

# Exploiting Schemas in Data Synchronization<sup>1,2</sup>

J. Nathan Foster<sup>a,3,\*</sup> Michael B. Greenwald<sup>b</sup>,  
Christian Kirkegaard<sup>c</sup>, Benjamin C. Pierce<sup>a</sup> and Alan Schmitt<sup>d</sup>

<sup>a</sup>*University of Pennsylvania, Philadelphia, PA 19104 USA*

<sup>b</sup>*Bell Labs, Lucent Technologies, Murray Hill, NJ 07974 USA*

<sup>c</sup>*BRICS, University of Aarhus, DK-8000 Aarhus C, Denmark*

<sup>d</sup>*INRIA Rhône-Alpes, 38330 Montbonnot, France*

---

## Abstract

Increased reliance on optimistic data replication has led to burgeoning interest in tools and frameworks for *synchronizing* disconnected updates to replicated data. We have implemented a generic synchronization framework, called Harmony, that can be used to build state-based synchronizers for a wide variety of tree-structured data formats. A novel feature of this framework is that the synchronization process—in particular, the recognition of conflicts—is driven by the schema of the structures being synchronized. We formalize Harmony’s synchronization algorithm, prove that it obeys a simple and intuitive specification, and illustrate, using simple address books as a case study, how it can be used to synchronize trees representing a variety of specific forms of application data, including sets, records, tuples, and relations.

*Key words:* synchronization, optimistic reconciliation, XML, Harmony

---

---

\* Corresponding author.

Email addresses: jnfoster@cis.upenn.edu (J. Nathan Foster), greenwald@lucent.com (Michael B. Greenwald), ck@brics.dk (Christian Kirkegaard), bcpierce@cis.upenn.edu (Benjamin C. Pierce), alan.schmitt@inrialpes.fr (Alan Schmitt).

<sup>1</sup> Extended version of a paper originally presented at the Symposium on Database Programming Languages (DBPL) in Trondheim, Norway, August 2005.

<sup>2</sup> Partially supported by the NSF under grant ITR-0113226.

<sup>3</sup> Partially supported by an NSF GRF.

## 1 Introduction

Optimistic replication strategies are attractive in a growing range of settings where weak consistency guarantees can be accepted in return for higher availability and the ability to update data while disconnected. These uncoordinated updates must later be *synchronized* (or *reconciled*) by automatically combining non-conflicting updates while detecting and reporting conflicting updates.

Our long-term aim is to develop a generic framework that can be used to build high-quality synchronizers, with minimal effort, for a wide variety of application data formats. As a step toward this goal, we have designed and built a prototype synchronization framework called Harmony, focusing on the important special cases of unordered and rigidly ordered data (including sets, relations, tuples, records, feature trees, etc.), with only limited support for list-structured data such as structured documents. An instance of Harmony that synchronizes multiple calendar formats (Palm Datebook, Unix ical, and iCalendar) has been deployed within our group; we are currently developing Harmony instances for bookmark data (handling the formats used by several common browsers, including Mozilla, Safari, and Internet Explorer), address books, application preference files, drawings, and bibliographic databases.

The Harmony system has two main components: (1) a domain-specific programming language for writing *lenses*—bi-directional transformations on trees—which we use to convert low-level (and possibly heterogeneous) concrete data formats into a high-level *synchronization schema*, and (2) a generic synchronization algorithm, whose behavior is controlled by the synchronization schema.

The synchronization schema actually guides Harmony’s behavior in two ways. First, by choosing an appropriate synchronization schema (and the lenses that transform concrete structures into this form and back), users of Harmony can control the *alignment* of the information being synchronized: the same concrete format might be transformed to different synchronization schemas (for example, making different choices of keys) to yield quite different synchronization semantics; this process is illustrated in detail in Section 7. Second, during synchronization, the synchronization schema is used to identify *conflicts*—situations where changes in one replica must *not* be propagated to the other because the resulting combined structure would be ill-formed. To our knowledge, Harmony is the first state-based synchronizer to preserve structural invariants in this way. (By contrast, *operation-based* synchronization frameworks preserve application invariants by working in terms of a high-level, application-specific algebra of operations rather than directly manipulating replicated data. See Section 8.)

Our language for lenses has been described in detail elsewhere [1]; in the present paper, our focus is on the synchronization algorithm and the way it uses schema information. The basic behavior of this algorithm is simple: we try to propagate changes from each replica to the other, validate the resulting trees according to the expected schema, and signal a conflict if validation fails. However, the details are somewhat subtle: there may be many changes to propagate from each replica to the others, leading to many possible choices of *where* to signal conflicts (i.e., which subset of the changes to propagate). To ensure progress, we want synchronization to propagate as many changes as possible while respecting the schema; at the same time, to avoid surprising users, we need the results of synchronization to be predictable; for example, small variations in the inputs should not produce large variations in the set of changes that are propagated. A natural way of combining these design constraints is to demand that the results of synchronization be *maximal*: if there is *any* way to validly propagate a given modification from one replica to the other, then that change *must* be propagated. Our main technical contribution is a simple one-pass, recursive tree-walking algorithm that does indeed yield results that are maximal in this sense (and hence also unique) for schemas satisfying a locality constraint called *path consistency* (a semantic variant of the *consistent element declaration* condition in W3C Schema).

After establishing some notation in Section 2, we explore the design space further, beginning in Section 3 with some simple synchronization examples. Section 4 focuses on difficulties that arise in a schema-aware algorithm. Section 5 presents the algorithm itself. Section 6 gives its formal specification (in particular, the definition of maximality) and proves that the algorithm meets the specification. Section 7 illustrates the behavior of the algorithm using a simple address book schema. Related and future work are discussed in Sections 8 and 9.

## 2 Data Model

Internally, Harmony manipulates structured data in an extremely simple form: unordered, edge-labeled trees; richer external formats such as XML are encoded in terms of unordered trees. We chose this simple data model on pragmatic grounds: experience shows that the reduction in the overall complexity of the Harmony system far outweighs the cost of manipulating ordered data in encoded form (see [1]). We write  $\mathcal{N}$  for the set of character strings and  $\mathcal{T}$  for the set of unordered, edge-labeled trees whose labels are drawn from  $\mathcal{N}$  and where labels of the immediate children of nodes are pairwise distinct. We draw trees sideways to save space: in text, each pair of curly braces denotes a tree node, and each “ $\mathbf{x} \mapsto \dots$ ” denotes a child labeled  $\mathbf{x}$ —e.g.,  $\{\text{Pat} \mapsto 111-1111, \text{Chris} \mapsto 222-2222\}$ . To avoid clutter, when an edge leads

to an empty tree, we omit the braces, the  $\mapsto$  symbol, and the final childless node—e.g., “111-1111” above actually stands for “ $\{111-1111 \mapsto \{\}\}$ .”

A tree can be viewed as a partial function from names to trees; we write  $t(n)$  for the immediate subtree of  $t$  labeled with the name  $n$  and  $\text{dom}(t)$  for its domain—i.e. the set of the names of its children. The concatenation operator,  $\cdot$ , is only defined for trees with disjoint domains;  $t \cdot t'$  is the tree mapping  $n$  to  $t(n)$  for  $n \in \text{dom}(t)$  and to  $t'(n)$  for  $n \in \text{dom}(t')$ . When  $n \notin \text{dom}(t)$ , we define  $t(n)$  to be  $\perp$ , the “missing tree.” By convention, we take  $\text{dom}(\perp) = \emptyset$ . To represent conflicts during synchronization, we enrich the set of trees with a special pseudo-tree  $\mathcal{X}$ , pronounced “conflict.” If  $S$  is a set of ordinary trees (i.e.,  $\perp \notin S$  and  $\mathcal{X}$  is not a subtree of any tree in  $S$ ), we write  $S_\perp$  for  $S \cup \{\perp\}$ , and  $S_{\mathcal{X}\perp}$  for the set obtained from  $S_\perp$  by allowing arbitrary subtrees to be replaced by  $\mathcal{X}$ . We prove many properties by induction on the height of trees. We define  $\text{height}(\perp) = \text{height}(\mathcal{X}) = 0$  and the height of an ordinary tree as  $\text{height}(t) = 1 + \max(\{\text{height}(t(k)) \mid k \in \text{dom}(t)\})$ . Note that  $\text{height}(\{\}) = 1$ , to avoid confusing it with  $\perp$  and  $\mathcal{X}$ .

A *path* is a sequence of names. We write  $\bullet$  for the empty path and  $p/q$  for the concatenation of  $p$  and  $q$ ; the set of all paths is written  $\mathcal{P}$ . The *projection* of  $t$  along a path  $p$ , written  $t(p)$ , is defined as follows: (1)  $t(\bullet) = t$ ; (2)  $t(n/p) = (t(n))(p)$  if  $t \neq \mathcal{X}$  and  $n \in \text{dom}(t)$ ; (3)  $t(n/p) = \perp$  if  $t \neq \mathcal{X}$  and  $n \notin \text{dom}(t)$ ; and (4)  $t(p) = \mathcal{X}$  if  $t = \mathcal{X}$ . A tree is *included* in another tree, written  $t \sqsubset t'$ , iff any missing or conflicting path in  $t'$  is missing in  $t$ :  $\forall p \in \mathcal{P}. (t'(p) = \perp \vee t'(p) = \mathcal{X}) \implies t(p) = \perp$ .

Our algorithm is formulated using a semantic notion of *schemas*—a schema  $S$  is a set of ordinary trees  $S \subseteq \mathcal{T}$ . In Section 7 we define a particular syntax for writing down schemas, which is used in the examples (and in our prototype implementation); however, our algorithm does not rely on this particular notion of schema.

### 3 Basics

Harmony’s synchronization algorithm takes two<sup>4</sup> replicas  $a, b \in \mathcal{T}_\perp$  and a common ancestor  $o \in \mathcal{T}_\mathcal{X}$  and yields new replicas in which all non-conflicting updates are merged. The missing tree,  $\perp$ , represents replicas that have been completely deleted (note that  $\perp$  is different from the empty tree,  $\{\}$ );  $\mathcal{X}$  marks paths in  $o$  that a previous synchronization run left in conflict.

<sup>4</sup> We focus on the two-replica case. Our algorithm generalizes straightforwardly to  $n$  simultaneously connected replicas, but the more realistic case where only a subset of the replicas may be connected at any given moment poses additional challenges.

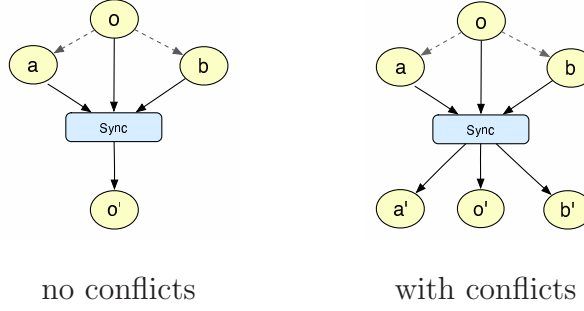


Fig. 1. Synchronizer architecture

Suppose, for example, that we have a tree  $o$  representing a phone book and that we make two replicas of this structure,  $a$  and  $b$ , and separately modify one phone number in each:

$$\begin{aligned} o &= \{\text{Pat} \mapsto 111-1111, \text{Chris} \mapsto 222-2222\} \\ a &= \{\text{Pat} \mapsto 111-1111, \text{Chris} \mapsto 888-8888\} \\ b &= \{\text{Pat} \mapsto 999-9999, \text{Chris} \mapsto 222-2222\} \end{aligned}$$

Synchronization takes these structures and produces structures

$$a' = b' = \{\text{Pat} \mapsto 999-9999, \text{Chris} \mapsto 888-8888\}$$

that reflect all the changes in  $a$  and  $b$  with respect to  $o$ . We also save the final merged state at the end of each synchronization, to use as the  $o$  input of the next synchronization.<sup>5</sup> Schematically, the synchronizer may be visualized like the diagram on the left of Figure 1.

In the remainder of this section we describe the fundamental choices in Harmony’s architecture and discuss the relative advantages and disadvantages of these choices, compared with alternatives found in other synchronization technologies.

**Loose Coupling** Harmony is a *state-based* synchronizer: only the current states of the replicas (plus the remembered common state  $o$ ) are supplied as inputs; the precise sequences of operations that produced  $a$  and  $b$  from  $o$  are not available to the synchronizer. Harmony is designed to require only *loose coupling* with applications: it manipulates application data in external, on-disk representations such as XML trees. The advantage of the loosely coupled approach is that we can use Harmony to synchronize off-the-shelf applications that were implemented without replication in mind. By contrast, synchronizers that work with the traces of operations that the application has performed on each replica and that propagate changes by undoing and/or replaying op-

<sup>5</sup> In a multi-replica system, an appropriate “last shared state” would instead be calculated from the causal history of the system.

erations require *tight coupling* between the synchronizer and application programs.

**Conflicts and Persistence** During synchronization, it is possible that some of the changes made to the two replicas are in conflict and cannot be merged. For example, suppose that, beginning from the same original  $o$  as in the example above, we change both Pat’s and Chris’s phone numbers in  $a$  and, in  $b$ , delete the record for Chris entirely, yielding replicas  $a = \{\text{Pat} \mapsto 123\text{-}4567, \text{Chris} \mapsto 888\text{-}8888\}$  and  $b = \{\text{Pat} \mapsto 111\text{-}1111\}$ . Clearly, there is no single phone book that incorporates both changes to Chris: we have a *delete/create conflict*—the subtree along the path **Chris** was deleted in  $b$  and a tree was created in  $a$  along the path **Chris/888-8888**. At this point, we must choose between two evils. On one hand, we can weaken users’ expectations of the *persistence* of their changes to the replicas—i.e., we can decline to promise that synchronization will never back out changes that the user has made to either replica. For example, here, we might back out the deletion of Chris, yielding  $a' = b' = \{\text{Pat} \mapsto 123\text{-}4567, \text{Chris} \mapsto 888\text{-}8888\}$ . The user would then be notified of the lost changes and given the opportunity to re-apply them if desired. Alternatively, we can keep persistence and instead give up *convergence*—i.e., we can allow the replicas to remain different after synchronization, propagating just the non-conflicting change to Pat’s phone number and leaving the conflicting information about Chris untouched in each replica— $a' = \{\text{Pat} \mapsto 123\text{-}4567, \text{Chris} \mapsto 888\text{-}8888\}$  and  $b' = \{\text{Pat} \mapsto 123\text{-}4567\}$ —and notifying the user of the conflict.<sup>6</sup>

In Harmony, we have chosen to favor persistence, for several pragmatic reasons. First, this choice makes it easier to ensure that unsupervised reconciliations are safe. (Unsupervised reconciliations seem extremely desirable from the point of view of system administration, facilitating automatic reconciliation upon dis/connection from/to networks and via nightly scripts.) Divergent systems are more likely to allow users to proceed with their work—the set of replicas may be globally inconsistent, but it is more likely that each replica is locally consistent. By contrast, convergent systems are more likely to force a user to resolve a conflict after a *remote* user initiated a synchronization attempt. For example, consider conflicting updates to a file with strict syntax requirements (e.g. LaTeX or C); the convergent system’s attempt to record both updates may result in a file that causes subsequent processing to fail. Second, divergent systems are less likely to hide conflicts for long periods of time. Instead, they remind users of the conflict at every synchronization at-

---

<sup>6</sup> An industrial-strength synchronization tool will not only report the conflict, but also assist in bringing the replicas back into agreement by providing graphical views of the differences, applying special heuristics, etc. We omit discussion of these aspects of synchronization, focusing on the synchronizer’s basic, “unattended” behavior.

tempt until the conflict is resolved. By contrast, convergent systems succeed, even if conflicts persist, because the replicas are made identical. To achieve this behavior, convergent systems typically resolve conflicts by simply discarding one update (perhaps saving it under a different name, so that the data can easily be retrieved for manual repair). For a more detailed survey of convergent systems, see Section 8.

With this refinement, the schematic view of the synchronizer looks like the diagram on the right side of Figure 1. (The new archive  $o'$ , is also an output of the synchronization process; see Section 5.)

**Local Alignment** Another fundamental consideration in the design of any synchronizer is *alignment*—i.e., the mechanism that identifies which parts of each replica represent “the same information” and should be synchronized with each other. Synchronization algorithms can be broadly grouped into two categories, according to whether they make alignment decisions *locally* or *globally*. Synchronizers that use global heuristics for alignment—e.g., the popular Unix tool `diff3`, Lindholm’s 3DM [2], the work of Chawathe et al [3], and FCDP [4]—make a “best guess” about what operations the user performed on the replicas by comparing the *entire* current states with the last common state. This works well in many cases (where the best guess is clear), but in boundary cases these algorithms can make surprising decisions—i.e., it can be difficult for a user to predict how data will be aligned, especially in cases where both replicas have changed significantly. (See Appendix B for a concrete example.) To avoid these issues, our algorithm employs a simple, local alignment strategy that associates the subtrees under children with the same name with each other. The behavior of this scheme should be easy for users to understand and predict. The cost of operating *completely* locally is that Harmony’s ability to deal with ordered data is limited, as we discuss in Section 7. An important avenue for future work is hybridizing local and global alignment techniques to combine their advantages.

**Lenses** The local alignment scheme described above works well when the replicas are represented in a format that naturally exposes the structure of the data to be synchronized. For example, if the replicas represent address books, then a good representation is as a bush where an appropriate *key field*, providing access to each contact, is at the root level. The key fields, which uniquely identify a contact, are often drawn from some underlying database:

$$\left\{ \begin{array}{l} 92373 \mapsto \left\{ \text{name} \mapsto \left\{ \begin{array}{l} \text{first} \mapsto \text{Megan}, \\ \text{last} \mapsto \text{Smith} \end{array} \right\}, \text{home} \mapsto 555-6666 \right\}, \\ 92374 \mapsto \left\{ \text{name} \mapsto \left\{ \begin{array}{l} \text{first} \mapsto \text{Pat}, \\ \text{last} \mapsto \text{Jones} \end{array} \right\}, \text{home} \mapsto 555-2222 \right\} \end{array} \right\}$$

Using the alignment scheme described above, the effect during synchronization

will be that entries from the two replicas with the same UID are synchronized with each other. Alternatively, if UIDs are not available, we might instead synthesize a UID by lifting information out of each record—e.g., by concatenating the `name` data and using it as the top-level key field:

$$\left\{ \begin{array}{l} \text{Megan:Smith} \mapsto \{\text{home} \mapsto 555-6666\}, \\ \text{Pat:Jones} \mapsto \{\text{home} \mapsto 555-2222\} \end{array} \right\}$$

It is unlikely, however, that the address book will be represented concretely (e.g., as an XML document) using either of these formats. To bridge this gap between actual, application-determined concrete formats and more synchronization-friendly representations, the Harmony system includes a domain-specific language for writing bi-directional transformations [1], which we call *lenses*. By passing each replica through a lens, we can transform the replicas from concrete formats into appropriately “pre-aligned” forms. After synchronization, our language guarantees that the updated replicas are transformed back into the appropriate concrete formats using the other side of the same lens (i.e., lenses can be thought of as *view update translators* [5]). Lenses also facilitate synchronization of heterogeneous formats. Since each replica is passed through a lens both before and after synchronization, it does not matter whether if the replicas are not represented in the same format. We can apply a different lens to each replica to transform the disparate concrete representations into the same format for synchronization.

## 4 The Role of Schemas

To formalize our intuitions about what a user may reasonably expect from a synchronizer, we impose two core requirements, which we call *safety* and *maximality*. We describe them informally in this section; Section 6 gives precise definitions.

The safety requirement encompasses four basic sanity checks: (1) The synchronizer must not “back out” any changes made at a replica since the last synchronization (because we favor persistence over convergence); (2) it must only copy data between replicas, never “making up” content; (3) it must halt at conflicting paths, leaving the replicas untouched below; (4) it must produce results that belong to the same schema as the originals.

Of course, safety alone is too weak: an algorithm that returns both replicas unchanged would be trivially safe! We therefore say that a run of a synchronizer is *maximal* just in case it propagates all the changes of every other safe run. Our specification, given formally in Section 6, is that every run must be both safe and maximal.

Our algorithm (unlike other state-based synchronizers) is designed to preserve structural invariants. As an example of how such invariants can be broken, consider a run of the algorithm sketched above where the inputs are as follows (we revert to the fully explicit notation for trees here, to remind the reader that “leaf values” are labels that lead to an empty subtree.):

$$\begin{aligned} o &= \{\text{Pat} \mapsto \{\text{Phone} \mapsto \{333-4444 \mapsto \{\}}\}\} \\ a &= \{\text{Pat} \mapsto \{\text{Phone} \mapsto \{111-2222 \mapsto \{\}}\}\} \\ b &= \{\text{Pat} \mapsto \{\text{Phone} \mapsto \{987-6543 \mapsto \{\}}\}\} \end{aligned}$$

The subtree labeled 333-4444 has been deleted in both replicas, and remains so in both  $a'$  and  $b'$ . The subtree labeled 111-2222 has been created in  $a$ , so we can propagate the creation to  $b'$ ; similarly, we can propagate the creation of 987-6543 to  $a'$ , yielding

$$a' = b' = \{\text{Pat} \mapsto \{\text{Phone} \mapsto \{111-2222 \mapsto \{\}, 987-6543 \mapsto \{\}\}\}\}$$

But this would be wrong. Pat’s phone number was *changed* in different ways in the two replicas: what’s wanted is a conflict. If the phonebook schema only allows a single number per person, then the new replica is not well formed! We avoid these situations by providing the schema as an input to the synchronizer. The synchronizer signals a conflict (leaving the replicas unchanged) if merging the changes at a particular point yields an ill-formed structure.

## Locality and Schemas

A simple way to ensure that the results produced by the synchronization algorithm are valid is to check for schema compliance at the end of each run, halting with a conflict if the results do not belong to the schema. This approach is not satisfying, however, since it discards every modification even if only a small part of the result is not compliant. For example, consider synchronizing replicas of a large address book where only one entry does not belong to the schema after synchronization. Using the approach described above the algorithm halts with a conflict on the *entire* addressbook because that single entry is ill-formed, and declines to propagate safe updates to any other entries. As this clearly conflicts with our goal of maximality, we have taken a more local approach to schema compliance, which entails some restrictions to the schemas that may be used. To motivate these restrictions, consider the following schema:  $\{\{\}, \{\mathbf{n} \mapsto \mathbf{x}, \mathbf{m} \mapsto \mathbf{x}\}, \{\mathbf{n} \mapsto \mathbf{y}, \mathbf{m} \mapsto \mathbf{y}\}, \{\mathbf{n} \mapsto \{\mathbf{x}, \mathbf{y}\}, \mathbf{m} \mapsto \mathbf{y}\}, \{\mathbf{n} \mapsto \mathbf{x}, \mathbf{m} \mapsto \{\mathbf{x}, \mathbf{y}\}\}\}$ . It expresses a non-local invariant: at most one of  $\mathbf{m}$  and  $\mathbf{n}$  has  $\{\mathbf{x}, \mathbf{y}\}$  as a subtree. Now, consider synchronizing two replicas belonging to this set with respect to an empty archive,  $o = \{\}$ , with  $a = \{\mathbf{n} \mapsto \mathbf{x}, \mathbf{m} \mapsto \mathbf{x}\}$  and  $b = \{\mathbf{n} \mapsto \mathbf{y}, \mathbf{m} \mapsto \mathbf{y}\}$ . An algorithm that aligns each replica by name will recursively synchronize the associated subtrees below  $\mathbf{n}$  and  $\mathbf{m}$ . However, it is not clear what *schema* to use for these recursive calls, because the set of trees that can validly appear under

$n$  depends on the subtree under  $m$  and vice versa. We might try the schema that consists of all the trees that can appear under  $n$  (and  $m$ ):  $\{x, y, \{x, y\}\}$ . With this schema, the synchronizer computes the tree  $\{x, y\}$  for both  $n$  and  $m$ , reflecting the fact that  $x$  and  $y$  were both added under  $n$  and  $m$ . However, these trees cannot be assembled into a well-formed tree:  $\{n \mapsto \{x, y\}, m \mapsto \{x, y\}\}$  does not belong to the schema. The “most synchronized” well-formed results are  $a' = \{n \mapsto x, m \mapsto \{x, y\}\}$  and  $b' = \{n \mapsto \{x, y\}, m \mapsto y\}$ , but there does not seem to be any way to find them without backtracking. The global invariant expressed by this schema—at most one of  $n$  or  $m$  may have  $\{x, y\}$  as a subtree—cannot easily be preserved by a local algorithm.

To avoid such situations, we impose a restriction on schemas, *path consistency*, that is analogous to the restriction on tree grammars embodied by W3C Schema. Intuitively, a schema is path consistent if any subtree that appears at some path in one tree can validly be “transplanted” to the same location in any other tree in the schema. This restriction ensures that the sub-schema used to synchronize each child is consistent across the schema; i.e., the set of trees that may validly appear under a child only depends on the path to the node and not the presence or absence of other information elsewhere in the tree.

To define path consistency precisely, we need a little new notation. First, the notion of projection at a path is extended pointwise to schemas—that is, if  $S \subseteq \mathcal{T}$  and  $p \in \mathcal{P}$ , we define  $S(p) = \{t(p) \mid t \in S \wedge t(p) \neq \perp\}$ . Note that a schema projection at a path is itself a schema. Next, we give definition of the transplant of a subtree from one tree to another at a given path. Let  $t \in \mathcal{T}$  and  $p \in \mathcal{P}$  with  $t(p) \in \mathcal{T}$ . The *update* of  $t$  at  $p$  with  $t'$ , written  $t[p \mapsto t']$ , is defined inductively on  $p$  as:  $t[\bullet \mapsto t'] = t'$ ,  $t[n/p \mapsto t'] = \{n \mapsto t(n)[p \mapsto t'], m \mapsto t(m) \mid m \in \text{dom}(t) \setminus \{n\}\}$ . A schema  $S$  is *path consistent* if, whenever  $t$  and  $t'$  are in  $S$ , for every path  $p$ , the result of updating  $t$  along  $p$  with  $t'(p)$  is also in the schema. Formally, a schema  $S$  is path consistent iff, for all  $t, t' \in S$  and  $p \in \mathcal{P}$ , we have  $t(p) \neq \perp \wedge t'(p) \neq \perp \implies t[p \mapsto t'(p)] \in S$ .

To complete the discussion of the role of schemas in synchronization, we must consider one final complication: for some inputs, there may be *no* safe, maximal runs belonging to the schema. Consider a run of the synchronizer on inputs  $o = \{v\}$ ,  $a = \{w, y, z\}$ , and  $b = \{w, x\}$ , with respect to the schema  $\{\{v\}, \{w, x\}, \{w, x, y\}, \{w, x, z\}, \{w, y, z\}\}$ . For the  $a$  replica, the only safe result that belongs to the schema is  $a' = a$ . However, on the  $b$  side, there are three safe results belonging to the schema:  $\{w, x\}$ ,  $\{w, x, y\}$ , and  $\{w, x, z\}$  but none is maximal. Notice that, since  $\{w, x, y, z\}$  does not belong to the schema, we cannot include both  $y$  and  $z$  in  $b'$  (without backing out the addition of  $x$ ). Indeed, for every safe  $b'$ , there is a path  $p$  where  $b'(p) \neq a'(p)$ , but, for a different choice of  $b'$ , the trees at that path are equal. To ensure that synchronization always has a maximal result, we stipulate that a *schema domain conflict*

```

sync( $S, o, a, b$ ) =
  if  $a = b$  then  $(a, a, b)$                                 – equal replicas: done
  else if  $a = o$  then  $(b, b, b)$                              – no change to  $a$ : propagate  $b$ 
  else if  $b = o$  then  $(a, a, a)$                              – no change to  $b$ : propagate  $a$ 
  else if  $o = \mathcal{X}$  then  $(o, a, b)$                        – unresolved conflict
  else if  $a = \perp$  and  $b \sqsubset o$  then  $(a, a, a)$            –  $a$  deleted more than  $b$ 
  else if  $a = \perp$  and  $b \not\sqsubset o$  then  $(\mathcal{X}, a, b)$          – delete/create conflict
  else if  $b = \perp$  and  $a \sqsubset o$  then  $(b, b, b)$            –  $b$  deleted more than  $a$ 
  else if  $b = \perp$  and  $a \not\sqsubset o$  then  $(\mathcal{X}, a, b)$          – delete/create conflict
  else                                                       – proceed recursively
    let  $(o'(k), a'(k), b'(k)) = \text{sync}(S(k), o(k), a(k), b(k))$ 
     $\forall k \in \text{dom}(a) \cup \text{dom}(b)$ 
    in if  $(\text{dom}(a') \not\subseteq \text{doms}(S))$  or  $(\text{dom}(b') \not\subseteq \text{doms}(S))$ 
      then  $(\mathcal{X}, a, b)$                                 – schema domain conflict
      else  $(o', a', b')$ 

```

Fig. 2. Synchronization Algorithm

occurs whenever propagating *all* of the (otherwise non-conflicting) additions and deletions of children at a node yields an ill-formed result (see Definition 3 below). On the above trees, our algorithm yields a schema domain conflict at the root since it cannot add  $y$  and  $z$  to  $a'$ .

Note, that the approach presented in this section is not the only way to ensure maximality in a schema-directed algorithm; we have considered several alternatives. First, we could do away with the notion of a schema domain conflict and require instead that schemas be closed under the “shuffling” of their domains with the domains of other trees in the schema. This approach amounts to declaring, by fiat, that the maximal result of every possible synchronization must be present in the schema. For example, in the schema above, we would need to include the tree  $\{\mathbf{w}, \mathbf{x}, \mathbf{y}, \mathbf{z}\}$  in the schema. We have not pursued this idea because it does not appear that shuffled schemas would be flexible enough to express the kinds of invariants needed by applications. Second, we could recognize schema domain conflicts, but, instead of requiring that the replicas remain unchanged, only require that the *domains* be unchanged. This approach would allow “deep” synchronization of subtrees, which has some obvious advantages. However, finding a natural notion of maximality for this variant is difficult. For these reasons, our first—simplest—proposal seems best.

## 5 Algorithm

The synchronization algorithm is depicted in Figure 2. Its structure is as follows: we first check for trivial cases (replicas being equal to each other or unmodified), then we check for *delete/create conflicts*, and in the general

case we recurse on each child label and check for *schema domain conflicts* before returning the results. In practice, synchronization will be performed repeatedly, with additional updates applied to one or both of the replicas between synchronizations. To support this, the algorithm constructs a new archive: we use the synchronized version at every path where the replicas agree and insert a conflict marker  $\mathcal{X}$  at conflicting paths.

Formally, the algorithm takes as inputs a path-consistent schema  $S$ , an archive  $o$ , and two current replicas  $a$  and  $b$ ; it outputs a new archive  $o'$  and two new replicas  $a'$  and  $b'$ . We require that both  $a$  and  $b$  belong to  $S_{\perp}$ . The input archive may contain the conflict marker  $\mathcal{X}$ . The algorithm also relies on one piece of new notation:  $\text{doms}(S)$  stands for the *domain-set* of  $S$ , the set of all domains of trees in  $S$ —i.e.,  $\text{doms}(S) = \{\text{dom}(t) \mid t \in S\}$ .

In the case where  $a$  and  $b$  are identical (i.e., the same tree or  $\perp$ ), they are immediately returned, and the new archive is set to their value. If one of the replicas is unchanged (equal to the archive), then all the changes in the other replica can safely be propagated, so we simply return three copies of it as the result replicas and archive. Otherwise, both replicas have changed, in different ways. If there was a conflict in the previous run, then it has not been resolved. If one replica is missing, then we check whether all the changes in the other replica are also deletions; formally, we check if the replica is included in the archive. If so, we consider the larger deletion (discarding the entire tree at this path) as subsuming the smaller; otherwise, we have a *delete/create conflict* and we return the original replicas.

Finally, in the general case, the algorithm recurses: for each  $k$  in the domain of either current replica, we call *sync* with the corresponding subtrees,  $o(k)$ ,  $a(k)$ , and  $b(k)$  (any of which may be  $\perp$ ), and the sub-schema  $S(k)$ ; we collect up the results of these calls to form new trees  $o'$ ,  $a'$ , and  $b'$ . If either of the new replicas is ill-formed (i.e., its domain is not in the domain-set of the schema), then we have a schema domain conflict and the original replicas are returned unmodified. Otherwise, the synchronized results are returned.

(Careful readers may note that a naive implementation of this algorithm can perform many redundant equality checks in the common case when the replicas are almost equal—checking the topmost nodes for equality, failing, recursing on their immediate children, checking these again for equality, etc. A combinatorial blowup in running time can be avoided by caching the results of earlier equality checks to speed up later ones.)

## 6 Specification

We now give a formal specification of the properties we want our synchronization algorithm to satisfy. Our presentation follows the basic approach used for specifying the Unison file synchronizer [6]. We start by with a few auxiliary definitions which are needed in the formal statements of safety and maximality. Both safety and maximality are based on a notion of *local similarity*, which relates two trees if their top-level nodes are similar—i.e., intuitively, if both are present, both are missing, or both represent conflicts.

**Definition 1 (Local similarity)** *Two elements of  $\mathcal{T}_{\mathcal{X}\perp}$  are locally similar, written  $t \sim t'$ , iff (1)  $t = t' = \mathcal{X}$  or (2)  $t = t' = \perp$  or (3)  $t \neq \perp \wedge t \neq \mathcal{X} \wedge t' \neq \perp \wedge t' \neq \mathcal{X}$ .*

In the following we silently rely on the fact that  $\sim$  is an equivalence relation.

The definition of safety relies on the notion of conflict. We use local similarity to capture all the simple notions of conflicts that consider only the presence or absence of a single node.

**Definition 2 (Local conflict)** *An archive  $o$  and replicas  $(a, b)$  have a local conflict, written  $\text{localconflict}(o, a, b)$ , iff either (1)  $(o = \mathcal{X}) \wedge (a \neq b)$  or (2)  $(o \neq a) \wedge (o \neq b) \wedge (a \not\sim b) \wedge (a = \perp \implies b \not\subseteq o) \wedge (b = \perp \implies a \not\subseteq o)$ .*

Intuitively, replicas  $(a, b)$  have a local conflict if there is an unresolved conflict recorded in the archive  $o$ , or if they have both changed since the state recorded in the archive but are not locally similar and are in a delete/create conflict. Schema domain conflicts are defined via a local test on domains.

**Definition 3 (Schema domain conflict)** *Let  $S$  be a schema. An archive  $o$  and replicas  $(a, b)$  have a schema domain conflict, written  $\text{sdomconflict}(S, o, a, b)$ , iff  $(a \neq \perp) \wedge (b \neq \perp) \wedge (\text{edom}(o, a, b) \not\subseteq \text{doms}(S) \vee \text{edom}(o, b, a) \not\subseteq \text{doms}(S))$  where  $\text{edom}(o, a, b) = \{k \in \text{dom}(a) \mid a(k) \not\subseteq o(k)\} \cup (\text{dom}(b) \setminus \text{dom}(o)) \cup (\text{dom}(a) \cap \text{dom}(b))$ .*

The set of names  $\text{edom}(o, a, b)$  is the expected domain of the tree that results from safely propagating into replica  $a$  the changes from replica  $b$ . It contains all names under which there is a change that is not simply a deletion in replica  $a$ , all names added by replica  $b$ , and all the names preserved in both  $a$  and  $b$ .

**Definition 4 (Conflict)** *Let  $S$  be a schema. An archive  $o$  and replicas  $(a, b)$  have a conflict, written  $\text{conflict}(S, o, a, b)$ , iff  $\text{localconflict}(o, a, b)$  or  $\text{sdomconflict}(S, o, a, b)$ .*

Safety and maximality are both expressed as properties of *runs*.

**Definition 5 (Synchronizer run)** A run is a tuple  $(S, o, a, b, o', a', b')$  comprising a schema  $S$  and trees representing the archive and replicas supplied as inputs,  $(o, a, b)$ , and computed as outputs,  $(o', a', b')$ , by the synchronizer.

A safe run of a synchronizer satisfies the following safety properties: the result of synchronization must reflect all user changes, it must not include changes that do not come from either replica, trees under a conflicting node should remain untouched, and the results should belong to the schema.

**Definition 6 (Locally safe run)** A run is locally safe, written  $\text{locallysafe}(S, o, a, b, o', a', b')$ , iff:

- (1) It does not discard changes:  $(o \approx a \implies a' \sim a)$  and  $(o \approx b \implies b' \sim b)$ ,
- (2) It does not “make up” content:  $(a \approx a' \implies b \sim a')$  and  $(b \approx b' \implies a \sim b')$  and  $(o' \neq \mathcal{X} \implies o' \sim a' \wedge o' \sim b')$ ,
- (3) It stops at conflicting paths (leaving replicas in their current states and recording the conflict):  $\text{conflict}(o, a, b) \implies (a', b', o') = (a, b, \mathcal{X})$ ,
- (4) The results belong to the schema (or are missing):  $a' \in S_\perp$  and  $b' \in S_\perp$ .

**Definition 7 (Safe run)** A run  $(S, o, a, b, o', a', b')$  is safe, written  $\text{safe}(S, o, a, b, o', a', b')$  iff it is locally safe on every path  $p$ :

$$\forall p \in \mathcal{P}. \text{locallysafe}(S(p), o(p), a(p), b(p), o'(p), a'(p), b'(p))$$

**Definition 8 (Maximal run)** A run  $(S, o, a, b, o', a', b')$  is maximal iff it is safe and propagates at least the same changes as any other safe run, i.e.

$$\begin{aligned} \forall o'', a'', b''. \text{safe}(S, o, a, b, o'', a'', b'') \implies \\ \left\{ \begin{array}{l} \forall p \in \mathcal{P}. a''(p) \sim b''(p) \implies a'(p) \sim b'(p) \\ \forall p \in \mathcal{P}. o''(p) \neq \mathcal{X} \implies o'(p) \neq \mathcal{X}. \end{array} \right. \end{aligned}$$

Next we give a precise specification of Harmony’s synchronization algorithm.

**Theorem 9** Let  $S$  be a path-consistent schema. If  $a, b \in S_\perp$  and  $\text{sync}(S, o, a, b)$  evaluates to  $(o', a', b')$ , then  $(S, o, a, b, o', a', b')$  is maximal.

Before proving Theorem 9, we prove a few technical lemmas. First we note that path consistency is stable under projection on a single name.

**Lemma 10** Let  $n$  be a name and  $S$  a path-consistent schema; the schema  $S(n)$  is also path consistent.

**Proof.** Let  $t_n$  and  $t'_n$  be trees in  $S(n)$ , and  $p$  a path with  $t_n(p) \in \mathcal{T}$  and  $t'_n(p) \in \mathcal{T}$ . As  $t_n$  and  $t'_n$  are in  $S(n)$ , there exist  $t$  and  $t'$  in  $S$  with  $t(n) = t_n$  and  $t'(n) = t'_n$ . As  $S$  is path consistent and  $t(n/p) = t_n(p) \in \mathcal{T}$  and  $t'(n/p) =$

$t'_n(p) \in \mathcal{T}$ , we have

$$\begin{aligned}
t[n/p \mapsto t'(n/p)] &= t[n/p \mapsto t'_n(p)] \\
&= \left\{ \begin{array}{l} n \mapsto t(n)[p \mapsto t'_n(p)] \\ m \mapsto t(m) \quad \text{for } m \in \text{dom}(t) \setminus \{n\} \end{array} \right\} \\
&= \left\{ \begin{array}{l} n \mapsto t_n[p \mapsto t'_n(p)] \\ m \mapsto t(m) \quad \text{for } m \in \text{dom}(t) \setminus \{n\} \end{array} \right\} \in S
\end{aligned}$$

hence  $t_n[p \mapsto t'_n(p)] \in S(n)$ .  $\square$

The key property enjoyed by path-consistent schemas is that we can assemble well-formed trees (belonging to the schema) out of well-formed subtrees.

**Lemma 11** *Let  $S$  be a path-consistent schema. If  $\text{dom}(t) \in \text{doms}(S)$  and for each  $k \in \text{dom}(t)$  we have  $t(k) \in S(k)$  then  $t \in S$ .*

**Proof.** Let  $\text{dom}(t) = \{n_1, \dots, n_j\}$ . As  $t(k) \in S(k)$  for every  $k \in \text{dom}(t)$ , for each  $k \in \{n_1 \dots n_j\}$  there exists a tree  $t'_k \in S$  with  $t'_k(k) = t(k)$ . Moreover, as  $\text{dom}(t) \in \text{doms}(S)$ , there is a tree  $t'' \in S$  with  $\text{dom}(t'') = \text{dom}(t)$ . We now show by a finite induction on  $i \leq j$  that  $t''[n_1 \mapsto t'_{n_1}(n_1)] \dots [n_{n_i} \mapsto t'_{n_i}(n_i)] \in S$ .

**Case ( $i = 1$ ):** follows immediately from path consistency, as both  $t''$  and  $t'_{n_1}$  are in  $S$ , and as both  $t''(n_1) \in \mathcal{T}$  and  $t'_{n_1}(n_1) \in \mathcal{T}$ .

**Case ( $i > 1$ ):** By IH,  $t''' = t''[n_1 \mapsto t'_{n_1}(n_1)] \dots [n_{i-1} \mapsto t'_{n_{i-1}}(n_{i-1})] \in S$ . We also have  $t'_{n_i} \in S$  by assumption. Moreover, both  $t'''(n_i) = t''(n_i) \in \mathcal{T}$  and  $t'_{n_i}(n_i) \in \mathcal{T}$ . By path consistency for  $S$ , we hence have  $t'''[n_i \mapsto t'_{n_i}(n_i)] \in S$ . We conclude that  $t = t''[n_1 \mapsto t'_{n_1}(n_1)] \dots [n_j \mapsto t'_{n_j}(n_j)] \in S$ .  $\square$

Next we lemma record some examples of safe runs (both to illuminate the definitions and to shorten the proof of Theorem 9).

**Lemma 12** *If  $a \in S_\perp$  and  $b \in S_\perp$  then the following runs are safe:*

- |                                       |  |
|---------------------------------------|--|
| (1) $(S, o, a, b, \mathcal{X}, a, b)$ | (4) $(S, o, a, o, a, a, a)$  |
| (2) $(S, o, a, a, a, a, a)$           | (5) $(S, o, \perp, b, \perp, \perp, \perp)$ , assuming $b \sqsubseteq o$ |
| (3) $(S, o, o, b, b, b, b)$           | (6) $(S, o, a, \perp, \perp, \perp, \perp)$ , assuming $a \sqsubseteq o$ |

**Proof.** Let  $p$  be a path.

(1) We must prove  $\text{locallysafe}(S(p), o(p), a(p), b(p), \mathcal{X}(p), a(p), b(p))$ . Conditions (1,2) are satisfied because each replica output is identical to the corresponding input (and because, as  $\sim$  is reflexive,  $a(p) \sim a(p)$  and  $b(p) \sim b(p)$ ), and as  $\mathcal{X}(p) = \mathcal{X}$ . Condition (3) is also trivially satisfied as  $a(p) = a(p)$ ,  $b(p) = b(p)$ , and  $\mathcal{X}(p) = \mathcal{X}$ . The final condition (4) is satisfied as, by the definition of schema projection, we have  $a(p) \in S(p)_\perp$  and  $b(p) \in S(p)_\perp$ .

(2) We must prove  $\text{locallysafe}(S(p), o(p), a(p), a(p), a(p), a(p), a(p))$ . Conditions (1,2) are satisfied because both input replicas and all the outputs are equal to  $a(p)$ . Condition (3) is satisfied because there is no conflict at  $p$ . Both clauses in the definition of local conflict are not satisfied because the replicas are identical. We argue that there cannot be a schema domain conflict by cases on  $a(p)$ . If  $a(p) = \perp$  then by definition there is no schema domain conflict. Otherwise,  $a(p) \neq \perp$  and we calculate  $\text{edom}(o(p), a(p), a(p))$  as follows:

$$\begin{aligned} \text{edom}(o(p), a(p), a(p)) = & \\ & \{k \in \text{dom}(a(p)) \mid a(p/k) \not\sqsubseteq o(p/k)\} \quad \left( \subseteq \text{dom}(a(p)) \right) \\ \cup & (\text{dom}(a(p)) \setminus \text{dom}(o(p/k))) \quad \left( \subseteq \text{dom}(a(p)) \right) \\ \cup & (\text{dom}(a(p)) \cap \text{dom}(a(p))) \quad \left( = \text{dom}(a(p)) \right) \end{aligned}$$

Hence,  $\text{edom}(o(p), a(p), a(p))$  is  $\text{dom}(a(p)) \in \text{doms}(S(p)_\perp)$  and there is no schema domain conflict at  $p$ . The final condition (4) holds as  $a(p) \in S(p)_\perp$ .

(3) We must prove  $\text{locallysafe}(S(p), o(p), o(p), b(p), b(p), b(p), b(p))$ . Condition (1) is satisfied because the first replica is equal to the archive,  $o(p)$ , and because the output computed for the second replica output is equal to the input for that replica. Condition (2) is satisfied because additionally we have that the output for the first replica is equal to the second replica input,  $b(p)$ . Condition (3) is satisfied because there is no conflict at  $p$ . The clauses of local conflict are not satisfied because the first replica is  $o(p)$ , and because  $o(p) \in S(p)_\perp$  implies  $o(p) \neq \mathcal{X}$ . As before, we argue that there is no schema domain conflict by cases. If  $b(p) = \perp$  then by definition, there is no schema domain conflict at  $p$ . Otherwise,  $b(p) \neq \perp$  and we calculate  $\text{edom}(o(p), o(p), b(p))$  directly:

$$\begin{aligned} \text{edom}(o(p), o(p), b(p)) = & \\ & \{k \in \text{dom}(o(p)) \mid o(p/k) \not\sqsubseteq o(p/k)\} \quad \left( = \emptyset \right) \\ \cup & (\text{dom}(b(p)) \setminus \text{dom}(o(p))) \\ \cup & (\text{dom}(o(p)) \cap \text{dom}(b(p))) \end{aligned}$$

In the first clause we use the fact that if  $t \neq \perp$ , then we have  $t \not\sqsubseteq \perp$ . Hence if  $k \in \text{dom}(b(p)) \setminus \text{dom}(o(p))$ , then  $b(p/k) \neq \perp$  and  $o(p/k) = \perp$  for any child

k. We calculate  $\text{edom}(o(p), b(p), o(p))$  similarly:

$$\begin{aligned} \text{edom}(o(p), b(p), o(p)) = & \\ & \{k \in \text{dom}(b(p)) \mid b(p/k) \not\sqsubseteq o(p/k)\} \left( \subseteq \text{dom}(b(p)) \setminus \text{dom}(o(p)) \right) \\ & \cup (\text{dom}(o(p)) \setminus \text{dom}(o(p))) \quad \left( = \emptyset \right) \\ & \cup (\text{dom}(o(p)) \cap \text{dom}(b(p))) \end{aligned}$$

It is easy to check that both domains are equal to  $\text{dom}(b(p)) \in \text{doms}(S(p)_\perp)$ . Hence, there is no schema domain conflict at path  $p$ . The final local safety condition (4) is satisfied as  $b(p) \in S(p)_\perp$  by definition.

- (4) Symmetric to the previous case, inverting the roles of  $a$  and  $b$ .
- (5) We must prove  $\text{locallysafe}(S(p), o(p), \perp(p), b(p), \perp(p), \perp(p), \perp(p))$ , assuming  $b \sqsubseteq o$ . Condition (1) for the first replica is satisfied because the output and input are both  $\perp$ . We argue that the second replica satisfies condition (1) by cases. If  $o(p) = \perp$ ,  $o(p) = \mathcal{X}$ , or  $b(p) = \perp$  then  $b(p) = \perp$  either by  $b \sqsubseteq o$  or assumption. Otherwise,  $o(p) \neq \perp$  and  $o(p) \neq \mathcal{X}$ . From  $b(p) \in S(p)_\perp$  we have  $b(p) \neq \mathcal{X}$  and by assumption we have  $b(p) \neq \perp$ . Hence,  $b(p) \sim o(p)$ . Local safety condition (2) is satisfied because all three outputs are  $\perp$ , which is the first replica input. Local safety condition (3) is satisfied because there is no conflict at  $p$ . The first clause of a local conflict is not satisfied because if  $o(p) = \mathcal{X}$ , then by  $b \sqsubseteq o$  we calculate that the second replica,  $b(p)$  is  $\perp$ , which is equal to the first replica,  $\perp(p)$ . The second clause is not satisfied because the first replica is  $\perp$  and  $b \sqsubseteq o$ . Moreover there cannot be a schema domain conflict as the first replica is  $\perp$ . The final local safety condition (4) is trivially satisfied as  $a' = b' = \perp \in \text{doms}(S(p)_\perp)$ .
- (6) Symmetric to the previous case, inverting the roles of  $a$  and  $b$ .  $\square$

**Proof of Theorem 9.** By induction on the sum of the depth of  $o$ ,  $a$ , and  $b$ , with a case analysis according to the first rule in the algorithm that applies.

**Case ( $a = b$ ):**  $\text{sync}(S, o, a, b) = (a, a, b)$

By Lemma 12(2) the run  $(S, o, a, a, a, a, a)$  is safe. We must show that  $(S, o, a, a, a, a, a)$  is maximal. The first condition for maximality is immediate as for all paths  $p$ , we have  $a'(p) \sim b'(p)$ . The second condition is also satisfied, since  $o' = a$ , hence we have  $o'(p) \neq \mathcal{X}$  for all paths  $p$ .

**Case ( $a = o$ ):**  $\text{sync}(S, o, a, b) = (b, b, b)$

By Lemma 12(3) the run  $(S, o, o, b, b, b, b)$  is safe. We must show that  $(S, o, o, b, b, b, b)$  is maximal. The first condition for maximality is immediate as for all paths  $p$ , we have  $a'(p) \sim b'(p)$ . The second condition is also satisfied, since  $o' = b$ , hence we have  $o'(p) \neq \mathcal{X}$  for all paths  $p$ .

**Case ( $b = o$ ):**  $\text{sync}(S, o, a, b) = (a, a, a)$

Symmetric to the previous case, inverting the roles of  $a$  and  $b$ .

**Case ( $o = \mathcal{X}$ ):**  $\text{sync}(S, o, a, b) = (\mathcal{X}, a, b)$

By Lemma 12(1) the run  $(S, \mathcal{X}, a, b, \mathcal{X}, a, b)$  is safe. We must now show that

the run is maximal. The predicate  $\text{localconflict}(\mathcal{X}, a, b)$  is satisfied because we know that  $o = \mathcal{X}$  and  $a \neq b$  (as the first case of the algorithm did not apply). Therefore, by safety condition (3), the only safe run is  $(S, \mathcal{X}, a, b, \mathcal{X}, a, b)$ , hence it is maximal.

**Case** ( $a = \perp$ ): there are two subcases to consider.

**Subcase** ( $b \sqsubset o$ ):  $\text{sync}(S, o, a, b) = (\perp, \perp, \perp)$

By Lemma 12(5) the run  $(S, o, \perp, b, \perp, \perp, \perp)$  is safe if  $b \sqsubset o$ . The first maximality condition is immediate as for all paths  $p$ , we have  $a'(p) \sim b'(p)$ . The second condition is also satisfied since  $o' = \perp$ , hence we have  $o'(p) \neq \mathcal{X}$  for all paths  $p$ .

**Subcase** ( $b \not\sqsubset o$ ):  $\text{sync}(S, o, a, b) = (\mathcal{X}, \perp, b)$

By Lemma 12(1) the run  $(S, o, \perp, b, \mathcal{X}, \perp, b)$  is safe. We must now prove that the run is maximal. None of the previous rules apply, so we must have  $b \neq a = \perp$ ,  $o \neq a = \perp$ , and  $b \neq o$ . Since  $a = \perp$  and  $b \neq \perp$ , we have  $a \approx b$ . Moreover, we have  $b \not\sqsubset o$  and  $b \neq \perp$ . Hence the predicate  $\text{localconflict}(o, a, b)$  is satisfied. As before, by safety condition (3), the only safe run is  $(S, \mathcal{X}, a, b, \mathcal{X}, a, b)$ , hence it is maximal.

**Case** ( $b = \perp$ ): symmetric to the previous case, inverting the roles of  $a$  and  $b$  and using Lemma 12(6) instead of Lemma 12(5) in the first subcase.

**Recursive case:** Since previous cases of the algorithm do not apply, we have  $a \neq b$ ,  $o \neq \mathcal{X}$ ,  $a \neq \perp$ ,  $b \neq \perp$ , and  $a \sim b$ .

By Lemma 10, each of the schemas  $S(k)$  are path consistent for  $k \in \text{dom}(a) \cup \text{dom}(b)$ . By the definition of schema projection, for each  $k$  we have  $a(k), b(k) \in S(k)_\perp$ . Thus, by the IH (which applies as  $a \neq \perp \neq b$ , hence the sum of the depths decreases), each recursive sub-run is maximal. In particular, each sub-run is safe and gives results in  $S(k)_\perp$ .

Let  $o'_r$ ,  $a'_r$ , and  $b'_r$  be trees obtained by the recursive calls, i.e.  $(o'_r(k), a'_r(k), b'_r(k)) = \text{sync}(S(k), o(k), a(k), b(k))$  for all  $k \in \text{dom}(a) \cup \text{dom}(b)$ . From the facts obtained by IH and the definition of safety, it follows that the run  $(S, o, a, b, o'_r, a'_r, b'_r)$  is locally safe at every non-empty path. We now check that it is locally safe at the empty path and maximal.

We first consider local safety. We start by showing that  $\text{dom}(a'_r) = \text{edom}(o, a, b)$ . For this let  $k \in \text{dom}(a) \cup \text{dom}(b)$  and consider the shape of possible tuples  $(o(k), a(k), b(k))$ . We let the variables  $t_o, t_a, t_b$  range over elements of  $\mathcal{T}$ :

**Subcases**  $(\mathcal{X}, \perp, \perp)$ ,  $(\perp, \perp, \perp)$ ,  $(t_o, \perp, \perp)$ : Cannot happen because the given child  $k \in \text{dom}(a) \cup \text{dom}(b)$ .

**Subcases**  $(\perp, t_a, \perp)$ ,  $(\perp, t_a, t_b)$ : Since  $o(k) \approx a(k)$  we get by local safety condition (1) that  $a'_r(k) \sim t_a$  and in particular  $a'_r(k) \neq \perp$ , hence  $k \in \text{dom}(a'_r)$ . On the other hand, since  $k \in \{n \in \text{dom}(a) \mid a(n) \not\sqsubset o(n)\}$  we also have that  $k \in \text{edom}(o, a, b)$ .

**Subcases**  $(\mathcal{X}, t_a, \perp)$ ,  $(\mathcal{X}, t_a, t_b)$ : Same as previous case, as  $t_a \not\sqsubset \mathcal{X}$ .

**Subcases**  $(t_o, \perp, t_b)$ ,  $(\mathcal{X}, \perp, t_b)$ : Since  $o(k) \approx a(k)$  we get by local safety condition (1) that  $a'_r(k) \sim \perp$  and therefore  $a'_r(k) = \perp$ , so  $k \notin \text{dom}(a'_r)$ . On the other hand, since  $a(k) = \perp$  we have that  $k \notin \text{dom}(a)$  so especially

$k \notin \{n \in \text{dom}(a) \mid a(n) \not\sqsubseteq o(n)\}$ ,  $k \notin \text{dom}(a) \cap \text{dom}(b)$ , and since  $o(k) \neq \perp$  we have that  $k \notin \{n \in \text{dom}(b) \mid o(n) = \perp\}$ , hence also  $k \notin \text{edom}(o, a, b)$ .

**Subcase**  $(\perp, \perp, t_b)$ : By Lemma 12(3) the run  $(S(k), \perp, \perp, t_b, t_b, t_b, t_b)$  is safe. Letting  $a''$  and  $b''$  stand for the replica outputs of this run, which are both  $t_b$ , we have that  $a'' \sim b''$ . By the maximality of the recursive sub-run we get that  $a'_r(k) \sim b'_r(k)$ . As  $t_b \approx \perp$ , we have  $b'_r(k) \sim t_b$  thus  $b'_r(k) \neq \perp$  by local safety (1), hence  $a'_r(k) \neq \perp$ . From this, we have  $k \in \text{dom}(a'_r)$ . On the other hand, since  $k \in \{n \in \text{dom}(b) \mid o(n) = \perp\}$  we also have that  $k \in \text{edom}(o, a, b)$ .

**Subcase**  $(t_o, t_a, \perp)$ : We first consider the case  $t_a \not\sqsubseteq t_o$ , which immediately implies  $t_a \neq t_o$ . In this case, as we have  $t_a \neq t_o$ ,  $t_o \neq \perp$ ,  $t_a \approx \perp$ ,  $t_a \neq \perp$ , and  $t_a \not\sqsubseteq t_o$ , the predicate  $\text{localconflict}(t_o, t_a, \perp)$  is satisfied. By local safety condition (3), we have  $a'_r(k) = t_a$  so  $k \in \text{dom}(a'_r)$ . On the other hand, since  $a(k) \not\sqsubseteq o(k)$  we have that  $k \in \{n \in \text{dom}(a) \mid a(n) \not\sqsubseteq o(n)\}$ , hence also  $k \in \text{edom}(o, a, b)$ .

We now consider the other case,  $t_a \sqsubseteq t_o$ . By Lemma 12(6) the run  $(S(k), t_o, t_a, \perp, \perp, \perp, \perp)$  is safe as  $t_a \sqsubseteq t_o$ . Letting  $a''$  and  $b''$  stand for the replica outputs of this run, which are both  $\perp$ , we have that  $a'' \sim b''$ . Therefore, by maximality of the recursive sub-run we get that  $a'_r(k) \sim b'_r(k)$ . Next we show that  $a'_r(k) = \perp$ . Since  $t_o \approx \perp$  we have  $b'_r(k) \sim \perp$  by local safety (1) and hence  $a'_r(k) = \perp$ . It follows that  $k \notin \text{dom}(a'_r)$ . On the other hand, since  $a(k) \sqsubseteq o(k)$  we have  $k \notin \{n \in \text{dom}(a) \mid a(n) \not\sqsubseteq o(n)\}$ , and since  $k \notin \text{dom}(b)$  we have  $k \notin \text{dom}(a) \cap \text{dom}(b)$  and  $k \notin \{n \in \text{dom}(b) \mid o(n) = \perp\}$ , hence also  $k \notin \text{edom}(o, a, b)$ .

**Subcase**  $(t_o, t_a, t_b)$ : We first show that  $a'_r(k) \neq \perp$ , by contradiction. Assume  $a'_r(k) = \perp \approx t_a$ , hence by local safety condition (2) we have  $t_b \sim a'_r(k) = \perp$ , a contradiction. Thus  $k \in \text{dom}(a'_r)$ . As  $k \in \text{dom}(a) \cap \text{dom}(b)$ , we conclude by  $k \in \text{edom}(o, a, b)$ .

By a symmetric argument we can show that  $\text{dom}(b'_r) = \text{edom}(o, b, a)$ . Returning now to our original argument that the run is locally safe at the root, we first consider the case where  $\text{dom}(a'_r) \notin \text{doms}(S)$  or  $\text{dom}(b'_r) \notin \text{doms}(S)$ . In this case  $o' = \mathcal{X}$ ,  $a' = a$ , and  $b' = b$  and it follows immediately from Lemma 12(1) that the run  $(\mathcal{X}, a, b, \mathcal{X}, a, b)$  is safe. To show that it is maximal we show that there is a schema domain conflict at the root. We have that  $\text{edom}(o, a, b) = \text{dom}(a'_r) \notin \text{doms}(S)$  or  $\text{edom}(o, b, a) = \text{dom}(b'_r) \notin \text{doms}(S)$ , so  $\text{sdomconflict}(S, o, a, b)$  holds. As above, by safety condition (3), the only safe run is  $(\mathcal{X}, a, b, \mathcal{X}, a, b)$ , hence it is maximal.

Finally, we consider the case where  $\text{dom}(a'_r) \in \text{doms}(S)$  and  $\text{dom}(b'_r) \in \text{doms}(S)$ . In this case  $o' = o'_r$ ,  $a' = a'_r$ , and  $b' = b'_r$  and since these are assembled as the concatenation of the results of each recursive call, we have  $o' \neq \perp$ ,  $a' \neq \perp$ , and  $b' \neq \perp$  (recall the difference between the empty tree and the missing tree). It follows directly that  $a \sim a'$ ,  $b \sim b'$ ,  $a' \sim b'$ ,  $o' \sim a'$ , and  $o' \sim b'$ , which immediately satisfies local safety conditions (1, 2). Local safety condition (3) is satisfied because there cannot be a conflict at the root: there is no local conflict because  $o \neq \mathcal{X}$  and  $a \sim b$ , and there is no

schema domain conflict because  $\text{edom}(o, a, b) = \text{dom}(a'_r) \in \text{doms}(S)$  and  $\text{edom}(o, b, a) = \text{dom}(b'_r) \in \text{doms}(S)$ . For the final safety condition (4), we show that  $a', b' \in S_\perp$ . As each sub-run is maximal (and hence, safe), for every  $k \in \text{dom}(a')$  we have  $a'(k) \in S(k)_\perp$ . Also, since  $k \in \text{dom}(a')$  we have  $a'(k) \neq \perp$  and so  $a'(k) \in S(k)$ . We also have that  $\text{dom}(a') \in \text{doms}(S)$ . By Lemma 11,  $a' \in S$  and hence  $a' \in S_\perp$ . By a symmetric argument, we have  $b' \in S_\perp$ . We conclude that the run is locally safe at the root.

To finish the proof, let  $(S, o, a, b, o'', a'', b'')$  be another safe run and let  $p$  be a path. We show that both maximality conditions are satisfied by cases on  $p$ :

**Subcase  $p = \bullet$ :** Both conditions trivially hold as  $a' \sim b'$  and  $o' \neq \mathcal{X}$ .

**Subcase  $p = k/p'$ :** The sub-run  $(S(k), o(k), a(k), b(k), o''(k), a''(k), b''(k))$  is also safe by definition. We have  $a''(p) = a''(k/p') = (a''(k))(p')$ , and  $b''(p) = b''(k/p') = (b''(k))(p')$ .

- To show the first maximality condition, observe that if  $a''(p) \sim b''(p)$ , then  $a''(p) = (a''(k))(p') \sim (b''(k))(p') = b''(p)$ . By IH, the run  $(o(k), a(k), b(k), o'(k), a'(k), b'(k))$  is maximal. It follows that  $(a'(k))(p') \sim (b'(k))(p')$ , hence  $a'(p) \sim b'(p)$ .
- For the second condition, observe that if  $o''(p) \neq \mathcal{X}$ , then  $(o''(k))(p') \neq \mathcal{X}$ . By IH, the run  $(o(k), a(k), b(k), o'(k), a'(k), b'(k))$  is maximal and so  $(o'(k))(p') \neq \mathcal{X}$ . Hence, we have  $o'(p) \neq \mathcal{X}$ .  $\square$

## 7 Case Study: Address Books

We now present a brief case study, illustrating how schemas can be chosen so as to obtain desirable behavior from our generic synchronizer on structures of realistic complexity. The examples use an address book schema loosely based on the vCard standard [7], which embodies some of the tricky issues that can arise when synchronizing larger structures with varied substructure. We begin with a concrete notation for writing down schemas.

**Schemas** Schemas are given by mutually recursive equations of the form  $\mathbf{x} = \mathbf{S}$ , where  $\mathbf{S}$  is an expression generated by the following grammar:

$$\mathbf{S} ::= \{\} \mid \mathbf{n}[\mathbf{S}] \mid !(\mathbf{F})[\mathbf{S}] \mid *(\mathbf{F})[\mathbf{S}] \mid \mathbf{S}, \mathbf{S} \mid \mathbf{S} \mid \mathbf{S}$$

The symbols  $\mathbf{n}$  and  $\mathbf{F}$  range over names in  $\mathcal{N}$  and finite sets in  $\mathcal{P}(\mathcal{N})$ , respectively. The first form of schema,  $\{\}$ , denotes the singleton set containing the empty tree;  $\mathbf{n}[\mathbf{S}]$  denotes the set of trees with a single child named  $\mathbf{n}$  where the subtree under  $\mathbf{n}$  is in  $\mathbf{S}$ ; the wildcard schema  $!(\mathbf{F})[\mathbf{S}]$  denotes the set of trees with *any* single child whose name is not in  $\mathbf{F}$ , where the subtree under that child is in  $\mathbf{S}$ ; the other wildcard schema,  $*(\mathbf{F})[\mathbf{S}]$  denotes the set of trees with *any number of* children with names not in  $\mathbf{F}$  where the subtree under

```

C = name[N],work[V],home?[V],org[O],email[E]
  | name[N],work?[V],home[V],org?[O],email[E]
E = !(pref,alts)[{}]|pref[V],alts[VS]
N = first[V],other?[VL],last[V]
O = name[V],unit[V]
V = ![{}]
VL = head[V],tail[VL]|nil[{}]
VS = *[{}]

```

Fig. 3. Address book schemas

each child is in  $S$ . The set described by  $S_1|S_2$  is the union of the sets described by  $S_1$  and  $S_2$ , while  $S_1,S_2$  denotes the set of trees of the form  $t_1 \cdot t_2$  where  $t_1$  belongs to  $S_1$  and  $t_2$  to  $S_2$ . Note that, as trees are unordered, the “,” operator is commutative (e.g.,  $n[X],m[Y]$  and  $m[Y],n[X]$  are equivalent). We abbreviate  $n[S]|[{}]$  as  $n?[S]$ , and likewise  $!(\emptyset)[S]$  as  $![S]$  and  $*(\emptyset)[S]$  as  $*[S]$ . In the following, all the schemas we write in this section are path consistent, a fact that can be verified syntactically: if a name appears twice in a node, like  $m$  in  $m[X],n[Y]|m[X],o[Z]$ , the subschemas associated with each occurrence of the name are textually identical.

Note that our schema formalism, like our generic synchronization algorithm, does not distinguish between “structure” edges and “data” edges in trees—this distinction is just a matter of convention in each particular application. This choice makes our technical work easier by simplifying the formal objects (trees and schemas) that we manipulate.

**Address Book Schema** Here is a typical contact (the notation  $[t_1;\dots;t_n]$ , which represents a list encoded as a tree, is explained below):

$$o = \left\{ \begin{array}{l} \text{name} \mapsto \{\text{first} \mapsto \text{Meg}, \text{other} \mapsto [\text{Liz}; \text{Jo}], \text{last} \mapsto \text{Smith}\}, \\ \text{email} \mapsto \{\text{pref} \mapsto \text{ms@c.edu}, \text{alts} \mapsto \text{meg@s.com}\}, \\ \text{org} \mapsto \{\text{name} \mapsto \text{City U}, \text{unit} \mapsto \text{CS Dept}\}, \\ \text{home} \mapsto 555-6666, \\ \text{work} \mapsto 555-7777 \end{array} \right\}$$

There are two sorts of contacts—“professional” contacts, which contain mandatory work phone and organization entries, plus, optionally, a home phone, and “personal” ones, which have a mandatory home phone and, optionally, a work phone and organization information. Contacts are not explicitly tagged with sorts, so some contacts, like the one for **Meg**, belong to both. Each contact also has fields representing name and email data.

The schema  $C$ , displayed in Figure 3, describes the record-like structure of both sorts of contacts. One level down, the subtrees that may validly appear below the **home** and **work** children are simple string values—i.e., trees with a single child leading to the empty tree—that belong to the  $V$  schema. The subtree be-

low the **name** child in a valid contact belongs to the **N** schema, which describes trees with a record-like structure containing mandatory **first** and **last** children, and optionally a child **other**. The **first** and **last** fields lead to values belonging to the **V** schema. The **other** field leads to a *list* of alternate names (e.g., middle names and nicknames) stored (for the sake of the example) in some particular order. Because our actual trees are unordered, we use a standard “cons cell” representation (as in Lisp) to encode ordered lists:  $[t_1; \dots; t_n]$  is encoded as  $\{\text{head} \mapsto t_1, \text{tail} \mapsto \{\dots \mapsto \{\text{head} \mapsto t_n, \text{tail} \mapsto \text{nil}\} \dots\}\}$ . The schema **VL** describes lists of values encoded like this. The email data for a contact is either a value, or a set of addresses with one distinguished “preferred” address. The **E** schema describes these structures using a union of a wildcard to represent single values (excluding **pref** and **alts** to ensure path consistency) and a record-like structure with fields **pref** and **alts** to represent sets of addresses. The schema **VS** describes bushes with any number of children, each leading to the empty tree, which are a natural way of encoding sets of values. Finally, organization data is represented by trees with two children, **name** and **unit**, each leading to values, as described by the **O** schema.

**The Need For Schemas** To illustrate how and where schema conflicts can occur, let us see what can go wrong when *no* schema information is used. We consider four runs of the synchronizer using the universal schema **Any** =  $*[\text{Any}]$  (whose denotation is  $\mathcal{T}$ ); each run shows a different example where schema-ignorant synchronization produces mangled results. In each case, the archive,  $o$ , is the tree above.

Suppose, first, that the  $a$  replica is obtained by deleting the **work** and **org** children, making the entry personal, and that the  $b$  replica is obtained by deleting the **home** child, making the entry professional:

$$a = \left\{ \begin{array}{l} \text{name} \mapsto \{\text{first} \mapsto \text{Meg}, \text{other} \mapsto [\text{Liz}; \text{Jo}], \text{last} \mapsto \text{Smith}\}, \\ \text{email} \mapsto \{\text{pref} \mapsto \text{ms@c.edu}, \text{alts} \mapsto \text{meg@s.com}\}, \\ \text{home} \mapsto 555-6666 \end{array} \right\}$$

$$b = \left\{ \begin{array}{l} \text{name} \mapsto \{\text{first} \mapsto \text{Meg}, \text{other} \mapsto [\text{Liz}; \text{Jo}], \text{last} \mapsto \text{Smith}\}, \\ \text{email} \mapsto \{\text{pref} \mapsto \text{ms@c.edu}, \text{alts} \mapsto \text{meg@s.com}\}, \\ \text{org} \mapsto \{\text{name} \mapsto \text{City U}, \text{unit} \mapsto \text{CS Dept}\}, \\ \text{work} \mapsto 555-7777 \end{array} \right\}$$

Although  $a$  and  $b$  are both valid address book contacts, the trees that result from synchronizing them with respect to the **Any** schema are not, since they have the structure neither of personal nor of professional contacts:

$$a' = b' = \left\{ \begin{array}{l} \text{name} \mapsto \{\text{first} \mapsto \text{Meg}, \text{other} \mapsto [\text{Liz}; \text{Jo}], \text{last} \mapsto \text{Smith}\}, \\ \text{email} \mapsto \{\text{pref} \mapsto \text{ms@c.edu}, \text{alts} \mapsto \text{meg@s.com}\} \end{array} \right\}$$

Now suppose that the replicas are obtained by updating the trees along the path **name/first**, replacing **Meg** with **Maggie** in  $a$  and **Megan** in  $b$  (from now

on, for the sake of brevity we only show the parts of the tree that are different from  $o$  and elide the rest):  $o(\text{name/first}) = \text{Meg}$ ,  $a(\text{name/first}) = \text{Maggie}$ , and  $b(\text{name/first}) = \text{Megan}$ . Synchronizing with respect to the **Any** schema yields results where *both* names appear under **first**:  $a'(\text{name/first}) = b'(\text{name/first}) = \{\text{Maggie}, \text{Megan}\}$ . These results are ill-formed because they do not belong to the **V** schema, which describes trees that have a *single* child.

Next consider updates to the email information where the  $a$  replica replaces the set of addresses in  $o$  with a single address, and  $b$  updates both **pref** and **alts** children in  $b$ :  $o(\text{email}) = \{\text{pref} \mapsto \text{ms@c.edu}, \text{alts} \mapsto \text{meg@s.com}\}$ ,  $a(\text{email}) = \{\text{meg@s.com}\}$ , and  $b(\text{email}) = \{\text{pref} \mapsto \text{meg.smith@cs.c.edu}, \text{alts} \mapsto \text{ms@c.edu}\}$ . Synchronizing these trees with respect to **Any** propagates the addition of the edge labeled **meg@s.com** from  $a$  to  $b'$  and yields conflicts on both **pref** and **alts** children, since both have been deleted in  $a$  but modified in  $b$ . The results after synchronizing are thus:  $a'(\text{email}) = \text{meg@s.com}$  and  $b'(\text{email}) = \{\text{meg@s.com}, \text{pref} \mapsto \text{meg.smith@cs.c.edu}, \text{alts} \mapsto \text{ms@c.edu}\}$ . The second result,  $b'$ , is ill-formed because it contains three children, whereas all the trees in the email schema **E** have either one or two children.

Next consider changes to the list of names along the path **name/other**. Suppose that  $a$  removes both **Liz** and **Jo**, but  $b$  only removes **Jo**:  $o(\text{name/other}) = [\text{Liz}; \text{Jo}]$ ,  $a(\text{name/other}) = []$ , and  $b(\text{name/other}) = [\text{Liz}]$ . Comparing the  $a$  replica to  $o$ , both **head** and **tail** are deleted and **nil** is newly added. Examining the  $b$  replica, the tree under **head** is identical to corresponding tree in  $o$  but deleted from  $a$ . The tree under **tail** is different from  $o$  but deleted from  $a$ . Collecting all of these changes, the algorithm yields these results:  $a'(\text{name/other}) = \text{nil}$  and  $b'(\text{name/other}) = \{\text{tail} \mapsto \text{nil}, \text{nil}\}$ . Here again, the second result,  $b'$ , is ill-formed: it has children **tail** and **nil**, which is not a valid encoding of a list.

Situations like these—invalid records, multiple children where a single value is expected, and mangled lists—provided the initial motivation for equipping a straightforward “tree-merging” synchronization algorithm with schema information. Fortunately, in all of these examples, the step that breaks the structural invariant can be detected by a simple, local, domain test. In the first example, where the algorithm removed the **home**, **work**, and **org** children, the algorithm tests if  $\{\text{name}, \text{email}\}$  is in  $\text{doms}(\mathbf{C})$ . Similarly, in the second example, where both replicas changed the **first** name to a different value, the algorithm tests if  $\{\text{Maggie}, \text{Megan}\}$  is in  $\text{doms}(\mathbf{V})$ . In the example involving the tree under **email**, the algorithm tests if the domain  $\{\text{meg@s.com}, \text{pref}, \text{alts}\}$  is in  $\text{doms}(\mathbf{E})$ . Finally, in the example where both replicas updated the list of **other** names, it tests whether  $\{\text{tail}, \text{nil}\}$  is in  $\text{doms}(\mathbf{VL})$ . All of these local tests fail and so the synchronizer halts with a schema domain conflict at the appropriate path in each case, ensuring that the results are valid according to the schema.

Next we further explore the strengths (and weaknesses) of our algorithm by studying its behavior on the structures used in address books.

**Values** The simplest structures in our address books, string values, are represented as trees with a single child that leads to the empty tree and described by  $![\{\} ]$ . When we synchronize two non-missing trees using this schema, there are three possible cases: (1) if either of  $a$  or  $b$  is identical to  $o$  then the algorithm set the results equal to the other replica; (2) if  $a$  and  $b$  are identical to each other but different to  $o$  then the algorithm preserves the equality; (3) if  $a$  and  $b$  are both different from  $o$  and each other then the algorithm reaches a schema domain conflict, setting  $o' = \mathcal{X}$ ,  $a' = a$  and  $b' = b$ . These behaviors follow from the local alignment mechanism employed in the algorithm—children with the same name are identified across replicas and so identical values are aligned with each other and distinct values synchronized separately. In the first two scenarios, the differences in  $a$  and  $b$  (with respect to  $o$ ) can be assembled into a value; in the third scenario, they cannot. That is, the algorithm enforces *atomic* updates to values, propagating updates from one side to the other only if  $o$  and either  $a$  or  $b$  are identical and otherwise leaving the replicas unchanged and signalling a schema domain conflict.

**Sets** Sets can be represented as bushes—nodes with many children, each labeled with the key of an element in the set— e.g., for sets of values, this structure is described by the schema  $*[\{\} ]$ . When synchronizing two sets of values, the synchronization algorithm *never* reaches a schema conflict; it always produces a valid result, combining the additions and deletions of values from  $a$  and  $b$ . For example, given these three trees representing value sets:  $o = \{\text{meg@s.com}\}$ , with  $a = \{\text{ms@c.edu}, \text{meg.smith@cs.c.edu}\}$  and  $b = \{\text{meg@s.com}, \text{meg.smith@cs.c.edu}\}$ . The synchronizer propagates the deletion of `meg@s.com` and the addition of two new children, `ms@c.edu` and `meg.smith@cs.c.edu`, yielding  $a' = b' = \{\text{ms@c.edu}, \text{meg.smith@cs.c.edu}\}$ , as expected.

**Records** Two sorts of record structures appear in the address book schema. The simplest records, like the one for organization data (`name[V],unit[V]`), have a fixed set of mandatory fields. Given two trees representing such records, the synchronizer aligns the common fields, which are *all* guaranteed to be present, and synchronizes the nested data one level down. It never reaches a schema domain conflict at the root of a tree representing such a record. Other records, which we call *sparse*, allow some variation in the names of their immediate children. For example, the contact schema uses a sparse record to represent the structure of each entry; some fields, like `org`, may be mandatory or optional (depending on the presence of other fields). As we saw in the preceeding section, on some inputs—namely, when the updates to the replicas cannot be combined into a tree satisfying the constraint expressed by the

sparse record schema—the synchronizer yields a schema conflict but preserves the sparse record structure.

**Lists** The address book schema uses lists of values to represent the ordered collection of optional **other** names for a contact. We included this example to show how our synchronization algorithm behaves when applied to structures that are a little beyond its intended scope—i.e., mostly consisting of unordered or rigidly ordered data, but also containing small amounts of list-structured data. Lists present special challenges, because we would like the algorithm to detect updates both to elements and to their relative position. On lists, our local alignment strategy matches up list elements by *absolute* position, leading to surprising results on some inputs. We illustrate the problem and propose a more sophisticated encoding of lists that reduces the chances of confusion.

On many runs of the synchronizer, updates to lists can be successfully propagated from one replica to the other. If either replica is identical to the archive, or if each replica modifies a disjoint subset of the elements of the list (leaving the list spine intact), then the synchronizer merges the changes successfully. At each step in a recursive tree walk only one of the elements will have changed, as in the following example (in this section, we consider lists of values, and write out the low-level tree representation for each tree):

$$\begin{aligned} o &= \{\text{head} \mapsto \text{Liz}, \text{tail} \mapsto \{\text{head} \mapsto \text{Jo}, \text{tail} \mapsto \text{nil}\}\} \\ a &= \{\text{head} \mapsto \text{Elizabeth}, \text{tail} \mapsto \{\text{head} \mapsto \text{Jo}, \text{tail} \mapsto \text{nil}\}\} \\ b &= \{\text{head} \mapsto \text{Liz}, \text{tail} \mapsto \{\text{head} \mapsto \text{Joanna}, \text{tail} \mapsto \text{nil}\}\} \end{aligned}$$

The changes under **head** are propagated from  $a$  to  $b'$ , and the changes to **tail** from  $b$  to  $a'$ , yielding results:

$$a' = b' = \{\text{head} \mapsto \text{Elizabeth}, \text{tail} \mapsto \{\text{head} \mapsto \text{Joanna}, \text{tail} \mapsto \text{nil}\}\}$$

There are some inputs, however, where synchronizing lists using the local alignment strategy and simple cons cell encoding produces strange results. Consider a run on the following inputs (where  $o$  is the same tree as above):  $o = [\text{Liz}; \text{Jo}]$ ,  $a = [\text{Jo}]$  and  $b = [\text{Liz}; \text{Joanna}]$ . Considering the changes that were made to each list from a high-level— $a$  removed the head and  $b$  renamed the second element—the result calculated for  $b'$  is surprising:  $[\text{Jo}; \text{Joanna}]$ . The algorithm does not recognize that  $\text{Jo}$  and  $\text{Joanna}$  should be aligned (because  $\text{Joanna}$  was obtained by renaming  $\text{Jo}$ ). Instead, it aligns pieces of the list by absolute position, matching  $\text{Jo}$  with  $\text{Liz}$  and  $\text{nil}$  with  $[\text{Joanna}]$ .

It is not surprising that our algorithm doesn't have an intuitive behavior when its inputs are lists. In general, detecting changes in relative position in a list requires global reasoning but our algorithm is essentially local. In order to avoid these problematic cases, we can use an alternative schema, which we

call the *keyed list schema*, for lists whose relative order matters. Keyed lists combine an alternate representation of lists with an appropriate schema to ensure that the local algorithm has reasonable (though conservative) behavior. Rather than embedding the elements under a spine of cons cells, one can lift up the value at each position into the spine of the list. For example, in the extended encoding, the list  $a$  from above is represented as the tree  $\{Jo \mapsto \{\text{head} \mapsto \{\}, \text{tail} \mapsto \text{nil}\}\}$ .<sup>7</sup> The schema for keyed lists of values is:  $\text{KVL} = !(\text{head}, \text{tail}, \text{nil})[\text{head}[\{\}], \text{tail}[\text{KVL}]] \mid \text{nil}$ . During synchronization, elements of the list are identified by the value above each cons cell; synchronization proceeds until a trivial case applies (unchanged or identical replicas), or when the two replicas disagree on the domain of an element, resulting in a schema domain conflict. In the problematic example, the algorithm terminates with a conflict at the root.

**Conclusion** The examples in this section demonstrate that schemas are a valuable addition to a synchronization algorithm: (1) we are guaranteed valid results in situations where a schema-blind algorithm would yield mangled results; (2) by selecting an appropriate encoding and schema for application data (moving keys high in the tree, etc.), we can tune the behavior of the generic algorithm to work well with a variety of structures. While building demos using our prototype implementation, we have found that this works well with rigidly structured data (e.g., values and records) and unstructured data (e.g., sets of values), but so far has limited utility when used with ordered and semi-structured data (e.g., lists and documents). In the future, we hope to extend our algorithm to better handle ordered data.

## 8 Related Work

Harmony combines a synchronization component with a view update component for dealing with alignment and heterogeneous replicas. The view update language and related work in that area are described in another paper [1]. Here, we focus on work related to optimistic replication and synchronization. For further information on this area, we recommend an excellent article by Saito and Shapiro [8] surveying the area of optimistic replication. In the taxonomy of the survey, Harmony is a multi-master state-transfer system, recognizing sub-objects and manually resolving conflicts. Harmony is further distinguished by some distinctions not covered in that survey: it is generic, loosely coupled to applications, able to synchronize heterogeneous representations, and supports

---

<sup>7</sup> For keyed lists of values, we could drop the child `head`, which always maps to the empty tree. However, we can also form keyed lists of arbitrary trees, not just values. For trees with structure richer than values, we would use some key field at the top of each cons cell and store the rest of the tree under `head`.

*unsupervised* operation. Supporting unsupervised runs (where Harmony does as much work as it can, and leaves conflicts for later) leads us to a design where conflicts result in divergent replicas, and also requires that the synchronizer has behavior that is intuitive and easy to predict.

**Synchronization Architectures** Harmony’s state-based architecture only requires loose coupling between the synchronizer and applications. By contrast, in an operation-based approach, applications must be capable of providing the synchronizer with a log of the operations performed on each replica. The behavior of the synchronizer, then, is to merge logs so that, after quiescence, when each replica applies its merged log to the last synchronized state, all the replicas share a uniform state. There are, broadly speaking, three alternatives when merging logs: (1) reorder operations on all replicas to achieve an identical schedule (c.f. Bayou [9]), (2) partially reorder operations, exploiting semantic knowledge to leave equivalent sequences unordered (c.f. IceCube [10]) or (3) perform no reordering, but transform the operations themselves, so that the different schedules on each different replica all have a uniform result (c.f. Molli et al [11]). Although we favor the state-based approach, because it facilitates loose coupling, there are some advantages to working with operations instead. State-based architectures have less information available at synchronization time; they cannot exploit the knowledge of temporal sequencing available in operation logs. When operation logs are available, systems can sometimes determine that two modifications are not in conflict if one is in the operation history of the other. Further, in an operation-based system, the designer can often choose operations that encode high-level application semantics. The synchronizer then manipulates operations that are close to the actual user operations. Treating user-level operations as primitives makes it likely from the perspective of the user that, even under conflict, the state computed by the synchronizer will be well-formed.

Other state-based synchronizers include FCDP [4] and most file system synchronizers. FCDP is a generic, state-based reconciler parameterized by ad-hoc translations from heterogeneous concrete representations to XML and back again. Broadly speaking, FCDP can be considered an instance of the Harmony architecture—but without the formal underpinnings. FCDP is less generic (our lens language is easier to extend to new applications), but it is better able to deal with certain edits to documents than Harmony. FCDP achieves this by fixing a specific semantics for ordered lists that is particularly suited for document editing. File system synchronizers (such as [12–17]) and PDA synchronizers (such as Palm’s HotSync), are generally state-based but not generic. An interesting exception is DARCS [18], a hybrid state-/operation-based revision control system built on a “theory of patches.”

On the other side, IceCube [10,19] is a generic, operation-based reconciler that is parameterized over an algebra of operations specific to the application data

being synchronized and by a set of syntactic/static and semantic/dynamic ordering constraints on these operations. Molli et al [20,21,11], have also implemented a generic operation-based reconciler, using the technique of operational transformation. Their synchronizer is parameterized on transformation functions for all operations, which must obey certain conditions, and also obeys a formal specification.

It is worth noting that the distinction between state-based and operation-based synchronizers is not black and white: various hybrids are possible. For example, we can build a state-based system with an operation-based core by comparing previous and current states to obtain a hypothetical (typically, minimal) sequence of operations. But this involves complex heuristics, which can conflict with our goal of presenting predictable behavior to the user. Similarly, some loosely-coupled systems can build an operation-based system with a state-transfer core by using an operation log in order to determine what part of the state to transfer. Bengal [22] is an interesting example of a hybrid approach; it is a loosely-coupled synchronizer that uses operation logs as an optimization to avoid scanning the entire replica during update detection. Bengal supports updates to a wide range of databases via OLE/COM hooks, but is not heterogeneous because reconciliation may only occur between replicas of the same database.

**Conflicts and Convergence** Harmony, unlike many reconcilers, does not guarantee convergence in the case of conflicts. Of course, a successful run of any synchronizer *aims* to converge; that is, all of the replicas in the system should eventually reach a uniform state. In the case of conflicts, reconcilers can choose one of three broad strategies.

- They can settle all conflicts by fiat, for example by choosing the variant with the latest modification time and discarding the other.
- They can converge without resolving the conflict. In other words, they can keep enough information to record *both* conflicting updates, and converge to a single state containing the union of the conflicting updates.
- They can choose to diverge. They can maintain the conflicting updates locally, only, and not converge until the conflicts are (manually) resolved.

Most modern reconcilers will not simply discard updates (they obey a “no lost updates” policy). Harmony chooses the third option (divergence) over the second (unconditional convergence).

Systems such as Ficus [23], Rumor [14], Clique [15], Bengal [22], and TAL/S5 [20,21,11] converge by making additional copies of primitive objects that conflict and renaming one of the copies. Additionally, conflicts in Bengal may be handled by user-programmed *conflict resolvers*. CVS embeds markers in the bodies of files where conflicts occurred. By contrast, systems such

as Harmony and IceCube [10] do not reconcile objects affected by conflicting updates. Systems that allow reconciliation to end with divergent replicas have a further choice. They must choose whether to leave the replicas completely untouched by reconciliation, or to try to achieve *partial convergence*. Harmony aims for partial convergence; in Section 6 we showed that Harmony is a *maximal synchronizer*, propagating as many safe changes as possible.

Like Harmony, the synchronizer of Molli et al [20,21,11] uses formal specifications to ensure safety, but unlike Harmony it chooses convergence over persistence of user changes. As a result, Molli’s synchronizer is satisfied with recording multiple conflicting versions in the reconciled replicas, and its specification is limited to describing the correctness of its transformation functions.

Operational transforms resolve conflicting schedules by transforming local operations to undo the local operation, performing the remote operation, and finally redoing the local operation. Understanding the correct behavior of “undo” in a collaborative environment is an additional prerequisite to the correct behavior of operational transformation. Munson and Dewan [24] note that group “undo” may remove the need for a merge capability in optimistic replication. Prakash and Knister [25] provide formal properties that individual primitive operations in a system must satisfy in order to be “undo”able in a groupware setting. Abowd and Dix [26] formally describe the *desired* behavior of undo (and hence of conflict resolution) in “groupware”, and identify cases in which undo is fundamentally ambiguous. In such ambiguous cases—even if the primitive operations are defined to have unique undo functions—the user’s intention cannot be preserved and it is preferable to report conflict than to lose a user’s modification.

**Heterogeneity** Answering queries from heterogeneous data sources is a well-studied area in the context of data integration [27–30]. If we consider the problem of augmenting a data integration system with view update (another well-studied area—see [1] for a survey), then the result can be used to implement an optimistic replication system that reconciles conflicts between heterogeneous data sources. However, to the best of our knowledge, no other generic synchronizers support reconciliation over truly heterogeneous replicas. FCDP [4] is designed to be generic, but the genericity is limited to using XML as the internal representation, and currently only reconciles documents. Some file synchronizers do support diversity in small ways. For example, file synchronizers often grapple with different representations of file names and properties when reconciling between two different system types. Some map between length-limited and/or case insensitive names and their less restrictive counterparts (c.f. [16,15]). Others map complex file attributes (e.g. the Macintosh resource fork) into directories, rather than files, on the remote replicas.

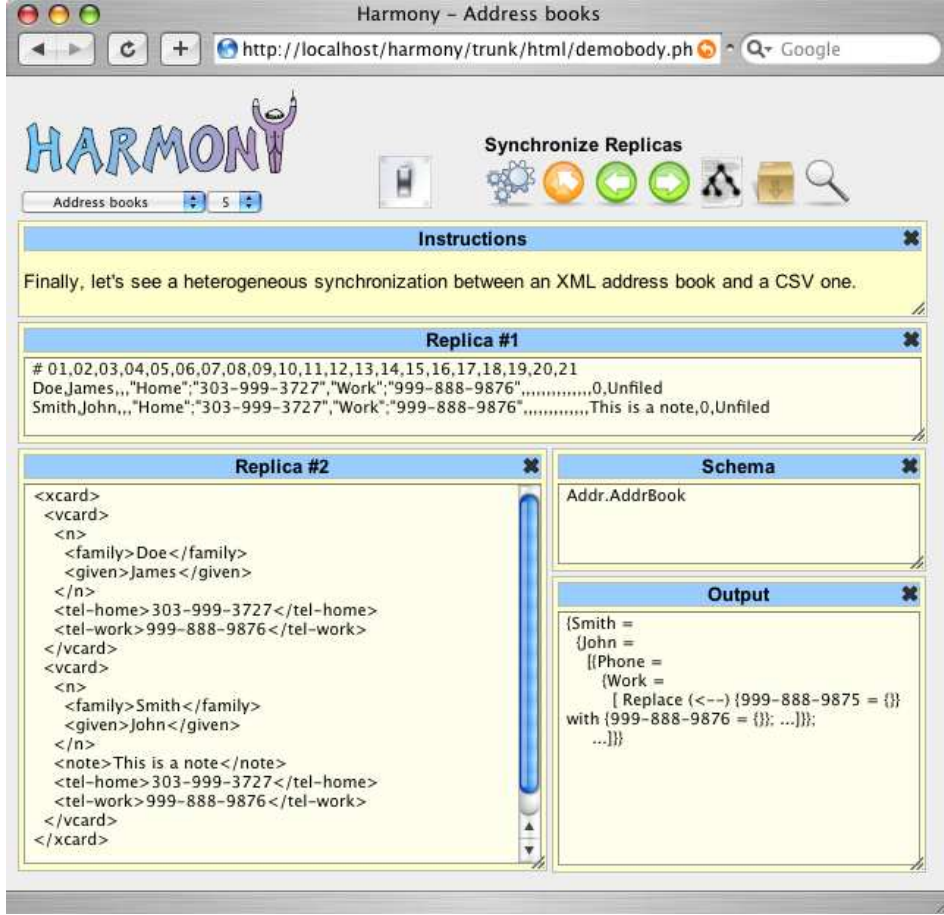


Fig. 4. Web demo screenshot

**Alignment** Harmony's emphasis on schema-based pre-alignment is influenced by examples we have found in the context of data integration where heterogeneity is a primary concern. Alignment, in the form of schema-mapping, has been frequently used to good effect (c.f. [31–35]). The goal of alignment, there, is to construct views over heterogeneous data, much as we transform concrete views into abstract views with a shared schema to make alignment trivial for the reconciler.

Some synchronizers differ mainly in their treatment of alignment. For example, the main difference between Unison [16,6] (which has almost trivial alignment) and CVS, is the comparative alignment strategy (based on the standard Unix tool `diff3`) used by CVS. At this stage, Harmony's core synchronization algorithm is deliberately simplistic, particularly with respect to ordered data. As we develop an understanding of how to integrate more sophisticated alignment algorithms in a generic and principled way, we hope to incorporate them into Harmony. Of particular interest are `diff3` and its XML based descendants: Lindholm's 3DM [2], the work of Chawathe et al [3], and FCDP [4].

## 9 Conclusions and Future Work

We have described the design of Harmony, a generic synchronization framework a tree-structured data based around the idea of *schema-directed synchronization*. A prototype implementation is available for download from the Harmony web site, <http://www.cis.upenn.edu/~bcpierce/harmony>. The web site also includes a live demo that allows experimenting with synchronization of small tree structures; Figure 4 shows a screenshot of the live demo, operating on heterogeneous address book formats. (The syntax for trees and schemas used in our implementation is actually a slight variant of the stylized version used in this paper; the precise definition is given in the programmer’s manual, which is available from the Harmony web site.)

The prototype has been used to construct specific synchronizers for a number of data formats, including several calendar formats (Palm Datebook, Unix ical, and iCalendar), bookmarks (Mozilla, Safari, and Internet Explorer), address books, application preference files, drawings, and bibliographic databases. Some of these instances are now in daily use by a small number of people.

One ongoing project is combining the core features of Harmony with a more sophisticated treatment of ordered structures. (More speculatively, although the Harmony framework has been designed with tree synchronization in mind, it may be generalizable to richer structures such as DAGs.) Another ongoing effort in our group aims to apply the ideas in Harmony (both lenses and synchronization) to the domain of relational data. Some preliminary results are described in [36].

## Acknowledgements

The Harmony project was begun in collaboration with Zhe Yang; Zhe contributed numerous insights whose genetic material can be found (generally in much-recombined form) in this paper. Trevor Jim provided the initial push to start the project by observing that the next step beyond the Unison file synchronizer (of which Trevor was a co-designer) would be synchronizing XML. Jonathan Moore, Owen Gunden, Malo Denielou, and Stéphane Lescuyer have collaborated with us on many aspects of Harmony’s design and implementation. Ongoing work with Sanjeev Khanna and Keshav Kunal on extensions of Harmony’s synchronization algorithm has deepened our understanding of the core mechanisms presented here. Conversations with Martin Hofmann, Zack Ives, Nitin Khandelwal, William Lovas, Kate Moore, Cyrus Najmabadi, Stephen Tse, Steve Zdancewic, and comments from the anonymous referees helped sharpen our ideas.

## References

- [1] J. N. Foster, M. B. Greenwald, J. T. Moore, B. C. Pierce, A. Schmitt, Combinators for bi-directional tree transformations: A linguistic approach to the view update problem, in: ACM SIGPLAN–SIGACT Symposium on Principles of Programming Languages (POPL), Long Beach, California, 2005, pp. 233–246.
- [2] T. Lindholm, XML three-way merge as a reconciliation engine for mobile data, in: ACM Workshop on Data Engineering for Wireless and Mobile Access (MobiDE), San Diego, California, 2003, pp. 93–97.
- [3] S. S. Chawathe, A. Rajamaran, H. Garcia-Molina, J. Widom, Change detection in hierarchically structured information, in: ACM SIGMOD Symposium on Management of Data (SIGMOD), Montreal, Quebec, 1996, pp. 493–504.
- [4] M. Lanham, A. Kang, J. Hammer, A. Helal, J. Wilson, Format-independent change detection and propagation in support of mobile computing, in: Brazilian Symposium on Databases (SBBD), Gramado, Brazil, 2002, pp. 27–41.
- [5] F. Bancilhon, N. Spyrtos, Update semantics of relational views, ACM Transactions on Database Systems 6 (4) (1981) 557–575.
- [6] B. C. Pierce, J. Vouillon, What’s in Unison? A formal specification and reference implementation of a file synchronizer, Tech. Rep. MS-CIS-03-36, Dept. of Computer and Information Science, University of Pennsylvania (2004).
- [7] T. Howes, M. Smith, F. Dawson, RFC 2425: A MIME content-type for directory information (Sep. 1998).
- [8] Y. Saito, M. Shapiro, Replication: Optimistic approaches, Technical Report HPL-2002-33, HP Laboratories Palo Alto (Feb. 8 2002).
- [9] W. K. Edwards, E. D. Mynatt, K. Petersen, M. J. Spreitzer, D. B. Terry, M. M. Theimer, Designing and implementing asynchronous collaborative applications with Bayou, in: ACM Symposium on User Interface Software and Technology (UIST), Banff, Alberta, 1997, pp. 119–128.
- [10] A.-M. Kermarrec, A. Rowstron, M. Shapiro, P. Druschel, The IceCube approach to the reconciliation of diverging replicas, in: ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing (PODC), Newport, Rhode Island, 2001, pp. 210–218.
- [11] P. Molli, G. Oster, H. Skaf-Molli, A. Imine, Using the transformational approach to build a safe and generic data synchronizer, in: ACM SIGGROUP Conference on Supporting Group Work (GROUP), Sanibel Island, Florida, 2003, pp. 212–220.
- [12] M. Satyanarayanan, J. J. Kistler, P. Kumar, M. E. Okasaki, E. H. Siegel, D. C. Steere, Coda: A highly available file system for a distributed workstation environment, IEEE Transactions on Computers C-39 (4) (1990) 447–459.

- [13] T. W. P. Jr., R. G. Guy, J. S. Heidemann, D. Ratner, P. L. Reiher, A. Goel, G. H. Kuenning, G. J. Popek, Perspectives on optimistically replicated, peer-to-peer filing., *Softw., Pract. Exper.* 28 (2) (1998) 155–180.
- [14] R. G. Guy, P. L. Reiher, D. Ratner, M. Gunter, W. Ma, G. J. Popek, Rumor: Mobile data access through optimistic peer-to-peer replication, in: *Proceedings of the ER Workshop on Mobile Data Access*, 1998, pp. 254–265.
- [15] B. Richard, D. M. Nioclais, D. Chalon, Clique: a transparent, peer-to-peer collaborative file sharing system, in: *International Conference on Mobile Data Management (MDM)*, Melbourne, Australia, 2003.
- [16] S. Balasubramaniam, B. C. Pierce, What is a file synchronizer?, in: *Fourth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom '98)*, 1998, full version available as Indiana University CSCI technical report #507, April 1998.
- [17] N. Ramsey, E. Csirmaz, An algebraic approach to file synchronization, in: *Proceedings of the 8th European Software Engineering Conference*, ACM Press, 2001, pp. 175–185.
- [18] D. Roundy, The DARCS system, <http://abridgegame.org/darcs/> (2004).
- [19] N. Preguia, M. Shapiro, C. Matheson, Efficient semantics-aware reconciliation for optimistic write sharing, *Technical Report MSR-TR-2002-52*, Microsoft Research (May 2002).
- [20] P. Molli, G. Oster, H. Skaf-Molli, A. Imine, Safe generic data synchronizer, *Rapport de recherche*, LORIA France (May 2003).
- [21] A. Imine, P. Molli, G. Oster, M. Rusinowitch, Proving correctness of transformation functions in real-time groupware, in: *ACM Conference on Computer Supported Cooperative Work (CSCW)*, Helsinki, Finland, 2003.
- [22] T. Ekenstam, C. Matheny, P. L. Reiher, G. J. Popek, The Bengal database replication system, *Distributed and Parallel Databases* 9 (3) (2001) 187–210.
- [23] P. L. Reiher, J. S. Heidemann, D. Ratner, G. Skinner, G. J. Popek, Resolving file conflicts in the ficus file system, in: *USENIX Summer Conference Proceedings*, 1994, pp. 183–195.
- [24] J. P. Munson, P. Dewan, A flexible object merging framework, in: *ACM Conference on Computer Supported Cooperative Work (CSCW)*, Chapel Hill, North Carolina, 1994, pp. 231–242.
- [25] A. Prakash, M. J. Knister, A framework for undoing actions in collaborative systems, *ACM Transactions on Computer-Human Interaction* 1 (4) (1994) 295–330.
- [26] G. D. Abowd, A. J. Dix, Giving undo attention, *Interacting with Computers* 4 (3) (1992) 317–342.

- [27] D. Florescu, A. Y. Levy, A. O. Mendelzon, Database techniques for the world-wide web: A survey, *ACM SIGMOD Record* 27 (3) (1998) 59–74.
- [28] S. Abiteboul, Querying semi-structured data, in: *International Conference on Database Theory (ICDT)*, Delphi, Greece, 1997, pp. 1–18.
- [29] A. Y. Halevy, Theory of answering queries using views, *ACM SIGMOD Record* 29 (4) (2000) 40–47.
- [30] I. Tatarinov, Z. G. Ives, A. Y. Halevy, D. S. Weld, Updating XML, in: *ACM SIGMOD Symposium on Management of Data (SIGMOD)*, Santa Barbara, California, 2001.
- [31] E. Rahm, P. A. Bernstein, A survey of approaches to automatic schema matching, *VLDB Journal* 10 (4) (2001) 334–350.
- [32] T. Milo, S. Zohar, Using schema matching to simplify heterogeneous data translation, in: *International Conference on Very Large Data Bases (VLDB)*, New York, New York, 1998.
- [33] C. Beeri, T. Milo, Schemas for integration and translation of structured and semi-structured data, in: *International Conference on Database Theory (ICDT)*, Jerusalem, Israel, 1999.
- [34] A. Doan, P. Domingos, A. Y. Halevy, Reconciling schemas of disparate data sources: A machine-learning approach, in: *ACM SIGMOD Symposium on Management of Data (SIGMOD)*, Santa Barbara, California, 2001.
- [35] J. Madhavan, P. A. Bernstein, E. Rahm, Generic schema matching with Cupid, in: *International Conference on Very Large Data Bases (VLDB)*, Roma, Italy, 2001, pp. 49–58.
- [36] A. Bohannon, J. A. Vaughan, B. C. Pierce, Relational lenses: A language for updateable views, Tech. Rep. MS-CIS-05-27, Dept. of Computer and Information Science, University of Pennsylvania (Dec. 2005).

## A Other Encodings

Here we record some additional encodings that we have developed for tuples, relations, and XML structures.

**Tuples** We can encode fixed-width tuples as record-like structures whose labels are in the set  $1 \dots n$ . Hence the tuple  $(t_1, t_2, t_3)$  may be encoded as the tree  $\{1 \mapsto t_1; 2 \mapsto t_2; 3 \mapsto t_3\}$ . The schema follows straightforwardly from this encoding; for the above triple, it is  $1[\mathbf{T}_1], 2[\mathbf{T}_2], 3[\mathbf{T}_3]$ . Since each tree in a tuple schema has every child in  $1 \dots n$ , the synchronizer aligns elements at the same positions in the tuple and never produces a schema domain conflict at the root.

**Relations** Relations, and other data structures consisting of a set of structured elements, present interesting alignment challenges. Considering a relation as a set of tuples, one needs to identify which tuple of the archive,  $o$ , should be associated to which tuple of each replica,  $a$  and  $b$ . Consider for instance the following relations:

	$o =$	Pat	333-4444		
		Chris	888-9999		
$a =$		Pat	111-2222		
		Chris	888-9999		
				$b =$	
				Pat	123-4567
				Jo	888-9999

If one chooses a schema where tuples are unstructured values, encoding  $(\text{Pat}, 333-4444)$  as a single string such as  $\text{Pat}:333-4444$  (using  $:$  as a separator so that it is easy to extract the components later),

$$o = \{\text{Pat}:333-4444 \text{ Chris}:888-9999\}$$

$$a = \{\text{Pat}:111-2222 \text{ Chris}:888-9999\}$$

$$b = \{\text{Pat}:123-4567 \text{ Jo}:888-9999\}$$

then the result of synchronization, using the schema for sets,  $*[\{\} ]$ , is the tree

$$\{\text{Pat}:111-2222, \text{Pat}:123-4567, \text{Jo}:888-9999.\}$$

The surprising duplication of the entry for Pat results from the fact that the tuple was considered unstructured, hence  $\text{Pat}:111-2222$  and  $\text{Pat}:123-4567$  are treated as two independent, non-conflicting additions.

A better way to encode this relation is to choose one attribute that is a key and encode a relation as a tree whose children are the key values pointing

$K$	$A$	$B$
$k_1$	$a_1$	$b_1$
$k_2$	$a_2$	$b_2$

can be represented as:

$$\left\{ \begin{array}{l} k_1 \mapsto \left\{ \begin{array}{l} A \mapsto a_1 \\ B \mapsto b_1 \end{array} \right\} \\ k_2 \mapsto \left\{ \begin{array}{l} A \mapsto a_2 \\ B \mapsto b_2 \end{array} \right\} \end{array} \right\}$$

Fig. A.1. Representing relations as trees

to a record containing the other attributes, as depicted in Figure A.1. This has the effect of aligning identical parts by keys, which can be chosen in an application-appropriate way.

A satisfying representation for our simple example would thus be (e.g., for  $o$ ):

$$\left\{ \begin{array}{l} \text{Pat} \mapsto \{ \text{Phone} \mapsto \{ 333-4444 \} \} \\ \text{Chris} \mapsto \{ \text{Phone} \mapsto \{ 888-9999 \} \} \end{array} \right\}$$

This representation of relations ensures that the correct components of each replica will be identified and aligned during synchronization. The schema that we use to synchronize this relation,  $*[\text{Phone}[V]]$ , which denotes bushes of keyed trees, each pointing to a single **Phone** edge that points to a value, implies that the result of synchronization is a proper encoding of an entry in the relation. In particular, it ensures that no entry contains two values under the **Phone** edge.

For this example, synchronizing  $o$ ,  $a$ , and  $b$  successfully removes **Chris**, adds **Jo**, and leaves **Pat** with a deeper schema domain conflict, because adding both 111-2222 and 123-4567 produces a tree that is not a value in  $![\{\}]$ .

**XML** Building on either the cons cell or keyed encoding for lists, it is easy to find an encoding for XML data. The XML element

```
<tag attr1="val1" ... attrm="valm">
  subelt1 ... subeltn
</tag>
```

is encoded into a tree of this form:

$$\left\{ \text{tag} \mapsto \left\{ \begin{array}{l} \text{attr1} \mapsto \text{val1} \\ \dots \\ \text{attrm} \mapsto \text{valm} \\ * \text{subelts} \mapsto \left[ \begin{array}{l} \text{subelt1} \\ \dots \\ \text{subeltn} \end{array} \right] \end{array} \right\} \right\}$$

The sub-elements **subelt1** to **subeltn** are placed in a *list* under a distinguished child named **\*subelts**, preserving their ordering. Attributes are en-

o1.txt	o2.txt	a.txt	b.txt
GOOD VERSION: Long ago, in a galaxy far, far away... OLD VERSION: Long ago, in a galaxy very far away... The End	GOOD VERSION: Long ago, in a galaxy far, far away... OLD VERSION: Long ago, in a galaxy very far away...	Long ago, in a galaxy far, far away...	Long ago, in a galaxy far away...

Fig. A.2. Files for `diff3` Example.

coded as unordered children, reflecting their treatment in XML. A leaf of an XML document—a “parsed character data” element containing a text string `str`—is converted to a tree of the form `{PCDATA -> str}`.

The reader may wonder why, since our goal is to handle XML data, we did not use a form of trees closer to XML’s data model in the other sections of the paper, avoiding the need for such complex encodings. Indeed, an early version of Harmony took *ordered* trees as primitive, allowing us to bypass this encoding. However, we discovered that this extra structure greatly increased the complexity of formalizing and implementing our lens programming language, FOCAL. On balance, the overall complexity of the system was minimized by making the core data structure as simple as possible and doing some extra programming *in* FOCAL to deal with XML structures in encoded form.

## B Alignment in `diff3`

The standard UNIX utility `diff3` identifies differences in a pair of replicas generated from a common ancestor. Unlike Harmony’s synchronization algorithm, which uses local information to align data, `diff3` (and related tools) align data by minimizing the “edit distance” of the replicas. This global heuristic can also lead to counterintuitive results on some inputs, as the following example shows.

Suppose that a pair of writers are working on a common draft of a movie script, as shown in the file `o1.txt` in Figure A.2. George Lucas and a staff writer independently edit the file. They both decide that the “good version” is better than the “old version”, so they delete the “OLD VERSION” and the text that follows. The staff writer creates the file `a.txt`, but Lucas goes

even further and removes the line containing the extra “far”, resulting in file `b.txt`.

Merging these files using `diff3` yields `b.txt`. The algorithm aligns the data in `a.txt` and `b.txt` in the obvious way—matching up the start of each file. Since all the updates (the deletions) to `o1.txt` in `a.txt` are a subset of the deletions in `b.txt`, the changes can be merged without conflict.

If, however, we started with `o2.txt`, which was almost identical to `o1.txt` but did not include the last line (“The End”), then merging produces a conflict. Rather than returning `b.txt`, `diff3` produces:

```
Long ago,  
in a galaxy  
<<<<<< a.txt  
far,  
far away...  
=====  
far away...  
>>>>>> b.txt
```

(The precise conflict differs slightly on different versions of `diff3`. The version shown ran on Linux, while `diff3` on Solaris does not include “far away...” as part of the conflict.)

Thus, in the global alignment scheme utilized in `diff3`, a modification at the end of one input affects the alignment of the entire file. Here, a non-conflicting run is turned into a conflicting one. It is not that one version is right and one version is wrong. The point is that it is extremely unlikely that a typical user could have predicted the behavior of `diff3` in these two situations.