[BT94] R. Backofen and R. Treinen, "How to Win a Game with Features," in *Constraints in Computational Logics, Proc. CCL'94*, J.-P. Jouannaud (Ed.), Springer LNCS **845**, 1994, pp. 320-335.

[Co84] A. Colmerauer, "Equations and inequations on finite and infinite trees," in: *Proc. 2nd Int. Conf. on Fifth Generation Computer Systems*, 1984, pp.85–99.

[Co90] A. Colmerauer, "An introduction to PROLOG III," *C. ACM* **33**, 1990, pp. 69–90.

[Co83] B. Courcelle, "Fundamental Properties of Infinite Trees," *Theoretical Computer Science* **25**, 1983, pp. 95–169.

[Ma88] M.J. Maher, "Complete axiomatizations of the algebras of finite, rational and infinite trees," in: *Proceedings of Third Annual Symposium on Logic in Computer Science, LICS'88*, pp. 348–357, Edinburgh, Scotland, 1988. IEEE Computer Society.

[Ma71] A.I. Mal'cev, "The Metamathematics of Algebraic Systems," volume 66 of *Studies in Logic and the Foundation of Mathematics*, North Holland, Amsterdam, London, 1971.

[Ro88] W.C. Rounds, "Set values for unification based grammar formalisms and logic programming," *Research Report* CSLI-88-129, Stanford, 1988.

[ST94] G. Smolka, R. Treinen, "Records for Logic Programming," *J. of Logic Programming* **18**(3) (1994), pp. 229-258.

and the rich amount of interesting relationships between them, are now well-understood (e.g., [Co83, Ma88]). We believe that free amalgamation, rational amalgamation and a further construction called "infinite amalgamation" (still to be investigated) reflect this role on the higher level of amalgamation constructions. Many of the results that we have obtained for free and rational amalgamation can be interpreted in this sense:

- The universality-property of the free amalgamated product (see [BS95]) reflects the status of the free term algebra as the absolutely free Σ-algebra.
- We have seen that the free amalgamated product is always a substructure of the rational amalgamated product. This reflects the fact that the free term algebra is always a substructure of the algebra of rational trees.
- It is well-known that the unification algorithm for the algebra of rational trees can be considered as the variant of the unification algorithm for the free term algebra where we omit the occur-check. Similarly, the decomposition scheme for rational amalgamation as given here is—essentially—the decomposition scheme for free amalgamation where we omit the "inter-structural" occur-check that is provided by the choice of a linear ordering in the latter scheme.

We would not be surprised if much more principles, techniques and theorems, well-known on the level of tree constructions, could be lifted to the level of combining structures. Our experience with rational amalgamation seems to indicate that this is a difficult, but promising line of research if we want to understand the scale of possibilities, and the limitations for combining solution domains and constraint solvers.

References

[Ac88] P. Aczel, "Non-well-founded Sets," *CSLI* Lecture Notes **14**, Stanford University, 1988.

[AP94] H. Ait-Kaci, A. Podelski, and G. Smolka, "A feature-based constraint system for logic programming with entailment," *Theoretical Comp. Science* **122**, 1994, pp.263–283.

[BS95] F. Baader and K.U. Schulz, "On the Combination of Symbolic Constraints, Solution Domains, and Constraint Solvers," in: Proceedings CP'95, U.Montanari, F.Rossi (Eds.), Springer LNCS 976, pp. 380-397.

[BT94] R. Backofen and G. Smolka, "A Complete and Recursive Feature Theory," *Theoretical Comp. Science* **146**, 1995, pp. 243–268.

where the sequence $\vec{y}'$ is obtained from $\vec{y}$ by removing y_i. Since the elements in the sequence $\vec{y}'$ are distinct atoms it follows as above that $\mathcal{B}^\Delta \models \delta'$.

In this second case we have seen that we can construct a new output pair (σ', δ') of Algorithm 2 such that $\mathcal{A}^\Sigma \models \sigma'$ and $\mathcal{B}^\Delta \models \delta'$. Moreover, the number of variables in (σ', δ') is strictly smaller than the number of variables in (σ, δ). We may now use the same subcase analysis as above, replacing (σ, δ) by (σ', δ'), and iterate this contraction of formulae, if necessary. After a finite number of steps we reach an output pair that satisfies all the assumptions that we made for (σ, δ) in the first subcase. As we have seen, this shows that the input formula γ has a solution in $\mathcal{A}^\Sigma \odot \mathcal{B}^\Delta$. □

As the last step, we show completeness of Algorithm 2.

Lemma 6.13 *If the input constraint γ has a solution in $\mathcal{A}^\Sigma \odot \mathcal{B}^\Delta$, then there exists an output pair (σ, δ) of Algorithm 2 such that $\mathcal{A}^\Sigma \models \sigma$ and $\mathcal{B}^\Delta \models \delta$.*

Proof. Lemma 6.10 shows that Algorithm 1 has an output pair $(\langle \gamma_1^\Sigma, U, W\rangle, \langle \gamma_1^\Delta, W, U\rangle)$ such that $\langle \gamma_1^\Sigma, U, W\rangle$ has a solution in $\mathcal{A}^\Sigma$ and $\langle \gamma_1^\Delta, W, U\rangle$ has a solution in $\mathcal{B}^\Delta$. In $\mathcal{A}^\Sigma$, variables of U are interpreted as distinct atoms in X under the given solution. Lemma 3.6 shows that $\mathcal{A}^\Sigma \models \forall\vec{u}\exists\vec{w}\exists\vec{v}_{1,\Sigma}\, \gamma_1^\Sigma$. In $\mathcal{B}^\Delta$, variables of W are interpreted as distinct atoms in Y under the given solution. By Lemma 3.6, $\mathcal{B}^\Delta \models \forall\vec{w}\exists\vec{u}\exists\vec{v}_{1,\Delta}\, \gamma_1^\Delta$. This shows that the sentences $\sigma := \forall\vec{u}\exists\vec{w}\exists\vec{v}_{1,\Sigma}\, \gamma_1^\Sigma$ and $\delta := \forall\vec{w}\exists\vec{u}\exists\vec{v}_{1,\Delta}\, \gamma_1^\Delta$ of the corresponding output pair (σ, δ) of Algorithm 2 are valid in $\mathcal{A}^\Sigma$ and $\mathcal{B}^\Delta$ respectively. □

7 Conclusion

In this paper we have introduced rational amalgamation, a general methodology for combining constraint systems. The present work, in connection with the discussion of free amalgamation in [BS95], seems to suggest a new view of the problem of combining solution domains and constraint solvers. There is now strong evidence that the situation considered in [BS95] and in this paper—the construction of "mixed" elements of a combined domain, given the "pure" elements of two component structures as construction units—is quite similar to the process of building the elements of a single structure, given the symbols of a fixed signature as construction units. We are confident that this analogy will help to isolate the most important methods for combining structures, and to understand the relationship and the differences between different amalgamation constructions.

When we compose elements, given the symbols of a fixed signature, three different structures may be obtained in a direct way, depending on the composition principle, namely the free term algebra, the algebra of rational trees, and the algebra of infinite trees. The privileged role of these three algebras, *Proof.* Assume that $\mathcal{A}^\Sigma \models \forall\vec{u}\exists\vec{w}\exists\vec{v}_{1,\Sigma}\ \gamma_1^\Sigma$ and $\mathcal{B}^\Delta \models \forall\vec{w}\exists\vec{u}\exists\vec{v}_{1,\Delta}\ \gamma_1^\Delta$. Let $\vec{u} = u_1, \ldots, u_m$, let $\vec{w} = w_1, \ldots, w_n$. For each variable u_i we select a distinct atom $x_i \in X$ of A $(1 \leq i \leq m)$, and for each variable w_j we select a distinct atom $y_j \in Y$ of B $(1 \leq j \leq n)$. Then there are elements $a_1, \ldots, a_n \in A$ and $b_1, \ldots, b_m \in B$ such that

$$\begin{array}{lcl} \mathcal{A}^\Sigma & \models & \exists\vec{v}_{1,\Sigma}\ \gamma_1^\Sigma(\vec{u}/\vec{x}, \vec{w}/\vec{a}) \\ \mathcal{B}^\Delta & \models & \exists\vec{v}_{1,\Delta}\ \gamma_1^\Delta(\vec{u}/\vec{b}, \vec{w}/\vec{y}). \end{array}$$

We distinguish two cases.

First case: $x_i \neq a_j$ and $b_i \neq y_j$, for all $1 \leq i \leq n$ and $1 \leq j \leq m$. Since $\mathcal{A}^\Sigma$ is non-trivial, we may choose an endomorphism $m_1 \in \mathcal{M}$ that maps all atoms in the set $\{a_1, \ldots, a_n\}$ to a non-atomic element $a \in A$ and fixes all other atoms. In particular, m_1 leaves the atoms $x_1, \ldots, x_m$ fixed, by assumption. Since $\mathcal{A}^\Sigma$ is non-collapsing, all elements in the set $\{m_1(a_1), \ldots, m_1(a_n)\}$ are non-atomic. Since γ_1^Σ is a positive formula we have

$$\mathcal{A}^\Sigma \models \exists\vec{v}_{1,\Sigma}\ \gamma_1^\Sigma(\vec{u}/\vec{x}, \vec{w}/\vec{m}_1(a)),$$

by Lemma 3.5. It follows that the Σ-constraint with A/N declaration, (γ_1^Σ, U, W), has a solution in $\mathcal{A}^\Sigma$.

Symmetrically we may choose an endomorphism $n_1 \in \mathcal{N}$ such that all elements in $\{n_1(b_1), \ldots, n_1(b_m)\}$ are non-atomic and

$$\mathcal{B}^\Delta \models \exists\vec{v}_{1,\Delta}\ \gamma_1^\Delta(\vec{u}/n_1(\vec{b}), \vec{w}/y).$$

It follows that the Δ-constraint with A/N declaration, (γ_1^Δ, W, U), has a solution in $\mathcal{B}^\Delta$. Now Lemma 6.9 shows that the input formula γ has a solution in $\mathcal{A}^\Sigma \odot \mathcal{B}^\Delta$.

Second case: Without loss of generality, $x_i = a_j$, for some $1 \leq i \leq m$ and $1 \leq j \leq n$. We consider the new formula $\gamma'_{1,\Sigma}$ $(\gamma'_{1,\Delta})$ that is obtained by replacing all occurrences of w_i in γ_1^Σ (resp. γ_1^Δ) by u_j. Consider the pair with the formulae $\sigma' = \forall\vec{u}\exists\vec{w}'\exists\vec{v}_{1,\Sigma}\ \gamma'_{1,\Sigma}$ and $\delta' = \forall\vec{w}'\exists\vec{u}\exists\vec{v}_{1,\Delta}\ \gamma'_{1,\Delta}$, where the sequence $\vec{w}'$ is obtained from $\vec{w}$ by removing w_i. Obviously, (σ', δ') is again an output pair Algorithm 2. We claim that $\mathcal{A}^\Sigma \models \sigma'$ and $\mathcal{B}^\Delta \models \delta'$.

We have

$$\mathcal{A}^\Sigma \models \exists\vec{v}_{1,\Sigma}\ \gamma'_{1,\Sigma}(\vec{u}/\vec{x}, \vec{w}'/\vec{a}'),$$

where $\vec{a}'$ denotes the sequence $a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n$. Since X is an $\mathcal{M}$-atom set, for each sequence $\vec{c} = c_1, \ldots, c_m$ of elements of A there exists an endomorphism $m_2 \in M$ such that $m_2(x_i) = c_i$, for $1 \leq i \leq m$. Now Lemma 3.5 shows that $\mathcal{A}^\Sigma \models \sigma'$.

Since $(\mathcal{B}^\Delta, \mathcal{N}, Y)$ is rational, there exists an endomorphism $n_2 \in \mathcal{N}$ that leaves all atoms but y_i fixed such that $n_2(y_i) = n_2(b_j)$. By Lemma 3.5,

$$\mathcal{B}^\Delta \models \exists\vec{v}_{1,\Delta}\ \gamma'_{1,\Delta}(\vec{u}/n_2(\vec{b}), \vec{w}'/\vec{y}'),$$

is a non-trivial braid of type B. Consequently, W contains all variables w of V_1 such that $\mu_{A\odot B}(w)$ is a trivial braid or a non-trivial braid of type A. The definition of $root_A$ implies that $\mu_{A\odot B} \circ root_A(u)$ is an open atom of $\mathcal{A}^{\Sigma}_{*}$, for all $u \in U$, and $\mu_{A\odot B} \circ root_A(w)$ is a non-atomic element or a bottom atom of $\mathcal{A}^{\Sigma}_{*}$, for all $w \in W$. Let $m_1 \in \mathcal{M}$ be an endomorphism that maps all the bottom atoms of the set $\{\mu_{A\odot B} \circ root_A(w) \mid w \in W\}$ to a non-atomic element of A and leaves all other atoms fixed. Since $\mathcal{A}^{\Sigma}$ is non-collapsing, all elements of the set $\{\mu_{A\odot B} \circ root_A \circ m_1(w) \mid w \in W\}$ are non-atomic. Since σ is a positive formula, Lemma 3.5 implies that $\nu_A := \mu_{A\odot B} \circ root_A \circ m_1$ is a solution of $\langle \sigma, U, W\rangle$ in $\mathcal{A}^{\Sigma}_{*}$.

On the other hand the definition of $root_B$ implies that $\mu_{A\odot B} \circ root_B(w)$ is an atom of $\mathcal{B}^{\Delta}_{*}$, for all $w \in W$, and $\mu_{A\odot B} \circ root_B(u)$ is a non-atomic element of $\mathcal{B}^{\Delta}_{*}$, for all $u \in U$. This shows that $\langle \delta, W, U\rangle$ has a solution in $\mathcal{B}^{\Delta}_{*}$.

But then, by Lemma 4.34, $\langle \sigma, U, W\rangle$ has a solution in $\mathcal{A}^{\Sigma}$ and $\langle \delta, W, U\rangle$ has a solution in $\mathcal{B}^{\Delta}$. □

6.1 Proof of Theorem 6.5

In order to proof Theorem 6.5 we shall use the following variant of Algorithm 1, which we call

Algorithm 2

The input constraint γ, and Steps 1 and 2, remain as above. The output of Algorithm 2 consists of the two positive universal-existential sentences

$$\sigma = \forall \vec{u} \exists \vec{w} \exists \vec{v}_{1,\Sigma}\ \gamma_1^{\Sigma}$$

and

$$\delta = \forall \vec{w} \exists \vec{u} \exists \vec{v}_{1,\Delta}\ \gamma_1^{\Delta}$$

where $\vec{u}$ ($\vec{w}$) represent the variables in U (resp. V), $\vec{v}_{1,\Sigma}$ represents the non-shared variables in γ_1^{Σ}, and $\vec{v}_{1,\Delta}$ represents the non-shared variables in γ_1^{Δ}.

Proposition 6.11 *The input formula γ has a solution in $\mathcal{A}^{\Sigma} \odot \mathcal{B}^{\Delta}$ if and only if there exists an output pair (σ, δ) of Algorithm 2 such that $\mathcal{A}^{\Sigma} \models \sigma$ and $\mathcal{B}^{\Delta} \models \delta$.*

Theorem 6.5 is an immediate consequence. In order to prove Proposition 6.11 we shall first show that Algorithm 2 is sound. As above we shall assume that the two components $\mathcal{A}^{\Sigma}$ and $\mathcal{B}^{\Delta}$ have the form $(\mathcal{A}^{\Sigma}, \mathcal{M}, X)$ and $(\mathcal{B}^{\Delta}, \mathcal{N}, Y)$ respectively.

Lemma 6.12 *If, for some output pair (σ, δ) of Algorithm 2, $\mathcal{A}^{\Sigma} \models \sigma$ and $\mathcal{B}^{\Delta} \models \delta$, then γ is solvable in $\mathcal{A}^{\Sigma} \odot \mathcal{B}^{\Delta}$.*

2. Let $\pi_A := \{\langle x_i, b_i\rangle \mid i = 1, \ldots, m\}$, let $\pi_B := \{\langle y_i, a_i\rangle \mid i = 1, \ldots, n\}$. Properties (1)–(3) and (5)–(7) show that for each $e \in \vec{a}$ ($e \in \vec{b}$), the tuple $\mathcal{K}_e := \langle e, \{a_1, \ldots, a_n\}, \{b_1, \ldots, b_m\}, \pi_A, \pi_B\rangle$ is a prebraid of type A (B).

3. Fix some $e \in \vec{a} \cup \vec{b}$. Let (m_3, n_3) be the standard normalizer for $\mathcal{K}_e$. By Lemma 3.5, (4), and (8),

$$\begin{array}{lcl} \mathcal{A}^\Sigma & \models & \exists \vec{v}_\Sigma\ \gamma_1^\Sigma(\vec{u}/m_3(\vec{x}), \vec{w}/m_3(\vec{a})), \\ \mathcal{B}^\Delta & \models & \exists \vec{v}_\Delta\ \gamma_1^\Delta(\vec{u}/n_3(\vec{b}), \vec{w}/n_3(\vec{y})). \end{array}$$

It follows easily from Lemma 4.34 that

$$\begin{array}{lcl} \mathcal{A}_*^\Sigma & \models & \exists \vec{v}_\Sigma\ \gamma_1^\Sigma(\vec{u}/m_3(\vec{x}), \vec{w}/m_3(\vec{a})), \\ \mathcal{B}_*^\Delta & \models & \exists \vec{v}_\Delta\ \gamma_1^\Delta(\vec{u}/n_3(\vec{b}), \vec{w}/n_3(\vec{y})). \end{array}$$

Now Theorem 5.5 shows that

$$\begin{array}{lcl} \mathcal{A}^\Sigma \odot \mathcal{B}^\Delta & \models & \exists \vec{v}_\Sigma\ \gamma_1^\Sigma(\vec{u}/root_A^{-1}(m_3(\vec{x})), \vec{w}/root_A^{-1}(m_3(\vec{a}))), \\ \mathcal{A}^\Sigma \odot \mathcal{B}^\Delta & \models & \exists \vec{v}_\Delta\ \gamma_1^\Delta(\vec{u}/root_B^{-1}(n_3(\vec{b})), \vec{w}/root_B^{-1}(n_3(\vec{y}))). \end{array}$$

Consider an element x_i of $\vec{x}$. Assume that x_i points in $\mathcal{K}_e$ to the subbraid $\mathcal{K}'$ with root b_i. Then $m_3(x_i) = o_{[\mathcal{K}']}$. Let $\mathcal{K}_i$ be the subbraid of $\mathcal{K}_e^{(m_3,n_3)}$ with root $n_3(b_i)$. By Lemma 4.18, $\mathcal{K}'$ and $\mathcal{K}_i$ are equivalent. It follows that $m_3(x_i) = o_{[\mathcal{K}_i]}$. The braid $\mathcal{K}_i$ is non-trivial and of type B, and it is the unique braid in standard normal form with root $n_3(b_i)$ (Prop. 4.37, Lemma 4.41). Hence $root_A^{-1}(m_3(x_i)) = \mathcal{K}_i$. The element $n_3(b_i)$ is a non-atomic element of B. Hence $root_B^{-1}(n_3(\vec{b})) = \mathcal{K}_i$ is the unique braid in standard normal form with root $n_3(b_i)$. Thus we have seen that $root_A^{-1}(m_3(\vec{x})) = root_B^{-1}(n_3(\vec{b}))$. Similarly it follows that $root_B^{-1}(n_3(\vec{y})) = root_A^{-1}(m_3(\vec{a}))$. This shows that the formula $\gamma_1^\Sigma \wedge \gamma_1^\Delta$ obtained after Step 1 has a solution in $\mathcal{A}^\Sigma \odot \mathcal{B}^\Delta$. Obviously this implies that the input constraint γ has a solution in $\mathcal{A}^\Sigma \odot \mathcal{B}^\Delta$. □

Next, we show completeness of the Algorithm 1.

Proposition 6.10 *If the input constraint γ has a solution in $\mathcal{A}^\Sigma \odot \mathcal{B}^\Delta$, then there exists an output pair $(\langle\sigma, U, W\rangle, \langle\delta, W, U\rangle)$ of Algorithm 1 such that $\langle\sigma, U, W\rangle$ has a solution in $\mathcal{A}^\Sigma$ and $\langle\delta, W, U\rangle$ has a solution in $\mathcal{B}^\Delta$.*

Proof. Assume that γ has a solution $\mu_{A\odot B}$ in $\mathcal{A}^\Sigma \odot \mathcal{B}^\Delta$.

In Step 1 of Algorithm 1 we identify two shared variables v and v' of V_1 if, and only if, $\mu_{A\odot B}(v) = \mu_{A\odot B}(v')$. With this choice, $\mu_{A\odot B}$ is a solution of the formula $\gamma_1^\Sigma \wedge \gamma_1^\Delta$ that is reached after Step 1, and $\mu_{A\odot B}$ assigns distinct values in $\mathcal{A}^\Sigma \odot \mathcal{B}^\Delta$ to all variables of V_1.

By Theorem 5.5, $\mu_{A\odot B} \circ root_A$ (resp. $\mu_{A\odot B} \circ root_B$) is a solution of $\sigma = \gamma_1^\Sigma$ in $\mathcal{A}_*^\Sigma$ (resp. of $\delta = \gamma_1^\Delta$ in $\mathcal{B}_*^\Delta$) that does not identify two variables of V_1.

By assumption, one of the two component structures, $\mathcal{A}^\Sigma$, say, is non-trivial. In Step 2, we choose as U the set of all variables u of V_1 such that $\mu_{A\odot B}(u)$

Lemma 6.9 *If, for some output pair $(\langle\sigma, U, W\rangle, \langle\delta, W, U\rangle)$ of Algorithm 1, $\langle\sigma, U, W\rangle$ has a solution in $\mathcal{A}^\Sigma$ and $\langle\delta, W, U\rangle$ has a solution in $\mathcal{B}^\Delta$, then the input constraint γ is solvable in $\mathcal{A}^\Sigma \odot \mathcal{B}^\Delta$.*

Proof. The output formulae σ and δ may be written in the form $\gamma_1^\Sigma(\vec{u}, \vec{w}, \vec{v}_\Sigma)$ and $\gamma_1^\Delta(\vec{u}, \vec{w}, \vec{v}_\Delta)$, where $\vec{u} = u_1, \ldots, u_m$ denotes the sequence of all elements of U, where $\vec{w} = w_1, \ldots, w_n$ denotes the sequence of all elements of W, and where $\vec{v}_\Sigma$ (resp. $\vec{v}_\Delta$) stands for the non-shared variables occurring in γ_1^Σ and γ_1^Δ respectively. The proof has now three steps. In the first step, the given solutions of the output constraints are used to construct similar solutions of a more specific form. In the second step, these latter solutions are used to define suitable braids. In the third step we apply standard normalization to these braids. This will yield a solution of the input constraint.

1. By assumption, there exists a solution μ_A of γ_1^Σ in $\mathcal{A}^\Sigma$ such that the elements $\mu_A(u_1), \ldots, \mu_A(u_m)$ are distinct atoms of $\mathcal{A}^\Sigma$, and the elements $\mu_A(w_1), \ldots, \mu_A(w_n)$ are non-atomic elements of $\mathcal{A}^\Sigma$. If some of the atoms $\mu_A(u_1), \ldots, \mu_A(u_m)$ are bottom atoms, then we apply an automorphism $m_1 \in \mathcal{M}$ such that the elements in $\{m_1(\mu_A(u_1)), \ldots, m_1(\mu_A(u_m))\}$ are distinct open atoms. In the other case, let $m_1 := Id$. If the stabilizers of the elements $m_1(\mu_A(w_1)), \ldots, m_1(\mu_A(w_n))$ contain open atoms $o_1, \ldots, o_k$ that do not belong to $\{m_1(\mu_A(u_1)), \ldots, m_1(\mu_A(u_m))\}$, then we apply an endomorphism m_2 that maps the atoms $o_1, \ldots, o_k$ to some bottom atom z and leaves the atoms $\{m_1(\mu_A(u_1)), \ldots, m_1(\mu_A(u_m))\}$ fixed. In the other case, let $m_2 := Id$. Since γ_1^Σ is a positive formula, $\nu_A := \mu_A \circ m_1 \circ m_2$ is a solution of γ_1^Σ, by Lemma 3.5. We have

(1) the elements $x_1 := \nu_A(u_1), \ldots, x_m := \nu_A(u_m)$ are distinct *open* atoms,

(2) the elements $a_1 := \nu_A(w_1), \ldots, a_n := \nu_A(w_n)$ are non-atomic,

(3) the open atoms occurring in the stabilizers of the elements $a_1, \ldots, a_n$ are in $\{x_1, \ldots, x_m\}$, and

(4) $\mathcal{A}^\Sigma \models \exists \vec{v}_\Sigma\ \gamma_1^\Sigma(\vec{u}/\vec{x}, \vec{w}/\vec{a})$.

(2) follows from the fact $\mathcal{A}^\Sigma$ is non-collapsing, (3) follows from Lemma 3.10, and (4) follows from the fact that ν_A solves γ_1^Σ. Symmetrically we can show that there exists a solution ν_B of γ_1^Δ in $\mathcal{B}^\Delta$ such that

(5) the elements $y_1 := \nu_B(w_1), \ldots, y_n := \nu_B(w_n)$ are distinct *open* atoms,

(6) the elements $b_1 := \nu_B(u_1), \ldots, b_m := \nu_B(u_m)$ are non-atomic,

(7) the open atoms occurring in the stabilizers of the elements $b_1, \ldots, b_m$ are in $\{y_1, \ldots, y_n\}$, and

(8) $\mathcal{B}^\Delta \models \exists \vec{v}_\Delta\ \gamma_1^\Delta(\vec{u}/\vec{b}, \vec{w}/\vec{y})$.

Proof of Theorem 6.3

To prove Theorem 6.3 we shall give an algorithm that reduces a mixed constraint γ in the signature $(\Sigma \cup \Delta)$ non-deterministically to a pair of constraints with A/N declarations over the "pure" signatures Σ and Δ respectively. We shall assume that the input formula γ has the form $\gamma = \gamma_0^\Sigma \wedge \gamma_0^\Delta$ where γ_0^Σ is a conjunction of atomic Σ-formulae, and γ_0^Δ is a conjunction of atomic Δ-formulae. Moreover we assume that γ does not contain any equation between variables. These assumptions do not really restrict the generality of the approach: simple techniques like "variable abstraction", now standard in this area, may be used to transform an arbitrary $(\Sigma \cup \Delta)$-constraint φ into a constraint γ of the form given above, preserving solvability in both directions.

Algorithm 1

The *input* is mixed a constraint $\gamma = \gamma_0^\Sigma \wedge \gamma_0^\Delta$ of the form described above. Let $V_0 = Var(\gamma_0^\Sigma) \cap Var(\gamma_0^\Delta)$ denote the set of *shared* variables of γ. The algorithm has two steps, both are nondeterministic.

Step 1: Variable identification. *Consider all possible partitions of the set of all shared variables, V_0. Each of these partitions yields one of the new constraints as follows. The variables in each class of the partition are "identified" with each other by choosing an element of the class as representative, and replacing in the input formula all occurrences of variables of the class by this representative.*

Step 2: Choose signature labels. *Let $\gamma_1^\Sigma \wedge \gamma_1^\Delta$ denote one of the formulae obtained by Step 1, let V_1 denote the set of representants of shared variables. The set V_1 is partitioned in two subsets U and W in some arbitrary way.*

Let $\sigma = \gamma_1^\Sigma$, let $\delta = \gamma_1^\Delta$. For each of the choices made in Step 1 and 2, the algorithm yields an *output pair* $(\langle \sigma, U, W\rangle, \langle \delta, W, U\rangle)$, each component representing a constraint with A/N declaration. □

Correctness of Algorithm 1

We shall prove that Algorithm 1 is correct in the following sense.

Proposition 6.8 *The input formula γ has a solution in $\mathcal{A}^\Sigma \odot \mathcal{B}^\Delta$ if and only if there exists an output pair $(\langle \sigma, U, W\rangle, \langle \delta, W, U\rangle)$ of Algorithm 1 such that $\langle \sigma, U, W\rangle$ has a solution in $\mathcal{A}^\Sigma$ and $\langle \delta, W, U\rangle$ has a solution in $\mathcal{B}^\Delta$.*

Note that Theorem 6.3 is an immediate consequence. In order to prove Proposition 6.8 we shall assume that the two components $\mathcal{A}^\Sigma$ and $\mathcal{B}^\Delta$ are SC-structures of the form $(\mathcal{A}^\Sigma, \mathcal{M}, X)$ and $(\mathcal{B}^\Delta, \mathcal{N}, Y)$ respectively. First we show soundness.

Definition 6.4 A non-collapsing SC-structure $(\mathcal{A}^\Sigma, \mathcal{M}, X)$ is called *rational* if for every atom $x \in X$ and every element $a \in A$ there exists an endomorphism $m \in \mathcal{M}$ that leaves all atoms $x' \neq x$ fixed such that $m(x) = m(a)$.[10]

The algebra of rational trees over a given signature is always a rational SC-structure. The same holds for feature structures, feature structures with arity, and domains with nested, rational lists (as described in 3.4). For rational SC-structures we obtain the following refinement and reformulation of Theorem 6.3.

Theorem 6.5 *Let $\mathcal{A}^\Sigma$ and $\mathcal{B}^\Delta$ be two non-trivial rational SC-structures over disjoint signatures, let $\mathcal{A}^\Sigma \odot \mathcal{B}^\Delta$ denote their rational amalgam. Then solvability of $(\Sigma \cup \Delta)$-constraints in $\mathcal{A}^\Sigma \odot \mathcal{B}^\Delta$ is decidable if if the positive universal-existential theory is decidable for both components $\mathcal{A}^\Sigma$ and $\mathcal{B}^\Delta$.*

Since existential quantification distributes over disjunction, the theorem may be slightly strengthened.

Theorem 6.6 *Let $\mathcal{A}^\Sigma$ and $\mathcal{B}^\Delta$ be two non-trivial rational SC-structures over disjoint signatures, let $\mathcal{A}^\Sigma \odot \mathcal{B}^\Delta$ denote their rational amalgam. Then the positive existential theory of $\mathcal{A}^\Sigma \odot \mathcal{B}^\Delta$ is decidable if the positive universal-existential theory is decidable for both components $\mathcal{A}^\Sigma$ and $\mathcal{B}^\Delta$.*

It is interesting to contrast this formulation with the corresponding combination result for free amalgamation (Theorem 22 of [BS95]) which needs *stronger* assumptions on the components: *Let $\mathcal{A}^\Sigma$ and $\mathcal{B}^\Delta$ be two strong SC-structures over disjoint signatures, let $\mathcal{A}^\Sigma \otimes \mathcal{B}^\Delta$ denote their free amalgam. Then the positive existential theory of $\mathcal{A}^\Sigma \otimes \mathcal{B}^\Delta$ is decidable if the positive theory is decidable for both components $\mathcal{A}^\Sigma$ and $\mathcal{B}^\Delta$.*

One application of Theorem 6.6 is the following

Corollary 6.7 *Rational amalgamated products $\mathcal{A}_1^{\Sigma_1} \odot \cdots \odot \mathcal{A}_k^{\Sigma_k}$ have decidable positive existential theory if the nontrivial components $\mathcal{A}_i^{\Sigma_i}$ are rational tree algebras, or nested, rational lists, or feature structures[11], or feature-structures with arity, for $i = 1, \ldots, k$, and if the signatures of the components are pairwise disjoint.*

Proof. For all these structures it has been shown that even the full positive theory is decidable, see [BS95]. □

In the rest of this section we prove Theorem 6.3 and Theorem 6.5.

[10] The existence of such an endomorphism is trivial if $x \notin Stab_{\mathcal{M}}^{\mathcal{A}}(a)$. In this case we may always take, e.g., the endomorphism $m = m_{x-a}$ of $\mathcal{M}$ that maps x to a and leaves all other atoms fixed. The situation of interest is the case where $x \in Stab_{\mathcal{M}}^{\mathcal{A}}(a)$ and $x \neq a$.

[11] As in Examples 3.4 we refer to [AP94], for specificity.

$$f(\mathit{unfold}(\mathcal{K}_1), \ldots, \mathit{unfold}(\mathcal{K}_n)).$$

It follows that *unfold* is a Σ-homomorphism. In the same way it follows that *unfold* is a Δ-homomorphism. It is then trivial to see that *fold* is $(\Sigma \cup \Delta)$-homomorphic, too. With the previous claims it follows that *fold* and *unfold* are $(\Sigma \cup \Delta)$-isomorphisms. This completes the proof of Theorem 5.2. □

6 Combination of Constraint Solvers

Our last aim is to show how constraint solvers for two component structures can be combined to a constraint solver for their rational amalgamated product. Constraint solvers, as considered here, are essentially algorithms that decide solvability of quantifier-free positive formulae in a given solution domain. We (mostly) disregard disjunction since its integration is a triviality.

Definition 6.1 Let Γ be a signature. A Γ-*constraint* is a conjunction of atomic Γ-formulae.

In order to decide solvability of a "mixed" $(\Sigma \cup \Delta)$-constraint in a rational amalgamated product $\mathcal{A}^\Sigma \odot \mathcal{B}^\Delta$ we shall decompose it into two pure constraints over the signatures Σ and Δ respectively. These output constraints are equipped with additional restrictions of a particular type.

Definition 6.2 An *A/N (atom/non-atom) declaration* for a constraint γ is a pair (U, W) such that $U \uplus W \subseteq Var(\gamma)$ is a disjoint union. Both U and W may be empty. A solution ν_A of a constraint γ in an SC-structure $(\mathcal{A}^\Sigma, \mathcal{M}, X)$ is called a *solution* of $\langle \gamma, U, W \rangle$ if ν_A assigns distinct atoms to the variables in U, and arbitrary non-atomic elements of A to the variables in W.

In order to avoid some ballast in proofs we shall assume that at least one of the two components is a *non-trivial* SC-structure, which means that it has at least one non-atomic element. We may now formulate our main result concerning combination of constraint solvers in the case of rational amalgamation.

Theorem 6.3 *Let $\mathcal{A}^\Sigma$ and $\mathcal{B}^\Delta$ be two non-collapsing SC-structures over disjoint signatures, let $\mathcal{A}^\Sigma \odot \mathcal{B}^\Delta$ denote their rational amalgam. Assume that at least one of the two components is a non-trivial SC-structure. Then solvability of $(\Sigma \cup \Delta)$-constraints in $\mathcal{A}^\Sigma \odot \mathcal{B}^\Delta$ is decidable if solvability of (Σ- resp. Δ-) constraints with A/N declarations is decidable for $\mathcal{A}^\Sigma$ and $\mathcal{B}^\Delta$.*

There seems to be no general way to characterize solvability of Γ-constraints with A/N declarations in purely logical terms. But for a restricted class of component structures—a class which is of particular interest in the context of rational amalgamation—a logical characterization of the problems that we have to solve in the two component structures can be given.

It is now simple to prove that that (m_1, n_1) is a simplifier for $\mathcal{K}$: assume that $m_1(o_1) = m_1(o_2)$ for $o_1, o_2 \in \mathcal{O}_A(\mathcal{K})$. Let o_i point in $\mathcal{K}$ to the subbraid $\mathcal{K}_i$ with root d_i, say. Then $op^A(\mathit{unfold}(\mathcal{K}_1)) = op^A(\mathit{unfold}(\mathcal{K}_2))$. Since the mapping op^A is injective, $\mathit{unfold}(\mathcal{K}_1) = \mathit{unfold}(\mathcal{K}_2)$. The first of the two equalities given above shows that $n_1(d_1) = \mathit{pur}^\Delta(\mathit{unfold}(\mathcal{K}_1)) = \mathit{pur}^\Delta(\mathit{unfold}(\mathcal{K}_2)) = n_1(d_2)$. With a symmetrical argument, as usual, it follows that (m_1, n_1) is a simplifier for $\mathcal{K}$.

In order to show that the step from $\mathcal{K}$ to $\mathit{unfold} \circ \mathit{fold}^*(\mathcal{K})$ can be described as a simplification of $\mathcal{K}$ it remains to prove that $\mathit{fold}^*(\mathit{unfold}(\mathcal{K}) = \mathcal{K}^{\langle m_1, n_1 \rangle}$. Without loss of generality, $\mathcal{K}$ has type A. Let a be the root of $\mathcal{K}$. The root of $\mathit{fold}^*(\mathit{unfold}(\mathcal{K}))$ is $\mathit{pur}^\Sigma(\mathit{unfold}(\mathcal{K})$. The root of $\mathcal{K}^{\langle m_1, n_1 \rangle}$ is $m_1(a)$. Both roots are identical, by the second of the two equalities given above. Assume that $o' \in \mathcal{O}_A(\mathit{fold}^*(\mathit{unfold}(\mathcal{K}))) \cap \mathcal{O}_A(\mathcal{K}^{\langle m_1, n_1 \rangle})$. Since $o \in \mathcal{O}_A(\mathcal{K}^{\langle m_1, n_1 \rangle})$ we know that o has the form $m_1(o)$, where $o \in \mathcal{O}_A(\mathcal{K})$. Assume that o points in $\mathcal{K}$ to the subbraid $\mathcal{K}_1$ with root d_1. Then $o' = m_1(o)$ points in $\mathcal{K}^{\langle m_1, n_1 \rangle}$ to the subbraid $\mathcal{K}_1^{\langle m_1, n_1 \rangle}$ with root $n_1(d_1)$. On the other hand $m_1(o) = op^A(\mathit{unfold}(\mathcal{K}_1))$, and this atom points in $\mathit{fold}^*(\mathit{unfold}(\mathcal{K}))$ to the subbraid with root $\mathit{pur}^\Delta(\mathit{unfold}(\mathcal{K}_1))$. The first of the two equalities given above shows that o is linked to the same element $n_1(d_1) = \mathit{pur}^\Delta(\mathit{unfold}(\mathcal{K}_1))$ in $\mathit{fold}^*(\mathit{unfold}(\mathcal{K}))$ and $\mathcal{K}^{\langle m_1, n_1 \rangle}$ respectively. The same holds of course for the open atoms in $\mathcal{O}_B(\mathit{fold}^*(\mathit{unfold}(\mathcal{K}))) \cap \mathcal{O}_B(\mathcal{K}^{\langle m_1, n_1 \rangle})$. Now Lemma 4.6 shows that $\mathit{fold}^*(\mathit{unfold}(\mathcal{K})) = \mathcal{K}^{\langle m_1, n_1 \rangle}$.

Summarizing, we have seen that the step from $\mathcal{K}$ to $\mathit{unfold} \circ \mathit{fold}^*(\mathcal{K})$ can be described as a simplification of $\mathcal{K}$, which proves Claim 3 as we have seen already. □

Once we know that both *fold* and *unfold* are bijections we may assume without loss of generality that op^A and op^B are the mappings that assigns to each rational tree $t \in R(\Sigma \cup \Delta, Z)$ the open atom $o_{[\mathit{fold}(t)]}$. It is easy to see that in this case $\mathit{fold}^* = \mathit{fold}$, hence fold^* and *unfold* are inverse bijections.

Let $\mathcal{K} \in R(\Sigma, X) \odot R(\Delta, Y)$ be a trivial braid, or a nontrivial braid of type A. The element $\mathit{unfold} \circ \mathit{pur}^\Sigma(\mathcal{K})$ is the root of $\mathit{unfold} \circ \mathit{fold}(\mathcal{K})$, i.e., the root of $\mathcal{K}$. Hence $\mathit{unfold} \circ \mathit{pur}^\Sigma(\mathcal{K}) = \mathit{root}_A(\mathcal{K})$, by the definition of root_A. Next assume that $\mathcal{K} \in R(\Sigma, X) \odot R(\Delta, Y)$ is a nontrivial braid of type B. Then $\mathit{unfold} \circ \mathit{pur}^\Sigma(\mathcal{K}) = op^A(\mathit{unfold}(\mathcal{K})) = o_{[\mathit{fold}(\mathit{unfold}(\mathcal{K}))]} = o_{[\mathcal{K}]} = \mathit{root}_A(\mathcal{K})$. We have seen that $\mathit{root}_A = \mathit{unfold} \circ \mathit{pur}^\Sigma$. Similarly it follows that $\mathit{root}_B = \mathit{unfold} \circ \mathit{pur}^\Delta$. Hence $\mathit{pur}^\Sigma = \mathit{fold} \circ \mathit{root}_A$ and $\mathit{pur}^\Delta = \mathit{fold} \circ \mathit{root}_B$. Let $f \in \Sigma$ be n-ary, let $\mathcal{K}_1, \ldots, \mathcal{K}_n \in R(\Sigma, X) \odot R(\Delta, Y)$. Then

$$
\begin{aligned}
&\mathit{unfold}(f_{\mathcal{A} \odot \mathcal{B}}(\mathcal{K}_1, \ldots, \mathcal{K}_n)) = \\
&\mathit{unfold}(\mathit{root}_A^{-1}(f_{\mathcal{A}_*}(\mathit{root}_A(\mathcal{K}_1), \ldots, \mathit{root}_A(\mathcal{K}_n)))) = \\
&\mathit{unfold}(\mathit{root}_A^{-1}(f(\mathit{pur}^\Sigma(\mathit{unfold}(\mathcal{K}_1)), \ldots, \mathit{pur}^\Sigma(\mathit{unfold}(\mathcal{K}_n))))) = \\
&\mathit{unfold}(\mathit{root}_A^{-1}(\mathit{pur}^\Sigma(f(\mathit{unfold}(\mathcal{K}_1), \ldots, \mathit{unfold}(\mathcal{K}_n))))) = \\
&\mathit{unfold}(\mathit{root}_A^{-1}(\mathit{root}_A(\mathit{fold}(f(\mathit{unfold}(\mathcal{K}_1), \ldots, \mathit{unfold}(\mathcal{K}_n)))))) =
\end{aligned}
$$

4. $\pi_A := \{\langle op^A(t'_i), pur^\Delta(t'_i)\rangle \mid i = 1, \ldots, l\}$,

5. $\pi_B := \{\langle op^B(t_i), pur^\Sigma(t_i)\rangle \mid i = 1, \ldots, k\}$.

If t is a Δ-tree, then $fold^*(t)$ is defined symmetrically, using op^B and pur^Δ instead of op^A and pur^Σ. The mapping

$$fold : R(\Sigma \cup \Delta, Z) \to R(\Sigma, X) \odot R(\Delta, Y)$$

assigns to each rational tree t the unique element of $R(\Sigma, X) \odot R(\Delta, Y)$ that represents the standard normal form of the braid $fold^*(t)$.

Claim 2 *fold is injective and unfold is surjective.*

Proof of Claim 2. Clearly we obtain t back again by unfolding $fold^*(t)$. Moreover, since no pending atoms can occur when we simplify $fold^*(t)$, it is easy to see that the result of the unfolding process is not influenced by simplification. Hence $t = unfold(fold(t))$, for each $t \in R(\Sigma \cup \Delta, Z)$, which implies that *fold* is injective and *unfold* is surjective. □

Claim 3 *unfold is injective and fold is surjective.*

Proof of Claim 3. Unfortunately, the proof of this claim is not simple. We shall proceed as follows. Let $\mathcal{K} \in R(\Sigma, X) \odot R(\Delta, Y)$ be an arbitrary braid in standard normal form. We shall show that the step from $\mathcal{K}$ to $unfold \circ fold^*(\mathcal{K})$ can be described as a simplification of $\mathcal{K}$. Since $unfold \circ fold(\mathcal{K})$ is obtained from $unfold \circ fold^*(\mathcal{K})$ by an additional (standard) simplification step, this shows that $\mathcal{K}$ and $unfold \circ fold(\mathcal{K})$ are equivalent braids. But both braids are in standard normal form. By Lemma 4.39, $\mathcal{K} = unfold \circ fold(\mathcal{K})$. Since $\mathcal{K}$ was arbitrary, $unfold \circ fold$ is the identity on $R(\Sigma, X) \odot R(\Delta, Y)$, which implies that *unfold* is injective and *fold* is surjective.

In order to show that $unfold \circ fold$ is a simplification, we define, given a braid $\mathcal{K} \in R(\Sigma, X) \odot R(\Delta, Y)$, an admissible pair of endomorphisms (m_1, n_1) as follows. Assume that $o' \in \mathcal{O}_A(\mathcal{K})$ points in $\mathcal{K}$ to the subbraid $\mathcal{K}'$ with root d, say. Then $m_1(o') := op^A(unfold(\mathcal{K}'))$. Similarly, let $o'' \in \mathcal{O}_B(\mathcal{K})$ point in $\mathcal{K}$ to the subbraid $\mathcal{K}''$ with root c, say. Then $n_1(o'') := op^A(unfold(\mathcal{K}''))$. We may extend these partial mappings to an admissible pair of endomorphisms (m_1, n_1). Now note that

$$\begin{aligned} n_1(d) &= pur^\Delta(unfold(\mathcal{K}')), \\ m_1(c) &= pur^\Sigma(unfold(\mathcal{K}'')). \end{aligned}$$

To see the first equality, recall that $n_1(d)$ is obtained by replacing each open atom o_r of the root d of $\mathcal{K}'$—pointing in $\mathcal{K}'$, say, to the subbraid $\mathcal{K}_r$—by the open atom $op^B(unfold(\mathcal{K}_r))$. But when we unfold $\mathcal{K}'$, then the maximal Σ-subtrees are exactly the trees of the form $unfold(\mathcal{K}_r)$. Purification replaces these subtrees by the open atoms $op^B(unfold(\mathcal{K}_r))$, which shows that $pur^\Delta(unfold(\mathcal{K}')) = n_1(d)$. The second equality follows in the same way.

(Later, we shall see that *unfold* and *fold* are inverse bijections. Eventually we shall show that both mappings are $(\Sigma \cup \Delta)$-homomorphisms.) We have to introduce some terminology. A tree $t \in R(\Sigma \cup \Delta, Z)$ is called a Σ-tree (Δ-tree) if the topmost function symbol of t belongs to Σ (resp. Δ). Suppose that we follow a path of the rational tree $t \in R(\Sigma \cup \Delta, Z)$, starting from the root. Each node of the path defines an occurrence of a unique subtree t' of t in the obvious way. Such an occurrence is called *relevant* if the topmost function symbol of the subtree belongs to another signature than the label of the predecessor node on the path. A subtree t' of t is called relevant if t' has at least one relevant occurrence in t.

The following claim gives a first connection between a braid and the rational tree that is obtained by unfolding the braid.

Claim 1 *For each braid $\mathcal{K} \in R(\Sigma, X) \odot R(\Delta, Y)$ the relevant subtrees of $unfold(\mathcal{K})$ are exactly the subtrees of the form $unfold(\mathcal{K}_i)$, where $\mathcal{K}_i$ is a subbraid of $\mathcal{K}$.*

Proof of Claim 1. If $\mathcal{K}$ is a trivial braid, then the root of $\mathcal{K}$ is a bottom atom z. Unfolding $\mathcal{K}$ yields the rational tree "z", which means that the claim holds trivially. If $\mathcal{K}$ is nontrivial, then the claim follows directly from the fact that all elements of $\mathcal{K}$ are non-atomic, by Definition 4.3. □

One preparation is needed before we can give the definition of *fold*. Let op^A (resp. op^B) be a 1-1 mapping that assigns to each Σ-tree (resp. Δ-tree) $t \in R(\Sigma \cup \Delta, Z)$ an open atom $op^A(t) \in \mathcal{O}_A$ (resp. $op^B(t) \in \mathcal{O}_B$). These mappings can be used to define "purifying" 1-1 functions

$$\begin{aligned} pur^\Sigma : R(\Sigma \cup \Delta, Z) &\rightarrow R(\Sigma, X) \\ pur^\Delta : R(\Sigma \cup \Delta, Z) &\rightarrow R(\Delta, Y) \end{aligned}$$

as follows. Both pur^Σ and pur^Δ fix all atoms $z \in Z$. Moreover, $pur^\Sigma(t) := op^A(t)$ for each Δ-tree t, and conversely $pur^\Delta(t) := op^B(t)$ for each Σ-tree t. If t is a Σ-tree, then $pur^\Sigma(t)$ is obtained from t by replacing all the outermost (= topmost) Δ-subtrees t_1 of t by $op^A(t_1)$. Symmetrically, if t is a Δ-tree, then $pur^\Delta(t)$ is obtained from t by replacing all the outermost Σ-subtrees t_1 of t by $op^B(t_1)$.

We may now obtain a braid representation $fold^*(t)$ of a rational tree $t \in R(\Sigma \cup \Delta, Z)$ as follows. If t is an atom $z \in Z$, then $fold^*(t)$ is the trivial braid with root z. In the other case, assume first that t is a Σ-tree. Then $fold^*(t)$ is the braid $\langle a, C, D, \pi_A, \pi_B \rangle$ with the following components:

1. $a := pur^\Sigma(t)$,
2. $C = \{pur^\Sigma(t_1), \ldots, pur^\Sigma(t_k)\} \cup \{pur^\Sigma(t)\}$, where $t_1, \ldots, t_k$ are the relevant Σ-subtrees of t,
3. $D = \{pur^\Delta(t'_1), \ldots, pur^\Delta(t'_l)\}$, where $t'_1, \ldots, t'_l$ are the relevant Δ-subtrees of t,

5.3 Proof of Theorem 5.2

In this subsection we shall prove Theorem 5.2. We consider the special situation where $\mathcal{A}^\Sigma = R(\Sigma, X)$ (resp. $\mathcal{B}^\Delta = R(\Delta, Y)$) is the non-ground algebra of rational trees for signature Σ (resp. Δ) [9]. From the introduction of Section 4 recall that the set of bottom atoms $Z = X \cap Y$ is infinite and $X = Z \uplus \mathcal{O}_B, Y = Z \uplus \mathcal{O}_B$.

The two structures $\mathcal{A}^\Sigma$ and $\mathcal{B}^\Delta$ are SC-structures of the form $(\mathcal{A}^\Sigma, \mathcal{M}, X)$ and $(\mathcal{B}^\Delta, \mathcal{N}, Y)$ where $\mathcal{M} := End^\Sigma_\mathcal{A}$ and $\mathcal{N} := End^\Delta_\mathcal{B}$. For a rational tree $t \in R(\Sigma, X)$, the set $Stab^\mathcal{A}_\mathcal{M}(t)$ is just the set of all atoms that label a leaf of t. If m is an endomorphism, then $m(t)$ is obtained from t just by replacing each atomic leaf x in t by the subtree $m(x)$. Consequently, if $m \in \mathcal{M}$ is an admissible endomorphism of $R(\Sigma, X)$, then $Stab^\mathcal{A}_\mathcal{M}(m(t)) = \{m(x) \mid x \in Stab^\mathcal{A}_\mathcal{M}(t)\}$. This equality, and the corresponding equality for $R(\Delta, Y)$ show that there are no pending atoms when we simplify braids over $R(\Sigma, X)$ and $R(\Delta, Y)$. It follows easily that $\mathcal{K}^{(m,n)} = \mathcal{K}^{\langle m,n \rangle}$, for all braids $\mathcal{K}$ over $R(\Sigma, X)$ and $R(\Delta, Y)$ and all admissible pairs of endomorphisms (m, n).

In the first step of the proof of Theorem 5.2 we show how braids in $R(\Sigma, Y) \odot R(\Delta, Z)$ can be mapped naturally to rational trees in $R(\Sigma \cup \Delta, Z)$. Let $\mathcal{K} = \langle a, C, D, \pi_A, \pi_B \rangle \in R(\Sigma, X) \odot R(\Delta, Y)$. Then $\pi := \pi_A \cup \pi_B$ can be considered as a mapping $R(\Sigma \cup \Delta, X \cup Y) \to R(\Sigma \cup \Delta, X \cup Y)$ that replaces open atom leaves by rational trees, similar to a substitution. The process where we apply π to the root a of $\mathcal{K}$ in an iterative way, obtaining a sequence $a_0 = a, a_1 = \pi(a), \ldots, a_{n+1} = \pi(a_n), \ldots$ is called *unfolding* of $\mathcal{K}$. Since, by Definition 4.3, π replaces open atoms by non-atomic elements, the above sequence converges (w.r.t. the usual metrics on infinite trees) to a unique limit tree $t_\mathcal{K}$. Obviously $t_\mathcal{K} \in R(\Sigma \cup \Delta, Z)$. We define the mapping

$$unfold : R(\Sigma, X) \odot R(\Delta, Y) \to R(\Sigma \cup \Delta, Z) : \mathcal{K} \mapsto t_\mathcal{K}.$$

Example 5.9 The following figure gives the rational tree that is obtained by unfolding the braid from Example 4.17:

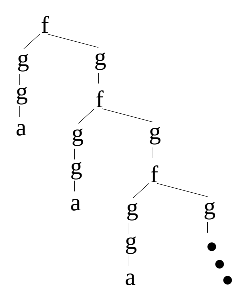

In the second step of the proof we shall define a mapping

$$fold : R(\Sigma \cup \Delta, Z) \to R(\Sigma, Y) \odot R(\Delta, Z).$$

[9] See [Co84] for the definition of the algebra of rational trees over a given signature.

length $n \geq 2$, such that $e_1 = e_n$ and every element e_i is directly linked to e_{i+1} via $\pi' := \pi'_A \cup \pi'_B$, for $i = 1, \ldots, n-1$. An element a of $\mathcal{K}$ is called interesting if its image $m(a)$ resp. $n(a)$ occurs in the sequence $e_1 \ldots, e_n$. An element b' of $\mathcal{K}$ is called a daughter of an element b of $\mathcal{K}$ if b' is directly linked to b via $\pi := \pi_A \cup \pi_B$.

Since $\mathcal{K}$ is acyclic, there has to be a interesting element $a \in C \cup D$ such that no daughter of a is interesting. Without loss of generality we assume that $a \in C$. Hence $m(a)$ occurs in $e_1 \ldots, e_n$, say, as element e_i (we may assume that $i > 1$). Since e_{i-1} is directly linked to $m(a) = e_i$ in $\mathcal{K}^{(m,n)}$, there exists a link $\langle o, e_{i-1} \rangle \in \pi'_A$ where $o \in \mathcal{O}_A(m(a))$. Let $\mathcal{O}_A(a) = \{o_1, \ldots, o_k\}$ and $b_i := \pi_A(o_i)$, for $i = 1, \ldots, k$. Thus $\{b_1, \ldots, b_k\}$ is the set of daughters of a in $\mathcal{K}$. From Lemma 3.10 it follows that o has the form $o = m(o_i)$, for some $1 \leq i \leq k$. Hence, since π'_A is a function, it follows from Definition 4.16 that $\langle o, e_{i-1} \rangle = \langle m(o_i), n(b_i) \rangle$ and $n(b_i) = e_{i-1}$. But this implies that the daughter b_i of a is interesting, which contradicts our choice of a. □

Theorem 5.8 *The set of all acyclic braids of $A \odot B$ forms a substructure of $\mathcal{A}^\Sigma \odot \mathcal{B}^\Delta$.*

Proof. Let $f \in \Sigma$ be an n-ary function symbol, let $\mathcal{K}_1, \ldots, \mathcal{K}_n$ be acyclic elements of $\mathcal{A}^\Sigma \odot \mathcal{B}^\Delta$. We have to show that the braid

$$f_{\mathcal{A}\odot\mathcal{B}}(\mathcal{K}_1, \ldots, \mathcal{K}_n) = root_A^{-1}(f_{\mathcal{A}_*}(root_A(\mathcal{K}_1), \ldots, root_A(\mathcal{K}_n)))$$

is acyclic. The elements of $\mathcal{O}_A(\{root_A(\mathcal{K}_1), \ldots, root_A(\mathcal{K}_n)\}$ represent—in the sense of Definition 4.40—acyclic subbraids. By Lemma 3.9,

$$\mathcal{O}_A(f_{\mathcal{A}_*}(root_A(\mathcal{K}_1), \ldots, root_A(\mathcal{K}_n)) \subseteq \mathcal{O}_A(\{root_A(\mathcal{K}_1), \ldots, root_A(\mathcal{K}_n)\}$$

and the open atoms in $a^* := f_{\mathcal{A}_*}(root_A(\mathcal{K}_1), \ldots, root_A(\mathcal{K}_n))$ represent acyclic subbraids. If $a^* \in \mathcal{A}_* \setminus \mathcal{O}_A^*$, then $\mathcal{K} := root_A^{-1}(a^*)$ is the unique braid in standard normal form with root a^*. Since all open atoms of the root of $\mathcal{K}$ represent acyclic subbraids, $\mathcal{K}$ itself is acyclic. In the other case, if $a^* = o \in \mathcal{O}_A(\{root_A(\mathcal{K}_1), \ldots, root_A(\mathcal{K}_n)\}$ is an atom, then it represents an acylic braid in standard normal form $\mathcal{K}$ of type B. But $\mathcal{K} := root_A^{-1}(a^*)$. We have seen that the set of all acyclic braids represents a Σ-substructure of $\mathcal{A}^\Sigma \odot \mathcal{B}^\Delta$. Symmetrically it follows that this set represents a Δ-substructure of $\mathcal{A}^\Sigma \odot \mathcal{B}^\Delta$. □

The proof that the set of all acyclic braids, considered as a $(\Sigma \cup \Delta)$-structure, is isomorphic to the free amalgamated product of $\mathcal{A}^\Sigma$ and $\mathcal{B}^\Delta$ (as introduced in [BS95]) cannot be given here. For readers that are familiar with the latter notion we mention that induction on the depth of an acyclic braid may be used to construct the factorising homomorphisms that characterize the free amalgamated product up to isomorphism.

2. Let $p \in \Sigma$ be an n-ary predicate symbol, let $\mathcal{K}_1, \ldots, \mathcal{K}_n \in A \odot B$. We define $\mathcal{A}^\Sigma \odot \mathcal{B}^\Delta \models p(\mathcal{K}_1, \ldots, \mathcal{K}_n)$ iff $\mathcal{A}^\Sigma_* \models p(root_A(\mathcal{K}_1), \ldots, root_A(\mathcal{K}_n))$.

The interpretation of the function symbols $g \in \Delta$ and the predicate symbols $q \in \Delta$ in $\mathcal{A}^\Sigma \odot \mathcal{B}^\Delta$ is defined symmetrically, using $root_B$.

Theorem 5.5 *As a Σ-structure, $\mathcal{A}^\Sigma \odot \mathcal{B}^\Delta, \mathcal{A}^\Sigma$ and $\mathcal{A}^\Sigma_*$ are isomorphic, and $root_A : \mathcal{A}^\Sigma \odot \mathcal{B}^\Delta \to \mathcal{A}^\Sigma_*$ is a Σ-isomorphism. As a Δ-structure, $\mathcal{A}^\Sigma \odot \mathcal{B}^\Delta, \mathcal{B}^\Delta$, and $\mathcal{B}^\Delta_*$ are isomorphic, and $root_B : \mathcal{A}^\Sigma \odot \mathcal{B}^\Delta \to \mathcal{B}^\Delta_*$ is a Δ-isomorphism.*

Proof. Recall that A^Σ and $\mathcal{A}^\Sigma_*$ are isomorphic, and similarly for $\mathcal{B}^\Delta$ and $\mathcal{B}^\Delta_*$. Lemma 5.3 and Definition 5.4 imply that $root_A : \mathcal{A}^\Sigma \odot \mathcal{B}^\Delta \to \mathcal{A}^\Sigma_*$ is a Σ-isomorphism and $root_B : \mathcal{A}^\Sigma \odot \mathcal{B}^\Delta \to \mathcal{B}^\Delta_*$ is a Δ-isomorphism. □

Theorem 5.5 makes clear that rational amalgamation is *not* a construction that can be used, say, to construct a rational tree algebra for a given signature Σ out of the finite tree algebra for Σ. Even if $\mathcal{B}^\Delta$ consists of atoms only, the rational amalgam $\mathcal{A}^\Sigma \odot \mathcal{B}^\Delta$, considered as a Σ-structure, is isomorphic to $\mathcal{A}^\Sigma$.

5.2 Free amalgamation and rational amalgamation

In this subsection we shall sketch a proof for Theorem 5.1. We define the notion of an *acyclic* braid and show that the set of all acyclic braids in standard normal form is a substructure of the rational amalgamated product. It is possible to prove that the free amalgamated product of the two component structures is isomorphic to this substructure.[7]

Definition 5.6 A prebraid $\mathcal{K} = \langle a, C, D, \pi_A, \pi_B \rangle$ is called *acyclic* if there is no sequence $e_1, e_2 \ldots, e_n$ of elements in $C \cup D$, of length $n \geq 2$, such that $e_1 = e_n$ and every element e_i is directly linked[8] via $\pi = \pi_A \cup \pi_B$ to e_{i+1}, for $i = 1, \ldots, n-1$. If $\mathcal{K}$ is acyclic, the *depth* of $\mathcal{K}$ is the largest number n such that there is a sequence $e_1, \ldots, e_n$ of elements of $\mathcal{K}$ where each element e_i is directly linked to e_{i+1} via π, for $i = 1, \ldots, n-1$.

Lemma 5.7 *Let (m, n) be a simplifier for the acyclic braid $\mathcal{K}$. Then the braid image $\mathcal{K}^{\langle m,n \rangle}$ is an acyclic braid.*

Proof. We may assume that $\mathcal{K} = \langle a, C, D, \pi_A, \pi_B \rangle$ is of type A. Let $\mathcal{K}^{(m,n)} = \langle m(a), C', D', \pi'_A, \pi'_B \rangle$. We show that the prebraid $\mathcal{K}^{(m,n)}$ is acyclic. Assume, to get a contradiction, that there is a sequence $e_1 \ldots, e_n$ of elements in $C' \cup D'$, of

[7] With the actual methods, the free amalgamated product can only be built for *strong* SC-structures over disjoint signatures. Hence we have to assume that the two components are strong and non-collapsing.

[8] compare Definition 4.2.

that interpret the symbols of the mixed signature $\Sigma \cup \Delta$. With this step, the definition of the rational amalgamated product is complete. In the following two subsections we add some evidence for the naturalness of rational amalgamation. First we consider the case where the two components are strong non-collapsing SC-structures over disjoint signatures. This is the situation where we can build both the free amalgam and the rational amalgam with our actual methods.

Theorem 5.1 *The free amalgamated product is–modulo isomorphism—a substructure of the rational amalgamated product.*

This shows that there are interesting relationships between distinct amalgamation constructions.

Eventually we consider a particular class of amalgamation components. We shall show

Theorem 5.2 *The rational amalgamated product of two algebras of rational trees over disjoint signatures is isomorphic to the algebra of rational trees over the combined signature.*

This shows that our general construction, complicated as it might appear, yields the expected result when we consider more concrete situations.

5.1 Functions and relations

Given the underlying domain of the rational amalgam of $\mathcal{A}^\Sigma$ and $\mathcal{B}^\Delta$ as constructed above, there is now a perfectly natural way to introduce functions and relations that interpret the symbols of the mixed signature $\Sigma \cup \Delta$. Consider the functions $root_A : A \odot B \to A_*$ and $root_B : A \odot B \to B_*$:

$$root_A(\mathcal{K}) \;:=\; \begin{cases} \text{the root of } \mathcal{K} & \text{if } \mathcal{K} \text{ is trivial or has type } A \\ o_{[\mathcal{K}]} \in \mathcal{O}_A^* & \text{if } \mathcal{K} \text{ is non-trivial and has type } B. \end{cases}$$

$$root_B(\mathcal{K}) \;:=\; \begin{cases} \text{the root of } \mathcal{K} & \text{if } \mathcal{K} \text{ is trivial or has type } B \\ o_{[\mathcal{K}]} \in \mathcal{O}_B^* & \text{if } \mathcal{K} \text{ is non-trivial and has type } A. \end{cases}$$

As a direct consequence of Lemma 4.41 we obtain

Lemma 5.3 *The functions $root_A$ and $root_B$ are bijections.*

Here is now the definition of the rational amalgamated product.

Definition 5.4 The *rational amalgamated product $\mathcal{A}^\Sigma \odot \mathcal{B}^\Delta$ of $\mathcal{A}^\Sigma$ and $\mathcal{B}^\Delta$* is the following $(\Sigma \cup \Delta)$-structure with carrier $A \odot B$:

1. Let $f \in \Sigma$ be an n-ary function symbol, let $\mathcal{K}_1, \ldots, \mathcal{K}_n \in A \odot B$. We define $f_{\mathcal{A}\odot\mathcal{B}}(\mathcal{K}_1, \ldots, \mathcal{K}_n) = root_A^{-1}(f_{\mathcal{A}_*}(root_A(\mathcal{K}_1), \ldots, root_A(\mathcal{K}_n)))$.

Definition 4.38 The process where we apply to a given (pre)braid $\mathcal{K}$ the simplifier (m,n) that maps each open atom $o \in \mathcal{O}(\mathcal{K})$, pointing in $\mathcal{K}$ to the subbraid $\mathcal{K}'$, to the open atom $o_{[\mathcal{K}']} \in \mathcal{O}_A^* \cup \mathcal{O}_B^*$ will be called *standard simplification* of $\mathcal{K}$. The prebraid $\mathcal{K}^{(m,n)}$ (braid $\mathcal{K}^{\langle m,n\rangle}$) will be called *the standard (braid) normal form* of $\mathcal{K}$.

Obviously all subbraids of a prebraid in standard normal form are again in standard normal form.

Lemma 4.39 *For each braid $\mathcal{K}$ there exists exactly one braid $\mathcal{K}'$ in standard normal form such that $\mathcal{K}$ and $\mathcal{K}'$ are equivalent.*

Proof. We have seen that standard normalization yields a braid in standard normal form that is equivalent to $\mathcal{K}$. If $\mathcal{K}'$ and $\mathcal{K}''$ are braids in standard normal form that are equivalent to $\mathcal{K}$, then $\mathcal{K}'$ and $\mathcal{K}''$ are irreducible (Lemma 4.36) and variants, by Lemma 4.33. It follows that there exists an admissible pair of automorphisms (m,n) such that $\mathcal{K}'' = \mathcal{K}'^{\langle m,n\rangle}$. Let $o \in \mathcal{O}_A(\mathcal{K}')$ point in $\mathcal{K}'$ to $\mathcal{K}_1$. Then $m(o)$ points in $\mathcal{K}''$ to $\mathcal{K}_1^{\langle m,n\rangle}$ and $\mathcal{K}_1$ and $\mathcal{K}_1^{\langle m,n\rangle}$ are equivalent. Since $\mathcal{K}'$ and $\mathcal{K}''$ are in standard normal form, $o = o_{[\mathcal{K}_1]} = o_{[\mathcal{K}_1^{\langle m,n\rangle}]} = m(o)$. Hence m coincides on the elements of $\mathcal{K}'$ of type A with identity. A symmetrical argument shows that n coincides on the elements of $\mathcal{K}'$ of type B with identity. □

Definition 4.40 Let $o \in \mathcal{O}_A^* \cup \mathcal{O}_B^*$. We say that o *represents* the unique braid $\mathcal{K}$ in standard normal form such that $o = o_{[\mathcal{K}]}$.

Lemma 4.41 *Given $e \in (A_* \cup B_*) \setminus (\mathcal{O}_A^* \cup \mathcal{O}_B^*)$ there exists a unique braid $\mathcal{K} \in A \odot B$ such that e is the root of $\mathcal{K}$.*

Proof. Let $e \in (A_* \cup B_*) \setminus (\mathcal{O}_A^* \cup \mathcal{O}_B^*)$. We may assume that $e \in A_* \setminus \mathcal{O}_A^*$. Let $\mathcal{O}_A(e) = \{o_1, \ldots, o_n\}$, and let o_i represent the braid in standard normal form $\mathcal{K}_i = \langle e_i, C_i, D_i, \pi_A^i, \pi_B^i\rangle$. Let $C := \bigcup_{i=1}^n C_i \cup \{e\}$, $D := \bigcup_{i=1}^n D_i$, $\pi_A := \bigcup_{i=1}^n \pi_A^i \cup \{\langle o_{[\mathcal{K}_i]}, e_i\rangle \mid i = 1,\ldots,n\}$, and $\pi_B := \bigcup_{i=1}^n \pi_B^i$. Then $\mathcal{K} = \langle e, C, D, \pi_A, \pi_B\rangle \in A \odot B$ has root e.

Conversely, let $\mathcal{K} = \langle e, C, D, \pi_A, \pi_B\rangle \in A \odot B$. Since each open atom o_i in $\mathcal{O}_A(e)$ represents a unique braid $\mathcal{K}_i$ to which it points in $\mathcal{K}$, the structure of $\mathcal{K}$ is completely determined by e. □

5 The rational amalgamated product

In this section we shall complete the construction of the rational amalgamated product, and we shall provide some evidence for the naturalness of this construction. In the first subsection we introduce functions and relations on $A \odot B$

Proof. Let $\mathcal{K} = \langle a, C, D, \pi_A, \pi_B \rangle$ be a prebraid in standard normal form. Assume, to get a contradiction, that (m, n) is a strict simplifier for $\mathcal{K}$. Without loss of generality we may assume that $m(o_1) = m(o_2)$ for distinct open atoms $o_1, o_2 \in \mathcal{O}_A(C)$. Let $d_i := \pi_A(o_i)$, and let $\mathcal{K}_i$ denote the subbraid of $\mathcal{K}$ with root d_i, for $i = 1, 2$. Since (m, n) is a simplifier, $n(d_1) = n(d_2)$. By Lemma 4.15, (m, n) is a simplifier for $\mathcal{K}_1$ and $\mathcal{K}_2$. By Lemma 4.18, $\mathcal{K}_i^{\langle m,n \rangle}$ is the unique subbraid of $\mathcal{K}^{(m,n)}$ with root $n(d_i)$, for $i = 1, 2$. Since $n(d_1) = n(d_2)$, also $\mathcal{K}_1^{\langle m,n \rangle} = \mathcal{K}_2^{\langle m,n \rangle}$ which implies that $\mathcal{K}_1$ and $\mathcal{K}_2$ are equivalent. Since $\mathcal{K}$ is in standard normal form it follows that $o_1 = o_{[\mathcal{K}_1]} = o_{[\mathcal{K}_2]} = o_2$, which contradicts our assumption. □

Proposition 4.37 *Let $\mathcal{K}$ be a prebraid. Let (m, n) denote the admissible pair of endomorphisms that maps each $o \in \mathcal{O}_A(\mathcal{K}) \cup \mathcal{O}_B(\mathcal{K})$ to $o_{[\mathcal{K}']}$ where $\mathcal{K}'$ is the unique subbraid of $\mathcal{K}$ such that o points to $\mathcal{K}'$. Then (m, n) is a simplifier for $\mathcal{K}$ and $\mathcal{K}^{(m,n)}$ is in standard normal form.*

Proof. Let (m_1, n_1) be a simplifier for $\mathcal{K} = \langle a, C, D, \pi_A, \pi_B \rangle$ such that $\mathcal{K}_1 := \mathcal{K}^{(m_1,n_1)}$ is irreducible. If m_1 identifies the open atoms $o, o' \in \mathcal{O}_A(\mathcal{C})$, then n_1 identifies $d := \pi_A(o)$ and $d' := \pi_A(o')$. It follows that o and o' point in $\mathcal{K}$ to subbraids that receive the same braid image under the simplification (m_1, n_1). Hence these subbraids are equivalent, which implies that $m(o) = m(o')$. It follows that the mapping $m_2 : m_1(o) \mapsto m(o)$ $(o \in \mathcal{O}_A(C))$ is well-defined. Symmetrically it follows that the mapping $n_2 : n_1(o) \mapsto n(o)$ $(o \in \mathcal{O}_B(D))$ is well-defined. Both mappings can be extended to admissible endomorphisms for which we shall use the same symbols.

Obviously $m_1 \circ m_2$ (resp. $n_1 \circ n_2$) and m (resp. n) coincide on $\mathcal{O}_A(C)$ (resp. $\mathcal{O}_B(D)$). Hence $m_1 \circ m_2$ (resp. $n_1 \circ n_2$) and m (resp. n) coincide on C (resp. D), by Lemma 4.1.

We shall now show that (m, n) is a simplifier for $\mathcal{K}$. Assume that $m(o) = m(o')$, for $o, o' \in \mathcal{O}_A(C)$. This means, by the definition of m, that o and o' point in $\mathcal{K}$ to equivalent subbraids $\mathcal{K}'$ and $\mathcal{K}''$, with roots $d := \pi_A(o)$ and $d' := \pi_A(o')$. We want to show that o and o' are already identified by m_1. Let $\mathcal{K}_1'$ and $\mathcal{K}_1''$ denote the braid images of $\mathcal{K}'$ and $\mathcal{K}''$ under the simplification (m_1, n_1) respectively. Then $\mathcal{K}'$ and $\mathcal{K}_1'$ are equivalent, and similarly for $\mathcal{K}''$ and $\mathcal{K}_1''$. Since $\mathcal{K}'$ and $\mathcal{K}''$ are equivalent, this implies that $\mathcal{K}_1'$ and $\mathcal{K}_1''$ are equivalent. But $\mathcal{K}_1'$ and $\mathcal{K}_1''$ are subbraids of the irreducible prebraid $\mathcal{K}_1$, by Lemma 4.18. Part (b) of Lemma 4.27 shows that $\mathcal{K}_1'$ and $\mathcal{K}_1''$ are irreducible. Since they are equivalent, both are variants, by Lemma 4.33. Part (c) of Lemma 4.27 shows that $\mathcal{K}_1' = \mathcal{K}_1''$. The root of $\mathcal{K}_1'$ is $n_1(d)$, and the root of $\mathcal{K}_1''$ is $n_1(d')$. Hence $n_1(d) = n_1(d')$ and $n(d) = n_2(n_1(d)) = n_2(n_1(d')) = n(d')$.

Symmetrically it follows that $n(o) = n(o')$ always implies that $m(\pi_B(o)) = m(\pi_B(o'))$, for all $o, o' \in \mathcal{O}_B(D)$. This shows that (m, n) is in fact a simplifier for $\mathcal{K}$. Obviously $\mathcal{K}^{(m,n)}$ is in standard normal form. □

Since two braids that are variants are obviously equivalent it is easy to see that we get in fact an equivalence relation. Let us also mention the following simple consequence of Theorem 4.31:

Lemma 4.33 *If two irreducible braids are equivalent, they are variants.*

4.4 Standard normalization

In order to define the underlying domain of the rational amalgam we shall now introduce a standard normal form for each braid. Let $\mathcal{O}_A^*$ be a subset of the set $\mathcal{O}_A$ of open atoms of A^Σ that has the same cardinality as the set of all equivalence classes of non-trivial[5] braids of type B. Similarly, let $\mathcal{O}_B^*$ be a subset of the set $\mathcal{O}_B$ of open atoms of B^Δ that has the same cardinality as the set of all equivalence classes of non-trivial braids of type A. Let $\mathcal{A}_*^\Sigma := SH_{\mathcal{M}}^{\mathcal{A}}(Z \cup \mathcal{O}_A^*)$, and let $\mathcal{B}_*^\Delta := SH_{\mathcal{N}}^{\mathcal{B}}(Z \cup \mathcal{O}_B^*)$. Lemma 10 of [BS95] shows

Lemma 4.34 *Every bijection between $Z \cup \mathcal{O}_A^*$ and $Z \cup \mathcal{O}_A$ extends to a Σ-isomorphism between $\mathcal{A}_*^\Sigma$ and $\mathcal{A}^\Sigma$. Similarly every bijection between $Z \cup \mathcal{O}_B^*$ and $Z \cup \mathcal{O}_B$ extends to a Δ-isomorphism between $\mathcal{B}_*^\Delta$ and $\mathcal{B}^\Delta$.*

We may now enumerate the elements of $\mathcal{O}_A^*$ and of $\mathcal{O}_B^*$ in the form

$$\begin{aligned}\mathcal{O}_A^* &= \{o_{[\mathcal{K}]} \mid \mathcal{K} \text{ is a nontrivial braid of type } B\},\\ \mathcal{O}_B^* &= \{o_{[\mathcal{K}]} \mid \mathcal{K} \text{ is a nontrivial braid of type } A\}.\end{aligned}$$

This means that $[\mathcal{K}] \mapsto o_{[\mathcal{K}]}$ establishes a bijection between the set of all equivalence classes of non-trivial braids of type A (B) and $\mathcal{O}_B^*$ ($\mathcal{O}_A^*$).

Let $\mathcal{K} = \langle a, C, D, \pi_A, \pi_B \rangle$ be a prebraid. For each open atom $o \in \mathcal{O}_A(C)$ ($o \in \mathcal{O}_B(D)$) we say that *o points in $\mathcal{K}$ to $\mathcal{K}'$* iff $\mathcal{K}'$ is the unique subbraid of $\mathcal{K}$ with root $\pi_A(o)$ ($\pi_B(o)$)[6].

Definition 4.35 A prebraid $\mathcal{K}$ is in *standard normal form* if $\mathcal{O}_A(\mathcal{K}) \cup \mathcal{O}_B(\mathcal{K}) \subseteq \mathcal{O}_A^* \cup \mathcal{O}_B^*$ and if every open atom $o \in \mathcal{O}_A(\mathcal{K}) \cup \mathcal{O}_B(\mathcal{K})$ points in $\mathcal{K}$ to a subbraid $\mathcal{K}'$ such that $o = o_{[\mathcal{K}']}$.

With $A \odot B$ we denote the set of all braids over $\mathcal{A}^\Sigma$ and $\mathcal{B}^\Delta$ in standard normal form. Note that trivial braids are always in standard normal form. Note also that the elements of a prebraid in standard normal form are in $\mathcal{A}_*^\Sigma \cup \mathcal{B}_*^\Delta$ (this follows from the remarks after Definition 3.8).

Lemma 4.36 *Every prebraid in standard normal form is irreducible.*

[5] compare Definition 4.3.

[6] compare Lemma 4.8.

Now $\mathcal{K}_1 = \mathcal{K}_0^{\langle m_1,n_1\rangle}$ is a subbraid of the prebraid $\mathcal{K}_0^{(m_1,n_1)}$ and both have the same root. Lemma 4.18 shows that $\mathcal{K}_1^{\langle m_2,n_2\rangle}$ is the unique subbraid of $(\mathcal{K}_0^{(m_1,n_1)})^{(m_2,n_2)} = \mathcal{K}_0^{(m_1\circ m_2, n_1\circ n_2)}$ given by its root, namely $\mathcal{K}_0^{\langle m_1\circ m_2, n_1\circ n_2\rangle}$. □

Corollary 4.30 *Let (m_1, n_1) be a simplifier for the braid $\mathcal{K}_0$, let (m_2, n_2) be a simplifier for the braid image $\mathcal{K}_1 = \mathcal{K}_0^{\langle m_1,n_1\rangle}$. Then there exists a simplifier (m, n) for $\mathcal{K}_0$ such that $\mathcal{K}_0^{\langle m,n\rangle} = \mathcal{K}_1^{\langle m_2,n_2\rangle}$.*

Proof. It follows from Lemma 4.1 that there exists a simplifier (m'_2, n'_2) of $\mathcal{K}_1$ such that (m'_2, n'_2) does not identify any pending atom of the simplification step leading from $\mathcal{K}_0$ to the braid image $\mathcal{K}_1$ with another atom, and $\mathcal{K}_1^{\langle m_2,n_2\rangle} = \mathcal{K}_1^{\langle m'_2,n'_2\rangle}$. Let $(m, n) := (m_1 \circ m'_2, n_1 \circ n'_2)$. Then, by the previous lemma, $\mathcal{K}_0^{\langle m,n\rangle} = \mathcal{K}_1^{\langle m'_2,n'_2\rangle} = \mathcal{K}_1^{\langle m_2,n_2\rangle}$. □

Theorem 4.31 *Let $\mathcal{K} = \mathcal{K}_0, \mathcal{K}_1, \ldots, \mathcal{K}_k$ be a sequence of braids such that each braid $\mathcal{K}_{i+1}$ is the braid image of $\mathcal{K}_i$ under a strict simplification, for $i = 0, \ldots, k-1$. Then $k \leq |\mathcal{O}(\mathcal{K})|$. If $\mathcal{K}'$ is an irreducible braid that is reached from $\mathcal{K}$ by a sequence of consecutive simplification steps (always taking braid images), then there exists a simplifier (m, n) for $\mathcal{K}$ such that $\mathcal{K}^{\langle m,n\rangle} = \mathcal{K}'$. If two irreducible braids $\mathcal{K}_1$ and $\mathcal{K}_2$ can be reached from $\mathcal{K}$ by sequences of consecutive simplification steps (always taking braid images), then $\mathcal{K}_1$ and $\mathcal{K}_2$ are variants.*

Proof. The first statement is trivial. The second statement follows from Corollary 4.30 by a simple induction. Assume that $\mathcal{K}_1$ and $\mathcal{K}_2$ are two irreducible braids that can be reached from $\mathcal{K}$ by sequences of consecutive simplification steps, always taking braid images. Then there exist simplifiers (m_1, n_1) and (m_2, n_2) of $\mathcal{K}$ such that $\mathcal{K}_1 = \mathcal{K}^{\langle m_1,n_1\rangle}$ and $\mathcal{K}_2 = \mathcal{K}^{\langle m_2,n_2\rangle}$. The prebraids $\mathcal{K}^{(m_1,n_1)}$ and $\mathcal{K}^{(m_2,n_2)}$ are not necessarily irreducible. But we may add further simplification steps (m'_1, n'_1) and (m'_2, n'_2) such that $(\mathcal{K}^{(m_1,n_1)})^{(m'_1,n'_1)}$ and $(\mathcal{K}^{(m_2,n_2)})^{(m'_2,n'_2)}$ are irreducible. By Lemma 4.15, (m'_1, n'_1) and (m'_2, n'_2) are—obviously non-strict—simplifiers for $\mathcal{K}_1$ and $\mathcal{K}_2$ respectively. It follows that $\mathcal{K}_1$ and $\mathcal{K}_1^{\langle m'_1,n'_1\rangle}$ are variants, and similarly for $\mathcal{K}_2$ and $\mathcal{K}_2^{\langle m'_2,n'_2\rangle}$. By Theorem 4.26, the two prebraids $(\mathcal{K}^{(m_1,n_1)})^{(m'_1,n'_1)}$ and $(\mathcal{K}^{(m_2,n_2)})^{(m'_2,n'_2)}$ are variants. By Lemma 4.10, the two subbraids given by their roots—which are, by Lemma 4.18, $\mathcal{K}_1^{\langle m'_1,n'_1\rangle}$ and $\mathcal{K}_2^{\langle m'_2,n'_2\rangle}$—are variants. Hence $\mathcal{K}_1$ and $\mathcal{K}_2$ are variants, by Lemma 4.11. □

On the basis of Theorem 4.31 we may introduce the following equivalence relation on the set of all braids.

Definition 4.32 Two braids are called *equivalent* if they can be simplified to the same irreducible braid image. If $\mathcal{K}$ is a braid, $[\mathcal{K}]$ denotes the set of all braids that are equivalent to $\mathcal{K}$.

We want to show that (m_∞, n_∞) is a simplifier for $\mathcal{K}$. Suppose that $m_\infty(o) = m_\infty(o')$ for open atoms $o \neq o' \in \mathcal{O}_A(\mathcal{K})$. We may assume that there exists a sequence $o = o_0^A, o_1^A, \ldots, o_r^A = o'$ of elements of $\mathcal{O}_A(C_1) \cup \mathcal{O}_A(C_2)$, where at least the elements $o_0^A, \ldots, o_{r-1}^A$ are in $\mathcal{O}_A(C_1)$, such that $o_i^A = m^i(o_0^A)$, for $0 \leq i \leq r$. Let $b_i := \pi_A(o_i^A)$ $(0 \leq i \leq r)$. Note that at least the elements $b_0, \ldots, b_{r-1}$ are in D_1 since $\mathcal{K}_1$ is a subbraid of $\mathcal{K}$ and π_A and π_A^1 coincide on $\mathcal{O}_A(C_1)$. Since $\langle o_0^A, b_0 \rangle \in \pi_A^1$ we know, by choice of (m, n), that $\langle m(o_0^A), n(b_0) \rangle \in \pi_A^2 \subseteq \pi_A$, which means that $b_1 = n(b_0)$. Similarly we see that $b_i = n^i(b_0)$ for $i = 0, \ldots, r$. But then we have

$$n_\infty(b_0) = n_\infty(n(b_0)) = n_\infty(b_1) = \ldots = n_\infty(b_{k-1}) = n_\infty(n(b_{r-1})) = n_\infty(b_r)$$

Thus n_∞ identifies the π_A-images of $o = o_0^A$ and $o' = o_r^A$. Symmetrically, if $n_\infty(o) = n_\infty(o')$ for atoms $o, o' \in \mathcal{O}_B(\mathcal{K})$, then m_∞ identifies the π_B-images of o and o'. Therefore (m_∞, n_∞) is in fact a simplifier for $\mathcal{K}$. But (m_∞, n_∞) is strict, by $(\diamond)$. This is a contradiction. Thus $\mathcal{K}_1 = \mathcal{K}_2$. $\square$

We shall now turn to simplification of braids. First we shall show that the result of two consecutive simplification steps may be obtained by a single simplification, similarly as for prebraids. We have to adapt the notion of a pending atom to the new situation.

Definition 4.28 Let (m, n) be a simplifier for the braid $\mathcal{K} = \langle a, C, D, \pi_A, \pi_B \rangle$. Let $\mathcal{K}^{\langle m,n \rangle} = \langle a', C', D', \pi'_A, \pi'_B \rangle$. Then the set

$$(\{m(o) \mid o \in \mathcal{O}_A(C)\} \setminus \mathcal{O}_A(C')) \;\cup\; (\{n(o) \mid o \in \mathcal{O}_B(D)\} \setminus \mathcal{O}_B(D'))$$

is called the set of *pending atoms* of the simplification step from $\mathcal{K}$ to the braid image $\mathcal{K}^{\langle m,n \rangle}$.

Note that this is really a new notion. The set of pending atoms of the simplification step from $\mathcal{K}$ to the *braid* image $\mathcal{K}^{\langle m,n \rangle}$ is a *superset* of the set of pending atoms of the simplification step from $\mathcal{K}$ to the image $\mathcal{K}^{(m,n)}$, but both sets are not necessarily identical.

Lemma 4.29 *Let (m_1, n_1) be a simplifier for the braid $\mathcal{K}_0$, let (m_2, n_2) be a simplifier for its braid image $\mathcal{K}_1 := \mathcal{K}_0^{\langle m_1,n_1 \rangle}$. Assume that m_2 and n_2 do not identify any pending atom of the simplification step leading from $\mathcal{K}_0$ to the braid image $\mathcal{K}_1$ with another atom. Then $(m_1 \circ m_2, n_1 \circ n_2)$ is a simplifier for $\mathcal{K}_0$ and $\mathcal{K}_0^{\langle m_1 \circ m_2, n_1 \circ n_2 \rangle} = \mathcal{K}_1^{\langle m_2,n_2 \rangle}$.*

Proof. Exactly as in the corresponding proof of Lemma 4.21 it follows that $(m_1 \circ m_2, n_1 \circ n_2)$ is a simplifier for $\mathcal{K}_0$. Our assumptions guarantee that (m_2, n_2) is also a simplifier for $\mathcal{K}_0^{(m_1,n_1)}$ such that m_2 and n_2 do not identify any pending atom of the simplification step leading from $\mathcal{K}_0$ to the image $\mathcal{K}_0^{(m_1,n_1)}$ with another atom. Hence Lemma 4.21 implies that $(\mathcal{K}_0^{(m_1,n_1)})^{(m_2,n_2)} = \mathcal{K}_0^{(m_1 \circ m_2, n_1 \circ n_2)}$.

admissible endomorphism that maps o_1 to o_2 and leaves all other atoms fixed, let n be the identity on B. Now (m,n) is a strict simplifier for $\mathcal{K}$, thus we get a contradiction.

(b) Assume, to get a contradiction, that (m,n) is a strict simplifier for $\mathcal{K}'$. Let $\mathcal{K} = \langle a, C, D, \pi_A, \pi_B \rangle$, let $\mathcal{K}' = \langle a', C', D', \pi'_A, \pi'_B \rangle$. Let $X_A = \mathcal{O}_A(C) \setminus \mathcal{O}_A(C')$, let $Y_B = \mathcal{O}_B(D) \setminus \mathcal{O}_B(D')$. By Lemma 4.1 we may assume that m (n) leaves the elements of X_A (Y_B) fixed. If $\{m(o) \mid o \in \mathcal{O}_A(C')\} \cap X_A = \emptyset = \{n(o) \mid o \in \mathcal{O}_B(D')\} \cap Y_B$, then m (n) only identifies open atoms of $\mathcal{O}(\mathcal{K}')$ and it is easy to see that (m,n) is a strict simplifier for $\mathcal{K}$, which yields a contradiction. In the other case, let m' be an admissible automorphism such that $\{m'(m(o)) \mid o \in \mathcal{O}_A(C')\} \cap X_A = \emptyset$, let n' be an admissible automorphism such that $\{n'(n(o)) \mid o \in \mathcal{O}_B(D')\} \cap Y_B = \emptyset$. Let m^* denote the endomorphism that coincides with $m \circ m'$ on $\mathcal{O}_A(C')$ and leaves all other open atoms fixed. Let n^* denote the endomorphism that coincides with $n \circ n'$ on $\mathcal{O}_B(D')$ and leaves all other open atoms fixed. Then (m^*, n^*) is a strict simplifier for $\mathcal{K}$, which yields a contradiction. □

(c) Let $\mathcal{K} = \langle a, C, D, \pi_A, \pi_B \rangle$, let $\mathcal{K}_i = \langle a_i, C_i, D_i, \pi^i_A, \pi^i_B \rangle$ $(i = 1,2)$. Assume that $\mathcal{K}_1$ and $\mathcal{K}_2$ are variants, but $\mathcal{K}_1 \neq \mathcal{K}_2$. There exists a pair of admissible automorphisms (m,n) such that $\mathcal{K}_2 = \mathcal{K}_1^{(m,n)}$. Without loss of generality we have $(\diamond)$: there exists an $o^* \in \mathcal{O}_A(C_1)$ such that $o^* \neq m(o^*) \in \mathcal{O}_A(C_2)$.

Consider an element $o^A_0 \in \mathcal{O}_A(C_1)$. If all elements of the "orbit"

$$o^A_0,\ o^A_1 := m(o^A_0),\ o^A_2 := m(o^A_1),\ o^A_3 := m(o^A_2) \ldots$$

are in $\mathcal{O}_A(C_1)$, then this sequence contains only a finite number of distinct elements, say, $o^A_0, \ldots, o^A_k$. In the other case, let k be the first index in the sequence $0, 1, \ldots$ such that $o^A_k \notin \mathcal{O}_A(C_1)$. This implies that $o^A_k \in \mathcal{O}_A(C_2)$. The set $\{o^A_0, \ldots, o^A_k\}$ is called the m-trace $tr_m(o^A_0)$ of o^A_0. Let $\sim_m$ be the smallest equivalence relation on $\mathcal{O}_A(C_1) \cup \mathcal{O}_A(C_2)$ such that $o \sim_m o'$ whenever o and o' both belong to the m-trace $tr(o^A)$ of the same element $o^A \in \mathcal{O}_A(C_1)$. Since m is an injective function, the equivalence classes of $\sim_m$ are just the maximal m-traces. For each equivalence class $[o^A]_{\sim_m}$, choose a representant $rep([o^A]_{\sim_m})$. Let $m_\infty \in \mathcal{M}$ be the admissible endomorphism that maps each $o^A \in \mathcal{O}_A(C_1) \cup \mathcal{O}_A(C_2)$ to the representant $rep([o^A]_{\sim_m})$ and leaves other atoms fixed. Since $o^A \in \mathcal{O}_A(C_1)$ implies that $m(o^A) \in \mathcal{O}_A(C_2)$, and since both atoms have the same representant, we know that $m_\infty(m(o^A)) = m_\infty(o^A)$ for all $o^A \in \mathcal{O}_A(C_1)$. This implies, by Lemma 4.1, that $m_\infty(m(a)) = m_\infty(a)$ for all $a \in C_1$.

Symmetrically, we may define the n-traces $tr_n(o^B)$ of elements $o^B \in \mathcal{O}_B(D_1)$, just by replacing $\mathcal{O}_A(C_i)$ by $\mathcal{O}_B(D_i)$ $(i = 1,2)$ and m by n. We obtain the equivalence relation $\sim_n$ by "identifying" all elements that belong to the same n-trace $tr_n(o^B)$, for some $o^B \in \mathcal{O}_B(D_1)$. For each equivalence class $[o^B]_{\sim_n}$, choose a representant $rep([o^B]_{\sim_n})$. Let $n_\infty \in \mathcal{N}$ be the admissible endomorphism that maps each $o^B \in \mathcal{O}_B(D_1) \cup \mathcal{O}_B(D_2)$ to the representant $rep([o^B]_{\sim_n})$ and leaves other atoms fixed. We have $n_\infty(n(o^B)) = n_\infty(o^B)$ for all $o^B \in \mathcal{O}_B(D_1)$, and $n_\infty(n(b)) = n_\infty(b)$ for all $b \in D_1$.

$o_1 = u_1, u_2 \ldots, u_k = o_2$ such that each pair $\langle u_i, u_{i+1}\rangle$ belongs either to $\sim_A^1$ or to $\sim_A^2$ $(1 \leq i < k)$. Let $d_i := \pi_A^0(u_i)$, for $i = 1, \ldots, k$. Thus $d_i \in D_0$ $(1 \leq i \leq k)$ and we have $b_1 = d_1$ and $b_2 = d_k$. Now (m_1, n_1) and (m_2, n_2) are simplifiers. Thus, if $\langle u_i, u_{i+1}\rangle \in\sim_A^1$, then $n_1(d_i) = n_1(d_{i+1})$, which implies $n_3(n_1(d_i)) = n_3(n_1(d_{i+1}))$, and if $\langle u_i, u_{i+1}\rangle \in\sim_A^2$, then $n_2(d_i) = n_2(d_{i+1})$, which implies $n_4(n_2(d_i)) = n_4(n_2(d_{i+1}))$ and, by $(**)$, $n_3(n_1(d_i)) = n_3(n_1(d_{i+1}))$. Therefore we obtain in fact

$$n_3(b_1') = n_3(n_1(b_1)) = n_3(n_1(b_2)) = n_3(b_2').$$

We have shown that (m_3, n_3) is a simplifier for $\mathcal{K}_1$. Symmetrically it follows that (m_4, n_4) is a simplifier for $\mathcal{K}_2$.

The two prebraids $\mathcal{K}_1^{(m_3,n_3)}$ and $\mathcal{K}_2^{(m_4,n_4)}$ have the same root $m_3(m_1(a_0))$, by $(**)$. It is trivial to see that $(**)$ also implies that $\{m_3(m_1(c)) \mid c \in C_0\}$ is the set of elements of type A of both prebraids, and $\{n_3(n_1(d)) \mid d \in D_0\}$ is the set of elements of type B of both prebraids. But then it follows easily that both prebraids have the same linking functions, which means that they are identical. □

Theorem 4.26 *Each sequence of iterated strict simplifications that starts from the prebraid $\mathcal{K}$ has length $\leq |\mathcal{O}(\mathcal{K})|$. If $\mathcal{K}'$ is an irreducible prebraid that is obtained from $\mathcal{K}$ by a sequence of simplifications, then there exists a simplifier (m, n) for $\mathcal{K}$ such that $\mathcal{K}^{(m,n)} = \mathcal{K}'$. If two irreducible prebraids $\mathcal{K}_1$ and $\mathcal{K}_2$ can be reached from $\mathcal{K}$ by sequences of simplifications, then $\mathcal{K}_1$ and $\mathcal{K}_2$ are variants.*

Proof. The first statement is Lemma 4.24. The second statement follows from Corollary 4.22 with a trivial induction. If $\mathcal{K}_1$ and $\mathcal{K}_2$ are two irreducible prebraids that are obtained from $\mathcal{K}$ by sequences of simplifications, then both prebraids can be obtained from $\mathcal{K}$ by a single simplification step, by Corollary 4.22. Lemma 4.25 shows that there exists a prebraid $\mathcal{K}_3$ that can be reached from $\mathcal{K}_1$ and $\mathcal{K}_2$ by simplification. Since $\mathcal{K}_1$ and $\mathcal{K}_2$ are irreducible, these simplification steps are not strict. Hence $\mathcal{K}_1$, $\mathcal{K}_2$ and $\mathcal{K}_3$ are variants, by Lemma 4.11 and Lemma 4.23. □

Before we treat simplification of braids, let us mention three properties of irreducible prebraids.

Lemma 4.27 *(a) If the prebraid $\mathcal{K} = \langle a, C, D, \pi_A, \pi_B\rangle$ is irreducible, then π_A and π_B are injective.*
(b) If $\mathcal{K}'$ is a subbraid of the irreducible prebraid $\mathcal{K}$, then $\mathcal{K}'$ is irreducible.
(c) If $\mathcal{K}_1$ and $\mathcal{K}_2$ are subbraids of the irreducible prebraid $\mathcal{K}$, and if $\mathcal{K}_1$ and $\mathcal{K}_2$ are variants, then $\mathcal{K}_1 = \mathcal{K}_2$.

Proof. (a) Assume that π_A, say, is not injective. Then there exist elements $\langle o_1, b_1\rangle$ and $\langle o_2, b_1\rangle$ in π_A where o_1 and o_2 are distinct. Let $m \in \mathcal{M}$ be an

Lemma 4.25 *Let (m_1, n_1) and (m_2, n_2) be two simplifiers for the prebraid $\mathcal{K}_0$, let $\mathcal{K}_1$ and $\mathcal{K}_2$ be the images of $\mathcal{K}_0$ under (m_1, n_1) and (m_2, n_2) respectively. Then there exist a simplifier (m_3, n_3) for $\mathcal{K}_1$ and a simplifier (m_4, n_4) for $\mathcal{K}_2$ such that $\mathcal{K}_1^{(m_3,n_3)} = \mathcal{K}_2^{(m_4,n_4)}$.*

Proof. Let $\mathcal{K}_i = \langle a_i, C_i, D_i, \pi_A^i, \pi_B^i \rangle$, for $i = 0, 1, 2$. The endomorphisms m_1 and n_1 define equivalence relations $\sim_A^1$ and $\sim_B^1$ on $\mathcal{O}_A(C_0)$ and $\mathcal{O}_B(D_0)$ respectively, where elements are equivalent with respect to $\sim_A^1$ ($\sim_B^1$) iff they have the same image under m_1 (n_1). The endomorphisms m_2 and n_2 define similar equivalence relations $\sim_A^2$ and $\sim_B^2$ on $\mathcal{O}_A(C_0)$ and $\mathcal{O}_B(D_0)$ respectively. Let $\sim_A := \sim_A^1 \sqcup \sim_A^2 = (\sim_A^1 \cup \sim_A^2)^*$ denote smallest equivalence relation on $\mathcal{O}_A(\mathcal{K}_0)$ that extends $\sim_A^1$ and $\sim_A^2$. Similarly, let $\sim_B$ denote the smallest equivalence relation on $\mathcal{O}_B(\mathcal{K}_0)$ that extends $\sim_B^1$ and $\sim_B^2$. Choose a system of representants for $\sim_A$ and a similar system for $\sim_B$. We shall write *rep*(o) for the representant of $[o]$ with respect to $\sim_A$ ($\sim_B$), for $o \in \mathcal{O}_A(\mathcal{K})$ ($o \in \mathcal{O}_B(\mathcal{K})$).

The elements of $\mathcal{O}_A(C_1)$ have the form $m_1(o_A)$ for $o_A \in \mathcal{O}_A(C_0)$, by Lemma 3.10. If, for $o_A, o'_A \in \mathcal{O}_A(C_0)$, $m_1(o_A) = m_1(o'_A) \in \mathcal{O}_A(C_1)$, then $o_A \sim_A^1 o'_A$ and *rep*$(o_A) =$ *rep*(o'_A). Thus

- the mapping $m_1(o_A) \mapsto rep(o_A)$ ($o_A \in \mathcal{O}_A(C_0)$) is welldefined. It can be extended to an admissible endomorphism $m_3 \in \mathcal{M}$. Similarly the mapping $n_1(o_B) \mapsto rep(o_B)$ ($o_B \in \mathcal{O}_B(D_0)$) is welldefined and can be extended to an admissible endomorphism $n_3 \in \mathcal{N}$.

Symmetrically we can show

- the mapping $m_2(o_A) \mapsto rep(o_A)$ ($o_A \in \mathcal{O}_A(C_0)$) is welldefined and can be extended to an admissible endomorphism $m_4 \in \mathcal{M}$, and the mapping $n_2(o_B) \mapsto rep(o_B)$ ($o_B \in \mathcal{O}_B(D_0)$) is welldefined and can be extended to an admissible endomorphism $n_4 \in \mathcal{N}$.

We have $(*)$

$$\begin{array}{rcll} m_3(m_1(o_A)) & = rep(o_A) = & m_4(m_2(o_A)) & (o_A \in \mathcal{O}_A(C_0)) \\ n_3(n_1(o_B)) & = rep(o_B) = & n_4(n_2(o_B)) & (o_B \in \mathcal{O}_B(D_0)) \end{array}$$

and, by Lemma 4.1, $(**)$

$$\begin{array}{rcll} m_3(m_1(c)) & = & m_4(m_2(c)) & (c \in C_0) \\ n_3(n_1(d)) & = & n_4(n_2(d)) & (d \in D_0). \end{array}$$

Clearly (m_3, n_3) and (m_4, n_4) are admissible. We shall now show that (m_3, n_3) is a simplifier for $\mathcal{K}_1$. Let $\langle o'_1, b'_1 \rangle, \langle o'_2, b'_2 \rangle \in \pi_A^1$ and suppose that $m_3(o'_1) = m_3(o'_2)$. We have to verify that $n_3(b'_1) = n_3(b'_2)$. For $i = 1, 2$, there exists $o_i \in \mathcal{O}_A(C_0)$ and $b_i := \pi_A^0(o_i) \in D_0$ such that $o'_i = m_1(o_i)$ and $b'_i = n_1(b_i)$. Since m_3 identifies o'_1 and o'_2 we know that $o_1 \sim_A o_2$. Thus there exists a sequence

$n_1(\pi_A^0(o')) = \pi_A^1(m_1(o'))$. Since $m_2(m_1(o)) = m_2(m_1(o'))$ and (m_2, n_2) is a simplifier for $\mathcal{K}_1$, this implies that

$$n_2(n_1(\pi_A^0(o))) = n_2(\pi_A^1(m_1(o))) = n_2(\pi_A^1(m_1(o'))) = n_2(n_1(\pi_A^0(o'))).$$

Hence in both cases $n(\pi_A^0(o)) = n(\pi_A^0(o'))$. Symmetrically it follows that $n(o) = n(o')$ implies $m(\pi_B^0(o)) = m(\pi_B^0(o'))$, for all $o, o' \in \mathcal{O}_A(D_0)$. Hence (m, n) is a simplifier for $\mathcal{K}_0$.

The prebraids $\mathcal{K}_0^{(m_1 \circ m_2, n_1 \circ n_2)}$ and $\mathcal{K}_1^{(m_2,n_2)}$ have the same root $m_2(m_1(a_0))$. It is trivial that they have the same elements. But then it follows easily that they have the same linking functions. □

Corollary 4.22 *Let (m_1, n_1) be a simplifier for the prebraid $\mathcal{K}_0$ and (m_2, n_2) be a simplifier for the prebraid $\mathcal{K}_1 = \mathcal{K}_0^{(m_1,n_1)}$. Then there exists a simplifier (m, n) for $\mathcal{K}_0$ such that $\mathcal{K}_0^{(m,n)} = \mathcal{K}_1^{(m_2,n_2)}$.*

Proof. It follows from Lemma 4.1 that there exists a simplifier (m_2', n_2') of $\mathcal{K}_1$ such that (m_2', n_2') does not identify any pending atom of the simplification step leading from $\mathcal{K}_0$ to $\mathcal{K}_1$ with another atom, and $\mathcal{K}_1^{(m_2,n_2)} = \mathcal{K}_1^{(m_2',n_2')}$. Let $(m, n) := (m_1 \circ m_2', n_1 \circ n_2')$. Then, by the previous lemma, $\mathcal{K}_0^{(m,n)} = \mathcal{K}_1^{(m_2',n_2')} = \mathcal{K}_1^{(m_2,n_2)}$. □

Let $\mathcal{K}$ be a prebraid. We have seen that a simplifier (m, n) that yields a permutation of $\mathcal{O}_A$ and $\mathcal{O}_B$ leads to the variant $\mathcal{K}^{(m,n)}$ of $\mathcal{K}$ (Lemma 4.12). The same is true under weaker assumptions. By Lemma 4.1, the image $\mathcal{K}^{(m,n)}$ is completely determined by the images of the elements in $\mathcal{O}_A(\mathcal{K})$ and $\mathcal{O}_B(\mathcal{K})$ under the endomorphisms m and n respectively. Hence we obtain

Lemma 4.23 *Let (m, n) be a simplifier for the prebraid $\mathcal{K}$. If the restrictions of m and n on $\mathcal{O}_A(\mathcal{K})$ and $\mathcal{O}_B(\mathcal{K})$ respectively are injective, then $\mathcal{K}^{(m,n)}$ is a variant of $\mathcal{K}$.*

Call a simplifier (m, n) for $\mathcal{K}$ *strict* if the restriction of m on $\mathcal{O}_A(\mathcal{K})$ or the restriction of n on $\mathcal{O}_B(\mathcal{K})$ is *not* injective. Lemma 3.10 shows that $|\mathcal{O}(\mathcal{K}^{(m,n)})| < |\mathcal{O}(\mathcal{K})|$ if (m, n) is strict. It follows that

Lemma 4.24 *$|\mathcal{O}(\mathcal{K})|$ gives an upper bound on the length of every sequence of strict simplifications for the prebraid $\mathcal{K}$.*

A prebraid $\mathcal{K}'$ is called *irreducible* if $\mathcal{K}'$ does not have a strict simplifier. We want to show that all irreducible prebraids that can be reached from a prebraid $\mathcal{K}$ by simplification are variants. For this purpose, the following lemma is needed that shows that simplification of prebraids is "locally confluent".

There is one technical point behind the definition of a simplifier that will cause some difficulties in the further development. Assume, in the situation of Definition 4.16, that $\mathcal{O}_A(C) = \{o_1, \ldots, o_k\}$ and $\mathcal{O}_B(D) = \{u_1, \ldots, u_l\}$. Then there is no guarantee that $\mathcal{O}_A(C') = \{m(o_1), \ldots, m(o_k)\}$ and $\mathcal{O}_B(D') = \{n(u_1), \ldots, n(u_l)\}$. In fact, Lemma 3.10 only shows the inclusion $\mathcal{O}_A(C') \subseteq \{m(o_1), \ldots, m(o_k)\}$.

Definition 4.19 The set

$$(\{m(o) \mid o \in \mathcal{O}_A(C)\} \setminus \mathcal{O}_A(C')) \ \cup \ (\{n(o) \mid o \in \mathcal{O}_B(D)\} \setminus \mathcal{O}_B(D'))$$

is called the set of *pending atoms* of the simplification step leading from $\mathcal{K}$ to $\mathcal{K}^{(m,n)}$.

As we shall see, pending atoms complicate the treatment of simplification. In principle we could restrict the amalgamation construction to a class of structures for which we can replace the inclusion from Lemma 3.10 mentioned above by an equality. In this case pending atoms cannot occur, image and braid image always coincide, and we could dispense with prebraids at all. However, our motivation was to give a general construction. For this reason we shall not follow this line.

Lemma 4.20 *Assume, in the situation of Definition 4.16, that $o \in \mathcal{O}_A(C)$ and $m(o)$ is not a pending atom of the simplification step leading from $\mathcal{K}$ to $\mathcal{K}^{(m,n)}$. Then $\pi'_A(m(o)) = n(\pi_A(o))$.*

Proof. A trivial consequence of the definition of π'_A as given in 4.16. □

While we are mainly interested in simplification of braids, it turns out to be simpler to treat simplification of prebraids in advance.

Lemma 4.21 *Let (m_1, n_1) be a simplifier for the prebraid $\mathcal{K}_0$ and (m_2, n_2) be a simplifier for the prebraid $\mathcal{K}_1 = \mathcal{K}_0^{(m_1,n_1)}$. Assume that m_2 and n_2 do not identify any pending atom of the simplification step leading from $\mathcal{K}_0$ to $\mathcal{K}_1$ with another atom. Then $(m_1 \circ m_2, n_1 \circ n_2)$ is a simplifier for $\mathcal{K}_0$ and $\mathcal{K}_0^{(m_1 \circ m_2, n_1 \circ n_2)} = \mathcal{K}_1^{(m_2,n_2)}$.*

Proof. Let $\mathcal{K}_i = \langle a_i, C_i, D_i, \pi^i_A, \pi^i_B \rangle$, for $i = 0, 1$. We may assume that both are of type A. Let $(m, n) := (m_1 \circ m_2, n_1 \circ n_2)$. If $m(o) = m(o')$ for $o, o' \in \mathcal{O}_A(C_0)$, then either $m_1(o) = m_1(o')$, or $m_1(o) \neq m_1(o')$ and $m_2(m_1(o)) = m_2(m_1(o'))$. In the former case we know that $n_1(\pi^0_A(o)) = n_1(\pi^0_A(o'))$ since (m_1, n_1) is a simplifier for $\mathcal{K}_0$. Hence $n(\pi^0_A(o)) = n(\pi^0_A(o'))$. In the latter case, neither $m_1(o)$ nor $m_1(o')$ can be pending, by assumption. Hence $m_1(o)$ and $m_1(o')$ are in $\mathcal{O}_A(\mathcal{K}_1)$. By Lemma 4.20, $n_1(\pi^0_A(o)) = \pi^1_A(m_1(o))$ and

Proof. Let $\mathcal{K}' = \langle a', C', D', \pi'_A, \pi'_B \rangle$ be a subprebraid of $\mathcal{K} = \langle a, C, D, \pi_A, \pi_B \rangle$. Then $C' \subseteq C$ and $D' \subseteq D$. Hence $\mathcal{O}_A(C') \subseteq \mathcal{O}_A(C)$ and $\mathcal{O}_B(D') \subseteq \mathcal{O}_B(D)$. Moreover, the functions π_A and $\pi'_A \subseteq \pi_A$ (π_B and $\pi'_B \subseteq \pi_B$) coincide on $\mathcal{O}_A(C')$ (resp. $\mathcal{O}_B(D')$). The rest is obvious. □

Definition 4.16 Let (m, n) be a simplifier for the prebraid $\mathcal{K} = \langle a, C, D, \pi_A, \pi_B \rangle$. The *image of* $\mathcal{K}$ *with respect to* (m, n) is the prebraid $\mathcal{K}^{(m,n)} := \langle a', C', D', \pi'_A, \pi'_B \rangle$ with the following components:[4]

1. $a' := m(a)$,
2. $C' := \{m(c) \mid c \in C\}$ and
   $D' := \{n(d) \mid d \in D\}$,
3. $\pi'_A := \{\langle m(o), n(d) \rangle \mid \langle o, d \rangle \in \pi_A, m(o) \in \mathcal{O}_A(C')\}$, and
   $\pi'_B := \{\langle n(o), m(c) \rangle \mid \langle o, c \rangle \in \pi_B, n(o) \in \mathcal{O}_B(D')\}$.

Now assume that $\mathcal{K}$ is a braid. The *braid-image of* $\mathcal{K}$ *with respect to* (m, n), $\mathcal{K}^{\langle m,n \rangle}$, is the unique subbraid of $\mathcal{K}^{(m,n)}$ with root a'.

Example 4.17 The following figure represents the braid-image of the braid from Example 4.4 under the simplification (m, n) where m maps o_3 to o_2 and n maps u_3 to u_2:

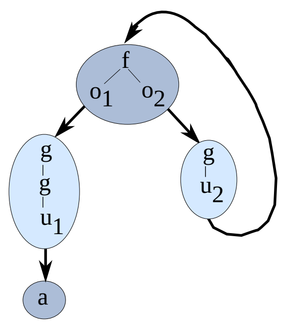

The next lemma gives a refinement of Lemma 4.15.

Lemma 4.18 *Let (m, n) be a simplifier for the prebraid $\mathcal{K}$, let $\mathcal{K}_1$ be the unique subbraid of $\mathcal{K}$ with root e, where e is an element of $\mathcal{K}$ of type A (resp. B). Then $\mathcal{K}_1^{\langle m,n \rangle}$ is the unique subbraid of $\mathcal{K}^{(m,n)}$ with root $m(e)$ (resp. $n(e)$).*

Proof. It follows directly from Definition 4.16 that $\mathcal{K}_1^{\langle m,n \rangle}$ is a subbraid of the prebraid $\mathcal{K}^{(m,n)}$. Obviously $m(e)$ (resp. $n(e)$) is the root of $\mathcal{K}_1^{\langle m,n \rangle}$. Now use Corollary 4.7. □

[4]Using Lemma 3.10 and the fact that both $\mathcal{A}^\Sigma$ and $\mathcal{B}^\Delta$ are non-collapsing it is trivial to verify that $\mathcal{K}^{(m,n)}$ is a prebraid.

Lemma 4.12 *Let* (m, n) *be an admissible pair of automorphisms. Let* $\mathcal{K}$*,* a'*,* C'*,* D'*,* π'_A*, and* π'_B *be defined as in Definition 4.9, 1.-3. Then* $\mathcal{K}' := \langle a', C', D', \pi'_A, \pi'_B \rangle$ *is a prebraid and a variant of* $\mathcal{K}$*.*

Proof. Since $a \in A \setminus \mathcal{O}_A$ it follows that $a' = m(a) \in A \setminus \mathcal{O}_A$, by choice of m. Thus $\mathcal{K}'$ satisfies Condition 1 of Definition 4.3. Since m and n are admissible automorphisms, all elements of $C' \cup D'$ that are images of non-atomic elements of $C \cup D$ under m and n respectively are non-atomic. Hence $\mathcal{K}'$ satisfies Condition 2 of Definition 4.3. Since m and n define permutations of $\mathcal{O}_A$ and $\mathcal{O}_B$ respectively, the first component o of each pair $\langle o, e \rangle$ in $\pi'_A \cup \pi'_B$ is an open atom. Lemma 3.10 shows that $o \in \mathcal{O}_A(C') \cup \mathcal{O}_B(D')$. Obviously, if $o \in \mathcal{O}_A(C')$, then $e \in D'$ and if $o \in \mathcal{O}_A(D')$, then $e \in C'$. Moreover, e is always non-atomic, by admissibility of m and n. Since m and n are automorphisms, π'_A and π'_B are functions. By Lemma 3.10, the domains of π'_A and π'_B are $\mathcal{O}_A(C')$ and $\mathcal{O}_B(D')$ respectively, which shows that $\mathcal{K}'$ satisfies Condition 3 of Definition 4.3. Thus $\mathcal{K}'$ is a prebraid. Clearly it is a variant of $\mathcal{K}$. □

Lemma 4.13 *Each variant of a braid is a braid.*

Proof. It suffices to verify that each variant of a braid satisfies Condition 4 of Definition 4.3. Let $\mathcal{K}$, $\mathcal{K}'$, and (m, n) as in Definition 4.9. Suppose that $d \in D$ is directly linked to $c \in C$ via π_A. Thus, for some $o \in \mathcal{O}_A(c)$ we have $\langle o, d \rangle \in \pi_A$. Lemma 3.10 shows that $m(o) \in \mathcal{O}_A(m(c))$. Clearly $m(c) \in C'$. Since $\langle m(o), n(d) \rangle \in \pi'_A$, the element $n(d) \in D'$ is directly linked to $m(c) \in C'$. Now a simple induction shows that all elements of $C' \cup D'$ are π'-descendants of the new root $a' = m(a)$, where $\pi' = \pi'_A \cup \pi'_B$. □

4.3 Simplification of braids

Two (pre)braids that are variants of each other are meant to denote the same object. But then we should not distinguish between two subbraids of one and the same (pre)braid if they are variants. In order to identify such subbraids, we shall use admissible pairs of endomorphisms of a particular type.

Definition 4.14 The admissible pair (m, n) is a *simplifier* for the prebraid $\mathcal{K} = \langle a, C, D, \pi_A, \pi_B \rangle$ if the following conditions hold:

- $\forall o_1, o_2 \in \mathcal{O}_A(C)$: $m(o_1) = m(o_2)$ implies $n(\pi_A(o_1)) = n(\pi_A(o_2))$,
- $\forall o_1, o_2 \in \mathcal{O}_B(D)$: $n(o_1) = n(o_2)$ implies $m(\pi_B(o_1)) = m(\pi_B(o_2))$.

Lemma 4.15 *Let* (m, n) *be a simplifier for the prebraid* $\mathcal{K}$*. Then* (m, n) *is a simplifier for each subprebraid of* $\mathcal{K}$*.*

Proof. Let $\mathcal{K} = \langle a, C, D, \pi_A, \pi_B \rangle$ and $e \in C \cup D$. Then e cannot be an open atom. Let $C' \subseteq C$ ($D' \subseteq D$) be the set of π-descendants of e in C (resp. D), where $\pi = \pi_A \cup \pi_B$. Note that all elements of $(C' \cup D') \setminus \{e\}$ are non-atomic since $\mathcal{K}$ is a prebraid. Let $\pi'_A \subseteq \pi_A$ (resp. $\pi'_B \subseteq \pi_B$) contain all ordered pairs $\langle o, c \rangle$ of π_A (resp. π_B) where $o \in \mathcal{O}_A(C')$ (resp. $o \in \mathcal{O}_B(D')$). For each such pair $\langle o, c \rangle$ the element c is in D' (C') since c is a π-descendant of e. Since $\langle o, c \rangle \in \pi$, the element c is non-atomic. It follows that $\langle e, C', D', \pi'_A, \pi'_B \rangle$ is a subbraid of $\mathcal{K}$. By Corollary 4.7 it is the unique subbraid of $\mathcal{K}$ with root e. □

4.2 Variants

The concrete open atoms that are used to organize links between elements of distinct type in a given braid should be regarded as irrelevant. This motivates the following definition.

Definition 4.9 Let $\mathcal{K} = \langle a, C, D, \pi_A, \pi_B \rangle$ and $\mathcal{K}' = \langle a', C', D', \pi'_A, \pi'_B \rangle$ be two prebraids, say, of type A. $\mathcal{K}'$ is called a *variant* of $\mathcal{K}$ if there exists an admissible pair of automorphisms (m, n) such that

1. $a' = m(a)$,
2. $C' = \{m(c) \mid c \in C\}$, and $D' = \{n(d) \mid d \in D\}$,
3. $\pi'_A := \{\langle m(o), n(d) \rangle \mid \langle o, d \rangle \in \pi_A\}$, and
   $\pi'_B := \{\langle n(o), m(c) \rangle \mid \langle o, c \rangle \in \pi_B\}$.

Lemma 4.10 *If two prebraids are variants, then the two subbraids given by their roots are variants.*

Proof. Let $\mathcal{K}$ and $\mathcal{K}'$ be variants, of the form as in the previous definition. Let $\mathcal{K}_1 = \langle a, C_1, D_1, \pi^1_A, \pi^1_B \rangle$ be the unique subbraid of $\mathcal{K}$ with root a, and let $\mathcal{K}_2 = \langle a, C_2, D_2, \pi^2_A, \pi^2_B \rangle$ be the unique subbraid of $\mathcal{K}'$ with root $a' = m(a)$. The elements in $C_1 \cup D_1$ are the π-descendants of a, for $\pi = \pi_A \cup \pi_B$. The elements in $C_2 \cup D_2$ are the π'-descendants of a', for $\pi' = \pi'_A \cup \pi'_B$. From the definition of π'_A and π'_B and from Lemma 3.10 (second part) it follows easily that the π'-descendants of $a' = m(a)$ are the m resp. n -images of the π-descendants of a. This shows that $C_2 = \{m(c) \mid c \in C_1\}$ and $D_2 = \{n(d) \mid d \in D_1\}$. The rest is obvious. □

The following lemma shows that the notion of a variant gives rise to an equivalence relation on the set of all (pre)braids. Since the set of admissible automorphisms of $\mathcal{A}^\Sigma$ (resp. $\mathcal{B}^\Delta$) defines (with composition) a group, the proof is obvious.

Lemma 4.11 *Each prebraid $\mathcal{K}_1$ is a variant of $\mathcal{K}_1$. If $\mathcal{K}_2$ is a variant of the prebraid $\mathcal{K}_1$, then $\mathcal{K}_1$ is a variant of $\mathcal{K}_2$. If $\mathcal{K}_2$ is a variant of the prebraid $\mathcal{K}_1$, and if $\mathcal{K}_3$ is a variant of $\mathcal{K}_2$, then $\mathcal{K}_3$ is a variant of $\mathcal{K}_1$.*

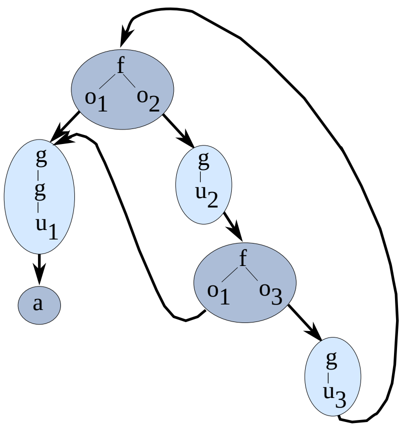

We sometimes write $\mathcal{O}_A(\mathcal{K})$ and $\mathcal{O}_B(\mathcal{K})$ for $\mathcal{O}_A(C)$ and $\mathcal{O}_B(D)$ respectively, and $\mathcal{O}(\mathcal{K})$ denotes the union $\mathcal{O}_A(\mathcal{K}) \cup \mathcal{O}_B(\mathcal{K})$. A quintuple $\mathcal{K} = \langle a, C, D, \pi_A, \pi_B \rangle$ that satisfies Conditions 1-3 of Definition 4.3 will be called a *prebraid*.

Definition 4.5 Let $\mathcal{K} = \langle a, C, D, \pi_A, \pi_B \rangle$ be a braid. The braid $\mathcal{K}' := \langle a', C', B', \pi'_A, \pi'_B \rangle$ (of type A or B) is a *subbraid* of $\mathcal{K}$ if $a' \in C \cup D$, $C' \subseteq C$, $D' \subseteq D$, $\pi'_A \subseteq \pi_A$, and $\pi'_B \subseteq \pi_B$.

Sub(pre)braids of prebraids are defined in the same way. We write $\mathcal{K}' \subseteq \mathcal{K}$ if $\mathcal{K}'$ is a sub(pre)braid of $\mathcal{K}$.

Lemma 4.6 *Let $\mathcal{K}_i = \langle a_i, C_i, D_i, \pi^i_A, \pi^i_B \rangle$ be a braid, and let $\pi^i = \pi^i_A \cup \pi^i_B$, for $i = 1, 2$. If $a_1 = a_2$, and if $\pi_1(o) = \pi_2(o)$ for all $o \in \mathcal{O}(\mathcal{K}_1) \cap \mathcal{O}(\mathcal{K}_2)$, then $\mathcal{K}_1 = \mathcal{K}_2$.*

Proof. A simple induction shows that each π_1-descendant of a_1 is a π_2-descendant of a_2, and vice versa. Hence both braids have the same elements. The second condition given in the lemma implies that both braids have the same linking functions. Hence $\mathcal{K}_1 = \mathcal{K}_2$. □

Corollary 4.7 *Let $\mathcal{K}_1$ and $\mathcal{K}_2$ be two prebraids and $\mathcal{K}_1 \subseteq \mathcal{K}_2$. Let $\mathcal{K}'_1$ be a subbraid of $\mathcal{K}_1$ and let $\mathcal{K}'_2$ be a subbraid of $\mathcal{K}_2$ such that $\mathcal{K}'_1$ and $\mathcal{K}'_2$ have the same root. Then $\mathcal{K}'_1 = \mathcal{K}'_2$.*

Proof. Since $\mathcal{K}'_1$ and $\mathcal{K}'_2$ are subbraids of $\mathcal{K}_2$ it is obvious that $\mathcal{K}'_1$ and $\mathcal{K}'_2$ satisfy the conditions of Lemma 4.6. Hence $\mathcal{K}_1 = \mathcal{K}_2$. □

Lemma 4.8 *For each element e of a prebraid $\mathcal{K}$ there exists a unique subbraid of $\mathcal{K}$ with root e.*

4.1 Braids and subbraids

Braids represent finite descriptions of possibly infinite objects that interweave a finite number of elements of two given SC-structures. Technically, suitable mappings, defined on atoms and pointing to non-atomic elements, are used to organize links between elements of distinct components that have to be interwoven.

Definition 4.2 Let $\mathcal{O}'_A \subseteq \mathcal{O}_A$, $\mathcal{O}'_B \subseteq \mathcal{O}_B$, let $\pi_A : \mathcal{O}'_A \to B$, $\pi_B : \mathcal{O}'_B \to A$, let $\pi := \pi_A \cup \pi_B$. An element $a \in A$ is *directly linked to* $b \in B$ via π if there is an $o \in \mathcal{O}_B(b)$ such that $a = \pi_B(o)$. Analogously $b \in B$ is directly linked to $a \in A$ via π if there exists an $o \in \mathcal{O}_A(a)$ such that $b = \pi_A(o)$. An element $a \in A \cup B$ is a *π-descendant of* $b \in A \cup B$ if there exists a sequence $a = a_0, a_1, \ldots, a_n = b$ $(n \geq 0)$ such that each a_i is directly linked to a_{i+1} via π, for $0 \leq i \leq n - 1$.

Definition 4.3 A *braid* of type A over $\mathcal{A}^\Sigma, \mathcal{B}^\Delta$ is a quintuple $\mathcal{K} = \langle a, C, D, \pi_A, \pi_B \rangle$, where

1. $a \in A \setminus \mathcal{O}_A$,
2. C is a finite subset of A containing a. All elements of $C \setminus \{a\}$ are *non-atomic*. D is a finite set of *non-atomic* elements of B,
3. $\pi_A : \mathcal{O}_A(C) \to D$ and $\pi_B : \mathcal{O}_B(D) \to C$ are mappings. For $\langle o, e \rangle \in \pi_A \cup \pi_B$, e is always a *non-atomic* element,
4. each element in $C \cup D$ is a π-descendant of a, for $\pi := \pi_A \cup \pi_B$.

The element a is called the *root* of $\mathcal{K}$. The elements in the sets C and D are called the *elements of $\mathcal{K}$ of type A and B* respectively. The functions π_A and π_B are called the *linking functions of $\mathcal{K}$ of type A and B* respectively.

Braids of type B, with root in $B \setminus \mathcal{O}_B$, are defined symmetrically. A braid $\mathcal{K}$ is called *trivial* if the root of $\mathcal{K}$ is a bottom atom $z \in Z$. In this case, z is the only element of the braid. It does not make sense to distinguish between the trivial braid $\langle z, \{z\}, \emptyset, \emptyset, \emptyset \rangle$ of type A and the trivial braid $\langle z, \emptyset, \{z\}, \emptyset, \emptyset \rangle$ of type B. We identify both braids. Hence, trivial braids have mixed type.

Example 4.4 The following figure represents a braid over two termalgebras, for signatures $\Sigma = \{f, a\}$ and $\Delta = \{g\}$ respectively.

E.g., quotient term algebras for collapse-free equational theories, rational tree algebras, feature structures, feature structures with arity, and the domains with nested, finite or rational lists (as mentioned in 3.4) are always non-collapsing.

4 The Domain of the Rational Amalgam

In this section we shall define the underlying domain of the rational amalgam of two non-collapsing SC-structures over disjoint signatures. This is the most complicated step of the rational amalgamation construction. The description would be much simpler if we would restrict the construction to components where the elements have a particular form (e.g., the form of trees). But such a restriction would contradict our motivation to describe general constructions. To define the elements of the rational amalgam, we shall first introduce the notion of a "braid" and its standard normal form. The set of braids in standard normal form will represent the carrier of the rational amalgam. We shall describe the rational amalgamation of two component structures. There are, however, no difficulties to interweave any finite number of components in the same way.

Throughout this section $(\mathcal{A}^{\Sigma}, \mathcal{M}, X)$ and $(\mathcal{B}^{\Delta}, \mathcal{N}, Y)$ denote two fixed non-collapsing SC-structures over disjoint signatures. We assume that the atom sets X and Y have the form $X = Z \uplus \mathcal{O}_A$ and $Y = Z \uplus \mathcal{O}_B$, where the sets $Z, \mathcal{O}_A$, and $\mathcal{O}_B$ are all infinite, and where $\mathcal{O}_A \cap \mathcal{O}_B = \emptyset$. The atoms in Z will be called *bottom atoms*, the atoms in $\mathcal{O}_A$ ($\mathcal{O}_B$) will be called *open atoms*. In the braid construction, the bottom atoms will play the role of ordinary atoms, or leaves. Open atoms, in contrast, can be considered as "named holes" that are only used to link elements of both structures. With $\mathcal{O}_A(a)$ and $\mathcal{O}_A(A')$ we denote the set of open atoms occurring in the stabilizer of $a \in A$ ($A' \subseteq A$) with respect to $\mathcal{M}$. Similarly expressions $\mathcal{O}_B(b)$ ($\mathcal{O}_B(B')$) are used to denote the set of open atoms occurring in the stabilizer of $b \in B$ ($B' \subseteq B$) with respect to $\mathcal{N}$.

An endomorphism $m \in \mathcal{M}$ ($n \in \mathcal{N}$) is called *admissible* if m (n) leaves all bottom atoms $z \in Z$ fixed and if $m(o) \in \mathcal{O}_A$ ($n(o) \in \mathcal{O}_B$) for all $o \in \mathcal{O}_A$ ($o \in \mathcal{O}_B$).[3] Automorphisms are called admissible if they define a permutation of the set of open atoms while leaving bottom atoms fixed. A pair $(m, n) \in \mathcal{M} \times \mathcal{N}$ is called admissible if both m and n are admissible.

Lemma 4.1 *Let $A' \subseteq A$. If the admissible endomorphisms $m_1, m_2 \in \mathcal{M}$ coincide on $\mathcal{O}_A(A')$, then m_1 and m_2 coincide on A'. Similarly, let $B' \subseteq B$. If the admissible endomorphisms $n_1, n_2 \in \mathcal{N}$ coincide on $\mathcal{O}_B(B')$, then n_1 and n_2 coincide on B'.*

[3] Intuitively, admissible endomorphisms cause a "renaming" of open atoms, compare Lemma 3.10. They may identify distinct open atoms.

A fundamental property of SC-structures is the following ([BS95], Lemma 13).

Lemma 3.7 *For each $a \in A$ there exists a* unique minimal *finite subset Y of X such that $\{a\} \in SH^{\mathcal{A}}_{\mathcal{M}}(Y)$.*

Definition 3.8 The *stabilizer of $a \in A$ with respect to $\mathcal{M}$*[2], $Stab^{\mathcal{A}}_{\mathcal{M}}(a)$ is the unique minimal finite subset Y of X such that $a \in SH^{\mathcal{A}}_{\mathcal{M}}(Y)$. The stabilizer of $A' \subseteq A$ is the set $Stab^{\mathcal{A}}_{\mathcal{M}}(A') := \bigcup_{a \in A'} Stab^{\mathcal{A}}_{\mathcal{M}}(a)$.

For the mathematical treatment of SC-structures, the concept of the stabilizer turns out to be extremely useful. It might give a good intuition to imagine that the stabilizer of an element a is the set of atoms "occurring" in a. Note, however, that "the" set of atoms (variables) occurring, e.g., in distinct terms that represent the same element of a quotient term algebra is *not* unique in general. It is trivial to see that $SH^{\mathcal{A}}_{M}(Y) = \{a \in A \mid Stab^{\mathcal{A}}_{M}(a) \subseteq Y\}$, for each $Y \subseteq X$. In the sequel, further properties of stabilizers will be used. The first lemma is a trivial consequence of the fact that stable hulls are Σ-substructures.

Lemma 3.9 *Let $f \in \Sigma$ be an n-place operator and $a_1, \ldots, a_n \in A$. Then $Stab^{\mathcal{A}}_{\mathcal{M}}(f_{\mathcal{A}}(a_1, \ldots, a_n)) \subseteq Stab^{\mathcal{A}}_{\mathcal{M}}(\{a_1, \ldots a_n\})$.*

The next lemma plays a crucial role in the rational amalgamation construction. It will be used in many proofs.

Lemma 3.10 *Let $m \in \mathcal{M}$ be an endomorphism of the SC-structure $(\mathcal{A}^{\Sigma}, \mathcal{M}, X)$ such that the restriction of m on X is a mapping $X \to X$. If $Stab^{\mathcal{A}}_{\mathcal{M}}(a) = \{x_1, \ldots, x_k\}$, then $Stab^{\mathcal{A}}_{\mathcal{M}}(m(a)) \subseteq \{m(x_1), \ldots, m(x_k)\}$. If m is an automorphism, then $Stab^{\mathcal{A}}_{\mathcal{M}}(m(a)) = \{m(x_1), \ldots, m(x_k)\}$.*

Proof. Let m_1 and m_2 be two endomorphisms in $\mathcal{M}$ that coincide on $\{m(x_1), \ldots, m(x_k)\} \subseteq X$. Then $m \circ m_1$ and $m \circ m_2$ are endomorphisms in $\mathcal{M}$ that coincide on $\{x_1, \ldots, x_k\}$. By assumption, $m \circ m_1$ and $m \circ m_2$ coincide on a. But then m_1 and m_2 coincide on $m(a)$. Hence $Stab^{\mathcal{A}}_{\mathcal{M}}(m(a)) \subseteq \{m(x_1), \ldots, m(x_k)\}$. Assume that m is an automorphism, and that $Stab^{\mathcal{A}}_{\mathcal{M}}(m(a))$ is a proper subset of $\{m(x_1), \ldots, m(x_k)\}$. The first part of the lemma, applied to m^{-1}, yields a proper subset of $\{x_1, \ldots, x_k\}$ that stabilizes a, which is impossible, by choice of $\{x_1, \ldots, x_k\}$. ☐

We may now characterize the subclass of SC-structures for which we can use the rational amalgamation construction.

Definition 3.11 An SC-structure $(\mathcal{A}^{\Sigma}, \mathcal{M}, X)$ is *non-collapsing* if every endomorphism $m \in \mathcal{M}$ maps non-atoms to non-atoms (i.e., $m(a) \in A \setminus X$ for all $a \in A \setminus X$ and all $m \in \mathcal{M}$).

[2]Whenever the monoid $\mathcal{M}$ is clear from the context, we shall not mention it.

$A_0 \subseteq SH^{\mathcal{A}}_{\mathcal{M}}(A_0)$. In general, however, the stable hull can be larger than the generated subalgebra. For example, if $\mathcal{A}^\Sigma := R(\Sigma, X)$ denotes the algebra of rational trees over signature Σ, if $\mathcal{M} = End^\Sigma_{\mathcal{A}}$, and if $Y \subseteq X$ is a subset of the set of variables, X, then $SH^{\mathcal{A}}_{\mathcal{M}}(Y)$ contains all *rational* trees with variables in Y, while Y generates all *finite* trees with variables in Y only.

Definition 3.2 The set $X \subseteq A$ is an *$\mathcal{M}$-atom set for $\mathcal{A}^\Sigma$* if every mapping $X \rightarrow A$ can be extended to an endomorphism in $\mathcal{M}$.

Definition 3.3 A countably infinite Σ-structure $\mathcal{A}^\Sigma$ is an *SC-structure* (simply combinable structure) iff there exists a submonoid $\mathcal{M}$ of $End^\Sigma_{\mathcal{A}}$ such that $\mathcal{A}^\Sigma$ has an infinite $\mathcal{M}$-atom set X where every $a \in A$ is stabilized by a finite subset of X with respect to $\mathcal{M}$. We denote this SC-structure by $(\mathcal{A}^\Sigma, \mathcal{M}, X)$. If $\mathcal{M} = End^\Sigma_{\mathcal{A}}$, then $(\mathcal{A}^\Sigma, End^\Sigma_{\mathcal{A}}, X)$ is called a *strong SC-structure*.

Examples 3.4 The class of SC-structures contains, e.g., all free structures (see, e.g., [Ma71]), rational tree algebras ([Co84, Ma88]), feature structures (for specificity, we refer to [AP94, BT94]), feature structures with arity ([ST94, BT94]), domains with nested, finite or rational lists (rational lists are used in Prolog III, see [Co90]), and domains with nested, finite or rational sets (as introduced in [Ac88] and used in [Ro88]). In each case, we have to take the non-ground variant since we assume the existence of a countably infinite set of atoms. With the exception of feature structures, all these structures are strong SC-structures. For details we refer to [BS95].

In the rest of this section, $(\mathcal{A}^\Sigma, \mathcal{M}, X)$ denotes a fixed SC-structure with carrier A.

Lemma 3.5 *Let $\varphi(v_1, \ldots, v_k)$ be a positive Σ-formula, let $m \in \mathcal{M}$, and let $a_1, \ldots, a_k$ be elements of A. Then $\mathcal{A}^\Sigma \models \varphi(v_1/a_1, \ldots, v_k/a_k)$ implies $\mathcal{A}^\Sigma \models \varphi(v_1/m(a_1), \ldots, v_k/m(a_k))$.*

Proof. It is simple to see that there exists a surjective endomorphism $m' \in \mathcal{M}$ that coincides with m on $\{a_1, \ldots, a_k\}$. The result follows from the well-known fact that validity of positive formulae is preserved under surjective homomorphisms. □

Lemma 3.6 *Let $\varphi(v_1, \ldots, v_k)$ be a positive Σ-formula, let $m \in \mathcal{M}$, and let $x_1, \ldots, x_k$ be distinct atoms in X. Then $\mathcal{A}^\Sigma \models \varphi(v_1/x_1, \ldots, v_k/x_k)$ implies $\mathcal{A}^\Sigma \models \forall v_1, \ldots \forall x_k\ \varphi$.*

Proof. Let $a_1, \ldots, a_k$ be arbitrary elements of A. Since X is an $\mathcal{M}$-atom set, there exists an $m \in \mathcal{M}$ such that $m(x_i) = a_i$, for $i = 1, \ldots, k$. Hence $\mathcal{A}^\Sigma \models \varphi(v_1/x_1, \ldots, v_k/x_k)$ implies $\mathcal{A}^\Sigma \models \varphi(v_1/a_1, \ldots, v_k/a_k)$, by Lemma 3.5. It follows that $\mathcal{A}^\Sigma \models \forall v_1 \ldots \forall x_k\ \varphi$. □

tion, and a related construction called "infinite amalgamation" seem to be the most important combination principles among a whole spectrum of related methods.

2 Preliminaries

A signature Σ consists of a set Σ_F of function symbols and a disjoint set Σ_P of predicate symbols (not containing "="), each of fixed arity. Expressions $\mathcal{A}^\Sigma$ denote Σ-structures over the carrier set A, and $f_\mathcal{A}$ ($p_\mathcal{A}$) stands for the interpretation of $f \in \Sigma_F$ ($p \in \Sigma_P$) in $\mathcal{A}^\Sigma$. Σ-terms ($t, t_1, \ldots$) and atomic Σ-formulas (of the form $t_1 = t_2$, or of the form $p(t_1, \ldots, t_n)$) are built as usual. A Σ-formula φ is written in the form $\varphi(v_1, \ldots, v_n)$ in order to indicate that the free variables of φ are in $\{v_1, \ldots, v_n\}$. We write $\mathcal{A}^\Sigma \models \varphi(v_1/a_1, \ldots, v_n/a_n)$ if φ becomes true in $\mathcal{A}^\Sigma$ under all assignments that map v_i to $a_i \in A$, for $1 \leq i \leq n$.

A Σ-homomorphism is a mapping h between two structures $\mathcal{A}^\Sigma$ and $\mathcal{B}^\Sigma$ such that $h(f_\mathcal{A}(a_1, \ldots, a_n)) = f_\mathcal{B}(h(a_1), \ldots, h(a_n))$ and $p_\mathcal{A}[a_1, \ldots, a_n]$ implies that $p_\mathcal{B}[h(a_1), \ldots, h(a_n)]$ for all $f \in \Sigma_F$, $p \in \Sigma_P$, and $a_1, \ldots, a_n \in A$. A Σ-isomorphism is a bijective Σ-homomorphism $h : \mathcal{A}^\Sigma \to \mathcal{B}^\Sigma$ such that $p_\mathcal{A}[a_1, \ldots, a_n]$ if, and only if, $p_\mathcal{B}[h(a_1), \ldots, h(a_n)]$, for all $a_1, \ldots, a_n \in A$. A Σ-endomorphism of $\mathcal{A}^\Sigma$ is a homomorphism $h^\Sigma : \mathcal{A}^\Sigma \to \mathcal{A}^\Sigma$. With $End_\mathcal{A}^\Sigma$ we denote the monoid of all endomorphisms of $\mathcal{A}^\Sigma$, with composition as operation.

If $g : A \to B$ and $h : B \to C$ are mappings, then $g \circ h : A \to C$ denotes their composition. Expressions like $\vec{v}, \vec{a}$ are used to denote finite sequences. If $\vec{a} = a_1, \ldots, a_n$ is a sequence of elements of A and if m is a mapping with domain A, then $m(\vec{a})$ denotes the sequence $m(a_1), \ldots, m(a_n)$. If $\vec{v} = v_1, \ldots, v_n$, then $\mathcal{A}^\Sigma \models \varphi(\vec{v}/\vec{a})$ is shorthand for $\mathcal{A}^\Sigma \models \varphi(v_1/a_1, \ldots, v_n/a_n)$. The symbol "$\uplus$" denotes disjoint set union.

3 Non-collapsing SC-structures

In this section we shall introduce the class of structures for which we can use the rational amalgamation construction (Definition 3.11). First we recall the definition given in [BS95]. In the sequel, we consider a fixed Σ-structure $\mathcal{A}^\Sigma$, and $\mathcal{M}$ denotes a submonoid of $End_\mathcal{A}^\Sigma$.

Definition 3.1 Let A_0, A_1 be subsets of $\mathcal{A}^\Sigma$. Then A_0 *stabilizes* A_1 *with respect to* $\mathcal{M}$ iff all elements m_1 and m_2 of $\mathcal{M}$ that coincide on A_0 also coincide on A_1. For $A_0 \subseteq A$ the *stable hull of* A_0 *with respect to* $\mathcal{M}$ is the set

$$SH_\mathcal{M}^\mathcal{A}(A_0) := \{a \in A \mid A_0 \text{ stabilizes } \{a\} \text{ with respect to } \mathcal{M}\}.$$

The stable hull of a set A_0 has properties that are similar to those of the subalgebra generated by A_0: $SH_\mathcal{M}^\mathcal{A}(A_0)$ is always a Σ-substructure of $\mathcal{A}^\Sigma$, and represents a most general object among all structures that can be considered as a reasonable combination of the two components. For a large class of component structures—so-called SC-structures—an explicit construction of the free amalgamated product of two components was given and it was shown how given constraint solvers for the component structures can be combined to a constraint solver for the free amalgam.

In the present paper we introduce a second systematic way to combine constraint systems over SC-structures, called *rational amalgamation*. Free and rational amalgamation both yield a combined structure with "mixed" elements that interweave a finite number of "pure" elements of the two components in a particular way. The difference between both constructions becomes transparent when we ignore the interior structure of these pure subelements and consider them as construction units with a fixed arity, similar to "complex function symbols". Under this perspective, and ignoring details, mixed elements of the free amalgam can be considered as finite trees, whereas mixed elements of the rational amalgam are like rational trees.

Mixed element of free amalgam (1) and of rational amalgam (2).

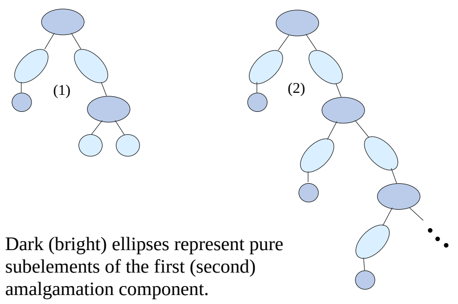

On this background it should not be surprising that in praxis rational amalgamation appears to be the preferred combination principle in situations where the two solution structures to be combined are themselves "rational" or "cyclic" domains: for example, it represents the way how rational trees and rational lists are interwoven in the solution domain of Prolog III ([Co90]), and a variant of rational amalgamation has been used to combine feature structures with non-wellfounded sets in a system introduced by W. Rounds [Ro88].

We introduce rational amalgamation as a general construction that can be used to combine so-called non-collapsing SC-structures over disjoint signatures. It is then shown how constraint solving in the rational amalgam can be reduced to constraint solving in the components. The decomposition scheme that is used is closely related to the decomposition algorithm for free amalgamation, but it avoids one highly non-deterministic step that is needed in the latter scheme. Hence, when matters of efficiency become important, rational amalgamation might be the better choice.

In the conclusion we sketch a more general view of the problem of combining constraint systems. From our present perspective, free and rational amalgama-

Combination of Constraint Systems II: Rational Amalgamation*

Klaus U. Schulz, Stephan Kepser
CIS, Universität München
Wagmüllerstraße 23, 80538 München, Germany
e-mail: schulz/kepser@cis.uni-muenchen.de
WWW: http://www.cis.uni-muenchen.de

Abstract

In a recent paper[1], the concept of "free amalgamation" has been introduced as a general methodology for interweaving solution structures for symbolic constraints, and it was shown how constraint solvers for two components can be lifted to a constraint solver for the free amalgam. Here we discuss a second general way for combining solution domains, called *rational amalgamation*. In praxis, rational amalgamation seems to be the preferred combination principle if the two solution structures to be combined are "rational" or "non-wellfounded" domains. It represents, e.g., the way how rational trees and rational lists are interwoven in the solution domain of Prolog III, and a variant has been used by W. Rounds for combining feature structures and hereditarily finite non-wellfounded sets. We show that rational amalgamation is a general combination principle, applicable to a large class of structures. As in the case of free amalgamation, constraint solvers for two component structures can be combined to a constraint solver for their rational amalgam. From this algorithmic point of view, rational amalgamation seems to be interesting since the combination technique for rational amalgamation avoids one source of non-determinism that is needed in the corresponding scheme for free amalgamation.

1 Introduction

The present paper, as its predecessor [BS95], marks one step in a program where we try to characterize the most important *general constructions* for combining solution domains and constraint solvers for symbolic constraints. In [BS95] the notion of the *free amalgamated product* of two component structures was introduced. This product is characterized by a universality-property: it

*This work was supported by a DFG grant (SSP "Deduktion") and by the EC Working Group CCL, EP6028.

[1]see [BS95].