow Y$ is a *deck transformation* of φ if $\varphi \circ \alpha = \varphi$. The set of deck transformations of φ forms a group under composition. A covering map $\varphi : Y \rightarrow X$ is called *normal* (or *regular*) if for each pair of lifts $y, y' \in Y$ of $x \in X$ there is a deck transformation mapping y to y' . If there is a normal covering map $\varphi : Y \rightarrow X$, then Y is called a *normal cover* of X . Every universal cover is normal, with deck transformation group being isomorphic to the fundamental group $\pi_1(X)$ of X (see [25, Proposition 1.39]). Since the special cube complex X_N of a net system N is finite and its universal cover \tilde{X}_N is normal, this implies that the automorphism group of \tilde{X}_N has a finite number of orbits. Therefore, the automorphism group of the directed labeled graph $\vec{G}(\tilde{X}_N)$ also has a finite number of orbits. Consequently, to $\vec{G}(\tilde{X}_N)$ we can apply Theorem 9.13 of Kuske and Lohrey [31] and deduce Proposition 9.2.

9.9. The MSO theory of hairings of event structures. The goal of this subsection is to prove Proposition 9.3 and Theorem 9.4.

Let $\dot{\mathcal{E}} = (\dot{E} = E \cup E_C, \dot{\leq}, \dot{\#})$ be the hairing of a trace regular event structure $\mathcal{E} = (E, \leq, \#)$ with a trace labeling λ over the trace alphabet $M = (\Sigma, I)$. Let h be a letter that does not belong to Σ and consider the trace alphabet $\dot{M} = (\Sigma \cup \{h\}, I)$. Note that since I is not modified, it means that $(h, a) \notin I$ for every $a \in \Sigma$. Let $\dot{\lambda}$ be the labeling of $\dot{\mathcal{E}}$ extending λ by setting $\dot{\lambda}(e_c) = h$ for any hair event $e_c \in E_C$.

9.9.1. Proof of Proposition 9.3. Assume that λ is a regular trace labeling of \mathcal{E} over $M = (\Sigma, I)$ and let $\dot{\lambda}$ be the labeling of $\dot{\mathcal{E}}$ defined as above. We first show that $\dot{\lambda}$ is a trace labeling of $\dot{\mathcal{E}}$, i.e., that $\dot{\lambda}$ satisfies (LES1), (LES2), and (LES3).

For any two events $e, e' \in E$, the properties (LES1)-(LES3) are satisfied because λ is a regular trace labeling of \mathcal{E} . Suppose now that e' is a hair event. Hence $\dot{\lambda}(e') = h$ and for any $a \in \Sigma \cup \{h\}$, $(\dot{\lambda}(e'), a) \notin \dot{I}$. Consequently (LES2) trivially holds. Since a hair event is not concurrent with any other event, (LES3) also trivially holds. If e is in minimal conflict with e' , then e cannot be a hair event and thus $\dot{\lambda}(e) \neq \dot{\lambda}(e')$, establishing (LES1).

We now show that $\dot{\lambda}$ is a regular labeling of $\dot{\mathcal{E}}$. Consider a configuration \dot{c} of $\dot{\mathcal{E}}$ and observe that if \dot{c} contains a hair event $e_{c'}$ associated with a configuration c' of \mathcal{E} , then $\dot{c} = c' \cup \{e_{c'}\}$. Consequently, for any such configuration \dot{c} , $\dot{\mathcal{E}}^{\dot{\lambda}} \setminus \dot{c}$ is empty. Therefore, all such configurations are equivalent for $R_{\dot{\mathcal{E}}^{\dot{\lambda}}}$. Observe that any configuration c_0 of $\dot{\mathcal{E}}$ that does not contain a hair event is also a configuration of \mathcal{E} . Consider two configurations $c_0, c'_0 \in \dot{\mathcal{E}}$ such that $c_0 R_{\mathcal{E}} c'_0$ and let f be an isomorphism from $\mathcal{E} \setminus c_0$ to $\mathcal{E} \setminus c'_0$. We define an isomorphism \dot{f} from $\dot{\mathcal{E}} \setminus c_0$ to $\dot{\mathcal{E}} \setminus c'_0$ as follows. For any event $e \in \mathcal{E} \setminus c_0$, let $\dot{f}(e) = f(e)$. For any configuration c of $\mathcal{E} \setminus c_0$, $f(c) = \{f(e) : e \in c\}$ is a configuration of $\mathcal{E} \setminus c'_0$ and we let $\dot{f}(e_c) = e_{f(c)}$. Observe that in any case, $\dot{\lambda}(\dot{f}(e)) = \dot{\lambda}(e)$. Consider any two events $e_1, e_2 \in \dot{\mathcal{E}} \setminus \{c_0\}$. If $e_1, e_2 \in \mathcal{E} \setminus c_0$, then $\dot{f}(e_1) \dot{\leq} \dot{f}(e_2)$ iff $e_1 \dot{\leq} e_2$ and $\dot{f}(e_1) \dot{\#} \dot{f}(e_2)$ iff $e_1 \dot{\#} e_2$ since f is an isomorphism from $\mathcal{E} \setminus c_0$ to $\mathcal{E} \setminus c'_0$. Suppose now that e_2 is a hair event e_c associated to a configuration c of $\mathcal{E} \setminus c_0$. Then $e_1 \dot{\leq} e_c$ if $e_1 \in c$ and $e_1 \dot{\#} e_c$ otherwise. In the first case, $f(e_1) \in f(c)$ and consequently $\dot{f}(e_1) = f(e_1) \dot{\leq} e_{f(c)} = \dot{f}(e_c)$. In the second case, $f(e_1) \notin f(c)$ and thus $\dot{f}(e_1) = f(e_1) \dot{\#} e_{f(c)} = \dot{f}(e_c)$. Since f is bijective, \dot{f} is also bijective and thus \dot{f} is an isomorphism from $\dot{\mathcal{E}} \setminus c_0$ to $\dot{\mathcal{E}} \setminus c'_0$. Consequently, since $R_{\mathcal{E}^{\lambda}}$ has

finite index, so does $R_{\dot{\mathcal{E}}\lambda}$, showing that $\dot{\lambda}$ is a regular trace labeling of $\dot{\mathcal{E}}$. This ends the proof of Proposition 9.3.

9.9.2. *Proof of Theorem 9.4.* The proof of Theorem 9.4 is based on Theorem 9.1 and Propositions 9.23 and 9.24.

Proposition 9.23. *For a trace regular event structure $\mathcal{E} = (E, \leq, \#, \lambda)$, if $\text{MSO}(\dot{\mathcal{E}})$ is decidable, then $\text{MSO}(\vec{G}(\mathcal{E}))$ is decidable.*

Proof. We transform any formula $\varphi_G(\mathbf{v}, \mathbf{X}) \in \text{MSO}(\vec{G}(\mathcal{E}))$, where $\mathbf{v} = \{v_1, \dots, v_n\}$ and $\mathbf{X} = \{X_1, \dots, X_m\}$, into a formula $\varphi_{\dot{E}S}(\mathbf{e}, \mathbf{X}) \in \text{MSO}(\dot{\mathcal{E}})$ where $\mathbf{e} = \{e_{v_1}, \dots, e_{v_n}\}$. The variables representing the vertices of $\vec{G}(\mathcal{E})$ will be replaced by variables representing the corresponding hair events of $\dot{\mathcal{E}}$.

We proceed by induction on the structure of $\varphi_G(v_1, \dots, v_n, X_1, \dots, X_m)$. We need to explain the transformations for atomic formulas, Boolean combinations of formulas, and existential quantifications over vertices and sets of vertices.

We first consider atomic formulas.

If $\varphi_G(v, X) = (v \in X)$, then we set $\varphi_{\dot{E}S}(e_v, X) := (e_v \in X)$.

If $\varphi_G(u, v) = (u = v)$, then we set $\varphi_{\dot{E}S}(e_u, e_v) := (e_u = e_v)$.

If $\varphi_G(u, v) = ((u, v) \in E_a)$ for some letter $a \in \Sigma$, then we set

$$\varphi_{\dot{E}S}(e_u, e_v) := \exists e((\dot{\lambda}(e) = a) \wedge (e \#_{\mu} e_u) \wedge (e \dot{\prec} e_v)).$$

If $\varphi_G(\mathbf{v}, \mathbf{X})$ is a boolean combination of formulas of $\text{MSO}(\vec{G}(\mathcal{E}))$, then $\varphi_{\dot{E}S}(\mathbf{e}, \mathbf{X})$ is the same boolean combination of the corresponding formulas in $\text{MSO}(\dot{\mathcal{E}})$.

If $\varphi_G(\mathbf{v}, \mathbf{X}) = \exists v \varphi'_G(\{v\} \cup \mathbf{v}, \mathbf{X})$ and $\varphi'_{\dot{E}S}(\{e_v\} \cup \mathbf{e}, \mathbf{X})$ is the formula obtained from $\varphi'_G(\{v\} \cup \mathbf{v}, \mathbf{X})$, then we set

$$\varphi_{\dot{E}S}(\mathbf{e}, \mathbf{X}) := \exists e_v \left((\dot{\lambda}(e_v) = h) \wedge \varphi'_{\dot{E}S}(\{e_v\} \cup \mathbf{e}, \mathbf{X}) \right).$$

If $\varphi_G(\mathbf{v}, \mathbf{X}) = \exists X \varphi'_G(\mathbf{v}, \{X\} \cup \mathbf{X})$ and $\varphi'_{\dot{E}S}(\mathbf{e}, \{X\} \cup \mathbf{X})$ is the formula obtained from $\varphi'_G(\mathbf{v}, \{X\} \cup \mathbf{X})$, then we set

$$\varphi_{\dot{E}S}(\mathbf{e}, \mathbf{X}) := \exists X \left((\forall e_v \in X (\dot{\lambda}(e_v) = h)) \wedge \varphi'_{\dot{E}S}(\mathbf{e}, \{X\} \cup \mathbf{X}) \right).$$

For every sentence φ_G in $\text{MSO}(\vec{G}(\mathcal{E}))$, the sentence $\varphi_{\dot{E}S}$ obtained by our construction is a sentence in $\text{MSO}(\dot{\mathcal{E}})$. Moreover, by induction on the structure of the sentence, it is obvious to see that for any trace regular event structure \mathcal{E} , $\vec{G}(\mathcal{E})$ satisfies φ_G if and only if $\dot{\mathcal{E}}$ satisfies $\varphi_{\dot{E}S}$. \square

Proposition 9.24. *For a trace regular event structure $\mathcal{E} = (E, \leq, \#, \lambda)$, if $\text{MSO}(\vec{G}(\mathcal{E}))$ is decidable, then $\text{MSO}(\mathcal{E})$ is decidable.*

Proof. Consider any formula $\varphi(\mathbf{e}, \mathbf{X}) = \varphi_{ES}(\mathbf{e}, \mathbf{X}) \in \text{MSO}(\mathcal{E})$, where $\mathbf{e} = \{e_1, \dots, e_n\}$ and $\mathbf{X} = \{X_1, \dots, X_m\}$. We first transform this formula into another formula of $\text{MSO}(\mathcal{E})$ as follows. Since \leq is the transitive closure of $\dot{\prec}$ and since the transitive closure of any binary relation expressible in MSO can also be expressed in MSO (see [22, Section 5.2.2]), we can assume that the atomic formulas of $\varphi(\mathbf{e}, \mathbf{X})$ are of the type $e \in X$, $e_1 = e_2$, $R_a(e)$ for $a \in \Sigma$, and $e_1 \dot{\prec} e_2$.

We now transform the formula $\varphi(\mathbf{e}, \mathbf{X})$ in such a way that each event variable (respectively, each set variable) has a label $a \in \Sigma$, i.e., it can be interpreted only by an event labeled by a (respectively, by a subset of events labeled by a). Assume in the following that $\Sigma = \{a_1, \dots, a_n\}$.

We transform $\varphi(\mathbf{e}, \mathbf{X})$ into $\varphi'(\mathbf{e}, \mathbf{X})$ in an inductive way as follows.

If $\varphi(x) = R_a(x)$ for $a \in \Sigma$, then $\varphi'(x) := R_a(x)$.

If $\varphi(x_1, x_2) = x_1 \dot{\prec} x_2$, then $\varphi'(x_1, x_2) := x_1 \dot{\prec} x_2$.

If $\varphi(\mathbf{e}, \mathbf{X}) = \neg \varphi_1(\mathbf{e}, \mathbf{X})$ and if $\varphi'_1(\mathbf{e}', \mathbf{X}')$ is the formula obtained from $\varphi_1(\mathbf{e}, \mathbf{X})$, then $\varphi'(\mathbf{e}', \mathbf{X}') := \neg \varphi'_1(\mathbf{e}', \mathbf{X}')$.

If $\varphi(\mathbf{e}, \mathbf{X}) = \varphi_1(\mathbf{e}_1, \mathbf{X}_1) \vee \varphi_2(\mathbf{e}_2, \mathbf{X}_2)$ with $\mathbf{e} = \mathbf{e}_1 \cup \mathbf{e}_2$ and $\mathbf{X} = \mathbf{X}_1 \cup \mathbf{X}_2$ and if $\varphi'_1(\mathbf{e}'_1, \mathbf{X}'_1)$ and $\varphi'_2(\mathbf{e}'_2, \mathbf{X}'_2)$ are the formulas obtained respectively from $\varphi_1(\mathbf{e}_1, \mathbf{X}_1)$ and $\varphi_2(\mathbf{e}_2, \mathbf{X}_2)$, then $\varphi'(\mathbf{e}' = \mathbf{e}'_1 \cup \mathbf{e}'_2, \mathbf{X}' = \mathbf{X}'_1 \cup \mathbf{X}'_2) := \varphi'_1(\mathbf{e}'_1, \mathbf{X}'_1) \vee \varphi'_2(\mathbf{e}'_2, \mathbf{X}'_2)$.

If $\varphi(e, X) = (e \in X)$, then let $\mathbf{X} = \{X_a : a \in \Sigma\}$ and set

$$\varphi'(e, \mathbf{X}) := \bigwedge_{a \in \Sigma} \left(R_a(e) \iff e \in X_a \right).$$

If $\varphi(\mathbf{e}, \mathbf{X}) = \exists e \varphi_1(\{e\} \cup \mathbf{e}, \mathbf{X})$ and if $\varphi'_1(\{e\} \cup \mathbf{e}', \mathbf{X}')$ is the formula obtained from $\varphi_1(\{e\} \cup \mathbf{e}, \mathbf{X})$, then

$$\varphi'(\mathbf{e}', \mathbf{X}') := \bigvee_{a \in \Sigma} \left(\exists e_a \left(R_a(e_a) \wedge \varphi'_1(\{e_a\} \cup \mathbf{e}', \mathbf{X}') \right) \right).$$

If $\varphi(\mathbf{e}, \mathbf{X}) = \exists Y \varphi_1(\mathbf{e}, \{Y\} \cup \mathbf{X})$ and if $\varphi'_1(\mathbf{e}', \mathbf{Y} \cup \mathbf{X}')$ is the formula obtained from $\varphi_1(\mathbf{e}, \{Y\} \cup \mathbf{X})$ where $\mathbf{Y} = \{Y_{a_1}, \dots, Y_{a_n}\}$, then

$$\varphi'(\mathbf{e}', \mathbf{X}') := \exists Y_{a_1} \dots \exists Y_{a_n} \bigwedge_{a_i \in \Sigma} \left((\forall e \in Y_{a_i} (R_{a_i}(e))) \right) \wedge \varphi'_1(\mathbf{e}', \mathbf{Y} \cup \mathbf{X}').$$

Consequently, from now on, we consider only formulas $\varphi_{ES}(\mathbf{e}, \mathbf{X})$ in which every variable e_a or S_a is indexed by a letter $a \in \Sigma$, meaning that it can only be interpreted by an event (or a set of events) labeled by a .

Given such a formula $\varphi_{ES}(\mathbf{e}, \mathbf{X}) \in \text{MSO}(\mathcal{E})$, we construct a formula $\varphi_G(\mathbf{S}, \mathbf{X}) \in \text{MSO}(\vec{G}(\mathcal{E}))$ where each event variable e is replaced by a second order variable representing a set of vertices S . The idea of the transformation is that an event variable e can be interpreted in \mathcal{E} by an event f if and only if the set S can be interpreted in $\vec{G}(\mathcal{E})$ by the set of sources of precisely those edges which are dual to the hyperplane \mathcal{H}_f . Similarly, a set of events will be represented by the set of sources of the edges dual to the corresponding hyperplanes.

We proceed by induction on the structure of $\varphi_{ES}(\mathbf{e}, \mathbf{X}) \in \text{MSO}(\vec{G}(\mathcal{E}))$. We first consider atomic formulas.

If $\varphi_{ES}(\{e_a\}, \emptyset) = (R_b(e_a))$, then we set $\varphi_G(\{S_{e_a}\}) := \perp$ if $a \neq b$ and $\varphi_G(\{S_{e_a}\}) := \top$ otherwise.

If $\varphi_{ES}(\{e_a\}, \{X_b\}) = (e_a \in X_b)$, then we set $\varphi_G(\{S_{e_a}\}, \{X_b\}) := \perp$ if $a \neq b$ and $\varphi_G(\{S_{e_a}\}, \{X_b\}) := (S_{e_a} \subseteq X_b)$ otherwise.

If $\varphi_{ES}(\{e_a, e_b\}, \emptyset) = (e_a = e_b)$, then we set $\varphi_G(\{S_{e_a}, S_{e_b}\}) := \perp$ if $a \neq b$ and $\varphi_G(\{S_{e_a}, S_{e_b}\}) := (S_{e_a} = S_{e_b})$ otherwise.

If $\varphi_{ES}(\{e_a, e_b\}, \emptyset) = (e_a \prec e_b)$, then we set

$$\varphi_G(\{S_{e_a}, S_{e_b}\}) := \left(\exists s \in S_{e_a} \exists t \in S_{e_b} ((s, t) \in E_a) \right) \wedge \neg \left(\exists s \in S_{e_b} \exists t \in S_{e_a} ((s, t) \in E_b) \right).$$

If $\varphi_{ES}(\mathbf{e}, \mathbf{X})$ is a boolean combination of formulas of $\text{MSO}(\mathcal{E})$, then $\varphi_G(\mathbf{S}, \mathbf{X})$ is the same boolean combination of the corresponding formulas in $\text{MSO}(\vec{G}(\mathcal{E}))$.

In the following, we need to ensure that the set of vertices representing the edges dual to a hyperplane (or to a set of hyperplanes) labeled by a are indeed the sources of edges labeled by a . Given a second order variable Y_a , this is ensured by the following formula $H_a(Y_a)$:

$$H_a(Y_a) := \forall s \in Y_a \exists t ((s, t) \in E_a).$$

We also need to ensure that the variables representing the edges dual to a hyperplane (or to a set of hyperplanes) labeled by a represent sets of edges that are closed in the directed median graph $\vec{G}(\mathcal{E})$ under the parallelism relation. Given a second order variable Y_a , this is ensured by the following formula $PC_a(Y_a) \in \text{MSO}(\vec{G}(\mathcal{E}))$:

$$\begin{aligned} PC_a(Y_a) &:= \bigwedge_{b \in \Sigma} \left(\forall s_1 \forall s_2 \forall t_1 \forall t_2 \left((s_1, t_1), (s_2, t_2) \in E_a \right) \wedge \left((s_1, s_2), (t_1, t_2) \in E_b \right) \right. \\ &\quad \left. \implies \left((s_1 \in Y_a) \iff (s_2 \in Y_a) \right) \right). \end{aligned}$$

We first consider second order existential quantification in $\text{MSO}(\mathcal{E})$. We want to replace a variable representing a set of events by a variable representing the set of sources of the edges dual to these events. If $\varphi_{ES}(\mathbf{e}, \mathbf{X}) = \exists X_a \psi_{ES}(\mathbf{e}, \{X_a\} \cup \mathbf{X})$ and $\psi_G(\mathbf{S}, \{X_a\} \cup \mathbf{X})$ is the formula obtained from $\psi_{ES}(\mathbf{e}, \{X_a\} \cup \mathbf{X})$, then we set

$$\varphi_G(\mathbf{S}, \mathbf{X}) := \exists X_a \left(H_a(X_a) \wedge PC_a(X_a) \wedge \psi_G(\mathbf{S}, \{X_a\} \cup \mathbf{X}) \right).$$

We now consider first order existential quantification in $\text{MSO}(\mathcal{E})$ and we use the previous transformation and the fact that an event is a minimal non-empty subset of events. The following formula ensures that S_a is a non-empty set of sources of edges dual to a set of events labeled by a :

$$NSH(S_a) := \left(H_a(S_a) \wedge PC_a(S_a) \wedge (\exists s \in S_a) \right).$$

If $\varphi_{ES}(\mathbf{e}, \mathbf{X}) = \exists e_a \psi_{ES}(\{e_a\} \cup \mathbf{e}, \mathbf{X})$ and $\psi_G(\{S_a\} \cup \mathbf{S}, \mathbf{X})$ is the formula obtained from $\psi_{ES}(\{e_a\} \cup \mathbf{e}, \mathbf{X})$, then we set:

$$\varphi_G(\mathbf{S}, \mathbf{X}) := \exists S_a \left(NSH(S_a) \wedge \left(\forall S'_a (NSH(S'_a) \wedge (S'_a \subseteq S_a)) \implies (S_a = S'_a) \right) \wedge \psi_G(\{S_a\} \cup \mathbf{S}, \mathbf{X}) \right).$$

For every sentence φ_{ES} in $\text{MSO}(\mathcal{E})$, the sentence φ_G obtained by our construction is a sentence in $\text{MSO}(\vec{G}(\mathcal{E}))$. Moreover, by induction on the structure of the sentence, it can be shown that for any trace regular event structure \mathcal{E} , \mathcal{E} satisfies φ_{ES} if and only if $\vec{G}(\mathcal{E})$ satisfies φ_G . \square


The “if” implication of Theorem 9.4 is the content of Proposition 9.23. To prove the converse implication, consider a trace regular event structure $\mathcal{E} = (E, \leq, \#, \lambda)$, such that $\text{MSO}(\vec{G}(\mathcal{E}))$ is decidable. By Theorem 9.1, $G(\mathcal{E})$ has finite treewidth. Obviously, this implies that $G(\dot{\mathcal{E}})$ has also finite treewidth. By Theorem 9.1, $\text{MSO}(\vec{G}(\dot{\mathcal{E}}))$ is decidable, and thus, by Proposition 9.24, $\text{MSO}(\dot{\mathcal{E}})$ is decidable.

Remark 9.25. Notice that the converse of Proposition 9.24 is not true: the MSO theory of trace conflict-free event structures is decidable [32], however the graphs of their domains may have infinite treewidth and thus an undecidable MSO theory. For example, the event structure $\mathcal{E} = (E, \leq, \#)$ consisting of two pairwise disjoint sets $X = \{x_0, x_1, x_2, \dots\}, Y = \{y_0, y_1, y_2, \dots\}$ of events, such that $x_0 < x_1 < x_2 < \dots$ and $y_0 < y_1 < y_2 < \dots$, and all events of X are concurrent with all events of Y , is conflict-free and trace regular, but its domain $\mathcal{D}(\mathcal{E})$ is the infinite square grid.

10. COUNTEREXAMPLES TO CONJECTURE 3.4

In this section, we use the general results obtained in Section 9 to construct a counterexample to Thiagarajan’s Conjecture 3.4. In view of Theorem 9.4, it suffices to find a trace regular event structure \mathcal{E} whose graph $G(\mathcal{E})$ has unbounded treewidth (i.e., it contains arbitrarily large square grid minors) and whose hairing $\dot{\mathcal{E}}$ is grid-free (as an event structure). To build such an example, as in [13], we start by constructing a finite NPC square complex. Namely, we consider an NPC square complex Z with one vertex, four edges, and three squares, and we show that Z is virtually special. This implies that the principal filter of the universal cover \tilde{Z} of Z is the domain $\mathcal{D}(\mathcal{E}_Z)$ of a trace regular event structure (i.e., \mathcal{E}_Z is the event structure unfolding of a net system N_Z). We prove that the median graph $G(\mathcal{E}_Z)$ of the domain has unbounded treewidth. On the other hand, to prove that $\dot{\mathcal{E}}_Z$ is grid-free we show that it is enough to prove that the graph $G(\mathcal{E}_Z)$ of the domain has bounded hyperbolicity (this correspond to bounded isometric square grids). In conclusion, we obtain the following result:

Theorem 10.1. *There exists a virtually special NPC square complex Z such that for the trace regular event structure \mathcal{E}_Z having the principal filter of \tilde{Z} as the domain, the hairing $\dot{\mathcal{E}}_Z$ is grid-free but the median graph $G(\mathcal{E}_Z)$ of \mathcal{E}_Z has unbounded treewidth. Consequently, $\text{MSO}(\dot{\mathcal{E}}_Z)$ is undecidable and thus Thiagarajan’s Conjecture 3.4 is false.*

FIGURE 4. The three squares defining the VH-complex Z

Badouel et al. [6, pp. 144–146] described a trace regular event structure that has a domain that is not context-free. Using the results of Section 9, we show that the hairing of this event structure is also a counterexample to Conjecture 3.4.

10.1. Proof of Theorem 10.1. The square complex Z consists of three squares Q_1, Q_2, Q_3 , one vertex v_0 , and four edges, colored and directed as in Figure 4. The four edges of Z are colored orange (color a), black (color b), blue (color x), and red (color y) as indicated in the figure. Since Z is a VH-complex, Z is nonpositively curved. Let $\tilde{Z} = (\tilde{Z}, \tilde{o}, \tilde{c})$ denote the directed and colored universal cover of Z . Pick any vertex \tilde{v}_0 of \tilde{Z} (\tilde{v}_0 is a lift of v_0) and let \mathcal{E}_Z denote the event structure whose domain is the principal filter $\mathcal{D}_Z = (\mathcal{F}_{\tilde{o}}(\tilde{v}_0, \tilde{Z}^{(1)}), \prec_{\tilde{o}})$ of (\tilde{Z}, \tilde{o}) . Let also $\vec{G}(\mathcal{E}_Z)$ and $G(\mathcal{E}_Z)$ denote the directed and the undirected 1-skeletons of \mathcal{D}_Z . Finally, denote by $\hat{\mathcal{E}}_Z$ the hairing of \mathcal{E}_Z .

First we investigate the properties of the complexes Z and \tilde{Z} , of the graphs \vec{G}_Z and G_Z , and of the event structure \mathcal{E}_Z . First, even if Z is not special, we show that it is virtually special:

Lemma 10.2. *The NPC square complex Z is virtually special. Consequently, the event structures \mathcal{E}_Z and $\hat{\mathcal{E}}_Z$ are trace regular.*

Proof. Let Z' be the square complex represented in Figure 5. As in Figure 3, one has to merge the left and right sides, as well as the lower and the upper sides. Consider the map φ sending all vertices of Z' to the unique vertex of Z , and each edge of Z' to the unique edge of Z with the same color.

The complex Z' has 8 vertices, 32 edges, and 24 squares. In Z' , a 4-cycle is the boundary of a square if opposite edges have the same label (and direction) and if the colors of the boundary of this square correspond to the colors of the boundary of one of the three squares of Z . In the figure, the number (2 or 4) in the middle of each 4-cycle represent the number of squares of Z' on the vertices of this 4-cycle. This implies that φ is a covering map from Z' to Z .


Observe that two edges are dual to the same hyperplane of Z' if and only if they have the same label. Using this, it is easy to check that Z' is special.

By Theorem 7.5 and Proposition 7.6, we have that for any vertex $\tilde{v} \in \tilde{Z}'$, $\mathcal{F}(\tilde{v}, \tilde{Z}')$ is the domain of a regular trace event structure $\mathcal{E}_{Z'}$. Since \tilde{Z} and \tilde{Z}' coincide and since all vertices of $\tilde{Z} = \tilde{Z}'$ are lifts of the unique vertex of Z , $\mathcal{F}(\tilde{v}, \tilde{Z}')$ is independent of the choice of \tilde{v} . Consequently, $\mathcal{E}_Z = \mathcal{E}_{Z'}$ is a trace regular event structure. The fact that $\hat{\mathcal{E}}_Z$ is trace regular follows from Proposition 9.3. \square

Remark 10.3. Observe that Z' coincides with the special cube complex X_{N^*} of the net system N^* from Examples 3.1 and 8.3. Consequently, \mathcal{E}_Z coincides with the event structure \mathcal{E}_{N^*} obtained as the unfolding of N^* .

Lemma 10.4. *The graph $G(\mathcal{E}_Z)$ is hyperbolic. Consequently, the event structures \mathcal{E}_Z and $\hat{\mathcal{E}}_Z$ are grid-free.*

Proof. Since Z is a square complex, its universal cover is a CAT(0) square complex. Thus any isometric grid of $G(\mathcal{E}_Z)$ is a flat grid. Suppose by way of contradiction that $G(\mathcal{E}_Z)$ contains a

FIGURE 5. A finite special cover Z' of the complex Z .

large $n \times n$ flat grid Λ . Since Λ is flat, Λ is a convex and thus a gated subgraph of $G(\mathcal{E})$. Let \tilde{v} denotes the gate of \tilde{v}_0 in Λ . By Lemma 6.1 the direction of the edges of the graph $\vec{G}(\mathcal{E}_Z)$ coincide with the basepoint order $\leq_{\tilde{v}_0}$. This implies that the direction of the edges of the grid Λ in $\vec{G}(\mathcal{E}_Z)$ coincides with the basepoint order of Λ with \tilde{v} as the basepoint. In particular, this implies that \tilde{v} is the unique source of Λ . The vertex \tilde{v} together with one of the four corners of Λ span an $n' \times n''$ subgrid Λ' of Λ , where $n' \geq \frac{n}{2}$ and $n'' \geq \frac{n}{2}$.

Let $\vec{v}\tilde{u}$ and $\vec{v}\tilde{u}'$ be the two outgoing from \tilde{v} edges of Λ' . Consider the square Q of Λ' containing those two edges. Suppose without loss of generality that $\vec{v}\tilde{u}$ is upward vertical and $\vec{v}\tilde{u}'$ is horizontal and to the right. The vertex \tilde{v} is the unique source of Q . Denote by \tilde{w} the vertex of Q opposite to \tilde{v} . We will analyze in which way one can now extend the square Q to the grid Λ' . Notice that the square Q as well as any other square of Λ is one of the three squares Q_1, Q_2, Q_3 of the complex Z .

First suppose that Q is the square Q_1 of Z . Since Λ' is directed according to $\leq_{\tilde{v}}$, one can extend Q horizontally only by adding a new square Q_1 to the right. Also we can extend Q vertically only by adding the square Q_3 on the top of Q . But then we cannot extend the resulting union of three squares to a 2×2 grid because we have to set a square with source \tilde{w} , one orange (color a) outgoing edge and another red (color y) outgoing edge, however the tile-set $\{Q_1, Q_2, Q_3\}$ does not contain such a square.

Now suppose that Q is the square Q_3 of Z . Then we can extend Q only by setting Q_1 to the right. From the case when $Q = Q_1$ we know that we cannot extend Q_1 to a 2×2 grid. This show that we cannot extend Q to the 2×3 grid.

Finally, suppose that Q is the square Q_2 of Z . The single possibility to extend Q vertically is to set a copy of Q_3 on top. From the case when $Q = Q_3$, we know that we cannot extend Q_3 to a 2×3 grid. This show that we cannot extend Q to the 3×3 grid.

We deduce that in all cases we have $n' \leq 2$ or $n'' \leq 2$, establishing that if $G(\mathcal{E}_Z)$ contains an $n \times n$ isometric grid, then $n \leq 4$. This proves that $G(\mathcal{E}_Z)$ is hyperbolic. The graph $G(\dot{\mathcal{E}}_Z)$ is also hyperbolic because any grid of $G(\dot{\mathcal{E}}_Z)$ comes from a grid of $G(\mathcal{E}_Z)$. By Proposition 9.17, the event structures \mathcal{E}_Z and $\dot{\mathcal{E}}_Z$ are thus grid-free, since \mathcal{E}_Z and $\dot{\mathcal{E}}_Z$ have respectively degrees 4 and 5. \square

Lemma 10.5. *The graph $G(\mathcal{E}_Z)$ has infinite treewidth, i.e., the directed graph $\vec{G}(\mathcal{E}_Z)$ is not context-free. Consequently, the theories $\text{MSO}(\vec{G}(\mathcal{E}_Z))$, $\text{MSO}_2(G(\mathcal{E}_Z))$, and $\text{MSO}(\dot{\mathcal{E}}_Z)$ are undecidable.*

Proof. The proof of this assertion in some sense is similar to the proof of implication (4) \Rightarrow (5) of Theorem 9.1. As in the proof of the implication we will show that the graph $G(\mathcal{E}_Z)$ has the infinite half-grid $\frac{1}{2}\Gamma$ as a minor. We will also denote by $z_{i,j}, i, j \geq 0$, the vertices of $\frac{1}{2}\Gamma_n$ and by $Z_{i,j}, i, j \geq 0$, the connected subgraph of $G(\mathcal{E}_Z)$ which will be mapped (contracted) to $z_{i,j}$.

The subgraphs $Z_{i,j}$ are also paths laying in two consecutive spheres $S(\tilde{v}_0, k-1) \cup S(\tilde{v}_0, k)$. The difference is that in the proof of implication (4) \Rightarrow (5) of Theorem 9.1 we first constructed the union F^* of all fences in a downward way and then constructed the paths $Z_{i,j} \subset F^*$ in an upward way. For the current claim, we will build the paths $Z_{i,j}$ level-by-level, in an upward manner.

For this we use the fact that $\vec{G}(\mathcal{E}_Z)$ is the graph of the principal filter $\mathcal{D}_Z = (\mathcal{F}_{\tilde{o}}(\tilde{v}_0, \vec{Z}^{(1)}), \prec_{\tilde{o}})$ of the universal cover (\vec{Z}, \tilde{o}) of Z (here \tilde{v}_0 is an arbitrary but fixed lift of v_0). Since Z has one vertex v_0 , all vertices \tilde{v} of $\vec{G}(\mathcal{E}_Z)$ are lifts of v_0 . Analogously to v_0 , each such vertex \tilde{v} is incident to four outgoing and to four incoming colored edges in \vec{Z} . However, in the graph $\vec{G}(\mathcal{E}_Z)$ of the domain, each vertex \tilde{v} has at most two incoming edges (otherwise, there exists a 3-cube in the interval $I(\tilde{v}_0, \tilde{v})$, but this is impossible since \vec{Z} is 2-dimensional). The four outgoing edges define three squares Q_1, Q_2, Q_3 having \tilde{v} as the source (for an illustration, see the leftmost figure in Fig. 6). Moreover, (\vec{Z}, \tilde{o}) satisfies the following determinism property: if two edges $\vec{e}^j, \vec{e}^{j'}$ outgoing from a vertex \tilde{v} of \vec{Z} have the same color as the edges outgoing from the source of a square Q_i of Z , then \vec{e}^j and $\vec{e}^{j'}$ belong in (\vec{Z}, \tilde{o}) to a Q_i -square. Using this fact, one can see that there exists an infinite directed path P_a with \tilde{v}_0 as the origin and in which all edges have color orange (color a). Analogously, there exists an infinite directed path P_y with \tilde{v}_0 as the origin and in which all edges have color red (color y). Since Z is a VH-complex, the paths P_a and P_y are locally-convex paths of \vec{Z} . Since $G(\mathcal{E})$ is median, by Lemma 4.2, P_a and P_y are convex paths, thus shortest paths, of $G(\mathcal{E})$. Let $P_a = (\tilde{u}_0 = \tilde{v}_0, \tilde{u}_1, \tilde{u}_2, \dots)$ and $P_y = (\tilde{v}_0, \tilde{v}_1, \tilde{v}_2, \dots)$ (recall again that all vertices of these paths as well as all vertices of $G(\mathcal{E})$ are lifts of v_0).

We continue with the following auxiliary claim, which will be used in the definition of paths $Z_{i,j}$:

Claim 10.6. *For any vertex $\tilde{v} \in S(\tilde{v}_0, k-1)$, for any outgoing edges $\vec{v}\tilde{u}, \vec{v}\tilde{u}'$, there exist $0 < p \leq 4$ distinct vertices $\tilde{u}_1 = \tilde{u}, \tilde{u}_2, \dots, \tilde{u}_p = \tilde{u}' \in S(\tilde{v}_0, k)$ and $p-1$ distinct vertices $\tilde{w}_1, \dots, \tilde{w}_{p-1} \in S(\tilde{v}_0, k+1)$ such that for every i , $\vec{u}_i\tilde{w}_i$ and $\vec{u}_i\tilde{w}_{i-1}$ are directed edges of $\vec{G}(\mathcal{E}_Z)$, and such that the following holds:*

- if $\vec{v}\tilde{u}$ and $\vec{v}\tilde{u}'$ are colored with the same colors as the two outgoing edges from the source of a square $Q \in \{Q_1, Q_2, Q_3\}$, then $p = 2$ and $\vec{u}\tilde{w}_1$ and $\vec{u}'\tilde{w}_1$ are colored as the corresponding edges of Q .
- if $\vec{v}\tilde{u}$ and $\vec{v}\tilde{u}'$ are colored respectively blue (color x) and red (color y), then $p = 3$ and $\vec{u}\tilde{w}_1, \vec{u}'\tilde{w}_2, \vec{v}\tilde{u}_2$ are colored respectively orange (color a), black (color b), and black (color b);
- if $\vec{v}\tilde{u}$ and $\vec{v}\tilde{u}'$ are colored respectively black (color b) and orange (color a), then $p = 3$ and $\vec{u}\tilde{w}_1, \vec{u}'\tilde{w}_2, \vec{v}\tilde{u}_2$ are colored respectively red (color y), red (color y), and blue (color x);
- if $\vec{v}\tilde{u}$ and $\vec{v}\tilde{u}'$ are colored respectively orange (color a) and red (color y), then $p = 4$ and $\vec{u}\tilde{w}_1, \vec{u}_2\tilde{w}_1, \vec{u}_3\tilde{w}_3, \vec{u}'\tilde{w}_3, \vec{v}\tilde{u}_2, \vec{v}\tilde{u}_3$ are colored respectively red (color y), orange (color a), blue (color x), orange (color a), blue (color x), and black (color b).

Proof. In the universal cover (\vec{Z}, \tilde{o}) of Z , all vertices of $\vec{G}(\mathcal{E}_Z)$ are lifts of the unique vertex v_0 of Z . Consequently, each of them has four outgoing edges colored with the four different colors. Taking this into account, the proof of the claim follows from Figure 6. \square

For each k , we construct iteratively a simple path $P_k = P(\tilde{u}_k, \tilde{v}_k) = (\tilde{u}_k = \vec{p}_{k,1}, \vec{q}_{k,1}, \dots, \vec{p}_{k,\ell-1}, \vec{q}_{k,\ell-1}, \vec{p}_{k,\ell} = \tilde{v}_k)$ such that $\vec{q}_{k,1}\vec{p}_{k,1}$ is colored orange (color a), $\vec{q}_{k,\ell-1}\vec{p}_{k,\ell}$ is colored red (color y), and for each i , $\vec{p}_{k,i} \in S(\tilde{v}_0, k)$ and $\vec{q}_{k,i} \in S(\tilde{v}_0, k-1)$. This path plays a role similar to the one of the path $P_2(u_k, v_k)$ in the proof of Theorem 9.1.

Let $P_1 = (\tilde{u}_1, \tilde{v}_0, \tilde{v}_1)$ and suppose that the simple path $P_k = P(\tilde{u}_k, \tilde{v}_k)$ has been defined. We define the path $P_{k+1} = P(\tilde{u}_{k+1}, \tilde{v}_{k+1})$ in two steps. First, let P'_{k+1} be the path obtained by concatenating the paths obtained by applying Claim 10.6 to each vertex $q_{k,i}$ of $P_k \cap S(\tilde{v}_0, k)$ and




FIGURE 6. To the proof of Claim 10.6

its two outgoing edges in P_k . Note that the first edges of P_k and P'_{k+1} are consecutive edges in a square Q of $\vec{G}(\mathcal{E}_Z)$. Since the first edge of P_k is orange (color a), necessarily $Q = Q_1$ and the first edge of P'_{k+1} is red (color y). Analogously, the last edges of P_k and P'_{k+1} are consecutive edges in a square Q' of $\vec{G}(\mathcal{E}_Z)$. Since the last edge of P_k is red (color y), then necessarily $Q' = Q_3$ and the last edge of P'_{k+1} is orange (color a).


Claim 10.7. P'_{k+1} is a simple path.

Proof. Let $P'_{k+1} = (\tilde{u}_{k+1} = \tilde{q}_{k+1,1}, \tilde{p}_{k+1,1}, \dots, \tilde{q}_{k+1,\ell'-1}, \tilde{p}_{k+1,\ell'-1}, \tilde{q}_{k+1,\ell'} = \tilde{v}_{k+1})$. Suppose first that there exists $i < j$ such that $\tilde{q}_{k+1,i} = \tilde{q}_{k+1,j}$. By convexity of P_a and P_y , we have $2 \leq i < j \leq \ell' - 1$. Since the path P_k is simple, $\tilde{q}_{k+1,i}$ and $\tilde{q}_{k+1,j}$ cannot both belong to P_k . First suppose that one of them belongs to P_k , say $\tilde{q}_{k+1,i}$. By construction, $\tilde{q}_{k+1,j}$ has a neighbor $\tilde{q}_{k,j'} \in P_k$ and $\tilde{q}_{k+1,i}$ has two distinct neighbors $\tilde{q}_{k,i'}, \tilde{q}_{k,i'+1} \in P_k$. Since $\tilde{q}_{k+1,j} = \tilde{q}_{k+1,i}$ has at most two incoming edges in $\vec{G}(\mathcal{E}_Z)$, we get that $\tilde{q}_{k,j'} = \tilde{q}_{k,i'}$ or $\tilde{q}_{k,j'} = \tilde{q}_{k,i'+1}$. Since the path P_k is simple, it means that $i' \leq j' \leq i' + 1$, but this is impossible by the construction of P'_{k+1} and Claim 10.6. Now suppose that both vertices $\tilde{q}_{k+1,i}$ and $\tilde{q}_{k+1,j}$ do not belong to P_k . Then, by construction, $\tilde{q}_{k+1,i}$ has a neighbor $\tilde{q}_{k,i'} \in P_k$ and $\tilde{q}_{k+1,j}$ has a neighbor $\tilde{q}_{k,j'} \in P_k$. Moreover, by Claim 10.6, the arcs $\tilde{q}_{k,i'} \tilde{q}_{k+1,i}$ and $\tilde{q}_{k,j'} \tilde{q}_{k+1,j}$ are blue (color x) or black (color b). Since the path P_k is simple, these edges are distinct and thus have distinct colors. By the quadrangle condition, these two edges are incident to the sink $\tilde{q}_{k+1,i} = \tilde{q}_{k+1,j}$ of a square Q . However, there is no square in Z (or in \tilde{Z}) where the sink is incident to a black and a blue edge (See Figure 4).

Assume now that there exist $i < j$ such that $\tilde{p}_{k+1,i} = \tilde{p}_{k+1,j}$. By the construction, $\tilde{p}_{k+1,i}$ is adjacent to $\tilde{q}_{k+1,i}, \tilde{q}_{k+1,i+1}$ and $\tilde{p}_{k+1,j}$ is adjacent to $\tilde{q}_{k+1,j}, \tilde{q}_{k+1,j+1}$. By the previous case, these four vertices are distinct. Consequently, $\tilde{p}_{k+1,i} = \tilde{p}_{k+1,j}$ has four incoming edges, which is impossible since \tilde{Z} is 2-dimensional. \square

The path $P_{k+1} = P(\tilde{u}_{k+1}, \tilde{v}_{k+1})$ is obtained from P'_{k+1} by concatenating the orange (color a) edge $\tilde{u}_k \tilde{u}_{k+1}$ at the beginning of P'_{k+1} and the red (color y) edge $\tilde{v}_k \tilde{v}_{k+1}$ at the end of P'_{k+1} . Since $\tilde{u}_{k+1} \in P_a$ and P_a is a convex path, \tilde{u}_{k+1} cannot coincide with any vertex of P'_{k+1} . For the same reason, \tilde{v}_{k+1} is different from any vertex of P'_{k+1} and different from \tilde{u}_{k+1} . Consequently, the path P_{k+1} is a simple path.

Now, for each k , we construct iteratively the paths $Z_{i,j}$ with $i + j = k$ by selecting subpaths of $P(u_k, v_k)$. We require that the paths $Z_{i,j}$ satisfy the following properties (See Figures 7 and 8):

FIGURE 7. Construction of the path $Z_{i+1,j+1}$.

- (1) $Z_{0,j} = \{\tilde{u}_j\}$ and $Z_{i,0} = \{\tilde{v}_i\}$ for each $i, j \in \{0, \dots, n\}$;
- (2) for each k , if $i + j = k$, then $Z_{i,j}$ is a subpath of P_k ;
- (3) for each i, j with $i + j = k - 1$, the last vertex $\tilde{v}_{i,j+1}^r$ of the path $Z_{i,j+1}$ appears in P_k before the first vertex $\tilde{v}_{i+1,j}^l$ of $Z_{i+1,j}$;
- (4) each $Z_{i,j}$ with $i, j \geq 1$ has its two end-vertices in $S(\tilde{v}_0, k)$ and its first edge is orange (color a) and its last edge is blue (color x);
- (5) for each pair (i, j) with $i + j = k$, the leftmost vertex $\tilde{v}_{i,j}^l$ of the path $Z_{i,j}$ is adjacent to the rightmost vertex $\tilde{v}_{i,j+1}^r$ of the path $Z_{i,j+1}$ by an orange (color a) edge belonging to P_{k+1} and the rightmost vertex $\tilde{v}_{i,j}^r$ of $Z_{i,j}$ is adjacent to the leftmost vertex $\tilde{v}_{i+1,j}^l$ of the path $Z_{i+1,j}$ by a red (color y) edge belonging to P_{k+1} ;
- (6) any two distinct paths $Z_{i,j}$ and $Z_{i',j'}$ are disjoint.

Recall that the path P'_2 is obtained by applying Claim 10.6 to $P_1 = (\tilde{u}_1, \tilde{v}_0, \tilde{v}_1)$ and that P'_2 is a $(\tilde{u}_1, \tilde{v}_1)$ -path starting with a red edge (color y) and ending by an orange edge (color a). Let $Z_{1,1}$ be the path obtained from P'_2 by removing these two edges. It is easy to see that Conditions (1)-(3) and (5)-(6) hold, and Condition (4) holds by the construction of P'_2 and Claim 10.6.

Suppose now that the paths $Z_{i,j}$ satisfying the previous conditions have been defined for all pairs (i, j) such that $i + j \leq k + 1$ and we have to define the paths $Z_{i,j}$ with $i + j = k + 2$ (See Figure 7 for an illustration of the construction described below).

By induction hypothesis, the edge $\tilde{v}_{i,j}^l \tilde{v}_{i,j+1}^r$ is orange (color a) and the edge $\tilde{v}_{i,j}^r \tilde{v}_{i+1,j}^l$ is red (color y). These two edges belong to P_{k+1} . There exists \tilde{u}^l such that $\tilde{v}_{i,j}^l \tilde{v}_{i,j+1}^r$ and $\tilde{v}_{i+1,j}^l \tilde{u}^l$ are consecutive in a square Q^l of \tilde{Z} . Consequently, $\tilde{v}_{i,j+1}^r \tilde{u}^l$ is red (color y) and the opposite edge of $\tilde{v}_{i,j}^l \tilde{v}_{i,j+1}^r$ in Q^l is orange (color a). Note that by construction, this edge also belongs to P_{k+1} . Analogously, there exists \tilde{u}^r such that $\tilde{v}_{i,j}^r \tilde{v}_{i+1,j}^l$ and $\tilde{v}_{i+1,j}^l \tilde{u}^r$ are consecutive in a square Q^r of \tilde{Z} . Consequently, $\tilde{v}_{i+1,j}^l \tilde{u}^r$ is orange (color a) and the opposite edge of $\tilde{v}_{i,j}^r \tilde{v}_{i+1,j}^l$ in Q^r is blue (color x). Note that by construction, this edge also belongs to P_{k+1} . We let $Z_{i+1,j+1}$ be the subpath of P_{k+2} comprised between \tilde{u}^l and \tilde{u}^r . By the construction of P_{k+2} and the properties of Q^l and Q^r , the first edge of $Z_{i+1,j+1}$ is orange and the last one is blue, i.e., $Z_{i+1,j+1}$ satisfies Conditions (2) and (4). Observe that $Z_{i,j+1}$ and $Z_{i+1,j}$ also satisfy Condition (5).

We continue with Condition (3). Consider any path $Z_{i',j'}$ with $i' + j' = i + j + 2 = k + 2$ and without loss of generality assume that $i + 1 < i'$. From the construction, $Z_{i+1,j+1}$ is comprised in the $(\tilde{u}_{k+2}, \tilde{v}_{i+1,j}^l)$ -subpath of P_{k+2} while $Z_{i',j'}$ is comprised in the $(\tilde{v}_{i+1,j}^r, \tilde{v}_{k+2})$ -subpath of P_{k+2} . Since these two subpaths are disjoint by Claim 10.7, $Z_{i+1,j+1}$ is disjoint from $Z_{i',j'}$. This establishes Condition (3).

It remains to show that Condition (6) holds, i.e., that the path $Z_{i+1,j+1}$ is disjoint from any other path $Z_{i',j'}$ with $i' + j' \leq i + j + 2 = k + 2$. If $i' + j' = i + j = k + 2$, this follows from Condition (3). If $i' + j' \leq k$, then this is trivially true since $Z_{i',j'} \subseteq P_k \subseteq S(\tilde{v}_0, k - 1) \cup S(\tilde{v}_0, k)$







FIGURE 8. The half-grid resulting from the contraction of the paths $Z_{i,j}$.

and $Z_{i+1,j+1} \subseteq P_{k+2} \subseteq S(\tilde{v}_0, k+1) \cup S(\tilde{v}_0, k+2)$. Assume finally that $i' + j' = k+1$. In this case if $Z_{i',j'}$ and $Z_{i+1,j+1}$ have a common vertex \tilde{w} , then $\tilde{w} \in S(\tilde{v}_0, k+1)$. If $i' \leq i$, by induction hypothesis, $Z_{i',j'}$ is comprised in the $(\tilde{u}_{k+1}, \tilde{v}_{i',j+1}^r)$ -subpath of P_{k+1} . Note that by construction all the vertices of $P_{k+1} \cap S(\tilde{v}_0, k+1)$ that appear before $\tilde{v}_{i',j+1}^r$ in P_{k+1} also appear before $\tilde{v}_{i',j+1}^r$ in P_{k+2} . Since all vertices of $Z_{i+1,j+1}$ appear after $\tilde{v}_{i',j+1}^r$ in P_{k+2} and since P_{k+2} is a simple path, necessarily $Z_{i+1,j+1}$ is disjoint from $Z_{i',j'}$. If $i' \geq i+2$, we obtain the same result by a symmetric argument. \square

Consequently, by Lemma 10.4, the event structure $\dot{\mathcal{E}}_Z$ is grid-free and by Lemma 10.5, $\text{MSO}(\dot{\mathcal{E}}_Z)$ is undecidable. This concludes the proof of Theorem 10.1.

Remark 10.8. By construction, the event structure $\dot{\mathcal{E}}_Z$ is strongly regular, and $\dot{\mathcal{E}}_Z$ is hyperbolic by Lemma 10.4. However, $\dot{\mathcal{E}}_Z$ is not strongly regular hyperbolic because \tilde{Z} (and thus $\tilde{\mathcal{Z}}$) is not hyperbolic. Indeed, in \tilde{Z} , it is possible to build an infinite grid by repeating the pattern described in Figure 9. Due to the orientation of the edges of this grid, it is easy to see that this grid cannot appear in any principal filter of (\tilde{Z}, \tilde{o}) . Consequently, \tilde{Z} is not hyperbolic, but any principal filter of \tilde{Z} is hyperbolic.

This leads to the following open question: Can one construct a finite directed special complex X such that \tilde{X} is hyperbolic and some principal filter of \tilde{X} is not context-free?

FIGURE 9. Part of a infinite grid in \tilde{Z} FIGURE 10. The tile that we recursively insert to build the domain of \mathcal{E}_{BDR} and the first four steps of the construction.FIGURE 11. The hyperplanes obtained during the first four steps of the construction of the domain of \mathcal{E}_{BDR} .

10.2. Another counterexample to Conjecture 3.4. Another counterexample to Conjecture 3.4 can be derived from the hairing $\hat{\mathcal{E}}_{BDR}$ of the trace regular event structure \mathcal{E}_{BDR} described by Badouel et al. [6, pp. 144–146 and Fig. 5–9]. The domain of \mathcal{E}_{BDR} is a plane graph defined recursively as a tiling of the quarterplane with origin v_0 by tiles consisting of two squares sharing an edge (see Figure 10, left). Namely, we start with two infinite directed paths with common origin v_0 , and at each step, we insert the tile in each free angle (see Figure 10, right for the tiling obtained after the first four steps). As observed in [6], the hyperplanes of $G(\mathcal{E}_{BDR})$ can be represented by an arrangement of axis-parallel pseudolines in the plane (see Figure 11).




FIGURE 12. The dependence relation of the the trace alphabet used to label \mathcal{E}_{BDR} .




FIGURE 13. The trace labeling of the events (hyperplanes) of \mathcal{E}_{BDR} obtained during the first five steps of the construction.

Badouel et al. [6] showed that the directed graph $\vec{G}(\mathcal{E}_{BDR})$ is not context-free. Indeed, for each k , there is a unique level k cluster coinciding with the sphere $S(v_0, k)$ of radius k and the diameters of spheres increase together with their radius. By Theorem 9.1, this shows that the graph $G(\mathcal{E}_{BDR})$ has infinite treewidth. On the other hand one can easily show that the planar graph $G(\mathcal{E}_{BDR})$ has bounded hyperbolicity. Indeed, suppose by way of contradiction that $G(\mathcal{E}_{BDR})$ contains a 3×3 isometric square grid Γ . Since the cube complex of $G(\mathcal{E}_{BDR})$ is 2-dimensional, Γ is a convex and thus gated subgraph of $G(\mathcal{E}_{BDR})$. Let v be the gate of v_0 in Γ . Then Γ contains a 2×2 directed grid Γ' having v as a source. Let v' be the center of Γ' and observe that v' has two incoming and two outgoing arcs. Since $G(\mathcal{E}_{BDR})$ is planar, the four squares of Γ' around v' are the unique faces of the planar graph $G(\mathcal{E}_{BDR})$ incident to v' . Consequently, v' is the source of only one square in Γ' and thus in $G(\mathcal{E}_{BDR})$. But in $G(\mathcal{E}_{BDR})$, each inner vertex is the source of two distinct squares (defined by the three outgoing edges at v'), a contradiction. By Proposition 9.17, the event structure \mathcal{E}_{BDR} is grid-free since \mathcal{E}_{BDR} has degree 3.

Finally, the fact that \mathcal{E}_{BDR} admits a regular trace labeling was established in [6]. The trace alphabet has 10 letters $\{1, \dots, 5, 6, \dots, 10\}$ and the independence relation is defined as the complement of the reflexive Petersen's graph represented in Figure 12.¹ The labeling of the events (hyperplanes) of \mathcal{E}_{BDR} is given in Figure 13 for the events obtained during first five steps of the construction. The idea is that the events constructed at step $4i + 1$ are labeled consecutively from left to right $1, 5, 4, 3, 2, 1, \dots, 1, 5$, those constructed at step $4i + 2$ are labeled $6, 10, 9, 8, 7, 6, \dots, 10, 9$, those constructed at step $4i + 3$ are labeled $1, 2, 3, 4, 5, \dots, 4, 5$, and those

¹In [6], the graph of the dependency relation is the disjoint union of two 5-cycles. It gives a regular nice labeling, but not a trace labeling. However, once we add the edges $1 - 6, 2 - 8, 3 - 10, 4 - 7$ and $5 - 9$ to the dependency graph, the construction of Badouel et al. [6] gives a trace regular labeling of \mathcal{E}_{BDR} .

constructed at step $4i + 4$ are labeled $6, 7, 8, 9, 10, \dots, 8, 9$. A tedious check of the construction shows that this labeling gives 40 types of labeled principal filters².

Consequently, \mathcal{E}_{BDR} is a grid-free trace regular event structure whose graph $G(\mathcal{E}_{BDR})$ has infinite treewidth. By Theorem 9.4, the MSO theory $\text{MSO}(\dot{\mathcal{E}}_{BDR})$ of the hairing of \mathcal{E}_{BDR} is undecidable.

Remark 10.9. By Corollary 8.6, the domain of \mathcal{E}_{BDR} is the principal filter of the universal cover of some finite (virtually) special cube complex. However, we do not even have an explicit construction of a small NPC square complex X_{BDR} such that the domain of \mathcal{E}_{BDR} is a principal filter of the universal cover of X_{BDR} .

In view of Remark 10.8, one can ask whether there exists such an X_{BDR} that has a hyperbolic universal cover.

We also do not know if the hairing operation is necessary in order to obtain grid-free trace regular event structures with undecidable MSO theories. In particular, we wonder whether $\text{MSO}(\mathcal{E}_Z)$ and $\text{MSO}(\mathcal{E}_{BDR})$ are decidable. If this is not the case, this would provide counterexamples to Conjecture 3.4 that are not based on encoding MSO formulas over the domain by MSO formulas over the hair events.

11. CONCLUSION

The three Thiagarajan’s conjectures were a driving force in authors research for a long time. Our motivation to work on those conjectures was their intrinsic beauty and fundamental nature (finite versus infinite and decidability versus undecidability, both expressible in combinatorial way) and also our expertise in median graphs and $\text{CAT}(0)$ cube complexes. This expertise allowed us to work with the domain of the event structure instead of the event structure itself and perform geometric operations on the domain which preserve the property to be median or $\text{CAT}(0)$. This also allowed us to use the rich and deep theory of median graphs, $\text{CAT}(0)$ cube complexes, and, more importantly, of special cube complexes. We strongly believe that those three ingredients are essential in the understanding of Thiagarajan’s conjectures.

Even if we found counterexamples to the three Thiagarajan’s conjectures, the work on them raised many interesting open questions and lead to a better understanding of trace regularity and to a surprising link between 1-safe Petri nets and finite special cube complexes. We think that the characterization of trace regular event structures provided by this bijection can be viewed as a positive answer to Thiagarajan’s Conjecture 3.3. The open questions related to the first two conjectures are presented in the papers [13] and [19]. The Questions 3 and 4 from [19] are related to the nice labeling conjecture and to the embedding question (which seems to be easier than the nice labeling question). The questions and conjectures from [13] describe several conjectured classes of event structures for which Conjecture 3.3 is true (hyperbolic and confusion-free event structures) and relate the decidability of existence of finite regular nice labelings with that of decidability of the question of whether a finite cube complex is virtually special, an open question formulated in [2] and [10].

We conclude this paper with a speculation about Question 2.3 and Conjecture 3.4. Our counterexample to Conjecture 3.4 shows that grid-freeness of a trace regular event structure \mathcal{E}_N does not implies the decidability of $\text{MSO}(\mathcal{E}_N)$. On the other hand, we proved that decidability of $\text{MSO}(\vec{G}(\mathcal{E}_N))$ (or of $\text{MSO}(G(\mathcal{E}_N))$) is equivalent to finite treewidth and implies the decidability of $\text{MSO}(\mathcal{E}_N)$. We also showed that the decidability of $\text{MSO}(\vec{G}(\mathcal{E}_N))$ is equivalent with the decidability of $\text{MSO}(\dot{\mathcal{E}}_N)$, where $\dot{\mathcal{E}}_N$ is the hairing of \mathcal{E}_N . On the other hand, we know that conflict-free event structures (which may have infinite treewidth) have decidable MSO theory. Therefore, in the attempt to correct the formulation of Conjecture 3.4, we think that it is necessary to define the “haired” version of the treewidth of the domain of the event structure \mathcal{E}_N . We know that bounded treewidth is characterized by bounded square grid minors. In a similar way, a *haired grid minor* of $G(\mathcal{E}_N)$ is a minor of $G(\mathcal{E}_N)$ which is a haired square grid. A *haired square grid* is a square grid in which to each vertex is added a pendant edge (hair). Notice

²In [6], only 20 types of labeled principal filters are mentioned

that the hairs are edges of $G(\mathcal{E}_N)$ and thus correspond to events of \mathcal{E}_N . We require additionally that in a haired square grid minor, the hairs correspond to pairwise conflicting events. The *haired treewidth* of $G(\mathcal{E}_N)$ is the supremum of the sizes of haired square grid minors of $G(\mathcal{E}_N)$. We wonder whether the MSO theory of a trace regular event structures \mathcal{E}_N is decidable whenever the graph $G(\mathcal{E}_N)$ has finite haired treewidth.

Note. After the completion of this paper, Didier Caucal informed us that he also has constructed a counterexample to Conjecture 3.4.

Acknowledgements. The work on this paper was supported by ANR project DISTANCIA (ANR-17-CE40-0015).

REFERENCES

- [1] I. Agol. The virtual Haken conjecture. *Doc. Math.*, 18:1045–1087, 2013. With an appendix by Ian Agol, Daniel Groves, and Jason Manning.
- [2] I. Agol. Virtual properties of 3-manifolds. *Proceedings of the International Congress of Mathematicians, Seoul*, pages 157–186, 2014.
- [3] F. Ardila, M. Owen, and S. Sullivant. Geodesics in CAT(0) cubical complexes. *Adv. Appl. Math.*, 48(1):142 – 163, 2012.
- [4] M.R. Assous, V. Bouchitté, C. Charretton, and B. Rozoy. Finite labelling problem in event structures. *Theor. Comput. Sci.*, 123:9–19, 1994.
- [5] S.P. Avann. Metric ternary distributive semi-lattices. *Proc. Amer. Math. Soc.*, 12(3):407 – 414, 1961.
- [6] E. Badouel, Ph. Darondeau, and J.-C. Raoult. Context-free event domains are recognizable. *Inf. Comput.*, 149(2):134–172, 1999.
- [7] H.-J. Bandelt and V. Chepoi. Metric graph theory and geometry: a survey. In J. E. Goodman, J. Pach, and R. Pollack, editors, *Surveys on Discrete and Computational Geometry: Twenty Years Later*, volume 453 of *Contemp. Math.*, pages 49–86. AMS, Providence, RI, 2008.
- [8] H.-J. Bandelt and M. van de Vel. Embedding topological median algebras in products of dendrons. *Proc. London Math. Soc.*, s3-58(3):439 – 453, 1989.
- [9] J.-P. Barthélemy and J. Constantin. Median graphs, parallelism and posets. *Discr. Math.*, 111(1-3):49–63, 1993.
- [10] M. Bridson and H. Wilton. On the recognition problem for vitrually special cube complexes. *arXiv preprint*, 2017. [arXiv:1408.2325v2](https://arxiv.org/abs/1408.2325v2).
- [11] M.R. Bridson and A. Haefliger. *Metric Spaces of Non-Positive Curvature*, volume 319 of *Grundlehren der mathematischen Wissenschaften*. Springer-Verlag, Berlin, 1999.
- [12] J.P. Burling. *On coloring problems of prototypes*. PhD thesis, Univ. of Colorado, 1965.
- [13] J. Chalopin and V. Chepoi. A counterexample to Thiagarajan’s conjecture on regular event structures. *arXiv preprint*, 2016. [arXiv:1605.08288](https://arxiv.org/abs/1605.08288).
- [14] J. Chalopin and V. Chepoi. A counterexample to Thiagarajan’s conjecture on regular event structures. In *ICALP*, volume 80 of *LIPICs*, pages 101:1–101:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017.
- [15] V. Chepoi. Classification of graphs by means of metric triangles. *Metody Diskret. Analiz.*, (49):75–93, 96.
- [16] V. Chepoi. Graphs of some CAT(0) complexes. *Adv. Appl. Math.*, 24(2):125–179, 2000.
- [17] V. Chepoi. Nice labeling problem for event structures: a counterexample. *SIAM J. Comput.*, 41(4):715–727, 2012.
- [18] V. Chepoi, F.F. Dragan, B. Estellon, M. Habib, and Y Vaxès. Diameters, centers, and approximating trees of delta-hyperbolic geodesic spaces and graphs. In *Symposium on Computational Geometry*, pages 59–68. ACM, 2008.
- [19] V. Chepoi and M. F. Hagen. On embeddings of CAT(0) cube complexes into products of trees via colouring their hyperplanes. *J. Comb. Theory, Ser. B*, 103(4):428–467, 2013.
- [20] B. Courcelle. The monadic second-order logic of graphs ii: Infinite graphs of bounded width. *Math. Syst. Theory*, 21:187–221, 1989.
- [21] B. Courcelle. The monadic second-order logic of graphs vi: On several representations of graphs by relational structures. *Discr. Appl. Math.*, 54:117–149, 1994.
- [22] B. Courcelle and J. Engelfriet. *Graph Structure and Monadic Second-Order Logic*, volume 138 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, Cambridge, 2012.
- [23] D.Ž. Djoković. Distance-preserving subgraphs of hypercubes. *J. Comb. Theory, Ser. B*, 14(3):263–267, 1973.
- [24] M. Gromov. Hyperbolic groups. In S. M. Gersten, editor, *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987.
- [25] M. F. Hagen. Weak hyperbolicity of cube complexes and quasi-arboreal groups. *J. Topology*, 7(2):385–418, 2014.
- [26] F. Haglund. *Aspects combinatoires de la théorie géométrique des groupes*. HDR, U. Paris Sud, 2008.
- [27] F. Haglund and D.T. Wise. Special cube complexes. *Geom. Funct. Anal.*, 17(5):1551–1620, 2008.

- [28] F. Haglund and D.T. Wise. A combination theorem for special cube complexes. *Annals Math.*, 176(3):1427–1482, 2012.
- [29] A. Hatcher. *Algebraic Topology*. Cambridge University Press, Cambridge, 2002.
- [30] D. Kuske. Distributive lattices with a decidable monadic second order theory. *Algebra Univers.*, 48:183–207, 2002.
- [31] D. Kuske and M. Lohrey. Logical aspects of cayley-graphs: the group case. *Ann. Pure Appl. Logic*, 131:263–286, 2005.
- [32] P. Madhusudan. Model-checking trace event structures. In *LICS*, pages 371–380. IEEE Computer Society, 2003.
- [33] R. Morin. Concurrent automata vs. asynchronous systems. In *MFCS 2005*, volume 3618 of *LNCS*, pages 686–698. Springer, 2005.
- [34] H.M. Mulder. *The Interval Function of a Graph*, volume 132 of *Mathematical Centre tracts*. Mathematisch Centrum, Amsterdam, 1980.
- [35] D.M. Müller and P. Schupp. The theory of ends, pushdown automata, and second order logic. *Theor. Comput. Sci.*, 37:51–75, 1985.
- [36] M. Nielsen, G.D. Plotkin, and G. Winskel. Petri nets, event structures and domains, part I. *Theor. Comput. Sci.*, 13:85–108, 1981.
- [37] M. Nielsen, G. Rozenberg, and P.S. Thiagarajan. Transition systems, event structures and unfoldings. *Inf. Comput.*, 118(2):191–207, 1995.
- [38] M. Nielsen and P.S. Thiagarajan. Regular event structures and finite Petri nets: the conflict-free case. In *ICATPN 2002*, volume 2360 of *LNCS*, pages 335–351. Springer, 2002.
- [39] N. Robertson and P. Seymour. Graph minors ii. algorithmic aspects of tree-width. *J. Algorithms*, 78:309–322, 1986.
- [40] N. Robertson and P. Seymour. Graph minors. v. excluding a planar graph. *J. Combin. Th., Ser. B*, 41:92–114, 1986.
- [41] M. Roller. Poc sets, median algebras and group actions. Technical report, Univ. of Southampton, 1998.
- [42] B. Rozoy and P.S. Thiagarajan. Event structures and trace monoids. *Theor. Comput. Sci.*, 91(2):285–313, 1991.
- [43] M. Sageev. Ends of group pairs and non-positively curved cube complexes. *Proc. London Math. Soc.*, s3-71(2):585–617, 1995.
- [44] L. Santocanale. A nice labelling for tree-like event structures of degree 3. *Inf. Comput.*, 208:652–665, 2010.
- [45] D. Seese. The structure of models of decidable monadic theories of graphs. *Ann. Pure Appl. Logic*, 53:169–195, 1991.
- [46] P.S. Thiagarajan. Regular trace event structures. Technical Report BRICS RS-96-32, Computer Science Department, Aarhus University, Aarhus, Denmark, 1996.
- [47] P.S. Thiagarajan. Regular event structures and finite Petri nets: A conjecture. In *Formal and Natural Computing*, volume 2300 of *LNCS*, pages 244–256. Springer, 2002.
- [48] P.S. Thiagarajan and S. Yang. Rabin’s theorem in the concurrency setting: a conjecture. *Theor. Comput. Sci.*, 546:225–236, 2014.
- [49] M. van de Vel. *Theory of convex structures*, volume 50 of *North-Holland Mathematical Library*. North-Holland Publishing Co.
- [50] P.M. Winkler. Isometric embedding in products of complete graphs. *Discr. Appl. Math.*, 7(2):221–225, 1984.
- [51] G. Winskel. *Events in computation*. PhD thesis, Edinburgh Univ., 1980.
- [52] G. Winskel and M. Nielsen. Models for concurrency. In S. Abramsky, Dov M. Gabbay, and T. S. E. Maibaum, editors, *Handbook of Logic in Computer Science (Vol. 4)*, pages 1–148. Oxford University Press, 1995.
- [53] D.T. Wise. *Non-positively curved squared complexes, aperiodic tilings, and non-residually finite groups*. PhD thesis, Princeton University, 1996.
- [54] D.T. Wise. Complete square complexes. *Comment. Math. Helv.*, 82(4):683–724, 2007.

LABORATOIRE D’INFORMATIQUE ET SYSTÈMES, CNRS, AIX-MARSEILLE UNIVERSITÉ, AND UNIVERSITÉ DE TOULON

Email address: jeremie.chalopin@lis-lab.fr

LABORATOIRE D’INFORMATIQUE ET SYSTÈMES, AIX-MARSEILLE UNIVERSITÉ, CNRS, AND UNIVERSITÉ DE TOULON

Email address: victor.chepoi@lis-lab.fr