

FUNCTORS ON RELATIONAL STRUCTURES WHICH ADMIT BOTH LEFT AND RIGHT ADJOINTS

Víctor Dalmau

Universitat Pompeu Fabra, Barcelona, Spain
victor.dalmau@upf.edu

Andrei Krokhin

Durham University, UK
andrei.krokhin@durham.ac.uk

Jakub Opršal

University of Birmingham, UK
j.oprsal@bham.ac.uk

Abstract. This paper describes several cases of adjunction in the homomorphism pre-order of relational structures. We say that two functors Λ and Γ between thin categories of relational structures are *adjoint* if for all structures $\mathbf{A}$ and $\mathbf{B}$, we have that $\Lambda(\mathbf{A})$ maps homomorphically to $\mathbf{B}$ if and only if $\mathbf{A}$ maps homomorphically to $\Gamma(\mathbf{B})$. If this is the case, Λ is called the left adjoint to Γ and Γ the right adjoint to Λ . In 2015, Foniok and Tardif described some functors on the category of digraphs that allow both left and right adjoints. The main contribution of Foniok and Tardif is a construction of right adjoints to some of the functors identified as right adjoints by Pultr in 1970. We generalise results of Foniok and Tardif to arbitrary relational structures, and coincidentally, we also provide more right adjoints on digraphs, and since these constructions are connected to finite duality, we also provide a new construction of duals to trees. Our results are inspired by an application in promise constraint satisfaction — it has been shown that such functors can be used as efficient reductions between these problems.

Key words and phrases. relational structure, digraph, homomorphism, homomorphism duality, constraint satisfaction problem.

1. INTRODUCTION

The study of homomorphisms between (di)graphs and general relational structures plays an important role in combinatorics and theoretical computer science [HN04, FV98, KŽ17]. The importance in theoretical computer science stems, in particular, from the fact that the Constraint Satisfaction Problem (CSP) and its relatives can be cast as the problem of existence of a homomorphism from one relational structure to another. The class of all relational structures of a given signature (e.g., all digraphs) admits the homomorphism preorder, where $\mathbf{A} \leq \mathbf{B}$ for two structure $\mathbf{A}$ and $\mathbf{B}$ if and only if there exists a homomorphism from $\mathbf{A}$ to $\mathbf{B}$. This preorder is interesting in its own right [HN04], and, moreover, important computational problems such as non-uniform CSPs and promise CSPs can be stated in terms of a relative position of an input structure with respect to a fixed structure or a pair of fixed structures. Specifically, the CSP with a template $\mathbf{A}$ (which is a fixed structure) asks whether a given structure $\mathbf{I}$ satisfies $\mathbf{I} \leq \mathbf{A}$ [FV98, KŽ17], and the *promise CSP* with a template $\mathbf{A}, \mathbf{B}$ (which is a pair of structures such that $\mathbf{A} \leq \mathbf{B}$) asks to distinguish between the cases $\mathbf{I} \leq \mathbf{A}$ and $\mathbf{I} \not\leq \mathbf{B}$ [AGH17, KO22]; the promise being that input falls into one of the two (disjoint) cases.

A (*thin*) *functor* from the class of structures of some signature to the class of structures of a possibly different signature is a mapping that is monotone with respect to the homomorphism preorder. Many well-known examples of constructions in graph

Victor Dalmau was supported by the MICIN under grants PID2019-109137GB-C22 and PID2022-138506NB-C22, and the Maria de Maeztu program (CEX2021-001195-M). Andrei Krokhin and Jakub Opršal were supported by the UK EPSRC grant EP/R034516/1. Andrei Krokhin was also supported by the UK EPSRC grant EP/X033201/1. This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie Grant Agreement No 101034413.

©V. Dalmau, A. Krokhin, J. Opršal; licensed under Creative Commons License CC-BY 4.0.

theory, such as the arc graph construction, are functors [FT13, FT15]. In this paper we investigate a kind of homomorphism duality of such functors called *adjunction* — two functors Λ and Γ form an adjoint pair if $\Lambda(\mathbf{H})$ maps homomorphically to $\mathbf{G}$ if and only if $\mathbf{H}$ maps homomorphically to $\Gamma(\mathbf{G})$ for all relational structures $\mathbf{G}, \mathbf{H}$ of the corresponding signatures. In this case, Λ is called a left adjoint of Γ and Γ is a right adjoint of Λ . If, for some Γ , there exists such Λ we also say that Γ has (or admits) a right adjoint.

Adjunction is a core concept of category theory, although we should note that here we work with posetal categories of relational structures, i.e., only existence of a homomorphism matters, rather than with the category of relational structures together with homomorphisms. Adjoints between categories of relational structures in the stricter categorical sense have been completely described by Pultr [Pul70]. Here, we call these functors *left and central Pultr functors* (following the nomenclature of Foniok and Tardif [FT15]) — a left Pultr functor is a left adjoint to a central Pultr functor. It is clear, that each pair of left and central Pultr functors is adjoint in our sense as well. Nevertheless, there are more functors that admit a right adjoint than left Pultr functors, e.g., the afore-mentioned arc graph construction. In this paper, we ask the question which central Pultr functors (i.e., the functors which admit a left adjoint by Pultr’s characterisation) admit a right adjoint?

A necessary condition for a central Pultr functor to have a right adjoint was given by Foniok and Tardif [FT15]. They also provided explicit constructions of the right adjoints for some of such functors mapping digraphs to digraphs. Loosely speaking, they gave explicit constructions of right adjoints to functors that either do not change the domain (i.e., the set of vertices) of the digraph, or such that the new domain is the set of edges (arcs) of the input digraph. In the present paper, we extend their results by proving an explicit construction that works with general relational structures. In particular, we provide adjoints to some of the functors between classes of relational structures satisfying the above-mentioned necessary condition, and that either do not change the domain of the structure, or the new domain is one of the relations of the input structure. Hence we prove a direct generalisation of [FT15] to arbitrary relational signatures (see Sections 5 and 6). Furthermore, by composing right adjoints constructed in this way, we obtain more adjunctions than [FT15] even for the digraph case (see Section 7). Finally, we believe that our constructions are more intuitive than those provided by [FT15] — in particular, the above mentioned necessary condition relates adjunction with *finite duality*, and therefore a construction of a right adjoint is related to a construction of a *dual of a tree*. Moreover, every finite tree can be built via natural inductive process. Our constructions, and the proofs that they work, reflect this inductive process for certain trees used in the definition of the functor.

2. PRELIMINARIES

We recall some basic definitions and notation. We use the symbol $\mathcal{P}(X)$ to denote the power set of a set X .

2.1. Structures and homomorphisms

A directed graph can be defined as a pair $\mathbf{G} = (G, E^{\mathbf{G}})$ where G is the set of vertices of $\mathbf{G}$ and $E^{\mathbf{G}} \subseteq G \times G$ is the set of edges of $\mathbf{G}$. This is a special case of a relational structure with a single relation of arity 2 as defined below.

Definition 2.1. A *relational signature* τ is a tuple of relational symbols $R, S, \dots$ where each symbol is assigned a positive integer, called an arity and denoted by $\text{ar } R, \text{ar } S, \dots$

A *relational τ -structure* is a tuple $\mathbf{A} = (A; R^{\mathbf{A}}, S^{\mathbf{A}}, \dots)$, where A is a set called the *domain (or universe)* of $\mathbf{A}$, and $R^{\mathbf{A}} \subseteq A^{\text{ar } R}, S^{\mathbf{A}} \subseteq A^{\text{ar } S}, \dots$ are relations on this domain of the corresponding arity.

The relational symbols $R, S, \dots$ in τ are also referred to as *τ -symbols*. When no confusion can arise, we say a τ -structure (dropping “relational”), or even simply a structure,

when τ is clear from the context. Two structures of the same signature are said to be *similar*. We will call elements from the domain of some structure *vertices*, and tuples in a relation R of some structure *R-edges*, or simply *edges* if the symbol R is either irrelevant or clear from the context.

In the rest of the paper, we will not use the symbol V as a relational symbol — we restrict its use to refer to the domain of the structure (e.g., in Definition 2.5 and Section 3 below).

Loosely speaking, a homomorphism between two similar structures is a map that preserves relations, e.g., a graph homomorphism would be a map between the two graphs that preserves edges. Formally, a homomorphism is defined as follows.

Definition 2.2. Let $\mathbf{A}$ and $\mathbf{B}$ be two structures of the same signature τ . A *homomorphism* $f: \mathbf{A} \rightarrow \mathbf{B}$ is defined to be a mapping $f: A \rightarrow B$ such that, for each relational symbol R in τ and each $(a_1, \dots, a_{\text{ar } R}) \in R^{\mathbf{A}}$, we have

$$(f(a_1), \dots, f(a_{\text{ar } R})) \in R^{\mathbf{B}}.$$

We will write $\mathbf{A} \rightarrow \mathbf{B}$ if there exists a homomorphism from $\mathbf{A}$ to $\mathbf{B}$. The set of all homomorphisms from $\mathbf{A}$ to $\mathbf{B}$ is denoted by $\text{hom}(\mathbf{A}, \mathbf{B})$.

The above definition can be rephrased by saying that the function $f: A \rightarrow B$ has a coordinate-wise action on each of the relations R , i.e., for each relational symbol R in the signature, the expression

$$f^R((a_1, \dots, a_{\text{ar } R})) = (f(a_1), \dots, f(a_{\text{ar } R}))$$

defines a function $f^R: R^{\mathbf{A}} \rightarrow R^{\mathbf{B}}$. We use the symbol f^R throughout this paper.

Finally, since we will be extensively working with the homomorphism preorder, this in particular means that we will often work with structures up to homomorphic equivalence — we say that two structures $\mathbf{A}$ and $\mathbf{B}$ are *homomorphically equivalent* if we have $\mathbf{A} \rightarrow \mathbf{B}$ and $\mathbf{B} \rightarrow \mathbf{A}$. Such structures would be identified if we followed the standard procedure to turn the homomorphism preorder into a proper partial order. Note that two structures $\mathbf{A}$ and $\mathbf{B}$ are homomorphically equivalent if and only if for every structure $\mathbf{C}$, we have $\mathbf{C} \rightarrow \mathbf{A}$ if and only if $\mathbf{C} \rightarrow \mathbf{B}$, i.e., they allow homomorphisms from the same structures. The same is also true for allowing homomorphisms to the same structures, i.e., $\mathbf{A}$ and $\mathbf{B}$ are homomorphically equivalent if and only if, for all $\mathbf{C}$, we have $\mathbf{A} \rightarrow \mathbf{C}$ if and only if $\mathbf{B} \rightarrow \mathbf{C}$. A structure is called a *core*, if it is not homomorphically equivalent to any of its proper substructures.

Certain structures called *trees* play a special role in this paper. We use a definition of a tree equivalent to the one given in [NT00, Section 3]. Loosely speaking, a relational structure is a *tree* if it is connected, contains no cycles, and none of its relations has tuples with repeated entries. This is more precisely defined by using the incidence graph of a structure.

Definition 2.3. The *incidence graph* of a structure $\mathbf{A} = (A; R^{\mathbf{A}}, \dots)$ is a bipartite multi-graph whose vertex set is the disjoint union of A and $R^{\mathbf{A}}$ for each relational symbol R . There is an edge connecting every tuple $(a_1, \dots, a_k) \in R^{\mathbf{A}}$ with every one of its coordinates $a_i \in A$. In particular, if some element appears multiple times in this tuple, then the edge connecting it to the tuple appears with the same multiplicity. In this case, the incidence graph contains a cycle of length 2 (i.e., two parallel edges).

A τ -structure is a (τ) -*tree* if its incidence graph is a tree, i.e., an acyclic connected digraph.

Remark 2.4. An undirected graph can be viewed as a relational structure with a single relation E whose universe is the set of vertices V . The relation $E \subseteq V \times V$ contains for every edge two tuples (u, v) and (v, u) . This means that no undirected graph with at least one edge is a tree according to the above definition since if (u, v) is an edge, then $u, (u, v), v, (v, u)$ is a 4-cycle of the incidence graph. Intuitively, relational structures with binary relations are directed graphs; an undirected graph is encoded as directed

by including both orientations of each edge which results in a directed cycle of length 2. Under the above definition, a directed graph is a tree if it is an oriented tree.

We now fix notation for some small structures that will be used later.

Definition 2.5. Fix a relational signature σ . We define the following structures:

- $\mathbf{V}_1$ is the structure with a single vertex, i.e., $V_1 = \{1\}$, and empty relations, i.e., $R^{V_1} = \emptyset$ for all σ -symbols R .
- Let S be a relational symbol, $\mathbf{S}_1$ is a structure with $\text{ar } S$ vertices related by S and all other relations empty. More precisely, $S_1 = \{1, \dots, k\}$ where $k = \text{ar } S$, $S^{\mathbf{S}_1} = \{(1, \dots, k)\}$, and $R^{\mathbf{S}_1} = \emptyset$ for all σ -symbols R except S .

Since adjunction is a form of duality, we will often mention homomorphism duality of structures.

Definition 2.6. A pair $(\mathbf{T}, \mathbf{D})$ of similar structures is called a *duality pair* if, for all structures $\mathbf{A}$ similar to $\mathbf{T}$, either $\mathbf{T} \rightarrow \mathbf{A}$ or $\mathbf{A} \rightarrow \mathbf{D}$. In this case, $\mathbf{D}$ is called a dual to $\mathbf{T}$.

It was shown in [NT00] that a structure has a dual if and only if this structure is homomorphically equivalent to a tree. We will also give an alternative proof of one of the implications in Section 4.

2.2. Pultr functors: gadget replacement and pp-constructions

Traditionally, in the CSP literature, pp-constructions would be described in the language of logic using so called *primitive positive formulae* (logical formulae that use only $\exists$, $\wedge$, and $=$). We refer to [BKW17, Definition 19] for details. In this paper, we define “pp-constructions” using a language similar to [FT15, Definitions 2.1–2.3].

Definition 2.7. Let σ and τ be two relational signatures. A (σ, τ) -Pultr template is a tuple of σ -structures consisting of $\mathbf{P}$, and $\mathbf{Q}_R$, one for each τ -symbol R , together with homomorphisms $\epsilon_{i,R}: \mathbf{P} \rightarrow \mathbf{Q}_R$ for each τ -symbol R and all $i \in \{1, \dots, \text{ar } R\}$.

Definition 2.8. Given a (σ, τ) -Pultr template as above, we define two functors Λ and Γ , called (*left and central*) Pultr functors.

- Given a τ -structure $\mathbf{A}$, we define a σ -structure $\Lambda(\mathbf{A})$ in the following way: For each $a \in A$, introduce to $\Lambda(\mathbf{A})$ a copy of $\mathbf{P}$ denoted by $\mathbf{P}_a$, and for each τ -symbol R and each $(a_1, \dots, a_k) \in R^{\mathbf{A}}$, introduce to $\Lambda(\mathbf{A})$ a copy of $\mathbf{Q}_R$ with the image of $\mathbf{P}$ under $\epsilon_{i,R}$ identified with $\mathbf{P}_{a_i}$ for all $i \in \{1, \dots, k\}$.
- Given a σ -structure $\mathbf{B}$, we define a τ -structure $\Gamma(\mathbf{B})$ whose universe consists of all homomorphisms $h: \mathbf{P} \rightarrow \mathbf{B}$. The relation $R^{\Gamma(\mathbf{B})}$ where R is a τ -symbol is then defined to contain all tuples $(h_1, \dots, h_k)$ of such homomorphisms for which there is a homomorphism $g: \mathbf{Q}_R \rightarrow \mathbf{B}$ such that $h_i = g \circ \epsilon_{i,R}$ for all $i \in \{1, \dots, k\}$.

Once the definitions are settled, it is not too hard to show that, for any Pultr template, the corresponding Pultr functors Λ and Γ are left and right adjoints. This statement is attributed to [Pul70], although it was rediscovered on numerous occasions, and it is considered folklore in category theory.

The fact that both Λ and Γ are (thin) functors follows from the adjunction. It can be also easily proved directly, for example, a homomorphism $f^\Gamma: \Gamma(\mathbf{A}) \rightarrow \Gamma(\mathbf{B})$ can be obtained from a homomorphism $f: \mathbf{A} \rightarrow \mathbf{B}$ by setting $f^\Gamma(h) = f \circ h$ for each $h: \mathbf{P} \rightarrow \mathbf{A}$.

Example 2.9. Consider the Pultr template consisting of structures $\mathbf{P} = (\{0, 1\}; \{(0, 1)\})$ and $\mathbf{Q}_E = (\{0, 1, 2\}; \{(0, 1), (1, 2)\})$ with $\epsilon_1(0) = 0$ and $\epsilon_1(1) = 1$, and $\epsilon_2(0) = 1$ and $\epsilon_2(1) = 2$.

The corresponding central Pultr functor Γ is the arc-graph construction that is usually denoted by δ ; given a (directed) graph $\mathbf{G} = (G, E^G)$, the digraph $\delta(\mathbf{G})$ is defined as the graph with the vertex set E^G and edges $((u, v), (v, w)) \in E^G \times E^G$.

The left Pultr functor Λ corresponding to this template provides a left adjoint to the arc-graph construction, and can be explicitly described as follows: Given a digraph $\mathbf{G}$,

the digraph $\Lambda(G)$ is constructed by replacing each vertex $v \in G$ with a pair of vertices v_0, v_1 connected by an edge, i.e., with $(v_0, v_1) \in E^{\Lambda(G)}$. Furthermore, for each edge $(u, v) \in E^G$, we identify the vertices u_1 and v_0 . For example, the image of a path of length n , i.e., the digraph $P_n = (\{0, \dots, n\}; \{(i, i+1) \mid i \in \{0, n-1\}\})$ is the path of length $n+1$ consisting of vertices $0_0, 0_1 = 1_0, 1_1 = 2_0, \dots, n_1$.

Observe that, in this case, we have that $\Lambda(\Gamma(P_n))$ is isomorphic to P_n . This does not need to happen for any digraph in place of P_n , nevertheless, the existence of a homomorphism $\Lambda(\Gamma(G)) \rightarrow G$ for all digraphs G follows from the adjunction in a straightforward way.

We remark that the connection between the adjunction of left and central Pultr functors and the algebraic reductions between CSPs has been described in [KOWŻ23, Section 4.1]. Let us briefly mention that the central Pultr functors are *pp-constructions* (more precisely, a structure A is pp-constructible from B if it is homomorphically equivalent to the structure $\Gamma(B)$ for some Pultr functor Γ), and the left Pultr functors are called *gadget replacements* in this context.

The main contribution of the present paper is an investigation of the cases of Pultr templates, for which Γ is also a left adjoint, i.e., it admits a right adjoint Ω . The following necessary condition for this was proved in [FT15] for the case of digraphs, but the proof goes through verbatim for general structures; we include it for completeness.

Theorem 2.10 ([FT15, Theorem 2.5]). *Let Λ and Γ be a pair of left and central Pultr functors defined by a (σ, τ) -Pultr template. If Γ has a right adjoint, then, for each τ -tree T , $\Lambda(T)$ is homomorphically equivalent to a σ -tree.*

Proof. A structure has a dual if and only if it is homomorphically equivalent to a tree [NT00]. Thus, it suffices to prove that if Γ has both a left adjoint Λ and a right adjoint Ω , and (T, D) is a duality pair, then $(\Lambda(T), \Omega(D))$ is a duality pair as well. To show that, observe that, for any A , $\Lambda(T) \not\rightarrow A$ is equivalent to $T \not\rightarrow \Gamma(A)$ since Λ is a left adjoint to Γ . The latter condition is equivalent to $\Gamma(A) \rightarrow D$, since D is a dual of T , which in turn is equivalent to $A \rightarrow \Omega(D)$, since Ω is a right adjoint to Γ . $\square$

Corollary 2.11. *If a central Pultr functor Γ has a right adjoint, then all the structures in its Pultr template are homomorphically equivalent to trees.*

Proof. As shown in [FT15] (and not hard to see), if the template consists of structures P and Q_R for τ -symbols R , then P is isomorphic to $\Lambda(V_1)$ and Q_R is isomorphic to $\Lambda(R_1)$ for each τ -symbol R (where tree τ -structures V_1 and R_1 are as in Definition 2.5). $\square$

It is open whether the condition in Theorem 2.10 is also sufficient to have a right adjoint. In this paper, as well as Foniok-Tardif do in [FT15], we focus on the cases when P and Q_R 's are actually trees, and we give a concrete construction of the adjoint in two cases: $P = V_1$ (this is shown in Section 5), and $P = S_1$ is the tree with a single S -edge for some σ -symbol S (this is shown in Section 6). Finally, in the last section, we combine these two constructions to prove that if P and Q_R 's are trees, and, moreover, the images of P in Q_R under the maps $\epsilon_{i,R}$ are either disjoint or intersect in one vertex, then the corresponding central Pultr functor has a right adjoint (Theorem 7.3).

Finally, let us note a well-known fact from category theory that if Ω_1 and Ω_2 are both right adjoints to Γ , then $\Omega_1(A)$ and $\Omega_2(A)$ are homomorphically equivalent for all structures A . The proof is immediate from the definitions: both structures allow a homomorphism from B if and only if $\Gamma(B) \rightarrow A$.

3. AN INDUCTIVE CONSTRUCTION OF TREES

For our constructions, it will be convenient to describe τ -trees by certain formal terms. Our terms will correspond to trees that are rooted, either in a vertex or in an edge (i.e., a tuple in one of the relations). Each term will correspond to a unique inductive construction of a tree, but the same tree can be obtained by several inductive constructions.




FIGURE 1. Examples of terms and trees they represent.

We recall that we assume that none of the relational signatures uses V as a relational symbol. Each term will be either a V -term (V stands for *vertices*) or an R -term where R is a relational symbol in τ . The rules for the inductive construction of terms and their semantics are as follows:

- vertex is a V -term.
- If $t_1, \dots, t_k$ are V -terms and R is a k -ary symbol in τ , then $\text{edge}_R(t_1, \dots, t_k)$ is an R -term.
- If t is an R -term and $i \in \{1, \dots, \text{ar } R\}$, then $\text{pr}_i(t)$ is a V -term.


The rooted tree $(T(t), r_t)$, where r_t is the root, corresponding to a term t is defined in the following way:

- $T(\text{vertex})$ is the one-vertex tree, rooted at its only vertex, i.e., $T(\text{vertex}) = V_1$ and $r_t = 1$.
- If $t = \text{edge}_R(t_1, \dots, t_k)$ is an R -term, then $T(t)$ is the tree obtained by taking the disjoint union of the trees $T(t_i)$ (with roots r_{t_i}) and adding the tuple $(r_{t_1}, \dots, r_{t_k})$ to the relation R . The root r_t is defined to be this new tuple.
- If $t = \text{pr}_i(s)$ is a V -term and $r_s = (v_1, \dots, v_k)$, then we set $T(t) = T(s)$ and $r_t = v_i$.

The domain of $T(t)$ is denoted by $T(t)$.

We say that t *represents a tree* T if $T(t)$ is isomorphic to T , and that t *represents a tree* T *rooted in* r if $T(t)$ is isomorphic to T via an isomorphism mapping r_t to r . The same tree can in general be represented by several terms (see the following example), but each term represents a unique tree up to isomorphism.

Example 3.1. The following are examples of terms and the trees they represent in the relational signature of digraphs. The corresponding roots are highlighted. We drop the brackets around arguments of pr_i to ease readability.



Note that the difference between the second tree and the third tree is the root; the latter tree is obtained from the other by applying pr_2 , and thus only changing the root. Further, we present a few examples of more complicated trees in Fig. 1.

We remark that a tree can be represented by multiple terms even when we fix the root. For example, the last tree in Fig. 1 can be also represented by the term

$$\text{edge}_E(\text{pr}_2 \text{edge}_E(\text{vertex}, \text{pr}_1 \text{edge}_E(\text{vertex}, \text{vertex})), \text{pr}_2 \text{edge}_E(\text{vertex}, \text{vertex})).$$

Lemma 3.2. *Fix a relational signature τ . Any finite τ -tree can be represented by a term. More precisely,*

- *for every finite tree T and $r \in T$, there is a V -term t such that $T(t)$ is isomorphic to T via an isomorphism that maps r_t to r , and*
- *for every finite tree T and $r \in R^T$ for some τ -symbol R , there is an R -term t such that $T(t)$ is isomorphic to T via an isomorphism that maps r_t to r .*

The proof is a simple argument by induction on the number of edges of the tree. We include it in detail to provide more intuition about terms.

Proof. We prove this statement by induction on the number of edges of T . Each induction step is moreover split in two: we first prove that trees with n edges rooted in an edge can be represented, and then, assuming the above, we show that trees with n edges rooted in a vertex can be represented.

- (1) We start with $n = 0$. There is a single tree T with no edges, namely the tree with one vertex. It cannot be rooted in an edge, but it can be rooted in its only vertex. It is represented by the term `vertex`.
- (2) Assume $n > 0$, and we can represent all trees rooted in a vertex with less than n edges. Assume that T has n edges and is rooted in an edge $(r_1, \dots, r_k) \in R^T$ for some k -ary symbol R . Removing this edge from T splits T into k connected components $T_1, \dots, T_k$, where T_i contains r_i for each i . Each of these components is a tree. Assuming that T_i rooted in r_i is represented by the term t_i for each i , T rooted in $(r_1, \dots, r_k)$ is represented by $\text{edge}_R(t_1, \dots, t_k)$.
- (3) Assume that $n > 0$, and we can represent all trees rooted in an edge with at most n edges. Let T be a tree with n edges and $r \in T$. Since T is connected, r is involved in some edge, say $(v_1, \dots, v_k) \in R^T$ where $k = \text{ar } R$ and $r = v_i$ for some i . Now, T rooted in $(v_1, \dots, v_k)$ is represented by an R -term t by the inductive assumption, so T rooted in r is represented by $\text{pr}_i(t)$. $\square$

Finally, our constructions use the notion of a *subterm* of a term. Intuitively a subterm of a term t is any proper term that appears as a part of t . The set of all subterms of t encodes all intermediate byproducts of the inductive construction of $T(t)$. Formally, we define subterms as follows.

- The only subterm of `vertex` is itself.
- If $t = \text{edge}_R(t_1, \dots, t_k)$ is an R -term, where $t_1, \dots, t_k$ are V -terms and R is a k -ary symbol in τ , then the set of its subterms consists of the term itself and all subterms of $t_1, \dots, t_k$.
- If $t = \text{pr}_i(s)$ is a V -term, where s is an R -term, then the set of its subterms consists of the term itself and all subterms of s .

If s is a subterm of t , then we write $s \leq t$, moreover, if s is a proper subterm of t , i.e., it is a subterm and $s \neq t$, we write $s < t$. A V -subterm of a term t is a subterm which is a V -term, and similarly an R -subterm for a relational symbol R is a subterm which is an R -term.

For example, the term

$$t = \text{edge}_E(\text{pr}_1(\text{edge}_E(\text{vertex}, \text{vertex})), \text{pr}_1(\text{edge}_E(\text{vertex}, \text{vertex})))$$

has four distinct subterms: two E -terms, which are t and $\text{edge}_E(\text{vertex}, \text{vertex})$, and two V -terms, which are $\text{pr}_1(\text{edge}_E(\text{vertex}, \text{vertex}))$ and `vertex`.

We note that statements about terms can be proven by an inductive principle: showing the statement is true for the term `vertex`, and then showing that if it is true for all proper subterms of a term t , it is also true for t .

4. PRELUDE: DUALS TO TREES

Before we get to the main construction of adjoints, let us briefly discuss a simpler construction of a dual of a tree. There are a few similarities between the construction of duals and right adjoints to central Pultr functors: as we mentioned before (see the proof of Theorem 2.10), if Γ is a central Pultr functor that has a left adjoint Λ and a right adjoint Ω , and (T, D) is a duality pair, then $(\Lambda(T), \Omega(D))$ is also a duality pair. Moreover, our construction of the dual uses the inductive construction of trees from the previous section in a similar way as the constructions of right adjoints in Sections 5 and 6, but the construction of a dual is conceptually easier, so we present it to create some intuition that will be useful below. We also note that we show an explicit connection between constructions of a dual and the construction of a right adjoint in Section 5.3.

Again, we fix a relational signature. Recall that a pair (T, D) of similar structures is called a duality pair if, for all structures A similar to T , either $T \rightarrow A$ or $A \rightarrow D$. In this case, D is called a dual to T , and it was shown in [NT00] that a structure has a dual if and only if it is homomorphically equivalent to a tree. See [LLT07, NT05] for other characterisations of finite duality. Our construction is loosely inspired by the construction in [NT05].

Another way to look at duals is that for any structure A , a homomorphism $A \rightarrow D$ should correspond to a ‘proof’ that $T \not\rightarrow A$. How can one prove that a tree does not map to a structure A if that is the case? This can be done by an inductive argument. Let us outline this argument in the case T and A are digraphs. More precisely, we describe a procedure that shows that, for a fixed root $r \in T$ and some $a \in A$, there is no homomorphism from T to A that maps r to a . Pick a neighbour s of r in T , and assume that $(r, s) \in E^T$, the other orientation is dealt with symmetrically. We can show that there is no homomorphism from T to A mapping r to a by showing that, for none of the neighbours b of a in A , there is a homomorphism $T \rightarrow A$ that maps the edge (r, s) to (a, b) . In turn, removing the edge (r, s) from T splits the tree into two subtrees T_1 , containing r , and T_2 , containing s . A homomorphism $T \rightarrow A$ that maps (r, s) to (a, b) is equivalent to a pair of homomorphisms $T_1 \rightarrow A$, that maps r to a , and $T_2 \rightarrow A$, that maps s to b . We have thus reduced the claim to proving that there is no homomorphism from (at least) one of the two smaller trees, and can therefore recursively repeat our strategy for these two smaller trees. We further design a structure D , in which we can encode proofs of the above form. In particular, the image of an element $a \in A$ under a homomorphism $A \rightarrow D$ will contain answers to questions of the form ‘Is there a homomorphism $T' \rightarrow A$ that maps r to a ?’ for all trees T' and all roots $r \in T'$ that would appear in the above inductive argument.

Definition 4.1. Let t_Q be an S -term for some relational symbol S . Let $\mathcal{T}_V$ be the set of all V -subterms of t_Q , and let $\mathcal{T}_R$ be the set of all R -subterms of t_Q for each relational symbol R . We define a structure $D(t_Q)$.

The domain of this structure, denoted by $D(t_Q)$, is the set of all tuples $v \in \{\text{true}, \text{false}\}^{\mathcal{T}_V}$, indexed by V -subterms of t_Q , such that $v_{\text{vertex}} = \text{true}$.

To define edges, we introduce the following notation. For terms $t_1, \dots, t_k \in \mathcal{T}_V$, a relational symbol R of arity k , and $v^1, \dots, v^k \in D(t_Q)$, we let

$$\&_{\text{edge}_R(t_1, \dots, t_k)}(v^1, \dots, v^k) = v_{t_1}^1 \wedge \dots \wedge v_{t_k}^k.$$

A tuple $(v^1, \dots, v^k)$ of vertices is related in a relation R (of arity k) in $D(t_Q)$ if

(D1) For all $t \in \mathcal{T}_S$ and $i \in \{1, \dots, k\}$ such that $\text{pr}_i(t) \in \mathcal{T}_V$, we have

$$\&_t(v^1, \dots, v^k) \Rightarrow v_{\text{pr}_i(t)}^i.$$

(D2) If $R = S$, i.e., if t_Q is an R -term, then

$$\&_{t_Q}(v^1, \dots, v^k) = \text{false}.$$

If $e = (v_1, \dots, v_k)$, we will often write $\&_t(e)$ instead $\&_t(v_1, \dots, v_k)$. Note that if $e \in R^{D(t_Q)}$, then we have a tuple $e^* \in \{\text{true}, \text{false}\}^{\mathcal{T}_R}$ defined by $e_t^* = \&_t(e)$, which satisfies $e_{t_Q}^* = \text{false}$ (assuming t_Q is an R -term). This draws a parallel to how vertices are defined. Finally, note that item (D1) is essentially quantified by the V -subterms of t_Q , since all of such subterms t' , with the exception $t' = \text{vertex}$, are of the form $\text{pr}_i(t)$ for some t , and t and i are uniquely defined by t' .

Theorem 4.2. For any relational structure Q that is homomorphically equivalent to $T(t_Q)$ for some S -term t_Q where S is a relational symbol, $D(t_Q)$ is a dual of Q .

Proof. Without loss of generality, we may assume that $Q = T(t_Q)$. We need to show that, for all structures A , $A \rightarrow D(t_Q)$ if and only if $Q \not\rightarrow A$.

First, assume that $Q \not\rightarrow A$. We define $f: A \rightarrow D(t_Q)$ by putting, for each $u \in A$ and each $t \in \mathcal{T}_V$, $f(u)_t = \text{true}$ if there is a homomorphism $h: T(t) \rightarrow A$ that maps the root to u , and $f(u)_t = \text{false}$ otherwise. Clearly, $f(u)_{\text{vertex}} = \text{true}$. Now,

assume that $e = (u_1, \dots, u_k) \in R^A$. Observe that $\&_t(f^R(e))$ is true if and only if there is a homomorphism $h: T(t) \rightarrow A$ mapping the root edge to e . In other words, there is a homomorphism $h: T(t) \rightarrow A$ mapping the root edge to e if and only if there are homomorphisms $h_i: T(t_i) \rightarrow A$ with $h_i(r_{t_i}) = u_i$ for all i . Assuming that such a homomorphism h exists, the homomorphisms h_i are defined as restrictions of h to the corresponding subtrees. Assuming homomorphisms h_i exist, taking the union of h_i defines a mapping h on all vertices of $T(t)$. This mapping clearly preserves all edges different from the root, since h_i are homomorphisms, and it also preserves the root edge, since it is mapped to $e \in R^A$. This means that h is indeed a homomorphism. We need to show that $f^R(e)$ satisfies (D1) and (D2). For (D1), we want

$$\&_t(f^R(e)) \Rightarrow f(u_i)_{\text{pr}_i(t)},$$

i.e., if there is a homomorphism $T(t) \rightarrow A$ mapping the root to e then there is a homomorphism $T(\text{pr}_i(t)) \rightarrow A$ mapping the root to u_i . This is trivial, since $T(t) = T(\text{pr}_i(t))$ and a homomorphism $h: T(t) \rightarrow A$ that maps the root r_t to $(u_1, \dots, u_k)$ necessarily maps $r_{\text{pr}_i(t)}$ to u_i . Finally, (D2) is clear from the definition, since we assumed that $Q \not\vdash A$.

For the other implication, we first prove the following by induction on the term $t \leq t_Q$:

Claim 4.3. *Let $t \leq t_Q$ and let $h: T(t) \rightarrow D(t_Q)$ be a homomorphism. Then $h(r_t)_t = \text{true}$ if t is a V -term, and $\&_t(h^R(r_t)) = \text{true}$ if t is an R -term.*

- (1) The case $t = \text{vertex}$ is trivial.
- (2) Let $t = \text{edge}_R(t_1, \dots, t_{\text{ar}_R})$. We assume that $h: T(t) \rightarrow A$ is a homomorphism and $h(r_t) = (v^1, \dots, v^{\text{ar}_R})$. This in particular means that, for each i , h maps the root of $T(t_i)$ to v^i . Hence, we can apply the inductive assumption on the restrictions of h to $T(t_i)$'s to get that $v_{t_i}^i = \text{true}$ for all i , and consequently the claim by the definition of $\&_t(h^R(r_t))$.
- (3) Let $t = \text{pr}_{i,R}(t')$, and $h: T(t) \rightarrow A$. Note that $T(t) = T(t')$, so h is also a homomorphism from $T(t')$. Let $h^R(r_{t'}) = (v^1, \dots, v^k)$, and observe that $h(r_t) = v^i$. By the inductive assumption, this implies that $\&_{t'}(h^R(r_{t'})) = \&_{t'}(v^1, \dots, v^k) = \text{true}$. Consequently, we get that $h(r_t)_t = v_t^i = \text{true}$ from (D1).

This concludes the proof of the claim. Assume for a contradiction that $f: A \rightarrow D(t_Q)$ and $g: Q \rightarrow A$ are homomorphisms. Since $T(t_Q) = Q$, the above claim applied to $t = t_Q$ and the homomorphism $h = f \circ g: Q \rightarrow D(t_Q)$, would imply that $\&_{t_Q}(h^R(r_{t_Q})) = \text{true}$ which would contradict (D2). $\square$


4.1. Example: Dual to a directed path

It is well known (see, e.g., [HN04, Proposition 1.20]) that a directed graph maps homomorphically to the graph $L_k = (\{1, \dots, k\}; <)$, where the edge relation is given by the strict order on the domain, if and only if it does not allow a homomorphism from a directed path with k edges (and $k+1$ vertices) — we denote this path by P_k . This means that L_k is the dual of P_k . Let us compare this observation with our construction of the dual.

Fix $k > 0$. We pick the term

$$t_k = \text{edge}_E(\text{pr}_2(\text{edge}_E(\dots \text{pr}_2(\text{edge}_E(\text{vertex}, \text{vertex})), \dots, \text{vertex})), \text{vertex}),$$

where edge_E appears k times, to represent P_k rooted in its last edge, i.e., the following graph.



We also let $s_0 = \text{vertex}$, $t_i = \text{edge}_E(s_{i-1}, \text{vertex})$ and $s_i = \text{pr}_2(t_i)$ for $i \in \{1, \dots, k-1\}$. Note that both t_i and s_i represent a path of length i — the difference is the root which is either the last edge or the last vertex. Finally, note that $s_0, \dots, s_{k-1}$ and $t_1, \dots, t_k$ are the only subterms of t_k .

By definition, a vertex of $D(t_k)$ is a tuple

$$u \in \{\text{true}, \text{false}\}^{\{s_0, \dots, s_{k-1}\}}$$

such that $u_{s_0} = \text{true}$. This allows us to write them simply as ordered k -tuples whose i -th entry is the value corresponding to s_{i-1} .

To check whether u is connected by an edge to v or not, we consider the expressions $\&_{t_i}(u, v) = u_{s_{i-1}} \wedge v_{s_0}$ for all i . Using the definition, we get that (u, v) is an edge if

(D1) for all i , $\&_{t_i}(u, v) \Rightarrow v_{s_i}$ since $s_i = \text{pr}_2(t_i)$,¹ and

(D2) $\&_{t_k}(u, v) = \text{false}$.

Since $v_{s_0} = \text{true}$, we may simplify $\&_{t_i}(u, v) = u_{s_{i-1}}$, and substitute into the conditions above:

(1) $u_{s_{i-1}} \Rightarrow v_{s_i}$, for all $i < k$, and

(2) $u_{s_{k-1}} = \text{false}$.

In particular, observe that u has no out-edge (i.e., there is no edge of the form (u, v) for any v) if its last entry is true.

It is not hard to see that L_k maps to a dual constructed this way. We can construct a homomorphism h by mapping $i \in \{1, \dots, k\}$ to the tuple starting with i true's and followed by all false's, i.e.,

$$\begin{aligned} h(1) &= (\text{true}, \text{false}, \dots, \text{false}) \\ h(2) &= (\text{true}, \text{true}, \text{false}, \dots, \text{false}) \\ &\vdots \\ h(k) &= (\text{true}, \text{true}, \dots, \text{true}) \end{aligned}$$

Note that this is exactly the same homomorphism that is constructed in the proof of Theorem 4.2 (assuming $A = L_k$). And indeed, it is easy to check that if $i < j$, then $h(i)$ and $h(j)$ satisfy the conditions for an edge given above.

Naturally, there is also a homomorphism the other way. One such homomorphism maps a tuple u that begins with i true's followed by a false to i . Again, it is easy to check that if (u, v) is an edge then v has to begin with one more true. This establishes that our construction is homomorphically equivalent to L_k (as it should be).

Remark 4.4. Let us note that any homomorphism constructed according to the proof of Theorem 4.2 uses only the vertices in the image of h above, i.e., the tuples of the form

$$(\text{true}, \dots, \text{true}, \text{false}, \dots, \text{false}),$$

where true appears at least once and false does not need to appear at all. This is quite easy to see: $T(t_i)$ maps to $T(t_j)$ for all $i < j$ via a homomorphism preserving the roots, hence, for any u in the image, we get that if $u_{t_j} = \text{true}$ then $u_{t_i} = \text{true}$ for all $i < j$.

We could force similar implications in the definition of the dual by requiring that $u_t \Rightarrow u_s$ whenever there is a homomorphism $T(s) \rightarrow T(t)$ preserving roots. We did not include this condition in the definition because our goal is to get a simple construction and not necessarily that the construction results in the smallest graph possible. Nevertheless, this raises a question: Would it be possible by enforcing such implications on our general construction to produce a dual that would be a *core* (i.e., a structure that is not homomorphically equivalent to any of its proper substructures)?

Finally, we note that our construction of the dual of a tree can be naturally extended to any (finite) tree duality, i.e., given a finite set of finite trees $\mathcal{F} = \{T_1, \dots, T_n\}$, we can construct their dual D that will satisfy that, for any structure A , $A \rightarrow D$ if and only if for all $i = 1, \dots, n$, $T_i \not\rightarrow A$. This D is constructed as in Definition 4.1 with the following changes: Assume that a term t_i represents T_i rooted in an edge for each i (we are assuming that none of T_i 's consists of a single vertex). Let $\mathcal{T}$ be the set of all subterms of any of the t_i 's. We use notation $\mathcal{T}_V$ and $\mathcal{T}_R$ (where R is a relational symbol) for the V -terms and S -terms, respectively, that belong to this set. Finally, replace the condition (D2) with

¹ The implication $\&_{t_i}(u, v) \Rightarrow v_{\text{pr}_1(t_i)}$ is not needed since $\text{pr}_1(t_i)$ is not a subterm of t_k .

(D2') $\&_{t_i}(e) = \text{false}$ for all $i = 1, \dots, n$ (naturally, this only applies if t_i is an R -term).

We denote a structure constructed this way by $\mathbf{D}(t_1, \dots, t_n)$.

Theorem 4.5. *Let $q_1, \dots, q_n$ be terms for each $i \in \{1, \dots, n\}$, and assume that q_i is an R_i -term, where R_i is a relational symbol. For each structure $\mathbf{A}$, we have that either there exists i such that $\mathbf{T}(q_i) \rightarrow \mathbf{A}$ or $\mathbf{A} \rightarrow \mathbf{D}(q_1, \dots, q_n)$.*

Proof. The proof is essentially identical to the proof of Theorem 4.2. First, the implication ' $\mathbf{T}(q_i) \not\rightarrow \mathbf{A}$ for all $i \in \{1, \dots, n\}$ implies $\mathbf{A} \rightarrow \mathbf{D}(q_1, \dots, q_n)$ ' is proven by the same argument as before. In particular, we define $f: \mathbf{A} \rightarrow \mathbf{D}(q_1, \dots, q_n)$ by $f(u)_t = \text{true}$ if $\mathbf{T}(t) \rightarrow \mathbf{A}$ for each $t \in \mathcal{T}_V$ and $u \in A$, and $f(u)_t = \text{false}$ otherwise. The rest of the proof is identical except that instead of arguing that $\&_t(f^R(e))$ satisfies (D2) using $\mathbf{Q} \not\rightarrow \mathbf{A}$, we argue that $\&_t(f^R(e))$ satisfies (D2') for some i using that $\mathbf{T}(q_i) \not\rightarrow \mathbf{A}$.

For the other implication, it is enough to argue that $\mathbf{T}(q_i) \not\rightarrow \mathbf{D}(q_1, \dots, q_n)$ for each i , since then $\mathbf{T}(q_i)$ does not map homomorphically to any $\mathbf{A}$ that maps to $\mathbf{D}(q_1, \dots, q_n)$. Observe that Claim 4.3 and its proof is valid in this case. Consequently, if $h: \mathbf{T}(q_i) \rightarrow \mathbf{D}(q_1, \dots, q_n)$ for some i , then $\&_{q_i}(h^{R_i}(r_{q_i})) = \text{true}$ by the claim, which would contradict (D2'). $\square$

5. ADJOINTS TO FUNCTORS NOT CHANGING THE DOMAIN

In this section, we describe the simpler of the cases of our construction of an adjoint. We consider (σ, τ) -Pultr templates where $\mathbf{P} = \mathbf{V}_1$ (i.e., a vertex) and, for each τ -symbol R , $\mathbf{Q}_R$ is a σ -tree. The homomorphisms $\epsilon_{i,R}: \mathbf{P} \rightarrow \mathbf{Q}_R$ are given by picking elements $x_1, \dots, x_{\text{ar } R} \in \mathbf{Q}_R$ that are the images of the unique vertex of $\mathbf{P}$ under $\epsilon_{i,R}$ for the respective i 's. Note that some of the x_i 's might coincide, and $\mathbf{Q}_R$ can also have other vertices. Naturally, the elements $x_1, \dots, x_{\text{ar } R}$ depend on the symbol R which will be always clear from the context. This means that the structure $\Gamma(\mathbf{A})$ can be equivalently described in the following way: the universe of $\Gamma(\mathbf{A})$ coincides with the universe of $\mathbf{A}$, and for every τ -symbol R , we have

$$(1) \quad R^{\Gamma(\mathbf{A})} = \{(r(x_1), \dots, r(x_{\text{ar } R})) \mid r: \mathbf{Q}_R \rightarrow \mathbf{A}\}.$$

Finally, in this case we have that every homomorphism $h: \mathbf{A} \rightarrow \mathbf{B}$ is also a homomorphism $h: \Gamma(\mathbf{A}) \rightarrow \Gamma(\mathbf{B})$.

We will construct the right adjoint Ω in two steps: First, we define the set of vertices of a structure $\Omega(\mathbf{B})$, and second, we define the edges on these vertices. Before defining edges, we introduce new notation.

Throughout these definitions, we fix the following setting: Fix a (σ, τ) -Pultr template with $\mathbf{P}$ being a singleton structure with empty relations and with $\mathbf{Q}_R$ being a σ -tree for each τ -symbol R . First, for each R , we pick a term t_R representing $\mathbf{Q}_R$, and define $\mathcal{T}$ to be the set of all subterms of any of the t_R 's. We use notation $\mathcal{T}_V$ and $\mathcal{T}_S$ (where S is a σ -symbol) for the V -terms and S -terms, respectively, that belong to this set.

Definition 5.1 (Vertices of $\Omega(\mathbf{B})$). Let $\mathbf{B}$ be a τ -structure. We define $\Omega(\mathbf{B})$ to be the set of all tuples

$$U \in \prod_{t \in \mathcal{T}_V} \mathcal{P}(\text{hom}(\Gamma(\mathbf{T}(t)), \mathbf{B}))$$

such that U_{vertex} is a singleton set, i.e., vertices of $\Omega(\mathbf{B})$ are tuples U indexed by V -terms in $\mathcal{T}$, where the t -th entry is a set of homomorphisms from $\Gamma(\mathbf{T}(t))$ to $\mathbf{B}$ such that U_{vertex} contains exactly one homomorphism.

Let us remark on the above definition. Observe that the domain of $\mathbf{T}(\text{vertex})$ is $\{r_{\text{vertex}}\}$, hence $U_{\text{vertex}} = \{r_{\text{vertex}} \mapsto u\}$ for some $u \in B$, i.e., U_{vertex} selects an element u of B . Also note that the relations in $\Gamma(\mathbf{T}(t))$ will be empty whenever none of the trees $\mathbf{Q}_R$ map to $\mathbf{T}(t)$, in which case U_t is a set of functions from $\mathbf{T}(t)$ to B .

Intuitively, elements of $\Omega(B)$ are elements u of B together with some extra information. The additional information stores data about mappings from the subtrees of each Q_R to B and how they interact.

To define edges, we need to introduce some notation. Let S be a relational symbol of arity k , $t_1, \dots, t_k \in \mathcal{T}_V$, and $U^1, \dots, U^k \in \Omega(B)$. Consider the term $t = \text{edge}_S(t_1, \dots, t_k)$ and the tree $\mathbf{T}(t)$. Since its domain $T(t)$ is the disjoint union of $T(t_1), \dots, T(t_k)$, we may associate to every tuple $(f_1, \dots, f_k) \in U_{t_1}^1 \times \dots \times U_{t_k}^k$ a unique mapping $f: T(t) \rightarrow B$ defined as $f = f_1 \cup \dots \cup f_k$. We define $\times_t(U^1, \dots, U^k)$ as the set of all such mappings, i.e., we let

$$\times_t(U^1, \dots, U^k) = \{f_1 \cup \dots \cup f_k \mid f_1 \in U_{t_1}^1, \dots, f_k \in U_{t_k}^k\}.$$

Note that $\times_t(U^1, \dots, U^k)$ is bijective to $U_{t_1}^1 \times \dots \times U_{t_k}^k$. Often, we will write simply $\times_t(E)$ instead of $\times_t(U^1, \dots, U^k)$ if $E = (U^1, \dots, U^k)$.

Definition 5.2 (Edges of $\Omega(\mathbf{B})$). Let $\mathbf{B}$ be a τ -structure, we define a σ -structure $\Omega(\mathbf{B})$ with universe $\Omega(B)$. For each σ -symbol S of arity k , let $S^{\Omega(\mathbf{B})}$ consist of all tuples $(U^1, \dots, U^k) \in \Omega(B)^k$ that satisfy:

(A1) For all $t \in \mathcal{T}_S$ and $i \in \{1, \dots, k\}$ such that $\text{pr}_i(t) \in \mathcal{T}_V$, we have

$$\times_t(U^1, \dots, U^k) \subseteq U_{\text{pr}_i(t)}^i.$$

(A2) For all $t \in \mathcal{T}_S$,

$$\times_t(U^1, \dots, U^k) \subseteq \text{hom}(\Gamma(\mathbf{T}(t)), \mathbf{B}).$$

Let us comment on the definition. In (A1), the trees represented by t and $\text{pr}_i(t)$ coincide, and hence both sets in the condition consist of mappings with the same domain. Item (A2) requires that the mappings that we defined by taking union of homomorphisms $\mathbf{T}(t_i) \rightarrow U^i$, where $t = \text{edge}_R(t_1, \dots, t_k)$, are actually homomorphisms. Consequently, given that E is an edge in $S^{\Omega(\mathbf{B})}$, we may thus define a tuple

$$E^* \in \prod_{t \in \mathcal{T}_S} \mathcal{P}(\text{hom}(\Gamma(\mathbf{T}(t)), \mathbf{B}))$$

with $E_t^* = \times_t(E)$. The intuition behind requiring condition (A2) is that the union of homomorphisms is a homomorphism as long as the root edge of t is mapped to an edge. Observe that condition (A2) does not need to be checked if $\text{pr}_i(t) \in \mathcal{T}_V$ for some i since then it follows from (A1). Nevertheless, it will be useful to talk about it separately.

We claim that this construction Ω yields a right adjoint to the considered cases of central Pultr functors Γ .

Theorem 5.3. Assume a (σ, τ) -Pultr template with $\mathbf{P}$ being the σ -structure with a single vertex and empty relations, and Q_R being a σ -tree for all τ -symbols R . Further, assume Γ is the central Pultr functor defined by this template, and Ω is defined as in Definition 5.2.

For every σ -structure $\mathbf{A}$ and τ -structure $\mathbf{B}$, there is a homomorphism $\Gamma(\mathbf{A}) \rightarrow \mathbf{B}$ if and only if there is a homomorphism $\mathbf{A} \rightarrow \Omega(\mathbf{B})$.

We now proceed to prove the above theorem in several steps. The following lemma proves one of the implications and gives further insights to why U 's and $\times_t(E)$'s are defined as above.

Lemma 5.4. If there is a homomorphism $f: \Gamma(\mathbf{A}) \rightarrow \mathbf{B}$, then there is a homomorphism $g: \mathbf{A} \rightarrow \Omega(\mathbf{B})$.

Proof. We define a mapping $g: \mathbf{A} \rightarrow \Omega(B)$ by

$$g(u)_t = \{f \circ h \mid h: \mathbf{T}(t) \rightarrow \mathbf{A}, h(r_t) = u\}.$$

We claim that this mapping is a homomorphism. First, we show that $g(u)$ is well-defined, i.e., that $g(u)_{\text{vertex}}$ is a singleton set, and the elements of $g(u)_t$ are homomorphisms from $\Gamma(\mathbf{T}(t))$ to $\mathbf{B}$. For the former, observe that $g(u)_{\text{vertex}} = \{r_{\text{vertex}} \mapsto f(u)\}$, since there is a single homomorphism $h: \mathbf{T}(\text{vertex}) \rightarrow \mathbf{A}$ which maps the root (and the

only vertex) r_{vertex} to u . For the latter, assume t is a V -term. Since $h: T(t) \rightarrow A$ is also a homomorphism $\Gamma(T(t)) \rightarrow \Gamma(A)$, and f is a homomorphism $\Gamma(A) \rightarrow B$, we get that $f \circ h: \Gamma(T(t)) \rightarrow B$ is homomorphism by composition.

To prove that g preserves the relations, assume that $e = (u_1, \dots, u_k) \in S^A$ is an edge. We first show that

$$\bigtimes_t (g^S(e)) = \{f \circ h \mid h: T(t) \rightarrow A, h^S(r_t) = e\}$$

for each $t \in \mathcal{T}_S$, $t = \text{edge}_S(t_1, \dots, t_k)$. This is true because, for any homomorphism $h: T(t) \rightarrow A$ that maps the root edge to $e = (u_1, \dots, u_k)$, its restriction $h_i: T(t_i) \rightarrow A$ maps the roots to the respective u_i for all i , and, for any tuple of homomorphisms $h_i: T(t_i) \rightarrow A$ which maps the roots to the respective u_i 's, their union is a homomorphism $T(t) \rightarrow A$.

To prove property (A1), we need to check that $\bigtimes_t (g^S(e)) \subseteq g(u_i)_{\text{pr}_i(t)}$. This is easy to see, since any homomorphism $h: T(t) \rightarrow A$ that maps r_t to e maps $r_{\text{pr}_i(t)}$, which is the i -th component of r_t , to u_i . Finally, the property (A2), that each mapping in $\bigtimes_t (g^S(e))$ is a homomorphism, is proved in the same way as the analogous statement for $g(u)$, i.e., it follows from the above claim. $\square$

The above lemma concludes one of the implications that we need for the adjunction. We turn to the other implication which we prove in two steps, each provided by one of the following two lemmas. The first lemma proves the adjunction in the special case when $A = T(t)$ for some $t \in \mathcal{T}$. This will be used in the proof of the general case.

Lemma 5.5. *Let $t \in \mathcal{T}$ and let $h: T(t) \rightarrow \Omega(B)$ be a homomorphism. Then the mapping $d: T(t) \rightarrow B$, defined so that $d(v) = f(r_{\text{vertex}})$ where f is the unique element of $h(v)_{\text{vertex}}$, is a homomorphism $\Gamma(T(t)) \rightarrow B$.*

Proof. We prove by induction on t that $d \in h(r_t)_t$ if t is a V -term, and $d \in \bigtimes_t (h^S(r_t))$ if t is an S -term.

Case $t = \text{vertex}$: This is a trivial case.

Case $t = \text{edge}_S(t_1, \dots, t_k)$: Note that restrictions of h to subtrees $T(t_i)$ are homomorphisms, so we know that, for all i , $h(r_{t_i})_{t_i}$ contains the restrictions of d by the inductive assumption. The claim then immediately follows from the definition of $\bigtimes_t (h^S(r_t))$.

Case $t = \text{pr}_i(t')$: Since h is a homomorphism from $T(t) = T(t')$ to B , we know that $d \in \bigtimes_{t'} (h^S(r_{t'}))$, and the claim subsequently follows by (A1).

The lemma then immediately follows either by the definition, if t is a V -term, or by (A2), otherwise. $\square$

Note that the induction in the above proof alternates between V -terms and S -terms. We use the above lemma to prove the general case.

Lemma 5.6. *If $g: A \rightarrow \Omega(B)$ is a homomorphism, then there is a homomorphism $f: \Gamma(A) \rightarrow B$.*

Proof. Recall that, for each τ -symbol R , t_R is a fixed term representing Q_R , and hence $T(t_R)$ and Q_R are isomorphic. We further assume (without loss of generality) that $Q_R = T(t_R)$. We define f by setting $f(u)$ to be the unique value attained by the single map in $g(u)_{\text{vertex}}$. This is a well-defined mapping on the vertices of $\Gamma(A)$. We need to show that it preserves the relations of $\Gamma(A)$. To this end, assume that R is a τ -symbol of arity k and $(u_1, \dots, u_k) \in R^{\Gamma(A)}$. This means that there is a homomorphism $h: Q_R \rightarrow A$, s.t., $h(x_i) = u_i$ for all $i \in [k]$. Observe that $g \circ h: T(t_R) \rightarrow \Omega(B)$ is a homomorphism since it is obtained as a composition of two homomorphisms, so Lemma 5.5 applies to $g \circ h$ in place of h and $f \circ h$ in place of d (since $fh(v)$ is the unique value attained by the single map in $gh(v)_{\text{vertex}}$). Consequently, $f \circ h$ is a homomorphism from $\Gamma(T(t_R))$ to B , and therefore

$$(f(u_1), \dots, f(u_k)) = (fh(x_1), \dots, fh(x_k)) \in R^B,$$

since $h(x_i) = u_i$ and $(x_1, \dots, x_k) \in R^{\Gamma(T(t_R))}$, where the latter follows from the definition of $R^{\Gamma(T(t_R))}$ as phrased in (1) by letting r be the identity homomorphism on $T(t_R)$. $\square$

Lemmas 5.4 and 5.6 together yield Theorem 5.3.


Remark 5.7. Let us briefly compare our general construction to that of Foniok and Tardif [FT15, Theorem 7.1] assuming that both signatures are digraphs and $\mathbf{P} = \mathbf{V}_1$ (denoted by $\vec{P}_0$ in [FT15]). An analogous comparison also applies for the construction presented in the next section. Also, examples below compare the two constructions in several concrete cases of central Pultr functors.

First, both the domain of our right adjoint Ω and Foniok-Tardif right adjoint Ω_{FT} consist of some tuples indexed by some rooted subtrees of $\mathbf{Q}_E$. The first difference is in how we choose the subtrees. Foniok and Tardif choose a vertex of $\mathbf{Q}_E$ called *middle vertex* which disconnects the two copies of $\mathbf{P}$ in $\mathbf{Q}_E$, and define the set of subtrees using this vertex. This is similar to choosing a term representing $\mathbf{Q}_E$ rooted in the middle vertex, although not every subtree used by Foniok and Tardif needs to be represented by a subterm. Furthermore, note that the requirement that the middle vertex separates the two copies of $\mathbf{P}$ creates an obstacle to generalising Foniok and Tardif's construction to relational structures with higher arity.

Second, the elements of tuples comprising the universe of Ω_{FT} are subsets of vertices of $\mathbf{B}$ while elements of our tuples are sets of mappings to B . The main reason that Foniok and Tardif are able to do this is that the middle vertex separates the two copies of $\mathbf{P}$, and hence they only need to track the value of the mappings on one distinguished vertex. Again, since our aim is to provide adjoints in a more general setting, we cannot afford to do that.

5.1. Example: An oriented path

In this example, we compare our construction to the construction introduced in [FT15, Definition 4.1]. Our goal is to construct the adjoint to the digraph Pultr functor Γ defined by the Pultr template where $\mathbf{Q}_E$ is the following digraph:



The maps $\epsilon_{1,E}$ and $\epsilon_{2,E}$ map the singleton $\mathbf{P}$ to x_1 and x_2 , respectively.

Let us start by fixing a term t_E representing $\mathbf{Q}_E$. Namely, we let

$$t_E = \text{edge}_E(\text{pr}_1(\text{edge}_E(\text{vertex}, \text{vertex})), \text{pr}_1(\text{edge}_E(\text{vertex}, \text{vertex})))$$

which represents $\mathbf{Q}_E$ rooted in the middle edge. It has two V -subterms and two E -subterms (including itself) that represent the following trees:

$$\begin{array}{ll} s_0 = \text{vertex} & \bullet \\ s_1 = \text{pr}_1(t_1) & \bullet \longrightarrow \circ \end{array} \quad \begin{array}{ll} t_1 = \text{edge}_E(s_0, s_0) & \circ \xrightarrow{\text{red}} \circ \\ t_E = \text{edge}_E(s_1, s_1) & \circ \xleftarrow{\text{red}} \circ \xrightarrow{\text{red}} \circ \end{array}$$

For a directed graph $\mathbf{H}$, the definition of $\Omega(\mathbf{H})$ is spelled out as follows. The vertices of $\Omega(\mathbf{H})$ are pairs $U = (U_{s_0}, U_{s_1})$, where U_{s_0} is the set containing the map that sends the unique vertex of $T(\text{vertex})$ to some $u_0 \in H$ and $U_{s_1} \subseteq H^{T(s_1)}$; this is because $\Gamma(T(s_1))$ has no edges since $\mathbf{Q}_E \not\rightarrow T(s_1)$. There is an edge from $U = (U_{s_0}, U_{s_1})$ to $V = (V_{s_0}, V_{s_1})$ if

- (1) $\times_{t_1}(U, V) \subseteq U_{s_1}$, and
- (2) $\times_{t_E}(U, V) \subseteq \text{hom}(\Gamma(\mathbf{Q}_E), \mathbf{H})$

(the remaining condition (A2) for $t = t_1$ is trivial). Let us simplify this definition. First, we will write homomorphisms from the above paths as tuples, writing the values of such homomorphisms from left to right as the vertices appear on the picture above. In this way, we have $U_{s_0} \subseteq H$, $U_{s_1} \subseteq H \times H$ for each U . Since $t_1 = \text{edge}_E(s_0, s_0)$, we get a bijection $\times_{t_1} \simeq U_{s_0} \times V_{s_0}$. Furthermore, using the above ordering, we may simply say that $\times_{t_1} = U_{s_0} \times V_{s_0}$. Similarly, $\times_{t_E}(U, V) \simeq U_{s_1} \times V_{s_1}$, since $t_E = \text{edge}_E(s_1, s_1)$, more precisely

$$\times_{t_E}(U, V) = \{(u_1, u_0, v_0, v_1) \mid (u_0, u_1) \in U_{s_1} \text{ and } (v_0, v_1) \in V_{s_1}\}.$$

The two conditions are then rephrased as follows:

- (1) $U_{s_0} \times V_{s_0} \subseteq U_{s_1}$, and
- (2) $(u_1, v_1) \in E^H$ for every $(u_0, u_1) \in U_{s_1}$ and $(v_0, v_1) \in V_{s_1}$.

We claim that this construction results (on the same input) in a digraph that is homomorphically equivalent to the one obtained by [FT15, Definition 4.1]. Namely, the adjoint constructed there, let us call it Ω' , is as follows: The vertices of $\Omega'(\mathbf{H})$ are pairs (a, A) , where $a \in H$ and $A \subseteq H$, and there is an edge from (a, A) to (b, B) if $b \in A$ and $A \times B \subseteq E^H$.

We show that, for every graph $\mathbf{H}$, there is a homomorphism $\alpha: \Omega'(\mathbf{H}) \rightarrow \Omega(\mathbf{H})$ defined by $\alpha((a, A)) = U$ where $U_{s_0} = \{a\}$ and $U_{s_1} = \{a\} \times A$. To show that α preserves edges, assume that (a, A) and (b, B) are connected by an edge in $\Omega'(\mathbf{H})$, i.e., $b \in A$ and $A \times B \subseteq E^H$, and $U = \alpha((a, A))$, $V = \alpha((b, B))$. We have that

$$\begin{aligned} \times_{t_1}(U, V) &= \{(a, b)\} \\ \times_{t_E}(U, V) &= A \times \{(a, b)\} \times B, \end{aligned}$$

and claim that: (1) $\times_{t_1}(U, V) \subseteq U_{s_1}$; which is true because $b \in A$. (2) $\times_{t_E}(U, V) \subseteq \text{hom}(\Gamma(\mathbf{Q}_E), \mathbf{H})$; which is true, since the only edge of $\Gamma(\mathbf{Q}_E)$ is (x_1, x_2) and the projection of $\times_{t_E}(U, V)$ to x_1, x_2 (the first and the last coordinate) is $A \times B \subseteq E^H$.

A homomorphism $\beta: \Omega(\mathbf{H}) \rightarrow \Omega'(\mathbf{H})$ is given by $\beta(U) = (a, A)$ where a is the unique element of U_{s_0} , and

$$A = \{a' \mid (a, a') \in U_{s_1}\}.$$

To show that it is a homomorphism, assume $(U, V) \in E^{\Omega(\mathbf{H})}$, and let $\beta(U) = (a, A)$ and $\beta(V) = (b, B)$. Since $\times_{t_1}(U, V) = U_{s_0} \times V_{s_0} = \{(a, b)\}$ and $\times_{t_1}(U, V) \subseteq U_{s_1}$, we have $(a, b) \in U_{s_1}$, which implies that $b \in A$. Also since $\times_{t_E}(U, V) \subseteq \text{hom}(\Gamma(\mathbf{Q}_E), \mathbf{H})$, $a \times A \subseteq U_{s_1}$, $b \times B \subseteq V_{s_1}$, and $t_E = \text{edge}_E(s_1, s_1)$, we have that

$$A \times \{(a, b)\} \times B \subseteq \times_{t_E}(U, V) \subseteq \text{hom}(\Gamma(\mathbf{Q}_E), \mathbf{H}),$$

where the first inclusion follows from the definition of $\times_{t_E}(U, V)$, and the second inclusion follows from (A2). In particular, the inclusion above implies that $A \times B \subseteq E^H$. This concludes the proof of the homomorphic equivalence of $\Omega(\mathbf{H})$ and $\Omega'(\mathbf{H})$.

We note that our $\Omega(\mathbf{H})$ can be reduced to a smaller homomorphically equivalent structure by requiring that vertices $U \in \Omega(\mathbf{H})$ satisfy

- (A3) for all $s, s' \in \mathcal{T}_V$ and homomorphisms $h: \mathbf{T}(s) \rightarrow \mathbf{T}(s')$ with $h(r_s) = r_{s'}$, we have


$$\{f \circ h \mid f \in U_{s'}\} \subseteq U_s.$$

Note that the elements constructed in the proof of Lemma 5.4 satisfy this property. In this particular example, this requirement would force that $U_{s_1} = a \times A$ for some $A \subseteq H$ and $a \in U_{s_0}$, since s_0 is embedded to s_1 as the root. This would then make the two homomorphisms defined above isomorphisms.

5.2. Example: A 4-ary relation defined by an oriented path

Our definition works also for Pultr templates that are not just digraph templates. As an example for comparison, let us consider a Pultr template that is similar to the previous example, but in this case maps digraphs to structures over a signature containing one 4-ary relational symbol R .

Specifically, $\mathbf{P}$ is still a singleton with no edges, and $\mathbf{Q}_R$ is the same digraph as $\mathbf{Q}_E$ above, but we now have 4 homomorphisms $\epsilon_{i,R}: \mathbf{P} \rightarrow \mathbf{Q}_R$ for $i = 0, 1, 2, 3$ which map the vertex of $\mathbf{P}$ to 0, 1, 2, or 3 respectively. Pictorially, the digraph $\mathbf{Q}_R$ together with its distinguished vertices $x_0, \dots, x_3$ is as follows.



We let $t_R = t_E$ where t_E is as above, and we use the same notation as in the previous example.

For a structure $\mathbf{B}$ with a 4-ary relation $R^{\mathbf{B}}$, vertices of $\Omega(\mathbf{B})$ are defined in a similar way as above, i.e., they are pairs (U_{s_0}, U_{s_1}) where $U_{s_0} = \{b\}$ for some $b \in B$ and $U_{s_1} \subseteq B \times B$. Two such vertices (U, V) are connected by an edge if

- (1) $\times_{t_R}(U, V) \subseteq R^{\mathbf{B}}$, and
- (2) $\times_{t_1}(U, V) \subseteq U_{s_1}$;

where $t_R = \text{edge}_E(s_1, s_1)$ and $t_1 = \text{edge}_E(s_0, s_0)$. More precisely, the first condition can be written as

$$\{(u_1, u_0, v_0, v_1) \mid (u_0, u_1) \in U_{s_1}, (v_0, v_1) \in V_{s_1}\} \subseteq R^{\mathbf{B}},$$

and the second condition can be written as $U_{s_0} \times V_{s_0} \subseteq U_{s_1}$.

The only real difference from the previous example is that $\times_{t_R}(U, V) \subseteq R^{\mathbf{B}}$ instead of requiring that the projection of $\times_{t_E}(U, V)$ on the first and the last coordinates is a subset of $E^{\mathbf{H}}$. As above, we could create a homomorphically equivalent $\Omega'(\mathbf{B})$ whose vertices would be pairs (u, U) where $u \in B$ and $U \subseteq B$. Two such vertices (u, U) and (v, V) would be then connected by an edge if $v \in U$ and $U \times \{(u, v)\} \times V \subseteq R^{\mathbf{B}}$.

5.3. Duals from adjoints

In [FT15], the authors claim that, for digraph Pultr templates with the edge relation defined by $\mathbf{Q}_E$, the image of the digraph with a single vertex and no edges under the right Pultr functor is a dual to $\mathbf{Q}_E$. We may be slightly more precise when talking about our constructions, namely, we claim the following.

Proposition 5.8. *Let τ contain a single symbol R , and fix a (σ, τ) -Pultr template defined by $\mathbf{P} = \mathbf{V}_1$ and $\mathbf{Q}_R = \mathbf{T}(t_R)$ for some term t_R such that $\mathbf{Q}_R \not\rightarrow \mathbf{T}(t)$ for any $t < t_R$ (i.e., t_R is a minimal term among those representing a structure homomorphically equivalent to $\mathbf{Q}_R$). Let Ω be the right adjoint to Γ as defined in Definition 5.2. Then the image of the τ -structure $\mathbf{V}_1$ under Ω is isomorphic to $\mathbf{D}(t_R)$.*

Proof. Observe that there is only one function from any set T to $V_1 = \{1\}$. Hence, there is only one candidate function f for a homomorphism $\Gamma(\mathbf{T}(t)) \rightarrow \mathbf{V}_1$ for any term t . It is not hard to observe that f is a homomorphism if and only if $\Gamma(\mathbf{T}(t))$ has no edges. In particular, f is a homomorphism for all $t < t_r$, since $\mathbf{Q}_R \not\rightarrow \mathbf{T}(t)$ in that case. Consequently, each of the components of an element $U \in \Omega(V_1)$ is either the empty set or a singleton set. Furthermore, we have that $\text{hom}(\Gamma(\mathbf{T}(t_R)), \mathbf{V}_1) = \emptyset$, since $\Gamma(\mathbf{T}(t_R))$ has an edge (witnessed by the identity homomorphism).

The rest of the proof is based on the idea of treating $\emptyset$ as false and any singleton set as true. With this interpretation, we get that $X \Rightarrow Y$ is equivalent to $X \subseteq Y$ and $X \wedge Y$ is equivalent to $X \times Y$ whenever X, Y are sets with at most one element. Note that this draws an immediate parallel between the definitions of $\&_t$ and $\times_t$, (D1) and (A1), and also (D2) and (A2) if we take into account that $\text{hom}(\Gamma(\mathbf{T}(t_R)), \mathbf{V}_1) = \emptyset$.

We may now define an isomorphism $\Omega(\mathbf{V}_1) \simeq \mathbf{D}(t_R)$ by assigning to $u \in \mathbf{D}(t_R)$ an element $U \in \Omega(\mathbf{V}_1)$ where, for each $t \in \mathcal{T}_V$, $u_t = \text{false}$ if $U_t = \emptyset$, and $u_t = \text{true}$ if $U_t = \{x \mapsto 1\}$. The observations in the first paragraph of this proof show that this assignment is bijective.

To show that it preserves edges, observe that, for each $t \in \mathcal{T}_R$, $\&_t(u^1, \dots, u^k) = \text{false}$ if and only if $\times_t(U^1, \dots, U^k) = \emptyset$, since the former is defined as $\&_t(u^1, \dots, u^k) = u_{t_1}^1 \wedge \dots \wedge u_{t_k}^k$ and the latter satisfies $\times_t(U_1, \dots, U_k) \simeq U_{t_1}^1 \times \dots \times U_{t_k}^k$. Furthermore, (A1) is equivalent to (D1) since, for all $t \in \mathcal{T}_R$ with $\text{pr}_i(t) \in \mathcal{T}_V$, $\&_t(u^1, \dots, u^k) \Rightarrow u_{\text{pr}_i(t)}^i$ is equivalent to $\times_t(U^1, \dots, U^k) \subseteq U_{\text{pr}_i(t)}^i$. Finally, (A2) $\&_{t_Q}(u^1, \dots, u^k) = \text{false}$ is equivalent to

$$(D2) \quad \text{for all } t \in \mathcal{T}_R, \quad \times_t(U^1, \dots, U^k) \subseteq \text{hom}(\Gamma(\mathbf{T}(t)), \mathbf{B}).$$

This is because $\text{hom}(\Gamma(\mathbf{T}(t)), \mathbf{B}) = \emptyset$ if and only if $t = t_Q$, hence (D2) is trivially satisfied unless $t = t_Q$, in which case it says $\times_{t_Q}(U^1, \dots, U^k) = \emptyset$. $\square$

6. ADJOINTS TO FUNCTORS WITH DOMAINS DEFINED BY A RELATION

In this section, we construct a right adjoint to the central Pultr functor defined by a (σ, τ) -Pultr template for which $\mathbf{P} = \check{S}_1$, for some choice of σ -symbol $\check{S}$, i.e., $\mathbf{P}$ is an $\check{S}$ -edge. Again, for each k -ary τ -symbol R , we have a σ -tree $\mathbf{Q}_R$ with a k -tuple $\epsilon_{1,R}, \dots, \epsilon_{k,R}$ of homomorphisms $\mathbf{P} \rightarrow \mathbf{Q}_R$. Each of these homomorphisms selects an $\check{S}$ -edge of $\mathbf{Q}_R$, we denote these edges by $x_1, \dots, x_k$, respectively, i.e., $x_i = \epsilon_{i,R}^{\check{S}}(e)$ for each $i = 1, \dots, k$ where $e \in \check{S}^{\mathbf{P}}$ is the unique $\check{S}$ -edge of $\mathbf{P}$. Naturally, the edges $x_1, \dots, x_{\text{ar } R}$ depend on the symbol R which will be always clear from the context. Let us repeat the definition of $\Gamma(\mathbf{A})$ for a σ -structure $\mathbf{A}$ using this notation: The domain of $\Gamma(\mathbf{A})$ is $\check{S}^{\mathbf{A}}$, and for each τ -symbol R of arity k , the corresponding relation of $\Gamma(\mathbf{A})$ is defined as

$$R^{\Gamma(\mathbf{A})} = \{(h^{\check{S}}(x_1), \dots, h^{\check{S}}(x_k)) \mid h: \mathbf{Q}_R \rightarrow \mathbf{A}\}.$$

We also note that, for each homomorphism $h: \mathbf{A} \rightarrow \mathbf{B}$ between two σ -structures $\mathbf{A}$ and $\mathbf{B}$, $h^{\check{S}}: \Gamma(\mathbf{A}) \rightarrow \Gamma(\mathbf{B})$ is a homomorphism.

The construction of Ω in this case is almost identical to the construction in Section 5.3 (Definitions 5.1 and 5.2), and we present it in a similar way.

We fix the following setting: Fix a (σ, τ) -Pultr template with $\mathbf{P} = \check{S}_1$ for some σ -symbol $\check{S}$, and assume $\mathbf{Q}_R$ is a σ -tree for each τ -symbol R . For each τ -symbol R , we pick a term t_R representing $\mathbf{Q}_R$, and we let $\mathcal{T}$ be the set of all subterms of any of the t_R 's. We use notation $\mathcal{T}_V$ and $\mathcal{T}_S$ for the V -terms and S -terms, where S is a σ -symbol, that belong to this set.

Definition 6.1 (Vertices of $\Omega(\mathbf{B})$). Let $\mathbf{B}$ be a τ -structure. We define $\Omega(\mathbf{B})$ to be the set of all tuples

$$U \in \prod_{t \in \mathcal{T}_V} \mathcal{P}(\text{hom}(\Gamma(\mathbf{T}(t)), \mathbf{B}))$$

such that U_{vertex} is a singleton set, i.e., $U_{\text{vertex}} = \{\emptyset\}$.

As before, vertices of $\Omega(\mathbf{B})$ are tuples U indexed by V -terms in $\mathcal{T}$, where the t -th entry is a set of homomorphisms from $\Gamma(\mathbf{T}(t))$ to $\mathbf{B}$ such that U_{vertex} contains exactly one homomorphism. A difference here is that this homomorphism is the map $\emptyset \rightarrow B$, since $\Gamma(\mathbf{T}(\text{vertex}))$ has no vertices as $\mathbf{P} \not\rightarrow \mathbf{T}(\text{vertex})$. Thus U_{vertex} does not contain any information; it serves a similar purpose as $u_{\text{vertex}} = \text{true}$ in Definition 4.1.

To define edges, we use similar notation $\times_t$ as before. Let $t_1, \dots, t_k \in \mathcal{T}_V$, S a relational symbol of arity k , and $U^1, \dots, U^k \in \Omega(\mathbf{B})$. Again, we consider the term $t = \text{edge}_S(t_1, \dots, t_k)$ and the tree $\mathbf{T}(t)$. Unlike in Section 5, the domain of $\Gamma(\mathbf{T}(t))$ is now $\check{S}^{\mathbf{T}(t)}$ and not $T(t)$. This means we have to distinguish two cases:

Case 1: $S = \check{S}$. In this case, the domain of $\Gamma(\mathbf{T}(t))$ is

$$\check{S}^{\mathbf{T}(t)} = \check{S}^{\mathbf{T}(t_1)} \cup \dots \cup \check{S}^{\mathbf{T}(t_k)} \cup \{r_t\}.$$

Hence, in order to define a mapping $f: \check{S}^{\mathbf{T}(t)} \rightarrow B$, we need to specify its value on r_t in addition to its restrictions to $\check{S}^{\mathbf{T}(t_i)}$'s. In detail, a k -tuple of mappings $f_1 \in U_{t_1}^1, \dots, f_k \in U_{t_k}^k$ together with an element $e_\bullet \in B$ uniquely defines a mapping $f: \check{S}^{\mathbf{T}(t)} \rightarrow B$ by

$$f = f_1 \cup \dots \cup f_k \cup \{r_t \mapsto e_\bullet\}.$$

Therefore, for each $e_\bullet \in B$, we denote by $\times_t(U^1, \dots, U_k; e_\bullet)$ the set of all such mappings. Again, we have that, for each $e_\bullet$, $\times_t(U^1, \dots, U_k; e_\bullet)$ is bijective to $U^1 \times \dots \times U^k$. We will simply write $\times_t(E; e_\bullet)$ for $\times_t(U^1, \dots, U^k; e_\bullet)$ if $E = (U^1, \dots, U^k)$.

Case 2: $S \neq \check{S}$. In this case, we have $\check{S}^{\mathbf{T}(t)} = \check{S}^{\mathbf{T}(t_1)} \cup \dots \cup \check{S}^{\mathbf{T}(t_k)}$. Thus the domain of $\Gamma(\mathbf{T}(t))$ is the disjoint union of domains of $\Gamma(\mathbf{T}(t_i))$, and we define $\times_t(U^1, \dots, U^k)$ the same way as before, i.e., as the set of all unions $f_1 \cup \dots \cup f_k$ where $f_i \in U_{t_i}^i$ for each i .

If $S \neq \check{S}$, S -edges are defined analogously to Definition 5.2. To define $\check{S}$ -edges, we have to take into account the new element $e_\bullet$.

Definition 6.2 (Edges of $\Omega(\mathbf{B})$). Let $\mathbf{B}$ be a τ -structure, we define a σ -structure $\Omega(\mathbf{B})$ with universe $\Omega(\mathbf{B})$.

We first define $\check{S}$ -edges. Let k be the arity of $\check{S}$. $\check{S}^{\Omega(\mathbf{B})}$ consists of all tuples $(U^1, \dots, U^k) \in \Omega(\mathbf{B})^k$ for which there exists $e_\bullet \in B$ such that

(B1) For all $t \in \mathcal{T}_S$ and $i \in \{1, \dots, k\}$ such that $\text{pr}_i(t) \in \mathcal{T}_V$, we have

$$\bigtimes_t (U^1, \dots, U^k; e_\bullet) \subseteq U_{\text{pr}_i(t)}^i.$$

(B2) For all $t \in \mathcal{T}_S$,

$$\bigtimes_t (U^1, \dots, U^k; e_\bullet) \subseteq \text{hom}(\Gamma(\mathbf{T}(t)), \mathbf{B}).$$

If $S \neq \check{S}$, the relation $S^{\Omega(\mathbf{B})}$ is defined in the same way as in Definition 5.2 except we use the meaning of $\bigtimes_t$ defined in this section.

Given an edge $E \in \check{S}^{\Omega(\mathbf{B})}$, we call the element $e_\bullet$, which satisfies the conditions (B1) and (B2), a *witness* of this edge. Note that this witness is the only significant difference between Definitions 5.2 and 6.2.

We claim that this definition indeed constructs a right adjoint to Γ .

Theorem 6.3. Assume a (σ, τ) -Pultr template with $\mathbf{P}$ being the σ -tree with $\text{ar } \check{S}$ vertices connected by an $\check{S}$ -edge for some σ -symbol $\check{S}$, and $\mathbf{Q}_R$ being a σ -tree for all τ -symbols R . Further, assume Γ is the central Pultr functor defined by this template, and Ω is defined as in Definition 6.2.

For every σ -structure $\mathbf{A}$ and τ -structure $\mathbf{B}$, there is a homomorphism $\Gamma(\mathbf{A}) \rightarrow \mathbf{B}$ if and only if there is a homomorphism $\mathbf{A} \rightarrow \Omega(\mathbf{B})$.

The proof is analogous to the proof of Theorem 5.3 with the following changes: we use $h^{\check{S}}$ in place of h whenever $h: \mathbf{A} \rightarrow \mathbf{B}$ was used as a homomorphism $\Gamma(\mathbf{A}) \rightarrow \Gamma(\mathbf{B})$. Furthermore, we use a witness $e_\bullet$ in place of the unique value of the homomorphism $f \in U_{\text{vertex}}$ throughout the proof. In particular, if $f: \mathbf{A} \rightarrow \Omega(\mathbf{B})$, we define $g: \Gamma(\mathbf{A}) \rightarrow \mathbf{B}$ by letting $g(e)$ be the witness of the edge $f^{\check{S}}(e)$. With these substitutions all the arguments of the previous section apply in the case of Theorem 6.3. For completeness and reference, we include the proof in full detail in the last subsection of this section.

6.1. Example: The arc graph construction

Recall the arc-graph construction from Example 2.9, which can be expressed as the central Pultr functor whose template consists of structures $\mathbf{P} = (\{0, 1\}; \{(0, 1)\})$ and $\mathbf{Q}_E = (\{0, 1, 2\}; \{(0, 1), (1, 2)\})$ with $\epsilon_1(0) = 0$ and $\epsilon_1(1) = 1$, and $\epsilon_2(0) = 1$ and $\epsilon_2(1) = 2$.

The right adjoint Ω according to our definition above would be constructed in the following way: First, we choose a term t_E representing $\mathbf{Q}_E$. We can pick t_2 as in Section 4.1.

$$\circ \longrightarrow \circ \xrightarrow{\text{red}} \circ$$

$$t_2 = \text{edge}_E(\text{pr}_2(\text{edge}_E(\text{vertex}, \text{vertex})), \text{vertex})$$

We also name all its subterms as in Section 4.1, i.e.,

$$\begin{aligned} s_0 &= \text{vertex} & t_1 &= \text{edge}_E(s_0, s_0) \\ s_1 &= \text{pr}_2(t_1) \end{aligned}$$

The vertices of $\Omega(\mathbf{B})$ are defined as pairs (U_{s_0}, U_{s_1}) such that $U_{s_0} = \{\emptyset\}$ and $U_{s_1} \subseteq \text{hom}(\Gamma(\mathbf{T}(s_1)), \mathbf{B})$. Two such vertices U, V are connected by an edge if there is $e_\bullet \in B$ such that

- (1) $\bigtimes_{t_1}(U, V; e_\bullet) \subseteq U_{s_1}$ — this is condition (B1) for $t = t_1$;
- (2) $\bigtimes_{t_i}(U, V; e_\bullet) \subseteq \text{hom}(\Gamma(\mathbf{T}(t_i)), \mathbf{B})$ for $i = 1, 2$ — this is condition (B2).

These conditions can be considerably simplified. First, since $U_{s_0} = V_{s_0} = \{\emptyset\}$, we have

$$\begin{aligned} \times_{t_1}(U, V; e_\bullet) &= \{r_{t_1} \mapsto e_\bullet\} \\ \times_{t_2}(U, V; e_\bullet) &= \{f \cup (r_{t_2} \mapsto e_\bullet) \mid f \in U_{s_1}\}. \end{aligned}$$

Further, since $\Gamma(T(s_1)) = \Gamma(T(t_1))$ is the graph with a single vertex and no edges, we can identify U_{s_1} and $\times_{t_1}(U, V; e_\bullet)$ with subsets of B . Connecting this observation with the comment about U_{s_0} , we can identify U with its coordinate $U_{s_1} \subseteq B$. Finally, if $e_\bullet$ is a witness, the elements of $\times_{t_2}(U, V; e_\bullet)$ are homomorphisms from a directed edge to $\mathbf{B}$ which correspond to the edges of $\mathbf{B}$. Using this correspondence, we can identify $\times_{t_2}(U, V; e_\bullet)$ with a subset of $E^{\mathbf{B}}$. Taking all of these into account the conditions above simplify to

- (1) $\times_{t_1}(U, V; e_\bullet) = \{e_\bullet\} \subseteq V$, and
- (2) $\times_{t_2}(U, V; e_\bullet) = U \times \{e_\bullet\} \subseteq E^{\mathbf{B}}$.

So, we can say that a vertex of $\Omega(\mathbf{B})$ is a subset U of B , and (U, V) is an edge of $\Omega(\mathbf{B})$ if

$$\exists e_\bullet \in V \text{ such that } U \times \{e_\bullet\} \subseteq E^{\mathbf{B}}.$$

We compare this construction with the functor δ_R described in [FT15, Definition 3.1] as a right adjoint to δ . For a digraph $\mathbf{B}$, the vertices of the digraph $\delta_R(\mathbf{B})$ are the complete bipartite subgraphs of $\mathbf{B}$, i.e., pairs (U^-, U^+) of subsets of vertices of $\mathbf{B}$ such that $U^- \times U^+ \subseteq E^{\mathbf{B}}$. There is an edge from (U^-, U^+) to (V^-, V^+) if $U^+ \cap V^- \neq \emptyset$. Below, we show that $\delta_R(\mathbf{B})$ and $\Omega(\mathbf{B})$ are homomorphically equivalent.

We start by constructing a homomorphism $h: \delta_R(\mathbf{B}) \rightarrow \Omega(\mathbf{B})$. We let

$$h(U^-, U^+) = U^-.$$

To show that it preserves edges, assume $U^+ \cap V^- \neq \emptyset$, i.e., there exists $e_\bullet \in U^+ \cap V^-$. We claim that this $e_\bullet$ witnesses that U^- and V^- is an edge in $\Omega(\mathbf{B})$. Clearly, $e_\bullet \in V^-$. Also, we have

$$U^- \times \{e_\bullet\} \subseteq U^- \times U^+ \subseteq E^{\mathbf{B}}.$$

A homomorphism $g: \Omega(\mathbf{B}) \rightarrow \delta_R(\mathbf{B})$ is a bit harder to construct. Guided by the above, it is natural to choose the first component of $g(U)$ to be $g(U)^- = U$. We need to define the second component $g(U)^+$. We let $g(U)^+$ be the largest set such that $g(U)^- \times g(U)^+ \subseteq E^{\mathbf{B}}$, i.e.,

$$g(U)^+ = \{v \in B \mid \forall u \in U, (u, v) \in E^{\mathbf{B}}\}.$$

Now, assume that U and V are connected by an edge in $\Omega(\mathbf{B})$ witnessed by $e_\bullet$. We claim that $e_\bullet \in g(U)^+ \cap g(V)^-$. By definition of $\Omega(\mathbf{B})$, we have $e_\bullet \in V = g(V)^-$, and $U \times \{e_\bullet\} \subseteq E^{\mathbf{B}}$, which implies that $e_\bullet \in g(U)^+$. Altogether, $e_\bullet \in g(U)^+ \cap g(V)^-$, and hence $(g(U), g(V)) \in E^{\delta_R(\mathbf{B})}$. This completes the proof.

This example shows how U_{vertex} can be eliminated from the definition of $\Omega(\mathbf{B})$. Let us repeat again, that the only purpose of U_{vertex} is to avoid case distinction between some E -terms, e.g., between $\text{edge}_E(t, s)$ and $\text{edge}_E(t, \text{vertex})$ for $s \neq \text{vertex}$. We may ignore it in this example, since we expanded every single case when it is used.

6.2. Example: Arc structure

In this subsection, we consider a certain variant of the arc graph construction, which we will call an *arc structure* and which encodes more information than the arc graph: The domain of the arc structure coincides with the domain of the arc graph, i.e., the set of all edges of the input graph, and we extend the signature with two more binary symbols that will relate those pairs of edges that are incident in a different sense. The goal of this example is two-fold: first, to show how the construction of right adjoints works in a more general signature, and second, to show how the right adjoint changes if we change the central Pultr functor in such a way that it encodes more information about the input structure. Note, for example, that any tree can be recovered from its





FIGURE 2. A digraph, its arc graph and its arc structure. The O and I relations of $\partial(G)$ are symmetric, and O and I loops on all vertices of $\partial(G)$ are omitted for readability.

image under the arc structure construction, but there are trees (e.g., the one presented in Fig. 2) that cannot be recovered from their arc graph.

We fix ϕ to be a signature with three binary relations D , I , and O , and we let γ be the signature of digraphs. We define a central Pultr functor ∂ using the (γ, ϕ) -Pultr template defined as follows: The digraph defining vertices is the digraph with a single directed edge, i.e., $\mathbf{P} = \mathbf{E}_1$, and the digraphs $\mathbf{Q}_D$, $\mathbf{Q}_I$, $\mathbf{Q}_O$ are the following where the images of ϵ_i 's are highlighted and labelled by x_i .



Note that the digraph $\mathbf{Q}_D$ with the two distinguished edges defines the arc-graph functor. This means that for each digraph G , the reduct (A, D^A) where $A = \partial(G)$ is the arc-graph of G (if D is interpreted as E), see Figure 2.

Now, to construct the right adjoint to ∂ , which we denote by ω_∂ , we fix the following terms representing the graphs $\mathbf{Q}_D$, $\mathbf{Q}_I$, and $\mathbf{Q}_O$.

$$t_D = \text{edge}_E(s_2, \text{vertex}) \quad t_I = \text{edge}_E(\text{vertex}, s_2) \quad t_O = \text{edge}_E(s_1, \text{vertex})$$

where $s_i = \text{pr}_i(\text{edge}_E(\text{vertex}, \text{vertex}))$ for $i = 1, 2$. We name all remaining subterms as follows, $s_0 = \text{vertex}$, $t_E = \text{edge}_E(\text{vertex}, \text{vertex})$. The set of terms defining ω_∂ is $\mathcal{T} = \{s_0, s_1, s_2, t_E, t_D, t_I, t_O\}$. Of which, the V -terms are $\mathcal{T}_V = \{s_0, s_1, s_2\}$. Which means that the vertices of $\omega_\partial(\mathbf{B})$ are triples $(U_{s_0}, U_{s_1}, U_{s_2})$ where $U_{s_0} = \{\emptyset\}$, and U_{s_1}, U_{s_2} are sets of functions from a 1-element set to B . We identify such a triple with a pair (U^+, U^-) where $U^+ \subseteq B$ is the set of images of function in U_{s_1} and $U^- \subseteq B$ the set of images of functions in U_{s_2} . Using a similar simplification of the definition of edges as in the previous example, we get that two such pairs (U^+, U^-) and (V^+, V^-) are connected by an edge if there exists $e_\bullet \in B$, so that the sets

$$\begin{aligned} \bigtimes_{t_E} (U, V; e_\bullet) &= \{e_\bullet\} & \bigtimes_{t_D} (U, V; e_\bullet) &= U^- \times \{e_\bullet\} \\ \bigtimes_{t_I} (U, V; e_\bullet) &= \{e_\bullet\} \times V^- & \bigtimes_{t_O} (U, V; e_\bullet) &= U^+ \times \{e_\bullet\} \end{aligned}$$

satisfy (B1) $\bigtimes_{t_E} (U, V; e_\bullet) \subseteq U^+ \cap V^-$ and (B2) $\bigtimes_{t_R} (U, V; e_\bullet) \subseteq R^B$ for each $R \in \{D, I, O\}$. This means that (U^+, U^-) and (V^+, V^-) are connected by an edge if there exists $e_\bullet \in U^+ \cap V^-$ such that

$$U^- \times \{e_\bullet\} \subseteq D^B, \{e_\bullet\} \times V^- \subseteq I^B, U^+ \times \{e_\bullet\} \subseteq O^B.$$

This completes the definition, although we can further refine it by requiring that for each vertex, the sets U^+ and U^- satisfy additional properties that would automatically imply the conditions above. Namely, we require

$$U^- \times U^+ \subseteq D^B, U^- \times U^- \subseteq I^B, U^+ \times U^+ \subseteq O^B.$$

To sum up the refined definition, we let $\omega_\partial(\mathbf{B})$ be the digraph with vertex set

$$\{(U^+, U^-) \mid U^- \times U^+ \subseteq D^B, U^- \times U^- \subseteq I^B, U^+ \times U^+ \subseteq O^B\}$$

where (U^+, U^-) and (V^+, V^-) form an edge if $U^+ \cap V^- \neq \emptyset$. It is not hard to check that even after the refinements, we still get the right adjoint to ∂ , i.e., that indeed there is a homomorphism $\mathbf{A} \rightarrow \omega_\partial(\mathbf{B})$ if and only if there is a homomorphism $\partial(\mathbf{A}) \rightarrow \mathbf{B}$ for any digraph $\mathbf{A}$ and ϕ -structure $\mathbf{B}$. Note how this compares to the right adjoint δ_R of the arc digraph functor as defined in the previous subsection.

6.3. Proof of Theorem 6.3

As noted above, the proof closely follows the proof of Theorem 5.3, and we present it with the same structure starting with the easier of the two implications.

The following lemma is an analogue to Lemma 5.4. The statement is identical, although the meaning of Γ and Ω has changed. In the proof, h is no longer a homomorphism from $\Gamma(\mathbf{T}(t)) \rightarrow \Gamma(\mathbf{A})$, hence we replace every instance of h being used as such a homomorphism with $h^{\check{S}}$. Furthermore, to show that g preserves an edge $e \in \check{S}^{\mathbf{A}}$, we need to provide a witness, which is $w_\bullet = f(e)$.

Lemma 6.4. *If there is a homomorphism $f: \Gamma(\mathbf{A}) \rightarrow \mathbf{B}$, then there is a homomorphism $g: \mathbf{A} \rightarrow \Omega(\mathbf{B})$.*

Proof. Assuming $f: \Gamma(\mathbf{A}) \rightarrow \mathbf{B}$, we define a mapping $g: \mathbf{A} \rightarrow \Omega(\mathbf{B})$ by setting, for all $u \in \mathbf{A}$ and all $t \in \mathcal{T}_V$,

$$g(u)_t = \{f \circ h^{\check{S}} \mid h: \mathbf{T}(t) \rightarrow \mathbf{A}, h(r_t) = u\},$$

and claim that it is a homomorphism from $\mathbf{A}$ to $\Omega(\mathbf{B})$. We need to check that $g(u)$ is well-defined, i.e., that $f \circ h^{\check{S}}$ is a homomorphism from $\Gamma(\mathbf{T}(t))$ to $\mathbf{B}$ for all $t \in \mathcal{T}_V$ and that $g(u)_{\text{vertex}} = \{\emptyset\}$. Indeed, $f \circ h^{\check{S}}$ is a homomorphism since it is a composition of homomorphisms $h^{\check{S}}: \Gamma(\mathbf{T}(t)) \rightarrow \Gamma(\mathbf{A})$ and $f: \Gamma(\mathbf{A}) \rightarrow \mathbf{B}$. And, $g(u)_{\text{vertex}} = \{\emptyset\}$ because $\mathbf{T}(\text{vertex})$ has no $\check{S}$ -edges, which implies that the unique mapping $h: \mathbf{T}(\text{vertex}) \rightarrow \mathbf{B}$ such that $h(r_{\text{vertex}}) = u$ is a homomorphism, and, moreover, this h satisfies $h^{\check{S}} = \emptyset$. Consequently, $f \circ h^{\check{S}} = \emptyset$, as we wanted to show.

To show that g is a homomorphism, assume first that $S \neq \check{S}$ is a σ -symbol of arity k , and $e = (u_1, \dots, u_k) \in S^{\mathbf{A}}$. We claim that, for each $t \in \mathcal{T}_S$,

$$\bigtimes_t (g(e)) = \{f \circ h^{\check{S}} \mid h: \mathbf{T}(t) \rightarrow \mathbf{A}, h^{\check{S}}(r_t) = e\}.$$

Let $t = \text{edge}_S(t_1, \dots, t_k)$, and observe that homomorphisms $h: \mathbf{T}(t) \rightarrow \mathbf{A}$ such that $h^{\check{S}}(r_t) = e$ are in 1-to-1 correspondence with k -tuples of homomorphisms $h_1, \dots, h_k$, such that $h_i: \mathbf{T}(t_i) \rightarrow \mathbf{A}$ and $h_i(r_{t_i}) = u_i$ for all $i \in [k]$, obtained as their restrictions to the respective subtrees. If h is the union of h_i 's then also $h^{\check{S}}$ is the union of $h_i^{\check{S}}$'s. The claim then easily follows. For (A1), we want to check that $\bigtimes_t (g(e)) \subseteq g(u_i)_s$ where $s = \text{pr}_i(t)$. Observe that if $h: \mathbf{T}(t) \rightarrow \mathbf{A}$ and $h^{\check{S}}(r_t) = e$, then $h(r_s) = u_i$, since the i -th component of r_t is r_s . The condition then follows from the claim. Finally, (A2) follows directly from the claim by the same argument as $g(u)_t \subseteq \text{hom}(\mathbf{T}(t), \mathbf{B})$.

Second, for the case $S = \check{S}$ and $e = (u_1, \dots, u_k) \in \check{S}^{\mathbf{A}}$, we need to pick an element $w_\bullet \in \mathbf{B}$ witnessing that $g^{\check{S}}(e) \in \check{S}^{\Omega(\mathbf{B})}$ — we pick $w_\bullet = f(e)$. We claim that, for each $t \in \mathcal{T}_{\check{S}}$,

$$\bigtimes_t (g^{\check{S}}(e); f(e)) = \{f \circ h^{\check{S}} \mid h: \mathbf{T}(t) \rightarrow \mathbf{A}, h^{\check{S}}(r_t) = e\}.$$

The inclusion ' $\supseteq$ ' is clear, since the restriction of h to $\mathbf{T}(t_i)$ is a homomorphism $h_i: \mathbf{T}(t_i) \rightarrow \mathbf{A}$ such that $h_i(r_{t_i}) = u_i$. The other inclusion also follows, since, given homomorphisms $h_i: \mathbf{T}(t_i) \rightarrow \mathbf{A}$ with $h_i(r_{t_i}) = u_i$, their union is a homomorphism $h: \mathbf{T}(t) \rightarrow \mathbf{A}$ with $h^{\check{S}}(r_t) = e$. Moreover,

$$h^{\check{S}} = h_1^{\check{S}} \cup \dots \cup h_k^{\check{S}} \cup (r_t \mapsto e),$$

and hence

$$f \circ h^{\check{S}} = (f \circ h_1^{\check{S}}) \cup \dots \cup (f \circ h_k^{\check{S}}) \cup (r_t \mapsto f(e)).$$

The conditions (B1) and (B2) then follow from the claim similarly as above. $\square$

The other implication is proved in the following two lemmas. As in the case of Lemma 5.5, the first lemma provides the adjunction in the special case $\mathbf{A} = \mathbf{T}(t)$ for $t \in \mathcal{T}$, and we derive the general case, which is covered by the second lemma, using the special case.

The statement of Lemma 5.5 is changed to account for differences in the definitions. The mapping d is now defined to map s to a witness of the edge $h^{\check{S}}(s)$ instead of the unique image of the element of $h(v)$. The proof is analogous, by induction on the term t , although now we have to distinguish four cases.

Lemma 6.5. *Let $t \in \mathcal{T}$, $h: \mathbf{T}(t) \rightarrow \Omega(\mathbf{B})$, and $d: \check{S}^{\mathbf{T}(t)} \rightarrow B$ be a map such that, for all $s \in \check{S}^{\mathbf{T}(t)}$, $d(s)$ is a witness for the edge $h^{\check{S}}(s)$. Then d is a homomorphism $\Gamma(\mathbf{T}(t)) \rightarrow \mathbf{B}$.*

Proof. The proof is similar to the proof of Lemma 5.5. We first show by induction on the term t that either $d \in h(r_t)_t$ if t is a V -term, or $d \in \times_t(h^{\check{S}}(r_t); d(r_t))$ if t is an $\check{S}$ -term, or $d \in \times_t(h^S(r_t))$ if t is an S -term for a symbol $S \neq \check{S}$.

Case $t = \text{vertex}$: Since $\check{S}^{\mathbf{T}(\text{vertex})} = \emptyset$, the only map $d: \check{S}^{\mathbf{T}(\text{vertex})} \rightarrow B$ is the empty map, which is in $h(r_{\text{vertex}})_{\text{vertex}}$ by definition.

Case $t = \text{edge}_{\check{S}}(t_1, \dots, t_k)$: Note that restrictions of h to subtrees $\mathbf{T}(t_i)$'s are homomorphisms, hence, for all $i \in \{1, \dots, k\}$, $h(r_{t_i})_{t_i}$ contains the corresponding restrictions of d by the inductive assumption. The claim then follows from the definition of $\times_t(h^{\check{S}}(r_t); d(r_t))$.

Case $t = \text{edge}_S(t_1, \dots, t_k)$ where $S \neq \check{S}$: This is proved in the same way as the above case, with the exception that we use the definition of $\times_t(h^S(r_t))$ instead of $\times_t(h^{\check{S}}(r_t); d(r_t))$.

Case $t = \text{pr}_i(s)$: Since h is a homomorphism from $\mathbf{T}(t) = \mathbf{T}(s)$ to $\mathbf{B}$, we know that either $\times_s(h^S(r_s))$, if s is an S -term for $S \neq \check{S}$, or $\times_s(h^{\check{S}}(r_s); d(r_s))$, if s is an $\check{S}$ -term, contains d by the inductive assumption. The claim then follows from either (A1), or (B1).

The lemma then immediately follows by the definition if t is a V -term, by (B2) if t is an $\check{S}$ -term, or by (A2) otherwise. $\square$

Finally, in the proof of Lemma 5.6, we define f to map $e \in \check{S}^{\mathbf{A}}$ to a witness of the edge $g(e) \in \check{S}^{\Omega(\mathbf{B})}$ instead of letting $f(u)$ be the unique value attained by the map in $g(u)_{\text{vertex}}$. The rest of the proof is then a straightforward application of the above, i.e., completely analogous to the proof of Lemma 5.6 using Lemma 5.5.

Lemma 6.6. *If there is a homomorphism $g: \mathbf{A} \rightarrow \Omega(\mathbf{B})$, then there is a homomorphism $f: \Gamma(\mathbf{A}) \rightarrow \mathbf{B}$.*

Proof. As before, we assume that, for each τ -symbol R , $\mathbf{Q}_R$ and $\mathbf{T}(t_R)$ are equal and not just isomorphic. We define a mapping $f: \Gamma(\mathbf{A}) \rightarrow B$ by setting, for all $s \in \check{S}^{\mathbf{A}}$, $f(s) = e_{\bullet}$ for some witness $e_{\bullet}$ of the edge $g^{\check{S}}(s) \in \check{S}^{\Omega(\mathbf{B})}$, and claim that this f is a homomorphism from $\Gamma(\mathbf{A})$ to $\mathbf{B}$.

We need to show that f preserves each relation R . Assume that $(u_1, \dots, u_k) \in R^{\Gamma(\mathbf{A})}$, i.e., there is a homomorphism $h: \mathbf{Q}_R \rightarrow \mathbf{A}$ such that

$$(h^{\check{S}}(x_1), \dots, h^{\check{S}}(x_k)) = (u_1, \dots, u_k).$$

The previous lemma applied to the homomorphism $g \circ h: \mathbf{T}(t_R) \rightarrow \Omega(\mathbf{B})$ in place of h and the map $f \circ h^{\check{S}}$ in place of d implies that $f \circ h^{\check{S}}: \Gamma(\mathbf{Q}_R) \rightarrow \mathbf{B}$ is a homomorphism, which in turn implies

$$(f(u_1), \dots, f(u_k)) = (fh^{\check{S}}(x_1), \dots, fh^{\check{S}}(x_k)) \in R^{\mathbf{B}}$$

since $(x_1, \dots, x_k) \in R^{\Gamma(\mathbf{Q}_R)}$, as we wanted to show. $\square$

This concludes the proof of Theorem 6.3.

7. COMPOSITION OF ADJOINTS

In this section, we give an example of what can be achieved by composing functors defined in Sections 5 and 6. The power of composing two adjoints to obtain more complicated constructions was observed in [FT15, Section 5], where the authors considered composition of digraph functors with adjoints. This section gives several examples that show that we can obtain adjoints to more digraph functors by composing functors that go outside of the scope of digraphs into general relational structures. Naturally, our constructions also give more adjoints between general relational structures.

We start with a few general observations. The key fact that makes composition of adjoints useful is the following well-known category-theoretical observation.

Lemma 7.1. *Assume that Λ_1, Γ_1 and Λ_2, Γ_2 are two pairs of (thin) adjoint functors, then $\Lambda_1 \circ \Lambda_2$ is a left adjoint to $\Gamma_2 \circ \Gamma_1$.*

Proof. The proof is straightforward. We get the following string of equivalences from the two adjoints: $\Lambda_1 \Lambda_2(\mathbf{A}) \rightarrow \mathbf{B}$ if and only if $\Lambda_2(\mathbf{A}) \rightarrow \Gamma_1(\mathbf{B})$ if and only if $\mathbf{A} \rightarrow \Gamma_2 \Gamma_1(\mathbf{B})$ for any two structures $\mathbf{A}$ and $\mathbf{B}$ of the right signatures. $\square$

We note that a composition of Pultr functors gives a Pultr functor [Pul70]. We include a sketch of a proof of a slightly weaker statement.

Lemma 7.2. *Let Λ_1, Γ_1 and Λ_2, Γ_2 be two pairs of left and central Pultr functors such that $\Lambda_1 \circ \Lambda_2$ and $\Gamma_2 \circ \Gamma_1$ are well-defined. Then there is a pair of Pultr functors Λ and Γ such that $\Lambda(\mathbf{A})$ and $\Lambda_1 \circ \Lambda_2(\mathbf{A})$ are isomorphic for all $\mathbf{A}$, and $\Gamma(\mathbf{B})$ and $\Gamma_2 \circ \Gamma_1(\mathbf{B})$ are homomorphically equivalent for all $\mathbf{B}$.*

Proof sketch. The template of the composition can be obtained as the Λ_1 -image of the Pultr template defining Λ_2 and Γ_2 , i.e., if the template of Λ_2 and Γ_2 is composed of structures $\mathbf{P}$ and $\mathbf{Q}_R$, and homomorphisms $\epsilon_{i,R}$, then the template of the composition consists of structures $\Lambda_1(\mathbf{P})$ and $\Lambda_1(\mathbf{Q}_R)$, and homomorphisms $\epsilon_{i,R}^{\Lambda_1}: \Lambda_1(\mathbf{P}) \rightarrow \Lambda_1(\mathbf{Q}_R)$ that are induced by Λ_1 from $\epsilon_{i,R}: \mathbf{P} \rightarrow \mathbf{Q}_R$.

The composition $\Lambda_1 \circ \Lambda_2$ is, by definition, a two-step process: In the first step (applying Λ_2), we replace each vertex with a copy of $\mathbf{P}$, and in the second step (applying Λ_1), each of these copies is replaced with a copy of $\Lambda_1(\mathbf{P})$. Analogously, an R -edge is replaced with a copy of $\Lambda_1(\mathbf{Q}_R)$. It is also not hard to observe that the identification corresponds to the maps $\epsilon_{i,R}^{\Lambda_1}$, which concludes that, for all $\mathbf{A}$, $\Lambda_1 \circ \Lambda_2(\mathbf{A})$ is isomorphic to $\Lambda(\mathbf{A})$. We get that $\Gamma_2 \circ \Gamma_1(\mathbf{B})$ and $\Gamma(\mathbf{B})$ are homomorphically equivalent for all $\mathbf{B}$, since both functors are right adjoints to Λ . $\square$

The following theorem is obtained by composing adjoints constructed in Sections 5 and 6. Although it is not an exhaustive list of adjunctions that can be constructed by such compositions, it provides more adjoints to digraph functors on top of those provided in [FT15]. The theorem concerns a relatively general case of Pultr templates where $\mathbf{P}$ is an arbitrary tree. We only require that copies of $\mathbf{P}$ in the respective $\mathbf{Q}_R$'s intersect in at most one vertex. We also note that this theorem covers the cases of central Pultr functors whose adjoints are provided by Theorems 5.3 and 6.3.

Theorem 7.3. *Assume a (σ, τ) -Pultr template with $\mathbf{P}$ and all $\mathbf{Q}_R$'s being σ -trees such that, for each τ -symbol R ,*

- (1) $\epsilon_{i,R}$ is injective for all $i \in \{1, \dots, \text{ar } R\}$, and
- (2) for each $i \neq j$, $i, j \in \{1, \dots, \text{ar } R\}$, the images of $\mathbf{P}$ under $\epsilon_{i,R}$ and $\epsilon_{j,R}$ intersect in at most one vertex.

Then the corresponding central Pultr functor Γ has a right adjoint.

Proof. Assume that $P = \{1, \dots, p\}$. The goal is to decompose the functor Γ into two Pultr functors Γ_1 and Γ_2 . The intermediate step is to construct a structure of a new signature. This new signature ν is obtained from σ by adding a new relational symbol S of arity p (the size of the domain of $\mathbf{P}$) while retaining all symbols in σ .

We define the first functor, Γ_1 that maps σ -structures to v -structures. Essentially, this functor simply adds a new relation S that is defined by $\mathbf{P}$, i.e., $\Gamma_1(\mathbf{B})$ is the v -structure with domain B , where the relations are defined as

$$\begin{aligned} R^{\Gamma_1(\mathbf{B})} &= R^{\mathbf{B}} \text{ for each } \sigma\text{-symbol } R, \text{ and} \\ S^{\Gamma_1(\mathbf{B})} &= \{(h(1), \dots, h(p)) \mid h: \mathbf{P} \rightarrow \mathbf{B}\}. \end{aligned}$$

It is clear that this functor is defined by a (σ, v) -Pultr template that satisfies the assumptions of Theorem 5.3.

The second functor, Γ_2 is defined by altering the original Pultr template for Γ . The new template consists of $\mathbf{P}'$ and $\mathbf{Q}'_R$'s defined as follows. First, $\mathbf{P}' = \mathbf{S}_1$. For each τ -symbol R , we obtain $\mathbf{Q}'_R$ from $\mathbf{Q}_R$ by the following procedure: For each $i \in \{1, \dots, \text{ar } R\}$, remove from $\mathbf{Q}_R$ all edges in the image of $\epsilon_{i,R}$, and add the edge $(\epsilon_{i,R}(1), \dots, \epsilon_{i,R}(p)) \in S^{\mathbf{Q}'_R}$. Finally, we let $\epsilon'_{i,R} = \epsilon_{i,R}$, for all i and R , which is a homomorphism $\mathbf{P}' \rightarrow \mathbf{Q}'_R$ since we added the corresponding S -edge into $\mathbf{Q}'_R$. Observe that, since $\epsilon_{i,R}$'s are injective, the above procedure does not introduce reflexive tuples (i.e., tuples with repeated entries), and, since the images of $\epsilon_{i,R}$ and $\epsilon_{j,R}$ intersect in at most one vertex for $i \neq j$, this does not introduce cycles into $\mathbf{Q}'_R$ using S -edges, and, since all other edges have been removed in the image, it does not introduce any other cycles. Hence, $\mathbf{Q}'_R$ is still a tree for each τ -symbol R , and therefore Theorem 6.3 applies.

The above two paragraphs show that Γ_1 and Γ_2 have adjoints Ω_1 and Ω_2 . And it is straightforward to check that $\Gamma = \Gamma_2 \circ \Gamma_1$. This concludes that $\Omega_1 \circ \Omega_2$ is a right adjoint to Γ by Lemma 7.1. $\square$

8. CONCLUSION

We have studied the problem of characterising central Pultr functors for arbitrary relational structures that admit a right adjoint, and, for those that do, giving an explicit construction for such an adjoint. There is a necessary condition for the existence of such an adjoint (cf. Theorem 2.10 and comments after it). We gave a sufficient condition in Theorem 7.3. These two conditions do not match, there is a gap between them, and it is not quite clear what the necessary and sufficient condition should be (even in the case of digraphs). Apart from the requirement that $\mathbf{P}$ and all $\mathbf{Q}_R$'s are trees, Theorem 7.3 has two additional assumptions. We believe that the second assumption (about intersection of images of $\mathbf{P}$ in $\mathbf{Q}_R$) is a technicality that can be removed with some extra work. How essential is the first assumption (about injectivity of homomorphisms $\epsilon_{i,R}$)? For example, is it true that, for every central Pultr functor Γ that has a right adjoint, there is another central Pultr functor Γ' such that (a) for every structure $\mathbf{A}$ of appropriate signature, $\Gamma(\mathbf{A})$ and $\Gamma'(\mathbf{A})$ are homomorphically equivalent, and (b) the Pultr template corresponding to Γ' has all homomorphisms $\epsilon_{i,R}$ injective?

Finally, let us discuss possible applications of our results in the complexity of (promise) CSPs. Left and central Pultr functors can alternatively be described as “gadget replacements” and “pp-constructions”, respectively. These two constructions are central to the algebraic theory of complexity of (promise) CSPs where gadget replacements are used as log-space reductions between (promise) CSPs whose templates are related via pp-constructions; see, e.g., [BBKO21] or [KOWŽ23, Section 4.1]. The opposite type of reductions, where gadgets are replaced by individual constraints, is less common, but has already been shown to be important; see, e.g., [KOWŽ23, BK22]. It was discovered recently, that the right adjoints to some functor can be used to characterise when the functor is a valid reduction between two (promise) CSPs [KOWŽ23, Theorem 4.6]. Moreover, recent developments in the theory of promise CSPs have called for characterisation of reductions which include central Pultr functors [BK22, KO22, DO23]. Our results provide a characterisation of when a reduction from a general class is a valid reduction between two promise CSPs. Although we are not aware of a specific interesting case of a promise CSP whose hardness is proven by providing a right adjoint to

a central Pultr functor, apart from the use of the arc graph construction in [WŽ20], our general results can be used as a foundation for further investigation of these reductions.

Acknowledgements

We would like to thank the reviewers for many insightful comments on the manuscript.

References

- [AGH17] Per Austrin, Venkatesan Guruswami, and Johan Håstad. $(2+\epsilon)$ -Sat is NP-hard. *SIAM J. Comput.*, 46(5):1554–1573, 2017. <https://eccc.weizmann.ac.il/report/2013/159/>, doi:10.1137/15M1006507.
- [BBKO21] Libor Barto, Jakub Bulín, Andrei Krokhin, and Jakub Opršal. Algebraic approach to promise constraint satisfaction. *J. ACM*, 68(4):28:1–66, August 2021. doi:10.1145/3457606.
- [BK22] Libor Barto and Marcin Kozik. Combinatorial gap theorem and reductions between promise CSPs. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1204–1220, Virtual Conference, 2022. SIAM. doi:10.1137/1.9781611977073.50.
- [BKW17] Libor Barto, Andrei Krokhin, and Ross Willard. Polymorphisms, and how to use them. In Andrei Krokhin and Stanislav Živný, editors, *The Constraint Satisfaction Problem: Complexity and Approximability*, volume 7 of *Dagstuhl Follow-Ups*, pages 1–44. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 2017. https://drops.dagstuhl.de/opus/frontdoor.php?source_opus=6959, doi:10.4230/DFU.Vol7.15301.1.
- [DO23] Victor Dalmau and Jakub Opršal. Local consistency as a reduction between constraint satisfaction problems. *preprint*, 2023. arXiv:2301.05084, doi:10.48550/arXiv.2301.05084.
- [FT13] Jan Foniok and Claude Tardif. Adjoint functors in graph theory. *preprint*, 2013. arXiv:1304.2215.
- [FT15] Jan Foniok and Claude Tardif. Digraph functors which admit both left and right adjoints. *Discrete Math.*, 338(4):527–535, 2015. arXiv:1304.2204, doi:10.1016/j.disc.2014.10.018.
- [FV98] Tomás Feder and Moshe Y. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: A study through datalog and group theory. *SIAM J. Comput.*, 28(1):57–104, February 1998. doi:10.1137/S0097539794266766.
- [HN04] Pavol Hell and Jaroslav Nešetřil. *Graphs and Homomorphisms*. Oxford University Press, 2004.
- [KO22] Andrei Krokhin and Jakub Opršal. An invitation to the promise constraint satisfaction problem. *ACM SIGLOG News*, 9(3):30–59, 2022. arXiv:2208.13538, doi:10.1145/3559736.3559740.
- [KOWŽ23] Andrei Krokhin, Jakub Opršal, Marcin Wrochna, and Stanislav Živný. Topology and adjunction in promise constraint satisfaction. *SIAM J. Comp.*, 52(1):38–79, 2023. arXiv:2003.11351, doi:10.1137/20M1378223.
- [KŽ17] Andrei Krokhin and Stanislav Živný, editors. *The Constraint Satisfaction Problem: Complexity and Approximability*, volume 7 of *Dagstuhl Follow-Ups*. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017. <http://www.dagstuhl.de/dagpub/978-3-95977-003-3>.
- [LLT07] Benoît Larose, Cynthia Loten, and Claude Tardif. A characterisation of first-order constraint satisfaction problems. *Logical Methods in Computer Science*, 3(4), 2007. doi:10.2168/LMCS-3(4:6)2007.
- [NT00] Jaroslav Nešetřil and Claude Tardif. Duality theorems for finite structures (characterising gaps and good characterisations). *Journal of Combinatorial Theory, Series B*, 80(1):80–97, 2000. doi:10.1006/jctb.2000.1970.
- [NT05] Jaroslav Nešetřil and Claude Tardif. Short answers to exponentially long questions: Extremal aspects of homomorphism duality. *SIAM J. Discret. Math.*, 19(4):914–920, August 2005. doi:10.1137/S0895480104445630.
- [Pul70] Aleš Pultr. The right adjoints into the categories of relational systems. In *Reports of the Midwest Category Seminar IV*, volume 137 of *Lecture Notes in Mathematics*, pages 100–113. Springer, 1970. <http://www.springer.com/us/book/9783540049265>, doi:10.1007/BFb0060437.
- [WŽ20] Marcin Wrochna and Stanislav Živný. Improved hardness for H -colourings of G -colourable graphs. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA’20)*, pages 1426–1435, 2020. arXiv:1907.00872, doi:10.1137/1.9781611975994.86.