

On the Definition of Operation-based Merging

1 Informal Abstract/Intro Outline

1.1 Background/Motivation

Existing definitions suffer from one or more of the following problems¹:

A) They are implicit. Most of the time they are “definition by solution” of the form: “here is my algorithm that has postcondition X (often not stated at all), therefore X is the definition of the (merging) problem”.

B) They are unrealizable (algorithmically), such as Sun et al. [1998, def. 3 and 4.3]²: “the *intention* of each operation must be preserved at every replica”. Equivalent to “the computer(s) must read my mind”.

C) They lead to one or both of the following paradoxes:

(C.1) The (general) paradox of operation-based merging: more information (about intermediate operations/states) results in more, not fewer, conflicts. It is implied by those definitions that require commutations of the exact operations involved.

(C.2) The paradox of lost *future* updates. A (past) operation guarantees that a certain subset of future operations will be lost upon a merge³. Concrete example: darcs.

The *weave* concept of Lippe and van Oosterom [1992] suffers from both paradoxes.

D) They are intricately tied to a particular setting, making generalization not immediate. We will invoke J. D. H. Smith’s often cited quote here: “What looks messy and complicated in a particular framework may turn out to be simple and obvious in the proper general one” [Smith, 1976, p. iv].

A number of optimistic replication systems attempt to provide a generic framework for merging. Some systems allow the user to formulate a merging algorithm (for a particular type of input) in a dedicated language [Terry et al., 1995]. Other systems allow the user to state some postconditions for the merge, which are then used as solution selectors for a generic algorithm that searches the state space [Kermarrec et al., 2001].

Debating the merits of these types of frameworks goes beyond the scope of this paper. Instead, we note that the problem

¹I realize that some if not all of this stuff is very “bare knuckles”, but nevertheless true. I need to “preen” it so it won’t sound so shocking (w.r.t. what gets published as research in some reputable places).

²Not an exact quote. Their original definition is far more convoluted, but means exactly this. Also, they assume that an observed ordering at a replica implies a dependency (def. 1 in their paper), which is simply false.

³We take the liberty of using “merge” as a noun, synonym with reconciliation. This is a common practice in the revision control community.

addressed in this paper is orthogonal⁴ to the existence of such frameworks. Before the user can hope to formulate either an algorithm or even a postcondition for a merging problem (on a particular type of input), the user must have a notion (definition) of what merging actually means.

1.2 Contributions

Since a definition by definition (sic) cannot be proved, we hope to convince that our definition of operation-based merging should be accepted by showing that it avoids the shortcomings previously enumerated.

We formulate our definition in the general setting of a transformation monoid⁵, also known as an M -act, which we subclass to model conflicts — the *merge monoid*. In order appearance, the strengths of our definition are:

1. We avoid the (general) paradox of operation-based merging by giving the definition of a conflict-free merge in terms of *safe embeddings* that commute. The necessity of the safe embedding notion has already been observed by Ramsey and Csirmaz [2001] in the context of filesystem merging. They use the rather unfortunate term “approximation” for it. We merely generalize their notion.

2. We avoid the paradox of lost future updates by making the following key observations about unification — the elimination of duplicate operations:

- (i) Unification is an automated *resolution* mechanism for *false-conflicts*. It is *not* applicable if there exists a conflict-free merge.
- (ii) In turn, false-conflicts can only be defined in terms of *zeropotent* operations. Roughly speaking, a zeropotent operation gives “conflict everywhere” when composed with itself.

Examples and counterexamples for the above, as well as their precise formulation, are provided in section 2.4.

3. Rather surprisingly, a useful definition of dependency can be given without making any additional assumption about the nature of the states. In our framework we are able to prove the intuitive truths that dependencies preclude commutation and vice versa.

Our notion of dependency is more discriminating than the event-based happened-before relation of Lamport [1978], in the sense that an operation β may happen after another operation α everywhere, but this need not imply that β depends

⁴Frankly, it is a precondition for their usefulness. But, I want this paper to get published at PODC.

⁵For the reader unfamiliar with abstract algebra, we introduce the notions used in this paper at the beginning of the next section.

on α . Although this distinction may be plain obvious to the reader, some previous research on merging applied Lamport’s celebrated result inappropriately [Sun et al., 1998, def. 1].

4. We introduce the δ -act algebraic structure, which captures the basic behavior of the text merging programs patch and diff.

We also introduce an algebraic three-way merge operator that needs to satisfy two simple axioms⁶, which can be trivially verified to hold for important algorithms. E.g., they can be verified for the widely-used⁷ diff3 text merging algorithm based on the excellent analysis due to Khanna, Kunal, and Pierce [2007].

We show that a three-way merge operator can be used to define a subclass of δ -act, which we call 3-act. In other words, a three-way operator implicitly defines a patch and delta functions. Although we formalize it, this fact has been known for quite some time. It is for instance exploited by RCS [Tichy, 1985] to provide selective merges via the rcsmerge command, equivalent (except in rare circumstances) to a diff followed by a patch.

5. A 3-act can be embedded in a merge monoid. Furthermore, if a three-way merge succeeds (the result is not a conflict), then it meets our operation-based definition of a conflict-free merge.

This result may seem surprising, since a three-way merge is presented by most authors as the prototypical state-based approach to merging, in (sharp) contrast to operation-based merging. The insight here is that many three-way merge algorithms can also be interpreted as a reconciling two sequences of operations which are automatically derived from the three states. Key to this connection is the notion of safe embedding. This connection extends the “cover” of our definition to an important class of merges typically described as state-based.

2 Main Body

2.1 Preliminaries

We recall some basic notions and facts from abstract algebra.

A monoid (M, \cdot) is a set together with an associative binary operation that has an identity element.

Given a set S , the set of functions mapping S to itself, $S^S := \{f: S \rightarrow S\}$ forms a monoid under function composition. These functions are called transformations or, from a category theoretic perspective, endomorphisms. Although the term (unary) operation is also justifiable from the perspective of universal algebra, we shall prefer the term transformation for the rest of this paper.

As matter of notation, we will write $(f; g)(x) := g(f(x))$ for function composition since it captures more naturally the historical order of a sequence of transformations.

Given a monoid (M, \cdot) , a submonoid is a subset $N \subseteq M$ that is closed under the monoid’s operation: $\forall x, y \in N \Rightarrow x \cdot y \in N$.

⁶It should not be construed that these are sufficient for a *useful* three-way merge algorithm. They are only required properties.

⁷Used in RCS [Tichy, 1985], CVS [Berliner, 1990] and many other revision control systems.

For our purposes, it suffices to define a transformation monoid (S, M) as a set S together with a submonoid $M \subseteq S^S$, in other words, a subset of the transformations on S that is closed under function composition. For a more general definition see Barr and Wells [1995, chap. 3.2].

2.2 The merge monoid

Let \mathcal{Q} be a set of states, not necessarily finite. Let $\mathcal{Q}_\perp := \mathcal{Q} \cup \{\perp\}$ be the extension of \mathcal{Q} with the “conflict” element $\perp \notin \mathcal{Q}$. Let $\mathcal{Q}_\perp^\mathcal{Q} := \{f: \mathcal{Q}_\perp \rightarrow \mathcal{Q}_\perp \mid f(\perp) = \perp\}$ be the set of all conflict-preserving transformations – endomorphisms on \mathcal{Q}_\perp for which \perp is a fixed point. The rationale for imposing this restriction is that we do not want “normal” transformations to be able to solve conflicts. Since $\mathcal{Q}_\perp^\mathcal{Q}$ is closed under function composition, it is a submonoid of the monoid of all transformations on \mathcal{Q}_\perp . Therefore, $(\mathcal{Q}_\perp, \mathcal{Q}_\perp^\mathcal{Q})$ meets the definition of a transformation monoid, and we shall call it *the full transformation monoid with conflicts*.

In applications we are typically interested only in certain classes of transformation, e.g. for text files we may be interested only in line insertions and deletions. Thus, it makes sense to consider only the transformations from a submonoid \mathcal{T} of $\mathcal{Q}_\perp^\mathcal{Q}$. Let $(\mathcal{Q}_\perp, \mathcal{T})$ be known as a *transformation monoid with conflicts* or, more succinctly, a *merge monoid*.

2.3 Conflict-free merges

Let $\alpha = (\alpha_m)$ and $\beta = (\beta_n)$ be strings (finite sequences) of functions from \mathcal{T} , formally elements of \mathcal{T}^* . A string of transformations is applied to an element $X \in \mathcal{Q}_\perp$ using the expected semantics of function composition: $\alpha(X) := (\alpha_1; \dots; \alpha_m)(X)$. By convention the empty string ε is applied as the identity function on $\mathcal{Q}_\perp^\mathcal{Q}$.

We can attempt a definition of a conflict-free merge in terms of commuting transformation sequences. Let O, A and B be three states from \mathcal{Q} , such that $A = \alpha(O)$ and $B = \beta(O)$. We can claim that $\{\alpha, \beta\}$ form a conflict-free merge at O with the result M iff $M = \alpha\beta(O) = \beta\alpha(O)$, or more explicitly

$$M = (\alpha_1; \dots; \alpha_m; \beta_1; \dots; \beta_n)(O) = (\beta_1; \dots; \beta_n; \alpha_1; \dots; \alpha_m)(O)$$

This tentative definition is unsatisfactory. The following example will illustrate its deficiency.

Example 1. Let $U = \{a, b, \dots\}$ be a universe of “atoms”, and let the states be subsets of U , i.e. $\mathcal{Q} = \wp(U)$. We define a family⁸ of transformations, inspired from the behavior of the patch program, that successfully insert and respectively remove an atom a from a state $S \in \mathcal{Q}$ iff a is and respectively is not included in S , producing a conflict otherwise:

$$i_a(S) := \begin{cases} S \cup \{a\} & \text{if } a \notin S \\ \perp & \text{if } a \in S \end{cases} \quad d_a(S) := \begin{cases} S \setminus \{a\} & \text{if } a \in S \\ \perp & \text{if } a \notin S \end{cases}$$

The above two transformations also need to map \perp to \perp in order to be elements of $\mathcal{Q}_\perp^\mathcal{Q}$. Unless otherwise specified, assume this mapping to happen in all examples.

⁸A family of this kind can be defined as the set of all partial applications of a two-argument function.

Observe that if $a \notin S$, then $(i_a; \mathfrak{d}_a)(S) = S$. Now consider merging the sequences of operations $\alpha = (i_a; \mathfrak{d}_a)$ and $\beta = i_a$ starting from the initial state $O = \emptyset$. We obtain:

$$\alpha\beta(O) = (i_a; \mathfrak{d}_a; i_a)(\emptyset) = \{a\} \neq \beta\alpha(O) = (i_a; i_a; \mathfrak{d}_a)(\emptyset) = \perp$$

It is unsatisfactory that we have to declare a conflict when the sequence $i_a; \mathfrak{d}_a$ is equivalent to the identity function when applied to O . One can interpret the above result as a paradox⁹ in the following way: more information about the operations performed (or equivalently about intermediate states) leads to more conflicts, not fewer as we would hope.

In order to avoid this paradox, we attempt to give a definition that does not use the exact sequences α and β , but rather any other pair of transformations $\hat{\alpha}$ and $\hat{\beta}$ from \mathcal{T} which have the same observable effect¹⁰ on O , more precisely $\hat{\alpha}(O) = \alpha(O)$ and $\hat{\beta}(O) = \beta(O)$, but unlike the original sequences have the advantage that they commute at O producing the result of the merge $M = \hat{\alpha}\hat{\beta}(O) = \hat{\beta}\hat{\alpha}(O)$. To continue our example, we can choose $\hat{\alpha} = \varepsilon$ and $\hat{\beta} = i_a$ and get result we'd expect $M = i_a(O)$.

Unfortunately, by fixing only two points for $\hat{\alpha}$ and $\hat{\beta}$, we've relaxed the definition too much, as we shall argue in the following example:

Example 2. Let $\mathcal{T} = \mathcal{Q}_\perp^{\mathcal{Q}}$ and we want to merge $\alpha = i_a$ with $\beta = i_b$ starting from $O = \emptyset$. We can choose the following two transformations as “solution”:

$$\varphi(S) = \begin{cases} \{a\} & \text{if } S = \emptyset \\ \emptyset & \text{otherwise} \end{cases} \quad \psi(S) = \begin{cases} \{b\} & \text{if } S = \emptyset \\ \emptyset & \text{otherwise} \end{cases}$$

Trivially $\alpha(O) = \varphi(O) = \{a\}$ and $\beta(O) = \psi(O) = \{b\}$. Since $\psi(\varphi(O)) = \varphi(\psi(O)) = \emptyset$, according to our tentative definition, we'd have to accept that \emptyset is a valid result for the merge. Clearly, we'd be hard pressed to sell this to a discerning user.

Intuitively, we need to coerce the $\hat{\alpha}$ and $\hat{\beta}$ to have the same effect as α and β not only at O but also for any argument that is “ O -like”. With this additional requirement, φ would not be able to “impersonate” α at $\{b\}$ in the above example.

We now formalize “ O -like” in a way that allows the useful choices for $\hat{\alpha}$ and $\hat{\beta}$ that we made in example 1, but disallowing the “solution” for example 2.

Definition 1 (safe embedding). Let $D \subseteq \mathcal{Q}_\perp$ be a subset of the domain of transformations. A transformation f is safely embedded in another transformation g over D , and we write $f \sqsubseteq_D g$, iff

$$\forall x \in D, f(x) \neq \perp \Rightarrow g(x) = f(x)$$

Informally, g safely embeds f over D if g is a “rewriting” of f that “breaks less often” over D . If g safely embeds f on the entire \mathcal{Q} , we can unambiguously write $f \sqsubseteq g$.

⁹The (general) paradox of operation-based merging.

¹⁰Depending on how loose of an analogy you wish to admit, you can see these functions as approximations, or even “weak bisimulations”.

We have already seen both examples and counterexamples of safe embeddings. In example 1, $(i_a; \mathfrak{d}_a) \sqsubseteq \varepsilon$. In example 2 however, $i_a(\{b\}) = \{a, b\} \neq \varphi(\{b\}) = \emptyset$, thus $i_a \not\sqsubseteq \varphi$.

Recall that given a function $f: X \rightarrow Y$, the preimage of a subset of the codomain $B \subseteq Y$, written¹¹ $f^{-1}[B]$ is the subset of the domain $\{x \in X \mid f(x) \in B\}$. Also recall that the equaliser of two functions f and g both defined on the same domain X and codomain Y is the subset of the domain $\text{Eq}(f, g) := \{x \in X \mid f(x) = g(x)\}$. Using these notions, we can give the following trivial characterization of safe embedding:

Proposition 1. For all transformations α, β :

$$\begin{aligned} \alpha \sqsubseteq_D \beta &\iff \alpha^{-1}[\mathcal{Q}] \cap D \subseteq \text{Eq}(\alpha, \beta) \\ \alpha \sqsubseteq \beta &\iff \alpha \sqsubseteq_{\alpha^{-1}[\mathcal{Q}]} \beta \iff \alpha^{-1}[\mathcal{Q}] \subseteq \text{Eq}(\alpha, \beta) \\ \alpha \sqsubseteq_D \beta \text{ and } \beta \sqsubseteq_D \alpha &\iff D \subseteq \text{Eq}(\alpha, \beta) \end{aligned}$$

Now observe the following: we are only interested in merging states from \mathcal{Q} , in other words elements of \mathcal{Q}_\perp other than \perp . If $\alpha(O) = A$ and $\beta(O) = B$, then if we impose the condition that $A, B \in \mathcal{Q}$, we obtain a restriction on what “ O -like” states can be, which can be expressed in terms of preimages as $O \in \alpha^{-1}[\mathcal{Q}] \cap \beta^{-1}[\mathcal{Q}]$.

Definition 2 (conflict-free merge). Let $\alpha, \beta \in \mathcal{T}$. A conflict-free merge of α and β is a pair of transformations $\hat{\alpha}, \hat{\beta} \in \mathcal{T}$ with the following properties:

- i) $\alpha \sqsubseteq \hat{\alpha}$ and $\beta \sqsubseteq \hat{\beta}$
- ii) $\forall X \in \alpha^{-1}[\mathcal{Q}] \cap \beta^{-1}[\mathcal{Q}], \hat{\beta}(\hat{\alpha}(X)) = \hat{\alpha}(\hat{\beta}(X)) \neq \perp$

The first requirement in the above definition cannot be relaxed “even one bit” without admitting solutions such as the one from example 2. To illustrate this, assume that we'd replace (i) with $\alpha \sqsubseteq_C \hat{\alpha}$ and $\beta \sqsubseteq_C \hat{\beta}$, where $C = \alpha^{-1}[\mathcal{Q}] \cap \beta^{-1}[\mathcal{Q}]$. This may seem innocuous enough because C is exactly the set over which X varies in the second requirement.

explain me better

Example 3. Let $\mathcal{T} = \mathcal{Q}_\perp^{\mathcal{Q}}$ and we want to merge $\alpha = i_a$ with $\beta = i_b$ starting from $O = \emptyset$ (same as example 2). The proposed solution is:

$$\varphi(S) = \begin{cases} S \cup \{a\} & \text{if } \{a, b\} \not\subseteq S \\ \emptyset & \text{otherwise} \end{cases} \quad \psi(S) = \begin{cases} S \cup \{b\} & \text{if } \{a, b\} \not\subseteq S \\ \emptyset & \text{otherwise} \end{cases}$$

Observe that $C = \alpha^{-1}[\mathcal{Q}] \cap \beta^{-1}[\mathcal{Q}] = \{S \mid \{a, b\} \not\subseteq S\}$ and $\alpha \sqsubseteq_C \varphi$ and also $\beta \sqsubseteq_C \psi$. But $\psi(\varphi(\emptyset)) = \psi(\{a\}) = \emptyset$ and $\varphi(\psi(\emptyset)) = \varphi(\{b\}) = \emptyset$, which is again difficult to accept as a merge of two insertions. On the other hand $\alpha \not\sqsubseteq \varphi$ (and $\beta \not\sqsubseteq \psi$), so φ and ψ do not constitute a merge according to definition 2.

¹¹We use the less often encountered notation with square brackets because most examples involve sets as function arguments.

Open Question 1. What happens if we replace (ii) in definition 2 with a “point-based” property: “Given α, β and $O \in \alpha^{-1}[\mathcal{Q}] \cap \beta^{-1}[\mathcal{Q}]$ a conflict-free merge is $\dot{\alpha}$ and $\dot{\beta}$ s.t. (i) holds and $\dot{\beta}(\dot{\alpha}(O)) = \dot{\alpha}(\dot{\beta}(O)) \neq \perp$ ”. Is this equivalent to definition 2 or not?

2.4 Solving false conflicts

Relying just on conflict-free merges is not entirely satisfactory. For instance, the merge of $\alpha = i_a$ with $\beta = i_a$ does not admit a conflict-free solution because we cannot conceivably embed i_a in anything but itself, and $(i_a; i_a)(S) = \perp$ for any S .

The solution may seem blatantly obvious in this case: keep just one i_a . Unfortunately, we cannot easily generalize this approach. Consider the following counterexample where $\mathcal{Q} = \mathbb{N}$ and $\text{succ}(n) := n + 1$. If $\alpha = \beta = \text{succ}$, then there is a conflict-free merge, namely $(\text{succ}; \text{succ})$. Eliminating the duplicate succ here does not seem justifiable. We note here that using Lippe and Oosterom’s weave concept, the merge would be just succ .

Let \perp be the transformation that always results in \perp , i.e. $\forall X \in \mathcal{Q}, \perp(X) = \perp$. Note that \perp is an absorbing element¹², meaning that for any transformation α , $(\alpha; \perp) = (\perp; \alpha) = \perp$.

Definition 3 (zeropotent¹³ transformation). A transformation α is called zeropotent iff $(\alpha; \alpha) = \perp$.

As examples, both i_a and \mathfrak{d}_a are zeropotent, while succ is not zeropotent.

It may seem that if a transformation is zeropotent we may always eliminate duplicates. This is not exactly true. Even if α is zeropotent, nothing can be said in general about the sequence $(\alpha; \beta; \alpha)$. For instance if $\alpha = i_a$ and $\beta = \mathfrak{d}_a$ then $(i_a; \mathfrak{d}_a; i_a)(\emptyset) = \{a\} \neq \perp$.

Simply discarding duplicate zeropotent operations would result in a surprising change of result for the merge in example 1, where $\alpha = (i_a; \mathfrak{d}_a)$ and $\beta = i_a$. If it were permissible to discard a duplicate i_a , the merge result could be just $(i_a; \mathfrak{d}_a)(\emptyset) = \emptyset$.

We argue that merges such as the above are invalid, because they lead to the following paradox: a transformation such as \mathfrak{d}_a from β can “cancel out” transformations that do not necessarily precede it, i.e. i_a from β . In the typical scenario of lost updates, the *last* update wins (and the other are lost). In this incorrect merging example however, an update that may not even have happened, namely i_a from β , is guaranteed to be lost upon the merge. That’s why we call this scenario a lost *future* update.

Unfortunately, some theoretical papers and practical systems define merges in a way that leads to the paradox of lost future updates. Lippe and van Oosterom [1992] perform merges using the weave concept which, by performing unification first, results in the incorrect \emptyset solution for the merge in example 1. We have verified that the revision control system *darcs* also suf-

fers from this problem. Example 1 can be easily translated into a text merging problem by mapping each atom to a text line.

We can easily generalize the notion of zeropotency to a pair of transformations in the following way:

Definition 4 (mutually zeropotent transformations). Two transformations α and β are called mutually zeropotent iff $(\alpha; \beta) = (\beta; \alpha) = \perp$.

Perhaps counterintuitively, one can find a pair of transformations that are mutually zeropotent, but still can be merged conflict-free (using safe embeddings of course): $\alpha = (i_a; \mathfrak{d}_a; i_b)$ and $\beta = (i_b; \mathfrak{d}_b; i_a)$.

An immediate consequence of the above observation is that we cannot hope to give a rule for when it is “safe” to eliminate duplicate transformations solely based on the zeropotency relation between the initial sequences of transformations we want to merge. Let us further illustrate the difficulty of the matter with the following:

Example 4. Let $\alpha = (i_a; i_b; \mathfrak{d}_a; i_a)$ and $\beta = (i_b; i_a)$. Clearly, no conflict-free merge exists due to i_b . What is less clear is how to give definition of false conflicts that would allow a duplicate i_b to be removed, while disallowing any i_a from being removed.

Definition 5 (merge with safe prefix false-conflict elimination). Let α, β be transformations we want to merge. A merge with safe prefix false-conflict elimination is a triplet of transformations $\dot{\alpha}, \dot{\beta}$ and $\dot{\gamma}$ with the following properties:

- i) $\alpha \sqsubseteq (\dot{\gamma}; \dot{\alpha})$ and $\beta \sqsubseteq (\dot{\gamma}; \dot{\beta})$
- ii) $\forall X \in \alpha^{-1}[\mathcal{Q}] \cap \beta^{-1}[\mathcal{Q}], (\dot{\gamma}; \dot{\alpha}; \dot{\beta})(X) = (\dot{\gamma}; \dot{\beta}; \dot{\alpha})(X) \neq \perp$
- iii) $\forall \dot{\gamma}_1, \dot{\gamma}_2 \in \mathcal{T}, (\dot{\gamma}_1; \dot{\gamma}_2) = \dot{\gamma} \Rightarrow (\dot{\gamma}_2; \dot{\alpha})$ and $(\dot{\gamma}_2; \dot{\beta})$ do not admit a conflict free merge.

The purpose of (iii) in the above definition is to prevent “pulling the zipper too much”. In example 4 for instance, it allows us to reject the following solution $\dot{\gamma} = (i_a; i_b), \dot{\alpha} = (\mathfrak{d}_a; i_a)$ and $\dot{\beta} = \varepsilon$, because we can choose $\dot{\gamma}_1 = i_b$ and $\dot{\gamma}_2 = i_a$ with the result that $(\dot{\gamma}_2; \dot{\alpha}) = (i_a; \mathfrak{d}_a; i_a)$ and $(\dot{\gamma}_2; \dot{\beta}) = i_a$ admit a conflict free merge.

Open Question 2. One could give a definition where the duplicated $\dot{\gamma}$ lies “in the middle”, meaning that $\alpha \sqsubseteq (\dot{\alpha}_1; \dot{\gamma}; \dot{\alpha}_2)$ and similar for β . Can this be reduced to definition 5 or not?

2.5 Dependencies

Definition 6. We say that β depends on α over $D \subseteq \mathcal{Q}$ and write $\alpha <_D \beta$ iff the following three statements are all true:

- i) $\forall x \in D, \alpha(x) \neq \perp \Rightarrow \beta(\alpha(x)) \neq \perp$
- ii) $\forall x \in D, \beta(x) = \perp \Rightarrow \alpha(x) \neq \perp$
- iii) $\exists x \in D, \beta(x) = \perp$

The intuition is that α makes β “not break” over D if it is applied before it. The condition (iii) is necessary to prevent (ii) from holding vacuously. Without (iii) the definition would imply that the identity transformation ε depends on any other

fix spacing around \perp

this needs better explanations

cite it somehow

need a decent phrase here

¹²Some authors use the term zero element.

¹³For the algebraically interested reader, we note that if all transformations from \mathcal{T} are zeropotent, then function composition itself becomes a zeropotent operation in the sense used in abstract algebra, more precisely: $xx\gamma = \gamma\gamma x = xx$.

transformation α because $\alpha(x) \neq \perp$ implies that $\varepsilon(\alpha(x)) = \alpha(x) \neq \perp$, which satisfies (i), and since $\varepsilon(x) \neq \perp$ for any $D \subseteq \mathcal{Q}$ and any $x \in D$, (ii) holds vacuously. In intuitive terms, there must be some “danger” that β is exposed to before α can “claim to make it safe”.

Henceforth we will make the notation $\alpha \rightsquigarrow_D \beta$ iff $\forall x \in D, \alpha(\beta(x)) = \beta(\alpha(x))$.

Proposition 2. $\alpha \rightsquigarrow_D \beta$ and $\alpha <_D \beta$ are mutually exclusive.

Proof. Assume they are both true at the same time. Then by definition 6(iii) we can choose $z \in D$ s.t. $\beta(z) = \perp$ and by (ii) $\alpha(z) \neq \perp$, which by (i) entails that $\beta(\alpha(z)) \neq \perp$.

On the other hand, $\alpha \rightsquigarrow_D \beta$ implies that $\beta(\alpha(z)) = \alpha(\beta(z)) = \alpha(\perp) = \perp$. Contradiction! \square

Corollary 1. (dependencies preclude commutation and vice versa)

$$\begin{aligned} \alpha <_D \beta \text{ or } \beta <_D \alpha &\Rightarrow \alpha \not\rightsquigarrow_D \beta \\ \alpha \rightsquigarrow_D \beta &\Rightarrow \alpha \not<_D \beta \text{ and } \beta \not<_D \alpha \end{aligned}$$

3 Delta-acts

Let \mathcal{Q} and \mathcal{D} be a sets. The elements of \mathcal{Q} are interpreted as states and the elements of \mathcal{D} as deltas between states.

Definition 7. (δ -act) A δ -act is a quadruple $(\mathcal{Q}, \mathcal{D}, \cdot, +)$ where $\cdot: \mathcal{Q} \times \mathcal{Q} \rightarrow \mathcal{D}$ is the delta creation function (think diff), and $+: \mathcal{Q} \times \mathcal{D} \rightarrow \mathcal{Q} \cup \{\perp\}$ is the delta application function (think patch), subject to the following law: $\forall A, B \in \mathcal{Q}, A + AB = B$.

The only relationship imposed on \cdot and $+$ is that the delta from A to B , when applied to A must yield exactly B . In general we can impose no other requirement. In particular we cannot say anything about “how” a delta gets applied to a state other than the one it was based on (A). Note that delta application may “fail”, i.e. produce \perp .

Definition 8. (three-way merge operator) A three way merge operator is a ternary operator with signature $: \mathcal{Q} \times \mathcal{Q} \times \mathcal{Q} \rightarrow \mathcal{Q} \cup \{\perp\}$, written in infix notation as $A \perp O \perp B$, where O is the “reference” state, while A and B are the states “to be merged”, subject to the following two laws:

$$\begin{aligned} A \perp A \perp B &= B && \text{identity law} \\ A \perp O \perp B &= B \perp O \perp A && \text{tip-commutativity} \end{aligned}$$

We will show how a three-way merge operator implicitly defines a δ -act. First, the set of deltas \mathcal{D} is the set of transformations from \mathcal{Q} to \mathcal{Q}_\perp . The delta creation function maps any pair of states to a function. Naturally, this is defined by partial application of the three way merge operator: $AB \mapsto (X \mapsto X \perp A \perp B)$. The patch function is simply function application. We still need to verify the law of a δ -act. It holds due to the identity law of the three way merge operator:

$$A + AB = (X \mapsto X \perp A \perp B)(A) = A \perp A \perp B = B$$

Thus, we have shown that three-way merge operator defines a subclass of δ -act, which we call 3-act.

Although a 3-act is not exactly a merge monoid because the domain of the delta functions is just \mathcal{Q} , it can be trivially embedded in a merge monoid by adding a mapping from \perp to \perp for all delta functions.¹⁴

The question is now whether a three-way merge has any relationship to our definition of operation-based mering. Assume a three-way merge of $A \perp O \perp B = M$ produces a valid state, i.e. $M \neq \perp$. The three-way merge also defines two transformations $\alpha(X) = X \perp O \perp B$ and $\beta(X) = X \perp O \perp A$. Observe that:

$$\begin{aligned} \alpha(\beta(O)) &= (A \perp O \perp O) \perp O \perp B = (O \perp O \perp A) \perp O \perp B = A \perp O \perp B \\ \beta(\alpha(O)) &= (O \perp O \perp B) \perp O \perp A = B \perp O \perp A = A \perp O \perp B \end{aligned}$$

Thus, the two transformations commute, which means they satisfy definition 2 (conflict-free merge). In summary, we have shown that a conflict-free three-way merge satisfies our definition of an operation-based conflict-free merge. The “trick” here is that the two operations that are merged are implicitly defined.

¹⁴Alternatively, we could have defined the three way merge operator on \mathcal{Q}_\perp , and added an additional law that made the result \perp if any of the three arguments are \perp . This construction seems less natural.

References

- Michael Barr and Charles Wells. *Category Theory for Computing Science*. Prentice Hall, Englewood Cliffs, 1995. ISBN 0133238091.
- Brian Berliner. CVS II: Parallelizing software development. In *Proceedings of the USENIX Winter 1990 Technical Conference*, pages 341–352, Berkeley, CA, 1990. USENIX Association. URL citeseer.ist.psu.edu/berliner90cvs.html.
- Anne-Marie Kermarrec, Antony Rowstron, Marc Shapiro, and Peter Druschel. The icecube approach to the reconciliation of divergent replicas. In *PODC '01: Proceedings of the twentieth annual ACM symposium on Principles of distributed computing*, pages 210–218, New York, NY, USA, 2001. ACM. ISBN 1-58113-383-9. doi: <http://doi.acm.org/10.1145/383962.384020>.
- Sanjeev Khanna, Keshav Kunal, and Benjamin C. Pierce. A formal investigation of Diff3. In Arvind and Prasad, editors, *Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*, December 2007.
- Leslie Lamport. Time, clocks, and the ordering of events in a distributed system. *Commun. ACM*, 21(7):558–565, 1978. ISSN 0001-0782. doi: <http://doi.acm.org/10.1145/359545.359563>.
- Ernst Lippe and Norbert van Oosterom. Operation-based merging. In *SDE 5: Proceedings of the fifth ACM SIGSOFT symposium on Software development environments*, pages 78–87, New York, NY, USA, 1992. ACM. ISBN 0-89791-554-2. doi: <http://doi.acm.org/10.1145/142868.143753>.
- Norman Ramsey and Előd Csirmaz. An algebraic approach to file synchronization. *SIGSOFT Softw. Eng. Notes*, 26(5):175–185, 2001. ISSN 0163-5948. doi: <http://doi.acm.org/10.1145/503271.503233>.
- Jonathan D. H. Smith. *Mal'cev Varieties*, volume 554 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin/New York, 1976. ISBN 9783540079996. doi: [10.1007/BFb0095447](https://doi.org/10.1007/BFb0095447).
- Chengzheng Sun, Xiaohua Jia, Yanchun Zhang, Yun Yang, and David Chen. Achieving convergence, causality preservation, and intention preservation in real-time cooperative editing systems. *ACM Trans. Comput.-Hum. Interact.*, 5(1):63–108, 1998. ISSN 1073-0516. doi: <http://doi.acm.org/10.1145/274444.274447>.
- D. B. Terry, M. M. Theimer, Karin Petersen, A. J. Demers, M. J. Spreitzer, and C. H. Hauser. Managing update conflicts in bayou, a weakly connected replicated storage system. In *SOSP '95: Proceedings of the fifteenth ACM symposium on Operating systems principles*, pages 172–182, New York, NY, USA, 1995. ACM. ISBN 0-89791-715-4. doi: <http://doi.acm.org/10.1145/224056.224070>.
- Walter F. Tichy. RCS — a system for version control. *Software — Practice and Experience*, 15(7):637–654, 1985. doi: <http://dx.doi.org/10.1002/spe.4380150703>. URL citeseer.ist.psu.edu/tichy85rcs.html.