

OWebSync: Seamless Synchronization of Distributed Web Clients

Kristof Jannes, Bert Lagaisse and Wouter Joosen

Abstract—Many enterprise software services are adopting a fully web-based architecture for both internal line-of-business applications and for online customer-facing applications. Although wireless connections are becoming more ubiquitous and faster, mobile employees and customers are often offline due to expected or unexpected network disruptions. Nevertheless, continuous operation of the software is expected. This paper presents OWebSync: a web-based middleware for data synchronization in interactive groupware with fast resynchronization of offline clients and continuous, interactive synchronization of online clients. To automatically resolve conflicts, OWebSync implements a fine-grained data synchronization model and leverages state-based Conflict-free Replicated Data Types. This middleware uses three main tactics to achieve the required interactive performance: Merkle-trees embedded in the tree-structured data, virtual Merkle-tree levels, and message batching. Our comparative evaluation with available operation-based and delta-state-based middleware solutions shows that OWebSync is especially better in operating in and recovering from offline settings and network disruptions. In addition, OWebSync scales more efficiently over time, as it does not store version vectors or other meta-data for all past clients.

Index Terms—CRDTs, Groupware, Web browsers, Eventual Consistency

1 INTRODUCTION

WEB applications are the default architecture for many online software services, both for internal line-of-business applications such as Customer Relationship Management (CRM), billing, and Human Resources (HR); as well as for customer-facing services. Browser-based service delivery fully abstracts the heterogeneity of the clients, solving the deployment and maintenance problems that come with native applications. Nevertheless, native applications are still used when rich and highly interactive GUIs are required, or when applications must function offline for a long time. The former reason is disappearing as HTML5 and JavaScript are becoming more powerful. The latter reason should be disappearing too with the arrival of WiFi, 4G and 5G ubiquitous wireless networks. In reality, connectivity is often missing for minutes to hours. Mobile employees can be working in cellars or tunnels, and customers sometimes want to use a web-based service on an airplane.

Interactive groupware applications, such as collaborative web applications with concurrent edits on shared data, should offer prompt data *synchronization* with *interactive* performance when online. We use the term synchronization here to describe the process of keeping data of multiple replicas eventually consistent by means of replication.

This paper focuses on prompt and seamless synchronization when clients were offline due to network disruptions, while maintaining interactive synchronization in the online setting. The research of Nielsen on usability engineering [1] states that remote interactions should take only 1-2 seconds to keep the user experience seamless and interactive. Users are annoyed after a 5 second waiting period and 10 seconds is the absolute maximum before users leave the application.

Several client-side frameworks for synchronization of semi-structured data exist. They support fine-grained and concurrent updates on local copies of shared data and operate conflict-free in online and offline situations. However, there is no generic, fully web-based middleware solution that can be used by interactive web applications to:

- 1) achieve continuous and interactive synchronization for online clients and prompt resynchronization for offline clients,
- 2) scale to tens (20-30) of online clients that concurrently edit a document with interactive performance,
- 3) tolerate hundreds of clients over time without inflating the data with versioning metadata.

State-of-the-art data synchronization frameworks are either operation-based, state-based or delta-state-based. Operation-based approaches distribute the updates as operations to all replicas. Operational Transformation, as used in Google Docs [2], is a popular operation-based technique for real-time synchronization in web applications, but it is not resilient against message loss or out-of-order messages. It requires a central server transforming the operations for other clients to deal with concurrent changes. Commutative Replicated Data Types [3], [4], as used in SwiftCloud [5], [6], Yjs [7], [8], [9] and Automerge [10], [11], are also operation-based. Again, updates must be propagated, as operations, to all clients using a reliable, exactly-once, message channel. However, no transformation is needed because concurrent operations are commutative. State-based Convergent Replicated Data Types [4] are resilient against message loss, but have often been considered as problematic since the full state has to be transferred between all replicas each time. However, it is used for background synchronization between data centers, e.g. in Riak [12]. Merkle Search Trees [13]

• The authors are with imec-DistriNet, KU Leuven, 3001 Leuven, Belgium.
E-mail: {kristof.jannes, bert.lagaisse, wouter.joosen}@cs.kuleuven.be

are proposed as a solution to the high bandwidth usage. It uses Merkle-trees [14] to replicate a basic key-value store like in Dynamo [15]. The solution works in large systems with low rates of updates for asynchronous background synchronization between backend servers; it is not suited for interactive groupware. Delta-state-based Conflict-free Replicated Data Types [16], as used in Legion [17], [18], need less of the message channel than the operation-based approaches. However, they use vector clocks to calculate delta-updates, which require one entry per writing client per object in the server-side metadata. This does not integrate well with the dynamic nature of the web, where it is often uncertain if a client will ever connect to a server again.

In this paper, we present OWebSync, a generic web middleware for data synchronization in browser-based applications and interactive groupware. It supports offline usage with fast resynchronization, as well as continuous and interactive synchronization between online clients. OWebSync provides a generic, reusable data type, based on JSON [19], that web application developers can leverage to model their application data. One can nest several map structures into each other to build a complex tree-structured data model. These data types support fine-grained and conflict-free synchronization by leveraging state-based Conflict-free Replicated Data Types (CRDTs). OWebSync solves the scalability issue that comes with operation-based approaches, where server-side metadata will grow linearly over time with the number of clients present in the system at some point. It reduces the required bandwidth by combining several tactics such as Merkle-trees embedded in the tree-structured data, virtual Merkle-tree levels, and message batching. As such, OWebSync can achieve the interactive performance of operation-based approaches, while maintaining the inherent *robustness* of state-based approaches.

This paper is structured as follows. Section 2 provides two motivating case studies and then provides the rationale and more background on synchronization mechanisms such as CRDTs. Section 3 describes the underlying data model based on CRDTs and Merkle-trees. Section 4 presents the deployment and runtime synchronization architecture together with two performance optimization tactics. Section 5 compares and evaluates performance in online and offline situations using OWebSync and other state-of-the-art synchronization frameworks. We discuss related work in Section 6 and then we conclude.

2 MOTIVATION AND BACKGROUND

This section explains the motivation of the goal and approach of OWebSync. First, we present two case studies of online software services for mobile employees and customers that often encounter offline settings due to expected or unexpected network disruptions. We then provide background information on Operational Transformation, Conflict-free Replicated Data Types and Merkle-trees.

2.1 Case studies

The motivation and requirements have emerged from two case studies from our applied research projects with industry, that have also been used for the evaluation of the

middleware. The first case study is an online software service from eWorkforce, a company that provides technicians to install network devices for different telecom operators at their customers' premises. The second company is eDesigners, which offers a web-based design environment for graphical templates that are applied to mass customer communication.

2.1.1 eWorkforce

eWorkforce has two kinds of employees that use the online software service: the help desk operators at the office and the technicians on the road. The help desk operators accept customer calls, plan technical intervention jobs and assign them to a technician. The technicians can check their work plan on a mobile device and go from customer to customer. They want to see the details of their next job wherever they are and must be able to indicate which materials they used for a job. Since they are always on the road, a stable internet connection is not always available. Moreover, they often work in offline mode when they work in basements to install hardware. Writing off all used materials is crucial for correct billing and inventory afterwards.

This case study requires support for long term offline usage, with quick synchronization when coming online, especially for last-minute changes to the work plan of the technicians. The help desk software must be operational at all times, even without connection to the central database, as customers can call for support and schedule interventions.

2.1.2 eDesigners

The company eDesigners offers a customer-facing multi-tenant web application to create, edit and apply graphical templates for mass communication based on the customer's company style. Templates can be edited by multiple users at the same time, even when offline. When two users edit the same document, a conflict occurs, and the versions need to be merged. Edits that are independent of each other should both be applied to the template, e.g. one edit changes the color of an object, another edit changes the size. When two users edit the same property of the same object, only one value can be saved. This should be resolved automatically as to not interrupt the user.

This case study requires that the application is always available, updates must always be possible, even offline when working on an airplane. When coming back online, the updates should be synchronized promptly without requiring the user or the application to manually resolve conflicts. When working online, the performance should be interactive, especially when two users are working on the same template next to each other.

2.2 Background

The previous section described the overall goal. In this section, we discuss the advantages and problems of state-of-the-art techniques such as Operational Transformation and Conflict-free Replicated Data Types (CRDTs).

2.2.1 Operational Transformation

OT [20] is a technique that is often used to synchronize concurrent edits on a shared document. It works by sending

the operations to the other replicas. The operations are not necessarily commutative, which means they cannot be applied immediately on other replicas. A concurrent edit might conflict with another operation. Therefore, a central server is used to transform the operations for the different replicas so that the resulting operations maintain the original semantics. The problem is that the transformation of the incoming operations of other clients on their local state can get very complex. Messages can also get lost or can arrive in the wrong order. Hence, OT is not resilient against message loss, or long-lasting offline situations [21].

2.2.2 Conflict-free Replicated Data Types

CRDTs [4], [22] are data structures designed for replication that guarantee eventual consistency without explicit coordination with other replicas. Conflict-free means that conflicts are resolved automatically in a systematic and deterministic way, such that the application or user does not have to deal with conflicts manually. There are two kinds of CRDTs: operation-based or Commutative Replicated Data Types (CmRDT) and state-based or Convergent Replicated Data Types (CvRDT).

Commutative Replicated Data Types. CmRDTs [22] make use of operations to reach consistency, just like OT. Concurrent operations in CmRDTs must be commutative and can be applied in any order. This way, there is no central server necessary to apply a transformation on the operations. As with OT, CmRDTs need a reliable message broadcast channel so that every message reaches every replica exactly-once. Causally ordered delivery is required in some cases.

Convergent Replicated Data Types. CvRDTs [22] are based on the state of the data type. Updates are propagated to other replicas by sending the whole state and merging the two CvRDTs. For this merge operation, there is a monotonic join semi-lattice defined. This means that there is a partial order defined over the possible states and a least-upper-bound operation between two states. The least-upper-bound is the smallest state that is larger or equal to both states according to the partial order. To merge two states, the least-upper-bound is computed, which will be the new state. CvRDTs require little from the message channel: messages can get lost or arrive out-of-order without a problem since the whole state is always sent. However, this state can get large, and needs to be communicated every time.

Delta-state CvRDTs. δ -CvRDTs [23], [24], [25] are a variant of state-based CRDTs with the advantage that in some cases only part of the state (a delta) needs to be sent for correct synchronization. When a client performs an update, a new delta is generated which reflects the update. Each client keeps a list of deltas and remembers which clients have already acknowledged a delta. As soon as all clients have acknowledged a delta, it can be discarded because the update is now reflected in the state of all clients. If a client was offline and has missed too many deltas, then the full state must be sent, just like with normal state-based CRDTs.

δ -CRDTs have some problems when using them in web applications. Browser-based clients come and go with a large churn rate and it is often unclear if a client will come back online in the future (e.g. browser cache cleared). Keeping extra metadata for all those clients, to be able to synchronize only the required deltas, can result in a large

storage or memory overhead to keep track of them at the server. One can always discard the metadata for clients that were offline and send the full state if they do come back online eventually. But this is of course not efficient when the state is large and the client already had most of the updates.

A variant of δ -CRDTs, called Δ -CRDTs [16], is proposed as a solution to this problem. Δ -CRDTs are comparable to δ -CRDTs, but instead of keeping track of the clients at the server, it includes extra metadata about concurrent versions of all clients in the data model, as vector clocks, to calculate the deltas dynamically. This solves the problem of keeping track of the deltas for clients at the server, but it still needs client identifiers and version numbers inside the vector clocks for each object, and each client that made a change.

Another approach to optimize δ -CRDTs is using join decompositions [26], [27]. This approach does not extend CRDTs with additional metadata that needs to be garbage collected. Instead, it can efficiently calculate a minimal delta to synchronize. While this improves the network usage compared to normal δ -CRDTs, it still requires clients to keep track of their neighbors. When there is no such data available, e.g. after a network partition, it needs to fall back to a state-based approach. However, it only requires sending the full state in a single direction, compared to bidirectionally in normal state-based CRDTs. A digest-driven approach is also supported, which will send a smaller digest of the actual state. However, for many CRDTs, such digest does not exist and for large, nested data, this digest would still be large.

2.3 Principles

We now introduce two state-based CRDTs and Merkle-Trees. We will use these as building blocks in the next section for our data model.

2.3.1 LWWRegister

A Last-Write-Wins Register (LWWRegister) [4] is a CvRDT that contains exactly one value (string, number or boolean) together with a timestamp of the last change. This timestamp will be used to merge another replica of this LWWRegister. The value associated with the highest timestamp is kept, while the other value is discarded. This conflict resolution strategy boils down to a simple last-write-wins strategy.

2.3.2 ORSet

An Observed-Removed Set (ORSet), as described by Shapiro et al. [4], [28], is a set CvRDT. Internally, the ORSet contains two sets: the observed set and the removed set. When an item is added to the set, it is added to the observed set together with a unique tag. When that item is removed, the associated tag is added to the removed set, and the item itself is removed from the observed set. This allows an item to be removed and added multiple times. All items present in the observed set, but not in the removed set are currently present in the set. The conflict resolution of the ORSet boils down to an add-wins resolution, i.e. a concurrent add and remove operation will result in the item being present in the set since each add will get a new identifier. To merge a local replica of an ORSet with another replica, the union of the respective observed and removed set is taken, and removed items are removed from the observed set.

2.3.3 Merkle-trees

Merkle-trees [14] or hash-trees are used to quickly compare two large data structures. Merkle-trees are trees where each node contains a hash. The values of the leaf nodes are hashed and each hash in an internal node is the hash of the hashes of all its immediate children. Two data structures can now be compared starting from the two top-level hashes. If the top-level hashes match, the data structures are equal. Otherwise, the tree can be descended using the mismatching hashes to find the mismatching items. Sub-trees which are already equal will have equal hashes at their top nodes, so they do not need further verification.

3 THE OWEBSYNC DATA MODEL

This section describes the data model of OWebSync that will be used for synchronization. The data model is a CvRDT for the efficient replication of JSON data structures and applies Merkle-trees internally to quickly find data changes.

3.1 Approach

OWebSync uses state-based CRDTs, which require little from the message channel compared to operation-based approaches. No state about other clients or client-based versioning metadata needs to be stored, unlike delta-state approaches. And even after long offline periods, the missed updates can be computed and synchronized seamlessly. To limit the overhead of full state exchanges between clients and server, we adopt Merkle-trees in the data structure to find the items that need to be synchronized efficiently. The CvRDT consist of two types: a LWWRegister and an ORMap extended with a Merkle-tree. The LWWRegister is used to store values in the leaves of the tree, and is implemented as described by Shapiro et al. [4]. The ORMap is a recursive data structure that represents a map containing a mapping from strings to other ORMaps or LWWRegisters.

3.2 Observed-Removed Map

The ORMap is implemented starting from a state-based Observed-Removed Set [28]. The items in the observed set are key-value pairs, where a key is a simple string, and the value is a reference to another CvRDT. Concurrent edits to different keys can be made without a problem. Edits to the same key and tag will be delegated to the child CRDT: either another ORMap or a LWWRegister. If two different replicas add the same key, they will get a different tag. This situation is difficult to resolve, and we opted to merge the two values, and keeping only the lexicographical greatest tag. As a result, a single replica of an ORMap has always at most one single value for a key.

We made two extensions to this basic ORMap to make state-based synchronization more efficient. First, we extended this data structure with a Merkle-tree using the object's logical tree-structure. This means that we keep an extra hash for all items in the observed set. When the child is a LWWRegister, the hash is the hash of the value of that register. When the child is another ORMap, the hash of it is the combined hash of the hashes of all its children, lexicographically ordered on the unique tags. This way, when one value in a register changes, all the hashes of the

```

1:  $KV$  ▷ Key-value store
2:  $p_0..p_n$  ▷ Array representation of a path in the tree
3: state:
4:    $O \leftarrow \emptyset$  ▷ Observed set with tuples  $(key, tag, hash)$ 
5:    $R \leftarrow \emptyset$  ▷ Removed set with  $tags$ 
6:    $PATH$  ▷ The path of this ORMap
7: query: GET ( $p_0..p_n$ )
8:   if  $\exists o \in O \wedge o.key = p_0$  then
9:      $c \leftarrow KV.GET(PATH + o.key)$ 
10:    return  $c.GET(p_1..p_n)$ 
11:   return  $\perp$ 
12: update: SET ( $p_0..p_n, value$ )
13:   if  $\exists o \in O \wedge o.key = p_0$  then
14:      $c \leftarrow KV.GET(PATH + o.key)$ 
15:      $c.SET(p_1..p_n, value)$ 
16:      $o.hash \leftarrow c.hash$ 
17:   else
18:      $c \leftarrow NEW\_ORMAP(PATH + p_0)$ 
19:      $c.SET(p_1..p_n, value)$ 
20:      $O \leftarrow O \cup \{(p_0, UNIQUE(), c.hash)\}$ 
21:   update: REMOVE ( $p_0..p_n$ )
22:   if  $\exists o \in O \wedge o.key = p_0$  then
23:     if  $LEN(p_0..p_n) = 1$  then
24:        $O \leftarrow O \setminus \{o\}$ 
25:        $R \leftarrow R \cup \{o.tag\}$ 
26:     else
27:        $c \leftarrow KV.GET(PATH + o.key)$ 
28:        $c.REMOVE(p_1..p_n)$ 
29:        $o.hash \leftarrow c.hash$ 
30:   update: REMOVE_WITH_TAG ( $p_0..p_n, tag$ )
31:   if  $LEN(p_0..p_n) = 0$  then
32:     if  $\exists o \in O \wedge o.tag = tag$  then
33:        $O \leftarrow O \setminus \{o\}$ 
34:        $R \leftarrow R \cup \{o.tag\}$ 
35:   else if  $\exists o \in O \wedge o.key = p_0$  then
36:      $c \leftarrow KV.GET(PATH + o.key)$ 
37:      $c.REMOVE\_WITH\_TAG(p_1..p_n, tag)$ 
38:      $o.hash \leftarrow c.hash$ 
39:   (Continues in Alg. 2)

```

Alg. 1. Simplified implementation of the state, query- and update operations of an ORMap with a Merkle-tree for synchronization. Several details and optimization tactics are abstracted.

parents will also change, so that a change can be detected by only comparing the top-level hash. Second, we do not store a child CRDT inside the observed set, instead we only store the tag, key and hash of that CRDT. The child CRDTs can be stored elsewhere using its path as a unique key.

Alg. 1 and Alg. 2 show the specification of the OWebSync ORMap with our two extensions. The ORMap supports several operations to query, update and merge this data structure. The GET operation is equal to the one in a basic ORMap. There is always at most one single object in the observed set with a specific key. The SET and REMOVE operation are also similar, but require updating the hash to keep the Merkle-tree up-to-date. The REMOVE_WITH_TAG operation removes a single element, based on the tag instead of the key. The MERGE operation is modified to make use of the Merkle-tree. It returns all paths that are changed accord-

```

40: merge: MERGE ( $p_0..p_n$ , remote)
41:    $N \leftarrow \emptyset$   $\triangleright$  paths that need synchronization
42:   if LEN( $p_0..p_n$ ) = 0 then
43:      $R \leftarrow R \cup \text{remote}.R$ 
44:      $O \leftarrow \{o \in O : o.tag \notin R\}$ 
45:     for all  $ro \in \text{remote}.O \wedge ro.tag \notin R$  do
46:       if  $\exists o \in O \wedge o.key = ro.key$  then
47:         if  $o.tag \neq ro.tag$  then
48:           if  $ro.tag > o.tag$  then
49:              $o.tag \leftarrow ro.tag$ 
50:           else
51:              $N \leftarrow N \cup \{PATH\}$ 
52:           if  $o.hash \neq ro.hash$  then
53:              $N \leftarrow N \cup \{PATH + o.key\}$ 
54:           else
55:              $N \leftarrow N \cup \{PATH + ro.key\}$ 
56:         else
57:           if  $\exists o \in O \wedge o.key = p_0$  then
58:              $c \leftarrow KV.GET(PATH + o.key)$ 
59:              $N \leftarrow N \cup c.MERGE(p_1..p_n, \text{remote})$ 
60:              $o.hash \leftarrow c.hash$ 
61:           else
62:              $c \leftarrow NEW\_ORMAP(PATH + p_0)$ 
63:              $N \leftarrow N \cup c.MERGE(p_1..p_n, \text{remote})$ 
64:              $O \leftarrow O \cup \{(p_0, c.tag, c.hash)\}$ 
65:   return  $N$ 

```

Alg. 2. Simplified implementation of the merge operations of an ORMap with a Merkle-tree for synchronization.

ing to the Merkle-tree. In a next step, the synchronization protocol will descend in the tree and continue the MERGE operation in the next level. Only the returned paths are merged further, the other branches of the tree do not need further processing. By splitting up this operation per level in the tree, only the updated registers and parent ORMaps need to be sent over the network, improving both the bandwidth usage as well as saving computation power as not all CRDTs need to be merged. We explain this synchronization protocol in more detail in Section 4. We use a key-value store to store the CRDTs, called *KV* in the specification.

Proof sketch. A state-based object is a CvRDT when the states of that object forms a monotonic join semilattice. This means that there is partial order defined over the states, and a least upper bound (LUB) operation on two states which results in the smallest state that is larger or equal to the two given states according to the partial order [4], [22]. The partial order of the modified ORMap defined here is similar to the ORSet [28], which contains two grow-only sets. As an optimization, removed items are only present in the removed-set and are removed from the observed-set, however, conceptually they are still part of the observed-set when determining the partial order. The LUB operation, equal to the MERGE operation in Alg. 2, takes the union of the respective observed- and removed-set. Again using the optimization that removed items are not actually stored in the observed-set anymore. Two complications added here are the key and the hash. When a key is present in both ORMaps, with a different tag, the LUB operation will only

```

{
  "drawing1": {
    "object36": {
      "fill": "#f00",
      "height": 50,
      "left": 50,
      "top": 100,
      "type": "rect",
      "width": 80
    }
  }
}

```

Fig. 1. JSON data structure of the eDesigners case study.

```

- drawing1.object36:
  {
    tag: 0a2f7bc2-129f-11e9-ab14-d663bd873d93,
    hash: 7319eae53558516daafac19183f2ee34,
    observed: [
      {
        key: "top",
        tag: 23c1259a-129f-11e9-ab14-d663bd873d93,
        hash: 65bdd1b610f629e54d05459c00523a2b
      },
      {
        key: "left",
        tag: 0eac2a3a-546f-11e9-8647-d663bd873d93,
        hash: 67507876941285085484984080f5951e
      },
      ...
    ],
    removed: []
  }
- drawing1.object36.top:
  {
    tag: 23c1259a-129f-11e9-ab14-d663bd873d93,
    hash: 65bdd1b610f629e54d05459c00523a2b,
    timestamp: 789778800000,
    value: "100"
  }

```

Fig. 2. Internal structure of two key-values that represent `object36` and the property `top` of the JSON data structure in Fig. 1. We use a pseudo-JSON notation here, however, in practise these two key-values are stored in a binary format in a key-value store.

keep the largest tag, and merge the two values according to the rules of the child CRDT. The other tag is considered removed. When the hash differs, the two values are merged according to the rules of the child CRDT, after which the hashes will become equal.

Example. As an example, we illustrate the conceptual representation of an application data object in the eDesigners case study, as well as the resulting underlying CRDTs in the OWebSync data model. Fig. 1 present a JSON data structure of a drawing with one rectangle object. Fig. 2 represents the internal structure of two CRDTs in that JSON structure. First, the key under which the CRDT is stored in a key-value store is listed, then the internal value of the CRDT. There is an ORMap stored in the key-value store for key "" (the root of the tree), `drawing1` and `drawing1.object36`. Only `drawing1.object36` is shown in the figure. For all the leaf-values, there is a LWWRegister stored under the respective keys, for conciseness, only `drawing1.object36.top` is shown. The application developer only needs to know about the conceptual JSON representation, the middleware will automatically translate this data model and its operations to the underlying CRDTs and maintain the Merkle-tree and the internal CRDT structure.

3.3 Considerations and discussion

The data model is best suited for semi-structured data that is produced and edited by concurrent users. Any data that can be modeled in a tree-like structure such as JSON and that can tolerate eventual consistency, can use OWebSync for the synchronization. Examples are the data items in the case studies: graphical templates, a set of tasks or used materials for a task. This data model is less suited for applications such as online banking which requires constraints on the data such as: “your balance can never be less than zero”. Text-editing is also not a great fit, because there is not much structure in the data. If you would see text as a list of characters, it would result in a tree with one top-level node and one level with many child nodes: the characters. There is no benefit in using a Merkle-tree here. OWebSync also expects that no clients are malicious.

The timestamps in the LWWRegisters are provided by the clients and we do not consider malicious clients. We also assume loosely synchronized clocks. If this assumption is not met, an accidental fault resulting in a clock several years in the future, might make edits to this LWWRegister impossible. Users can resolve this manually by removing that register and creating a new one with the same key in the ORMap, loosing concurrent changes. We do not consider this an important drawback, as most personal devices these days automatically synchronize their time with the internet.

In the current data model, the ORMap keeps the tags of all removed children eternally, so-called tombstones. As a result, the size of an ORMap can accumulate over time and performance will degrade. With a modest usage of deletion, this will not be a problem. Even when you remove a large sub-tree of several levels deep, only the tag of its root is kept in the parent. One strategy to clean up tombstones could be to remove those older than e.g. one month. We then expect that a client will not be offline for more than a month while performing concurrent edits. This can be enforced by logging out the user after a month of no usage. Delta-state-based CRDTs can avoid tombstones by encoding the causal context as a compact version vector. However, this version vector grows in size with the number of clients that make changes to this ORMap. We opted for tombstones, which can be garbage collected after a sufficiently long time, because we target a dynamic environment such as the web. Web clients come and go, without long term presence.

Another kind of conflict occurs when two different types of CRDTs are assigned concurrently at the same position in the JSON structure. In this case, the merge-operation of the defined CRDTs cannot be used to resolve the conflict. This is solved by posing an order on the possible CRDTs, e.g. LWWRegister < ORMap. This means that when such a conflict occurs, the ORMap is selected as actual value, while the LWWRegister is discarded.

Another conflict is a concurrent remove and update of a child CRDT. The CRDT proposed here maintains a remove-wins semantic. This means that updates done to children are discarded when the parent is removed concurrently.

Beside primitive values and maps, the JSON specification also contains ordered lists. This is currently not supported by OWebSync. We focused on the initial key data structures: last-write-wins registers and maps. Keeping a

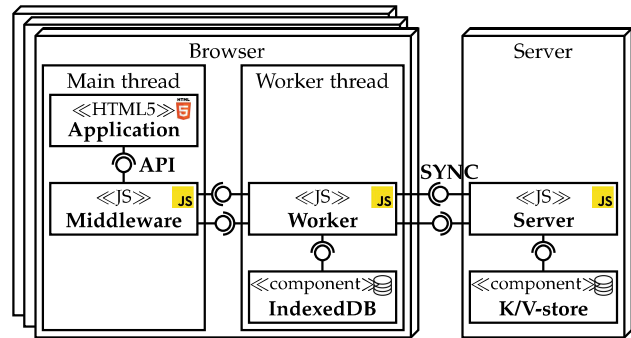


Fig. 3. Overall architecture of the OWebSync middleware

total numbered order, as lists do, is rarely needed. Unique IDs in a map are better suited in a distributed setting. In the case studies, the ordering of items in a set was based on application-specific properties such as dates, times or other values, instead of an auto-incremented number of a list. However, CvRDTs for ordered lists exist [4], [29] and could be added in future work. Adding new kinds of CRDTs to the data model is straight-forward. An existing CvRDT can be used as is, except for an extra hash to be part of the Merkle-tree. For a CRDT that represents a leaf value (e.g. a Multi-Value Register [4]), the hash is simply the hash of that value. For CRDTs that can contain other values, the hash must combine the hashes of all the children.

4 ARCHITECTURE AND SYNCHRONIZATION PROTOCOL

This section describes the deployment and execution architecture of OWebSync as well as the synchronization protocol. This middleware architecture is key to support the data- and synchronization model described in the previous section. We also elaborate on a set of key performance optimization tactics to achieve continuous, prompt synchronization for online interactive clients.

4.1 Overall architecture

The middleware architecture is depicted in Fig. 3 and consists of a client and a server subsystem. The client-tier middleware API is fully implemented in JavaScript and runs completely in the browser, without add-ins or plugins. The server is a light-weight process, which listens to incoming web requests. The server is only responsible for data synchronization, it does not run application logic. Both client and server have a key-value store to persist data, and they communicate using only web-based HTTP traffic and WebSockets [30]. All communication messages are sent and received inside the client and server subsystems using asynchronous workers. The tags in the ORMap are UUIDs [31] and we use the MD5 [32] algorithm for hashing. We first elaborate on the client-tier subsystem with the public middleware API for applications, and then describe the client-server synchronization protocol.

4.2 Client-tier middleware and API

The public programming API of the middleware is located completely at the client-tier, and web applications are de-

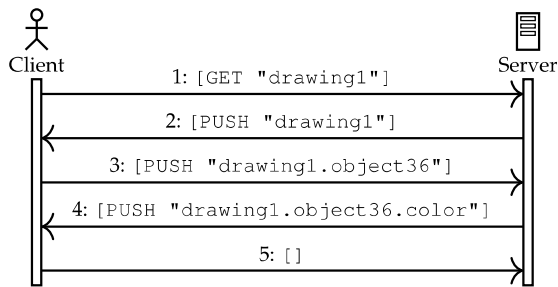


Fig. 4. Synchronization protocol when another client made an update to the color. A GET message only sends the path and hash value, a PUSH message also sends the respective CRDT. E.g. for message 3, the first CRDT in Figure 2 is sent.

veloped as client-side JavaScript applications that use the following API:

- GET (path): returns a JavaScript object or primitive value for a given path.
- LISTEN (path, callback): similar to GET, but every time the value changes, the callback is executed.
- SET (path, value): set or update a value at a given path.
- REMOVE (path): remove the value or sub-tree at the given path.

The OWebSync middleware is loaded as a JavaScript library in the client and the middleware is then available in the global scope of the web page. One can then load and edit data using typical JavaScript paths. An example from the eDesigners case study:

```

let d1 = await OWebSync.get("drawing1")
d1.object36.color = "#f00"
OWebSync.set("drawing1", d1)
  
```

The object at "drawing1" is fetched from disc and is represented as a JavaScript object in memory. If there would be other drawings (e.g. drawing2), they will not be loaded. The access to "d1.object36.color" is just a plain JavaScript object access and does not involve OWebSync. For performance reasons, it is best to always scope to the smallest possible object from the database, in this example that would be:

```
OWebSync.set("drawing1.object36.color", "#f00")
```

4.3 Synchronization protocol

The synchronization protocol between client and server consists of three key messages that the client can send to the server and vice versa:

- GET (path, hash): the receiver returns the CRDT at a given path if the hash is different from its own CRDT at the given path.
- PUSH (path, CRDT): the sender sends the CRDT data structure at a given path and the receiver will merge it at the given path.
- REMOVE (path, tag): removes the CRDT at a given path if the unique identifier of the value is matching the given tag. As such, a newer value with a different tag will not be removed.

```

1: KV
2: sync: SYNC (msgs)
3:   resp ← []
4:   for all msg ∈ msgs do
5:     if msg ≡ GET(path, hash) then
6:       c ← KV.GET(path)
7:       if c.hash ≠ hash then
8:         resp.PUSH(PUSH(path, c))
9:     else if msg ≡ PUSH(path, crdt) then
10:      c ← KV.GET("")
11:      paths ← c.MERGE(path.SPLIT(".", " "), crdt)
12:              ▷ Procedure from Alg. 2
13:      for all p ∈ paths do
14:        if KV.HAS(p) then
15:          c ← KV.GET(p)
16:          resp.PUSH(PUSH(p, c))
17:        else
18:          resp.PUSH(GET(p, ⊥))
19:      else if msg ≡ REMOVE(path, tag) then
20:        c ← KV.GET("")
21:        c.REMOVE_WITH_TAG(path.SPLIT(".", " "), tag)
22:                ▷ Procedure from Alg. 1
23:      return resp
  
```

Alg. 3. Specification of the synchronization protocol, using the ORMap specified in Section 3. Some details are abstracted for conciseness.

The protocol, depicted in Alg. 3, is initiated by a client, which will traverse the Merkle-tree of the CRDTs. The synchronization starts with the CRDT in the root of the tree. The client will send a GET message to the server with the given path and hash value of the CRDT. If the server concludes that the hash of the path matches the client's hash, the synchronization stops. An empty message is sent to signal this to the other side. All data is consistent at that time. If the hash does not match, the server returns a PUSH message with the CRDT that is located at the path requested by the client. This does not include the child CRDTs, only the metadata (key, tag, and hash) of the immediate children. The client must merge the new CRDT with the CRDT at its requested path. The specification is listed in Alg. 2. The MERGE operation returns a set with all changed paths. Those paths are the paths of the conflicting CRDTs that still need to be synchronized. The client will then PUSH the CRDTs belonging to those paths to the server. The server then needs to merge those CRDTs. If a child does not yet exist, an empty child is created and a GET message is sent. The process continues by traversing the tree and exchanging PUSH and GET messages until the leaves of the tree are reached. The CRDT in this leaf is a register and can be merged immediately. All parents of this leaf are now updated such that finally the top-level hash of client and server match. If the top-level hashes do not match, other updates have been done in the meantime, and the process is repeated. Per PUSH message that is sent, the process descends one level in the Merkle-tree. The length of the synchronization protocol is therefore limited to the maximum depth of the Merkle-tree.

The third type of message, REMOVE, is not strictly necessary, but can improve the bandwidth requirements. If

during this synchronization process between a client and the server, a child is removed at one side, but not at the other side, a REMOVE message is sent to the other party so that it can remove that value and add the tag to the removed set of the correct ORMap. Alternatively, this additional third message type of REMOVE could be avoided if a PUSH of the parent would be sent instead. However, the push of a parent with many children would cause a serious overhead compared to a REMOVE message with only a path and a tag.

Fig. 4 shows an example of the eDesigners case study where the client changed the color of an object. If the client had made multiple changes, e.g. he also changed the height, the start of the synchronization protocol would be the same, except that the height will also be included in message four.

4.4 Performance optimization tactics

The main optimization tactic to achieve prompt synchronization for interactive groupware is the reduction of network traffic by the Merkle-trees. However, there are additional tactics needed to further improve synchronization time. To reduce the overhead of the synchronization protocol between the many clients and the server, two optimization tactics are applied by both the client and the server.

4.4.1 Virtual Merkle-tree levels

When the number of child values in an ORMap increases, all the metadata for those children (key, tag, and hash) needs to be sent each time during the synchronization to check for changes. This leads to very high network usage since it cannot make use of the Merkle-tree efficiently. To solve this problem, we introduced extra, virtual, levels in the Merkle-tree. Whenever an ORMap needs to be transmitted which contains many children (i.e. hundreds), instead an extra Merkle-tree level is sent. This extra level combines the many children in groups of e.g. 10. This number can be adapted to the total number of children. As a result, 10 times fewer hashes will be sent, combined with the key-ranges the hashes belong to. The other party can verify the hashes and determine which ones are changed and then push the 10 children for which the combined hash did not match. This improvement leads to a slight delay in synchronization time since it adds one extra round-trip, but when only a small part of the children is updated, it uses much less bandwidth and reduces the computation time.

4.4.2 Message batching

In the basic protocol, all messages are sent to the other party as soon as a mismatch of a hash in the Merkle-tree is detected. This leads to lots of small messages (GET, PUSH, and REMOVE) being sent out, and as a consequence, many messages are coming in while still doing the first synchronization. This results in many duplicated messages and doing a lot of duplicated work on sub-trees since the top-level hash will only be up-to-date when the bottom of the tree is synchronized. To solve this problem, all messages are grouped in a list and are sent out in batch after a full pass of a whole level of the tree has occurred. At the other side, the messages are processed concurrently, and all resulting messages are again grouped in a list, and are only sent out after the incoming batch was fully iterated.

If no further messages are resulting from the processing of a batch, an empty list is sent to the other party to end the synchronization. As a result, fewer messages are sent between a client and server, and only one synchronization round per client is occurring at the same time, resulting in no duplicated messages and work on sub-trees.

5 PERFORMANCE EVALUATION

The performance evaluation will focus on situations where all clients are continuously online, as well as on situations where the network is interrupted. For online situations, we are especially interested in the time it takes to distribute and apply an update to all other clients that are editing the same data. For the offline situation, we are especially interested in the time it takes for all clients to get back in sync with each other after the network disruption, and in the time it takes to restore normal interactive performance.

The performance evaluation in this paper is performed using the eDesigners case study¹, as this scenario has the largest set of shared data and objects between users. The eWorkforce case study has fewer shared data with fewer concurrent updates as technicians typically work on their own data island and the data contains fewer objects with less frequent changes. To compare performance, we implemented the case study 5 times on 5 representative JavaScript technologies for web-based data synchronization: our OWebSync platform, which uses state-based CRDTs with Merkle-trees, Yjs [9] and Automerge [11] which use operation-based CRDTs, and ShareDB [33] which makes use of OT. We used Legion [17] for testing delta-state CRDTs. Both Yjs (2698 GitHub stars) and ShareDB (3768 GitHub stars) are widely used open source technologies available on GitHub. Automerge is the implementation of the JSON data type of Kleppmann and Beresford [34]. Legion is the implementation of Δ -CRDTs of van der Linde [16], [17]. We did not evaluate Google Docs, which uses OT, as it is text-based, and can not be used to synchronize the JSON-documents used in the test. Instead, we opted for ShareDB. We use Fabric.js [35] for the graphical interface.

Test setup. Both the clients and the server are deployed as separate Docker containers on a set of VMs in the Azure [36] public cloud. A VM has 4 vCPU cores and 8 GB of RAM (Standard A4 v2) and holds up to 3 client containers. A client container contains a browser that loads the client-side middleware from the server. The middleware server is deployed on a separate VM (Standard F4s v2). The monitoring server that captures all performance data is deployed on a separate VM. VMs in Azure have their clocks synchronized with the host machine, which is synchronized with the internal Microsoft time servers. The Linux `tc` tool [37] is used to artificially increase the latency between the containers to an average of 60 ms with 10 ms jitter, which resembles the latency of a 4G network in the US [38].

Our evaluation contains three benchmarks. The first benchmark represents the continuous online scenario where clients are making updates for 10 minutes and stay online

1. A try-out demo application on the middleware is available on an anonymous website (<http://owebSync.cloudapp.net>). One can open multiple Chrome browsers as concurrent clients. Use different profiles for each client: e.g. your normal profile and an incognito window.




Fig. 5. Aggregated boxplots containing the times to synchronize an update to all other clients in the online scenario. Each boxplot contains all 10 iterations for each of the 30 tests in the fully online situation. To compare technologies that have results of the same order of magnitude, as well as results in different orders of magnitude, we opted for a logarithmic Y-axis.

the whole time. The second benchmark is the offline scenario where the network connection between the clients and the server is disrupted during the test. The third and last benchmark is used to measure the total size of the data set over a longer time period².

5.1 Performance of continuous online updates

The first benchmark represents the continuous online scenario where clients are making updates for 10 minutes and stay online the whole time. In total, we executed 30 tests for this benchmark: 6 tests to be executed by each of the 5 technologies. These 6 tests vary in the number of clients and data size: 8, 16, or 24 clients are performing continuous concurrent updates on 100 or 1000 objects in a single shared data set. One such object was shown in Fig. 1 in Section 3 and has 7 attributes. The total depth of the tree is four (root – drawing1 – object36 – top). The root has one child, while drawing1 has either 100 or 100 children. And these children have itself each 7 children. This is the shape of the tree defined by the application, however, because drawing1 has many children, OWebSync will transparently add an extra layer in the tree to reduce the network usage. This increases the total depth of the tree to 5.

Each client edits one object, which leads to two random writes, the x and y position, on a shared object every second. We use at most 24 clients, which are editing the same document concurrently. In comparison, Google Docs, which is the most popular collaborative editing system today, supports a maximum of 100 concurrent users according to Google itself [2]. But in practice, latency starts to increase significantly when the number of users exceeds 10 [39]. Our performance results show the same problem for ShareDB, which uses the same technique. In our performance evaluation, one iteration of a test takes 10 minutes. Before each test, the database is populated and the initial synchronization is performed. The first minute is used for warm-up. To ensure the stability of the test results, all tests are repeated 10 times.

The following performance measurements quantify the statistical division of the time it takes to synchronize a single update to all other clients in the case of a continuous online situation. The synchronization times of the updates are illustrated in Fig. 5.

Analysis of the results. For the test with 8 clients and 100 objects, all operation-based approaches (ShareDB, Yjs, and Automerge) synchronize the updates faster than the state-based approaches (Legion and OWebSync). For these three operation-based approaches, 99% is below 0.3 seconds. Legion needs about 1.0 second for synchronizing the 99th percentile and OWebSync needs 1.3 seconds. The reason for this is that Legion and OWebSync do not keep track of which updates have been sent to which client. Hence, each time the data is synchronized, a few extra round-trips are required to calculate which updates are needed. ShareDB, Yjs, and Automerge can just send the operations. On a faster network, with less latency, both Legion and OWebSync will be able to synchronize faster than in this test, since the round-trip time will be less. But even with this high latency in this benchmark, OWebSync performs within the guidelines of 1-2 seconds for interactive performance. For the test with 24 clients and 1000 objects, ShareDB has raised to 7.7 seconds for the 99th percentile. The server cannot keep up with transforming the incoming operations. Since the operations in Yjs and Automerge are commutative and do not need a transformation, the server does not become a bottleneck. These tests show that state-based CRDTs, which are currently only used for background synchronization between servers, can also be used in interactive groupware. This improvement is due to the use of Merkle-trees embedded in the data structure, the use of virtual Merkle-tree levels for large objects, and message batching.

Network trade-off. The trade-off for this scalable, prompt synchronization, is that OWebSync has a rather large network usage compared to the other tested technologies (Fig. 6). Only Automerge requires more bandwidth because it stores the whole history and uses long text-based UUIDs as client identifiers, compared to just integers in Legion. The usage of Merkle-trees reduced the network usage of OWebSync with about a factor 8 in the worst case (1000

² Tables with the detailed performance results and the raw logs and data of all tests are available on an anonymous Azure storage account: <https://owebsyncdata.blob.core.windows.net/logs/data.zip>

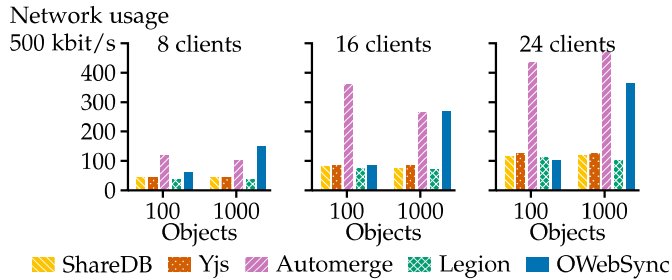


Fig. 6. Network usage per client for each test in the online scenario.

objects under a single node in the tree), compared to normal state-based CRDTs. Introducing extra, virtual, levels in the Merkle-tree for nodes with many children lowered the bandwidth with another factor 3. Even in the test with 24 clients and 1000 objects, the used bandwidth is only 360 kbit/s per client. This is much less than the available bandwidth, which is on average 27 Mbit/s on a mobile network in the US [40]. The server consumes about 8.7 Mbit/s, which is acceptable for a typical data center. The data structure has an important effect on the network usage. One might create a tree-structure with few nodes which have many children. This will make the Merkle-tree less useful since the metadata of all the children needs to be exchanged to be able to determine which children are updated. This can be seen in Fig. 6 by comparing the network usage of the tests with 100 objects to the tests with 1000 objects. The other possibility is that there are fewer children per node, but with an increased depth of the tree. This positively affects the network usage, as less metadata will need to be exchanged. However, synchronizing the whole tree will take more round-trips as there are more levels in the tree.

Interpretation and discussion. For interactive web applications and groupware, usability guidelines [1], [41] state that remote response times should be 1 to 2 seconds on average. 3 to 5 seconds is the absolute maximum before users are annoyed. The user is often leaving the web application after 10 seconds of waiting time. We start from these numbers to assess the update propagation time between users in a collaborative interactive online application with continuous updates. We are interested in the time for a user to receive an update from another online user. These numbers should be achieved not only for the average user (the mean synchronization time) but also for the 99th percentile (i.e. *most of the users* [15]). The 99th percentile for the synchronization time of the OWebSync test with 24 clients and 1000 objects is below 1.5 seconds. ShareDB operates with sub-second synchronization times when sharing 100 objects between 8 writers. But when the number of objects and writers increases, the synchronization time raises to 7.7 seconds for the 99th percentile. This is in line with the observations of Dang et al. [39] for Google Docs, which also uses OT. The other technologies stay well below 5 seconds in the online scenario and can be called interactive.

5.2 Performance in disconnected scenarios

We now present the performance analysis when the network between the clients and the server is disrupted. In these tests, we have an analogous test setup. However, during the

10-minute execution, we start dropping all messages after 3 minutes for 1 minute (shown at 2 minutes in the graphs as the first minute is used as a warm-up). This 1-minute network disruption will lead to many conflicting operations, which will automatically be resolved by the middleware. During the disruption, there will be 1440 offline updates in the largest experiment with 24 clients. A longer offline period will not change much for OWebSync since only the state is kept and the same client moving the same object twice will result in the same amount of state to be sent. Operation-based approaches will take longer when the time increases since they have to send all operations anyway.

We evaluate the time that is needed to achieve full bidirectional synchronization of all concurrent updates on all clients during the network disruption. We also evaluate the time that is needed to restore normal interactive performance in the online setting after the disruption.

Analysis of the results. The boxplots of these tests, shown in Fig. 7, show that OWebSync can synchronize all missed updates faster than ShareDB, Yjs, Automerge, and Legion. Note that these boxplots are different from the previous figure. At the median of these boxplots, only 50% of the missed updates are synchronized. Only at the upper whisker, most of the missed updates are fully synchronized. Whiskers have a maximum of 1.5 times the interquartile range.

Then, each user is fully up-to-date with everything that was updated during the network disruption. In the large scale scenario with 24 clients and 1000 updates, the time to synchronize all missed updates in case of network failure is 3.5 seconds for the 99th percentile for OWebSync, which is acceptable for interactive web applications. The other technologies need more than 5 seconds to only synchronize half of the missed updates, meaning that users will become annoyed. The operation-based approaches need several tens of seconds to synchronize all of the missed updates because they must replay all missed operations on the clients that were offline. This is due to their operation-based nature. OWebSync only needs to merge the new state, which it does in the same way as if the failure never happened. Legion could keep up with OWebSync in the online scenario, but now we see that resynchronization after network disruptions starts to take longer when the scale of the test or the size of the data set increases.

Timeline analysis of the tests. The timelines in Fig. 8 show the resynchronization times on the y-axis, without the offline time during the network disruption, for each update done at a given moment during the test timeline. This means that for an update done 20 seconds before the end of the disruption, and which got synchronized to all other clients 22 seconds later, the resynchronization time is 2 seconds.

In the test with 24 clients and 1000 objects, OWebSync quickly returns to the same performance as before the network disruption. Legion needs more time to synchronize the missed updates, but also quickly returns to the same performance. The operation-based approaches take much longer to synchronize missed updates and take tens of seconds to return to the original performance. ShareDB and Automerge need more than half a minute to return to the same interactive performance as before. This means that in a setting with frequent disconnections, the user will not be able to regain interactive performance. When coming

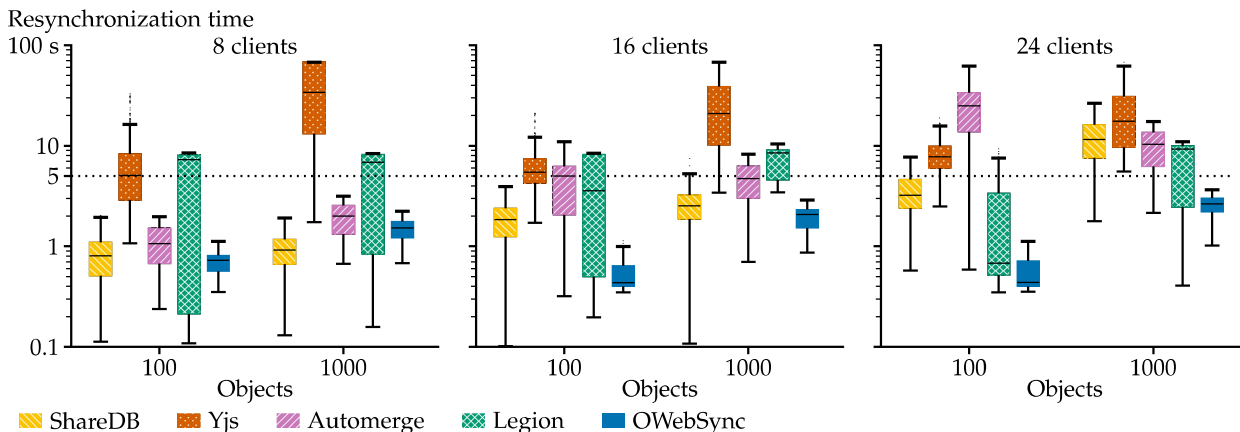


Fig. 7. Boxplots of the time it takes for an update during the failure scenario to be received by all clients. The time before a client notices the network connection is reestablished is not taken into account. Note that the median here means that only 50% of all missed updates are synchronized to all clients. Only at the upper whisker, most of the missed updates are synchronized.

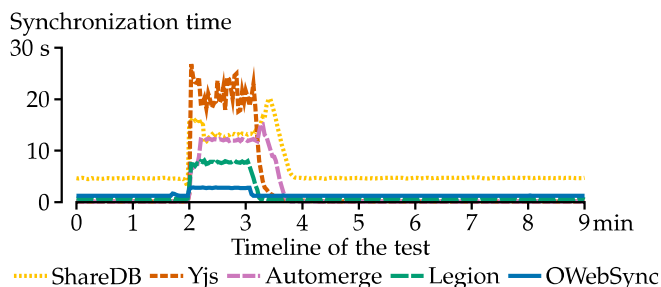


Fig. 8. Mean time to synchronize updates in case of a network disruption between minute 2 and 3 for the test with 24 clients, 1000 objects.

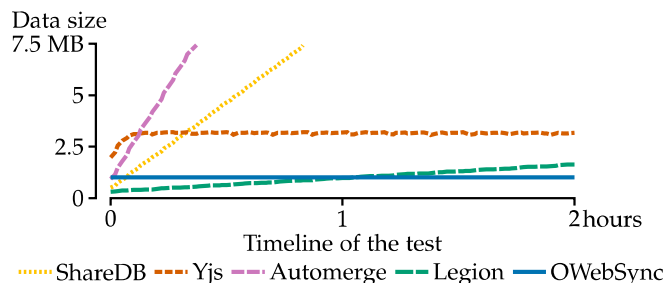


Fig. 9. Evolution of the total data size on the server.

back online, those technologies cannot achieve prompt and interactive synchronization immediately.

5.3 Total size of the data model

The third and last benchmark is used to measure the total size of the data set over a longer time (2 hours). Every 10 minutes, 5 new client browsers will start making changes. After those 10 minutes, the browsers are shut down and replaced by others. After 2 hours, about 60 browsers of clients are introduced into the system. This benchmark simulates the eDesigners case study over the course of a few years. Several employees and external consultants will have worked on the template using different browsers on their devices (desktop, laptop, tablet). In the meantime, they might have cleared their browser cache, used an incognito session or switched to a new device. This scenario is used to verify how well the 5 frameworks will perform over time.

All other technologies used in the evaluation use some form of client identifiers and version numbers to keep track of changes (e.g. vector clocks in Legion). This means that the size of the data set will grow over time, especially in highly dynamic settings like the web. Fig. 9 shows the total data size on the server over time while several users are joining and leaving. The size of the data set on the server remains constant over time when using OWebSync. Other techniques grow with the number of clients and operations. In the dynamic setting of the web, keeping track of all clients

with version vectors and client identifiers will eventually inflate and pollute the metadata. Users can clear the browser cache, browse incognito or visit the web application on multiple devices including someone else's device for one time. By storing those client identifiers in the data model on the server, the performance will decrease over time. Yjs is an exception and stops growing fast in size after a few minutes. This is because Yjs will garbage collect old operations after 100 seconds [9]. This operation is not safe and clients that were offline for a longer time might end up in an inconsistent state or lose data.

The first two benchmarks are performed on a clean data set, meaning that the size of the data on the server is still small. If we would start the tests after e.g. 5 hours of warm-up, the results for the other technologies would be worse. We evaluated a worst-case scenario for OWebSync, with clean data sets for the other frameworks.

5.4 Summary

Our evaluation shows that the operation-based approaches work well in continuous online situations with a limited number of users. Operational Transformation cannot be used with many clients as the server eventually becomes a bottleneck. Operation-based approaches can synchronize updates faster than state-based approaches like Legion and OWebSync. However, when network disruptions occur, these technologies cannot achieve acceptable performance and need tens of seconds to achieve synchronization. Delta-

TABLE 1
Synchronization times in seconds for 24 clients and 1000 objects.

	online		offline	
	50%	99%	50%	99%
ShareDB	4.45	7.69	12.67	25.10
Yjs	0.14	0.17	20.21	109.15
Automerger	0.14	0.20	11.59	18.90
Legion	0.64	1.03	7.61	8.56
OWebSync	1.34	1.49	2.87	3.53

state CRDTs, as used in Legion, can recover faster from network disruptions than operation-based approaches, but still need more than 8 seconds to synchronize missed updates, which cannot be called interactive anymore. Moreover, the size of the data set will increase with both the number of updates and the number of clients. OWebSync can achieve much better performance in the order of seconds, which is still acceptable for interactive groupware. In a setting with frequent offline situations, e.g. for mobile employees, OWebSync is the most appropriate technology and outperforms all other frameworks. Over time, OWebSync can continue to deliver the same interactive performance, as no client identifiers or version vectors are stored. Table 1 summarizes the results in seconds of the large scale test with 24 clients and 1000 objects for the average user (50th percentile) and most of the users (99th percentile) for both settings.

6 RELATED WORK

The related work consists of three types of work: 1) concepts and techniques such as CRDTs and OT, 2) NoSQL data systems such as Dynamo and Cassandra, as well as synchronization frameworks between data centers and 3) synchronization frameworks for replication to the client.

Concepts and techniques. The concepts and techniques like OT and CRDTs were discussed in Section 2. Other text-based versioning systems such as Git [42] are not made to manage data structures and do not always guarantee valid data structures after synchronization. Code, XML or JSON can end up malformed and often require user-level resolution.

We now discuss some other extensions to CRDTs. Conflict-free Partially Replicated Data Types [43] allow to replicate only part of a CRDT. This helps with bandwidth and memory consumption, as well as security and privacy [44]. OWebSync allows replicating any arbitrary subtree of the whole CRDT tree. Hybrid approaches combining operation-based and state-based CRDTs are also possible as demonstrated by Bendy [45]. For data that can tolerate staleness, one can make use of state-based CRDTs, while for data with interactive performance requirements, operation-based CRDTs can be used. This dynamic decision is only made between the servers, and not on the clients. For clients, only operation-based CRDTs are available. A garbage collection technique can be used to reduce the memory usage of operation-based CRDTs by defining a join-protocol for dynamic environments [46]. But this only treats transient network disruptions where clients will come back online eventually, which is not necessarily the case for web clients. Strong Eventually Consistent Replicated Objects (SECROs) [47] are similar to operation-based CRDTs, but

do not impose restrictions on commutativity of operations. However, by doing so, they need a global total order and cannot tolerate network disruptions.

Distributed data systems and NoSQL systems. Based on the Dynamo paper [15], many other open-source NoSQL systems have been developed for structured or semi-structured data, focusing on eventual consistency within or between data centers. Dynamo uses multi-value registers to maintain multiple versions of the data and expects application-level conflict resolution. Cassandra [48], [49] supports fine-grained versioning of cells in a wide-column store. It uses wall-clock timestamps for each row-column cell and adopts a last-write-wins strategy to merge two cells. CouchDB [50] and MongoDB [51] focus on semi-structured document storage, typically in a JSON format. CouchDB offers only coarse-grained versioning per document and stores multiple versions of the document. Applications need to resolve version conflicts manually. It also does not support fine-grained conflict detection within two documents.

Several commercial database systems allow to use CRDTs as the underlying data model: e.g. Riak [12], Akka [52] and Redis [53]. Besides those commercial products, several research projects have emerged. Merkle Search Trees (MSF) [13] implement a key-value store like Dynamo using a state-based CRDT and a Merkle-tree. It builds the Merkle-tree on top of the flat data structure, while OWebSync will make use of the tree-like structure of the data to build the Merkle-tree. MSF is targeted to asynchronous background synchronization between backend servers, and not for interactive groupware with replication to the clients. Antidote [54] is a research project to develop a geo-replicated database over world-wide data centers. It adopts operation-based commutative CRDTs for highly-available transactions and supports partial replication but assumes continuous online connections as the default operational situation for clients. SMAC [55] uses an operation-based CRDT storage system for state management tasks for distributed container deployments. DottedDB [56] uses node-wide dot-based clocks to find changes that need to be replicated, without the need for explicit tombstones. It does not support replication to the clients, or offline edits.

Client-tier frameworks for synchronization. Many client-side frameworks have appeared to enable synchronization between native clients. Cimbiosys [57] is an application platform that supports content-based partial replication and synchronization with arbitrary peers. While it shares some of the goals of OWebSync, it is best suited to synchronize collections of media data (e.g. pictures, movies) and not for JSON documents with fine-grained conflict resolution. SwiftCloud [5], [6], [58] is a distributed object database with fast reads and writes using a causally-consistent client-side local cache and operation-based CRDTs. Metadata used for causality in the form of vector clocks are assigned by the data centers. Hence, the size of the metadata is bound by the number of data centers, and not by the number of updates or the number of clients. The cache is limited in size and the data is only partially available, limiting what data can be read and updated during offline operation. Because it uses operation-based CRDTs, it needs a reliable exactly-once message channel, which is implemented by using a log.

Besides the frameworks for native clients, there are

several JavaScript frameworks made for synchronization between distributed web clients. Legion [17], [18] is a framework for extending web applications with peer-to-peer interactions. It also supports client-server usage and uses delta-state-based CRDTs for the synchronization. Automerge [10], [11] is a JavaScript library for data synchronization adopting the operation-based JSON data type of Kleppman [34]. It uses vector clocks which grow in size with the number of clients. PouchDB [59] is a client-side JavaScript library that can replicate data from and to a CouchDB server. Local data copies are stored in the browser for offline usage. PouchDB only supports conflict detection and resolution at the coarse-grained level of a whole document. ShareDB [33] is a client-server framework to synchronize JSON documents and adopts OT as synchronization technique between the different local copies. ShareDB can thus not be used in extended offline situations. In case of short network disruptions, it can store the operations on the data in memory and resend them when the connection is restored. The offline operations are lost when the browser session is closed. Yjs [7], [8], [9] is a JavaScript framework for synchronizing structured data and supports maps, arrays, XML and text documents. All data types also use operation-based CRDTs for synchronization. Swarm.js [60] is a JavaScript client library for the Swarm database, based on operation-based CRDTs with a partially ordered log for synchronization after offline situations. Swarm.js also focuses on peer-to-peer architectures like chat applications and decentralized CDNs, while OWebSync focuses on client-server line-of-business applications. In contrast with OWebSync, none of these JavaScript frameworks support all of the following: fine-grained conflict resolution, interactive updates when online and fast resynchronization after being offline, as well as being scalable to tens of concurrently online clients and hundreds of writers over time.

7 CONCLUSION

This paper presented a web middleware that supports seamless synchronization of both online and offline clients that are concurrently editing a shared data set. Our OWebSync middleware implements a generic data model, based on JSON, that combines state-based CRDTs with Merkle-trees. This allows to quickly find differences in the data set and synchronize them to other clients. Apart from the regular CRDT structure and the hashes of the Merkle-tree, no other metadata needs to be stored. Existing approaches use client identifiers and version numbers, or even the full history, to track updates, which will pollute the metadata and decrease performance over time.

The comparative evaluation shows that the operation-based approaches cannot achieve acceptable performance in case of network disruptions and need tens of seconds to achieve resynchronization. Current state-based approaches using delta-state-based CRDTs are faster to recover than the operation-based ones, but cannot achieve prompt resynchronization of missed updates. The state-based approach with Merkle-trees of OWebSync can achieve better performance in the order of seconds for both online updates and missed offline updates, making it suitable for interactive web applications and groupware.

REFERENCES

- [1] J. Nielsen, *Usability Engineering*. Nielsen Norman Group, 1993. [Online]. Available: <https://www.nngroup.com/books/usability-engineering/>
- [2] "Google docs," <https://support.google.com/docs/answer/2494822>, 2018.
- [3] N. Preguiça, J. M. Marques, M. Shapiro, and M. Letia, "A commutative replicated data type for cooperative editing," in *2009 29th IEEE International Conference on Distributed Computing Systems*, June 2009, pp. 395–403.
- [4] M. Shapiro, N. Preguiça, C. Baquero, and M. Zawirski, "A comprehensive study of convergent and commutative replicated data types," Inria – Centre Paris-Rocquencourt ; INRIA, Research Report RR-7506, Jan. 2011. [Online]. Available: <https://hal.inria.fr/inria-00555588>
- [5] M. Zawirski, A. Bieniusa, V. Balesgas, S. Duarte, C. Baquero, M. Shapiro, and N. Preguiça, "Swiftcloud: Fault-tolerant geo-replication integrated all the way to the client machine," INRIA, Research Report RR-8347, Oct. 2013. [Online]. Available: <https://hal.inria.fr/hal-00870225>
- [6] N. Preguiça, M. Zawirski, A. Bieniusa, S. Duarte, V. Balesgas, C. Baquero, and M. Shapiro, "Swiftcloud: Fault-tolerant geo-replication integrated all the way to the client machine," in *2014 IEEE 33rd International Symposium on Reliable Distributed Systems Workshops*. IEEE, 2014, pp. 30–33.
- [7] P. Nicolaescu, K. Jahns, M. Dertnl, and R. Klamma, "Yjs: A framework for near real-time p2p shared editing on arbitrary data types," in *Engineering the Web in the Big Data Era*. Cham: Springer International Publishing, 2015, pp. 675–678.
- [8] —, "Near real-time peer-to-peer shared editing on extensible data types," in *Proceedings of the 19th International Conference on Supporting Group Work*, ser. GROUP '16. New York, NY, USA: ACM, 2016, pp. 39–49.
- [9] "Yjs," <https://github.com/y-js/yjs>, 2014.
- [10] M. Kleppman and A. R. Beresford. (2018) Automerge: Real-time data sync between edge devices. [Online]. Available: <http://martin.kleppmann.com/papers/automerger-mobiuk18.pdf>
- [11] "Automerge," <https://github.com/automerger/automerger>, 2017.
- [12] "Riak," <http://docs.basho.com/riak/kv>, 2010.
- [13] A. Auvolat and F. Taïani, "Merkle Search Trees: Efficient State-Based CRDTs in Open Networks," in *SRDS 2019 - 38th IEEE International Symposium on Reliable Distributed Systems*. Lyon, France: IEEE, Oct. 2019. [Online]. Available: <https://hal.inria.fr/hal-02303490>
- [14] R. Merkle, "Method of providing digital signatures," 1982, uS patent 4309569. The Board Of Trustees Of The Leland Stanford Junior University.
- [15] G. DeCandia, D. Hastorun, M. Jampani, G. Kakulapati, A. Lakshman, A. Pilchin, S. Sivasubramanian, P. Voshall, and W. Vogels, "Dynamo: amazon's highly available key-value store," in *ACM SIGOPS operating systems review*, vol. 41(6). New York, NY, USA: ACM, 2007, pp. 205–220.
- [16] A. van der Linde, J. a. Leitão, and N. Preguiça, "Δ-crds: Making δ-crds delta-based," in *Proceedings of the 2Nd Workshop on the Principles and Practice of Consistency for Distributed Data*, ser. PaPoC '16. New York, NY, USA: ACM, 2016, pp. 12:1–12:4.
- [17] A. van der Linde, P. Fouto, J. a. Leitão, N. Preguiça, S. Castiñeira, and A. Bieniusa, "Legion: Enriching internet services with peer-to-peer interactions," in *Proceedings of the 26th International Conference on World Wide Web*, ser. WWW '17. Republic and Canton of Geneva, Switzerland: International World Wide Web Conferences Steering Committee, 2017, pp. 283–292.
- [18] "Legion," <https://github.com/albertlinde/Legion>, 2016.
- [19] T. Bray, "The javascript object notation (json) data interchange format," Internet Requests for Comments, IETF, RFC 7158, 2014. [Online]. Available: <https://www.rfc-editor.org/rfc/rfc7158.txt>
- [20] C. A. Ellis and S. J. Gibbs, "Concurrency control in groupware systems," *SIGMOD Rec.*, vol. 18, no. 2, pp. 399–407, Jun. 1989.
- [21] S. Kumawat and A. Khunteta, "A survey on operational transformation algorithms: Challenges, issues and achievements," *International Journal of Computer Applications*, vol. 3, no. 12, pp. 30–38, Jul. 2010.
- [22] M. Shapiro, N. Preguiça, C. Baquero, and M. Zawirski, "Conflict-free replicated data types," in *SSS 2011 - 13th International Symposium Stabilization, Safety, and Security of Distributed Systems*, ser. Lecture Notes in Computer Science, X. Défago, F. Petit, and

- V. Villain, Eds., vol. 6976. Berlin, Heidelberg: Springer Berlin Heidelberg, Oct. 2011, pp. 386–400.
- [23] P. S. Almeida, A. Shoker, and C. Baquero, “Efficient state-based crdts by delta-mutation,” in *Networked Systems*. Cham: Springer International Publishing, 2015, pp. 62–76.
- [24] —, “Delta state replicated data types,” *Journal of Parallel and Distributed Computing*, vol. 111, no. Supplement C, pp. 162 – 173, 2018.
- [25] “Delta crdts,” <https://github.com/peer-base/js-delta-crdts>, 2018.
- [26] V. Enes, C. Baquero, P. S. Almeida, and A. Shoker, “Join decompositions for efficient synchronization of crdts after a network partition: Work in progress report,” in *First Workshop on Programming Models and Languages for Distributed Computing*, ser. PMLDC ‘16. New York, NY, USA: ACM, 2016, pp. 6:1–6:3.
- [27] V. Enes, P. S. Almeida, C. Baquero, and J. Leitão, “Efficient synchronization of state-based crdts,” in *2019 IEEE 35th International Conference on Data Engineering (ICDE)*, April 2019, pp. 148–159.
- [28] M. Shapiro, *Replicated Data Types*. New York, NY: Springer New York, 2017, pp. 1–5.
- [29] H.-G. Roh, M. Jeon, J.-S. Kim, and J. Lee, “Replicated abstract data types: Building blocks for collaborative applications,” *Journal of Parallel and Distributed Computing*, vol. 71, no. 3, pp. 354–368, 2011.
- [30] I. Hickson, “The websocket api, w3c candidate recommendation,” Tech. Rep., 2012. [Online]. Available: <https://www.w3.org/TR/2012/CR-websockets-20120920/>
- [31] P. Leach, M. Mealling, and R. Salz, “A universally unique identifier (uuid) urn namespace,” Internet Requests for Comments, RFC 4122, 2005. [Online]. Available: <https://www.rfc-editor.org/rfc/rfc4122.txt>
- [32] R. Rivest, “The md5 message-digest algorithm,” Internet Requests for Comments, RFC 1321, 1992. [Online]. Available: <https://www.rfc-editor.org/rfc/rfc1321.txt>
- [33] “Sharedb,” <https://github.com/share/sharedb>, 2013.
- [34] M. Kleppmann and A. R. Beresford, “A conflict-free replicated json datatype,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 28, no. 10, pp. 2733–2746, 2017.
- [35] “Fabric.js,” <https://github.com/fabricjs/fabric.js>, 2011.
- [36] “Azure,” <https://azure.microsoft.com>, 2019.
- [37] W. Almesberger, “Linux network traffic control – implementation overview,” EPFL, Tech. Rep., 1999. [Online]. Available: <https://www.almesberger.net/cv/papers/tcio8.pdf>
- [38] “opensignal.com,” <https://www.opensignal.com/reports/2019/01/usa/mobile-network-experience>, 2019.
- [39] Q.-V. Dang and C.-L. Ignat, “Performance of real-time collaborative editors at large scale: User perspective,” in *Internet of People Workshop, 2016 IFIP Networking Conference*, ser. Proceedings of 2016 IFIP Networking Conference, Networking 2016 and Workshops. Vienna, Austria: IFIP, May 2016, pp. 548–553.
- [40] “Speedtest.net,” <http://www.speedtest.net/reports/united-states/2018/Mobile/>, 2018.
- [41] J. Nielsen. (2010) Website response times. [Online]. Available: <https://www.nggroup.com/articles/website-response-times/>
- [42] “Git,” <https://git-scm.com/>, 2005.
- [43] I. Briquemont, M. Bravo, Z. Li, and P. Van Roy, “Conflict-free partially replicated data types,” in *2015 IEEE 7th International Conference on Cloud Computing Technology and Science (CloudCom)*. IEEE, 2015, pp. 282–289.
- [44] S. A. Kollmann, M. Kleppmann, and A. R. Beresford, “Snapdoc: Authenticated snapshots with history privacy in peer-to-peer collaborative editing,” *Proceedings on Privacy Enhancing Technologies*, vol. 2019, no. 3, pp. 210 – 232, 2019.
- [45] C. Bartolomeu, M. Bravo, and L. Rodrigues, “Dynamic adaptation of geo-replicated crdts,” in *Proceedings of the 31st Annual ACM Symposium on Applied Computing*, ser. SAC ‘16. New York, NY, USA: ACM, 2016, pp. 514–521.
- [46] J. Bauwens and E. Gonzalez Boix, “Memory efficient crdts in dynamic environments,” in *Proceedings of the 11th ACM SIGPLAN International Workshop on Virtual Machines and Intermediate Languages*, ser. VMIL 2019. New York, NY, USA: ACM, 2019, pp. 48–57.
- [47] K. De Porre, F. Myter, C. De Troyer, C. Scholliers, W. De Meuter, and E. Gonzalez Boix, “Putting order in strong eventual consistency,” in *Distributed Applications and Interoperable Systems*. Cham: Springer International Publishing, 2019, pp. 36–56.
- [48] “Apache cassandra,” <https://cassandra.apache.org>, 2009.
- [49] A. Lakshman and P. Malik, “Cassandra: a decentralized structured storage system,” *ACM SIGOPS Operating Systems Review*, vol. 44, no. 2, pp. 35–40, 2010.
- [50] “Couchdb,” <https://couchdb.apache.org>, 2005.
- [51] “Mongodb,” <https://www.mongodb.com/>, 2009.
- [52] “Akka,” <https://doc.akka.io/docs/akka/current/distributed-data.html>, 2018.
- [53] C. Biyikoglu, “Under the hood: Redis crdts (conflict-free replicated data types),” Redis Labs, White paper, 2017. [Online]. Available: <https://redislabs.com/docs/under-the-hood/>
- [54] “Antidote,” <http://syncfree.github.io/antidote>, 2014.
- [55] J. Eberhardt, D. Ernst, and D. Bernbach, “Smac: State management for geo-distributed containers,” Technische Universitaet Berlin, Tech. Rep., 2016.
- [56] R. J. T. Gonçalves, P. S. Almeida, C. Baquero, and V. Fonte, “Dotteddb: Anti-entropy without merkle trees, deletes without tombstones,” in *2017 IEEE 36th Symposium on Reliable Distributed Systems (SRDS)*. IEEE, 2017, pp. 194–203.
- [57] V. Ramasubramanian, T. L. Rodeheffer, D. B. Terry, M. Walraed-Sullivan, T. Wobber, C. C. Marshall, and A. Vahdat, “Cimbiosys: A platform for content-based partial replication,” in *Proceedings of the 6th USENIX symposium on Networked systems design and implementation*, 2009, pp. 261–276.
- [58] M. Zawirski, N. Preguiça, S. Duarte, A. Bieniusa, V. Balesgas, and M. Shapiro, “Write fast, read in the past: Causal consistency for client-side applications,” in *Proceedings of the 16th Annual Middleware Conference*, ser. Middleware ‘15. New York, NY, USA: ACM, 2015, pp. 75–87.
- [59] “Pouchdb,” <https://pouchdb.com>, 2013.
- [60] “Swarm.js,” <https://github.com/gritzko/swarm>, 2013.



Kristof Janne is a Ph.D. candidate in the Department of Computer Science at KU Leuven in Belgium, and a member of the research group imec-DistriNet. His research activities are under the supervision of Prof. Dr. Wouter Joosen and Dr. Bert Lagaisse. He received his Master's degree in computer science from the KU Leuven in 2018. His main research interests are in the area of data synchronization, consensus and decentralization.



Bert Lagaisse is a senior industrial research manager at the imec-DistriNet research group in which he manages a portfolio of applied research projects on cloud technologies, distributed data management and security middleware in close collaboration with industrial partners. He has a strong interest in distributed systems, in enterprise middleware, cloud platforms and security services. He obtained his MSc in computer science at KU Leuven in 2003 and finished his Ph.D. in the same domain in 2009.



Wouter Joosen is full professor at the Department of Computer Science of the KU Leuven in Belgium, where he teaches courses on software architecture and component-based software engineering, distributed systems and the engineering of secure service platforms. His research interests are in aspect-oriented software development, focusing on software architecture and middleware, and in security aspects of software, including security in component frameworks and security architectures.

January 7, 2021

Manish Parashar, EIC
IEEE Transactions on Parallel and Distributed Systems

Dear Mr. Parashar,

Re: Major Revision Decision Re: TPDS-2020-02-0106

We would like to thank all reviewers for their detailed feedback on our manuscript. We feel their comments enabled us to improve the quality of our manuscript, and appreciate their time and effort.

Attached to this letter, you will find an improved version of our manuscript. This letter also provides a detailed response to the comments of the reviewers.

If you require any additional information, please do not hesitate to contact us.

Yours Faithfully,

Kristof Jannes
Bert Lagaisse
Wouter Joosen

Detailed answers to reviewer comments and questions

Below, we have provided detailed answers to the individual comments provided by the reviewers, and how we solved their concerns in the revised manuscript. We hope that these responses and improvements address the concerns the reviewers have about the original manuscript, and welcome any additional feedback that can improve the manuscript further.

Reviewer: 1

"This paper presents OWebSync, a system to synchronize state in collaborative web-based applications. The paper addresses an important problem, is of interest to the TPDS audience and in general it is easy to follow even though the organization could be improved."

Response: We have improved the organisation in the revised manuscript. Section 1 now describes the contributions of the paper only after presenting the existing approaches. Section 2 now only focuses on the motivation and background, and does not include the approach anymore. The background that was presented in Section 3 is now moved to Section 2. In Section 5, we have moved the description of each benchmark closer to the actual results and discussion of the respective benchmark.

Detailed comments:

- *"The intro motivates the paper well, but the organization could be improved. The contributions of the paper (briefly discussed in the paragraph "In this paper, we present..." would be better suited to appear after discussing the existing approaches, right before the last paragraph of the intro. that presents the roadmap"*
Response: We agree that this placement is more intuitive for the reader, and have moved the paragraph accordingly.
- *"When first an acronym such as CRM or HR, the full name should be given"*
Response: We thank the reviewer for pointing this out, and have written the acronyms in full in the revised manuscript.
- *"§2.1.2 when there are conflicting edits, what's the right decision from the application's perspective? How does this maps to the underlying data structures?"*
Response: A great question! We focus on usability, without user interaction or too much overhead for the developer. There are however two choices left for a developer. They can either pass the JSON object structure in full, and every JSON map will be mapped to an ORMap, and the leaf values to LWWRegisters. This results in a very fine-grained conflict resolution. The other option for a developer is to *stringify* an object to a String, so that full object is mapped to a LWWRegister. Now conflict resolution is less fine-grained, and instead of merging two changes to the same object, only one change is kept (the latest one). This allows developers to have a kind of atomic changes.
- *"§2.2.1 "Hence, OT is not resilient against message loss in offline situations". This is confusing as, by definition, when offline are messages are lost"*
Response: We agree that is wording is unfortunate and changed the sentence accordingly in the revised manuscript.
- *"§2.2.4 discussed the approach taken in the paper. This would be better suited for the next section, that introduces the system and data model"*
Response: We fully agree, and we have moved this section to the next section.
- *"§3.1 and §3.2 contains some general background about data structures, this would fit better in §2"*
Response: All background material in Section 3 has now been moved to Section 2.
- *"The discussion of the ORMap should provide some context (maybe some examples?) to make the paper more accessible to a wider audience"*
Response: We have extended the manuscript with a more elaborate description of the ORSet in Section 2 to make the paper more accessible, and extended the description of the ORMap in Section 3.
- *"First, we extended this data structure with a Merkle-tree that we pattern over the object's logical tree-structure." What's the meaning of "pattern over"?"*
Response: We agree that the wording is unclear. We mean that we extended this data structure with a Merkle-tree using the object's logical tree-structure. We have also changed the wording in the revised manuscript.
- *"A figure with an example of the data structures, for instance for the example already provided in the paper,*

would be rather useful to help the reader better understand the concepts being discussed"

Response: We agree that the current example in Fig. 1 and Fig. 2 is hard to follow, and we have improved the notation in the figures, and extended the description in the revised manuscript to provide a more precise example.

- "Algorithm 1 (should be an Alg. not a "Figure") needs a more detailed explanation both in the text and as comments in the pseudo-code itself. Understanding all the details of the Merge procedure, in particular, can be challenging"

Response: We have improved the notation of the pseudo-code and slightly extended it, as well as extended its explanation in Section 3.2. We also renamed the Fig. to Alg. .

- "In Alg. 1 the definition of the KV is missing (this is stated in the text but the alg. should be self-contained)"

Response: We have addresses this by adding the definition of the KV in Alg. 1.

- "In the line of a previous comment, a figure depicting what's happening in Fig.2 and Fig.3 would be rather useful. How does each entry in the JSON maps to the tree?"

Response: This indeed may require more elaboration in the paper. However, due to space considerations we did not include a figure, but we have extended Section 3.2, paragraph *Example* to describe this mapping in the text.

- "Even though the paper assumes non-malicious users, MD5 is rather weak as there are several semi-automated tools to generate collisions for arbitrary content. Why not select a stronger hash function?"

Response: We have opted for MD5 mostly for performance reasons, it is a very fast hash, and collisions without malicious users are very rare. It is possible to switch the MD5 implementation to a stronger hash, however, to defend against malicious users we would need to make significant changes to other parts as well.

- "\$4.3 the GET returns the CRDT if the hash is different from its own CRDT at the given path. What happens otherwise? How does the receiver distinguishes between no changes and a message lost/delayed? Same applies to the REMOVE"

Response: OWebSync does not send individual messages, but instead collects all messages in a list, and send them as one message. If, after processing the list of received messages, there are no messages to send in response, and empty list is send. This signals the end of the synchronization protocol to the other side. This is also discussed in the manuscript in Section 4.3.

- "The synchronization starts with the highest CRDT in the tree" Meaning the root?"

Response: We indeed mean the root, and have updated the manuscript accordingly.

- "\$4.3 "If during a merger process, a child..." This paragraph is rather confusing. What "sides" are being discussed? Client and server side? Two distinct clients? The text should be improved and the problem being addressed should be defined more precisely"

Response: With sides we mean the client and the server side. We have improved the wording in the text. We tried here to sketch a scenario where sending a REMOVE message is more beneficial for the bandwidth usage than sending a PUSH message.

- "In the evaluation, the text discussing the results and the plots themselves are quite far apart, reorganizing the text so that it's closer to the plots would improve the readability of the text. Also it would be more useful to introduce each benchmark and the results rather than presenting first all benchmarks and then all the results. I've found myself going back and forth in the paper trying to understand the results in light of the benchmark being used"

Response: We agree that this placement is more intuitive for the reader, and we have moved the description of each benchmark closer to the actual results and discussion of the respective benchmark, instead of presenting them all up front.

- "Regarding the results some things are not clear. "The synchronization times of the succeeded updates are illustrated in ...". Does this means that some updates might fail causing clients to diverge? If so on which techniques and what is the impact of this in the results and application?"

Response: Some update operations will conflict with each other because two different clients move the same object concurrently. Only one of the two updates will be fully propagated to all other clients, and the time for the other update is not taken into account, as there is never a time where that update is fully replicated to all clients. We have improved the wording in the revised manuscript.

- "What percentiles to the whiskers cover? The sentence "Only at the upper whisker, all of the missed updates are fully synchronized" seems to contradict the caption "Only at the upper whisker, most of the missed updates are synchronized". It is the 100th percentile? The 99th?"

Response: Whiskers cover a maximum of 1.5 times the interquartile range. The outliers are shown as dots. The sentences are indeed contradicting this statement, and we have fixed this in the revised manuscript.

- *“What is the shape of the tree, depth and fanout, being used in OWebSync? What’s the impact of different shapes in the results? How is the shape decide in the first place? Can the application have some control over this? This seems to have a significant impact in the results hence evaluating some possible configurations or at least discussing their impact would benefit the paper”*

Response: The total depth of the three is four (root – drawing1 – object36 – top). The root has one child, while drawing1 has either 100 or 100 children. And these children have itself each 7 children. This is the shape of the tree defined by the application, however, because drawing1 has many children, OWebSync will transparently add an extra layer in the tree to reduce the network usage. This increases the total depth of the tree to 5. We have clarified this in the revised manuscript in Section 5.1.

- *The key of Fig. 6 is not very legible, it’d be better to place the name of the system below each respective whisker plot. Same applies for the other plots*

Response: We agree that this plot was not very readable, so we removed the black border around the boxplots, this way the color stays visible, even for the very small boxplots.

- *“full synchronization to all clients in the online scenarios” What does this means? It is unclear how time is measured. Is it the time for all clients to sync after all the updates finish? Is this realistic? If so, what does this means for the case of offline clients? Do they all come back after the updates finish just to synchronize? If that’s the case I don’t think this is a very realistic measurement, as the measurement should be taken while the system is processing other updates.”*

Response: The time measured for an update in the online scenario, is the time between one client making that update, and all other clients receiving that update in the application. During the whole benchmark, each client will make one update per second, while synchronizing updates.

For the offline scenario, we only measure the time starting from the end of the offline period, until all other clients have delivered the update to the application. During the whole offline period, and also afterwards when they come back online, all clients make one update every second.

- *“What are the assumptions regarding clock synchrony to perform this measurement?”*

Response: We assume loosely synchronized clocks. VMs in Azure have their clocks synchronized with the host machine, which is synchronized with the internal Microsoft time servers. We have clarified this in the revised manuscript.

- *“What is the impact of each technique in the server resources? It is mentioned that OT techniques are more resource intensive, but what about the other systems? Is OWebSync more efficient on server resources than the other approaches? This is a key comparison to be able to fully access the practicality of OWebSync.”*

Response: We did not discuss the impact of OWebSync on server CPU in the manuscript. We give here the CPU results for the largest scale test scenario with 1000 objects and 24 clients. OWebSync has an average server-side CPU of 33%. Legion and Yjs do better with respectively 5% and 10%. Automerge and ShareDB need more resources: 54% and 101% respectively.

However, we do not think that server resources are the most important problem, in our opinion, the client-side CPU is much more important. The average client-side CPU usage in the same experiment is: 9% for OWebSync, 9% for Legion, 10% for ShareDB, 20% for Yjs and 22% for Automerge.

We have left out this discussion on CPU usage for space considerations.

Reviewer: 2

“The problem of the paper is well-motivated and the background is explained precisely and succinctly. It is convincing that the paper aims at making a significant contribution to the state of the art, with relevant practical applications, as illustrated by two use cases.”

“My main concerns, however, are about the Chapter 3 (data model). While probably none of the issues are a complete show-stopper, significant additional work is required to make that part of the paper unambiguous, convincing, and correct. The main issue are errors in the core ORMap CRDT implementation. It is under-specified, contains errors, and lacks any kind of (even informal) attempt to proof that it is actually correct.”

Response: We would like to thank the reviewer for pointing out the inconsistencies in the ORMap specification. We have corrected and extended the pseudo-code in the revised manuscript, as well as improved the notation. We also added a proof sketch near the end of Section 3.2, in which we argue that despite our changes, the presented ORMap is still a valid CvRDT.

*“A fundamental concern is that the authors claim that their ORMap is based on Shapiro’s ORSet, but in fact it is fundamentally different. A usual ORSet checks, in its GET function, if some key is in the Removed set. The authors’ version does not do so, without explanation. In an ORSet, a key may appear more than once (with different unique tags) in the Observed set (if added by multiple nodes independently). Checking for set inclusion is obviously easy, but if you implement a map instead of set, with key-value pairs, you would end up with multiple values for a single key (and different tags), and deciding which value is the “right” one is not obvious. The authors do not address that problem, may because their “ORMap” is fundamentally different from a usual ORSet, but fail to explain that part. Also, ORSet’s delete function deletes *all* instances of a key (all tags), whereas the code in Fig 1 apparently removes only one. The evident differences are significant, and it is certainly not obvious that the proposed ORMap data type is correct. At least some better justification why this data type is a correct CRDT should be provided”*

Response: We agree that we did not rigorously specify our ORMap. We have rewritten a large part of Section 3 to make it clearer and correct errors. The basic explanation of Shapiro’s ORSet is moved to Section 2. Our ORMap is based on Shapiro’s ORSet, with some differences, which are now described in detail in Section 3.2.

In our ORMap, there can only be one item of a key. If different replicas have different items with the same key, but a different tag, the child elements would be merged, and the lexicographic largest tag is assigned to the result.

“I was trying to understand the ORMap better by reading their pseudo code implementation (Fig. 1), but this a difficult task as well. They use a “JOIN” function in line 42 and line 46 that is never defined or explained. Possibly this is meant to be a recursive call to the MERGE function, but I am not sure about this. Maybe some other magic is happening in that undefined function that addresses some the issue I raise here? Hard to tell, if he paper lacks that information.”

Response: We thank the reviewer for pointing this out. We indeed mean a recursive call to the MERGE function, and we have corrected this in the revised manuscript.

“An important case to be handled by CRDTs is resolving conflicts. Lets consider a very simple case, two nodes concurrently add items with the same key A and different values (1, 2), and different tags. So node 1 has (tag1, A, 1) in its Observed set, and the remote node has (tag2, A, 2) and sends a state update. From what I get from Fig. 1, if a key already exists at node 1 (e.g., line 34, line 40), the local observed set is never changed (the only place this happens is in line 47, which is only reached for a key not locally existing), so based on the code in Fig. 1, tag2 will never make it into node 1’s Observed set. In the other direction, by the same argument, the remote node will never add tag1 to its Observed set. While the values (1,2) might be merged by the leaf-level LWWRegister based on their timestamps, the tag in the Observed sets will remain inconsistent after merge. If a node is deleted (by tag id) on one side, it will not be deleted on the other side, resulting in observable difference.”

Response: We agree that this is a scenario we should discuss and have added this to the revised manuscript. In our ORMap, there cannot be multiple items with the same key. If different replicas would have a different item with the same key, we will merge those children according to their respective rules (either LWWRegister or ORMap) and only keep the lexicographic largest tag for the result. This has some drawbacks when you add two items for the same on different replicas, and remove one of them quickly after adding it. The removed item can resurface again (with the tag of the other item), if the two different items are already merged on some replicas. However, eventually all replicas will have the same state.

“Some additional minor notes on Section 3:”

- *“The embedded Merkle tree requires calculating hashes over sets. These hashes needs to be precisely defined in a deterministic way (a set, on its own, has no order, so a simple way of calculating a hash over the (arbitrarily ordered) concatenation of the content is not deterministic. A deterministic hash function is a prerequisite, of course easy to implement, but I consider it important to specify this explicitly in the paper.”*

Response: When calculating hashes over sets, we order them lexicographically on the unique tags. We also added this explanation to the revised manuscript.

- *“One possibly critical assumption of authors is that there are no faulty clients. But not only malicious clients are an unsolved problem, also accidental faults (maybe a client has a local clock that is 10 years off into the future), resulting in a state that cannot be modified for years... This seems to be a real practical problem, and a discussion if such problem can somehow reasonably be mitigated would be a valuable contribution.”*

Response: We agree that this is a worthwhile discussion and we have added it to Section 3.3. We assume loosely synchronized clocks. If this assumption is not met, an accidental fault resulting in a clock several years in the future, might make edits to this LWWRegister impossible. Users can resolve this manually by removing that register and creating a new one with the same key in the ORMap, loosing concurrent

changes. We do not consider this an important drawback, as most personal devices these days automatically synchronize their time with the internet.

- “Sec. 3.2 explains that “when a value is removed, *it* is added to the Removed set”, which confused me (what is added? the value of the key/value pair? or rather the key? actually its a random ID associated with the key/value pair). This becomes clearer much later when looking at the code (Fig. 1), but I suggest being less ambiguous right at the beginning of this section.”

Response: When that item is removed, the associated tag is added to the removed set. We have made this clearer in the revised manuscript at the end of Section 2.

- Related to previous comments, also the SET function (in Fig. 1) does not consider the Removed set. If a key’s tag (id) is in the Removed set, its not correct to just update the value. Instead a new entry in the observed set with a new tag needs to be generated (at least that is also what the authors explain at the beginning of Section 3.2)

Response: In our implementation we do not keep the removed elements in the ORMap, only the tag is kept in the removed set. As a result, we do not need to consider the removed set in the SET function. We agree that this was not clear in the original manuscript and we have made clarifications in the revised manuscript.

- “p5 s/needs to be send/needs to be sent/”

Response: We thank the reviewer for pointing this out, and have corrected it in the revised manuscript.

“Section 4 (synchronization protocol):”

- “While the description appears to be valid, the protocol is presented in a very informal way only. A more rigorous specification should be provided.”

Response: We have addressed this by adding a pseudo-code specification of the protocol in Alg. 3.

- “I don’t agree with the statement “the number of messages, and thus the length of the synchronization protocol, is therefore limited to the maximum depth of the Merkle-tree”. This is obviously true if you consider a change to a single leaf in the tree, as you need to descent a single path in the tree from root to leaf. But the way the statement is made it appears as a statement about an arbitrary merge. If you have to merge two significantly modified trees, I am assuming you will have many more messages. For example, in Fig. 5, if in drawing1 all objects object1..object100 have changed, for each of these 100 objects a push message needs to be sent in step 3, and each of that messages can trigger a PUSH or GET message from client to server in step 4.”

Response: We would like to thank the reviewer for pointing out this mistake. It is indeed not the number of messages that is limited by the depth of the tree, but rather the length of the synchronization process. We have corrected this in Section 4.3 in the revised manuscript.

“Section 5 (evaluation): The evaluation is, over all, well explained and convincing. Fig. 6, however, is somewhat unfortunate because most boxes in the box-plot are so small that their color is not visible.”

Response: We have addressed this issue by removing the black borders around the box-plots, to make the colors better readable when they are small.

“References: Reference [35] is incomplete (lacks proceedings where it was published) Reference [51] is incomplete (also lacks information where it was published)”

Response: We have corrected these references and made it clear that these are online resources, rather than peer-reviewed publications.

Reviewer: 3

“In overall the paper is well written and does a good job of comparing with other similar systems that are available for JS based middleware. The authors try to show the merits and limitations of the chosen approach. One interesting contribution they have is to show the trade-offs among state based and operation based designs when considering the likelihood of partitions, and confirms that operation based approaches are lighter when communication is available but are less efficient when recovering from partitions.”

“My main issue with the proposal is the use of a set of tombstones for the removed IDs since this set will likely grow linearly in workloads where deletions are common. Notice that in state based delta CRDTs it is possible to track causality in a compact way for a recursive map and avoid tombstones. This choice in the present paper might be compensated since a remove wins semantics is also chosen, and deleting sub-trees could possibly get rid of inner

tombstones. In any case, it would be valuable for this paper to discuss this issue in more detail and possibly add an workload that is more heavy in deletions and tombstone creation."

Response: We would like to thank the reviewer for taking time and effort to present helpful feedback about the manuscript, and agree that this indeed requires more elaboration in the paper. We have extended Section 3.3 in the revised manuscript to include a more extensive discussion on tombstones. In short, we argue that tombstones are a better trade-off for the web, as the size of the metadata is not dependent on the number of clients that are making changes. Users can clear their browser cache, or use another browser, or another device; generating many client identifiers over time.

"Some details:"

- "1. Introduction.
"replicas eventual consistent" -> "replicas eventually consistent"
Response: We thank the reviewer for pointing this out, and have corrected it in the revised manuscript.
- "Its important to clarify the different approaches to delta CRDTs in [17,24] and [18] since although sharing some terminology they differ. As far as I can tell Legion uses [18] and the recursive map with common version vectors is in [24]."
Response: This is a good observation, and we have corrected the references. A discussion on the difference between Δ - and δ -CRDTs can be found in Section 2.
- "3.1. LWW
Last-writer-win register are not always consistent with causality when clocks get out of sync. A better design is probably a multi-value register (also observed remove) and use of clock timestamp only to choose the value to present among the possibly concurrent values. This is not an issue that requires change, bit its important to signal the drawback."
Response: We agree that a discussion on the drawbacks of timestamps valuable, and we have added it in Section 3.3 of the revised manuscript. In short, we focused on ease-of-use for end-users by automating the full process of merging replicas, with occasional causal errors which can be solved manually as a trade-of.
- "3.2. OO-Map
I found the use of path[0] and path[1..] hard to follow, maybe clarify notation or add a small concrete example when explaining the implementation."
Response: We agree that the notation could be clearer, and we updated the notation for paths to $p_0..p_n$ to present an array of the different parts of a path. For example for path `drawing1.object36.top`, this would map to $p_0 = \text{drawing1}$, $p_1 = \text{object36}$, $p_2 = p_n = \text{top}$.