

1 Authenticated Data Structures as Functors in 2 Isabelle/HOL

3 **Andreas Lochbihler** 

4 Digital Asset, Zurich, Switzerland

5 andreas.lochbihler@digitalasset.com, mail@andreas-lochbihler.de

6 **Ognjen Marić**

7 Digital Asset, Zurich, Switzerland

8 ognjen.marić@digitalasset.com

9 — Abstract —

10 Authenticated data structures allow several systems to convince each other that they are referring
11 to the same data structure, even if each of them knows only a part of the data structure. Using
12 inclusion proofs, knowledgeable systems can selectively share their knowledge with other systems
13 and the latter can verify the authenticity of what is being shared.

14 In this paper, we show how to modularly define authenticated data structures, their inclusion
15 proofs, and operations thereon as datatypes in Isabelle/HOL, using a shallow embedding. Modularity
16 allows us to construct complicated trees from reusable building blocks, which we call Merkle functors.
17 Merkle functors include sums, products, and function spaces and are closed under composition and
18 least fixpoints.

19 As a practical application, we model the hierarchical transactions of Canton, a practical inter-
20 operability protocol for distributed ledgers, as authenticated data structures. This is a first step
21 towards formalizing the Canton protocol and verifying its integrity and security guarantees.

22 **2012 ACM Subject Classification** Theory of computation → Logic and verification; Theory of
23 computation → Higher order logic; Theory of computation → Cryptographic primitives

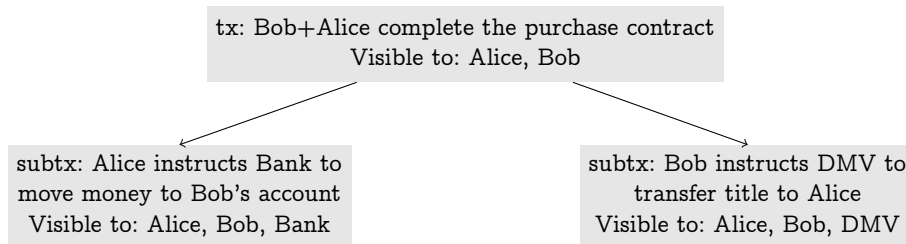
24 **Keywords and phrases** Merkle tree, functor, distributed ledger, datatypes, higher-order logic

25 **Supplement Material** The formalization is available in the Archive of Formal Proofs [15].

26 **1** Introduction

27 An authenticated data structure (ADS) allows several systems to use succinct digests to
28 convince each other that they are referring to the same data structure, even if each of
29 them knows only a part of the data structure. This has two main benefits. First, it saves
30 storage and bandwidth, as the systems only have to store parts of the entire structure that
31 they are interested in, and exchange just digests instead of the whole structure. This has
32 been exploited for a wide range of applications, e.g., logs in Certificate Transparency and
33 the blockchain structure and lightweight clients in Bitcoin. Second, ADSs allow parts of
34 the structure to be kept confidential to a subset of the systems involved in processing the
35 structure. For example, distributed ledger technology (DLT) promises to keep multiple
36 organizations synchronized about the state of their joint business workflows. Synchronization
37 requires transactions, i.e., atomic changes to the shared state. Yet organizations often do
38 not want to share all the changes with all involved parties. Some DLT protocols such
39 as the Canton interoperability protocol [6] and Corda [7] leverage ADSs to provide both
40 transactionality and varying levels of confidentiality. The formalization of Canton was the
41 starting point for this work.

42 Merkle trees [17] are the prime example of an ADS. The original Merkle tree is a binary
43 tree with data at the leaves, where every node is assigned a hash (serving as the digest)
44 using a cryptographic hash function h : a leaf with data d has hash $h\ d$ and an inner node



■ **Figure 1** A hierarchical Canton transaction. DMV is the department of motor vehicles.

45 has the hash $h(h_l, h_r)$ where h_l and h_r denote the hashes of the two children. If the hash
 46 of the root is known to all systems, then one system can convince another that a certain
 47 leaf stores data d . If π is the path from the root to the leaf, the inclusion proof consists
 48 of the sibling hashes of the nodes on the path. Given such an inclusion proof, the other
 49 system can recompute the hashes of the nodes on the path and check that the result matches
 50 the common root hash. This shows that the leaf indeed stores the given data if the hash
 51 function is collision-resistant. Moreover, the other system learns only hashes (of hashes) of
 52 the other data in the tree. So if h is preimage-resistant, then the inclusion proof does not
 53 leak information about the rest of the tree, provided that the hashed data contains sufficient
 54 entropy. This idea generalizes to finite tree data structures in general [18].

55 In this work, we consider authenticated data structures, which generalize Merkle trees to
 56 arbitrary shape, and we show how to modularly define them as datatypes in Isabelle/HOL.
 57 Modularity allows us to construct complicated trees from small reusable building blocks, for
 58 which properties are easy to prove. To that end, we consider authenticated data structures
 59 as functors and equip them with appropriate operations and their specifications. We show
 60 that this class of functors includes sums, products, and function spaces, and is closed under
 61 composition and least fixpoints. Concrete functors are defined as algebraic datatypes using
 62 Isabelle/HOL's datatype package [1]. This shallow embedding makes it possible to use
 63 Isabelle's rich infrastructure for datatypes.

64 As a practical application, we define ADSs over the hierarchical transactions [3] in the
 65 Canton protocol. To see an example of such a transaction, suppose that Alice wants to sell a
 66 car title to Bob. Transactionality allows Alice and Bob to exchange the title and the money
 67 atomically, which reduces their counterparty risks. Figure 1 shows the corresponding Canton
 68 transaction¹ for exchanging the money and the title. The transaction is generated from
 69 a smart contract that implements the purchase agreement. Such smart contracts can be
 70 conveniently written in the functional programming language DAML [8], which is built on
 71 the same hierarchical transactions as Canton.

72 Canton's hierarchical transactions offer three advantages over conventional flat trans-
 73 actions found in other DLT solutions. First, complex transactions can be composed from
 74 smaller building blocks. In the example above, the atomic swap transaction composes two
 75 transactions: the money transfer and the title transfer. Second, if a participant is involved
 76 only in a subtransaction, then the participant learns the contents of just this subtransaction,
 77 but not of other parts. In the example, the Bank only sees the money transfer, but not
 78 what Alice bought; similarly, the department of motor vehicles (DMV) does not see the
 79 amount the car was sold for. This also improves scalability as everyone must process only

¹Here and elsewhere in the paper, we take significant liberties in the presentation of Canton and focus on parts relevant for the construction of ADSs and for reasoning about them.

80 the data they are involved in. Third, they include mandatory authorization checks, which
81 are enforced even in the presence of Byzantine parties. Authorization flows from top to
82 bottom to enable delegation.

83 This hierarchy, enriched with some additional data, is encoded in ADSs and the protocol
84 exchanges inclusion proofs for such trees. More details will be given throughout the paper.
85 For now, it suffices to summarize the resulting requirements on the formalization:

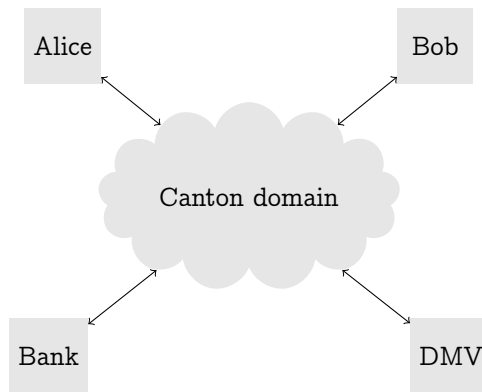
- 86 1. Hashes allow for checking whether two inclusion proofs refer to the same ADS. This
87 allows Canton to commit the example transaction atomically at all participants, even if
88 the Bank and the DMV see only a part of it.
- 89 2. Inclusion proofs allow us to prove inclusion for multiple leaves at the same time. Canton
90 sends such inclusion proofs to save bandwidth. Note that conventional inclusion proofs
91 are only for a single leaf.
- 92 3. Multiple inclusion proofs can be merged into one if they refer to the same ADS. This
93 is because Canton merges inclusion proofs only if they have the same set of recipients.
94 This reduces the load on the sender because it can multi-cast the same inclusion proof
95 to all recipients. Merging also simplifies the recipients' job: for example, Alice will
96 receive inclusion proofs for the entire transaction as well as both sub-transactions in
97 Figure 1. Merging them leaves her with just a single data structure representing the
98 entire transaction.

99 Our main contribution is a modular construction principle for ADSs as HOL datatypes,
100 i.e., functors. We also derive a variant of the ADSs that models inclusion proofs. To that end,
101 we introduce the class of Merkle functors, which are equipped with operations for hashing and
102 merging as required above. Our construction is modular in the sense that the class of Merkle
103 functors includes sums, products, and function spaces, and is closed under composition and
104 least fixpoints. Accordingly, the construction works for any inductive datatype (sums of
105 products and exponentials). Moreover, we show that the theory is suitable for constructing
106 concrete real-world instances such as Canton's transaction trees. The construction lives
107 in the symbolic models, i.e., we assume that no hash collisions occur. Our Isabelle/HOL
108 formalization is available in the Archive of Formal Proofs [15].

109 The rest of the paper is structured as follows. In Section 2, we describe our abstract
110 interface for ADSs. Section 3 shows how to construct such interfaces for tree-like structures
111 in a modular fashion. Section 4 demonstrates how to create inclusion proofs for general rose
112 trees and Canton transactions in particular. We discuss the related work in Section 5 and
113 conclude in Section 6.

114 **2** *Inclusion Proofs for Authenticated Data Structures*

115 We now present the operations and abstract interfaces for ADSs, motivated by their appli-
116 cation to Canton. Figure 2 shows a suitable a Canton-based deployment, where the Bank
117 and the DMV handle payments and car titles. The participants communicate with each
118 other using the Canton protocol. Unlike in most other DLT solutions, business data resides
119 with Canton primarily at the participants' nodes that share the data only on a need-to-know
120 basis [5]. Canton participants run a two-phase commit protocol to atomically update the
121 system state using transactions. The protocol is run over a Canton domain, which is operated
122 by a third party. The domain acts as the commit coordinator. While the participants may
123 be Byzantine, the domain is assumed to be honest-but-curious. That is, it is trusted to
124 correctly execute the protocol, but it should not learn the contents of a transaction (e.g.,



■ **Figure 2** Example topology of a Canton-based distributed ledger

125 how much Alice pays to Bob). Instead, it should only learn the minimal metadata that
 126 allows the protocol to tolerate Byzantine participants. Consequently, Canton sends business
 127 data through the domain only in encrypted or hashed form.

128 This motivates the *transaction tree* structure that Canton uses. The structure for the
 129 example transaction from Figure 1 is shown in Figure 3. Each (sub)-transaction of Figure 1
 130 is turned into a *view* in Figure 3, which consists of the view *data* and view *metadata*.
 131 For example, the node labeled by 1 in Figure 3 is the view corresponding to the top-level
 132 transaction in Figure 1. Its two children that are leaves contain the view's data and metadata.
 133 The metadata contains the information about who is affected by the view (here, Alice and
 134 Bob) and should therefore participate in the two-phase commit. The metadata is shared with
 135 Alice, Bob and the domain. The view data contains the confidential data with the actual
 136 state updates, and is shared only with Alice and Bob. This view also has two *subviews*, which
 137 correspond to the sub-transactions in Figure 1 as expected. A view can have an arbitrary
 138 number of subviews; the views labeled by 1.1 and 1.2 have no subviews, for example.

139 Additionally, the entire transaction is also equipped with metadata describing transaction-
 140 wide parameters, common to all views. Some of it is visible to all the involved participants,
 141 but not the domain, and some of it is visible to both the domain and the participants. The leaf
 142 children of the tree's root node store this metadata. Formally, the transaction tree can be mod-
 143 elled by the following datatypes, for some types *common-metadata*, *participant-metadata*,
 144 *view-metadata*, and *view-data* whose contents are not relevant for this paper.

145 **datatype** *view* =

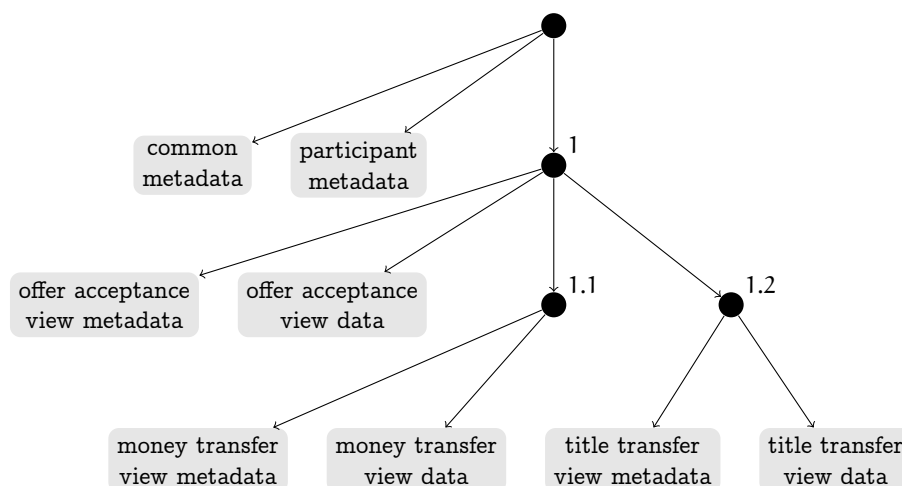
146 *View* \langle *view-metadata* \rangle \langle *view-data* \rangle (*subviews*: \langle *view list* \rangle)

147 **datatype** *transaction* =

148 *Transaction* \langle *common-metadata* \rangle \langle *participant-metadata* \rangle (*views*: \langle *view list* \rangle)

149 In Figure 3, the *Transaction* and *View* constructors become the inner nodes (black circles)
 150 and the data sits at the leaves (grey rectangles).

151 An ADS over this structure allows the participants and the domain use the root hash to
 152 refer to a transaction, and be sure that they are all referring to the same transaction tree.
 153 When constructing root hashes, it is useful to think of ADSs with multiple roots (i.e., forests)
 154 rather than just a single root like in a Merkle tree. For example, consider how the root node
 155 of a binary Merkle tree is constructed from two children. The two children themselves are
 156 Merkle trees, so we already have a forest of Merkle trees. More precisely, this forest has the
 157 shape of a pair. By adding the root node, we combine the whole forest into a larger Merkle



■ **Figure 3** Simplified Canton transaction tree for car title sale of Figure 1

158 tree. By the construction of Merkle trees, the new root hash is computed solely from the root
 159 hashes of the two child trees. Note that the concrete hash operation depends on the shape
 160 of the forest (a pair in this case). The new root is again a degenerate forest of a single tree
 161 with a single root hash. This view underlies our modular construction principle in Section 3.

162 In our construction, we will use the following conventions.

- 163 1. Raw data to be arranged in an ADSs is written as usual, e.g., $\langle a, 'a \text{ list} \rangle$.
- 164 2. Hashes and forests of hashes carry a subscript $_h$ as in $\langle a_h \rangle$. We leave hashes for now
 165 abstract as type variables and define them only in Section 3. Since the root hash identifies
 166 an ADS, we represent ADSs by their hashes.

167 Taking a root hash can make communication more efficient, but it is not enough for our
 168 purposes. For example, Bank does not know the contents of view 1.2 or even who is involved
 169 in view 1.2; the domain hides the latter. The views that are visible to a participant are
 170 called the participant's *projection* of the transaction. Canton aims to achieve the following
 171 integrity guarantee [3]: There exists a shared ledger that adheres to the underlying DAML
 172 smart contracts such that its projection to each honest participant consists exactly of the
 173 updates that have passed the participant's local checks. Achieving this guarantee for the
 174 Bank hinges on the Bank's ability to ensure that the view 1.1 is really included in the
 175 transaction tree. Thus, we also have to be able to prove that a substructure is included in a
 176 root hash.

177 Inclusion proofs are therefore the main workhorse in our formalization and the focus of
 178 this paper. We will denote the type of inclusion proofs over the source type with a subscript
 179 $_m$, e.g., $\langle a_m \rangle$, $\langle a_m, 'a_h \rangle$ *tree* $_m$. We need two operations on inclusion proofs:

- 180 1. Computing the (forest of) root hashes (which identifies the ADS to which the inclusion
 181 proof corresponds).
- 182 2. Merging two inclusion proof with the same root hash.

183 Accordingly, we introduce two type synonyms for these operations:

184 **type_synonym** $\langle a_m, 'a_h \rangle$ *hash* = $\langle a_m \Rightarrow 'a_h \rangle$

185 **type_synonym** a_m *merge* = $\langle a_m \Rightarrow 'a_m \Rightarrow 'a_m \text{ option} \rangle$

186 The merge operation returns *None* iff the inclusion proofs have different (forests of)
 187 root hashes. We require that merging is idempotent, commutative, and associative. The

188 locale *merkle-interface* below captures these properties. Associativity is expressed using
 189 the monadic ($\gg=$) on the *option* type. The merge operation makes inclusion proofs with
 190 the same hash into a semi-lattice. We fix the induced order as another parameter *bo* of the
 191 locale, where an inclusion proof is smaller than another if it reveals less. In that case, we say
 192 that the smaller is a *blinding* of the larger inclusion proof.

```

193 type_synonym 'am blinding-of = ⟨'am ⇒ 'am ⇒ bool⟩
194 locale merkle-interface =
195   fixes h :: ⟨('am, 'ah) hash⟩
196     and bo :: ⟨'am blinding-of⟩
197     and m :: ⟨'am merge⟩
198   assumes merge-respects-hashes: ⟨h a = h b ⟷ (∃ ab. m a b = Some ab)⟩
199     and idem: ⟨m a a = Some a⟩
200     and commute: ⟨m a b = m b a⟩
201     and assoc: ⟨m a b ⟹ m c = m b c ⟹ m a⟩
202     and bo-def: ⟨bo a b ⟷ m a b = Some b⟩

```

203 As expected for a semi-lattice, merging computes the least upper bound in the blinding
 204 relation:

```

205   (m a b = Some ab) = (bo a ab ∧ bo b ab ∧ (∀ u. bo a u ⟶ bo b u ⟶ bo ab u))

```

206 Also, the equivalence closure of the blinding relation gives the equivalence classes of the
 207 inclusion proofs under the hash function: $\text{equivclp } bo = \text{vimage2p } h \ h \ (=)$ where equivclp
 208 R denotes the equivalence closure of R and $\text{vimage2p } f \ g \ R = (\lambda x \ y. R (f \ x) (g \ y))$ the
 209 preimage of a relation under a pair of functions.

210 Our interface does not provide generic operations to build inclusion proofs for subtrees
 211 of tree-shaped data. This is because the construction depends on the exact shape of the tree.
 212 In Section 4, we will show how to create such proofs for the general shape of rose trees and
 213 Canton transactions in particular, using standard functional programming techniques.

214 **3** *Modularly Constructing Forests of Authenticated Data Structures*

215 In this section, we develop the theory to modularly construct ADSs and their inclusion
 216 proofs as HOL datatypes, including the operations for *merkle-interface*. We first introduce
 217 the concept of a blidable position (Section 3.1), which models a node in an ADS, and show
 218 how we obtain ADSs for Canton's transaction trees by introducing blidable positions in the
 219 right spots of the datatype definitions (Section 3.2).

220 The specification *merkle-interface* is not inductive and therefore not preserved by
 221 datatype constructions. We therefore generalize the specification and show that the general-
 222 ization is preserved under composition of functors and least fixpoints (Section 3.4). Finally,
 223 we show that sums, products and function spaces preserve the generalization (Section 3.5).

224 **3.1** *Blidable position*

225 A blidable position represents a node (inner node or leaf) in an ADS. Every node in an
 226 ADS comes with its root hash. In this work, we model such hashes symbolically. That is,
 227 we assume that no hash collisions occur, i.e., the hash function from values to the type of
 228 hashes is injective. We do not assume surjectivity though: some hashes do not correspond to
 229 any value. We model such values as garbage coming from a countable set (the naturals). A

countable set is large enough given that ADS are always finite in practice (since one cannot compute a hash of infinite amounts of data).

```
232 type_synonym garbage = ⟨nat⟩
233 datatype 'ah blindableh = Content ⟨'ah⟩ | Garbage ⟨garbage⟩
```

Since the hash function is injective, we can identify the values $'a$ with a subset of the hashes, namely those of form $Content _$. Accordingly, we could also have written $'a \text{ blindable}_h$ instead of $'a_h \text{ blindable}_h$. However, as an ADS contains hashes of hashes, it is more accurate to use $'a_h$ here.

For example, a degenerate Merkle tree with a single leaf, which stores some data x , has the root hash $Content \ x$. What does an inclusion proof for this tree look like? It can take two forms:

- 241 1. The inclusion proof proves inclusion of x , i.e., the leaf is not blinded. The inclusion proof thus contains x .
- 242 2. The inclusion proof does not prove inclusion of x , i.e., the leaf is blinded. So the inclusion proof contains only the hash of x .

In the second case, the recipients of such an inclusion proof cannot verify that the hash is meaningful (unless they already know the contents). So the hash could also be garbage. The following datatype formalizes these cases.

```
248 datatype ('am, 'ah) blindablem = Unblinded ⟨'am⟩ | Blinded ⟨'ah blindableh⟩
```

In general, inclusion proofs are nested, e.g., if a Merkle tree leaf contains another Merkle tree as data. We therefore use the inclusion proof type variable $'a_m$ instead of $'a$ for values, similar to $'a_h$ in $blindable_h$,

Note that our hashes are typed. Accordingly, the formalization cannot confuse hashes of ADSs that store *ints* in their leaves with hashes of ADSs that store some other data, say *string*. In the real world, this could happen as hashes are usually just bitstrings. However, for reasoning about inclusion proofs, the garbage hashes adequately model such confusion possibilities: If security best practices are followed, type flaw attacks lead to different hashes unless a hash collision occurs. So the hash of the *int* Leaf would be interpreted as garbage in the type of hashes for the ADS of *strings*. This is adequate for inclusion proofs because we care about the contents of a hash only if the position is unblinded, i.e., of shape $Content$.

Having introduced the types for blindable positions, we now define the corresponding operations and show that they satisfy the specification *merkle-interface*. The hash operation converts an inclusion proof into the root hash of the tree. We define it in two steps:

- 263 (i) *hash-blindable'* assumes that there are no nested inclusion proofs, i.e., $'a_m = 'a_h$.
- 264 (ii) *hash-blindable* generalizes *hash-blindable'* to nested inclusion proofs. It first converting nested inclusion proofs to their root hashes using the hash function that is given as a parameter. Here, *map-blindable_m* is the mapper generated by the `datatype` command.

```
267 primrec hash-blindable' :: ⟨⟨('ah, 'ah) blindablem, 'ah blindableh⟩ hash⟩ where
268   ⟨hash-blindable' (Unblinded x) = Content x⟩
269   | ⟨hash-blindable' (Blinded x) = x⟩
```

270
271 **definition** *hash-blindable*

```
272   :: ⟨('am, 'ah) hash ⇒ ⟨⟨('am, 'ah) blindablem, 'ah blindableh⟩ hash⟩ where
273   ⟨hash-blindable h = hash-blindable' ◦ map-blindablem h id⟩
```

Next, we define the blinding order *blinding-of-blindable*. Like *hash-blindable*, it is parametrized by the hash function and the blinding order for the nested inclusion proofs.

276 The first clause lifts the blinding order in case the inclusion proof unblinds the contents.
 277 The second clause, when the position on the left is blinded, checks that both positions have
 278 the same hash.

```
279 context fixes h :: ⟨('am, 'ah) hash⟩ and bo :: ⟨'am blinding-of⟩ begin
280 inductive blinding-of-blindable :: ⟨('am, 'ah) blindablem blinding-of⟩ where
281   ⟨blinding-of-blindable (Unblinded x) (Unblinded y)⟩ if ⟨bo x y⟩
282 | ⟨blinding-of-blindable (Blinded x) t⟩ if ⟨hash-blindable h t = x⟩
283 end
```

284 Merging of blindable positions works similarly. If both positions are unblinded, *merge-blindable*
 285 tries to merge the contents. If both are blinded, it succeeds iff the hashes are the same.
 286 Otherwise, it checks that the hashes are the same and, if so, returns the unblinded version.

```
287 context fixes h :: ⟨('am, 'ah) hash⟩ and m :: ⟨'am merge⟩ begin
288 fun merge-blindable :: ⟨('am, 'ah) blindablem merge⟩ where
289   ⟨merge-blindable (Unblinded x) (Unblinded y) = map-option Unblinded (m x y)⟩
290 | ⟨merge-blindable (Blinded t) (Blinded u) =
291   (if t = u then Some (Blinded u) else None)⟩
292 | ⟨merge-blindable (Blinded x) (Unblinded y) =
293   (if x = Content (h y) then Some (Unblinded y) else None)⟩
294 | ⟨merge-blindable (Unblinded y) (Blinded x) =
295   (if x = Content (h y) then Some (Unblinded y) else None)⟩
296 end
```

297 It is straightforward to show that these definitions preserve the specification *merkle-interface*.
 298 That is, if the operations for nested inclusion proofs satisfy *merkle-interface*, then so do the
 299 operations for *blindable_m*.

```
300 lemma merkle-blindable:
301   ⟨merkle-interface
302     (hash-blindable h)
303     (blinding-of-blindable h bo)
304     (merge-blindable h m)⟩
305   if ⟨merkle-interface h bo m⟩
```

306 3.2 Example: Canton transaction trees

307 We now illustrate how to use *blindable_h* and *blindable_m* to define the ADSs and
 308 inclusion proofs for the Canton transaction trees from Section 2. As shown in Figure 3, the
 309 transaction tree contains a node for the transaction tree as a whole, every view, and every
 310 leaf (*common-metadata*, *participant-metadata view-metadata*, and *view-data*). Yet, the
 311 datatype declarations do not contain the information what should become a separate node
 312 in the ADS. To make the construction systematic, we consider the blindable positions to be
 313 marked in the datatype with the type constructor *blindable*.

```
314 type_synonym 'a blindable = ⟨'a⟩
```

315 So we pretend in this section as if *views* and *transactions* were defined as follows:

```
316 datatype view = View
317   ⟨((view-metadata blindable × view-data blindable) × view list) blindable⟩
318 datatype transaction = Transaction
```



```

319  ⟨⟨(common-metadata blindable × participant-metadata blindable) × view list⟩
320  blindable⟩

```

To define the hashes and inclusion proofs, we simply replace each type constructor τ with its counterparts τ_h and τ_m . For views, this looks as follows. Here \times_h , \times_m , $list_h$, and $list_m$ are type synonyms for \times and $list$; Section 3.5 introduces them formally. We abuse notation by writing $view-metadata_h$ $view-metadata_m$ for the blindable position of $view-metadata$.

```

325 type_synonym view-metadata_h = ⟨view-metadata blindable_h⟩
326 type_synonym view-data_h = ⟨view-data blindable_h⟩
327 datatype view_h = View_h ⟨⟨(view-metadata_h ×_h view-data_h) ×_h view_h list_h⟩ blindable_h⟩
328 type_synonym view-metadata_m = ⟨(view-metadata, view-metadata) blindable_m⟩
329 type_synonym view-data_m = ⟨(view-data, view-data) blindable_m⟩
330 datatype view_m = View_m
331  ⟨⟨(view-metadata_m ×_m view-data_m) ×_m view_m list_m,
332  (view-metadata_h ×_h view-data_h) ×_h view_h list_h⟩ blindable_m⟩

```

These types nest hashes and inclusion proofs: A view node, e.g., nests hashes and inclusion proofs for the metadata, the data, and all the subviews. In particular, the $view_h$ and $view_m$ datatypes recurse through the $blindable_h$ and $blindable_m$ type constructors. This works because $blindable_h$ and $blindable_m$ are bounded natural functors (BNFs) [21]. In fact, this transformation works for any datatype declaration thanks to the compositionality of BNFs. The construction for transaction trees is accordingly:

```

339 type_synonym common-metadata_h = ⟨common-metadata blindable_h⟩
340 type_synonym common-metadata_m =
341  ⟨(common-metadata, common-metadata) blindable_m⟩
342 type_synonym participant-metadata_h = ⟨participant-metadata blindable_h⟩
343 type_synonym participant-metadata_m =
344  ⟨(participant-metadata, participant-metadata) blindable_m⟩
345 datatype transaction_h = Transaction_h
346  ⟨⟨(common-metadata_h ×_h participant-metadata_h) ×_h view_h list_h⟩ blindable_h⟩
347 datatype transaction_m = Transaction_m
348  ⟨⟨(common-metadata_m ×_m participant-metadata_m) ×_m view_m list_m,
349  (common-metadata_h ×_h participant-metadata_h) ×_h view_h list_h⟩ blindable_m⟩

```

3.3 Composition

Having defined the types of ADSs, we next must define the operations on ADSs and prove that they satisfy *merkle-interface*. Doing so directly is possible, but prohibitively cumbersome. Instead, we modularize the proofs following the structure of the types. We can derive preservation lemmas for all involved type constructors analogous to *merkle-blindable*.

The preservation lemmas are compositional by construction: if $'a_h \tau_h / 'a_m \tau_m$ and $'b_h \sigma_h / 'b_m \sigma_m$ preserve *merkle-interface*, then so does their composition $'a_h \tau_h \sigma_h / 'a_m \tau_m \sigma_m$. Moreover, every nullary functor also satisfies *merkle-interface* with the discrete ordering ($=$).

definition *merge-discrete* :: $\langle 'a \text{ merge} \rangle$ where

360 $\langle \text{merge-discrete } x \ y = (\text{if } x = y \text{ then } \text{Some } y \text{ else } \text{None}) \rangle$

361 **lemma** *merkle-discrete*: $\langle \text{merkle-interface } id \ (=) \ \text{merge-discrete} \rangle$

362 For *view-data*, for example, we compose the corresponding discrete functor with a blindable
363 position.

364 **abbreviation** *hash-view-data* :: $\langle (\text{view-data}_m, \text{view-data}_h) \ \text{hash} \rangle$ **where**

365 $\langle \text{hash-view-data} \equiv \text{hash-blindable } id \rangle$

366 **abbreviation** *blinding-of-view-data* :: $\langle \text{view-data}_m \ \text{blinding-of} \rangle$ **where**

367 $\langle \text{blinding-of-view-data} \equiv \text{blinding-of-blindable } id \ (=) \rangle$

368 **abbreviation** *merge-view-data* :: $\langle \text{view-data}_m \ \text{merge} \rangle$ **where**

369 $\langle \text{merge-view-data} \equiv \text{merge-blindable } id \ \text{merge-discrete} \rangle$

370

371 **lemma** *merkle-view-data*:

372 $\langle \text{merkle-interface } \text{hash-view-data} \ \text{blinding-of-view-data} \ \text{merge-view-data} \rangle$

373 **by** (*rule merkle-blindable*)(*rule merkle-discrete*)

374 If we do the same for *view-metadata* and consider the pair *view-metadata* \times *view-data*,
375 composition immediately gives us the following (the operations for products will be introduced
376 in Section 3.5).

377 **lemma** $\langle \text{merkle-interface}$

378 $\ (\text{hash-prod } \text{hash-view-metadata} \ \text{hash-view-data})$

379 $\ (\text{blinding-of-prod } \text{blinding-of-view-metadata} \ \text{blinding-of-view-data})$

380 $\ (\text{merge-prod } \text{merge-view-metadata} \ \text{merge-view-data}) \rangle$

381 3.4 Inductive generalization for least fixpoints

382 The *view* datatype is the least fixpoint of the functor

383 $\ 'a \ F = ((\text{view-metadata } \text{blindable} \times \text{view-data } \text{blindable}) \times 'a \ \text{list}) \ \text{blindable}$

384 and so are *view_h* and *view_m* of analogous functors F_h and F_m . Composition gives us a
385 preservation theorem for F , but we need more for least fixpoints.

386 In fact, *merkle-interface* is not inductive, so least fixpoints need not preserve it. The
387 problem is the following: In the inductive preservation proof, we get the induction hypoth-
388 esis only for smaller values. We therefore cannot use F 's preservation theorem because
389 *merkle-interface* requires the conditions to hold on *all* values, not just the smaller ones. So
390 we must generalize *merkle-interface* to make it inductive.

391 In our first attempt with a direct generalization, the proofs about the merge operation
392 turned out to be rather cumbersome. The associativity law in particular required many
393 case distinctions due to the *options*. We therefore present a different approach where the
394 focus is on the blinding relation and merge is merely characterized as the join. We abstractly
395 derive commutativity, idempotence, and associativity for merge once and for all from the
396 ordering properties and merge's characterization. This leads to simpler proofs where all case
397 distinctions dealt with by Isabelle's proof automation.

398 Our generalization splits *merkle-interface* into three locales (Figure 4):

- 399 1. The locale *blinding-respects-hashes* splits off the first assumption of *merkle-interface*.
400 No relativization is needed here because the (inductive) blinding order *bo* occurs only
401 once and in a negative position. The preservation proof can therefore use rule induction
402 rather than structural induction.

```

locale blinding-respects-hashes =
  fixes h :: ⟨('am, 'ah) hash⟩
  and bo :: ⟨'am blinding-of⟩
  assumes hash: ⟨bo ≤ vimage2p h h (=)⟩

locale blinding-of-on = blinding-respects-hashes ⟨h⟩ ⟨bo⟩ for A h bo +
  assumes refl: ⟨x ∈ A ⇒ bo x x⟩
  and trans: ⟨[[ bo x y; bo y z; x ∈ A ]] ⇒ bo x z⟩
  and antisym: ⟨[[ bo x y; bo y x; x ∈ A ]] ⇒ x = y⟩

locale merge-on = blinding-of-on ⟨UNIV⟩ ⟨h⟩ ⟨bo⟩ for A h bo m +
  assumes join: ⟨[[ h x = h y; x ∈ A ]]
    ⇒ ∃z. m x y = Some z ∧ bo x z ∧ bo y z ∧ (∀u. bo x u ⇒ bo y u ⇒ bo z u)⟩
  and undefined: ⟨[[ h x ≠ h y; x ∈ A ]] ⇒ m x y = None⟩

```

■ **Figure 4** Inductive generalization of *merkle-interface*

- 403 2. The locale *blinding-of-on* formalizes the order properties of the blinding relation *bo*
 404 (reflexivity, transitivity, antisymmetry). It fixes a set *A* in addition to the Merkle
 405 operations; the inductive proof for fixpoints instantiates *A* with the set of smaller terms
 406 for which the properties hold by the induction hypothesis. Accordingly, one of the
 407 variables in the properties is restricted to *A*. (Since the induction proof will be structural,
 408 it suffices to restrict one variable instead of all.) Unlike *hash* for *blinding-respects-hashes*,
 409 transitivity and antisymmetry cannot be shown by rule induction even though *bo* occurs
 410 as an assumption, because *bo* occurs multiple times, but rule induction acts only on
 411 one. Accordingly, *F*'s preservation theorem does not apply to the induction hypothesis
 412 because it assumes that all occurrences are the same.²
- 413 3. The locale *merge-on* augments *blinding-of-on* with the characterization for merge as the
 414 join. While *merge-on*'s assumptions are again restricted by *A*, the restriction is removed
 415 on the assumptions of the parent locale *blinding-of-on* by setting *A* to the type universe
 416 *UNIV*.
- 417 This change is crucial and the reason for introducing three locales: When we prove *join*
 418 for the least fixpoint, we can (and must) use that *bo* is an order *everywhere*. This is
 419 because *join* uses *bo* with many different arguments, in particular the result *z* of the
 420 merge. In a unified locale, we would have to prove that *z* stays within the set *A*, which
 421 incurred a lot more proof effort.

422 In the unrestricted case, *merge-on* is equivalent to *merkle-interface*:

423 **lemma** ⟨*merkle-interface* *h bo m* ↔ *merge-on UNIV h bo m*⟩

424 We are now ready to define the class of Merkle functors. For readability, we only spell
 425 out the case of unary functors. The generalization to *n*-ary functors is as expected.

426 ► **Definition 1** (Merkle functor). *A unary BNF* F_h *and binary BNF* F_m *constitute a*
 427 *unary Merkle functor if there exist operations* $\text{hash-}F' :: ((a_h, a_h) F_m, a_h F_h) \text{ hash}$

²Alternatively, we could have generalized the property such that different blinding relations are allowed. Preservation of transitivity becomes preservation of relation composition and antisymmetry transforms into preservation of intersections. For reflexivity, we would still have needed to the set *A* however.

428 and *blinding-of-F* :: ('a_m, 'a_h) hash ⇒ 'a_m blinding-of ⇒ ('a_m, 'a_h) F_m blinding-of
 429 and *merge-F* :: ('a_m, 'a_h) hash ⇒ 'a_m merge ⇒ ('a_m, 'a_h) F_m merge with the
 430 following properties

$$\begin{array}{l}
 \text{Monotonicity} \quad \frac{bo \leq bo'}{\text{blinding-of-F } h \text{ } bo \leq \text{blinding-of-F } h \text{ } bo'} \\
 \text{Congruence} \quad \frac{\forall a \in A. \forall b. m \ a \ b = m' \ a \ b}{\forall x \in \{y. \text{set}_1\text{-F}_m \ y \subseteq A\}. \forall b. \text{merge-F } h \ m \ x \ y = \text{merge-F } h \ m' \ x \ y} \\
 \text{Hashes} \quad \frac{\text{blinding-respects-hashes } h \ bo}{\text{blinding-respects-hashes } (\text{hash-F } h) \ (\text{blinding-of-F } h \ bo)} \\
 \text{Blinding order} \quad \frac{\text{blinding-of-on } A \ h \ bo}{\text{blinding-of-on } \{x. \text{set}_1\text{-F}_m \ x \subseteq A\} \ (\text{hash-F } h) \ (\text{blinding-of-F } h \ bo)} \\
 \text{Merge} \quad \frac{\text{merge-on } A \ h \ bo \ m}{\text{merge-on } \{x. \text{set}_1\text{-F}_m \ x \subseteq A\} \ (\text{hash-F } h) \ (\text{blinding-of-F } h \ bo) \ (\text{merge-F } h \ m)}
 \end{array}$$

432 where $\text{hash-F } h = \text{hash-F}' \circ \text{map-F}_m \ h \ \text{id}$.

433 Every Merkle functor preserves *merkle-interface*: set $A = \text{UNIV}$ in the merge property
 434 and use the above equivalence between *merkle-interface* and *merge-on*.

435 We are now ready to state and prove the main theoretical contribution of this paper.

436 ► **Theorem 2.** *Merkle functors of arbitrary arity are closed under composition and*
 437 *least fixpoints.*

438 **Proof.** Closure under composition is obvious from the shape of the properties and the fact
 439 that BNFs are closed under composition.

440 For closure under least fixpoints, we consider a functor F and its least fixpoint T through
 441 one of F 's arguments. say $\text{datatype } T = T \ "T \ F"$, and similarly for T_h and T_m . The
 442 operations are defined as follows, where we omit all Merkle operation parameters for type
 443 parameters that are not affected.

444 ■ The hash operation $\text{hash-T}'$ is defined by primitive recursion:

$$445 \quad \text{hash-T}' (T_m \ x) = T_h \ (\text{hash-F}' \ (\text{map-F}_m \ \text{hash-T}' \ x)).$$

446 ■ The blinding order blinding-of-T is defined inductively by the following rule:

$$447 \quad \frac{\text{blinding-of-F } \text{hash-T } \text{blinding-of-T } \ x \ y}{\text{blinding-of-T } (T_m \ x) \ (T_m \ y)}$$

448 Monotonicity ensures that blinding-of-T is well-defined.

449 ■ Merge merge-T is defined by well-founded recursion (over the subterm relation on T_m):

$$450 \quad \text{merge-T } (T_m \ x) \ (T_m \ y) = \text{map-option } T_m \ (\text{merge-F } \text{hash-T } \text{merge-T } \ x \ y)$$

451 Congruence ensures that merge-F calls merge-T recursively only on smaller arguments.
 452 We have not been able to define merge-T with primitive recursion, which allows pattern
 453 matches only on one argument, not two. Our attempts with `primrec` failed because the
 454 recursive call occurs under merge-F , which is not F_m 's mapper. The usual trick of using
 455 parametricity theorems to extract the recursive calls into map-F_m did not work because
 456 the parametricity theorem for merge-F is too weak. It is also not clear how it could be
 457 strengthened without excluding important examples of Merkle functors such as *blindable*.

458 Well-founded recursion works well, except that Isabelle has no automatic parametricity
 459 inference for well-founded recursion. We therefore manually proved the parametricity
 460 theorems that the transfer package needs.

461 Monotonicity and preservation of *blinding-respects-hashes* are proven by rule induction on
 462 *blinding-of-T*. Congruence, *blinding-of-on*, and *merge-on* are shown by structural induction
 463 on the argument that is constrained by *A*. ◀

464 It is not possible to formalize this theorem abstractly in Isabelle/HOL because it is not
 465 possible to abstract over type constructors. Instead, we have axiomatized a binary Merkle
 466 functor using the `bnf_axiomatization` command and carried out the construction and proofs
 467 for least fixpoints and composition. This approach is similar to how Blanchette et al. have
 468 formalized the theory of bounded natural functors [2]. The axiomatization also illustrates
 469 how the definition and proofs generalize to several functors with type arguments. Moreover,
 470 all the example ADS constructions in Section 3.6 merely adapt these proofs to the concrete
 471 functors at hand.

472 3.5 Concrete Merkle functors

473 We now present concrete Merkle functors. They show that the class of Merkle functors is
 474 sufficiently large to be of interest. In particular, it contains all inductive datatypes (least
 475 fixpoints of sums of products). We have formalized all of the following.

- 476 ■ The discrete functor from Section 3.3 with hash operation *id* and the discrete blinding
 477 order ($=$) is a nullary Merkle functor.
- 478 ■ Blindable positions $blindable_h$ and $blindable_m$ are a unary Merkle functor.³
- 479 ■ Sums and products are binary Merkle functors. We set $'a_h \times_h 'b_h = 'a_h \times 'b_h$ and
 480 $'a_m \times_m 'b_m = 'a_m \times 'b_m$ and similarly for $+_h$ and $+_m$. Formally, \times_m and $+_m$ should
 481 take four type arguments. However, as sums and products do not themselves contain
 482 blidable positions, the type arguments $'a_h$ and $'b_h$ are ignored and we therefore omit
 483 them. The hash operations *hash-prod* and *hash-sum* are the mappers *map-prod* and
 484 *map-sum*, respectively. The blinding orders *blinding-of-prod* and *blinding-of-sum* are
 485 the relators *rel-prod* and *rel-sum*. The merge operations are defined as follows:

486 $merge-prod\ ma\ mb\ (x,\ y)\ (x',\ y') =$
 487 $ma\ x\ x' \gg= (\lambda x''.\ map-option\ (Pair\ x'')\ (mb\ y\ y'))$

488 $merge-sum\ ma\ mb\ (Inl\ x)\ (Inl\ y) = map-option\ Inl\ (ma\ x\ y)$
 489 $merge-sum\ ma\ mb\ (Inr\ x)\ (Inr\ y) = map-option\ Inr\ (mb\ x\ y)$
 490 $merge-sum\ ma\ mb\ (Inr\ v)\ (Inl\ va) = None$
 491 $merge-sum\ ma\ mb\ (Inl\ va)\ (Inr\ v) = None$

- 492 ■ The function space $'a \Rightarrow 'b$ is a unary Merkle functor in the codomain. (Like for sums
 493 and products, $'a \Rightarrow_h 'b_h = 'a \Rightarrow 'b_h$ and $'a \Rightarrow_m 'b_m = 'a \Rightarrow 'b_m$ and we omit the
 494 ignored $'b_h$.) Hashing is function composition and the blinding order is pointwise. Merge
 495 is defined by

³The proof of transitivity preservation requires that the blinding order *bo* on $'a_m$ respects hashes everywhere, not only on *A*. This is the reason why we have split the locale *blinding-respects-hashes* from *blinding-of-on*.

```

496     merge-fun m f g =
497     (if  $\forall x. m (f x) (g x) \neq \text{None}$  then Some ( $\lambda x. \text{the } (m (f x) (g x))$ ) else None)

```

498 Proving the Merkle properties requires choice.

499 3.6 Case study: Merkle rose trees and Canton's transactions

500 Thanks to Theorem 2, all datatypes built from the Merkle functors in the previous section
501 are also Merkle functors. We now show the elegance and expressiveness of Merkle functors
502 using three datatypes: lists, rose trees and Canton transaction, where each builds on the
503 previous ones.

504 Lists are isomorphic to the datatype

```

505 datatype 'a list' = List' (unit + 'a × 'a list')

```

506 and therefore also a Merkle functor. We have carried out this construction as *list* occurs in
507 Canton transaction trees. Like sums, products, and function spaces, *lists* do not contain
508 blinding positions directly, so $\text{list}_h = \text{list}_m = \text{list}$. Hashing and the blinding order are the
509 mapper and the relator. Initially, we tried to prove *merkle-interface* for *lists* directly, but
510 the proofs about merge quickly got out of control. We therefore carried out the fixpoint
511 construction of Theorem 2 for *list'* and transferred the definitions and theorems to *list* using
512 the transfer package [13].

513 Rose trees are then given by the datatype

```

514 datatype 'a rose-tree = Tree (<'a × 'a rose-tree list> blinding)

```

515 Applying our construction, we obtain Merkle rose trees as

```

516 datatype 'ah rose-treeh = Treeh (<'ah ×h 'ah rose-treeh listh> blindingh)
517 datatype ('am, 'ah) rose-treem = Treem
518 (<'am ×m ('am, 'ah) rose-treem listm, 'ah ×h 'ah rose-treeh listh> blindingm)

```

519 with the corresponding operations and their properties.

520 From here, it is only a small step to transactions in Canton. Views are Merkle rose
521 trees where the data at the nodes is instantiated with the Merkle functor corresponding to
522 *view-metadata blinding* × *view-data blinding*. Then, transactions compose the Merkle
523 functor for *common-metadata blinding* × *participant-metadata blinding* × - *list* with
524 views. We have lifted our machinery from these raw Merkle functors to the datatypes *view_m*
525 and *transaction_m* using the lifting and transfer packages [13].

526 4 Creating Inclusion Proofs

527 So far, given a tree-like data type *t*, we showed how to systematically construct the
528 corresponding type of ADSs *t_h* and their inclusion proofs, *t_m*. To make use of this
529 construction in practice, we must also be able to create values of type *t_m* from values of type
530 *t*. As in the case of our composition and fixpoint theorem, HOL's lack of abstraction over
531 type constructors makes it impossible to express this process in HOL in its full generality.
532 Instead, we show how it works on rose trees, as these are the most general type of tree in
533 terms of branching. The construction can be easily adapted for other kinds of trees.

534 There are three basic operations:

- 535 ■ Hashing *hash-source-tree* returns the root hash for a source tree.

536 ■ Embedding *embed-source-tree* returns the inclusion proof that proves inclusion of the
537 whole tree.

538 ■ Fully blinding *blind-source-tree* returns the inclusion proof that proves no inclusion at
539 all.

540 Hashing and fully blinding conceptually do the same thing, but their return types ($'a_h$
541 *rose-tree_h* and ($'a_m, 'a_h$) *rose-tree_m*) differ. As rose trees are parameterized by their node
542 label type, hashing, embedding and fully blinding take parameters which hash or embed the
543 node labels. The expected properties hold: the embedded and fully blinded versions of the
544 same source tree have the same hash, namely the hashing of the source tree, and the former
545 is a blinding of the latter.

546 The more interesting operations concern creating an inclusion proof for a subtree of a
547 tree. For example, with Canton's hierarchical transactions, we would like to prove that a
548 subtransaction is really part of the entire transaction. Such a proof consists of the subtree
549 itself, together with a path connecting the tree's root to the subtree's root. As noticed by
550 Seefried [20], this corresponds to a zipper [12] focused on the subtree. This enables simple
551 manipulation of such proofs in a functional programming style, well-suited to HOL. The
552 zippers for rose trees are captured by the following types.

553 **type_synonym** $'a$ path-*elem* = $\langle 'a \times 'a$ rose-*tree list* $\times 'a$ rose-*tree list* \rangle

554 **type_synonym** $'a$ path = $\langle 'a$ path-*elem list* \rangle

555 **type_synonym** $'a$ zipper = $\langle 'a$ path $\times 'a$ rose-*tree* \rangle

556 Given a zipper that focuses on a node, we define the operations that turn rose trees into
557 zippers and vice-versa

558 $tree-of-zipper$ ($[], t$) = t

559 $tree-of-zipper$ ($(a, l, r) \cdot z, t$) = $tree-of-zipper$ ($z, Tree$ ($a, l @ t \cdot r$))

560 $zipper-of-tree$ $t \equiv ([], t)$

561 The zippers for inclusion proofs have the exact same shape, except that all the type
562 constructors are subscripted by $_m$ and have another type parameter capturing the type of
563 hashes (e.g., ($'a, 'a_h$) *zipper_m*). Like for source trees, we define operations that blind and
564 embed a path respectively, and define operations that convert between Merkle rose trees
565 and their zippers. As expected, given the same source zipper, blinded and embedding its
566 path yield a Merkle rose tree with the same hash. Furthermore, reconstructing a Merkle rose
567 tree constructed by embedding a source zipper gives the same result as first reconstructing
568 the source zipper, and then embedding it into a Merkle rose tree. Finally, we show that
569 reconstruction of trees from zippers respects the blinding relation if the Merkle operations
570 on the labels satisfy *merkle-interface*:

571 $blinding-of-tree$ h *bo* ($tree-of-zipper_m$ (p, t)) ($tree-of-zipper_m$ (p, t')) =

572 $blinding-of-tree$ h *bo* t t'

573 Inclusion proofs derived from zippers prove inclusion of a single subtree of the rose tree.
574 When we want to create an inclusion proof for several subtrees, we create an inclusion proof
575 for each individual subtree and then merge them into one. To that end, we have defined the
576 function *zippers-rose-tree* that enumerates the inclusion proof zippers for all nodes of a rose
577 tree.

578 For Canton's transactions, we have lifted the zippers and their theorems from rose trees
579 to *views*. We define the projection of the inclusion-proof embedded view for one participant
580 P as follows:

- 581 1. Enumerate all zippers for the views in the transaction using the lifted version of
582 *zippers-rose-tree*.
- 583 2. Each such zipper gives us direct access to the view's metadata. Use the metadata to
584 determine whether P is a recipient of the view. If not, filter out the zipper.
- 585 3. Convert the zippers into inclusion proofs for the view and compose each of them with
586 the transaction metadata inclusion proof.
- 587 4. Merge all these inclusion proofs into one.

588 This gives an inclusion proof for the recipient's projection of the transaction. At the end
589 of the two-phase commit protocol, the domains's commit message contains an inclusion
590 proof of the view common data for all the views that the participant should have received.
591 By comparing this inclusion proof against the projection using *blinding-of-transaction*,
592 the participant can decide whether it has received all views it was supposed to receive.
593 (Conversely, checking that it does not receive extraneous views is simple as it can be read
594 from the view metadata.)

595 **5** *Related Work*

596 Miller et al. developed a lambda calculus with authentication primitives for generic tree
597 structures [18]. The calculus was formalized in Isabelle/HOL by Brun and Traytel [4]. In the
598 calculus, the programmer annotates the structures with authentication tags. Given a value
599 of such a structure, and a function operating on it, their presented method automatically
600 creates a correctness proof accompanying a result. The proof allows a verifier that holds
601 only a digest of values with authentication tags (but not the values themselves) to check
602 the function's result for correctness. The proof is a stream of inclusion proofs, one for each
603 tagged value that the function operates on. Merging of inclusion proofs is not considered,
604 although the streams can be optimized by sharing. Unlike Brun and Traytel [4] who use
605 a deep embedding with the Nominal library, our embedding is shallow. Furthermore, our
606 ADSs can provide inclusion proofs for multiple sub-structures simultaneously. However, we
607 do not aim to derive correctness proofs for functions on the data structures.

608 White [22] designed a cryptographic ledger with lightweight proofs of transaction validity
609 and formalized the design in Coq. The ledger is a function from assets to addresses.
610 Transactions move assets between addresses and transform one ledger into another. The
611 transactions' plausibility can be proved by checking that the assets existed in the old ledger
612 and that the assets in the new ledger were moved to the correct place. Ledgers are represented
613 by a tree, where leaves list assets and a tree path encodes an address. A Merkle structure
614 over the tree and Merkle inclusion proofs of the assets' movement relieve the verifiers from
615 having to know the entire ledger. A merge operation allows a single Merkle tree to provide
616 several inclusion proofs. The Coq development is tailored to the use case: the Merkle trees
617 are binary and the leaves are restricted to fixed single type (either asset lists or sentinels that
618 mark empty subtrees). Our generic development can be instantiated to cover this structure.

619 Yu et al. [23] use Merkle constructions on different binary trees to implement logs with
620 inclusion and exclusion proofs. The constructions are proved correct using a pen-and-paper
621 approach. The proved properties are then used in the Tamarin verification tool to analyze a
622 security protocol.

623 Ogawa et al [19] formalize binary Merkle trees as used in a timestamping protocol. They
624 automatically verify parts of the protocol using the Mona theorem prover.

625 Seefried [20] observed that inclusion proofs in a Merkle tree correspond to the derivative
626 of the type, i.e., a Huet-style zipper [12], where the subtrees in zipper context have been

627 replaced by the Merkle root hashes. McBride showed that zippers represent one-hole contexts
 628 [16]. In this analogy, our inclusion proofs correspond to contexts with arbitrarily many holes
 629 where the subtrees without holes have been replaced by the corresponding hashes. These
 630 many-hole zippers must not be confused with Kiselyov’s zippers [14] and Hinze and Jeuring’s
 631 webs [11], which are derived from the traversal operation rather than the data structure .

632 **6 Conclusion and Future Work**

633 We have presented a modular construction principle for authenticated data structures over
 634 HOL datatypes (i.e., functors) that have a tree-like shape, and basic operations over these
 635 structures. The class of supported functors includes sums, products, and functions, and
 636 is closed under composition and least fixpoints. The supported operations are root hash
 637 computations and merging of inclusion proofs. We showed how to instantiate the construction
 638 to rose trees, as well as to a real-world structure used in Canton, a Byzantine fault tolerant
 639 atomic commit protocol.

640 The ongoing formalization of the Canton protocol will continue to test our abstractions
 641 and trigger further improvements. As noted earlier, ADSs cannot only improve storage
 642 efficiency, but also provide confidentiality. For example, Canton uses them to keep parts of a
 643 transaction confidential to a subset of the transaction’s participants. However, reasoning
 644 about confidentiality is not straightforward. As hashing is injective, we can simply write
 645 $inv\ h\ x$ in HOL to obtain the pre-image of a hash x . In fact, our current model does not
 646 even distinguish between the authenticated data structure and its root hash because of this.
 647 A sound confidentiality analysis must therefore restrict the adversary using an appropriate
 648 calculus, e.g., a Dolev-Yao style deduction relation [9].

649 In a system, if a source substructure S is unblinded somewhere in an inclusion proof ip ,
 650 then the confidentiality analysis of the structure should unblind all occurrences of *Blinded*
 651 (*Content* S), in ip , regardless of the position where they occur. Our blinding orders and
 652 the merge operation do not do this. For example, consider a binary Merkle tree of two
 653 leaves that both store a value x . So both leaves have the same hash, and the recipient of
 654 an inclusion proof for one leaf detects that the other leaf has the same hash. So they can
 655 deduce that the other leaf also contains the value x . Yet, in our blinding order, the inclusion
 656 proof for one leaf is strictly smaller than an inclusion proof for both leaves. For proving
 657 Canton’s integrity guarantees, this is not a problem because confidentiality is not a concern.
 658 Moreover, all leaves in the transaction tree contain nonces and the domain checks that all
 659 hashes in its inclusion proof are distinct. So the lack of unblinding might not be a problem
 660 for reasoning about confidentiality in Canton, even though a proper treatment would simplify
 661 the soundness argument.

662 A related issue is that our modular approach does not apply to commutative structures,
 663 such as multisets. The conceptual problem is that the issue with substructures and confiden-
 664 tiality also appears when merging inclusion proofs for commutative structures. One option is
 665 consider Merkle functors as quotients with respect to a normalization function that collects
 666 all unblinding information and propagates the unblinding across the whole inclusion proof.
 667 The normalized inclusion proofs then serve as the canonical representatives. We have not
 668 yet worked out whether such a construction can still be modular and whether the quotients
 669 are still BNFs [10].

670 Moreover, our representation of hashes as terms makes hashing injective. While this
 671 is "morally equivalent" to standard cryptographic assumptions, an alternative (followed
 672 by [4]) would be to prove results about authentication as a disjunction: either the result
 673 holds, or a hash collision was found. The advantage of such a statement would be that hash

674 collisions become explicit, which simplifies the soundness argument for the formalization. As
 675 is, nothing that prevents us from conceptually "evaluating" the hash function on arbitrarily
 676 many inputs, which would not be cryptographically sound. To make hash collisions explicit,
 677 we must make hashes explicit, i.e., use a type like *bitstrings* instead of terms. This can be
 678 done as additional step.

```
679 typedecl bitstring
680 class encode =
681   fixes encode :: ('a  $\Rightarrow$  bitstring)
682   assumes enj-encode: (<inj encode>
```

683 Encoding functions must be defined for all types used as arguments to *blindable*. For
 684 *blindable* itself, we then define the actual hash operation as follows.

```
685 primrec root-hash :: ('ah :: encode) blindableh  $\Rightarrow$  bitstring where
686   <root-hash (Garbage garbage) = encode-garbage garbage>
687   | <root-hash (Content x) = encode x>
```

688 This can be lifted to forests using the functorial structure of Merkle functors, similar
 689 to how $hash-F\ h = hash-F' \circ map-F_m\ h\ id$ first hashes the elements of F using h and
 690 then applies the actual function $hash-F'$. We do not expect problems with extending
 691 our constructions to such a model, but it is unclear how severely the indirection through
 692 *bitstrings* impacts our proofs, in particular the Canton formalization.

693 — References —

- 694 1 Jasmin Christian Blanchette, Johannes Hölzl, Andreas Lochbihler, Lorenz Panny, Andrei
 695 Popescu, and Dmitriy Traytel. Truly modular (co)datatypes for Isabelle/HOL. In Gerwin
 696 Klein and Ruben Gamboa, editors, *Interactive Theorem Proving (ITP 2014)*, volume 8558
 697 of *LNCS (LNAI)*, pages 93–110. Springer, 2014. doi:10.1007/978-3-319-08970-6_7.
- 698 2 Jasmin Christian Blanchette, Andrei Popescu, and Dmitriy Traytel. Operations on bounded
 699 natural functors. *Archive of Formal Proofs*, 2017. [http://isa-afp.org/entries/BNF_](http://isa-afp.org/entries/BNF_Operations.html)
 700 [Operations.html](http://isa-afp.org/entries/BNF_Operations.html), Formal proof development.
- 701 3 Sören Bleikertz, Andreas Lochbihler, Ognjen Marić, Simon Meier, Matthias Schmalz,
 702 and Ratko G. Veprek. A structured semantic domain for smart contracts. Computer
 703 Security Foundations poster session (CSF 2019), [https://www.canton.io/publications/](https://www.canton.io/publications/csf2019-abstract.pdf)
 704 [csf2019-abstract.pdf](https://www.canton.io/publications/csf2019-abstract.pdf), 2019.
- 705 4 Matthias Brun and Dmitriy Traytel. Generic authenticated data structures, formally. In
 706 John Harrison, John O’Leary, and Andrew Tolmach, editors, *Interactive Theorem Proving*
 707 *(ITP 2019)*, volume 141 of *LIPICs*, pages 10:1—10:18. Schloss Dagstuhl–Leibniz-Zentrum für
 708 Informatik, 2019. doi:10.4230/LIPICs.ITP.2019.10.
- 709 5 Canton: A private, scalable, and composable smart contract platform. [https://www.canton.](https://www.canton.io/publications/canton-whitepaper.pdf)
 710 [io/publications/canton-whitepaper.pdf](https://www.canton.io/publications/canton-whitepaper.pdf), 2019.
- 711 6 Canton: Global synchronization beyond blockchain. <https://www.canton.io/>, 2020.
- 712 7 Corda: Open source blockchain platform for business. <https://www.corda.net/>, 2020.
- 713 8 Digital Asset. Daml programming language. <https://daml.com>, 2020.
- 714 9 D. Dolev and A. Yao. On the security of public key protocols. *IEEE Transactions on*
 715 *Information Theory*, 29(2):198–208, 1983.
- 716 10 Basil Fürer, Andreas Lochbihler, Joshua Schneider, and Dmitriy Traytel. Quotients of bounded
 717 natural functors. In *IJCAR 2020*, LNCS. Springer, 2020. To appear.
- 718 11 Ralf Hinze and Johan Jeuring. Weaving a web. *J. Funct. Program.*, 11(6):681—689, 2001.
 719 doi:10.1017/S0956796801004129.
- 720 12 Gérard Huet. The zipper. *Journal of Functional Programming*, 7(5):549—554, 1997. doi:
 721 10.1017/S0956796897002864.

- 722 13 Brian Huffman and Ondřej Kunčar. Lifting and transfer: A modular design for quotients in
723 Isabelle/HOL. In *Certified Programs and Proofs*, volume 8307 of *LNCS*, pages 131—146.
724 Springer, 2013. doi:10.1007/978-3-319-03545-1_9.
- 725 14 Oleg Kiselyov. Zippers with several holes. <http://okmij.org/ftp/Haskell/Zipper2.lhs>,
726 2011.
- 727 15 Andreas Lochbihler and Ognjen Marić. Authenticated data structures as functors. *Archive of*
728 *Formal Proofs*, April 2020. http://isa-afp.org/entries/ADS_Functor.html.
- 729 16 Conor McBride. The derivative of a regular type is its type of one-hole contexts, 2001.
- 730 17 Ralph C. Merkle. A digital signature based on a conventional encryption function. In
731 *Advances in Cryptology (CRYPTO 1987)*, volume 293 of *LNCS*, pages 369–378. Springer,
732 1987. doi:10.1007/3-540-48184-2_32.
- 733 18 Andrew Miller, Michael Hicks, Jonathan Katz, and Elaine Shi. Authenticated data struc-
734 tures, generically. In *Principles of Programming Languages (POPL 2014)*, pages 411—423.
735 Association for Computing Machinery, 2014. doi:10.1145/2535838.2535851.
- 736 19 Mizuhito Ogawa, Eiichi Horita, and Satoshi Ono. Proving properties of incremental Merkle
737 trees. In Robert Nieuwenhuis, editor, *Automated Deduction (CADE 2005)*, volume 3632 of
738 *LNCS*, pages 424–440. Springer, 2005. doi:10.1007/11532231_31.
- 739 20 Sean Seefried. Automatically generating tamper resistant data structures. Functional Program-
740 ming Sydney, [http://code.ouroborus.net/fp-syd/past/2017/2017-04-Seefried-Merkle.](http://code.ouroborus.net/fp-syd/past/2017/2017-04-Seefried-Merkle.pdf)
741 pdf, 2017.
- 742 21 Dmitry Traytel, Andrei Popescu, and Jasmin C. Blanchette. Foundational, compositional
743 (co)datatypes for higher-order logic: Category theory applied to theorem proving. In *Logic*
744 *in Computer Science (LICS 2012)*, pages 596—605. IEEE Computer Society, 2012. doi:
745 10.1109/LICS.2012.75.
- 746 22 Bill White. A theory for lightweight cryptocurrency ledgers. [https://github.com/](https://github.com/input-output-hk/qeditas-ledgertheory)
747 [input-output-hk/qeditas-ledgertheory](https://github.com/input-output-hk/qeditas-ledgertheory), 2015.
- 748 23 Jiangshan Yu, Vincent Cheval, and Mark Ryan. DTKI: a new formalized PKI with no trusted
749 parties. *The Computer Journal*, 59(11):1695–1713, November 2016. arXiv: 1408.1023. URL:
750 <http://arxiv.org/abs/1408.1023>, doi:10.1093/comjnl/bxw039.