

Explicit Hopcroft’s Trick in Categorical Partition Refinement

Takahiro Sanada¹[0000–0003–3409–6963], Ryota Kojima¹, Yuichi Komorida^{2,3}[0000–0002–3371–524], Koko Muroya¹[0000–0003–1760–988], and Ichiro Hasuo^{2,3}[0000–0002–8300–4650]

¹ RIMS, Kyoto University, Japan

{tsanada, kojima, kmuroya}@kurims.kyoto-u.ac.jp

² National Institute of Informatics, Japan

³ SOKENDAI, Tokyo, Japan

komorin@nii.ac.jp i.hasuo@acm.org

Abstract. Algorithms for *partition refinement* are actively studied for a variety of systems, often with the optimisation called *Hopcroft’s trick*. However, the low-level description of those algorithms in the literature often obscures the essence of Hopcroft’s trick. Our contribution is twofold. Firstly, we present a novel formulation of Hopcroft’s trick in terms of general trees with weights. This clean and explicit formulation—we call it *Hopcroft’s inequality*—is crucially used in our second contribution, namely a general partition refinement algorithm that is *functor-generic* (i.e. it works for a variety of systems such as (non-)deterministic automata and Markov chains). Here we build on recent works on coalgebraic partition refinement but depart from them with the use of *fibrations*. In particular, our fibrational notion of *R-partitioning* exposes a concrete tree structure to which Hopcroft’s inequality readily applies. It is notable that our fibrational framework accommodates such algorithmic analysis on the categorical level of abstraction.

Keywords: Partition refinement · Category theory · Coalgebra · Fibration · Tree algorithm

1 Introduction

Partition refinement refers to a class of algorithms that computes behavioural equivalence of various types of systems—such as the language equivalence for deterministic finite automata (DFAs), bisimilarity for labelled transition systems (LTSs) and Markov chains, etc.—by a fixed-point iteration. Such algorithms also yield *quotients* of state spaces, making systems smaller and thus easier to analyse e.g. by model checking.

Since its original introduction by Moore [20] for DFAs, partition refinement has been actively studied for enhanced performance and generality. On the performance side, Hopcroft [11] introduced what is now called *Hopcroft’s trick* that greatly improves the asymptotic complexity. The original paper [11] is famously

hard to crack; works such as [6, 16] present its reformulation, again focusing on DFAs. On the generality side, partition refinement for systems other than DFAs has been pursued, such as LTSs [15, 24], probabilistic transition systems with non-determinism [7], weighted automata [18], and weighted tree automata [9, 10]. Hopcroft’s trick is used in many of these works for enhanced performance, too.

Such a variety of target systems is uniformly addressed by a recent body of work on *coalgebraic partition refinement* [4, 5, 14, 26]. Here, a target system is identified with a categorical construct called *coalgebra* $c: C \rightarrow FC$ (see e.g. [13]), where C represents the state space, the *functor* F specifies the type of the system, and c represents the dynamics. By changing the functor F as a parameter, the theory accommodates many different systems such as DFAs and weighted automata. The coalgebraic partition refinement algorithms in [4, 5, 14, 26] are *functor-generic*: they apply uniformly to such a variety of systems.

The current work is inspired by [14] which successfully exploits Hopcroft’s trick for generic coalgebraic partition refinement. In [14], their coalgebraic algorithm is described in parallel with its set-theoretic (or even binary-level) concrete representations, letting the latter accommodate Hopcroft’s trick. Their experimental results witnessed its superior performance, beating some existing tools that are specialised in a single type of systems.

However, the use of Hopcroft’s trick in [14] is formulated in low-level set-theoretic terms, which seems to obscure the essence of the algorithm as well as the optimisation by Hopcroft’s trick, much like in the original paper [11]. Therefore, in this paper, we aim at 1) an explicit formulation of Hopcroft’s trick, and 2) a categorical partition refinement algorithm that exposes an explicit data structure to which Hopcroft’s trick applies.

We achieve these two goals in this paper: 1) an explicit formulation that we call *Hopcroft’s inequality*, and 2) a categorical algorithm that uses a *fibration*. Here is an overview.

Hopcroft’s Inequality We identify *Hopcroft’s inequality* (Thm. 2.9) as the essence of Hopcroft’s trick. Working on general trees with a general notion of vertex weight, it uses the classification of edges into *heavy* and *light* ones and bounds a sum of weights in terms of (only) the root and leaf weights. This inequality can be used to bound the complexity of many tree generation algorithms, including those for partition refinement.

This general theory can accommodate different weights. We exploit this generality to systematically derive partition refinement algorithms with different complexities (§6.2).

A Fibrational Partition Refinement Algorithm Hopcroft’s inequality does not directly apply to the existing coalgebraic partition refinement algorithms [4, 5, 14, 26] since the latter do not explicitly present a suitable tree structure. To address this challenge, we found the categorical language of *fibrations* [12] to be a convenient vehicle: it allows us to speak about the relationship between 1) an equivalence relation (an object in a fibre category) and 2) a partitioning of a state space (a mono-sink in the base category). The outcome is a partition refinement algorithm that is both *abstract* (it is functor-generic and applies to a variety of

systems) and *concrete* (it explicitly builds a tree to which Hopcroft’s inequality applies.) Our development relies on the fibrational theory of bisimilarity [8, 17]; yet ours is the first fibrational partition refinement algorithm.

More specifically, in a fibration $p: \mathbb{E} \rightarrow \mathbb{C}$, an equivalence relation R on a set X is identified with an object $R \in \mathbb{E}_X$ in the fibre over X (consider the well-known fibration $\mathbf{EqRel} \rightarrow \mathbf{Set}$ of sets and equivalence relations over them). We introduce a categorical notion of R -partitioning; it allows $R \in \mathbb{E}_X$ to induce a mono-sink (i.e. a family of monomorphisms) $\{\kappa_i: C_i \rightarrow C\}_{i \in I}$. The latter is identified with the set of R -equivalence classes.

Fig. 1 illustrates one iteration of our fibrational partition refinement algorithm \mathbf{fPR}^H (Algo. 2). In the last step (Fig. 1c), the object C_{01} is divided into three parts C_{010} , C_{011} and C_{012} along $(c \circ \kappa)^* \overline{F}R$. We call the mono-sink $\{C_{01i} \rightarrow C_{01}\}_{i \in \{0,1,2\}}$ the $(c \circ \kappa)^* \overline{F}R$ -partitioning of C_{01} . In this manner, a tree structure explicitly emerges in the base category \mathbb{C} . Hopcroft’s inequality directly applies to this tree, allowing us to systematically present the Hopcroft-optimisation on the categorical level of abstraction.

We note that, at this moment, our fibrational framework (with a fibration $p: \mathbb{E} \rightarrow \mathbb{C}$) has only one example, namely the fibration $\mathbf{EqRel} \rightarrow \mathbf{Set}$ of equivalence relations over sets. While it is certainly desirable to have other examples, their absence does not harm the value of our fibrational framework: we do not use fibrations for additional generality (beyond functor-genericity);⁴ we use them to explicate trees in the base category (cf. Fig. 1).

Contributions Summarising, our main technical contributions are as follows.

- *Hopcroft’s inequality* that explicates the essence of Hopcroft’s trick.
- A fibrational notion of R -partitioning that turns a fibre object into a mono-sink (§4).
- A fibrational partition refinement algorithm \mathbf{fPR}^H that combines the above two (§6.1).
- Functor-generic partition refinement algorithms $\mathbf{fPR}_{w_C}^{H-ER}$, $\mathbf{fPR}_{w_P}^{H-ER}$, $\mathbf{fPR}_{w_R}^{H-ER}$, obtained as instances of \mathbf{fPR}^H but using different weights in Hopcroft’s inequality. The three achieve slightly different, yet comparable to the best known, complexity bounds (§6.2).

2 Hopcroft’s Inequality

We present our first contribution, *Hopcroft’s inequality*. It is a novel formalisation of Hopcroft’s trick in terms of rooted trees. It also generalises the trick, accommodating arbitrary *weights* (Def. 2.2) besides the particular one to count the number of items in classes that is typically and widely used (e.g. [11, 14, 15, 24]).

Notation 2.1. Let T be a rooted tree. We denote the set of leaves by $L(T)$, the set of vertices by $V(T)$, the set of edges in the path from v to u by $\text{path}(v, u)$,

⁴ In this sense, we can say that our use of fibrations is similar to some recent usages of string diagrams in *specific* monoidal categories, such as in [2, 22].




Fig. 1: An iteration in our algorithm fPR^H (Algo. 2). Fig. 1a shows an equivalence relation R over C , and the corresponding partitioning $C_{00}, C_{01}, C_1 \twoheadrightarrow C$ of the state space C . (The history of refinement is recorded as a tree; this is important for complexity analysis.) In Fig. 1b, the equivalence relation R is refined into $c^* \overline{FR}$ along the one-step transition of the system dynamics c , and is further restricted to the partition C_{01} . In Fig. 1c, the resulting equivalence relation $(c \circ \kappa)^* \overline{FR}$ over C_{01} yields a partitioning of C_{01} , expanding the tree.

the set of children of $v \in V(T)$ by $\text{ch}(v)$, and the subtree whose root is $v \in V(T)$ by $\text{tr}(v)$.

Definition 2.2 (weight function). Let T be a rooted finite tree. A *weight function* of T is a map $w: V(T) \rightarrow \mathbb{N}$ satisfying $\sum_{u \in \text{ch}(v)} w(u) \leq w(v)$ for each $v \in V(T)$. We call a weight function *tight* if $\sum_{u \in \text{ch}(v)} w(u) = w(v)$ for all $v \in V(T) \setminus L(T)$.

Definition 2.3 (heavy child choice). For a weight function w of a tree T , a *heavy child choice* (hcc for short) is a map $h: V(T) \setminus L(T) \rightarrow V(T)$ satisfying $h(v) \in \text{ch}(v)$ and $w(h(v)) = \max_{u \in \text{ch}(v)} w(u)$ for every $v \in V(T) \setminus L(T)$. We write $h(v)$ as h_v and call the vertex h_v a *heavy child* of v , and a non-heavy child a *light child*. We define $\text{lch}_h(v) = \text{ch}(v) \setminus \{h_v\}$. An edge (v, u) is a *light edge* if $u \in \text{lch}_h(v)$. We define $\text{lpath}(v, u) = \{e \in \text{path}(v, u) \mid e \text{ is a light edge}\}$.

Note that a heavy child choice always exists but is not unique in general.

Examples are in Fig. 2; the weight on the left is not tight while the right one is tight.




Fig. 2: Examples of rooted trees, each with a weight function and an hcc. The heavy children are indicated by thick edges. Thin edges are light edges.

In the rest of this section, our technical development is towards Hopcroft’s inequality in Thm. 2.9. It gives an upper bound for a sum of weights—we only count those for light children, which is the core of the optimisation in Hopcroft’s

partition refinement algorithm [11]—in terms of weights of the root and the leaves. This upper bound makes no reference to the tree’s height or internal weights, making it useful for complexity analysis of tree generation algorithms.

The following lemma crucially relies on the definition of weight function.

Lemma 2.4. *Let T be a finite tree with a root r , w be a weight function of T , and S be an arbitrary set of edges of T . Then $\sum_{v \in V(T)} \sum_{\substack{u \in \text{ch}(v) \\ (v,u) \notin S}} w(u) \geq \sum_{l \in L(T)} |\text{path}(r,l) \setminus S| \cdot w(l)$ holds. The equality holds when w is tight.*

Lem. 2.5 is our first key lemma; we use Lem. 2.4 in its proof. It relates the sum of weights of the light children—for which we aim to give an upper bound in Thm. 2.9—with the leaf weights and (roughly) the tree height.

Lemma 2.5. *Let T be a finite tree with a root r , w be a weight function of T , and h be an hcc for w . Then the following inequality holds. The equality holds when w is tight.*

$$\sum_{v \in V(T)} \sum_{u \in \text{ch}_h(v)} w(u) \geq \sum_{l \in L(T)} |\text{lpath}(r,l)| \cdot w(l). \quad (1)$$

For the right tree in Fig. 2, the left-hand side of (1) is $(14+7)+5+(2+3)+0+2 = 33$, and the right-hand side is $1 \times 5 + 0 \times 10 + 1 \times 9 + 2 \times 2 + 2 \times 3 + 2 \times 2 + 1 \times 5 = 33$.

The inequality in (1) is the opposite of what we want (namely an upper bound for the left-hand side). We thus force an equality using *tightening*.

Definition 2.6 (tightening). Let w be a weight function of a rooted finite tree T , and h be its heavy child choice. The *tightening* $w' : V(T) \rightarrow \mathbb{N}$ of w along h is defined recursively by

$$w'(u) = \begin{cases} w(u) & \text{if } u \text{ is the root of } T \\ w'(v) - \sum_{u' \in \text{ch}_h(v)} w(u') & \text{if } u = h_v \text{ for the parent } v \text{ of } u \\ w(u) & \text{otherwise.} \end{cases}$$

In Fig. 2, the weight function of the right tree is a tightening of that of the left tree. We observe that tightening maintains a heavy child choice:

Lemma 2.7. *Let T be a rooted finite tree, w be a weight function of T , h be an hcc for w , and w' be the tightening of w along h . The following hold.*

1. *The map w' is a tight weight function of T .*
2. *The map h is also a heavy child choice for w' .*
3. *For the root r , $w(r) = w'(r)$ holds, and for each $v \in V(T)$, $w(v) \leq w'(v)$ holds.*

Our second key lemma towards Thm. 2.9 is as follows, bounding $|\text{lpath}(r,v)|$ that occurs on the right in (1). Its proof is by what is commonly called *Hopcroft’s trick* [1, 11]: it observes that, along a light edge, weights decay at least by $1/2$.

Lemma 2.8. *Let T be a finite tree with a root r , w be a weight of T , and h be an hcc for w . For each vertex $v \in V(T)$ with $w(v) \neq 0$, the following inequality holds: $|\text{lpath}(r,v)| \leq \log_2 w(r) - \log_2 w(v)$.*

We combine Lem. 2.8 and Lem. 2.5 (its equality version; we can use it via tightening) to obtain Hopcroft’s inequality. It bounds a sum of weights by the root and leaf weights.

Theorem 2.9 (Hopcroft’s inequality). *Let T be a finite tree with root r , w be a weight function of T , and h be a heavy child choice for w . The following inequality holds.*

$$\sum_{v \in V(T)} \sum_{u \in \text{lch}_h(v)} w(u) \leq w(r) \log_2 w(r) - \sum_{\substack{l \in L(T) \\ w(l) \neq 0}} w(l) \log_2 w(l) \quad (2)$$

For complexity analysis, we use Hopcroft’s inequality in the following form. Assume that a tree generation algorithm takes $t(v)$ time to generate all the children (both heavy and light) of v . If there exists K such that $t(v)$ is bounded by K times the sum of all light children, then the time to generate the whole tree is bounded by $Kw(r) \log_2 w(r)$.

Corollary 2.10. *Let T be a rooted finite tree with root r , w be a weight function of T , and h be a heavy child choice for w . If a map $t: V(T) \rightarrow \mathbb{N}$ satisfies that there exists a constant $K \in \mathbb{N}$ such that $t(v) \leq K \sum_{u \in \text{lch}_h(v)} w(u)$ for every $v \in V(T)$, then the sum of $t(v)$ is bounded by $Kw(r) \log_2 w(r)$, that is $\sum_{v \in V(T)} t(v) \leq Kw(r) \log_2 w(r)$.*

Remark 2.11. Further adaptations of Hopcroft’s trick are pursued in the literature, e.g. in [25], where the notion of heavy child choice is relaxed with an extra parameter $\alpha \in [1/2, 1)$. Our theory can easily be extended to accommodate α , in which case the above description corresponds to the special case with $\alpha = 1/2$. Details are deferred to another venue.

3 Categorical Preliminaries

The rest of the paper is about our second contribution, namely a functor-generic partition refinement (PR) algorithm optimised by an explicit use of Hopcroft’s inequality (Thm. 2.9). It is given by our novel formulation of coalgebraic PR algorithms in fibrational terms. Here we shall review some necessary categorical preliminaries.

We use categorical formalisation of intersections and unions.

Definition 3.1. For monomorphisms $m: A \rightarrow C$ and $n: B \rightarrow C$ in \mathbb{C} , the *intersection* $m \cap n: A \cap B \rightarrow C$ and the *union* $m \cup n: A \cup B \rightarrow C$ are defined by the pullback and pushout, respectively:

$$\begin{array}{ccc} A \cap B & \xrightarrow{\pi_2} & B \\ \pi_1 \downarrow & \lrcorner & \downarrow n \\ A & \xrightarrow{m} & C \end{array} \quad \text{and} \quad \begin{array}{ccc} A \cap B & \xrightarrow{\pi_2} & B \\ \pi_1 \downarrow & \lrcorner & \downarrow \\ A & \longrightarrow & A \cup B \end{array} .$$

We say $m: A \rightarrow C$ and $m': A' \rightarrow C$ are *equivalent* if there is an isomorphism $\phi: A \rightarrow A'$ such that $m = m' \circ \phi$. The set $\mathbf{Sub}(\mathbb{C})_C$ of equivalence classes of monomorphisms whose codomains are C forms a lattice, assuming enough limits and colimits.

3.1 Fibrations

A fibration $p: \mathbb{E} \rightarrow \mathbb{C}$ is a functor satisfying some axioms. When $p(R) = C$ for an object $R \in \mathbb{E}$ and an object $C \in \mathbb{C}$, we see that R equips C with some information, e.g. a predicate, a relation, a topology, etc. The main example in this paper is the fibration $\mathbf{EqRel} \rightarrow \mathbf{Set}$ where C is a set and R is an equivalence relation over C .

Fibrational constructs that are the most relevant to us are the *inverse image* $f^*(R')$ and the *direct image* $f_*(R)$ along a morphism $f: S \rightarrow S'$ in \mathbb{C} . In the case of $\mathbf{EqRel} \rightarrow \mathbf{Set}$, these are computed as follows.

$$\begin{array}{ccc}
 f^*(R') = \{(x, y) \mid (fx, fy) \in R'\} & \xleftarrow{f^*} & R' \\
 \vdots & & \vdots \\
 S & \xrightarrow{f} & S'
 \end{array}
 \qquad
 \begin{array}{ccc}
 R \vdash f_* & f_*(R) = \{(fx, fy) \mid (x, y) \in R\} & \\
 \vdots & & \vdots \\
 S & \xrightarrow{f} & S'
 \end{array}$$

In what follows we introduce some basics of fibrations; they formalise the intuition above. For details, see e.g. [12].

Definition 3.2 (fibration). Let $p: \mathbb{E} \rightarrow \mathbb{C}$ be a functor. A morphism $f: P \rightarrow R$ in \mathbb{E} is *Cartesian* if for any $g: Q \rightarrow R$ in \mathbb{E} with $pg = pf \circ v$ for some $v: pQ \rightarrow pP$, there exists a unique $h: Q \rightarrow P$ in \mathbb{E} above v (i.e. $ph = v$) with $f \circ h = g$. The functor p is a *fibration* if for each $R \in \mathbb{E}$ and $u: C \rightarrow pR$ in \mathbb{C} , there are an object u^*R and a Cartesian morphism $\dot{u}(R): u^*R \rightarrow R$ in \mathbb{E} . See below.

$$\begin{array}{ccccc}
 & & Q & \xrightarrow{g} & R \\
 & & \downarrow h & \nearrow \dot{u}(R) & \downarrow \\
 \mathbb{E} & & u^*R & \xrightarrow{\dot{u}(R)} & R \\
 \downarrow p & & & & \\
 \mathbb{C} & & B & \xrightarrow{u \circ v} & D \\
 & & \downarrow v & \nearrow u & \\
 & & C & \xrightarrow{u} & D
 \end{array}$$

The category \mathbb{E} is called the *total category* and the category \mathbb{C} is called the *base category* of the fibration. For an object $C \in \mathbb{C}$, the objects in \mathbb{E} above C form a category \mathbb{E}_C , called the *fibre category* above C . The fibre category \mathbb{E}_C is the category of “equivalence relations” on C .

Definition 3.3 (fibre category). Let $p: \mathbb{E} \rightarrow \mathbb{C}$ be a fibration and $C \in \mathbb{C}$. The *fibre category* \mathbb{E}_C over C is the subcategory of \mathbb{E} whose objects are defined by $\text{ob}(\mathbb{E}_C) = \{R \in \mathbb{E} \mid pR = C\}$, and morphisms are defined by $\mathbb{E}_C(Q, R) = \{f \in \mathbb{E}(Q, R) \mid pf = \text{id}_C\}$ for $Q, R \in \text{ob}(\mathbb{E}_C)$.

Example 3.4 ($\mathbf{EqRel} \rightarrow \mathbf{Set}$ is a fibration). Let \mathbf{EqRel} be the category of equivalence relations. The objects of \mathbf{EqRel} are pairs (S, R) of a set S and an equivalence relation R on S . A morphism $f: (S, R) \rightarrow (S', R')$ in \mathbf{EqRel} is a function $f: S \rightarrow S'$ satisfying $(f(x), f(y)) \in R'$ for all $(x, y) \in R$. We sometimes write just R for (S, R) when no confusion arises. The functor $p: \mathbf{EqRel} \rightarrow \mathbf{Set}$ defined by $p(S, R) = S$ is a fibration.

For a morphism $u: C \rightarrow D$ in the base category of a fibration, the map $u^*: \text{ob}(\mathbb{E}_D) \rightarrow \text{ob}(\mathbb{E}_C)$ extends to a functor $u^*: \mathbb{E}_D \rightarrow \mathbb{E}_C$ between fibre categories. We call the functor $u^*: \mathbb{E}_D \rightarrow \mathbb{E}_C$ an *inverse image* functor.

Given a fibration $\mathbb{E} \rightarrow \mathbb{C}$ and an endofunctor $F: \mathbb{C} \rightarrow \mathbb{C}$ on the base category, if $R \in \mathbb{E}$ is above $C \in \mathbb{C}$, we would like to get an object in \mathbb{E} above FC . A lifting of F specifies the choice of an object above FC .

Definition 3.5 (lifting, fibred lifting). Let $p: \mathbb{E} \rightarrow \mathbb{C}$ be a fibration and $F: \mathbb{C} \rightarrow \mathbb{C}$ be a functor. A *lifting* of F is a functor $\bar{F}: \mathbb{E} \rightarrow \mathbb{E}$ with $p \circ \bar{F} = F \circ p$. A functor $\bar{F}: \mathbb{E} \rightarrow \mathbb{E}$ is a *fibred lifting* of F if \bar{F} is a lifting of F and preserves Cartesian morphisms.

When $\bar{F}: \mathbb{E} \rightarrow \mathbb{E}$ is a fibred lifting of $F: \mathbb{C} \rightarrow \mathbb{C}$, for $f: C \rightarrow D$ in \mathbb{C} and $R \in \mathbb{E}_D$, we have $\bar{F}(f^*R) = (Ff)^*(\bar{F}R)$ in \mathbb{E}_{FC} . An important example of a lifting is a relation lifting.

Definition 3.6 (relation lifting [13]). Let $F: \mathbf{Set} \rightarrow \mathbf{Set}$ be a weak pullback preserving functor. We define a lifting $\text{Rel}(F): \mathbf{EqRel} \rightarrow \mathbf{EqRel}$ of F along the fibration $p: \mathbf{EqRel} \rightarrow \mathbf{Set}$, called the *relation lifting* of F , as follows. For an object $(C, R) \in \mathbf{EqRel}$, there is the inclusion $\langle r_1, r_2 \rangle: R \hookrightarrow C \times C$. We define the relation lifting on the object R by $\text{Rel}(F)(R) = \text{Im}\langle Fr_1, Fr_2 \rangle$, where $\text{Im}\langle Fr_1, Fr_2 \rangle$ is the image factorisation. By the assumption that F preserves weak pullbacks, we can show that $\text{Rel}(F)(R)$ is an equivalence relation. $\text{Rel}(F)$ can be extended to a functor.

$$\begin{array}{ccc} FR & \xrightarrow{\langle Fr_1, Fr_2 \rangle} & FC \times FC \\ & \searrow & \nearrow \\ & \text{Im}\langle Fr_1, Fr_2 \rangle & \\ & = \text{Rel}(F)(R) & \end{array}$$

In this paper, we deal with a restricted class of fibrations, called \mathbf{CLat}_{\sqcap} -fibrations.

Definition 3.7 (\mathbf{CLat}_{\sqcap} -fibration). A fibration $p: \mathbb{E} \rightarrow \mathbb{C}$ is a \mathbf{CLat}_{\sqcap} -fibration if each fibre \mathbb{E}_C is a complete lattice and each inverse image functor $u^*: \mathbb{E}_D \rightarrow \mathbb{E}_C$ preserves meets \sqcap .

For a \mathbf{CLat}_{\sqcap} -fibration, there always exists the left adjoint $u_*: \mathbb{E}_C \rightarrow \mathbb{E}_D$ to an inverse image functor u^* , as is well-known (cf. Freyd's adjoint functor theorem). The functor u_* is defined by $u_*(P) = \sqcap \{R \in \mathbb{E}_D \mid P \sqsubseteq u^*(R)\}$ on objects. We call u_* a *direct image* functor.

Example 3.8 ($\mathbf{EqRel} \rightarrow \mathbf{Set}$ is a \mathbf{CLat}_{\sqcap} -fibration). The functor $p: \mathbf{EqRel} \rightarrow \mathbf{Set}$ from Example 3.4 is a \mathbf{CLat}_{\sqcap} -fibration. We describe the inverse image functor f^* and the direct image functor f_* for a function $f: S \rightarrow S'$. For an equivalence relation R' on S' , the inverse image $f^*(R')$ is the equivalence relation $\{(x, y) \in S \times S \mid (f(x), f(y)) \in R'\}$ on S . For an equivalence relation R on S , the direct image $f_*(R)$ is the *equivalence closure* of the relation $\{(f(x), f(y)) \in S' \times S' \mid (x, y) \in R\}$.

3.2 Coalgebras and Bisimulations

Coalgebras are widely used as a generalisation of state-based systems [13, 23].

Definition 3.9 (*F-coalgebra*). Let \mathbb{C} be a category and $F: \mathbb{C} \rightarrow \mathbb{C}$ be an endofunctor. An *F-coalgebra* is a pair (C, c) of an object $C \in \mathbb{C}$ and a morphism $c: C \rightarrow FC$.

For an *F-coalgebra* $c: C \rightarrow FC$, F specifies the type of the system, the carrier object C represents the “set of states” of the system, and c represents the transitions in the system. When $\mathbb{C} = \mathbf{Set}$, for an *F-coalgebra* $c: C \rightarrow FC$ and a state $x \in C$, the element $c(x) \in FC$ represents properties (e.g. acceptance) and successors of x .

A major benefit of coalgebras is that their theory is *functor-generic*: by changing a functor F , the same theory uniformly applies to a vast variety of systems.

Example 3.10. We describe some *F-coalgebras* for functors F on \mathbf{Set} .

1. For the powerset functor \mathcal{P} , a \mathcal{P} -coalgebra $c: C \rightarrow \mathcal{P}C$ is a *Kripke frame*. For a state $x \in C$, $c(x) \in \mathcal{P}C$ is the set of successors of x .
2. Let Σ be an alphabet and $N_\Sigma = 2 \times (\mathcal{P}\Sigma)^\Sigma$. An N_Σ -coalgebra $c: C \rightarrow N_\Sigma C$ is a *non-deterministic automaton* (NA). For a state $x \in C$, let $(b, t) = c(x) \in 2 \times (\mathcal{P}C)^\Sigma$. The state x is accepting iff $b = 1$, and there is a transition $x \xrightarrow{a} y$ in the NA iff $y \in t(a)$.
3. The distribution functor \mathcal{D} is defined on a set X to be $\mathcal{D}X = \{d: X \rightarrow [0, 1] \mid \{x \in X \mid d(x) \neq 0\} \text{ is finite and } \sum_{x \in X} d(x) = 1\}$. A \mathcal{D} -coalgebra $c: C \rightarrow \mathcal{D}C$ is a *Markov chain*. For a state x , $c(x) \in \mathcal{D}C$ is a probability distribution $C \rightarrow [0, 1]$, which represents the probabilities of transitions to successor states of x .

We are interested in how similar two states of a state-transition system are. We consider two states to be similar if one state can mimic the transitions of the other. *Bisimilarity* by Park [21] and Milner [19] is a notion that captures such behaviour of states. Hermida and Jacobs [8] formulated bisimilarity as a coinductive relation on a coalgebra, using a fibration.

Definition 3.11 (bisimulations and the bisimilarity). Let $p: \mathbb{E} \rightarrow \mathbb{C}$ be a \mathbf{CLat}_Γ -fibration, $F: \mathbb{C} \rightarrow \mathbb{C}$ be a functor, $c: C \rightarrow FC$ be an *F-coalgebra* and \overline{F} be a lifting of F . An (F, \overline{F}) -bisimulation is a $c^* \circ \overline{F}$ -coalgebra in \mathbb{E}_C , that is an object $R \in \mathbb{E}_C$ with $R \sqsubseteq c^*(\overline{F}(R))$ since a morphism in \mathbb{E}_C is a relation (\sqsubseteq). By the Knaster–Tarski theorem, there exists the greatest (F, \overline{F}) -bisimulation $\nu(c^* \circ \overline{F})$ with respect to the order of \mathbb{E}_C , and it is called the (F, \overline{F}) -bisimilarity.

In the above definition, the choice of \overline{F} determines a notion of bisimulation. The relation lifting $\text{Rel}(F)$ (Def. 3.6) is often used as a lifting of F . For all the functors we consider, the bisimilarity wrt. $\text{Rel}(F)$ coincides with the *behavioural equivalence*, another well-known notion of bisimilarity [13, §4.5].

Example 3.12 ($(F, \text{Rel}(F))$ -bisimilarities). We illustrate $(F, \text{Rel}(F))$ -bisimilarity (also called *logical F -bisimilarity* [13]) for F in Example 3.10. Let $C \in \mathbf{Set}$ and $R \in \mathbf{EqRel}_C$.

1. ($F = \mathcal{P}$). The $(\mathcal{P}, \text{Rel}(\mathcal{P}))$ -bisimilarity $\nu(c^* \circ \text{Rel}(\mathcal{P}))$ for a \mathcal{P} -coalgebra $c: C \rightarrow \mathcal{P}C$ is the maximum relation R on C such that if $(x, y) \in R$ then
 - for every $x' \in c(x)$, there is $y' \in c(y)$ such that $(x', y') \in R$, and
 - for every $y' \in c(y)$, there is $x' \in c(x)$ such that $(x', y') \in R$.
2. ($F = N_\Sigma$). The $(N_\Sigma, \text{Rel}(N_\Sigma))$ -bisimilarity $\nu(c^* \circ \text{Rel}(N_\Sigma))$ for an N_Σ -coalgebra $c: C \rightarrow N_\Sigma C$ is the ordinary bisimilarity for the NA c , that is the maximum relation R on C such that if $(x, y) \in R$ then $\pi_1(c(x)) = \pi_1(c(y))$ and
 - for each $a \in \Sigma, x' \in \pi_2(c(x))(a)$, there is $y' \in \pi_2(c(y))(a)$ such that $(x', y') \in R$, and
 - for each $a \in \Sigma$ and $y' \in \pi_2(c(y))(a)$, there is $x' \in \pi_2(c(x))(a)$ such that $(x', y') \in R$.
3. ($F = \mathcal{D}$) [3]. The $(\mathcal{D}, \text{Rel}(\mathcal{D}))$ -bisimilarity $\nu(c^* \circ \text{Rel}(\mathcal{D}))$ for a \mathcal{D} -coalgebra $c: C \rightarrow \mathcal{D}C$ is the maximum relation R such that if $(x, y) \in R$ then $\sum_{z \in K} c(x)(z) = \sum_{z \in K} c(y)(z)$ for every equivalence class $K \subseteq C$ of R .

4 Fibrational Partitioning

We introduce the notion of fibrational partitioning, one that is central to our algorithm that grows a tree using fibre objects (cf. Fig. 1).

Given an “equivalence relation” R over C —identified with an object $R \in \mathbb{E}_C$ over $C \in \mathbb{C}$ in a suitable fibration $p: \mathbb{E} \rightarrow \mathbb{C}$ —a fibrational R -partitioning is a mono-sink, shown on the right, that is subject to certain axioms. The notion allows us to explicate equivalence classes (namely $\{C_i\}_i$) in the abstract fibrational language.

$$\begin{array}{ccc}
 & & R \quad \mathbb{E} \\
 C_0 & \xrightarrow{\kappa_0} & \downarrow \downarrow p \\
 \vdots & & C \quad \mathbb{C} \\
 C_n & \xrightarrow{\kappa_n} &
 \end{array}$$

Definition 4.1 (R -partitioning). Let \mathbb{C} be a category with pullbacks and an initial object 0 , and $p: \mathbb{E} \rightarrow \mathbb{C}$ be a \mathbf{CLat}_Γ -fibration. Let $C \in \mathbb{C}$ and $R \in \mathbb{E}_C$. An R -partitioning is a mono-sink (i.e. a family of monomorphisms) $\{\kappa_i: C_i \rightarrow C\}_{i \in I}$ that satisfies:

1. $\kappa_i^*(R) = \top_{C_i}$ for all $i \in I$,
2. $\bigsqcup_{i \in I} (\kappa_i)_*(\top_{C_i}) = R$, and
3. $C_i \not\cong 0$ and $C_i \cap C_j \cong 0$ for each $i, j \in I$ with $i \neq j$.

We say a \mathbf{CLat}_Γ -fibration p admits partitioning if (1) for each $C \in \mathbb{C}$ and $R \in \mathbb{E}_C$, there is an R -partitioning; and moreover, (2) for each $C \in \mathbb{C}$, $R, R' \in \mathbb{E}_C$ such that $R' \sqsubseteq R$, and each R -partitioning $\{\kappa_i: C_i \rightarrow C\}_{i \in I}$, we have $\bigsqcup_{i \in I} (\kappa_i)_*(\kappa_i^* R') = R'$.

$$\begin{array}{ccc}
 \text{(Beck–Chevalley)} & & \text{(modularity)} \\
 \begin{array}{ccc}
 A \cap B \xrightarrow{n} B & & \mathbb{E}_{A \cap B} \xleftarrow{n^*} \mathbb{E}_B \\
 m \downarrow \lrcorner \quad \downarrow \lambda & \implies & m^* \downarrow \quad \downarrow \lambda^* \\
 A \xrightarrow{\kappa} C & & \mathbb{E}_A \xleftarrow{\kappa^*} \mathbb{E}_C
 \end{array} & & \begin{array}{ccc}
 \mathbb{E}_A \times \mathbb{E}_B \xrightarrow{\kappa^* \times \lambda^*} \mathbb{E}_C \times \mathbb{E}_C & \xrightarrow{\sqcup} & \mathbb{E}_C \\
 \lambda^* \times \lambda^* \downarrow & & \downarrow \lambda^* \\
 \mathbb{E}_B \times \mathbb{E}_B & \xrightarrow{\sqcup} & \mathbb{E}_B
 \end{array}
 \end{array}$$

Fig. 3: Conditions for Assum. 5.1.

Cond. 3 asserts that the components C_i are nontrivial and disjoint. Cond. 1 says the partitioning $\{C_i\}_i$ is *not too coarse*—the original equivalence R , when restricted to C_i , should relate all pairs of elements in C_i . Conversely, Cond. 2 means that $\{C_i\}_i$ is *not too fine*—if it were finer than R , then the relation $\bigsqcup_{i \in I} (\kappa_i)_*(\top_{C_i})$ over C would be finer than R . See the concrete description of $(\kappa_i)_*$ in Example 3.8.

Example 4.2 (**EqRel** \rightarrow **Set** admits partitioning). **EqRel** \rightarrow **Set** admits partitioning. Indeed, given an equivalence relation $R \in \mathbf{EqRel}_C$ over C , the monosink $\{\kappa_S: S \rightarrow C\}_{S \in C/R}$, where $S \in C/R$ is naturally identified with a subset of C , is an R -partitioning. Cond. 1–3 are easily verified following Example 3.8.

An R -partitioning is not necessarily unique. This happens when $R \in \mathbf{EqRel}_C$ has singleton equivalence classes. Let $A \subseteq C$ be an arbitrary subset such that each $x \in A$ composes a singleton R -equivalence class. Then $\{\kappa'_S: S \rightarrow C\}_{S \in I}$, where $I = (C/R) \setminus \{\{x\} \mid x \in A\}$, is also an R -partitioning. With this monosink (that is “narrower” than the original $\{\kappa_S\}_{S \in C/R}$), Cond. 2 is satisfied since the equivalence closure operation included in the direct images $(\kappa_i)_*(\top_{C_i})$ (see Example 3.8) compensates the absence of $x \in A$.

The fibration **EqRel** \rightarrow **Set** is our leading example, and unfortunately, the only example that we know admits partitioning. There are many other examples of **CLat** $_{\top}$ -fibrations (see [17]), but they fail to admit partitioning, typically due to the failure of Cond. 2 of Def. 4.1. This absence of examples does not harm the value of our fibrational framework: our goal is to explicate categorical essences of partition refinement; and we do not aim at new instances via categorical abstraction (although such are certainly desirable).

5 The Naive Fibrational Algorithm $\mathbf{fPR}^{\text{naive}}$

We introduce a naive fibrational partition refinement algorithm, called $\mathbf{fPR}^{\text{naive}}$, as a preparation step to our main algorithm \mathbf{fPR}^{H} (Algo. 2).

In what follows, a prefix-closed set $T \subseteq \mathbb{N}^*$ (where \mathbb{N}^* is the set of strings over \mathbb{N}) is identified with a rooted tree. We denote the leaves of T by $L(T)$.

We introduce further conditions that make fibrations “compatible” with partitioning. It is easy to see that **EqRel** \rightarrow **Set** satisfies the conditions on p in Assum. 5.1.

Assumption 5.1. Assume a **CLat** $_{\top}$ -fibration $p: \mathbb{E} \rightarrow \mathbb{C}$ that satisfies the following conditions.

1. For each $C \in \mathbb{C}$, the lattice $\mathbf{Sub}(\mathbb{C})_C$ of subobjects of C in \mathbb{C} is distributive.
2. (Beck–Chevalley) For every pullback diagram along monomorphisms in \mathbb{C} , shown in the first diagram in Fig. 3, the induced diagram, the second in Fig. 3, commutes.
3. For any monomorphisms $\kappa: A \rightarrow C$ and $\lambda: B \rightarrow C$, the third diagram in Fig. 3 is a fork, where the following diagram is a *fork* if $h_1 \circ g_1 \circ f = h_2 \circ g_2 \circ f$.

$$\begin{array}{ccccc} W & \xrightarrow{f} & X & \xrightarrow{g_1} & Y_1 \\ & & g_2 \downarrow & & \downarrow h_1 \\ & & Y_2 & \xrightarrow{h_2} & Z \end{array}$$

Definition 5.2. Let $p: \mathbb{E} \rightarrow \mathbb{C}$ be a \mathbf{CLat}_Γ -fibration that satisfies Assum. 5.1, $F: \mathbb{C} \rightarrow \mathbb{C}$ be a functor, and $\overline{F}: \mathbb{E} \rightarrow \mathbb{E}$ be its lifting along p (Def. 3.5). Algo. 1 shows our *naive fibrational partition refinement algorithm*. Given a coalgebra $c: C \rightarrow FC$, it computes a $\nu(c^*\overline{F})$ -partitioning of C , i.e. modulo the (F, \overline{F}) -bisimilarity of c (Def. 3.11).

Algorithm 1 The naive fibrational partition refinement algorithm $\text{fPR}^{\text{naive}}$.

Input: A coalgebra $c: C \rightarrow FC$ in \mathbb{C} .

Output: A mono-sink $\{\kappa_i: C_i \rightarrow C\}_{i \in I}$ for some I .

- 1: $T := \{\epsilon\} \subseteq \mathbb{N}^*$; $C_\epsilon := C$; $\kappa_\epsilon := \text{id}_C: C_\epsilon \rightarrow C$; $R := \top_C$ ▷ initialisation
 - 2: **while** $c^*\overline{F}R \neq R$ **do** ▷ the main loop
 - 3: $R := c^*\overline{F}R$; $L := L(T)$
 - 4: **for** $\sigma \in L$ **do**
 - 5: Take a $\kappa_\sigma^*(R)$ -partitioning $\{\lambda_{\sigma,k}: C_{\sigma k} \rightarrow C_\sigma\}_{k \in \{0, \dots, n_\sigma\}}$ of C_σ
 - 6: **for** $k = 0, \dots, n_\sigma$ **do** $\kappa_{\sigma k} := \kappa_\sigma \circ \lambda_{\sigma,k}: C_{\sigma k} \rightarrow C$
 - 7: $T := T \cup \{\sigma 0, \dots, \sigma n_\sigma\}$
 - 8: **return** $\{\kappa_\sigma: C_\sigma \rightarrow C\}_{\sigma \in L(T)}$
-




Fig. 4: The R -partitioning gets finer as the algorithm runs.

Algo. 1 starts with $R = \top_C \in \mathbb{E}_C$ and a singleton family of a monomorphism $\{\kappa_\epsilon: C_\epsilon \rightarrow C\}$. With each iteration, the object R on C gets smaller and closer to $\nu(c^*\overline{F})$ and R -partitioning $\{\kappa_\sigma: C_\sigma \rightarrow C\}_\sigma$ gets finer (see Fig. 4). When the algorithm terminates, R is equal to $\nu(c^*\overline{F})$ and a $\nu(c^*\overline{F})$ -partitioning is returned.

Combining the loop invariant (Lem. 5.3) and termination (Lem. 5.4), we can prove the correctness of the naive algorithm.

Lemma 5.3 (loop invariant). *At the beginning of each iteration of the main loop, the following hold.*

1. The mono-sink $\{\kappa_\sigma: C_\sigma \rightarrow C\}_{\sigma \in L(T)}$ is an R -partitioning.
2. $\nu(c^*\overline{F}) \sqsubseteq R$.

Lemma 5.4 (termination). *If \mathbb{E}_C is a well-founded lattice, Algo. 1 terminates.*

Proposition 5.5 (correctness). *If \mathbb{E}_C is well-founded, then Algo. 1 terminates and returns $\nu(c^*\overline{F})$ -partitioning $\{\kappa: C_i \twoheadrightarrow C\}_{i \in I}$.*

6 Optimised Algorithms with Hopcroft's Inequality

Recall that the naive algorithm grows a tree *uniformly* so that every leaf has the same depth (see Fig. 4; note that, even if C_σ is fine enough, we extend the node by a trivial partitioning). By selecting leaves in a smart way and generating a tree selectively, the time cost of each iteration can be reduced, so that Hopcroft's inequality is applicable.

In §6.1, we present a functor-generic and fibrational algorithm enhanced with the Hopcroft-type optimisation, calling it fPR^H . We use Hopcroft's inequality (§2) for complexity analysis.

In §6.2 we instantiate fPR^H to the fibration $\mathbf{EqRel} \rightarrow \mathbf{Set}$, obtaining three concrete (yet functor-generic) algorithms $\text{fPR}_{w_C}^{H-ER}$, $\text{fPR}_{w_P}^{H-ER}$, $\text{fPR}_{w_R}^{H-ER}$ that use different weight functions. $\text{fPR}_{w_C}^{H-ER}$ is essentially the algorithm in [14]. The other two ($\text{fPR}_{w_P}^{H-ER}$, $\text{fPR}_{w_R}^{H-ER}$) use the weight functions from the works [6, 11, 16] on DFA partition refinement. The three algorithms exhibit slightly different asymptotic complexities.

6.1 A Fibrational Algorithm fPR^H Enhanced by Hopcroft's Inequality

We fix a \mathbf{CLat}_\square -fibration $p: \mathbb{E} \rightarrow \mathbb{C}$, functors $F: \mathbb{C} \rightarrow \mathbb{C}$ and $\overline{F}: \mathbb{E} \rightarrow \mathbb{E}$, an F -coalgebra $c: C \rightarrow FC$, and a map $w: \text{ob}(\mathbf{Sub}(\mathbb{C})_C) \rightarrow \mathbb{N}$ (which we use for weights). We write $w(C')$ for $w(\lambda: C' \twoheadrightarrow C)$ when no confusion arises.

The following conditions clarify which properties of $\mathbf{EqRel} \rightarrow \mathbf{Set}$ are necessary to make our optimised fibrational algorithm fPR^H work: the last one (Assum. 6.1.10) is for complexity analysis; all the other ones are for correctness.

- Assumption 6.1.**
1. \mathbb{C} has pullbacks, pushouts along monos, and an initial object 0 .
 2. The fibre category \mathbb{E}_0 above an initial object 0 is trivial, that is $\top_0 = \perp_0$.
 3. \overline{F} is a fibred lifting of F along p .
 4. $F: \mathbb{C} \rightarrow \mathbb{C}$ preserves monomorphisms whose codomain is not 0 .
 5. The fibration p admits partitioning.
 6. The fibration $p: \mathbb{E} \rightarrow \mathbb{C}$ satisfies the three conditions in Assum. 5.1.
 7. The fibre category \mathbb{E}_C is a well-founded lattice.
 8. If $C' \twoheadrightarrow C$ and $R \in \mathbb{E}_{C'}$, every R -partitioning $\{\lambda_k: D_k \twoheadrightarrow C'\}_{k \in K}$ is finite ($|K| < \infty$).
 9. If $\kappa: A \twoheadrightarrow C$ and $\lambda: B \twoheadrightarrow C$ are monomorphisms and $A \cap B \cong 0$, then the functor $\mathbb{E}_A \times \mathbb{E}_B \xrightarrow{\kappa_* \times \lambda_*} \mathbb{E}_C \times \mathbb{E}_C \xrightarrow{\sqcup} \mathbb{E}_C$ is injective on objects.

Table 1: Correspondence between constructs in $\text{fPR}_{(F,\bar{F}),w}^{\text{H}}$ and the theory in §2

Constructs in Algo. 2	Notions in §2
The tree of subobjects $\{C_\sigma\}_{\sigma \in J}$ of C , cf. Fig. 5	A tree $T = (V, E)$, cf. Fig. 2
A set $\{C_\sigma\}_{\sigma \in J}$ of objects in \mathbb{C}	A set V of vertices
A set $\{\kappa_{\sigma,k}: C_{\sigma k} \rightarrow C_\sigma\}_{\sigma k \in J}$ of monomorphisms	A set E of edges
A map $w: \text{ob}(\mathbf{Sub}(\mathbb{C})_C) \rightarrow \mathbb{N}$ with Assum. 6.1.10	A weight function $w: V \rightarrow \mathbb{N}$ (Def. 2.2)
A choice of k_0 made in Line 10 of Algo. 2	An hcc $h: V \setminus L(T) \rightarrow V$ (Def. 2.3)
The complexity result of Prop. 6.9	The inequality of Cor. 2.10

10. For $C' \rightarrow C$, $R \in \mathbb{E}_{C'}$, and R -partitioning $\{\kappa_i: C_i \rightarrow C'\}$ of C' , $\sum_{i=1}^n w(C_i) \leq w(C')$.

Assum. 6.1.3 is not overly restrictive. Indeed, the following functors on \mathbf{Set} have a fibred lifting. The functors described in Example 3.10 are examples of the functor defined by (3).

Proposition 6.2. *Consider the endofunctors on \mathbf{Set} defined by the BNF below.*

$$F ::= \text{Id} \mid A \mid \coprod_{b \in B} F_b \mid \prod_{b \in B} F_b \mid \mathcal{P}F \mid \mathcal{D}F \quad \text{where } A \text{ and } B \text{ are sets. (3)}$$

The relation lifting $\text{Rel}(F): \mathbf{EqRel} \rightarrow \mathbf{EqRel}$ of F (Def. 3.6) is fibred.

Proposition 6.3. *The fibration $p: \mathbf{EqRel} \rightarrow \mathbf{Set}$ with the relation lifting $\text{Rel}(F)$ of a functor F defined by (3) and a coalgebra $c: C \rightarrow FC$ for a finite set C satisfies the assumptions 1–9 of Assum. 6.1.*

Definition 6.4 (fPR^{H}). Let the \mathbf{CLat}_\square -fibration $p: \mathbb{E} \rightarrow \mathbb{C}$, the map $w: \text{ob}(\mathbf{Sub}(\mathbb{C})_C) \rightarrow \mathbb{N}$, functors $F: \mathbb{C} \rightarrow \mathbb{C}$ and $\bar{F}: \mathbb{E} \rightarrow \mathbb{E}$, and the object C satisfy Assum. 6.1. Algo. 2 shows the algorithm $\text{fPR}_{(F,\bar{F}),w}^{\text{H}}$ ((F,\bar{F}) and w are omitted when clear from the context). Given a coalgebra $c: C \rightarrow FC$, it computes a $\nu(c^*\bar{F})$ -partitioning of C , like the naive algorithm.

The algorithm fPR^{H} exposes a tree structure to which Hopcroft’s inequality applies. Table 1 summarises how constructs in fPR^{H} fit §2.

Much like $\text{fPR}^{\text{naive}}$ (Algo. 1), fPR^{H} grows a tree, as shown in Fig. 5. We take the generated tree as $T = (V, E)$ in §2. Note that, whereas $\text{fPR}^{\text{naive}}$ expands the tree uniformly so that every leaf has the same depth (Fig. 4), fPR^{H} expands leaves selectively (Fig. 5).

fPR^{H} chooses k_0 so that $w(C_{\rho k_0})$ is maximised (Line 10). These choices constitute a heavy child choice (Def. 2.3), an essential construct in Hopcroft’s inequality (Thm. 2.9).

fPR^{H} starts with $R = \top_C \in \mathbb{E}_C$, the singleton family of a monomorphism $\{\kappa_\epsilon = \text{id}_C: C_\epsilon \rightarrow C\}$, and a marking $C_\epsilon^{\text{cl}} = 0$ of states. For each σ , $C_\sigma^{\text{cl}} \rightarrow C_\sigma$

Algorithm 2 An optimised fibrational partition refinement algorithm $\text{fPR}_{(F,\bar{F}),w}^H$.

Input: A coalgebra $c: C \rightarrow FC$ in \mathbb{C} .

Output: A mono-sink $\{\kappa_i: C_i \rightarrow C\}_{i \in I}$ for some I .

```

1:  $J := \{\epsilon\} \subseteq \mathbb{N}^*$ ;  $C_\epsilon := C$ ;  $C_\epsilon^{\text{cl}} := 0$  ▷ initialisation
2: while there is  $\rho \in L(J)$  such that  $C_\rho^{\text{cl}} \neq C_\rho$  do ▷ the main loop
3:    $R := \bigsqcup_{\sigma \in L(J)} (\kappa_\sigma)_*(\top_{C_\sigma})$  ▷ Partitioning (Line 3–9)
4:   Choose a leaf  $\rho \in L(J)$  such that  $C_\rho^{\text{cl}} \neq C_\rho$ 
5:    $R_\rho := (c \circ \kappa_\rho)^*(\bar{F}(R))$ 
6:   if  $R_\rho = \top_{C_\rho}$  then
7:      $C_\rho^{\text{cl}} := C_\rho$ 
8:   continue
9:   Take an  $R_\rho$ -partitioning  $\{\kappa_{\rho,k}: C_{\rho k} \rightarrow C_\rho\}_{k \in \{0, \dots, n_\rho\}}$ 
10:  Choose  $k_0 \in \{0, \dots, n_\rho\}$  s.t.  $w(C_{\rho k_0}) = \max_{k \in \{0, \dots, n_\rho\}} w(C_{\rho k})$  ▷ Relabelling
    (Line 10–12)
11:  MARKDIRTY( $\rho, k_0$ )
12:   $J := J \cup \{\rho 0, \dots, \rho n_\rho\}$ 
13: return  $\{\kappa_\sigma: C_\sigma \rightarrow C\}_{\sigma \in L(J)}$ 

14: procedure MARKDIRTY( $\rho, k_0$ )
15:   for  $k \in \{0, \dots, n_\rho\}$  do  $C_{\rho k}^{\text{cl}} := C_{\rho k}$ 

16:   Let  $B$  be the pullback of the diagram:
      
$$\begin{array}{ccc}
 B & \xrightarrow{\quad} & C \\
 \downarrow & \lrcorner & \downarrow c \\
 F(C_{\rho k_0} \cup (\bigcup_{\sigma \in L(J) \setminus \{\rho\}} C_\sigma)) & \xrightarrow{\quad} & FC
 \end{array}$$

17:   ▷ the bottom morphism is mono by Assum. 6.1.4
18:   for  $\tau \in L(J \cup \{\rho 0, \dots, \rho n_\rho\})$  do
19:      $C_\tau^{\text{cl}} := C_\tau^{\text{cl}} \cap B$  ▷ states not in  $B$  are marked as dirty

```

is a “subset” of C_σ consisting of *clean* states; the rest of C_σ consists of *dirty* states. Therefore, initially, all states of $C = C_\epsilon$ are marked dirty. The main loop, consisting of **Partitioning** (Line 3–9) and **Relabelling** (Line 10–12), iterates until there is no dirty state (Line 2).

The **Partitioning** part selects one leaf C_ρ whose states include at least one dirty state (Line 4). The tree is expanded at this selected leaf only. This selection makes Algo. 2 different from Algo. 1, which expands the tree at every leaf (cf. Fig. 4 and Fig. 5).

The **Relabelling** part then updates the clean/dirty marking. Firstly, it chooses one “heavy child” $C_{\rho k_0}$ (Line 10) from the leaves generated in **Partitioning**. Then the iteration calls the MARKDIRTY procedure (Line 14–19). It first collects states (B in Line 16) whose all “successors” with respect to the coalgebra $c: C \rightarrow FC$ are in the object $C_{\rho k_0} \cup (\bigcup_{\sigma \in L(J) \setminus \{\rho\}} C_\sigma)$; the latter intuitively consists of states “unaffected” by tree expansion. The procedure marks only states in B as clean (Line 19), which means that the rest of the states are marked dirty.




Fig. 5: At each iteration one leaf of the tree is selected and refined.

Towards the correctness theorem of our optimised fibrational algorithm fPR^{H} (Thm. 6.8), we first make a series of preliminary observations.

Notation 6.5. We write R_i for R defined at Line 3 of Algo. 2 at the i -th iteration. We write J_i for J at the beginning of the i -th iteration. We write $C_\sigma^{\text{cl},i}$ and $\kappa_\sigma^{\text{cl},i}$ for C_σ^{cl} and the monomorphism $\kappa_\sigma^{\text{cl}}: C_\sigma^{\text{cl}} \rightarrow C$, respectively, at Line 16 at the i -th iteration.

We identify loop invariants Prop. 6.6. Termination of fPR^{H} follows from Assum. 6.1.7 and 9 (Prop. 6.7). Combining these, we prove the correctness of fPR^{H} in Thm. 6.8.

Proposition 6.6 (loop invariants). *At the beginning of the i -th iteration, the following hold.*

1. $(c \circ \kappa_\sigma \circ \kappa_\sigma^{\text{cl},i})^* \overline{F}(R_i) = \top_{C_\sigma^{\text{cl},i}}$ for each leaf $\sigma \in L(J_i)$.
2. The mono-sink $\{\kappa_\sigma: C_\sigma \rightarrow C\}_{\sigma \in L(J_i)}$ is an R_i -partitioning.
3. $\nu(c^* \overline{F}) \sqsubseteq R_i$.

Therefore, after Algo. 2 terminates, $(c \circ \kappa_\sigma)^* \overline{F}R = \top_{C_\sigma}$ holds for each $\sigma \in L(J)$, $\{\kappa_\sigma \rightarrow C\}_{\sigma \in L(J)}$ is an R -partitioning, and $\nu(c^* \overline{F}) \sqsubseteq R$, for $R \in \mathbb{E}_C$ defined in Line 3.

Proposition 6.7 (termination). *Algo. 2 terminates.*

The key observation for the proof of termination is that in each iteration of the main loop either the partition is refined, or it is not but the number of dirty leaves decreases.

Theorem 6.8 (correctness). *Algo. 2 terminates and returns a $\nu(c^* \overline{F})$ -partitioning.*

The explicit correspondence between fPR^{H} and §2 (table 1) allows us to directly use Hopcroft's inequality. The following result, while it does not give a complexity bound for fPR^{H} itself, plays a central role in the amortised analysis of its concrete instances in §6.2.

Proposition 6.9. *If each call of MARKDIRTY in Algo. 2 takes $\mathcal{O}(K \sum_{k=0}^{n_\rho} w(C_{\rho k}))$ time for some K , the total time taken by the repeated calls of MARKDIRTY is $\mathcal{O}(Kw(C) \log w(C))$.*

Algorithm 3 The algorithm $\text{fPR}_w^{\text{H-ER}}$, obtained as an instance of fPR^{H} (Algo. 2) where p is $\mathbf{EqRel} \rightarrow \mathbf{Set}$, with a semantically equivalent formulation of MARKDIRTY (successor-centric in fPR^{H} ; predecessor-centric here in $\text{fPR}_w^{\text{H-ER}}$).

Input: A coalgebra $c: C \rightarrow FC$ in \mathbf{Set} .

Output: A mono-sink $\{\kappa_i: C_i \rightarrow C\}_{i \in I}$ for some I .

(the same as Line 1–13 of Algo. 2)

```

14: procedure MARKDIRTY( $\rho, k_0$ )
15:   for  $k \in \{0, \dots, n_\rho\}$  do  $C_{\rho k}^{\text{cl}} := C_{\rho k}$ 
16:   for  $k \in \{0, \dots, n_\rho\} \setminus \{k_0\}$  and  $y \in C_{\rho k}$  do
17:     for  $x$ : predecessor of  $y$  do
18:       Find  $\tau \in L(J \cup \{\rho 0, \dots, \rho n_\rho\})$  such that  $x \in C_\tau$ 
19:       If such  $\tau$  exists, then  $C_\tau^{\text{cl}} := C_\tau^{\text{cl}} \setminus \{x\}$  ▷ mark  $x$  as dirty

```

6.2 Concrete Yet Functor-Generic Algorithms $\text{fPR}_{w_C}^{\text{H-ER}}$, $\text{fPR}_{w_P}^{\text{H-ER}}$, $\text{fPR}_{w_R}^{\text{H-ER}}$

We instantiate the fibrational algorithm $\text{fPR}_{(F, \bar{F}), w}^{\text{H}}$ with $\mathbf{EqRel} \rightarrow \mathbf{Set}$ as a base fibration. In this situation, the functor F is an endofunctor on \mathbf{Set} and \bar{F} is an endofunctor on \mathbf{EqRel} which is a fibred lifting of F . This instantiation also enables a semantically equivalent reformulation of MARKDIRTY—its “implementation” is now “predecessor-centric” rather than “successor-centric”—and this aids more refined complexity analysis.

For a weight function w (a parameter of fPR^{H}), we introduce three examples w_C, w_P, w_R , leading to three functor-generic algorithms $\text{fPR}_{(F, \bar{F}), w_C}^{\text{H-ER}}$, $\text{fPR}_{(F, \bar{F}), w_P}^{\text{H-ER}}$ and $\text{fPR}_{(F, \bar{F}), w_R}^{\text{H-ER}}$.

Definition 6.10 ($\text{fPR}_w^{\text{H-ER}}$). Let $\mathbf{Set} \xrightarrow{F} \mathbf{Set}$ and $\mathbf{EqRel} \xrightarrow{\bar{F}} \mathbf{EqRel}$ be functors, $C \xrightarrow{c} FC$ be a coalgebra, and $w: \mathcal{P}(C) \rightarrow \mathbb{N}$ be a function (which amounts to $w: \text{ob}(\mathbf{Sub}(\mathbb{C})_C) \rightarrow \mathbb{N}$ in §6.1), all satisfying Assum. 6.1 (C must be finite, in particular). The algorithm $\text{fPR}_w^{\text{H-ER}}$ is shown in Algo. 3; it computes a $\nu(c^* \bar{F})$ -partitioning of C .

Line 14–19 of Algo. 3 uses this categorical notion of predecessor (Line 17), which is in Def. 6.11. Its equivalence to the original definition (Line 14–19 of Algo. 2) is easy; so $\text{fPR}_w^{\text{H-ER}}$ is correct by Thm. 6.8. The successor-centric description is more convenient in the correctness proof (Thm. 6.8), while the predecessor-centric one is advantageous for complexity analysis.

Definition 6.11 (predecessor [14]). Let $c: C \rightarrow FC$ be a coalgebra in \mathbf{Set} . For $x, y \in C$, we say x is a *predecessor* of y if $x \notin B$, where B is a subset of C defined

Table 2: Examples of partition refinement algorithms induced by $\text{fPR}_{(F,\overline{F}),w}^{\text{H-ER}}$

Functor $F(X)$	Weight function	System	Algorithm: a variation of
$2 \times X^A$	$w_C(A), w_R(A)$	DFA	Hopcroft's algorithm [6, 11, 16]
$\mathcal{P}(A \times X)$	$w_C(A)$	LTS	[24]
$\mathcal{D}(X)$	$w_C(A)$	Markov chain	[25]
$\mathcal{P}(\mathcal{D}(X))$	$w_C(A)$	Markov decision process	[7]

by the following pullback:

$$\begin{array}{ccc}
 B & \xrightarrow{\quad} & C \\
 \downarrow & \lrcorner & \downarrow c \\
 F(C \setminus \{y\}) & \xrightarrow{F(i_y)} & FC
 \end{array}$$

Here i_y is the canonical injection.

For w as a parameter of $\text{fPR}_w^{\text{H-ER}}$, we introduce three functions.

Definition 6.12 (weight functions w_C, w_P, w_R). We define $w_C, w_P, w_R: \mathcal{P}(C) \rightarrow \mathbb{N}$ —called the *cardinality*, *predecessor* and *reachability* weights, respectively—as follows: $w_C(A) = |A|$, $w_P(A) = \sum_{x \in A} |\{y \in C \mid y \text{ is a predecessor of } x\}|$, and $w_R(A) = |A \cap C'|$, where $C' = \{x \in C \mid x \text{ is a successor of some } y \in C\}$.

The cardinality weight is the most commonly used one in various partition refinement algorithms, including [14]. The latter two have been used in [6, 11, 16] for DFA partition refinement; we use them for the first time in categorical partition refinement algorithms.

Our algorithm $\text{fPR}_{(F,\overline{F}),w}^{\text{H-ER}}$ induces concrete partition refinement algorithms for various systems as shown in table 2.

The effect of different weight functions is illustrated in appendix D.

The predecessor-centric MARKDIRTY and concrete choices of w allow the following fine-grained complexity analysis.

Theorem 6.13 (complexity of $\text{fPR}_{w_C}^{\text{H-ER}}, \text{fPR}_{w_P}^{\text{H-ER}}, \text{fPR}_{w_R}^{\text{H-ER}}$).

1. Let $M = \max_{y \in C} |\{x \in C \mid x \text{ is a predecessor of } y\}|$, the “in-degree” of $c: C \rightarrow FC$; suppose it takes $\mathcal{O}(f)$ time to compute $c(x) \in FC$ for each $x \in C$. The time complexity of $\text{fPR}_{w_C}^{\text{H-ER}}$ is $\mathcal{O}(fM|C| \log |C|)$.
2. Let $m = w_P(C) = \sum_{y \in C} |\{x \in C \mid x \text{ is a predecessor of } y\}|$. The time complexity of $\text{fPR}_{w_P}^{\text{H-ER}}$ is $\mathcal{O}(fm \log m)$. Since $m \leq |C|^2$, it is also bounded by $\mathcal{O}(fm \log |C|)$.
3. The time complexity of $\text{fPR}_{w_R}^{\text{H-ER}}$ is $\mathcal{O}(fM|C'| \log |C'|)$, where f and M are from above and C' is from Def. 6.12.

We note that the complexity bound given in [14] is $\mathcal{O}(fm \log |C|)$ which is the same as 2: their algorithm is essentially $\text{fPR}_{w_C}^{\text{H-ER}}$, but their more fine-grained, element-wise analysis derived the aforementioned bound.

We sketch the proof of 1 for illustration. It uses Hopcroft’s inequality (Thm. 2.9) in its amortised analysis.

Proof (sketch). (of Thm. 6.13.1) We can check that Algo. 3 (with $w = w_C$) satisfies the premise of Prop. 6.9. By implementing this algorithm properly (preparing a table for x and τ in Line 18 of Algo. 3, as done in [14]), it takes $\mathcal{O}(1)$ time to execute each iteration of the loop at Line 17. Thus, the loop at Line 16 takes $\mathcal{O}(M|C_{\rho k}|)$ time for each k . Therefore the time taken for each call of MARKDIRTY is $\mathcal{O}(M \sum_{k \in \{0, \dots, n_\rho\} \setminus \{k_0\}} |C_{\rho k}|)$. By Prop. 6.9, the time taken for the repeated calls of MARKDIRTY in total is $\mathcal{O}(M|C| \log |C|)$.

The complexity of the other parts of the algorithm is also bounded. We write C_σ^{di} for $C_\sigma \setminus C_\sigma^{\text{cl}}$. The computation of R_ρ (Line 3–5 of Algo. 2) takes $\mathcal{O}(f|C_\rho^{\text{di}}|)$, and the computation of R_ρ -partitioning (Line 9 of Algo. 2) takes $\mathcal{O}(|C_\rho^{\text{di}}|)$, using appropriate data structures (see [16]). Hence, it takes $\mathcal{O}(f|C_\rho^{\text{di}}|)$ for each iteration of the main loop except for MARKDIRTY.

Therefore, the total time for Algo. 3 except for MARKDIRTY (let us write $T_{\setminus \text{MARKDIRTY}}$) is $\mathcal{O}(\sum_{\rho \text{ for each iteration}} f|C_\rho^{\text{di}}|)$. We use amortised analysis to bound this sum. Specifically, it is easy to see that the sum $\sum_{\rho} f|C_\rho^{\text{di}}|$ is bounded by the number of times that states are marked as dirty, multiplied by f . Throughout the algorithm, the number of times that states are marked as dirty (at Line 19 of Algo. 3) is at most the time consumed by MARKDIRTY, which is $\mathcal{O}(M|C| \log |C|)$. Therefore, $T_{\setminus \text{MARKDIRTY}}$ is $\mathcal{O}(fM|C| \log |C|)$; so is the total time.

Acknowledgments. The authors are supported by ERATO HASUO Metamathematics for Systems Design Project (No. JPMJER1603), JST. T. Sanada and R. Kojima are supported by JST Grant Number JPMJFS2123. Y. Komorida is supported by JSPS KAKENHI Grant Number JP21J13334.

References

1. Berkholz, C., Bonsma, P.S., Grohe, M.: Tight lower and upper bounds for the complexity of canonical colour refinement. *Theory Comput. Syst.* **60**(4), 581–614 (2017). <https://doi.org/10.1007/s00224-016-9686-0>
2. Bonchi, F., Holland, J., Piedeleu, R., Sobocinski, P., Zanasi, F.: Diagrammatic algebra: from linear to concurrent systems. *Proc. ACM Program. Lang.* **3**(POPL), 25:1–25:28 (2019). <https://doi.org/10.1145/3290338>, <https://doi.org/10.1145/3290338>
3. de Vink, E., Rutten, J.: Bisimulation for probabilistic transition systems: a coalgebraic approach. *Theoretical Computer Science* **221**(1), 271–293 (1999). [https://doi.org/10.1016/S0304-3975\(99\)00035-3](https://doi.org/10.1016/S0304-3975(99)00035-3)
4. Deifel, H.P., Milius, S., Schröder, L., Wißmann, T.: Generic partition refinement and weighted tree automata. In: ter Beek, M.H., McIver, A., Oliveira, J.N. (eds.) *Formal Methods – The Next 30 Years*. pp. 280–297. Springer International Publishing, Cham (2019)
5. Dorsch, U., Milius, S., Schröder, L., Wißmann, T.: Efficient coalgebraic partition refinement. In: Meyer, R., Nestmann, U. (eds.) *28th International Conference on*

- Concurrency Theory (CONCUR 2017). Leibniz International Proceedings in Informatics (LIPIcs), vol. 85, pp. 32:1–32:16. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany (2017). <https://doi.org/10.4230/LIPIcs.CONCUR.2017.32>
6. Gries, D.: Describing an algorithm by Hopcroft. *Acta Informatica* **2**(2), 97–109 (1973)
 7. Groote, J.F., Rivera Verduzco, J., De Vink, E.P.: An efficient algorithm to determine probabilistic bisimulation. *Algorithms* **11**(9) (2018). <https://doi.org/10.3390/a11090131>
 8. Hermida, C., Jacobs, B.: Structural induction and coinduction in a fibrational setting. *Information and Computation* **145**(2), 107–152 (1998). <https://doi.org/10.1006/inco.1998.2725>
 9. Höfberg, J., Maletti, A., May, J.: Backward and forward bisimulation minimisation of tree automata. In: Holub, J., Žďárek, J. (eds.) *Implementation and Application of Automata*. pp. 109–121. Springer Berlin Heidelberg, Berlin, Heidelberg (2007)
 10. Höfberg, J., Maletti, A., May, J.: Bisimulation minimisation for weighted tree automata. In: Harju, T., Karhumäki, J., Lepistö, A. (eds.) *Developments in Language Theory*. pp. 229–241. Springer Berlin Heidelberg, Berlin, Heidelberg (2007)
 11. Hopcroft, J.E.: An $n \log n$ algorithm for minimizing states in a finite automaton. In: *Theory of Machines and Computations*. pp. 189–196. Academic Press (1971)
 12. Jacobs, B.: *Categorical Logic and Type Theory*. Studies in logic and the foundations of mathematics, Elsevier Science (1999)
 13. Jacobs, B.: *Introduction to Coalgebra: Towards Mathematics of States and Observation*. Cambridge Tracts in Theoretical Computer Science, Cambridge University Press (2016). <https://doi.org/10.1017/CB09781316823187>
 14. Jacobs, J., Wißmann, T.: Fast coalgebraic bisimilarity minimization. *Proc. ACM Program. Lang.* **7**(POPL) (Jan 2023). <https://doi.org/10.1145/3571245>
 15. Kanellakis, P.C., Smolka, S.A.: CCS expressions, finite state processes, and three problems of equivalence. *Information and Computation* **86**(1), 43–68 (1990). [https://doi.org/10.1016/0890-5401\(90\)90025-D](https://doi.org/10.1016/0890-5401(90)90025-D)
 16. Knuutila, T.: Re-describing an algorithm by Hopcroft. *Theoretical Computer Science* **250**(1), 333–363 (2001). [https://doi.org/10.1016/S0304-3975\(99\)00150-4](https://doi.org/10.1016/S0304-3975(99)00150-4)
 17. Komorida, Y., Katsumata, S., Hu, N., Klin, B., Humeau, S., Eberhart, C., Hasuo, I.: Codensity games for bisimilarity. *New Gener. Comput.* **40**(2), 403–465 (2022). <https://doi.org/10.1007/s00354-022-00186-y>
 18. Lombardy, S., Sakarovitch, J.: Morphisms and minimisation of weighted automata. *Fundam. Informaticae* **186**(1-4), 195–218 (2022). <https://doi.org/10.3233/FI-222126>
 19. Milner, R.: *Communication and Concurrency*. Prentice-Hall, Inc., USA (1989)
 20. Moore, E.F.: *Gedanken-Experiments on Sequential Machines*, pp. 129–154. Princeton University Press, Princeton (1956). <https://doi.org/10.1515/9781400882618-006>
 21. Park, D.: Concurrency and automata on infinite sequences. In: Deussen, P. (ed.) *Theoretical Computer Science*. pp. 167–183. Springer Berlin Heidelberg, Berlin, Heidelberg (1981)
 22. Piedeleu, R., Kartsaklis, D., Coecke, B., Sadrzadeh, M.: Open system categorical quantum semantics in natural language processing. In: Moss, L.S., Sobocinski, P. (eds.) *6th Conference on Algebra and Coalgebra in Computer Science, CALCO 2015*, June 24–26, 2015, Nijmegen, The Netherlands. LIPIcs, vol. 35, pp. 270–289. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2015). <https://doi.org/10.4230/LIPIcs.CALCO.2015.270-289>

- 4230/LIPIcs.CALCO.2015.270, <https://doi.org/10.4230/LIPIcs.CALCO.2015.270>
23. Rutten, J.: Universal coalgebra: a theory of systems. *Theoretical Computer Science* **249**(1), 3–80 (2000). [https://doi.org/10.1016/S0304-3975\(00\)00056-6](https://doi.org/10.1016/S0304-3975(00)00056-6), *modern Algebra*
 24. Valmari, A.: Bisimilarity minimization in $O(m \log n)$ time. In: Franceschinis, G., Wolf, K. (eds.) *Applications and Theory of Petri Nets*. pp. 123–142. Springer Berlin Heidelberg, Berlin, Heidelberg (2009)
 25. Valmari, A., Franceschinis, G.: Simple $o(m \log n)$ time markov chain lumping. In: Esparza, J., Majumdar, R. (eds.) *Tools and Algorithms for the Construction and Analysis of Systems*. pp. 38–52. Springer Berlin Heidelberg, Berlin, Heidelberg (2010)
 26. Wißmann, T., Deifel, H., Milius, S., Schröder, L.: From generic partition refinement to weighted tree automata minimization. *Formal Aspects Comput.* **33**(4-5), 695–727 (2021). <https://doi.org/10.1007/s00165-020-00526-z>

A Proofs for Section 2 (Hopcroft's Inequality)

Lemma 2.4. *Let T be a finite tree with a root r , w be a weight function of T , and S be an arbitrary set of edges of T . Then $\sum_{v \in V(T)} \sum_{\substack{u \in \text{ch}(v) \\ (v,u) \notin S}} w(u) \geq \sum_{l \in L(T)} |\text{path}(r, l) \setminus S| \cdot w(l)$ holds. The equality holds when w is tight.*

Proof. Let r be the root of T . We prove by the induction on the size of T .

(**Base case**). T is a tree whose vertex is only r . Thus we have

$$\sum_{v \in V(T)} \sum_{\substack{u \in \text{ch}(v) \\ (v,u) \notin S}} w(u) = \sum_{\substack{u \in \emptyset = \text{ch}(r) \\ (v,u) \notin S}} w(u) = 0 = |\emptyset| \cdot w(r) = \sum_{l \in L(T)} |\text{path}(r, l) \setminus S| \cdot w(l).$$

(**Induction step**). Let $U = \text{ch}(r) \cap S$ and $W = \text{ch}(r) \setminus S$. The following calculation shows the inequality:

$$\begin{aligned} & \sum_{v \in V(T)} \sum_{\substack{u \in \text{ch}(v) \\ (v,u) \notin S}} w(u) \\ &= \sum_{r' \in W} w(r') + \sum_{r' \in U \cup W} \sum_{v \in V(\text{tr}(r'))} \sum_{\substack{u \in \text{ch}(v) \\ (v,u) \notin S}} w(u) \\ &\geq \sum_{r' \in W} w(r') + \sum_{r' \in U \cup W} \sum_{l \in L(\text{tr}(r'))} |\text{path}(r', l) \setminus S| \cdot w(l) \\ & \hspace{15em} \text{induction hypothesis} \\ &= \sum_{r' \in U} \sum_{l \in L(\text{tr}(r'))} |\text{path}(r', l) \setminus S| \cdot w(l) + \sum_{r' \in W} \left(w(r') + \sum_{l \in L(\text{tr}(r'))} |\text{path}(r', l) \setminus S| \cdot w(l) \right) \\ &\geq \sum_{r' \in U} \sum_{l \in L(\text{tr}(r'))} |\text{path}(r', l) \setminus S| \cdot w(l) + \sum_{r' \in W} \left(\sum_{l \in L(\text{tr}(r'))} w(l) + \sum_{l \in L(\text{tr}(r'))} |\text{path}(r', l) \setminus S| \cdot w(l) \right) \\ & \hspace{15em} \text{by } w(r') \geq \sum_{l \in L(\text{tr}(r'))} w(l) \\ &= \sum_{r' \in U} \sum_{l \in L(\text{tr}(r'))} |\text{path}(r', l) \setminus S| \cdot w(l) + \sum_{r' \in W} \sum_{l \in L(\text{tr}(r'))} (w(l) + |\text{path}(r', l) \setminus S| \cdot w(l)) \\ &= \sum_{r' \in U} \sum_{l \in L(\text{tr}(r'))} |\text{path}(r, l) \setminus S| \cdot w(l) + \sum_{r' \in W} \sum_{l \in L(\text{tr}(r'))} |\text{path}(r, l) \setminus S| \cdot w(l) \\ &= \sum_{l \in L(T)} |\text{path}(r, l) \setminus S| \cdot w(l). \end{aligned}$$

When w is tight, we have $w(r') = \sum_{l \in L(\text{tr}(r'))} w(l)$ for every $r' \in W$, and the equality holds.

Lemma 2.8. *Let T be a finite tree with a root r , w be a weight of T , and h be an hcc for w . For each vertex $v \in V(T)$ with $w(v) \neq 0$, the following inequality holds: $|\text{lpath}(r, v)| \leq \log_2 w(r) - \log_2 w(v)$.*

Proof. Let $\langle (v_1, u_1), \dots, (v_m, u_m) \rangle$ be the sequence of the light edges of $\text{lpath}(r, v)$ in the order from r to v . Since $u_i \neq h_{v_i}$, we have $2 \cdot w(u_i) \leq w(v_i)$. We also have $w(v_{i+1}) \leq w(u_i)$. Hence, the following inequality holds:

$$\begin{aligned} w(r) &\geq w(v_1) \geq 2 \cdot w(u_1) \\ &\geq 2 \cdot w(v_2) \geq 2^2 \cdot w(u_2) \\ &\dots \\ &\geq 2^{m-1} \cdot w(v_m) \geq 2^m \cdot w(u_m) \\ &\geq 2^m \cdot w(v). \end{aligned}$$

Taking the logarithm of both sides yields $|\text{lpath}(r, v)| = m \leq \log_2 w(r) - \log_2 w(v)$.

Theorem 2.9 (Hopcroft's inequality). *Let T be a finite tree with root r , w be a weight function of T , and h be a heavy child choice for w . The following inequality holds.*

$$\sum_{v \in V(T)} \sum_{u \in \text{lch}_h(v)} w(u) \leq w(r) \log_2 w(r) - \sum_{\substack{l \in L(T) \\ w(l) \neq 0}} w(l) \log_2 w(l) \quad (2)$$

Proof. Let w' be the tightening of w . By Lem. 2.7, w' is a tight weight function, and h is also a heavy child choice for w' . If we have

$$\sum_{v \in V(T)} \sum_{u \in \text{lch}_h(v)} w'(u) \leq w'(r) \log_2 w'(r) - \sum_{l \in L(T), w'(l) \neq 0} w'(l) \log_2 w'(l) \quad (4)$$

then the desired inequality holds:

$$\begin{aligned} \sum_{v \in V(T)} \sum_{u \in \text{lch}_h(v)} w(u) &\leq \sum_{v \in V(T)} \sum_{u \in \text{lch}_h(v)} w'(u) && \text{Lem. 2.7} \\ &\leq w'(r) \log_2 w'(r) - \sum_{\substack{l \in L(T) \\ w'(l) \neq 0}} w'(l) \log_2 w'(l) \\ &\leq w(r) \log_2 w(r) - \sum_{\substack{l \in L(T) \\ w(l) \neq 0}} w(l) \log_2 w(l). && \text{Lem. 2.7} \end{aligned}$$

Now, our goal is to show (4). It is proven by the following calculation:

$$\sum_{v \in V(T)} \sum_{u \in \text{lch}_h(v)} w'(u)$$

$$\begin{aligned}
&= \sum_{l \in L(T)} \sum_{(v,u) \in \text{lpath}(r,l)} w'(l) && w': \text{tight, and Lem. 2.5} \\
&= \sum_{\substack{l \in L(T) \\ w'(l) \neq 0}} \sum_{(v,u) \in \text{lpath}(r,l)} w'(l) \\
&\leq \sum_{\substack{l \in L(T) \\ w'(l) \neq 0}} (\log_2 w'(r) - \log_2 w'(l)) \cdot w'(l) && \text{Lem. 2.8} \\
&= \left(\sum_{\substack{l \in L(T) \\ w'(l) \neq 0}} w'(l) \right) \log_2 w'(r) - \sum_{\substack{l \in L(T) \\ w'(l) \neq 0}} w'(l) \log_2 w'(l) \\
&= w'(r) \log_2 w'(r) - \sum_{\substack{l \in L(T) \\ w'(l) \neq 0}} w'(l) \log_2 w'(l).
\end{aligned}$$

B Proofs for Section 5 (The Naive Fibrational Algorithm $\text{fPR}^{\text{naive}}$)

Lemma 5.3 (loop invariant). *At the beginning of each iteration of the main loop, the following hold.*

1. The mono-sink $\{\kappa_\sigma: C_\sigma \twoheadrightarrow C\}_{\sigma \in L(T)}$ is an R -partitioning.
2. $\nu(c^* \overline{F}) \sqsubseteq R$.

Proof. Firstly we prove the following two lemmas.

Lemma B.1. *Assume a \mathbf{CLat}_\square -fibration $p: \mathbb{E} \rightarrow \mathbb{C}$ satisfies Assum. 5.1. Now let $\{\kappa_i: A_i \twoheadrightarrow C\}_{i \in I}$ be a mono-sink in \mathbb{C} that is pairwise disjoint, $\lambda: B \twoheadrightarrow C$ be a monomorphism with $A_i \cap B \cong 0$ for each $i \in I$, $P_i \in \mathbb{E}_{A_i}$, and $R \in \mathbb{E}_B$. Then we have $\lambda^* \left(\left(\bigsqcup_{i \in I} (\kappa_i)_*(P_i) \right) \sqcup \lambda_*(R) \right) = R$.*

The above claim can be understood as follows: when a monomorphism $B \twoheadrightarrow C$ is added to a mono-sink $\{\kappa_i: A_i \twoheadrightarrow C\}_{i \in I}$, if B is disjoint from each A_i ($A_i \cap B \cong 0$), then the objects P_i above A_i do not interfere with R above B .

$$\begin{array}{ccc}
\begin{array}{c} \cdots \\ \cdot \\ \cdot \end{array} P_i & \xrightarrow{(\kappa_i)_*} & (\kappa_i)_* P_i \mathbb{E} \\
\cdot & \nearrow R & \xrightarrow{\lambda_*} \lambda_* R \\
\cdot & & \downarrow p \\
\begin{array}{c} \cdots \\ \cdot \\ \cdot \end{array} A_i & \xrightarrow{\kappa_i} & C \\
\cdot & \nearrow B & \xrightarrow{\lambda} C
\end{array}$$

Cond. 3 can be equationally expressed by $\lambda^*(\kappa_*(Q) \sqcup \lambda_*(R)) = (\lambda^* \kappa_*(Q)) \sqcup (\lambda^* \lambda_*(R))$; it can be thought of as a fibrational analogue of *modularity* of a lattice.


Proof. Let $A = \bigcup_{i \in I} A_i$, $\kappa: A \twoheadrightarrow C$, and $\iota_i: A_i \twoheadrightarrow A$ for each $i \in I$. Notice that the following diagram is a pullback since $\mathbf{Sub}(\mathbb{C})_C$ is distributive (Cond. 1):

$$\begin{array}{ccc}
 0 = B \cap \left(\bigcup_{i \in I} A_i\right) & \twoheadrightarrow & B \\
 \downarrow & \lrcorner & \downarrow \lambda \\
 \bigcup_{i \in I} A_i & \xrightarrow{\kappa} & C.
 \end{array} \tag{5}$$

We have

$$\begin{aligned}
 & \lambda^* \left(\left(\bigsqcup_{i \in I} (\kappa_i)_*(P_i) \right) \sqcup \lambda_*(R) \right) \\
 &= \lambda^* \left(\left(\bigsqcup_{i \in I} (\kappa \circ \iota_i)_*(P_i) \right) \sqcup \lambda_*(R) \right) \\
 &= \lambda^* \left(\kappa_* \left(\bigsqcup_{i \in I} (\iota_i)_*(P_i) \right) \sqcup \lambda_*(R) \right) && \kappa_*: \text{left adjoint} \\
 &= \left(\lambda^* \kappa_* \left(\bigsqcup_{i \in I} (\iota_i)_*(P_i) \right) \right) \sqcup (\lambda^* \lambda_*(R)) && \text{Cond. 3} \\
 &= \perp \sqcup R && \text{eq. (5) and Cond. 2} \\
 &= R.
 \end{aligned}$$

The following lemma intuitively says that a partitioning can be refined by another partitioning of one of its leaves, in the manner shown as follows.



Lemma B.2. Let $p: \mathbb{E} \rightarrow \mathbb{C}$ be a \mathbf{CLat}_{\sqcap} -fibration that satisfies the three conditions in Lem. B.1, $C \in \mathbb{C}$, $R \in \mathbb{E}_C$, $\{\kappa_j: C_j \twoheadrightarrow C\}_{j \in \{0, \dots, m\}}$ be an R -partitioning of C , R_0 be an object of \mathbb{E}_{C_0} , and $\{\lambda_k: D_k \twoheadrightarrow C_0\}_{k \in \{0, \dots, n\}}$ be an R_0 -partitioning of C_0 . The family of monomorphisms

$$\Pi = \{\kappa_0 \circ \lambda_k: D_k \twoheadrightarrow C\}_{k \in \{0, \dots, n\}} \cup \{\kappa_j: C_j \twoheadrightarrow C\}_{j \in \{1, \dots, m\}}$$

is a Q -partitioning of C where

$$Q = \left(\bigsqcup_{k=0}^n (\kappa_0 \circ \lambda_k)_*(\top_{D_k}) \right) \sqcup \left(\bigsqcup_{j=1}^m (\kappa_j)_*(\top_{C_j}) \right) = (\kappa_0)_*(R_0) \sqcup \left(\bigsqcup_{j=1}^m (\kappa_j)_*(\top_{C_j}) \right).$$

Proof. We check that Π satisfies the conditions 1–3 of Def. 4.1.

Π satisfies Def. 4.1.2 by the definition of Q .

For each $k \in \{0, \dots, n\}$ and $j \in \{1, \dots, m\}$, we have $D_k \cap C_j = 0$. Since $\{\kappa_j: C_j \twoheadrightarrow C\}_{j \in \{0, \dots, m\}}$ and $\{\lambda_k: D_k \twoheadrightarrow C_0\}_{k \in \{0, \dots, n\}}$ are partitioning, we have $C_j \cap C_{j'} = 0$ for $j, j' \in \{1, \dots, m\}$ with $j \neq j'$ and $D_k \cap D_{k'} = 0$ for $k, k' \in \{1, \dots, n\}$ with $k \neq k'$. Moreover, we have $C_j \not\cong 0$ and $D_k \not\cong 0$ for each j and k . Thus, Π satisfies Def. 4.1.3.

We have

$$\begin{aligned} \kappa_j^*(Q) &= \kappa_j^* \left((\kappa_0)_*(R_0) \sqcup \left(\bigsqcup_{j'=1}^m (\kappa_{j'})_*(\top_{C_{j'}}) \right) \right) \\ &= \top_{C_j} \end{aligned} \quad \text{Lem. B.1}$$

for $j \in \{1, \dots, m\}$, and

$$\begin{aligned} &(\kappa_0 \circ \lambda_k)^*(Q) \\ &= (\kappa_0 \circ \lambda_k)^* \left(\left(\bigsqcup_{k'=0}^n (\kappa_0 \circ \lambda_{k'})_*(\top_{D_{k'}}) \right) \sqcup \left(\bigsqcup_{j=1}^m (\kappa_j)_*(\top_{C_j}) \right) \right) \\ &= \top_{D_k} \end{aligned} \quad \text{Lem. B.1}$$

for $k \in \{0, \dots, n\}$. Hence, Π satisfies Def. 4.1.1. Therefore, Π is a Q -partitioning of C .

We go back to the proof of Lem. 5.3. We write R_i and T_i for R and T , respectively, at the beginning (line 2) of the i -th iteration of the main loop. We prove 1 and 2 by the induction on i .

(Base case). We have $R_0 = \top_C$ and $T_0 = \{\epsilon\}$. The mono-sink is $\{\kappa_\epsilon: C_\epsilon \twoheadrightarrow C\}$, and this is an R_0 -partitioning of C . We also have $\nu(c^*\overline{F}) \sqsubseteq \top_C = R_0$.

(Induction step). Assume that $\{\kappa_\sigma: C_\sigma \twoheadrightarrow\}_{\sigma \in L(T_i)}$ is an R_i -partitioning and $\nu(c^*\overline{F}) \sqsubseteq R_i$. By $\nu(c^*\overline{F}) \sqsubseteq R_i$, we have $\nu(c^*\overline{F}) = c^*\overline{F}(\nu(c^*\overline{F})) \sqsubseteq c^*\overline{F}(R_i) = R_{i+1}$. Hence, 2 holds for $i+1$. By Lem. B.2 and the induction hypothesis, we have $\{\kappa_\sigma: C_\sigma \twoheadrightarrow C\}_{\sigma \in L(T_{i+1})}$ is an R_{i+1} -partitioning. Therefore, 1 holds for $i+1$.

Lemma 5.4 (termination). *If \mathbb{E}_C is a well-founded lattice, Algo. 1 terminates.*

Proof. We write R_i for R at the beginning (line 2) of the i -th iteration of the main loop, and have

$$R_0 = \top_C, \quad R_1 = c^*\overline{F}R_0, \quad \dots, \quad R_{i+1} = c^*\overline{F}R_i, \quad \dots$$

We have $R_1 \sqsubseteq \top_C = R_0$. It is proven by the induction on i and the functoriality of $c^*\overline{F}$ that $R_i = c^*\overline{F}R_{i-1} \sqsubseteq R_{i-1}$ holds for each i such that R_i is defined. By this observation and the condition in line 2, we have the strictly descending sequence in \mathbb{E}_C : $R_0 \supset R_1 \supset \dots \supset R_n \supset \dots$. Since \mathbb{E}_C is well-founded, the length of the sequence is finite, that is Algo. 1 terminates.

Proposition 5.5 (correctness). *If \mathbb{E}_C is well-founded, then Algo. 1 terminates and returns $\nu(c^*\overline{F})$ -partitioning $\{\kappa: C_i \rightarrow C\}_{i \in I}$.*

Proof. By Lem. 5.4, Algo. 1 terminates. From the termination condition (line 2) of the main loop, $R = c^*\overline{F}R$ holds when the algorithm terminates. Hence, R is a fixed point. By Lem. 5.3.2, R is greater than or equal to the greatest fixed point $\nu(c^*\overline{F})$ of the functor $c^*\overline{F}$. Thus, we have $R = \nu(c^*\overline{F})$.

C Proofs for Section 6 (Optimised Algorithms with Hopcroft's Inequality)

Proposition 6.3. *The fibration $p: \mathbf{EqRel} \rightarrow \mathbf{Set}$ with the relation lifting $\text{Rel}(F)$ of a functor F defined by (3) and a coalgebra $c: C \rightarrow FC$ for a finite set C satisfies the assumptions 1–9 of Assum. 6.1.*

Proof. We only prove that $p: \mathbf{EqRel} \rightarrow \mathbf{Set}$ satisfies the premise 3 of Lem. B.1. The other conditions are easy to check.

Given injections $\kappa: A \rightarrow C$ and $\lambda: B \rightarrow C$ in \mathbf{Set} and equivalence relations R on A and S on B . We want to show $(\lambda^*\kappa_*R) \sqcup (\lambda^*\kappa_*S) = \lambda^*((\kappa_*R) \sqcup (\lambda_*S))$. For any $(x, y) \in \lambda^*((\kappa_*R) \sqcup (\lambda_*S))$, we have $(\lambda(x), \lambda(y)) \in (\kappa_*R) \sqcup (\lambda_*S)$. Hence, there exist $m \in \mathbb{N}$ and $z_0, \dots, z_m \in C$ such that $z_0 = \lambda(x)$, $z_m = \lambda(y)$, and $(z_i, z_{i+1}) \in \kappa_*R$ or $(z_i, z_{i+1}) \in \lambda_*S$ for each $i = 0, \dots, m-1$. We can assume that $z_i \neq z_{i+1}$ and $z_i \in \lambda(B)$ for each i without loss of generality since $\lambda(x), \lambda(y) \in \lambda(B)$. Thus, the sequence $x = \lambda^{-1}(z_0), \lambda^{-1}(z_1), \dots, \lambda^{-1}(z_m) = y$ in B satisfies $(\lambda^{-1}(z_i), \lambda^{-1}(z_{i+1})) \in \lambda^*\kappa_*R$ or $(\lambda^{-1}(z_i), \lambda^{-1}(z_{i+1})) \in \lambda^*\lambda_*S$ for each i . This means $(x, y) \in (\lambda^*\kappa_*R) \sqcup (\lambda^*\lambda_*S)$. Hence, $\lambda^*((\kappa_*R) \sqcup (\lambda_*S)) \sqsubseteq (\lambda^*\kappa_*R) \sqcup (\lambda^*\lambda_*S)$ holds. Conversely, we can show $\lambda^*((\kappa_*R) \sqcup (\lambda_*S)) \supseteq (\lambda^*\kappa_*R) \sqcup (\lambda^*\lambda_*S)$ by the similar argument.

Proposition 6.6 (loop invariants). *At the beginning of the i -th iteration, the following hold.*

1. $(c \circ \kappa_\sigma \circ \kappa_\sigma^{\text{cl}, i})^*\overline{F}(R_i) = \top_{C_\sigma^{\text{cl}, i}}$ for each leaf $\sigma \in L(J_i)$.
2. The mono-sink $\{\kappa_\sigma: C_\sigma \rightarrow C\}_{\sigma \in L(J_i)}$ is an R_i -partitioning.
3. $\nu(c^*\overline{F}) \sqsubseteq R_i$.

Therefore, after Algo. 2 terminates, $(c \circ \kappa_\sigma)^\overline{F}R = \top_{C_\sigma}$ holds for each $\sigma \in L(J)$, $\{\kappa_\sigma \rightarrow C\}_{\sigma \in L(J)}$ is an R -partitioning, and $\nu(c^*\overline{F}) \sqsubseteq R$, for $R \in \mathbb{E}_C$ defined in Line 3.*

Proof. We first prove lemmas to prove the loop invariants. The next technical lemma follows easily from the fibredness of \overline{F} .

Lemma C.1. *Let $m: A \rightarrow C$ be a morphism in \mathbb{C} and $P_1, P_2 \in \mathbb{E}_C$. If $m^*P_1 = m^*P_2$, then, for any $C' \in \mathbb{C}$, $\kappa: C' \rightarrow C$ and $\lambda: C' \rightarrow FA$ such that $c \circ \kappa = Fm \circ \lambda$, we have $\kappa^*(c^*\overline{F}P_1) = \kappa^*(c^*\overline{F}P_2)$.*

Proof. For $j = 1, 2$, we have

$$\begin{aligned} \kappa^*(c^*\overline{F}P_j) &= \lambda^*((Fm)^*\overline{F}P_j) & c \circ \kappa &= Fm \circ \lambda \\ &= \lambda^*(\overline{F}(m^*P_j)). & \text{Assum. 6.1.3} & \end{aligned}$$

Hence, by the assumption $m^*P_1 = m^*P_2$, we obtain $\kappa^*(c^*\overline{F}P_1) = \lambda^*(\overline{F}(m^*P_1)) = \lambda^*(\overline{F}(m^*P_2)) = \kappa^*(c^*\overline{F}P_2)$.

The next lemma (Lem. C.2) identifies the subobject of C that is “unaffected” by the refinement from R_i to R_{i+1} , that is, by the tree expansion at the i -th iteration. It is not hard to see that the untouched leaves in $\bigcup_{\sigma \in L(T) \setminus \{\rho\}} C_\sigma$ belong to that “unaffected” part. A crucial observation—central to the Hopcroft-type optimisation—is that *at most one* new child $C_{\rho k_0}$ can also be added, where we pick the heavy one for better complexity.

Lemma C.2. *Let ρ be the leaf $\rho \in L(T_i) \setminus L(T_{i+1})$ that was chosen in **Partitioning** of the i -th iteration, and $k_0 \in \{0, \dots, n_\rho\}$ be the index that was chosen in **Relabelling** of the i -th iteration. We have $m^*(R_i) = m^*(R_{i+1})$ for the morphism $m: C_{\rho k_0} \cup \left(\bigcup_{\sigma \in L(T_i) \setminus \{\rho\}} C_\sigma\right) \rightarrow C$.*

Proof. Let $P = \bigsqcup_{\sigma \in L(T_i) \setminus \{\rho\}} (\kappa_\sigma)_*(\top_{C_\sigma})$, and $R_\rho = \bigsqcup_{k \in \{0, \dots, n_\rho\}} (\kappa'_{\rho k})_*(\top_{C_{\rho k}})$. By Lem. B.1.1, the following diagram is pullback:

$$\begin{array}{ccc} C_{\rho k_0} & \xrightarrow{\kappa'_{k_0}} & C_\rho \\ \downarrow \iota' & \lrcorner & \downarrow \kappa_\rho \\ C_{\rho k_0} \cup \left(\bigcup_{\sigma \in L(T_i) \setminus \{\rho\}} C_\sigma\right) & \xrightarrow{m} & C. \end{array}$$

Thus, the following diagram commutes by Lem. B.1.2.

$$\begin{array}{ccc} \mathbb{E}_{C_{\rho k_0}} & \xleftarrow{(\kappa'_{k_0})^*} & \mathbb{E}_{C_\rho} \\ \downarrow \iota'_* & & \downarrow (\kappa_\rho)_* \\ \mathbb{E}_{C_{\rho k_0} \cup \left(\bigcup_{\sigma \in L(T_i) \setminus \{\rho\}} C_\sigma\right)} & \xleftarrow{m^*} & \mathbb{E}_C \end{array} \quad (6)$$

By chasing the diagram, we have

$$\begin{aligned} m^*((\kappa_\rho)_*(\top_{C_\rho})) &= \iota'_*((\kappa'_{k_0})^*(\top_{C_\rho})) & \text{eq. (6)} \\ &= \iota'_*(\top_{C_{\rho k_0}}), & (\kappa'_{k_0})^*: \text{right adjoint and } \top: \text{limit} \end{aligned} \quad (7)$$

and

$$\begin{aligned} m^*((\kappa_\rho)_*(R_\rho)) &= \iota'_*((\kappa'_{k_0})^*(R_\rho)) & \text{eq. (6)} \\ &= \iota'_*(\top_{C_{\rho k_0}}). & \{\kappa'_{\rho k}: C_{\rho k} \rightarrow C_\rho\}_k \text{ is } R_\rho\text{-partitioning} \end{aligned} \quad (8)$$

There are monomorphisms $\iota_\sigma : C_\sigma \rightarrow C_{\rho k_0} \sqcup \bigcup_{\sigma \in L(T_i) \setminus \{\rho\}} C_\sigma$ for each $\sigma \in L(T_i) \setminus \{\rho\}$ with $\kappa_\sigma = m \circ \iota_\sigma$. We have

$$\begin{aligned}
 & m^*(R_i) \\
 &= m^* \left(\bigsqcup_{\sigma \in L(T_i)} (\kappa_\sigma)_*(\top_{C_\sigma}) \right) && \text{Definition of } R_i \\
 &= m^* \left((\kappa_\rho)_*(\top_{C_\rho}) \sqcup \bigsqcup_{\sigma \in L(T_i) \setminus \{\rho\}} (\kappa_\sigma)_*(\top_{C_\sigma}) \right) \\
 &= m^* \left((\kappa_\rho)_*(\top_{C_\rho}) \sqcup \bigsqcup_{\sigma \in L(T_i) \setminus \{\rho\}} (m \circ \iota_\sigma)_*(\top_{C_\sigma}) \right) && \kappa_\sigma = m \circ \iota_\sigma \\
 &= m^* \left((\kappa_\rho)_*(\top_{C_\rho}) \sqcup m_* \left(\bigsqcup_{\sigma \in L(T_i) \setminus \{\rho\}} (\iota_\sigma)_*(\top_{C_\sigma}) \right) \right) && m_*: \text{ left adjoint} \\
 &= m^*(\kappa_\rho)_*(\top_{C_\rho}) \sqcup m^* m_* \left(\bigsqcup_{\sigma \in L(T_i) \setminus \{\rho\}} (\iota_\sigma)_*(\top_{C_\sigma}) \right) && \text{Lem. B.1.3} \\
 &= \iota'_*(\top_{C_{\rho k_0}}) \sqcup m^* m_* \left(\bigsqcup_{\sigma \in L(T_i) \setminus \{\rho\}} (\iota_\sigma)_*(\top_{C_\sigma}) \right) && (7)
 \end{aligned}$$

and

$$\begin{aligned}
 & m^*(R_{i+1}) \\
 &= m^* \left(\left(\bigsqcup_{k=0}^{n_\rho} (\kappa_\rho \circ \kappa'_{\rho k})_*(\top_{C_{\rho k}}) \right) \sqcup \left(\bigsqcup_{\sigma \in L(T_i) \setminus \{\rho\}} (\kappa_\sigma)_*(\top_{C_\sigma}) \right) \right) && \text{Definition of } R_{i+1} \\
 &= m^* \left((\kappa_\rho)_*(R_\rho) \sqcup \left(\bigsqcup_{\sigma \in L(T_i) \setminus \{\rho\}} (\kappa_\sigma)_*(\top_{C_\sigma}) \right) \right) && (\kappa_\rho)_*: \text{ left adjoint} \\
 &= m^* \left((\kappa_\rho)_*(R_\rho) \sqcup \left(\bigsqcup_{\sigma \in L(T_i) \setminus \{\rho\}} (m \circ \iota_\sigma)_*(\top_{C_\sigma}) \right) \right) && \kappa_\sigma = m \circ \iota_\sigma \\
 &= m^* \left((\kappa_\rho)_*(R_\rho) \sqcup m_* \left(\bigsqcup_{\sigma \in L(T_i)} (\iota_\sigma)_*(\top_{C_\sigma}) \right) \right) && m_*: \text{ left adjoint} \\
 &= m^*(\kappa_\rho)_*(R_\rho) \sqcup m^* m_* \left(\bigsqcup_{\sigma \in L(T_i)} (\iota_\sigma)_*(\top_{C_\sigma}) \right) && \text{Lem. B.1.3}
 \end{aligned}$$

$$= \iota'_*(\top_{C_{\rho k_0}}) \sqcup m^* m_* \left(\bigsqcup_{\sigma \in L(T_i) \setminus \{\rho\}} (\iota_\sigma)_*(\top_{C_\sigma}) \right). \quad (8)$$

Therefore, we obtain $m^*(R_i) = m^*(R_{i+1})$.

We illustrate the proof of Lem. C.2. See Fig. 6. A set C is the disjoint union $C_0 \cup C_1 \cup C_2$, where $C_0 = C_{00} \cup C_{01} \cup C_{02}$. An equivalence relation R on C corresponds to a partitioning $\{C_0, C_1, C_2\}$, and an equivalence relation R' on C corresponds to a partitioning $\{C_{00}, C_{01}, C_{02}, C_1, C_2\}$. As illustrated in Fig. 6, when restricted to $C_{00} \cup C_1 \cup C_2$, R and R' coincide. This restriction yields the same equivalence relation, and Lem. C.2 formalises this coincidence. An important observation is that, we cannot *add* one extra to the restriction. For example, if we restrict R and R' to $C_{00} \cup C_{01} \cup C_1 \cup C_2$ (with both C_{00} and C_{01} included), we have $m^*(R) \neq m^*(R')$, as depicted in Fig. 7.




Fig. 6: An example situation of Lem. C.2 in $\mathbf{EqRel} \rightarrow \mathbf{Set}$: $m^*(R) = m^*(R')$.

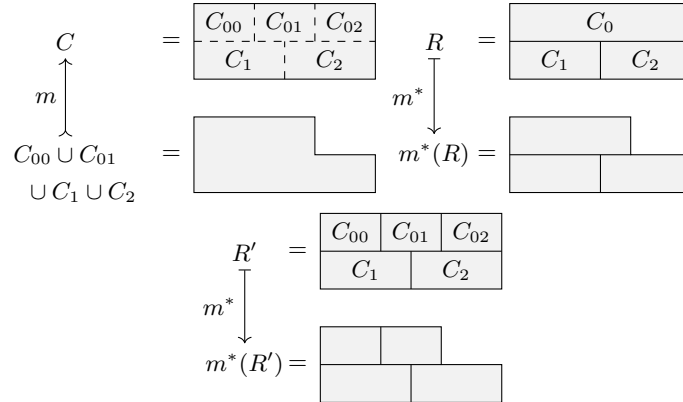


Fig. 7: We cannot add more than one child to the domain of m in Lem. C.2.

We go back to the proof of Prop. 6.6. The claim 2 follows from Lem. B.2 and the induction on i . The claim 3 follows from the assumption that p admits

partitioning (Assum. 6.1.5) and the induction on i . We prove 1 by the induction on i .

(**Base case**). We have $C_\epsilon^{\text{cl},0} = 0$. Hence, from Assum. 6.1.2, we have

$$(\kappa_\epsilon^{\text{cl},0})^* \kappa_\epsilon^* c^* \overline{F}(R_0) = \top_{C_\epsilon^{\text{cl},0}}.$$

(**Induction step**). We assume that 1 holds for i , and need to show 1 for $i + 1$. For each $\sigma \in L(T_{i+1})$, $C_\sigma^{\text{cl},i+1}$ is defined by the following pullback:

$$\begin{array}{ccc} C_\sigma^{\text{cl},i+1} & \xrightarrow{\kappa} & C_\sigma^{\text{cl},i} \\ \downarrow & \lrcorner & \downarrow \kappa_\sigma^{\text{cl},i} \\ & & C_\sigma \\ & & \downarrow \kappa_\sigma \\ B & \xrightarrow{l} & C \\ \downarrow h & \lrcorner & \downarrow c \\ F\left(C_{\rho k_0} \cup \left(\bigcup_{\sigma \in L(T) \setminus \{\rho\}} C_\sigma\right)\right) & \xrightarrow{Fm} & FC \end{array} .$$

Applying Lem. C.1 to the above diagram and the result of Lem. C.2, we have

$$(\kappa_\sigma^{\text{cl},i+1})^* \kappa_\sigma^* c^* \overline{F}R_{i+1} = (\kappa_\sigma^{\text{cl},i+1})^* \kappa_\sigma^* c^* \overline{F}R_i. \quad (9)$$

If $\sigma \in T_i$, the right-hand side is calculated as:

$$\begin{aligned} (\kappa_\sigma^{\text{cl},i+1})^* \kappa_\sigma^* c^* \overline{F}R_i &= \kappa^* (\kappa_\sigma^{\text{cl},i})^* \kappa_\sigma^* c^* \overline{F}R_i \\ &= \kappa^* (\top_{C_\sigma^{\text{cl},i}}) && \text{Induction hypothesis} \\ &= \top_{C_\sigma^{\text{cl},i+1}}. && \kappa_* \dashv \kappa^* \text{ and } \top_{C_\sigma^{\text{cl},i}} \text{ is a limit} \end{aligned}$$

Otherwise, we have $\sigma = \rho k$ for some $k \in \{0, \dots, n_\rho\}$ where $\rho \in L(T_i)$ is the chosen leaf at the **Partitioning** step. Then, we have

$$\begin{aligned} &(\kappa_{\rho k}^{\text{cl},i+1})^* \kappa_{\rho k}^* c^* \overline{F}R_i \\ &= (\kappa_{\rho k}^{\text{cl},i+1})^* \kappa_{\rho k,k}^* \kappa_\rho^* c^* \overline{F}R_i && \kappa_{\rho k} = \kappa_\rho \circ \kappa_{\rho k,k} \\ &= (\kappa_{\rho k}^{\text{cl},i+1})^* (\top_{C_{\rho k}}) && \text{Def. 4.1.1} \\ &= \top_{C_{\rho k}^{\text{cl},i+1}}. && (\kappa_{\rho k}^{\text{cl},i+1})_* \dashv (\kappa_{\rho k}^{\text{cl},i+1})^* \text{ and } \top_{C_{\rho k}} \text{ is a limit} \end{aligned}$$

Hence, we have $(\kappa_\sigma^{\text{cl},i+1})^* \kappa_\sigma^* c^* \overline{F}R_{i+1} = \top_{C_\sigma^{\text{cl},i+1}}$ for each $\sigma \in L(T_{i+1})$.

Proposition 6.7 (termination). *Algo. 2 terminates.*

Proof. To prove termination, we use the lexicographical order \trianglelefteq on $\mathbb{E}_C \times \mathbb{N}$, where the order of \mathbb{N} is the usual order \leq . The order \leq is well-founded, and \sqsubseteq is well-founded by Assum. 6.1.7. Thus, \trianglelefteq is also well-founded. Let $d_i = |\{\sigma \in L(T_i) \mid C_\sigma^{\text{cl},i} \neq C_\sigma\}|$.

We show that if $C_\rho^{\text{cl},i} \neq C_\rho^i$ holds for some $\rho \in L(T)$, then we have $(R_{i+1}, d_{i+1}) \triangleleft (R_i, d_i)$. Suppose that, at the line 4 of the i -th iteration of the main loop, a leaf $\rho \in L(T_i)$ such that $C_\rho^{\text{cl},i} \neq C_\rho^i$ is chosen. We have two cases.

Case $(c \circ \kappa_\rho)^*(\overline{FR}_i) \sqsubset \top_{C_\rho}$. In this case, we have $R_i \sqsubset R_{i+1}$ by Assum. 6.1.9. Thus, $(R_{i+1}, d_{i+1}) \triangleleft (R_i, d_i)$ holds.

Case $(c \circ \kappa_\rho)^*(\overline{FR}_i) = \top_{C_\rho}$. In this case, the line 7 is executed, and the number of $\sigma \in T$ with $C_\sigma^{\text{cl}} \neq C_\sigma$ reduces: $d_{i+1} < d_i$.

By the above argument shows that, after finite number of iteration, we have $C_\rho^{\text{cl}} = C_\rho$ for all $\rho \in L(T)$. This means that Algo. 2 terminates.

Theorem 6.8 (correctness). *Algo. 2 terminates and returns a $\nu(c^*\overline{F})$ -partitioning.*

Proof. Termination is ensured by Prop. 6.7. Let R, T and $\{C_\sigma \rightsquigarrow C\}_{\sigma \in L(T)}$ be as defined in the last iteration of the main loop. Our goal is to show that $\{C_\sigma \rightsquigarrow C\}_{\sigma \in L(T)}$ is a $\nu(c^*\overline{F})$ -partitioning of C .

We prove $R = c^*\overline{F}(R)$. By Prop. 6.6, $\{C_\sigma \rightsquigarrow C\}_{\sigma \in L(T)}$ is an R -partitioning. We also have $\kappa_\sigma^*(c^*\overline{F}(R)) = \top_{C_\sigma}$ for each $\sigma \in L(T)$. Thus, we have

$$\begin{aligned} R &= \bigsqcup_{\sigma \in L(T)} (\kappa_\sigma)_*(\top_{C_\sigma}) && \text{the definition of } R \\ &= \bigsqcup_{\sigma \in L(T)} (\kappa_\sigma)_*(\kappa_\sigma^*(c^*\overline{F}(R))) && \text{Def. 4.1.2} \\ &= c^*\overline{F}(R). && \text{Assum. 6.1.5} \end{aligned}$$

From Prop. 6.6.3, we have $\nu(c^*\overline{F}) \sqsubseteq R$. Hence, we have $R = \nu(c^*\overline{F})$.

D Illustration of Runs of Algo. 3

We illustrate Algo. 3 for non-deterministic automata.

Example D.1. Let $\Sigma = \{a, b\}$, $C = \{x, y, z, w, v\}$, N_Σ be the functor from Example 3.10, and $\overline{N}_\Sigma = \text{Rel}(N_\Sigma)$ (Def. 3.6). We define a coalgebra $c: C \rightarrow N_\Sigma C$ as




Fig. 8: The snapshots at the end of first and second iterations of $\text{fPR}_{(N_\Sigma, \overline{N}_\Sigma), w_C}^{\text{H-ER}}$ and $\text{fPR}_{(N_\Sigma, \overline{N}_\Sigma), w_P}^{\text{H-ER}}$ for c . Yellow areas depict partitions which are refined as the main loop repeats. Clean states are white and dirty states are black.

shown in Fig. 8. We compare the executions $\text{fPR}_{(N_\Sigma, \overline{N_\Sigma}), w_C}^{\text{H-ER}}$ and $\text{fPR}_{(N_\Sigma, \overline{N_\Sigma}), w_P}^{\text{H-ER}}$. At the both initialisations (Line 1), we have a tree whose sole leaf is equal to C , which represents the initial partition, and every state in C marked as dirty.

In the first iteration, both $\text{fPR}_{(N_\Sigma, \overline{N_\Sigma}), w_C}^{\text{H-ER}}$ and $\text{fPR}_{(N_\Sigma, \overline{N_\Sigma}), w_P}^{\text{H-ER}}$ split C_ϵ and obtain C_0 and C_1 . The k_0 chosen at Line 10 is 0 (i.e. C_0) in both algorithms, and predecessors of z are marked as dirty (the centre figure of Fig. 8).

In the second iteration, both $\text{fPR}_{(N_\Sigma, \overline{N_\Sigma}), w_C}^{\text{H-ER}}$ and $\text{fPR}_{(N_\Sigma, \overline{N_\Sigma}), w_P}^{\text{H-ER}}$ split C_0 and obtain C_{00} and C_{01} . The k_0 chosen at Line 10 of $\text{fPR}_{(N_\Sigma, \overline{N_\Sigma}), w_C}^{\text{H-ER}}$ is 1 (i.e. C_{01}) because $w_C(C_{01}) = 3 > 1 = w_C(C_{00})$, and predecessors of w are marked as dirty (the left figure of Fig. 8). The k_0 chosen at Line 10 of $\text{fPR}_{(N_\Sigma, \overline{N_\Sigma}), w_P}^{\text{H-ER}}$ is 0 (i.e. C_{00}) because $w_P(C_{00}) = 4 > 3 = w_P(C_{01})$, and predecessors of the states in C_{01} are marked as dirty (the right figure of Fig. 8).

$\text{fPR}_{(N_\Sigma, \overline{N_\Sigma}), w_P}^{\text{H-ER}}$ marks states as dirty less than $\text{fPR}_{(N_\Sigma, \overline{N_\Sigma}), w_C}^{\text{H-ER}}$. Hence, $\text{fPR}_{(N_\Sigma, \overline{N_\Sigma}), w_P}^{\text{H-ER}}$ terminates in fewer steps compared with $\text{fPR}_{(N_\Sigma, \overline{N_\Sigma}), w_C}^{\text{H-ER}}$.