

LASER: Level-based Asynchronous Scheduling and Execution Regime for Spatiotemporally Constrained Multi-Robot Timber Manufacturing

Zhenxiang Huang*, Lior Skoury, Tim Stark, Aaron Wagner,
Hans Jakob Wagner, Thomas Wortmann, and Achim Menges

Abstract—Automating large-scale manufacturing in domains like timber construction requires multi-robot systems to manage tightly coupled spatiotemporal constraints, such as collision avoidance and process-driven deadlines. This paper introduces LASER (Level-based Asynchronous Scheduling and Execution Regime), a complete framework for scheduling and executing complex assembly tasks, demonstrated on a screw-press gluing application for timber slab manufacturing. Our central contribution is to integrate a barrier-based mechanism into a constraint programming (CP) scheduling formulation that partitions tasks into spatiotemporally disjoint sets, which we define as “levels.” This structure enables robots to execute tasks in parallel and asynchronously within a level, synchronizing only at level barriers, which guarantees collision-free operation by construction and provides robustness to timing uncertainties. To solve this formulation for large problems, we propose two specialized algorithms: an iterative temporal-relaxation approach for heterogeneous task sequences and a bi-level decomposition for homogeneous tasks that balances workload. We validate the LASER framework by fabricating a full-scale 2.4 m × 6 m timber slab (Fig. 1) with a two-robot system mounted on parallel linear tracks, successfully coordinating 108 subroutines and 352 screws under tight adhesive time windows. Computational studies show our method scales steadily with size compared to a monolithic approach.

Index Terms—Multi-Robot Systems, Task Scheduling, Construction Automation, Constraint Programming, Timber Manufacturing.

I. INTRODUCTION

The increasing demand for automation in manufacturing and construction has positioned multi-robot systems (MRS) as a powerful solution for efficiency, scalability, and safety. In large assemblies, many steps can run in parallel if spatial interference and process windows are respected. Realizing this potential, however, requires addressing tightly coupled challenges in multi-robot task allocation and scheduling, motion planning in shared workspaces, and execution under uncertainty. These challenges are magnified in industrial settings, where complex spatiotemporal constraints, bespoke parts at scale, and strict quality criteria demand schedules that are both *collision-aware* and *deadline-aware* by construction.

This research has been supported by the Deutsche Forschungsgemeinschaft (DFG) under Germany’s Excellence Strategy EXC 2120/1 – 390831618.

The authors are with Institute for Computational Design and Construction (ICD) and Cluster of Excellence IntCDC at University of Stuttgart, Germany.

*zhenxiang.huang@icd.uni-stuttgart.de




Fig. 1. Timber slab prototype manufacturing using a two-robot system with our proposed framework.

Timber fabrication serves as a low-carbon alternative to steel and concrete, and its sustainability potential can be amplified through scalable, automated robotic manufacturing. We focus on *screw-press gluing*, a connection technique that applies localized bonding pressure via screw tapping, eliminating the need for massive and costly press machines. As a traditional crafting method, this process is labor-intensive due to its dense screwing pattern (one screw per 220 cm² in our system); naively automating with a single robot is often infeasible at scale because the number of screws can violate adhesive time windows. In a hollow timber slab system, two cross-laminated timber (CLT) plates sandwich discrete timber beams for material-efficient load transfer. To manufacture such a slab segment, efficient multi-robot collaboration must respect: (i) spatial collision-avoidance of robots and payloads, and (ii) adhesive *open* and *close* windows that couple glue application, placement, and all subsequent screws. They are the primary drivers of scheduling and execution design.

This paper presents LASER (Level-based Asynchronous Scheduling and Execution Regime), a complete framework (Fig. 2) to systematically address task scheduling and execution for spatiotemporally constrained multi-robot manufacturing, demonstrated here on the screw-press gluing process for timber slab assemblies. The main innovation of our scheduling formulation is to integrate *levels* as time–space partitions of tasks whose robot workspace occupancies do not conflict. Robots execute asynchronously within a level and synchronize only at *level barriers*. This lightweight mechanism yields collision-free operation by schedule design, robust to execution delays and timing uncertainty.

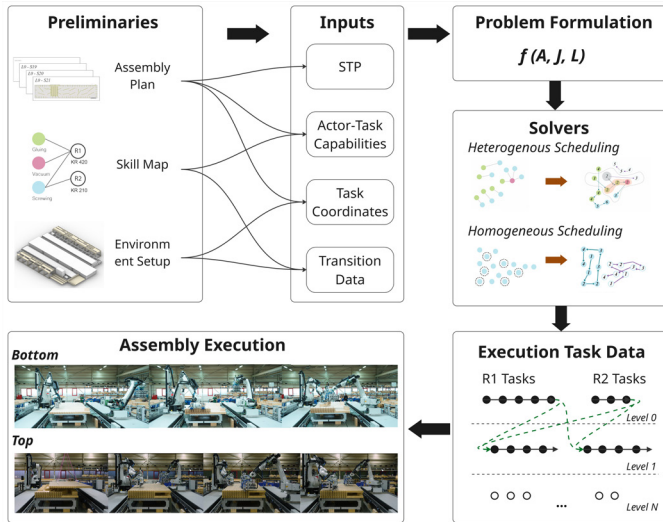


Fig. 2. The overall workflow of our work. We process preliminary data as inputs to our scheduling formulation, which is solved by two specific algorithms for heterogeneous and homogeneous problems. We convert the scheduled results into level-associated task objects in a semi-centralized control framework and validate our approach in a physical slab assembly.

We make the following specific contributions: **1) Constraint Programming (CP) scheduling formulation with level barriers** that embeds spatial safety and adhesive time windows as constraints while minimizing makespan and level compactness. **2) Two efficient scheduling algorithms:** an iterative temporal-relaxation algorithm with heuristic task clustering for the *heterogeneous* bottom assembly session; and a bi-level decomposition that translates the problem into set partitioning and vehicle routing problems (VRP) for the *homogeneous* top screw-pressing session. **3) Barrier-based, semi-centralized execution** that queues level-scoped subroutines to each robot, enabling parallel execution and sparse synchronization. **4) Physical validation and scalability study** via a full-scale 2.4 m × 6 m prototype featuring 34 internal beams, 15 levels, 108 subroutines, and 352 screws with a 2-robot setup, alongside a computational study over 8 design variants demonstrating scalability.

II. RELATED WORK

Our framework sits at the intersection of three major research threads within multi-robot systems—industrial man-

ufacturing, scheduling with spatiotemporal constraints, and asynchronous execution.

A. Multi-Robot Systems for Industrial Manufacturing

Early work established multi-robot assembly as a distinct industrial challenge with coordinated fixtures, tool sequencing, and collision-aware co-manipulation. The survey by Marvel *et al.* synthesizes strategies and metrics for industrial multi-robot assembly, covering fixtureless methods, synchronization schemes, and evaluation protocols [1]. In automotive spot-welding, Pellegrinelli *et al.* developed integrated approaches that jointly consider cell design and off-line motion planning [2], [3]. Beyond automotive, manufacturing tasks such as aircraft assembly [4] and brick layering [5] introduce heterogeneous collaboration with tool change and precedence requirements. Large multi-station industrial lines motivate hierarchical allocation and scheduling; Zhou *et al.* present a step-wise MS-MRTA method validated on a production case [6]. For broader assembly planning at the structure level, discrete optimization (ILP/IQP) has been used to co-design assembly sequences and resource usage for multi-robot teams [7].

B. Multi-Robot Task Scheduling with Spatiotemporal Constraints

Gerkey and Mataric differentiate allocation types by robot/task multiplicity and time extension [8]; Korsah *et al.* add interrelated utilities and constraints [9]; and Nunes *et al.* target precedence and time windows explicitly [10]. Exact/heuristic schedulers that couple space and time have been proposed, TERCIO [11] handles tight spatial interference and deadlines at scale. Under execution uncertainty, hierarchical methods decouple per-robot policies from conflict resolution while preserving time-window feasibility [12]. These works are grounded in temporal reasoning: Simple Temporal Networks (STNs) and their controllability/dispatchability properties [13], [14]. In practical settings, constraint programming (CP) offers compact encodings of precedence, resource sharing, and spatial safety using interval variables and global constraints. The CP-SAT solver [15], part of Google OR-Tools, supports scalable formulations for precedence, resource sharing, and spatial safety.

C. Asynchronous Multi-Robot Execution Frameworks

Bridging plans to execution, recent frameworks advocate partial-order, asynchronous dispatch to address execution uncertainty without dense time-stamped synchronization. APEX-MR post-processes task/motion plans to build a temporal plan graph that enables robust asynchronous execution for cooperative assembly and demonstrates substantial speedups on long-horizon tasks [16]. The Bidirectional Temporal Plan Graph (BTPG) further introduces switchable dependencies to reduce wait time while maintaining safety guarantees during execution [17]. At the middleware level, asynchronous multi-arm/multi-robot execution with continuous collision checking has been implemented in MoveIt 2, providing a practical substrate for independent yet safe trajectory execution under a

central scheduler [18]. While prior works like [16] enable asynchronous execution, our level-based barrier provides a novel, integrated mechanism that simplifies the execution framework by encoding collision avoidance directly into the schedule structure, thereby obviating the need for continuous collision checking or complex temporal plan graphs during execution.

III. PRELIMINARIES

The cooperative slab assembly for our multi-robot system provides a set of preliminaries for the scheduling model. They form the basis for the problem formulation in Sec. IV. They include the design of slab assembly plan, a library of manipulation skills and subprograms mapped to robots, and the environment setup with robot cell layout as a result of a co-design process to ensure the feasibility of fabrication.

A. Assembly Plan and Task Hierarchy

The timber slab assembly plan provides a list of elements $E = \{e_1, e_2, \dots, e_N\}$ for the bottom session tasks. Each element is associated with a set of task primitives, including gluing, picking, placing, and screw-pressing.

- **Element Types:** For the bottom session, elements can be either linear beams (manipulated by a screw effector with an integrated gripper) or planar crowns (handled by a vacuum gripper). Each element is associated with a list of glue lines and screws to be fastened.
- **Task Primitives:** The tasks for the bottom session are heterogeneous, involving a mix of picking, placing, gluing, and screw-pressing. For the top session, tasks are homogeneous, consisting solely of high-density screw-pressing to secure the top CLT plate.

The assembly plan specifies an unordered set of elements, subject to a small number of precedence constraints. For instance, a crown must be placed before its adjacent beams to facilitate proper slot-in placement. Additionally, for the top session, screws are divided into a **priority set** and a **reinforcement set**, where all screws in the priority set must be fastened before those in the reinforcement set.

B. Environment Setup and Cell Layout

Our fabrication platform consists of two industrial robots, a KUKA KR 420 R3330 (Robot 1) and a KUKA KR 210 R3100 (Robot 2), each mounted on a 10.7-meter linear axis. The robots are positioned on opposite sides of a central movable worktable to facilitate cooperative tasks. This setup specifies the kinematics and geometries of all robots and static obstacles, as well as the placement of the workpiece in the shared workspace. The cell co-design ensures every task pose is kinematically reachable by at least one robot, a critical assumption for our formulation.

C. Manipulation Skills and Subprograms

Each task primitive maps to a manipulation skill corresponding to a pre-defined, feedback-driven subprogram on the KUKA KRL controller. Our platforms feature a distinct

division of labor based on payload capacity: Robot 1 (KUKA KR 420) is equipped with a gluing effector, a heavy-payload vacuum gripper, and a screw effector with an integrated gripper, enabling it to perform all tasks including gluing, placing large crowns, and screw-pressing. In the meantime, Robot 2 (KUKA KR 210) is equipped solely with a screw effector with an integrated gripper; it is limited to placing linear beams and screw-pressing but utilizes a more powerful motor to reduce cycle time and maximize utilization. We collect average execution times for these skills and encode them into the scheduling model to minimize the estimated makespan.

IV. PROBLEM FORMULATION

In this section, we formulate the scheduling of multi-robot screw-press gluing as a **Constraint Programming (CP)** problem. Building upon the discussion in Sec. III, we derive the following key inputs for our CP model, as summarized in Fig. 3:

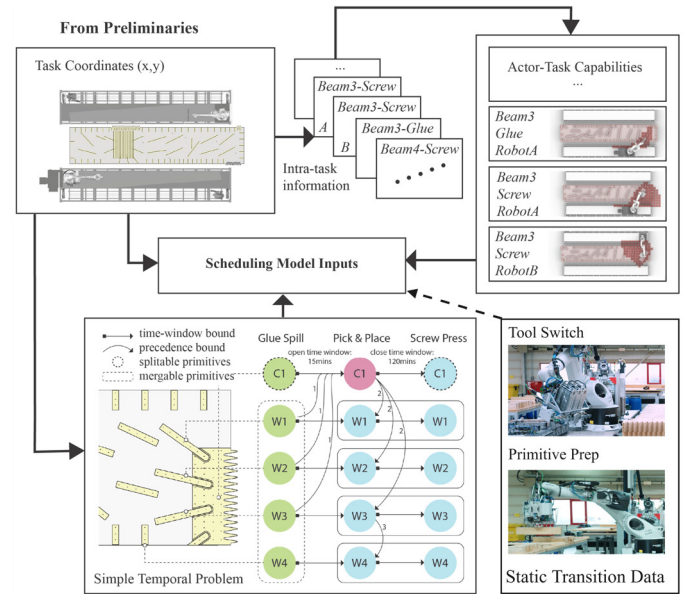


Fig. 3. Scheduling input summary.

- **Simple Temporal Problem (STP):** We model the temporal relationships between all tasks as a set of STPs, which define temporal constraints (lower and upper bounds) for task primitives. For each element e , the STP enforces that gluing must occur after placement but within the adhesive's open time, and that screwing must occur after placement but within the close time window. For elements i and j with a precedence requirement, a constraint is added to ensure that the placement of i precedes the placement of j .
- **Actor-Task Capabilities:** Each task is associated with a specific set of capabilities. This includes the estimated duration for each capable actor to complete the task, robot kinematic states at each task, and a binary representation of its spatial occupancy, crucial for collision avoidance.

- **Task Coordinate (x,y):** Each task primitive has one or two (only with gluing task) 2D coordinates in the shared workspace. This information is used to evaluate spatial proximity during the homogeneous screw scheduling, as discussed in Section III.
- **Transition and Tool Switch Information:** The model accounts for logistical overhead, including the time required for tool-switching and preparatory actions like screw loading and material picking. To simplify the spatial problem, these transition tasks are designed to occupy non-colliding areas. The total transition time (τ_k^{ij}) between two consecutive tasks is pre-calculated based on the chosen actor and the sequence of actions. For instance, the transition from a vacuum gripper placement task to a gluing task may include the time for storing the gripper, picking up the glue effector, and pressurizing the adhesive.

The problem is defined by the following inputs, variables, and constraints. Let T denote the set of task primitives and K the set of actors. We define a comprehensive set of temporal constraints \mathcal{T} that encapsulates both process precedence and adhesive time windows. Central to our formulation is the concept of a *level*, formally defined as a set of tasks spatiotemporally disjoint from other levels and synchronized by a global barrier. The decision variables include: the boolean assignments A_k^i for task i to actor k ; the boolean sequence variables J_k^{ij} (1 if task j immediately follows i on actor k); and the integer level index $\ell^i \in \mathbb{Z}_{\geq 0}$ for task i . The start/end times t_S^i, t_E^i are modeled as auxiliary variables whose values are propagated based on the main integer decisions. Parameters include the duration d_k^i and transition times ρ_k^{ij} . The model minimizes the makespan C_{\max} and the maximum level index L_{\max} , weighted by a user-defined factor λ , as follows:

$$\min C_{\max} + \lambda L_{\max}, s.t. \quad (1)$$

$$\sum_{k \in K} A_k^i = 1 \quad \forall i \in T \quad (2)$$

$$A_k^i = 1 \Rightarrow t_E^i \geq t_S^i + d_k^i \quad \forall i \in T, \forall k \in K \quad (3)$$

$$J_k^{ij} = 1 \Rightarrow t_S^j \geq t_E^i + \rho_k^{ij} \quad \forall i, j \in T, \forall k \in K \quad (4)$$

$$J_k^{ij} = 1 \Rightarrow \ell^j \geq \ell^i \quad \forall i, j \in T, \forall k \in K \quad (5)$$

$$L^{uv} \leq t_{\eta_v}^v - t_{\eta_u}^u \leq U^{uv} \quad \forall (u, v) \in \mathcal{T} \quad (6)$$

$$(A_a^i \wedge A_b^j \wedge R_{ab}^{ij} = 1) \Rightarrow \ell^i \neq \ell^j \quad \forall i, j \in T, \forall a, b \in K \quad (7)$$

$$(J_k^{ij} = 1 \wedge R_{\text{edge}}^{(ij)} = 1) \Rightarrow \ell^i \neq \ell^j \quad \forall i, j \in T, \forall k \in K \quad (8)$$

$$t_E^i \leq B_{\ell^i}, \quad B_k \leq B_{k+1} \quad (9)$$

$$\ell^i < \ell^j \Rightarrow B_{\ell^i} \leq t_S^j \quad (10)$$

$$C_{\max} \geq t_E^i, \quad L_{\max} \geq \ell^i, \quad \forall i \in T. \quad (11)$$

Constraints (2) and (3) encode assignment and processing duration. (4)–(5) enforce routing gaps and level monotonicity along the sequence J . (6) governs all temporal constraints

from \mathcal{T} , including precedence (where $L = 0, U = \infty$) and adhesive windows via bounded differences. Binary node (R_{ab}^{ij}) and edge ($R_{\text{edge}}^{(ij)}$) conflicts are precomputed via voxelized occupancy with safety margins, they parametrize (7)–(8) to lift spatial conflicts into level separation. (9)–(10) model timestamp synchronization via barrier variables B_{ℓ} . Since these barrier constraints directly bound task start times, Eq. (6) validates temporal window constraints inclusive of synchronization idles. To robustly accommodate execution uncertainty, we apply a safety buffer to the temporal bounds in Eq. (6) by lowering the upper bound U , ensuring the schedule remains valid even under stochastic delays. (11) defines the objectives.

V. SCHEDULING ALGORITHMS FOR SCREW-PRESS GLUING TIMBER SLAB ASSEMBLY

The formulation presented in the previous section specifies the feasibility and optimality criteria for our problem. However, a monolithic CP model faces significant scalability challenges, particularly for a complete slab assembly that may involve over 800 task primitives. Our key insight is that the timber slab assembly process consists of two fundamentally different scheduling problems.

The **bottom session** is a *sparse, heterogeneous, constraint-driven* problem. The number of decision nodes is moderate (under 80 in all test cases), but they are linked by complex temporal and logical constraints (e.g., glue windows, tool changes). Conversely, the **top session** is a *dense, homogeneous, routing-driven* problem, with hundreds of identical screw tasks where path efficiency and workload balance are dominant.

The computational study on Fig. 9(a) confirms that a direct CP model of this problem becomes intractable beyond 80-100 tasks. We therefore dictate a hybrid strategy (Fig. 4): for the constraint-rich bottom session that falls within this tractable range, we employ a direct CP method with primitive clustering and iterative temporal relaxation. For the dense top session, a decomposition approach is necessary and operable under a fixed number of levels.

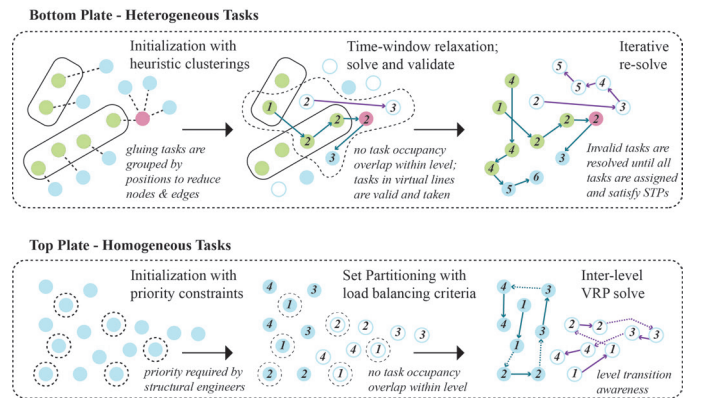


Fig. 4. Diagram illustration of our algorithms. Bottom Plate: Heterogeneous scheduling with clustering and iterative temporal relaxation. Top Plate: Homogeneous scheduling with decomposed Set Partitioning and Vehicle Routing Problem (VRP) solvers.

A. Bottom session: Primitive Clustering and Iterative Temporal Relaxation

The assembly of the bottom session involves a large number of discrete timber elements that must be glued, placed, and pressed onto a large CLT plate. The ideal sequence of operations to minimize tool changes would involve performing all gluing tasks first, followed by all placement tasks for large elements (using a vacuum gripper), and finally all screwing tasks for smaller linear elements (using a screw effector). Even the tight open-time window of the adhesive (e.g., 15 minutes at 20°C) makes a single round of such a sequence impossible, we can leverage the inherent efficiency of clustering tasks while systematically addressing the strict temporal windows as shown in Algorithm 1:

Algorithm 1 Iterative Clustering and Scheduling for Bottom Session

```

1: Input: Task Primitives  $T_{all}$ , Actor Set  $K$ , Temporal
   Constraints  $\mathcal{T}$ , Collision  $R$ 
2: Output: Schedule  $S_{final}$ 
3:  $(T_{active}, T_{deferred}) \leftarrow \text{InitializeBatches}(T_{all})$ 
4:  $\mathcal{T}_{relaxed} \leftarrow \text{ApplyRelaxation}(\mathcal{T})$ 
5:  $V \leftarrow \{\text{dummy}\}$ 
6: while  $V \neq \emptyset$  do
7:    $S_{skeleton} \leftarrow \text{SolveCP}(T_{active}, \mathcal{T}_{relaxed}, K, R)$ 
8:    $V \leftarrow \text{FindViolations}(S_{skeleton}, \mathcal{T})$ 
9:   if  $V \neq \emptyset$  then
10:     $T_{active} \leftarrow \text{SplitBatches}(T_{active}, V)$ 
11:   end if
12: end while
13:  $S_{final} \leftarrow \text{InsertDeferredTasks}(S_{skeleton}, T_{deferred}, R)$ 
14: Return  $S_{final}$ 

```

We first coalesce consecutive glue primitives into clustered batched nodes, reducing the number of decision variables and minimizing travel overhead. During this aggregation, we relax the temporal bounds associated with the batched glue tasks to permit a feasible initial route. Three critical screw primitives (the center and end screws) are merged into each placement task to ensure immediate pressing and prevent slippage from the glue line. All remaining screws are temporarily deferred. The relaxed instance is solved; any elements whose glue-place chains violate the strict temporal windows are identified, and their glue primitives are detached to form smaller batches. This process repeats on the unresolved subproblem until every element satisfies its exact (unrelaxed) bounds. Finally, we reinsert the deferred screws under the full set of assignment, routing, and level constraints. These deferred tasks are placed into idle slots between levels with fixed actor assignments, maintaining spatial safety.

B. Top session: Set Partitioning and Vehicle Routing Problem

The top session assembly involves applying glue to the sandwiched timber elements before a top CLT plate is placed via a crane, followed by a dense screw pattern to ensure a secure bond. The multi-robot collaboration must be highly

efficient to complete all tasks before the adhesive’s strict close time window (approximately 120 minutes under standard conditions).

To address this challenge, we leverage the homogeneous nature of the screwing tasks and reframe the problem as an assignment problem aimed at workload balance. Instead of minimizing makespan as the primary objective in (1), we adopt a load-balancing objective inspired by the **Multway Number Partitioning Problem**. This approach seeks to minimize the workload difference between the actors within each level while minimizing the total number of levels to maximize parallelism. For a system with N actors, the general objective is defined as:

$$g = \max_{i \in K} \sum_{j \in T_i} w_j - \min_{i \in K} \sum_{j \in T_i} w_j \quad (12)$$

where K is the set of actors, T_i is the set of tasks assigned to actor i , and w_j is the processing time of task j . This equilibrium is achievable for screwing tasks because each primitive’s duration is nearly invariant for a given actor, regardless of its sequence, as screw loading occurs in parallel with travel. For our setup in the top session, the KUKA KR210’s screwing cycle is 11s (6s screwing and 5s loading), and the KUKA KR420’s is 16s (11s screwing and 5s loading), indicating a distribution of 16:11 leads to temporal optimality. We allow a small time difference threshold, δ :

$$g \leq \delta \quad (13)$$

as a constraint for more search space for the solver.

We solve this problem using a two-stage approach, as detailed in Algorithm 2:

- 1) **Task Assignment (Multway Number Partitioning Problem):** We begin by attempting to assign all tasks with a minimal number of levels, $L_{max} = k_{min}$, where $k_{min} \geq 2$ to enable multi-robot cooperation. A CP model is built that satisfies the processing, sequencing, level-monotonicity, barrier, and STP constraints described in Sec. IV. To enlarge the search space, we temporarily relax the edge-conflict constraints while retaining the static occupancy constraints. We also introduce a heuristic to prevent isolated assignments: each task assigned to a level must have at least one other task on that level within a distance threshold δ_{dist} . This is enforced via additional binary variables linking proximity to assignment; we omit the detailed encoding for brevity. We relax constraint (2) to allow for partial assignment and maximize the number of assigned tasks. Starting with the priority screw set, we solve for $k_{min} = 2$. Any unassigned tasks are added to the reinforcement set and solved together with $k_{min} = 2$. We find in Sec. VII that this fixed 2-level solve yields at least 95% task assignment within a modest time budget.
- 2) **Routing Refinement (Vehicle Routing Problem):** With tasks assigned to specific actors and levels, we solve a multi-vehicle routing problem on each level to ensure

collision-free travel. A local search algorithm is implemented on each level, where the start points of each actor's path are fixed by the previous level's solution. This makes each sub-problem a series of fast and scalable vehicle routing problems with fixed candidates. The objective is to minimize the total travel distance, while satisfying the edge-based spatial conflict constraint (8) that was relaxed during the assignment phase. The unassigned screws from the first stage are integrated post-routing by relaxing the workload equilibrium constraint (13) and dynamically inserted into existing scheduled levels or added to a minimal number of new levels (e.g., 1–2) to minimize the final makespan.

Algorithm 2 Bi-Level Scheduling for Top Session Screwing

- 1: **Input:** Task sets $[T_{\text{prio}}, T_{\text{reinf}}]$, Actors K , Collision data R , Load-balance threshold δ
 - 2: **Output:** Completed Schedule S_{final} (Task Assignments, Orders, Levels)
 - 3: Initialize $S_{\text{assignments}} \leftarrow \emptyset$, $T_{\text{unassigned}} \leftarrow \emptyset$, $L_{\text{max}} \leftarrow 2$
 - 4: **for each** T_{set} in $[T_{\text{prio}}, T_{\text{reinf}}]$ **do**
 - 5: $T_{\text{current_solve}} \leftarrow T_{\text{set}} \cup T_{\text{unassigned}}$
 - 6: $S_{\text{partial}} \leftarrow \text{SetPartition}(T_{\text{current_solve}}, L_{\text{max}}, \delta)$
 - 7: $T_{\text{newly_assigned}} \leftarrow \text{GetAssignedTasks}(S_{\text{partial}})$
 - 8: $T_{\text{unassigned}} \leftarrow T_{\text{current_solve}} \setminus T_{\text{newly_assigned}}$
 - 9: $S_{\text{assignments}} \leftarrow S_{\text{assignments}} \cup S_{\text{partial}}$
 - 10: **end for**
 - 11: $S_{\text{initial_routes}} \leftarrow$ Initialize empty schedule
 - 12: $p_0 \leftarrow -1$ {Start index at level 0 is not forced}
 - 13: **for each level** l from 0 to $L_{\text{max}} - 1$ **do**
 - 14: $T_l \leftarrow$ Tasks in $S_{\text{assignments}}$ at level l
 - 15: $T_{l+1} \leftarrow$ Tasks in $S_{\text{assignments}}$ at level $l + 1$
 - 16: $S_l, p_{l+1} \leftarrow \text{SolveVRP}(T_l, T_{l+1}, p_l)$
 - 17: Add S_l to $S_{\text{initial_routes}}$
 - 18: **end for**
 - 19: $S_{\text{final}} \leftarrow \text{InsertUnassigned}(S_{\text{initial_routes}}, T_{\text{unassigned}})$
 - 20: **Return** S_{final}
-

VI. LEVEL-BASED TASK EXECUTION

We employ a semi-centralized control framework, *Twico* [19], to orchestrate the system. While general-purpose middlewares like MoveIt 2 [18] offer asynchronous capabilities, *Twico* is specifically tailored to dispatch scheduled primitives as parameterized KUKA subprogram calls with level-based barrier protocol for spatial safety, ensuring seamless alignment with our task scheduling model.

Before execution, the scheduled primitives are post-processed into executable task objects. To optimize efficiency, consecutive tasks of the same type (e.g., multiple screwing points or continuous glue lines) are batched into single subprogram calls, and global task poses are transformed into each robot's local coordinate frame. These objects, defined within a unified task schema [20] including task ID, level index, and subprogram parameters, are stored on the central server and dispatched to parallel execution threads.

In the physical layer, each robot's controller interprets the task parameters into KRL subroutines and executes the required motions and skills. The controller itself monitors primitive completion through internal signal-based commands and reports status back to the server. The server functions as a global barrier manager: it tracks the completion status of the current level and only releases tasks for level $l + 1$ once all actors have confirmed the completion of level l . This protocol acts as a robust control law that absorbs timing uncertainty and ensures collision-free operation without the need for dense timestamp synchronization.

VII. FRAMEWORK VALIDATION AND PERFORMANCE

We validate our framework through three studies: first, by successfully fabricating a full-scale, complex architectural prototype; second, by highlighting the robustness embedded in our framework's design; and third, by demonstrating the scalability and necessity of our tailored hybrid algorithm on a representative set of design variants.

A. Physical Prototype Fabrication

The 2.4 m × 6 m slab (34 elements; 352 screws) executed collision-free under adhesive windows using 15 levels (11 bottom, 4 top) and 108 subroutines. Realized makespans were 28/20 min vs. model estimates 25/19 min (both within 5%). The physical execution, while successful, highlighted the importance of our level-based barrier mechanism. Minor variations in material handling time and screw-driving duration (up to 5% deviation from the model's estimates) were naturally absorbed within each level's asynchronous execution window, without requiring re-planning or risking collisions. This real-world validation confirms the framework's inherent robustness against the uncertainties of physical fabrication.

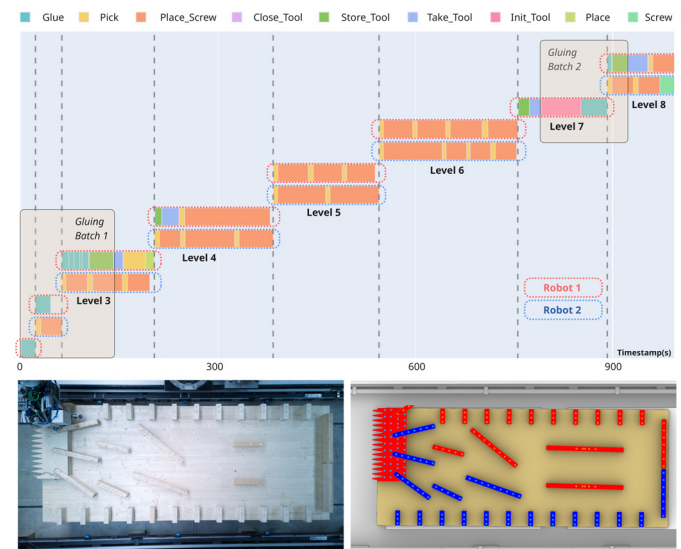


Fig. 5. Overview of bottom session of our prototype. Top: Execution Gantt Chart from the first 990s up until level 8. Bottom left: Picture of the full assembly. Bottom right: Task assignment diagram to the robots.

Fig. 5 summarizes the bottom assembly session involving heterogeneous task cooperation, such as gluing application, placement of large timber crowns and small timber beams, as well as screwing. Fig. 6 provides a detailed analysis on two scheduled levels from the bottom session. Fig. 7 demonstrates the scheduling and execution of the top screwing session, which contains 4 levels for 178 screws while maintaining workload balance within each level.




Fig. 6. Analysis of Level 2 (left) and Level 8 (right) of the bottom session. Specifically, Level 2 consists gluing, tool switching, and pick-and-place of a crown element for Robot 1, while Robot 2 is able to execute pick-and-place and screwing tasks for 3 smaller beams, leading to efficient task distribution.

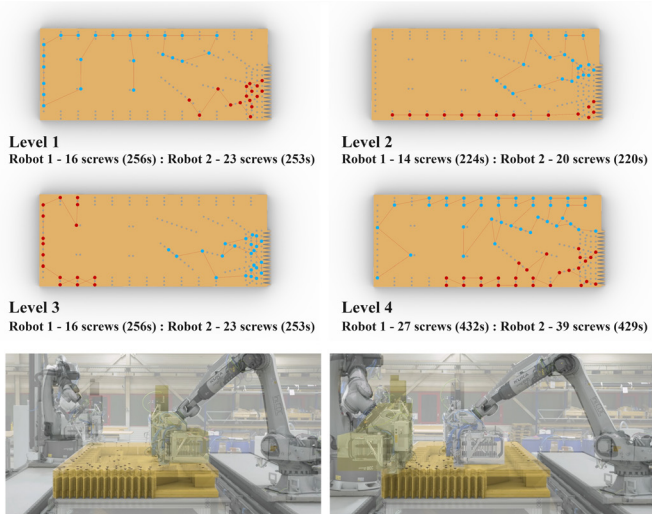


Fig. 7. Scheduling result (top) and execution process with simulation overlay (bottom) of the top session. A total of 4 levels are scheduled with an estimated 98.8% robot utilization ratio.

B. Robustness to Execution Variability

We tested schedule robustness by injecting zero-mean i.i.d. noise into per-primitive durations and tool-switch times (Gaussian and bounded uniform), re-simulating the controller-dispatcher loop in a virtual environment provided by KUKA.OfficeLite. Across noise levels up to 5%, all runs respected level barriers and adhesive windows. Beyond simulated robustness, our framework incorporates two key mechanisms at the design level to handle real-world uncertainties:




Fig. 8. Four of the eight representative slab geometries used for the computational study, illustrating the significant diversity in element count (15 to 74) and screw density (41 to 743). A detailed table of per-test metrics is available in the supplementary material.

a) *Proactive Safety Buffers*: The scheduler operates on conservative time windows. For instance, a hard 15-minute adhesive open-time is translated into a softer 13.5-minute scheduling constraint (a 10% safety buffer), ensuring minor delays do not compromise process integrity.

b) *Inherent Fault Tolerance*: The level-based barrier protocol provides a powerful fault-tolerance mechanism. A task failure or hardware malfunction on one robot simply pauses the system at a safe, collision-free state upon completion of the current level. This creates a natural checkpoint for human intervention without jeopardizing the other robot, a critical feature for long-horizon autonomous construction.

C. Scalability and Performance

To analyze the framework’s performance across a range of complexities, we evaluated it on 8 distinct slab geometries sourced from a real-world building project (Fig. 8). These cases represent a significant diversity in scale, with 15 to 74 internal elements and 41 to 743 screws.

Our central finding is that a monolithic CP formulation becomes sensitive to problem size, shown in Figure 9(a). The monolithic approach for the top session problem fails to find a solution within the 1800s time limit for instances with more than 300 tasks (Slabs E-H). This data-driven insight is the direct justification for our hybrid algorithm strategy for top and bottom sessions.

a) *For the Bottom Session*: After our primitive clustering, the number of decision nodes is identical to the element count (max 74 in Test H). This count falls within the tractable range for a direct CP solver. Therefore, our heterogeneous algorithm model assignment, sequence and levels as variables directly in the CP model with only relaxed temporal constraints, as it optimally minimizes the level count under the spatial disjunction of the sub-problem. Results from Figure 9(a) indicates that the problem is still solvable at its maximum size under the context of the studied cases.

b) *For the Top Session*: The number of tasks (screws) is an order of magnitude larger, far exceeding the CP solver’s tractable limit. This necessitates the decomposition into a set-partitioning problem for load-balancing and a VRP for routing. As shown in Figure 9(b), this tailored approach scales effectively, solving even the largest instances within 460 seconds.

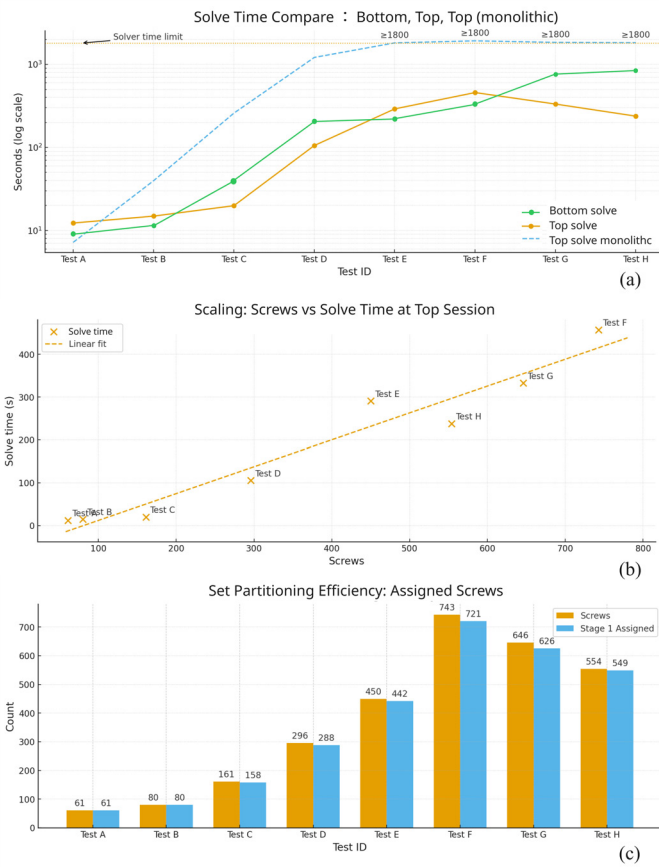


Fig. 9. Scaling, ablation and makespan overview: (a) solve time compare: bottom, top and top(monolithic); (b) top-session screws vs. solve time; (c) efficiency at set partition stage for top session solve.

We also notice from Figure 9(c) that the initial partitioning model manages to assign at least 98% of the screws to the load-balanced levels, making the insertion problem in the final step of Algorithm 2 easier.

VIII. CONCLUSION

We have introduced LASER, a level-based scheduling and execution framework that enables time-critical process of multi-robot screw-press gluing for large-scale timber slab assembly. The resulting schedule allows for asynchronous execution within levels and single synchronization between them. By transforming stringent adhesive windows and spatial safety requirements into a robust, level-based execution structure, our work paves the way for automated robotic fabrication of complex, sustainable timber architecture.

By embedding collision avoidance directly into the schedule, our system safely operates without the computational overhead of continuous collision checking. While such a mechanism efficiently coordinates sparse tasks, highly crowded workspaces may decrease the collaboration efficiency or cause levels to degenerate into serialized single-robot tasks and reduce parallelism. Future work could explore scalable

temporal-based decomposition methods for heterogeneous collaboration with larger element counts and hybridize level-barriers with temporal networks to reduce idle times.

REFERENCES

- [1] J. A. Marvel, R. Bostelman, and J. Falco, "Multi-robot assembly strategies and metrics," *ACM Computing Surveys*, vol. 51, no. 1, pp. 14:1–14:32, 2018.
- [2] S. Pellegrinelli, N. Pedrocchi, L. Molinari Tosatti, A. Fischer, and T. Tolio, "Multi-robot spot-welding cells: An integrated approach to cell design and motion planning," *CIRP Annals*, vol. 63, no. 1, pp. 17–20, 2014.
- [3] S. Pellegrinelli, N. Pedrocchi, L. M. Tosatti, and T. Tolio, "Multi-robot spot-welding cells for car-body assembly: Design and motion planning," *Robotics and Computer-Integrated Manufacturing*, vol. 44, pp. 97–116, 2017.
- [4] V. Tereshchuk, S. Devasia, and A. G. Banerjee, "A scheduling method for multi-robot assembly of aircraft structures with soft task precedence constraints," *Robotics and Computer-Integrated Manufacturing*, vol. 71, p. 102145, 2021.
- [5] M.-N. Stamatopoulos, S. Velhal, A. Banerjee, and G. Nikolakopoulos, "Safety-aware optimal scheduling for autonomous masonry construction using collaborative heterogeneous aerial robots," in *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2025, pp. 4889–4896.
- [6] B. Zhou, R. Zhou, Y. Gan, F. Fang, and Y. Mao, "Multi-robot multi-task cooperative spot welding task allocation based on stepwise optimization: An industrial case study," *Robotics and Computer-Integrated Manufacturing*, vol. 73, p. 102197, 2022.
- [7] P. Culbertson, S. Bandyopadhyay, and M. Schwager, "Multi-robot assembly sequencing via discrete optimization," Technical Report, Stanford MSL, 2019.
- [8] B. P. Gerkey and M. J. Mataric, "A formal analysis and taxonomy of task allocation in multi-robot systems," *The International Journal of Robotics Research*, vol. 23, no. 9, pp. 939–954, 2004.
- [9] G. A. Korsah, A. Stentz, and M. B. Dias, "A comprehensive taxonomy for multi-robot task allocation," *The International Journal of Robotics Research*, vol. 32, no. 12, pp. 1495–1512, 2013.
- [10] E. Nunes, M. Manner, H. Mitiche, and M. Gini, "A taxonomy for task allocation problems with temporal and ordering constraints," *Robot. Auton. Syst.*, vol. 90, no. C, p. 55–70, Apr. 2017.
- [11] M. C. Gombolay, R. J. Wilcox, and J. A. Shah, "Fast scheduling of robot teams performing tasks with temporospatial constraints," *IEEE Transactions on Robotics*, vol. 34, no. 1, pp. 220–239, 2018.
- [12] S. Choudhury, J. K. Gupta, M. J. Kochenderfer, D. Sadigh, and J. Bohg, "Dynamic multi-robot task allocation under uncertainty and temporal constraints," in *Robotics: Science and Systems (RSS)*, 2020.
- [13] R. Dechter, I. Meiri, and J. Pearl, "Temporal constraint networks," *Artificial Intelligence*, vol. 49, no. 1–3, pp. 61–95, 1991.
- [14] P. Morris, "Temporal dynamic controllability revisited," in *Proceedings of AAAI*, 2005, pp. 1193–1198.
- [15] L. Perron and F. Didier, "Cp-sat," Google. [Online]. Available: https://developers.google.com/optimization/cp/cp_solver/
- [16] P. Huang, R. Liu, S. Aggarwal, C. Liu, and J. Li, "Apex-mr: Multi-robot asynchronous planning and execution for cooperative assembly," in *Robotics: Science and Systems (RSS)*, 2025.
- [17] Y. Su, R. Veerapaneni, and J. Li, "Bidirectional temporal plan graph: Enabling switchable passing orders for more efficient multi-agent path finding plan execution," in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2024, pp. 17 559–17 566.
- [18] P. Stoop, T. Ratnayake, and G. Toffetti, "A method for multi-robot asynchronous trajectory execution in moveit2," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 2024, pp. 17 694–17 700.
- [19] L. Skoury, S. Leder, A. Menges, and T. Wortmann, "Digital twin architecture for the aec industry : A case study in collective robotic construction," *Proceedings of the ACM/IEEE 27th International Conference on Model Driven Engineering Languages and Systems*, 2024.
- [20] L. Skoury, S. Tremel, N. Opgenorth, F. Amtsberg, H. J. Wagner, A. Menges, and T. Wortmann, "Towards data-informed co-design in digital fabrication," *Automation in Construction*, vol. 158, p. 105229, 2024.