

# Collaborative Large-Scale Dense 3D Reconstruction with Online Inter-Agent Pose Optimisation

Stuart Golodetz, Tommaso Cavallari, Nicholas A. Lord,  
Victor A. Prisacariu, David W. Murray and Philip H. S. Torr



Fig. 1: Globally consistent reconstructions produced by our approach, based on the *Flat*, *House* and *Lab* subsets of our dataset.

**Abstract**—Reconstructing dense, volumetric models of real-world 3D scenes is important for many tasks, but capturing large scenes can take significant time, and the risk of transient changes to the scene goes up as the capture time increases. These are good reasons to want instead to capture several smaller sub-scenes that can be joined to make the whole scene. Achieving this has traditionally been difficult: joining sub-scenes that may never have been viewed from the same angle requires a high-quality camera relocaliser that can cope with novel poses, and tracking drift in each sub-scene can prevent them from being joined to make a consistent overall scene. Recent advances, however, have significantly improved our ability to capture medium-sized sub-scenes with little to no tracking drift: real-time globally consistent reconstruction systems can close loops and re-integrate the scene surface on the fly, whilst new visual-inertial odometry approaches can significantly reduce tracking drift during live reconstruction. Moreover, high-quality regression forest-based relocalisers have recently been made more practical by the introduction of a method to allow them to be trained and used online. In this paper, we leverage these advances to present what to our knowledge is the first system to allow multiple users to collaborate interactively to reconstruct dense, voxel-based models of whole buildings using only consumer-grade hardware, a task that has traditionally been both time-consuming and dependent on the availability of specialised hardware. Using our system, an entire house or lab can be reconstructed in under half an hour and at a far lower cost than was previously possible.

**Index Terms**—Collaborative, large-scale, dense 3D reconstruction, inter-agent relocalisation, pose graph optimisation

## 1 INTRODUCTION

Reconstructing dense, volumetric models of real-world 3D scenes is an important task in computer vision and robotics, with applications in content creation for films and games [28], augmented reality [53], cultural heritage preservation [57] and building information modelling [37]. Since the seminal KinectFusion work of Newcombe et al. [38], which demonstrated real-time voxel-based reconstruction of a desk-sized scene in real time using only a consumer-grade RGB-D sensor, huge progress has been made in increasing the size of scene we are able to reconstruct [39, 42, 43, 51] and in compensating for tracking drift [17, 20, 29, 54, 55]. However, even with the most sophisticated voxel-based approaches [17, 29], capturing the data needed to reconstruct a large scene (e.g. a whole building) can take significant time and planning, and require considerable concentration on the part of the user. Moreover, the risk of transient changes to the scene (e.g. people moving around) goes up as the capture time increases, corrupting the model

and forcing the user to restart the capture. There are thus good reasons to want to split the capture into several shorter sequences, which can be captured either over multiple sessions or in parallel (by multiple users) and then joined to make the whole scene.

Achieving this has traditionally been difficult: joining the sub-scenes requires the ability to accurately determine the relative transformations between them (a problem that can be expressed as camera relocalisation), even though the areas in which they overlap may never have been viewed from the same angles, and tracking drift in each sub-scene can prevent them from being joined to make a consistent overall scene. Recent advances, however, have significantly improved our ability to capture consistent, medium-sized sub-scenes, e.g. by closing loops and re-integrating the scene surface on the fly [17], which yields accurate poses for individual frames once loops have been closed, or by combining visual and inertial cues using an extended Kalman filter [27] to achieve accurate camera tracking during live reconstruction. Moreover, in RGB-D relocalisation, keyframe-based relocalisers such as random ferns [23], which were previously widely used for relocalisation in a single-user context but were unable to relocalise from novel poses, have recently been giving way to high-quality regression-based methods such as SCoRe forests [49, 52], driven by recent work [5] that showed how they could be trained and used online. Unlike keyframe-based methods, such approaches have been shown to be much better suited to relocalisation from novel poses, which is critical when aligning sub-scenes captured from different angles.

In this paper, we leverage these advances to present what to our knowledge is the first system to allow multiple users to collaborate

- 
- This work was done whilst all authors were at the University of Oxford.
  - S. Golodetz, T. Cavallari and N. Lord assert joint first authorship. They were previously with the University of Oxford, and are now with FiveAI Ltd. E-mail: {stuart,tommaso.cavallari,nick}@five.ai.
  - V. Prisacariu, D. Murray and P. Torr are with the University of Oxford. E-mail: {victor,dwm,phst}@robots.ox.ac.uk.

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxxx




Fig. 2: A globally consistent collaborative reconstruction of a three-storey house, as produced by our approach. The reconstruction involved relocalising the 6 *Priory* sequences from our dataset.

*interactively* to reconstruct dense, voxel-based models of whole buildings. Unlike previous multi-agent mapping approaches [6, 36], our approach is able to reconstruct detailed, globally consistent dense models. Moreover, using our system, an entire house or lab can be captured and reconstructed in under half an hour using only consumer-grade hardware (see §4.4), providing a low-cost, time-efficient and interactive alternative to existing panorama scanner-based methods for whole-building reconstruction. For example, capturing a building using a Matterport scanner, as done in [2], involves scanning each 2–3m segment of the building separately (taking around 1–2 minutes per segment) and combining the results, a process that can ultimately take many hours. Alternatives such as a Faro 3D laser range finder, as used in [56], are similarly laborious, with each (large) segment taking up to an hour to reconstruct. Moreover, in both cases, the scanners themselves are quite expensive in comparison to consumer-grade cameras. Our approach, by contrast, is inexpensive in terms of hardware, and allows multiple users to work together interactively to reconstruct the space, significantly reducing the time involved and obviating the need for users to repeatedly move a scanner to reconstruct different segments of a building. A comparison of our approach’s properties to those of current methods is shown in Table 1. We achieve a 16-fold increase in scale compared to state-of-the-art low-cost dense solutions such as InfiniTAM [29] and BundleFusion [17], making ours the first low-cost, interactive approach that can produce dense reconstructions of this size, and bringing such methods within range of the much more expensive and *offline* alternatives that are currently employed.

We have integrated our approach into the open-source SemanticPaint framework [24, 25], making it easy for existing SemanticPaint and InfiniTAM [43] users to benefit from our work, and constructed a new dataset of sub-scenes to demonstrate its effectiveness. We make both our source code and this dataset available online. Figures 1 and 2 show reconstructions produced by our approach.

## 2 RELATED WORK

Although previous work on multi-agent mapping has not focused on dense reconstruction, multi-agent mapping itself has a rich research history in computer vision and robotics, and several good surveys exist [46, 47]. Existing approaches occupy two main categories:

(i) *Decentralised* approaches eschew the use of a central server and instead produce a local map of the scene on each agent, often transmit-

	[2]	[56]	[29]	[17]	Ours
Collaborative	✓	✓	✗	✗	✓
Interactive	✗	✗	✓	✓	✓
Low-Cost Hardware	✗	✗	✓	✓	✓
Largest Reconstruction (m <sup>2</sup> )	1900	10370	50	50	820

Table 1: A comparison of our approach to state-of-the-art methods for indoor building reconstruction. ‘Largest Reconstruction’ refers to the size of the largest reconstruction demonstrated in the relevant papers. For our method, 820m<sup>2</sup> refers to a collaborative reconstruction of the *Lab* subset of our dataset, and was obtained by cross-referencing our reconstruction with an official floorplan.

ting these local maps between agents when they meet to share knowledge of parts of the scene that individual agents have not yet visited. For example, Cunningham et al. [15] proposed an approach called DDF-SAM, in which each robot produces a landmark-based map and shares compressed, timestamped versions of it with neighbouring robots. This was extended by [16], which registered landmark-based maps together using an approach based on Delaunay triangulation and RANSAC. The same lab later proposed DDF-SAM 2.0 [14], which avoids repeated and expensive recreation of the combined neighbourhood map through the use of ‘anti-factors’. Cieslewski et al. [11] presented a sophisticated decentralised collaborative mapping back-end based on distributed version control. Choudhary et al. [9] performed decentralised mapping using object-based maps, decreasing the bandwidth required for map sharing, but depending on the existence of pre-trained objects in the scene. Most recently, Cieslewski et al. [10] aimed to minimise the bandwidth each robot uses for inter-agent relocalisation by first establishing which other robots have relevant information, and then only communicating with the best robot found. Decentralised approaches have numerous applications, including search and rescue [35], agricultural robotics [7], planetary exploration [3] and underwater mapping [40, 50], but because of the limited computing power that tends to be available on mobile agents, most approaches target robustness to unreliable network connections and mechanical failures, rather than reconstructing detailed scene geometry, limiting their use for tasks like content creation, building information modelling or cultural heritage preservation.

(ii) *Centralised* approaches, by contrast, can take advantage of the computing power provided by one or more central servers and their ability to communicate with all agents at once to produce detailed, globally consistent maps. For example, Chebrolu et al. [6] described a semi-dense approach based on LSD-SLAM [18] in which monocular clients produced pose-graph maps of keyframes that were sent to a central server for optimisation. Riazuelo et al. [45] described a cloud-based distributed visual SLAM system inspired by PTAM [31]. Their clients performed tracking only, with the more expensive mapping steps being performed in the cloud. Mohanarajah et al. [36] described another cloud-based approach based on the authors’ Rapyuta robotics platform. The client robots estimated their local poses by running dense, keyframe-based visual odometry on RGB-D images from a PrimeSense camera. The keyframes were then centrally optimised using g<sup>2</sup>o [32]. Poiesi et al. [41] described yet another cloud-based approach that performs Structure from Motion (SfM) and local bundle adjustments on monocular videos from smartphones to reconstruct a consistent point cloud map for each client, and more costly periodic full bundle adjustments to align the maps of different clients. Forster et al. [22] demonstrated centralised, keyframe-based collaborative mapping from micro aerial vehicles (MAVs), each equipped with a monocular camera and an IMU. More recently, Schmuck and Chli [48] have shown how to incorporate server-to-client feedback into a multi-UAV collaborative approach, allowing agents to share information.

A few approaches do not fit cleanly into either category. Reid et al. [44] described a distributed approach in which multiple autonomous ground robots were controlled from a centralised ground control station (GCS), but were able to operate independently for a while if the connection to the GCS failed. McDonald et al. [34] described a stereo approach in which a single agent reconstructed a scene over multiple sessions, effectively collaborating with itself over time. Chen et

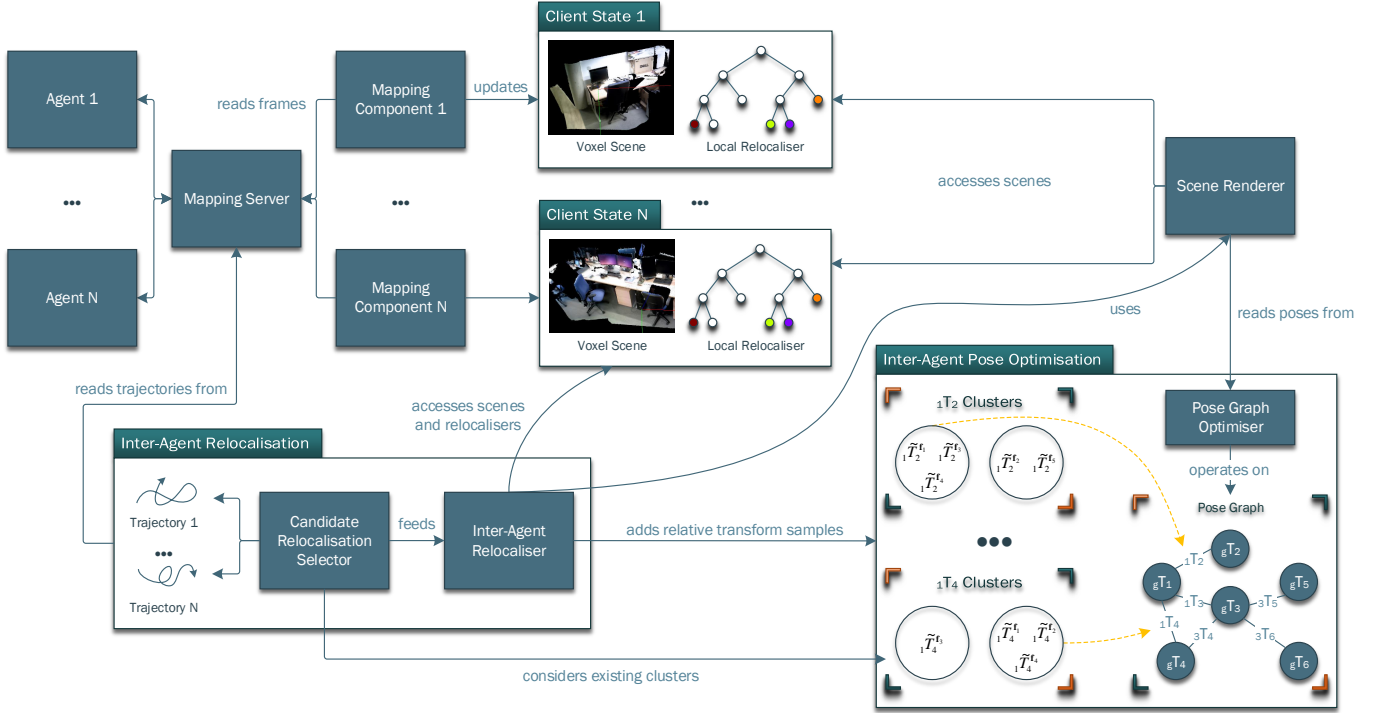


Fig. 3: The architecture of our system. Individual agents track their poses and feed posed RGB-D frames to the mapping server. A separate mapping component is instantiated for each agent, which reconstructs a voxel scene and trains a local relocaliser. Separately, the candidate relocalisation selector repeatedly selects a pose from one of the agents’ trajectories for relocalisation against another scene. The inter-agent relocaliser uses the scene renderer to render synthetic RGB and depth images of the corresponding scene from the selected pose, and passes them to the local relocaliser of the target scene. If a relocalisation succeeds (and is verified, see §3.3.1), a sample of the relative transform between the two scenes is recorded. The relative transform samples for each scene pair are clustered for robustness (see §3.3.2). Whenever the cluster to which a sample is added is sufficiently large, we construct a pose graph by blending the relative poses in the largest clusters, and trigger a pose graph optimisation. The optimised poses are then used for rendering the overall scene.

al. [8] described an approach that is initially peer-to-peer, with each robot building a pose graph independently, storing its sensor data, and exchanging information with other robots on rendezvous, but later client-server, with the robots transferring their pose graphs and sensor data to a central server for pose graph optimisation and the building of a coloured point cloud map. Fankhauser et al. [19] described an approach that allowed an Asctec Firefly hexacopter and a quadrupedal ground robot to work together to help the ground robot to navigate safely. Whilst no fixed-position central server was used, the ground robot had significant computing power and effectively played the role of a server, bundle adjusting data from the hexacopter to maintain a globally consistent map.

### 3 OUR APPROACH

Since we target scenarios such as indoor building reconstruction, where hardware failure is a minor concern, we adopt a centralised approach with multiple clients and a powerful central server (e.g. a laptop or desktop with one or more high-end GPUs): see Figure 3. This gives us the computational power to relocalise different clients’ sub-scenes against each other fast enough to support interactive reconstruction, and reduces the network bandwidth we require (since there is then no need for direct peer-to-peer communication). Moreover, since interactive reconstruction is a key target of our approach, we stipulate that each client must estimate accurate local poses for a sequence of RGB-D frames (see §3.1) and transmit both the frames and the poses to the central server (see §3.2). We do this for two reasons: (i) it is helpful to allow users to see a good approximation of the final reconstruction interactively to help guide the mapping process, and (ii) insisting that the local submaps are consistent during live reconstruction is essential if we are to generate a consistent overall map on the fly. The ability to provide a consistent live map has important implications for applica-

tions such as robotics, since a robot that cannot trust its map must waste computation on modelling this uncertainty when making decisions.

For interactive reconstruction, our insistence on accurate local poses implies the client-side use of a high-quality camera tracker (e.g. one based on visual-inertial odometry [27]) that can provide such poses on the fly. However, our approach is general enough that we can also support an *offline* (batch) reconstruction mode. For this, accurate live local poses are not strictly necessary, allowing more flexible schemes to be used. In particular, one option (see §4.5) is to reconstruct a submap for each client using BundleFusion [17] and then send the bundle-adjusted poses to the server. Alternatively, the clients could transmit inaccurate local poses and the bundle adjustment could be performed on the server, although that has some downsides (see §3.1) and would involve a different architecture.

In our architecture, regardless of how the clients are implemented, the server constructs a sub-scene and trains a relocaliser for each client based on the frames and poses it receives (however they are produced), and relocalises the sub-scenes from different clients against each other to establish a consistent global map (see §3.3). For a batch reconstruction, all of the sub-scenes are sent across to the server (or simply read directly from disk) before any relocalisation is performed. For an interactive reconstruction, the server will start relocalising clients against each other immediately, and new clients can join on the fly to contribute to the map. During the reconstruction process itself, the sub-scenes are kept separate to make it possible to continually optimise the relative poses between them (see §3.3.2); if desired, they can later be fused into a single consistent map at the end of the reconstruction process, as described in the supplementary material.

In practice, it can be helpful for collaborating clients to see which parts of the space have already been reconstructed by other users. For this reason, our approach allows each client to ask the server to render



the global scene from a specified camera pose (see §3.4). This server-to-client feedback is particularly helpful in situations in which users are out of each other’s visual range, since it can help them (i) plan how they will meet to join their sub-scenes together, and (ii) decide which parts of the space still need to be mapped.

### 3.1 Local Pose Estimation (Client-Side)

As mentioned above, each client in our system must estimate accurate local poses for a sequence of RGB-D frames and then transmit both the frames and the poses to the central server. This is non-trivial, since traditional pose estimation approaches, particularly those based purely on visual tracking, have tended to be subject to significant tracking drift, particularly at larger scales. For collaborative mapping, a common solution has been for clients to simply transmit inaccurate local poses and rely on the server to perform global optimisation (e.g. pose graph optimisation between keyframes [8, 36, 44] or bundle adjustment [6, 21, 48]) to achieve a globally consistent map (the optimised poses can then be sent back to the clients if desired [6, 36, 48]). However, (i) such global optimisations do not scale well, limiting the overall size of map that can be constructed, and (ii) the local poses on the clients are only corrected as global optimisations on the server finish, meaning that for much of the time they cannot be fully trusted. For these reasons, we choose in our approach to place the burden of accurate local pose estimation on the clients, since (i) that is a more natural fit for the interactive scenario we target, and (ii) it frees the server to focus on reconstructing a dense global map in real time. Two side benefits of this approach are that it significantly simplifies the design of the server (see §3.3) and reduces server-side memory consumption, making it possible to support more agents and larger global maps.

In practice, our approach is agnostic to how a client that can satisfy this requirement is implemented, and many popular solutions (generally based on some form of visual-inertial odometry) exist [1, 4, 26, 27] (for concrete examples, see §4). Moreover, our approach naturally supports collaborations between different types of client, making it possible for them to cooperate on the same reconstruction.

### 3.2 RGB-D Frame Transmission (Client → Server)

We maintain a TCP connection between each client and the server to guarantee in-order delivery of RGB-D frames. To minimise the network bandwidth, we only transmit frames that were tracked successfully, and we compress the depth images in PNG format (lossless) and the RGB images in JPG format (lossy). Moreover, to maintain a smooth, interactive user experience on the client, we transmit messages containing the frames (and their accompanying poses) to the server on a separate thread, which iteratively reads a frame message from a *pooled queue* of reusable messages, compresses it and sends it to the server.<sup>1</sup> The main thread writes uncompressed frame messages into this queue based on the current input. We discard messages that would overflow the queue (e.g. when the network cannot keep up with the client) to maintain interactivity and bound client memory usage; the way in which this interacts with our compression strategy is evaluated in §4.2.1.

At the server end, each client has a handler (running on a separate thread) that maintains a pooled queue of uncompressed frame messages. When a compressed frame message arrives, it is immediately uncompressed and pushed onto the queue (as on the client, we discard messages that would overflow the queue). On the main thread, we run a mapping component for each client (see §3.3), which reads RGB-D frames and their accompanying poses from its client’s queue as necessary and creates a local sub-scene.

### 3.3 Global Mapping (Server-Side)

The server has two jobs: (i) constructing a sub-scene and training a localiser for each client, and (ii) determining the relative transforms between these sub-scenes to establish a global coordinate system. To achieve the first, a separate mapping component for each client (a)

<sup>1</sup>Further details about the design of the pooled queue data structure can be found in the supplementary material.

runs the open-source InfiniTAM reconstruction engine [42, 43] on incoming (posed) RGB-D frames to construct a voxel-based map, and (b) trains a regression forest-based localiser for the sub-scene online as per Cavallari et al. [5]. To achieve the second, the server attempts to localise synthetic images of one agent’s sub-scene using another agent’s localiser to find estimates of the relative transform between the sub-scenes (see §3.3.1). These samples are clustered in transformation space to help suppress outliers, and a pose graph is constructed and optimised in the background to further refine the relative transforms between the sub-scenes (see §3.3.2). This optimisation is inspired by the single-agent sub-mapping approach of Kähler et al. [29], which showed how to build globally consistent voxel-based models by dividing a scene into small sub-scenes and optimising the relative poses between them. However, we construct a pose graph where each agent is represented by a single node, and the edges denote the relative transforms between different agents’ sub-scenes. This differs from [29], where all sub-scenes came from one agent and no relocalisation was needed to establish the transforms between them.

#### 3.3.1 Inter-Agent Relocalisation

To maintain a smooth, interactive experience on the server, we attempt relocalisations between the sub-scenes of different clients on a separate thread and (if available) a separate GPU. The scheduling of relocalisation attempts depends on the server mode. In *batch* mode, relocalisations are attempted only once all of the client sub-scenes have been fully created, at which point they are attempted as quickly as possible (i.e. a new relocalisation attempt is scheduled as soon as the previous attempt finishes). In *interactive* mode, we relocalise whilst the client sub-scenes are still being reconstructed, but at most once every 50 frames: we do this to space out relocalisation attempts and allow sufficient time in between attempts for the localisers to be trained.

To schedule an attempt, we first randomly generate a list of 10 candidate relocalisations. Each candidate  $\mathbf{k} = (a, \mathbf{f})$ , where  $\mathbf{f} = (b, i)$ , denotes an attempt to relocalise frame  $i$  of scene  $b$  in scene  $a$ ’s coordinate system. To balance different pairs of scenes, which may have been reconstructed from varying numbers of frames, we first uniformly sample an  $(a, b)$  scene pair from the set of all scene pairs, and then uniformly sample a frame index  $i$  from scene  $b$ . Each generated candidate  $\mathbf{k} = (a, (b, i))$  is then scored via

$$\Phi(\mathbf{k}) = \phi_{\text{new}}(\mathbf{k}) - \phi_{\text{conf}}(\mathbf{k}) - \phi_{\text{homog}}(\mathbf{k}). \quad (1)$$

In this,  $\phi_{\text{new}}$  aims to give a boost to candidates that might connect new nodes to the pose graph, defined as 1 if one of the sub-scenes  $a$  and  $b$  already has an optimised global pose and the other does not, and 0 otherwise.  $\phi_{\text{conf}}$  penalises candidates that can only add a relative transform sample between a pair of sub-scenes that are already confidently relocalised with respect to each other:

$$\phi_{\text{conf}}(a, (b, i)) = \max \left( 0, \max_{c \in \text{clusters}(a, b)} (|c| - N) \right). \quad (2)$$

Here,  $N = 2$  is an empirically-chosen threshold on the number of relocalisations needed before we become confident that the relative transform between a pair of sub-scenes is correct.  $\phi_{\text{homog}}$  penalises candidates whose frame  $i$  has a local pose in scene  $b$  that is too close to (within 5cm / 5° of) one that has already been tried against scene  $a$ ; we use a penalty of 5 for poses that are too close, and 0 otherwise.

Having scored the candidates, we schedule a relocalisation attempt for a candidate  $\mathbf{k} = (a, \mathbf{f})$  with maximum score, where  $\mathbf{f} = (b, i)$ . This proceeds as shown in Figure 4. Let  $\mathbf{f}^s$  denote the pose of frame  $\mathbf{f}$  in the coordinate system of a sub-scene  $s$ . First, we render synthetic RGB and depth raycasts  $\tilde{C}_b^{\mathbf{f}}$  and  $\tilde{D}_b^{\mathbf{f}}$  of sub-scene  $b$  from the known pose  $\mathbf{f}^b$ . Then, we try to relocalise them using sub-scene  $a$ ’s localiser to obtain an estimated pose  $\mathbf{f}^a$  for  $\mathbf{f}$  in  $a$ ’s coordinate system. The effectiveness of relocalising from synthetic images is evaluated in §4.2.2.

To verify the estimated pose, we render a synthetic depth raycast  $\tilde{D}_a^{\mathbf{f}}$  of sub-scene  $a$  from  $\mathbf{f}^a$ , and compare it to  $\tilde{D}_b^{\mathbf{f}}$ . To perform the comparison, we let  $\Omega$  be the common domain of  $\tilde{D}_a^{\mathbf{f}}$  and  $\tilde{D}_b^{\mathbf{f}}$  (i.e. the set of possible pixel coordinates in the two images),  $\Omega_s$  be the subset of  $\Omega$

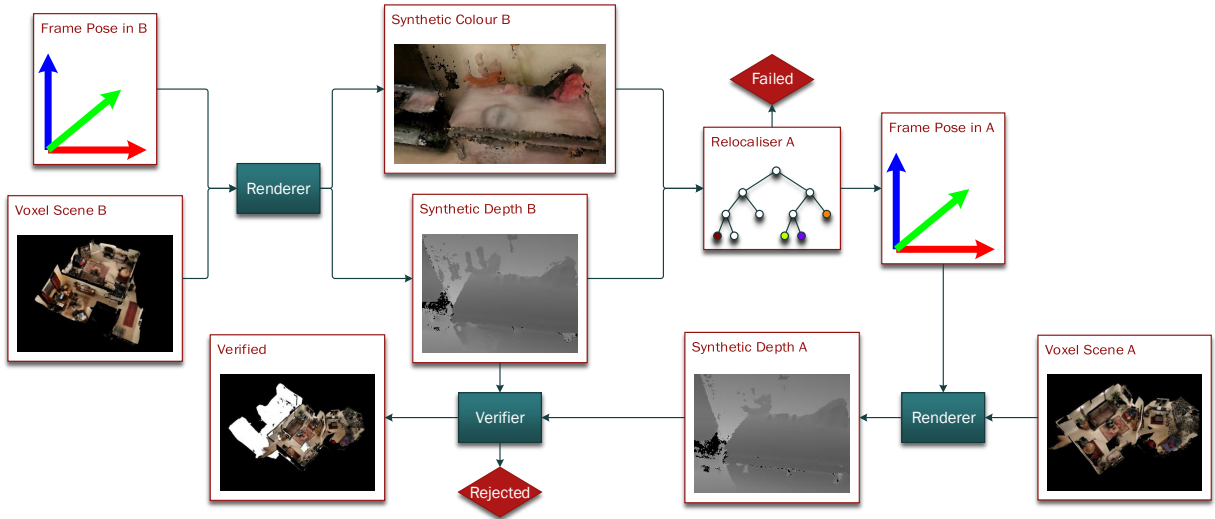


Fig. 4: To relocalise the scene of an agent  $b$  against that of another agent  $a$ , we first choose an arbitrary frame from  $b$ ’s trajectory and render synthetic RGB and depth raycasts of  $b$ ’s scene from the frame’s pose. Then, we try to relocalise them using  $a$ ’s relocaliser, which either fails, or produces an estimated pose for the frame in  $a$ ’s coordinate system. If a pose is proposed, we verify it by rendering a synthetic depth raycast of  $a$ ’s scene from the proposed pose and comparing it to the synthetic depth raycast of  $b$ ’s scene. We accept the pose if and only if the two depth raycasts are sufficiently similar (see §3.3.1).

for which  $\tilde{D}_s^f(\mathbf{x})$  is valid (where  $s$  is either  $a$  or  $b$ ), i.e.

$$\Omega_s = \{\mathbf{x} \in \Omega : \tilde{D}_s^f(\mathbf{x}) > 0\}, \quad (3)$$

and  $\Omega_{a,b} = \Omega_a \cap \Omega_b$ . Then, provided  $|\Omega_a|/|\Omega|$  is greater than a threshold (empirically set to 0.5), i.e. provided the estimated pose  ${}^f\tilde{T}_a$  in  $a$ ’s coordinate system is plausible, in the sense that it points sufficiently towards the reconstructed sub-scene, we compute a mean (masked) absolute depth difference between  $\tilde{D}_a^f$  and  $\tilde{D}_b^f$  via

$$\mu = \left( \sum_{\mathbf{x} \in \Omega_{a,b}} |\tilde{D}_a^f(\mathbf{x}) - \tilde{D}_b^f(\mathbf{x})| \right) / |\Omega_{a,b}|, \quad (4)$$

and add a relative transform sample  ${}_a\tilde{T}_b^f = ({}^f\tilde{T}_a)^{-1} {}^f\tilde{T}_b$  between  $a$  and  $b$  if and only if  $\mu$  is sufficiently small (empirically, we found 5cm to be a good threshold). The effectiveness of this step is evaluated in §4.2.3.

### 3.3.2 Inter-Agent Pose Optimisation

We incrementally cluster the relative transform samples  $\{{}_a\tilde{T}_b^f\}$  we add for each pair of sub-scenes  $(a, b)$  prior to performing pose graph optimisation to suppress the effect of outliers on the final result. After adding each sample, we look to see if the cluster to which the sample has just been added now has  $\geq N$  samples (i.e. the sample has contributed to a confident relative transform): if so, it is worthwhile to run pose graph optimisation, since the pose graph we construct may have changed since the last run.

To construct the pose graph  $G$ , we first construct an intermediate graph  $G'$  that has a node for every sub-scene and an edge connecting every pair of sub-scenes between which a confident relative transform exists (i.e. every pair whose largest cluster has size  $\geq N$ ). For each edge  $(a, b)$  in  $G'$ , we blend the relative transform samples  $\{{}_a\tilde{T}_b^f\}$  in the largest cluster using dual quaternion blending [30] to form an overall estimate  ${}_a\tilde{T}_b$  of the relative transform from  $b$  to  $a$ . Next, we compute the connected components of  $G'$ , such that two sub-scenes end up in the same connected component if and only if there is a chain of confident relative transforms between them. Finally, we set the pose graph  $G$  to be the connected component of  $G'$  containing the sub-scene corresponding to the first agent. Note that, by design,  $G$  is quite small, containing only one edge for each sub-scene pair: this is helpful in our context, since it allows the optimisation to be run repeatedly in the background to optimise the overall map.

The goal of optimising  $G$  is to find an optimised global pose  ${}_g\hat{T}_a$  for each sub-scene  $a$  for which  $G$  has a node. To perform the optimisation, we use the approach of Kähler et al. [29], as implemented in InfiniTAM [42, 43]. This uses Levenberg-Marquardt to minimise the error function

$$\varepsilon(G) = \sum_{(a,b) \in \text{edges}(G)} \|\mathbf{v}({}_g\hat{T}_b^{-1} {}_g\hat{T}_a {}_a\tilde{T}_b)\|_2, \quad (5)$$

where  $\mathbf{v}(T) = (\hat{\mathbf{q}}(T), \mathbf{t}(T))^T$  denotes the concatenation of  $\hat{\mathbf{q}}(T)$ , the three imaginary components of a quaternion representing the rotational part of  $T$ , and  $\mathbf{t}(T)$ , the translational part of  $T$ . Implicitly, this optimisation is trying to achieve  ${}_g\hat{T}_a^{-1} {}_g\hat{T}_b = {}_a\tilde{T}_b$  for every  $a$  and  $b$ , i.e. to ensure that the optimised global poses of all the scenes are consistent with the estimated relative poses between them.

### 3.4 Global Map Feedback (Server → Client)

To allow users to collaborate more effectively to reconstruct a scene (particularly in situations in which they can be out of each other’s visual range for at least some of the time), we allow each client to ask the server to render the global scene from a single specified camera pose: the rendered images can then be shown to the users to help them plan how to join their sub-scenes together and decide which parts of the space still need to be mapped. As when transmitting frames in the client-to-server direction, we compress the rendered images before sending them to the clients via TCP.

To achieve acceptable performance for multiple clients, we render low-resolution images and service only the most recent request from each client. Specifically, the server stores a pose from which to render for each client, and has a background thread that repeatedly renders an image for one client at a time (from its desired rendering pose), in a round-robin fashion to ensure fairness between the different clients. Each client can update its desired rendering pose at any point by sending a pose update message to the server (in practice, we send such messages once per frame). The reason for adopting this approach is to reduce network bandwidth and maintain interactivity on both the server and the clients, at the cost of some client-side rendering latency (specifically, the global scene image we show on each client will generally be slightly out of date in comparison to the most recent version of the global scene on the server). For live collaborative reconstruction, we found that this latency was not a major problem in practice: clients simply need to have an idea of what the global scene looks like so that they can plan what actions to take next, and those actions do not generally vary significantly in response to minor changes in the global scene.

## 4 EXPERIMENTS

We perform both quantitative and qualitative experiments to evaluate our approach. As mentioned in §3.1, our approach is able to work with any RGB-D client that can provide sufficiently accurate local poses to the server. To demonstrate this, we implemented two different types of client: (i) a straightforward Android client that captures and streams the RGB-D images and poses provided by the visual-inertial odometry solution in an Asus ZenFone AR augmented reality smartphone, and (ii) a PC-based client integrated into SemanticPaint [24]. By using both visual and inertial cues, the ZenFone client is able to provide high-quality poses, but camera limitations mean that the depth provided is fairly low resolution (only  $224 \times 172$ ). By contrast, the PC-based client can be used with conventional RGB-D cameras such as a Kinect, PrimeSense or Orbbec Astra, which can provide higher-resolution depth (e.g.  $640 \times 480$ ), but the visual-only tracking provided in the InfiniTAM framework [43] on which SemanticPaint is based can drift far more than visual-inertial odometry, particularly at larger scales.

One option for mitigating this is to capture data using a real-time globally consistent reconstruction system such as BundleFusion [17], which can close loops and re-integrate the scene surface on the fly, and then use our PC-based client to feed frames with bundle-adjusted poses to the server. However, BundleFusion can only provide accurate poses once loops have been closed (limiting its use to collaboration in batch mode), and its visual-only tracker provides poses that, even after loop closure resolution, are less accurate than those provided by visual-inertial odometry. Moreover, such trackers can sometimes fail before loops have been closed, e.g. due to untextured regions or a lack of interesting geometry, in which case reconstruction will fail. Since reliable tracking and accurate local poses are important in our large-scale context, we used the Android client for most of our experiments and to create our dataset. (On a practical level, smartphones are also more convenient for collaborative reconstruction than heavy laptops.) Nevertheless, to demonstrate our approach’s adaptability to other client implementations, we show a collaborative reconstruction using BundleFusion in §4.5.

In overview, we conduct the following experiments. In §4.1, we evaluate the quality of a reconstruction we obtain with our collaborative approach. In §4.2, we evaluate the effectiveness of various individual components of our system. In §4.3, we demonstrate the scalability of our approach in terms of both time and memory usage. In §4.4, we time how long our approach takes to produce consistent reconstructions for four different subsets of our dataset. We show in particular that using our system, it is possible to collaboratively reconstruct an entire three-storey research lab in under half an hour. Finally, in §4.5, we show the adaptability of our approach to different client implementations, as previously mentioned. Further details about our dataset can be found in the supplementary material.

### 4.1 Reconstruction Quality

To demonstrate our ability to achieve high-quality collaborative reconstructions, we combined 4 sequences from the *Flat* subset of our dataset to reconstruct a combined map of a two-bedroom flat (see Figure 5). Since a ground-truth reconstruction of the flat was not available (and we did not have access to a LiDAR scanner with which to obtain one), we validated our reconstruction by comparing a variety of measurements made on our combined map with ground-truth measurements made with a laser range finder (a Bosch Professional GLM 40) in the real world. To achieve this, we first converted the voxel-based maps of the sub-scenes produced by InfiniTAM [42, 43] into mesh-based maps using Marching Cubes [33]. We then applied the relative transforms we had estimated between the sub-scenes during the relocalisation process to these meshes to transform them into a common coordinate system. Finally, we imported the transformed meshes into MeshLab [12], and used its point-to-point measurement tool to make our measurements. As shown in Figure 5, we found that the measurements on our reconstructed model were consistently within 5cm of the ground truth, indicating that we are able to achieve reconstructions that correspond well to real-world geometry. Further examples, this time showing collaborative reconstructions of multi-storey houses and a three-storey




Fig. 5: An example of the high-quality reconstructions we can achieve using our collaborative approach. We joined the 4 *Flat* sequences from our dataset to make a combined map of a two-bedroom flat. The purple lines show the measurements we performed both using a laser range finder (in the real world) and on a combined mesh of our model (in MeshLab [12]) to validate our approach (the ordering of the numbers is range finder then mesh).

research lab, can be found in the supplementary material.

### 4.2 Effectiveness of Individual Components

#### 4.2.1 Frame Compression

To demonstrate the impact that compressing the RGB-D frames we transmit over the network has on the final reconstruction quality on the server, we used CloudCompare [13] to compare a locally-reconstructed reference model of an office with two reconstructions performed from frame data sent over a WiFi connection, one with frame compression disabled and the other with it enabled (see Figure 6). We measured the bandwidth of this WiFi connection to be roughly 8 Mbits/s, which is less than the roughly 25 Mbits/s that we needed to transmit the same sequence in compressed form without loss over a wired connection, and much less than the roughly 250 Mbits/s that would have been needed to transmit the sequence without loss uncompressed. In other words, with compression enabled, we are able to transmit around 1 in 3 of the frames in the sequence over WiFi whilst maintaining real-time rates; without compression, this drops to more than 1 in 30. As Figure 6 shows, this has a significant effect on the resulting reconstruction quality: with compression disabled (b), we lose some parts of the map completely and have a higher error rate across the map as a whole; with compression enabled (c), we manage to reconstruct more or less the entire map, and the error rate is greatly reduced.

#### 4.2.2 Relocalisation from Synthetic Images

As described in §3.3.1, we relocalise the scenes of different agents against each other using synthetic images, rather than the real RGB-D frames originally captured by the agents. This avoids the prohibitive memory cost of storing all RGB-D frames acquired by each agent in RAM (easily hundreds of MBs per agent). However, we might expect using synthetic images to lower relocalisation performance, since we train the relocalisers for each scene using the real input frames.

To verify that this is not a problem, we used the approach of Cavallari et al. [5] on the standard 7-Scenes dataset [49] to compare the results we are able to obtain with synthetic images to those we can obtain using real images. Unlike [5], which used at most 10 modes in each leaf of

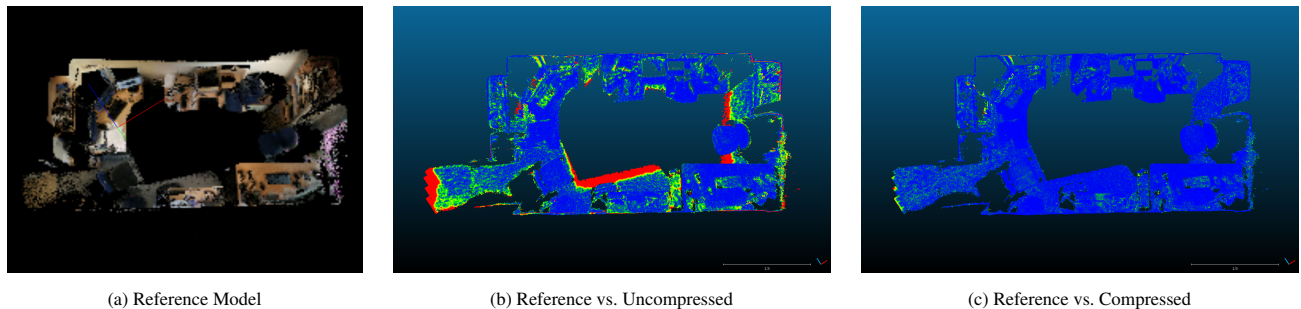


Fig. 6: An example showing the impact of frame compression on the reconstruction quality we are able to achieve on the server whilst maintaining real-time frame rates by discarding some frames: (a) a locally-reconstructed reference model of an office; (b) the differences between the reference model and a model reconstructed from the frames that we managed to transmit without using compression; (c) the differences between the reference model and a model reconstructed from the frames that we managed to transmit with compression. The comparisons were made using CloudCompare [13]. The errors range from blue = 0cm to red  $\geq 5$ cm (or missing). With compression enabled, we are forced to discard far fewer frames, allowing us to achieve a much lower error rate with respect to the reference model.

Sequence	Real Images	Synthetic Images
Chess	99.85%	<b>100.00%</b>
Fire	<b>99.20%</b>	98.85%
Heads	<b>100.00%</b>	<b>100.00%</b>
Office	99.80%	<b>99.95%</b>
Pumpkin	90.10%	<b>93.00%</b>
RedKitchen	91.34%	<b>99.98%</b>
Stairs	78.30%	<b>88.10%</b>

Table 2: Comparing relocalisation results obtained by rendering synthetic images of the scenes in the 7-Scenes dataset [49] from the test poses to those obtained using the real test images, by adapting a regression forest pre-trained on *Office* [5]. Percentages denote proportions of test frames with  $\leq 5$ cm translation error and  $\leq 5^\circ$  angular error.

the regression forests, we used at most 50 modes, since we found that this gave better results (in both cases). To test our synthetic approach, we first reconstructed each 7-Scenes sequence from the real training images as normal, then rendered synthetic frames from the testing poses, rather than using the testing images in the dataset. As Table 2 shows, our results using synthetic images were at least as good as, and in most cases actually higher than, the results using the real images, verifying that using synthetic images does not decrease relocalisation performance in practice. This is likely because by rendering synthetic images from the reconstructed scene, we implicitly remove noise from the frames to be relocalised, improving pose estimation accuracy.

#### 4.2.3 Depth Difference Verification

As mentioned in §3.3.1, we verify each proposed inter-agent relocalisation of a sub-scene  $b$  against a sub-scene  $a$  by rendering a synthetic depth image of scene  $a$  from the proposed pose and comparing it to the synthetic depth image from scene  $b$  that was passed to the relocaliser, masking out pixels that do not have a valid depth in both images. To evaluate the effectiveness of this approach, we first took 8 pairs of sequences  $\{a, b\}$  from our dataset that we were able to successfully relocalise with respect to each other during the normal operation of our approach, and recorded the relative transform between them as the ‘ground truth’ transform for later use. We then attempted to relocalise every frame from scene  $b$  using the relocaliser of scene  $a$ , and vice-versa. Next, for each frame for which the relocaliser proposed a relative transform, we ran our verification step on the proposed transform, thereby classifying it as either *Verified* or *Rejected*. We then compared the proposed transform to the ‘ground truth’ transform, classifying it as *Correct* if it was within  $5^\circ$  and 5cm of the ground truth, and as *Incorrect* otherwise. Finally, we counted Verified/Correct transforms as true positives (TP), Verified/Incorrect transforms as false positives (FP), and similarly for the true negatives (TN) and false negatives (FN), as shown in Table 3.

Our results show that our verifier has an extremely high average recall rate (99.2%), meaning that it largely manages to avoid rejecting correct transforms. It also has a reasonably good average specificity (75.9%), meaning that it is fairly good at pruning the number of incorrect transforms we need to deal with. However, a fairly large number of incorrect transforms still manage to pass the verification stage: as mentioned in §3.3.2, these are dealt with later by clustering the transforms and only making use of the transforms from the largest cluster. The additional effects of this clustering step are evaluated in §4.2.4.

#### 4.2.4 Relative Transform Clustering

As described in §3.3.2, we incrementally cluster the relative transform samples  $\{a \tilde{T}_b^f\}$  we add for each pair of sub-scenes  $(a, b)$  prior to performing pose graph optimisation to suppress the effect of outliers on the final result. This is achieved by checking each new relative transform sample to see if there is an existing cluster to which it can be added (we specify that this is possible if and only if it is within  $20^\circ$  and 10cm of any existing relative transform in the cluster). If so, we add the sample to the first such cluster we find; if not, we create a new cluster for it.

To evaluate the effectiveness of this approach, we took the same 8 pairs of sequences  $\{a, b\}$  we used to evaluate our depth difference verifier in §4.2.3, and again relocalised every frame from scene  $b$  using the relocaliser of scene  $a$ , and vice-versa. We then counted the number of relative transform samples that were added during this process, and examined the clusters into which they had been collected. In particular, we compared the size of the largest cluster in each case (i.e. the largest ‘correct’ cluster) with the size of the largest cluster whose blended transform (obtained by blending all of the relative transforms in the cluster using dual quaternion blending [30]) was not within  $20^\circ$  and 10cm of the blended transform of the correct cluster. We refer to this latter cluster as the largest ‘incorrect’ cluster. The difference between these two sizes gave us a measure of the *safety margin* of our approach in each case, i.e. the number of consistent erroneous samples that would need to have been added to the largest incorrect cluster to cause it to have been chosen instead of the correct cluster.

As our results in Table 4 show, for most pairs of scenes the size of the correct cluster was significantly larger than the size of the largest incorrect cluster, indicating that in practice we are very likely to accumulate  $\geq N$  samples from the correct cluster and become confident about it long before we accumulate  $\geq N$  samples from an incorrect cluster. For two pairs of scenes (*f/bathroom2study* & *f/study2sittingroom*, and *h/hall2frontbedroom* & *h/mainbedroom2studio*), the safety margins were much lower than in the other cases. However, in both cases, the pairs of scenes in question have comparatively low overlap (see the figures showing the green and yellow sequences for *Flat* and *House* in the supplementary material). Moreover, whilst the blended transforms of the correct cluster and the largest incorrect cluster in each case were not within  $20^\circ$  and 10cm of each other, a manual inspection of the relevant transforms showed that they were still comparatively close,



Scene 1	Scene 2	Total Frames	Relocalised Frames	Verifier Performance						
				TP	FP	TN	FN	Precision	Recall	Specificity
f/bathroom2study	f/kitchen2study	4901	2190	517	931	741	1	35.7%	99.8%	44.3%
f/bathroom2study	f/study2sittingroom	4989	519	24	59	436	0	28.9%	100.0%	88.1%
f/kitchen2study	f/study2sittingroom	5544	1254	325	317	609	3	50.6%	99.1%	65.8%
f/study2sittingroom	f/turret2sittingroom	5895	1350	555	211	579	5	72.5%	99.1%	73.3%
h/frontroom2study	h/hall2frontbedroom	7714	1229	245	90	894	0	73.1%	100.0%	90.9%
h/frontroom2study	h/hall2oldkitchen	8171	1307	373	175	757	2	68.1%	99.5%	81.2%
h/hall2frontbedroom	h/hall2oldkitchen	7141	653	95	106	451	1	47.3%	99.0%	81.0%
h/hall2frontbedroom	h/mainbedroom2studio	6942	839	90	131	615	3	40.7%	96.8%	82.4%
Average (all scene pairs)		—	—	—	—	—	—	52.1%	99.2%	75.9%

Table 3: Evaluating the effectiveness of the depth difference verification we perform on proposed relocalisations between pairs of sub-scenes (see §4.2.3). We attempt to relocalise every frame of each sub-scene against the other using [5], and record the number of frames that we were able to relocalise, together with statistics on how many correct/incorrect relocalisations were verified/rejected by the verifier.

Scene 1	Scene 2	Total Frames	Samples Added	Correct Cluster	Largest Incorrect Cluster	Safety Margin
f/bathroom2study	f/kitchen2study	4901	1472	1317 (89.5%)	18 (1.2%)	1299
f/bathroom2study	f/study2sittingroom	4989	82	21 (25.6%)	18 (22.0%)	3
f/kitchen2study	f/study2sittingroom	5544	589	506 (85.9%)	5 (0.8%)	501
f/study2sittingroom	f/turret2sittingroom	5895	765	759 (99.2%)	2 (0.3%)	757
h/frontroom2study	h/hall2frontbedroom	7714	360	249 (69.2%)	16 (4.4%)	233
h/frontroom2study	h/hall2oldkitchen	8171	592	442 (74.7%)	10 (1.7%)	432
h/hall2frontbedroom	h/hall2oldkitchen	7141	242	192 (79.3%)	12 (5.0%)	180
h/hall2frontbedroom	h/mainbedroom2studio	6942	260	102 (39.2%)	69 (26.5%)	33

Table 4: Evaluating the extent to which our approach of incrementally clustering the relative transform samples between different pairs of sub-scenes is able to remove outliers and find consistent relocalisations between the sub-scenes. We attempt to relocalise every frame of each sub-scene against the other using [5], and record the total number of samples that were added (equal to the number of relocalised frames that passed depth difference verification), and the sizes of the correct cluster and the largest incorrect cluster produced by our method in each case (together with the percentages of the samples added these sizes represent). The safety margin for each scene pair refers to the number of consistent erroneous samples that would need to have been added to the largest incorrect cluster to cause it to have been chosen instead of the correct cluster.




Fig. 7: How the time taken by our approach (i) for sub-scene fusion / overall scene rendering (on the main thread), and (ii) for attempting background relocalisations between different sub-scenes (on a separate thread), changed whilst collaboratively reconstructing an office using a server and 4 Android clients, connected via a WiFi router (see §4.3.1).

meaning that the safety margin before hitting a cluster with a grossly incorrect transform is in practice somewhat higher in each case.

### 4.3 Scalability

#### 4.3.1 Time

To evaluate how well the time taken by our approach scales when multiple clients are collaborating to reconstruct a scene, we timed an interactive collaborative reconstruction of an office using a server and four Android clients, connected via a WiFi router (the server was

directly connected to the router via Ethernet for improved bandwidth, whilst the clients were connected wirelessly). In particular, we timed how long it took to process each frame (i.e. fusing a sub-scene for each client and rendering the overall scene) on the main thread, and how long it took to perform relocalisations between the different sub-scenes on a separate thread.

The results of this process are shown in Figure 7. The four clients connected one at a time at suitably spaced intervals, before later disconnecting again (see the corresponding time spikes, caused by allocating / deallocating memory for the different clients). As Figure 7 shows, the cost of fusing the sub-scenes and rendering the scene was consistently low (generally less than 50ms) throughout the experiment, allowing the user to continue to view the collaborative reconstruction interactively from different angles in real time on the server even with multiple clients connected.

Relocalisations between the sub-scenes were started once the second client connected (roughly 19% of the way through the experiment). As expected, since our approach schedules a single relocalisation at a time, the cost of attempting relocalisations between the sub-scenes stayed bounded over time (generally under 300ms, including the time taken to render the synthetic images and perform depth difference verification). Interestingly, there was a drop in the time taken per relocalisation after each client connected, caused by the online relocaliser for the new client taking some time to train (initially, each new relocaliser has empty leaves, making relocalisation attempts faster but less likely to be successful; the time taken and the likelihood of a successful relocalisation increase as the relocaliser is trained). The spike in the relocalisation time roughly 73% of the way through the experiment was caused by the synthetic image rendering for the active relocaliser (which happens on the main GPU, where the sub-scenes are stored) blocking whilst the memory for the first client to disconnect was deallocated.

Overall, these results show that, from a time perspective, our approach should be able to scale to a much larger number of clients






Fig. 8: The GPU memory we use during a collaborative reconstruction of the 6 *Priory* sequences from our dataset.

	Flat	Priory	House	Lab
# Sequences	4	6	4	4
Longest Sequence: # Frames	3079	1640	4372	3188
Longest Sequence: Time (s)	615.8	328.0	874.4	637.6
Average Collaboration Time (s)	15.8	38.4	107.4	272.0
Average Total Time (s)	631.6	366.4	981.8	909.6

Table 5: The total times taken to collaboratively reconstruct four different subsets of our dataset (see §4.4).

without impeding interactivity (we were only able to test with the four Android clients we had available, but given that the time taken by our approach increases by only a small amount for each additional client added, it should in practice remain reasonable even when handling a much larger number of clients). This is to be expected, given the efficiency of modern voxel fusion frameworks such as InfiniTAM [43], and the way in which we schedule our relocalisation attempts in the background to avoid blocking the primary GPU. In practice, the main limitation of our approach is currently GPU memory usage, as we show in the next section.

#### 4.3.2 Memory Usage

To evaluate the scalability of our approach’s GPU memory usage, we performed an offline collaborative reconstruction of the 6 *Priory* sequences from our dataset (these represent a three-story house: see the supplementary material) and examined the memory usage on the server as the number of clients increased. To maximise the number of agents we were able to handle, we added new sub-scenes to the server one at a time and deleted the training data used by each sub-scene’s relocaliser once it had been fully trained: this limited the maximum memory used by the relocalisers, leaving us bounded primarily by the size of the reconstructed voxel scenes. Figure 8 shows that the final GPU memory we use for each forest-based relocaliser [5] is  $< 500\text{MB}$ , meaning that we can potentially handle around 24 relocalisers on a 12GB GPU such as an NVIDIA Titan X (the actual number may be slightly less than this due to driver overhead). The memory used by the voxel scenes is currently the bottleneck preventing the scaling of our approach to more agents: currently, each scene takes up just over 1GB of memory, limiting us to around 11 agents on a Titan X (assuming the relocalisers are stored on a secondary GPU). To scale further, we could reduce the memory used by meshing each scene with Marching Cubes [33] and discarding the voxel maps.

#### 4.4 Overall Timings for Collaborative Reconstruction

To evaluate how long it takes to produce consistent reconstructions using our approach, we computed the average times taken to collaboratively reconstruct four different subsets of our dataset (see Table 5). We computed the time taken to capture the sequences in each subset as  $\frac{1}{5}$  times the length of the longest sequence, assuming parallel capturing at 5Hz. The average collaboration time is the time taken to relocalise the agents against each other and compute optimised global poses for their maps. To account for the random selection of frames to relocalise, we globally mapped each subset 5 times and reported the average time in the table. The average total time is from the start of the capturing




Fig. 9: A collaborative reconstruction of a two-bedroom flat, produced by running our approach in batch mode on four sub-scenes captured using BundleFusion [17].

process to the output of a globally consistent map: this was under half an hour for all subsets we tested.

#### 4.5 Adaptability to Different Clients

To demonstrate our approach’s ability to work with clients implemented in different ways, we used BundleFusion [17], a state-of-the-art real-time globally consistent reconstruction approach that can close loops and re-integrate the scene surface on the fly, to recapture the two-bedroom flat that forms the *Flat* subset of our dataset. Since BundleFusion can only produce accurate camera poses once loops have actually been closed, we reconstructed each of the four sub-scenes involved ahead of time and then combined them using the batch mode of our system. The results of this process are shown in Figure 9. In comparison to reconstruction using the Android client, we found that reconstruction using BundleFusion required far more care to be taken during the capture process, since its visual-only tracker is far more prone to fail than one based on visual-inertial odometry. (If an IMU is available, BundleFusion’s tracker can in principle be replaced with one based on visual-inertial odometry, e.g. [4], but doing so is non-trivial and beyond the scope of the present paper.) Nevertheless, at *Flat* scale, we found BundleFusion to be a viable replacement for our Android client.

### 5 CONCLUSION

In this paper, we have shown how to enable multiple agents to collaborate interactively to reconstruct dense, volumetric models of 3D scenes. Existing collaborative mapping approaches have traditionally suffered from an inability to trust the local poses produced by their mobile agents, forcing them to perform costly global optimisations (e.g. on a server) to ensure a consistent map, and limiting their ability to perform collaborative dense mapping interactively.

By leveraging recent advances to construct rigid local sub-scenes that do not need further refinement, and joining them using a state-of-the-art regression-based relocaliser, we avoid expensive global optimisations, opting only to refine the relative poses between individual agents’ maps. Our system allows multiple users to collaboratively reconstruct consistent dense models of entire buildings in under half an hour using only consumer-grade hardware, providing a low-cost, time-efficient and interactive alternative to existing whole-building reconstruction methods based on panorama scanners, and making it easier than ever before for users to capture detailed 3D scene models at scale.

#### ACKNOWLEDGMENTS

This work was supported by Innovate UK/CCAV project 103700 (Street-Wise), the EPSRC, ERC grant ERC-2012-AdG 321162-HELIOS,

EPSRC grant Seebibyte EP/M013774/1 and EPSRC/MURI grant EP/N019474/1. We would also like to thank Manar Marzouk and Maria Anastasia Tsitsigkou for their help with the collaborative dataset collection, and Oscar Rahnama, Tom Joy, Daniela Massiceti, Nantas Nardelli, Mohammad Najafi and Adnane Boukhayma for their help with the experiments.

## REFERENCES

- [1] Apple. ARKit. <https://developer.apple.com/arkit>.
- [2] I. Armeni, O. Sener, A. R. Zamir, H. Jiang, I. Brilakis, M. Fischer, and S. Savarese. 3D Semantic Parsing of Large-Scale Indoor Spaces. In *CVPR*, pp. 1534–1543, 2016.
- [3] A. Bajpai, G. Burroughes, A. Shaikat, and Y. Gao. Planetary Monocular Simultaneous Localization and Mapping. *JFR*, 33(2):229–242, 2016.
- [4] M. Bloesch, S. Omari, M. Hutter, and R. Siegwart. Robust Visual Inertial Odometry Using a Direct EKF-Based Approach. In *IROS*, pp. 298–304, 2015.
- [5] T. Cavallari, S. Golodetz\*, N. A. Lord\*, J. Valentin, L. D. Stefano, and P. H. S. Torr. On-the-Fly Adaptation of Regression Forests for Online Camera Relocalisation. In *CVPR*, 2017.
- [6] N. Chebrolu, D. Marquez-Gamez, and P. Martinet. Collaborative Visual SLAM Framework for a Multi-Robot System. In *PPNIV*, 2015.
- [7] F. A. A. Cheein and R. Carelli. Agricultural Robotics: Unmanned Robotic Service Units in Agricultural Tasks. *IEEE Industrial Electronics Magazine*, 7(3), 2013.
- [8] H. Chen, J. Zhong, Y. Fu, and Y. Lou. Pose-graph Based 3D Map Fusion with Distributed Robot System. In *ROBIO*, 2014.
- [9] S. Choudhary, L. Carlone, C. Nieto, J. Rogers, H. I. Christensen, and F. Dellaert. Distributed Mapping with Privacy and Communication Constraints: Lightweight Algorithms and Object-based Models. *arXiv preprint arXiv:1702.03435*, 2017.
- [10] T. Cieslewski, S. Choudhary, and D. Scaramuzza. Data-Efficient Decentralized Visual SLAM. *arXiv preprint arXiv:1710.05772*, 2017.
- [11] T. Cieslewski, S. Lynen, M. Dymczyk, S. Magnenat, and R. Siegwart. Map API – Scalable Decentralized Map Building for Robots. In *ICRA*, pp. 6241–6247, 2015.
- [12] P. Cignoni, M. Callieri, M. Corsini, M. Dellepiane, F. Ganovelli, and G. Ranzuglia. MeshLab: an Open-Source Mesh Processing Tool. In *Eurographics Italian Chapter Conference*, pp. 129–136, 2008.
- [13] CloudCompare v2.9 (GPL software), 2017. Retrieved from <http://www.danielgm.net/cc>.
- [14] A. Cunningham, V. Indelman, and F. Dellaert. DDF-SAM 2.0: Consistent Distributed Smoothing and Mapping. In *ICRA*, pp. 5220–5227, 2013.
- [15] A. Cunningham, M. Paluri, and F. Dellaert. DDF-SAM: Fully Distributed SLAM using Constrained Factor Graphs. In *IROS*, pp. 3025–3030, 2010.
- [16] A. Cunningham, K. M. Wurm, W. Burgard, and F. Dellaert. Fully Distributed Scalable Smoothing and Mapping with Robust Multi-robot Data Association. In *ICRA*, pp. 1093–1100, 2012.
- [17] A. Dai, M. Nießner, M. Zollhöfer, S. Izadi, and C. Theobalt. BundleFusion: Real-time Globally Consistent 3D Reconstruction using Online Surface Re-integration. *TOG*, 36(3):24, 2017.
- [18] J. Engel, T. Schöps, and D. Cremers. LSD-SLAM: Large-Scale Direct Monocular SLAM. In *ECCV*, pp. 834–849, 2014.
- [19] P. Fankhauser, M. Bloesch, P. Krüsi, R. Diethelm, M. Wermelinger, T. Schneider, M. Dymczyk, M. Hutter, and R. Siegwart. Collaborative Navigation for Flying and Walking Robots. In *IROS*, pp. 2859–2866, 2016.
- [20] N. Fioraio, J. Taylor, A. Fitzgibbon, L. D. Stefano, and S. Izadi. Large-Scale and Drift-Free Surface Reconstruction Using Online Subvolume Registration. In *CVPR*, pp. 4475–4483, 2015.
- [21] P. Fleck, C. Arth, C. Pirchheim, and D. Schmalstieg. Tracking and Mapping with a Swarm of Heterogeneous Clients. In *ISMAR*, 2015.
- [22] C. Forster, S. Lynen, L. Kneip, and D. Scaramuzza. Collaborative Monocular SLAM with Multiple Micro Aerial Vehicles. In *IROS*, pp. 3962–3970, 2013.
- [23] B. Glocker, J. Shotton, A. Criminisi, and S. Izadi. Real-Time RGB-D Camera Relocalization via Randomized Ferns for Keyframe Encoding. *TVCG*, 21(5):571–583, May 2015.
- [24] S. Golodetz\*, M. Sapienza\*, J. P. C. Valentin, V. Vineet, M.-M. Cheng, A. Arnab, V. A. Prisacariu, O. Kähler, C. Y. Ren, D. W. Murray, S. Izadi, and P. H. S. Torr. SemanticPaint: A Framework for the Interactive Segmentation of 3D Scenes. Technical Report TVG-2015-1, Department of Engineering Science, University of Oxford, October 2015. Released as arXiv e-print 1510.03727.
- [25] S. Golodetz\*, M. Sapienza\*, J. P. C. Valentin, V. Vineet, M.-M. Cheng, V. A. Prisacariu, O. Kähler, C. Y. Ren, A. Arnab, S. L. Hicks, D. W. Murray, S. Izadi, and P. H. S. Torr. SemanticPaint: Interactive Segmentation and Learning of 3D Worlds. In *ACM SIGGRAPH Emerging Technologies*, p. 22, 2015.
- [26] Google. ARCore. <https://developers.google.com/ar>.
- [27] J. A. Hesch, D. G. Kottas, S. L. Bowman, and S. I. Roumeliotis. Camera-IMU-based localization: Observability analysis and consistency improvement. *IJRR*, 33(1):182–201, 2014.
- [28] J. Huang, A. Dai, L. Guibas, and M. Niessner. 3DLite: Towards Commodity 3D Scanning for Content Creation. *TOG*, 2017.
- [29] O. Kähler, V. A. Prisacariu, and D. W. Murray. Real-Time Large-Scale Dense 3D Reconstruction with Loop Closure. In *ECCV*, pp. 500–516, 2016.
- [30] L. Kavan, S. Collins, C. O’Sullivan, and J. Zara. Dual Quaternions for Rigid Transformation Blending. Technical Report TCD-CS-2006-46, Trinity College Dublin, 2006.
- [31] G. Klein and D. Murray. Parallel Tracking and Mapping for Small AR Workspaces. In *ISMAR*, pp. 225–234, 2007.
- [32] R. Kümmerle, G. Grisetti, H. Strasdat, K. Konolige, and W. Burgard. g<sup>2</sup>o: A General Framework for Graph Optimization. In *ICRA*, pp. 3607–3613, 2011.
- [33] W. E. Lorensen and H. E. Cline. Marching Cubes: A High Resolution 3D Surface Construction Algorithm. *ACM SIGGRAPH Computer Graphics*, 21(4):163–169, 1987.
- [34] J. McDonald, M. Kaess, C. Cadena, J. Neira, and J. J. Leonard. Real-time 6-DOF multi-session visual SLAM over large-scale environments. *RAS*, 61(10):1144–1158, 2013.
- [35] N. Michael, S. Shen, K. Mohta, Y. Mulgaonkar, V. Kumar, K. Nagatani, Y. Okada, S. Kiribayashi, K. Otake, K. Yoshida, K. Ohno, E. Takeuchi, and S. Tadokoro. Collaborative Mapping of an Earthquake-Damaged Building via Ground and Aerial Robots. *JFR*, 29(5):832–841, 2012.
- [36] G. Mohanarajah, V. Usenko, M. Singh, R. D’Andrea, and M. Waibel. Cloud-based Collaborative 3D Mapping in Real-Time with Low-Cost Robots. *TASE*, 12(2):423–431, 2015.
- [37] S. Murali, P. Speciale, M. R. Oswald, and M. Pollefeys. Indoor Scan2BIM: Building Information Models of House Interiors. In *IROS*, 2017.
- [38] R. A. Newcombe, S. Izadi, O. Hilliges, D. Molyneaux, D. Kim, A. J. Davison, P. Kohli, J. Shotton, S. Hodges, and A. Fitzgibbon. KinectFusion: Real-Time Dense Surface Mapping and Tracking. In *ISMAR*, pp. 127–136, 2011.
- [39] M. Nießner, M. Zollhöfer, S. Izadi, and M. Stamminger. Real-time 3D Reconstruction at Scale using Voxel Hashing. *TOG*, 32(6):169, 2013.
- [40] L. Paull, G. Huang, M. Seto, and J. J. Leonard. Communication-Constrained Multi-UAV Cooperative SLAM. In *ICRA*, pp. 509–516, 2015.
- [41] F. Poiesi, A. Locher, P. Chippendale, E. Nocerino, F. Remondino, and L. V. Gool. Cloud-based collaborative 3D reconstruction using smartphones. In *CVMP*, 2017.
- [42] V. A. Prisacariu, O. Kähler, M. M. Cheng, C. Y. Ren, J. Valentin, P. H. S. Torr, I. D. Reid, and D. W. Murray. A Framework for the Volumetric Integration of Depth Images. *arXiv preprint arXiv:1410.0925*, 2014.
- [43] V. A. Prisacariu, O. Kähler, S. Golodetz, M. Sapienza, T. Cavallari, P. H. S. Torr, and D. W. Murray. InfiniTAM v3: A Framework for Large-Scale 3D Reconstruction with Loop Closure. *arXiv preprint arXiv:1708.00783v1*, 2017.
- [44] R. Reid, A. Cann, C. Meiklejohn, L. Poli, A. Boeing, and T. Braunl. Cooperative Multi-Robot Navigation, Exploration, Mapping and Object Detection with ROS. In *IV*, 2013.
- [45] L. Riazuelo, J. Civera, and J. M. M. Montiel. C<sup>2</sup>TAM: A Cloud framework for cooperative tracking and mapping. *RAS*, 62(4):401–413, 2014.
- [46] W. Rone and P. Ben-Tzvi. Mapping, localization and motion planning in mobile multi-robotic systems. *Robotica*, 31:1–23, 2013.
- [47] S. Saeedi, M. Trentini, M. Seto, and H. Li. Multiple-robot Simultaneous Localization and Mapping – A Review. *JFR*, 33(1):3–46, 2016.
- [48] P. Schmuck and M. Chli. Multi-UAV Collaborative Monocular SLAM. In *ICRA*, pp. 3863–3870, 2017.
- [49] J. Shotton, B. Glocker, C. Zach, S. Izadi, A. Criminisi, and A. Fitzgibbon. Scene Coordinate Regression Forests for Camera Relocalization in RGB-D Images. In *CVPR*, pp. 2930–2937, 2013.
- [50] L. Silveira, F. Guth, P. Drews-Jr, P. Ballester, M. Machado, F. Codevilla,

- N. Duarte-Filho, and S. Botelho. An Open-source Bio-inspired Solution to Underwater SLAM. *IFAC-PapersOnLine*, 48(2):212–217, 2015.
- [51] F. Steinbrücker, J. Sturm, and D. Cremers. Volumetric 3D Mapping in Real-Time on a CPU. In *ICRA*, pp. 2021–2028, 2014.
  - [52] J. Valentin, M. Nießner, J. Shotton, A. Fitzgibbon, S. Izadi, and P. Torr. Exploiting Uncertainty in Regression Forests for Accurate Camera Relocalization. In *CVPR*, pp. 4400–4408, 2015.
  - [53] O. Wasenmüller, M. Meyer, and D. Stricker. Augmented Reality 3D Discrepancy Check in Industrial Applications. In *ISMAR*, pp. 125–134, 2016.
  - [54] T. Whelan, M. Kaess, H. Johannsson, M. Fallon, J. J. Leonard, and J. McDonald. Real-time large scale dense RGB-D SLAM with volumetric fusion. *IJRR*, 34(4-5):598–626, 2015.
  - [55] T. Whelan, S. Leutenegger, R. F. Salas-Moreno, B. Glocker, and A. J. Davison. ElasticFusion: Dense SLAM Without A Pose Graph. In *RSS*, 2015.
  - [56] E. Wijmans and Y. Furukawa. Exploiting 2D Floorplan for Building-scale Panorama RGBD Alignment. In *CVPR*, pp. 308–316, 2017.
  - [57] M. Zollhöfer, C. Siegl, M. Vetter, B. Dreyer, M. Stamminger, S. Aybek, and F. Bauer. Low-Cost Real-Time 3D Reconstruction of Large-Scale Excavation Sites. *JOCCH*, 9(1), 2016.