ISSN 1392-3730/eISSN 1822-3605

VILNIUS TECH
Vilnius Gediminas Technical University

JOURNAL of CIVIL ENGINEERING and MANAGEMENT

2026
Volume 32
Issue 7
Pages 973–991
https://doi.org/10.3846/jcem.2026.28639

TOPOLOGICALLY ADAPTIVE ANT COLONY OPTIMIZATION FOR SCHEDULING MEGA-STRUCTURE PILE FOUNDATIONS WITH ENDOGENOUS SPATIO-TEMPORAL CONSTRAINTS

Jingbo HAO[1], Mengxi ZHANG[1,2]✉, Qiaoling MIN[1], Renchao WANG[1], Xuesong XIANG[3], Libin HAN[3], Mingchao LI[1,2]

[1]*State Key Laboratory of Hydraulic Engineering Intelligent Construction and Operation, Tianjin University, Tianjin 300350, China*
[2]*School of Water Resources and Hydropower Engineering, Xinjiang University of Technology, Xinjiang 848023, China*
[3]*Yunnan Airport Group Co., Ltd., Kunming 650220, China*

Article History:
- received 11 April 2026
- accepted 1 June 2026

Abstract. Pile foundation scheduling of super-large engineering structure presents significant concerns and challenges due to the dynamically triggered and evolved vibration exclusion zones generated by impact boring. However, conventional resource-constrained project scheduling theories often exhibit spatial blindness, resulting in cascading construction critical path delays. Therefore, this study proposes a Topologically Adaptive Ant Colony Optimization (TA-ACO) algorithm embedded within a state-dependent simulation environment. The algorithm employs a dual-colony cooperative mechanism to effectively decouple the adversarial resource coupling between piling crews and grouting crews. Specifically, a novel look-ahead heuristic strategy is introduced, guiding crews to maximize spatial sparsity, which is analogous to solving the Maximum Independent Set problem and avoiding topological deadlocks. Computational experiments demonstrate that TA-ACO effectively mitigates spatial fragmentation and reduces the project makespan by 21.5% to 28.3% across different large-scale instances, compared to the industry-standard greedy strategy. The engineering application results of large airports show that the proposed algorithm compressed the construction period of a critical zone, which includes 97 pile foundations, from 47 days to 34 days, securing a valuable 13-day lead time for subsequent superstructure assembly.

Keywords: intelligent construction, project scheduling, pile foundation, ant colony optimization, decision-dependent uncertainty, spatio-temporal constraints.

✉Corresponding author. E-mail: ***zhangmx@tju.edu.cn***

1. Introduction

In mega-infrastructure projects, such as integrated airport hubs and high-speed rail terminals, project managers must shorten construction schedules while satisfying strict spatio-temporal restrictions in subsurface operations. Conventional project management theory, rooted in the classical Resource-Constrained Project Scheduling Problem (RCPSP) formally introduced in the 1960s (Mozhdehi et al., 2024; Hartmann & Briskorn, 2022; Ding et al., 2023), predominantly postulates that time is the primary decision variable while renewable resources constitute the limiting bottlenecks. These models are effective for representing precedence relations, crew capacities, and equipment limitations, where the feasibility of an activity is mainly determined by resource availability and temporal logic. However, they usually treat the physical workspace as a fixed background condition rather than as a decision-sensitive constraint that can alter the feasibility of subsequent activities. Furthermore, this theoretical framework frequently exhibits what has been termed spatial blindness, wherein physical space is regarded merely as a static container for activities rather than a constraint (Li et al., 2025; Zhang et al., 2024). Unlike superstructure assembly, in the domain of Pile Foundation Construction Scheduling Problem (PFCSP), physical space must be reconceptualized not as a passive vessel, but as a scarce, depletable, and dynamic resource that dictates the feasibility of execution sequences.

The problem undergoes a fundamental change when geotechnical vibration constraints are introduced. Specifically, impact-bored piling generates intense shear waves that can structurally compromise adjacent concrete during

Copyright © 2026 The Author(s). Published by Vilnius Gediminas Technical University

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (*http://creativecommons.org/licenses/by/4.0/*), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

its critical early-age curing phase (Hong & Park, 2015; Zhu et al., 2025). This physical phenomenon manifests managerially as a highly unusual constraint morphology: Endogenous Spatio-Temporal Constraints (ESTC). Under ESTC, forbidden zones are not pre-defined static areas; rather, they are dynamically triggered, positioned, and timed solely by the scheduling decisions themselves. Consequently, the scheduler is not merely navigating a static maze but is dynamically constructing the walls of the maze with every task assigned, creating a form of decision-dependent uncertainty where a current action instantaneously reshapes the feasible search space for all future tasks.

Despite the theoretical maturity of the classical RCPSP, its application to this domain is fundamentally limited by its reliance on idealized assumptions. As explicitly noted in the comprehensive survey by Hartmann and Briskorn (2022), classical models are typically predicated on exogenous constraints, limits that exist prior to project commencement, rendering them structurally unable to capture the decision-dependent dynamics of ESTC. More recent extensions, such as the Multi-Skilled RCPSP, have improved classical scheduling models by considering workforce flexibility, skill substitution, and collaborative resource allocation. For example, recent heuristic and meta-heuristic studies have focused on assigning heterogeneous workers to tasks more efficiently under skill-related constraints (Vermeire & Vanhoucke, 2025; Hu et al., 2025). However, these models mainly address the substitutability of resources, whereas the PFCSP is governed by a different type of coupling: piling crews generate vibration exclusion zones, while grouting crews help release these constraints by completing the subsequent curing-related process. The PFCSP is not defined by resource flexibility, but by a strict adversarial coupling of heterogeneous resources, a complexity that transcends standard multi-skilled modeling (Chen et al., 2026). In this context, piling crews act as constraint generators (creating vibration zones), while grouting crews act as constraint relievers (completing tasks to eventually lift restrictions). This bidirectional interaction breaks the logic of traditional priority rules, necessitating a specialized solver capable of managing dynamic topology rather than static assignment.

To precisely position the research gap, we categorize the spatio-temporal constraints found in the literature into a hierarchy based on their dependency on scheduling decisions.

The first level comprises Static Constraints, a category extensively studied in manufacturing and site layout planning. In these domains, such as block assembly in shipbuilding (Ge & Wang, 2021; Kwon & Lee, 2015) or precast site layout (Yang et al., 2024), exclusion zones are deterministic and fully defined prior to project commencement (Abotaleb et al., 2016; Lei et al., 2026). Although computationally demanding, the solution space in these problems remains fixed, allowing algorithms to optimize within a static geometric boundary. In such settings, the spatial conflict relationships can be largely identified before scheduling, and a candidate sequence is evaluated against a pre-defined conflict structure. This assumption is not held in PFCSP, because the conflict structure is activated only after specific piling and curing operations are scheduled.

The second level involves Exogenous Dynamic Constraints. Here, resource availability or spatial accessibility fluctuates over time, but these variations are dictated by external factors independent of the scheduler's actions. Quintessential examples include calendar-based unavailability (Pfitzner et al., 2025; Zhang et al., 2020), weather-induced shutdowns in offshore construction (Kerkhove & Vanhoucke, 2017), or tidal windows in port scheduling (Zhong et al., 2022). While the feasible domain is dynamic, its evolution follows a predictable or stochastic pattern that is exogenous to the optimization process itself. The scheduler therefore responds to a changing environment, but does not create the changing constraints. By contrast, in impact-bored piling, each scheduling decision may itself generate a temporary exclusion zone and thereby reshape the feasible domain for subsequent decisions.

The third level, which is the focus of this study, involves Endogenous Decision-Dependent Uncertainty. This represents the most complex constraint morphology, where the emergence, location, and duration of constraints are direct consequences of the scheduling decisions. While decision-dependent uncertainty has garnered attention in diverse fields, such as patient no-show probabilities in appointment scheduling (Cheng et al., 2024; Homem-de-Mello et al., 2022), risk profile changes in R&D portfolios (Han et al., 2024; Jing & Fu, 2026), and cascading failure scenarios in power grid hardening (He et al., 2025; Ma et al., 2025) – these studies mainly show that decisions can influence future uncertainty sets or system states through changes in probabilities, risks, capacities, or scenario structures. In contrast, the PFCSP involves decision-induced physical exclusion zones with explicit locations, durations, and hard safety implications. Existing approaches, such as those for construction scheduling (Pfitzner et al., 2025; Salhab et al., 2026), primarily focus on the resolution of stochastic parameters or non-anticipativity. In contrast, the PFCSP requires managing the deterministic generation of transient exclusion zones.

Within construction scheduling, the closest line of research concerns workspace interference and spatial conflict management (Tao et al., 2018), for example, considered location and congestion effects in space-time repetitive project scheduling, which provides an important basis for incorporating spatial interactions into scheduling models. However, such interference is generally formulated as a soft congestion cost or as a multi-objective trade-off between workspace utilization and schedule efficiency. This treatment is suitable when limited interference can be tolerated or penalized. In contrast, vibration interference in impact-bored piling is a hard safety constraint: once a pile enters the early-age curing period, nearby piling operations must be prohibited rather than simply penal-

ized. Similarly, BIM-based automated planning methods, such as proposed by Ostapska et al. (2025), are effective for detecting geometric collisions or sequencing operations based on known spatial objects, but they generally address pre-defined or directly observable spatial conflicts. They do not explicitly model exclusion zones that are endogenously triggered, propagated, and released by scheduling decisions.

Motivated by the above gap, this study formulates the PFCSP as a scheduling problem with heterogeneous crews and endogenously generated spatio-temporal exclusion zones. To solve this problem, we develop a Topologically Adaptive Ant Colony Optimization (TA-ACO) algorithm embedded within a state-dependent simulation environment, so that the feasible workspace can be updated after each scheduling decision. Unlike generic metaheuristics or standard ACO variants applied to RCPSP (Ding et al., 2023), TA-ACO is engineered with deep spatio-temporal insight: it employs a novel look-ahead factor that guides piling crews to maximize spatial sparsity, analogous to solving the Maximum Independent Set problem, while directing grouting crews to prioritize constraint relaxation. To validate the proposed approach, we implemented the algorithm within the scheduling management of a real-world mega-airport hub project in Southwest China. Computational experiments demonstrate that the TA-ACO reduces the project makespan by up to 28.3% against the industry-standard greedy strategy. Microscopic analysis reveals that this performance gain stems from the algorithm's capability to prevent spatial fragmentation, thereby maintaining high active workfront density throughout the project lifecycle. Furthermore, in a field application within a critical pilot zone, the optimized schedule successfully compressed the construction period from 47 days to 34 days, securing a valuable 13-day lead time for subsequent superstructure assembly.

The remainder of this paper is organized as follows. Section 2 formulates the PFCSP using a Mixed-Integer Linear Programming model. Section 3 details the proposed TA-ACO algorithm. Section 4 presents computational experiments and microscopic analysis based on a real-world project. Section 5 discusses the mechanism of spatial fragmentation and managerial implications.

2. Problem formulation

To provide a precise and unambiguous definition of the PFCSP, this chapter develops a comprehensive Mixed-Integer Linear Programming (MILP) model. This model formalizes the problem's variables, objective function, and the full set of operational constraints. By capturing the problem's structure with mathematical rigor, this formulation not only clarifies its inherent complexity but also provides the foundation for the exact solution method used for performance benchmarking in our computational experiments.

2.1. Model components and notation

The mathematical model is constructed upon a defined set of indices, parameters, and decision variables. These components are detailed in Tables 1–3.

2.2. Mathematical model

The primary goal of the PFCSP is to complete the entire construction project as quickly as possible. This is achieved by minimizing the makespan, which is defined as the finish time of the last-completed task in the project. The objective function is therefore:

$$Minimize\ C_{\max}. \tag{1}$$

The minimization of the objective function is subject to the following sets of constraints that govern task assignment, timing, precedence, crew routing, and the unique spatio-temporal interactions.

$$\sum_{c\in\mathcal{C}_p} x_{jc} = 1, \quad \forall j \in \mathcal{T}_p; \tag{2}$$

$$\sum_{c\in\mathcal{C}_g} x_{jc} = 1, \quad \forall j \in \mathcal{T}_g; \tag{3}$$

$$f_j = st_j + d_j, \ \ \forall j \in \mathcal{T}; \tag{4}$$

$$C_{\max} \geq f_j, \quad \forall j \in \mathcal{T}; \tag{5}$$

$$st_{g(j)} \geq f_j, \quad \forall j \in \mathcal{T}_p; \tag{6}$$

$$\sum_{j\in\mathcal{T}} y_{0jc} = 1, \quad \forall c \in \mathcal{C}; \tag{7}$$

$$\sum_{j\in\mathcal{T}} y_{j0c} = 1, \quad \forall c \in \mathcal{C}; \tag{8}$$

$$\sum_{i\in\mathcal{T}_0, j\neq i} y_{ijc} = x_{jc}, \ \forall j \in \mathcal{T}, \ \forall c \in \mathcal{C}; \tag{9}$$

$$\sum_{j\in\mathcal{T}_0, j\neq i} y_{ijc} = x_{ic}, \ \forall i \in \mathcal{T}, \ \forall c \in \mathcal{C}; \tag{10}$$

$$st_j \geq f_i + T_{travel}(i,j) - M\cdot\left(1 - y_{ijc}\right), \ \forall i \in \mathcal{T}_0, \ \forall j \in \mathcal{T}, \ i \neq j, \ \forall c \in \mathcal{C}; \tag{11}$$

$$f_i \leq st_{g(h)} + M\cdot\left(1 - z_{ih}\right), \ \forall i,h \in \mathcal{T}^p, \ \ i \neq h, \ D_{ih} = 1; \tag{12}$$

$$st_i \geq f_{g(h)} + T^{cure} - M(1 - z_{ih}), \ \forall i,h \in \mathcal{T}^p, \ i \neq h, \ D_{ih} = 1. \tag{13}$$

The objective function (1) seeks to minimize the project makespan, which represents the total time required to complete all construction tasks. Assignment constraints (2) and (3) ensure that each piling task is assigned to exactly one piling crew, and each grouting task to exactly one grouting crew. Task timing and makespan constraints are established by (4) and (5). Constraint (4) defines the relationship between a task's start time, duration, and finish time, while constraint (5) links the makespan to the finish time of all individual tasks. The technological precedence

Table 1. Set definitions

Symbol	Description
$\mathcal{P}$	Set of piles, indexed by p.
$\mathcal{T}_p$, $\mathcal{T}_g$	Sets of piling tasks and grouting tasks, respectively, indexed by i, j.
$\mathcal{T}$	Sets of all real tasks.
$\mathcal{T}_0$	Set of all tasks including the virtual depot.
$\mathcal{C}_p$, $\mathcal{C}_g$, $\mathcal{C}$	Sets of piling crews, grouting crews, and all crews, respectively, indexed by c.

Table 2. Parameters

Parameters	Description
d_j	The processing time of each task $j \in$ T.
T_{travel} (i, j)	The travel time for any crew between the location of task $i \in \mathcal{T}_0$ and task $j \in \mathcal{T}_0$.
T_{cure}	The mandatory concrete curing time that must elapse after a grouting task is completed.
D_{safe}	The minimum safety distance that must be maintained for any piling activity.
$g(j)$	A mapping function, where for each piling task $j \in \mathcal{T}_p$, $g(j)$ denotes its corresponding grouting task in $\mathcal{T}_g$.
D_{ih}	A pre-calculated binary parameter for each pair of a piling task $i \in \mathcal{T}_p$ and a grouting task $j \in \mathcal{T}_g$. D_{ih} = 1 if the physical distance between the piles associated with tasks i and j is less than or equal to D_{safe}, and D_{ih} = 0 otherwise.
M	A sufficiently large positive constant.

Table 3. Decision variables

Variables	Description
st_j	The start time of each task $j \in \mathcal{T}$.
f_j	The finish time of each task $j \in \mathcal{T}$.
$C_{\max}$	The makespan of the project, defined as the completion time of the last task.
x_{jc}	An assignment variable, where x_{jc} = 1 if task $j \in \mathcal{T}$ is assigned to crew $c \in \mathcal{C}$, and 0 otherwise.
y_{ijc}	A sequencing variable, where y_{ijc} = 1 if crew $c \in \mathcal{C}$ performs task $j \in \mathcal{T}$ immediately after $i \in \mathcal{T}_0$, and 0 otherwise.
z_{ih}	A binary variable used to linearize the disjunctive ESTC constraint. z_{ih} = 1 indicates that piling task i is scheduled after the curing-related exclusion period of pile h; otherwise, piling task i must be completed before that exclusion period starts.

constraint (6) enforces the fundamental construction logic: the grouting of any pile cannot commence until its corresponding piling task has been completed.

Crew routing and sequencing constraints for (7) to (11) manage the operational logic of the crews. Constraints (7) and (8) ensure that each crew begins its work tour from the virtual depot and returns to it upon completion of all its assigned tasks. Constraints (9) and (10) are flow conservation equations that guarantee the continuity of each crew's route. If a crew visits a task, it must have come from a preceding location and must depart to a subsequent one. Constraint (11) is the crucial non-overlapping condition for tasks performed by the same crew. It stipulates that if task j is performed immediately after task i by crew c, then the start time of j must be delayed enough to account for the finish time of i and the travel time between them. The big-M formulation ensures this constraint is active only for consecutively scheduled tasks.

Finally, the dynamic ESTC (12) and (13) as shown in Figure 1 model are the novel and most challenging aspects of the PFCSP. For any piling task i that is physically close to a pile undergoing its construction-and-curing cycle, a disjunctive condition applies. The piling work at location i must either be fully completed before the cycle at location j begins, or it must wait to start until after the cycle at j is entirely finished. The auxiliary binary variable z_{ij} and the big-M are used to linearize this "either-or" logic, rigorously capturing the dynamic, decision-dependent nature of the vibration exclusion zones.

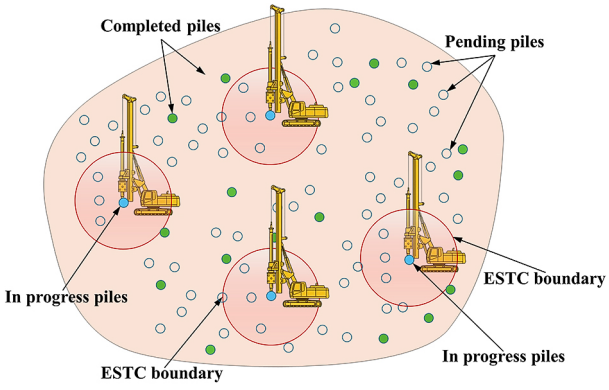


Figure 1. Schematic illustration of ESTC in pile foundation construction

2.3. Computational complexity analysis

The PFCSP, as formulated in this chapter, belongs to the class of NP-hard combinatorial optimization problems. Its computational intractability can be formally established by demonstrating that it is a generalization of the well-known NP-hard RCPSP.

We can establish this relationship by restriction. Consider a simplified special case of the PFCSP under the following conditions: (1) all crews are of a single, homogeneous type, effectively removing the distinction between piling and grouting crew; (2) all spatial constraints are nullified by setting the safety distance D_{safe} = 0, which deactivates the spatio-temporal constraints (12) and (13); (3) all travel times are disregarded by setting T_{travel} (i, j) = 0. The problem resulting from these restrictions is equivalent to the classical RCPSP, which involves scheduling a set of tasks subject to precedence and resource availability constraints to minimize makespan. Since the RCPSP is known to be NP-hard, and it exists as a special case of the PFCSP, the PFCSP is, by definition, also NP-hard.

Classifying the PFCSP as merely NP-hard significantly understates its complexity, which stems primarily from ESTC where the existence, location, and duration of vibration exclusion zones are direct consequences of scheduling decisions. Unlike standard exogenous constraints, this decision-dependent uncertainty creates a highly non-monotonic search landscape where locally optimal choices often lead to global suboptimality. Furthermore, this tight coupling renders problem representation inseparable from solution evaluation, making traditional permutation-based metaheuristics ill-suited since the feasibility and cost of any decision are intrinsically bound to the precise temporal context in which it is made.

3. TA-ACO algorithm

3.1. Overall framework of the TA-ACO

To effectively address the PFCSP, which is characterized by heterogeneous resources and endogenously generated vibration exclusion zones, we depart from traditional metaheuristic frameworks. Conventional algorithms typically employ an indirect, permutation-based decoding approach, where a solution is represented as a static sequence of tasks and evaluated by a separate decoder (Pan & Zhang, 2021). This separation proves critically inefficient for the PFCSP: the feasibility and cost of a task are profoundly dependent on the precise temporal execution of preceding tasks. In such a dynamic environment, a decision made early in a permutation can unknowingly render a vast portion of the subsequent solution space infeasible, leading to extensive backtracking or invalid solutions.

To overcome this blindness, we propose a state-dependent constructive simulation framework as shown in Figure 2. This architecture integrates the solution construction and evaluation processes, modeling the scheduling problem as a discrete-event simulation where decisions are made chronologically as resources become available.

Central to this framework is the shared global scheduling state (SGSS). The SGSS functions not merely as a memory structure but as a dynamic environment that evolves in real-time, capturing the instantaneous spatio-temporal topology of the site (Xu et al., 2014). The SGSS encapsulates four critical dimensions of the project snapshot:

1. Temporal State: The current Gantt chart, tracking start and finish times for all scheduled tasks.
2. Resource State: The real-time availability and status of every construction crew.

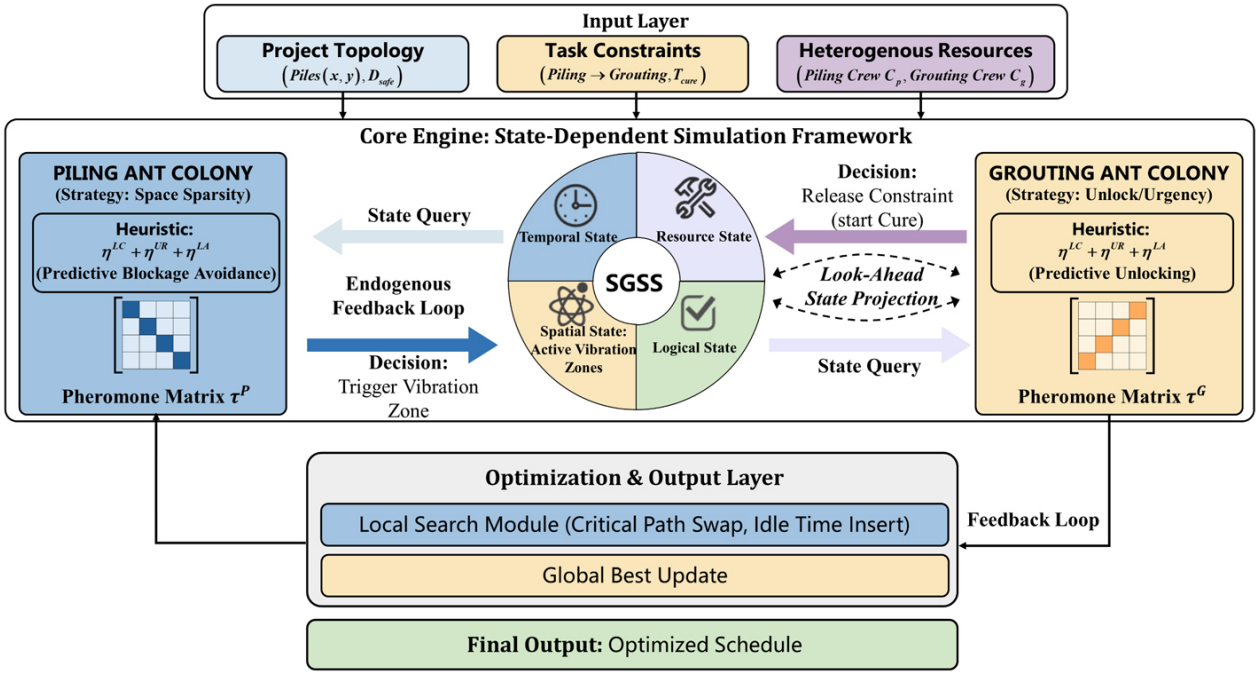


Figure 2. Topologically adaptive ACO framework

3. Spatial State: A continuously updated registry of all active vibration exclusion zones, defined by their center coordinates, radii, and expiration times.
4. Logical State: The completion status of all technological precedence constraints.

In this state-dependent environment, TA-ACO uses two ant colonies to reflect the two types of construction crews. The piling colony selects piling tasks while considering active and potential vibration exclusion zones. The grouting colony selects grouting tasks to release precedence-related constraints and reduce waiting time.

The solution construction unfolds as follows: At any discrete decision point, all newly idle crews simultaneously assess the current SGSS. Each ant evaluates its set of feasible next tasks based on a multi-factor heuristic function and its colony's specific pheromone matrix. After a task is selected, the SGSS is updated immediately. The selected task is fixed in the schedule, the corresponding crew status is updated, and a new vibration exclusion zone may be activated. Therefore, subsequent decisions are made based on the latest temporal, spatial, and logical state of the project. This reduces the chance of generating infeasible schedules.

3.2. Solution construction and multi-factor heuristic function

3.2.1. Solution construction

Within the time-driven direct scheduling framework, a complete schedule is constructed incrementally through a sequence of local decisions made by individual ants. This section details the mechanics of this constructive process, focusing on the sophisticated heuristic function that guides the ants' choices.

As established in Section 3.1, the scheduling process unfolds as a discrete-event simulation. A decision-making event is triggered whenever a crew completes its assigned task and becomes idle. At a given decision time t, every idle crew $c \in C_{idle}$ consults the SGSS to determine its candidate task set, denoted as $J_c(t)$. A task j is considered a candidate for crew c if it satisfies all of the following conditions: 1) task type compatibility: the task type matches the crew type; 2) precedence feasibility: all precedence constraints for task j have been met; 3) spatio-temporal feasibility: executing task j at its location must be possible. Specifically for a piling task, its location must not be within any active vibration exclusion zone at the projected start time.

If the candidate set $J_c(t)$ is not empty, crew c probabilistically selects its next task j from this set. The selection is governed by the standard ACO probabilistic transition rule, which balances the learned knowledge stored in the pheromone matrix (τ) with the problem-specific heuristic information (η):

$$P_c(j,t)=\frac{\left[\tau^{r(c)}_{i_c(t),j}\right]^{\alpha}\left[\eta_{i_c(t),j}(t)\right]^{\beta}}{\sum_{k\in J_c(t)}\left[\tau^{r(c)}_{i_c(t),k}\right]^{\alpha}\left[\eta_{i_c(t),k}(t)\right]^{\beta}},\quad j\in J_c(t), \tag{14}$$

where: $i_c(t)$ denote the last task completed by crew c before decision time t, and let $r(c)\in\{p,g\}$ denotes the beam subset of feasible candidate tasks. $P_c(j, t)$ is the probability of crew c choosing task j at time t. τ_{cj} is the pheromone value associated with assigning task j to a crew of the same type as c. We maintain two independent matrices: τ^P for the piling colony and τ^G for the grouting colony. η_j is the composite heuristic value for the task j, which is the focus of the subsequent section. α and β are parameters that control the relative importance of the pheromone trail versus the heuristic information.

3.2.2. Multi-factor heuristic function

The heuristic function η_j is the core intelligent component guiding an ant's task selection. To navigate the complex trade-offs in the PFCSP, we depart from a monolithic, short-sighted heuristic. Instead, we design a composite heuristic function that is a product of three distinct factors, each targeting a different strategic aspect of the problem:

$$\eta_j=\left(\eta_j^{LC}\right)^{\delta_1}\cdot\left(\eta_j^{UR}\right)^{\delta_2}\cdot\left(\eta_j^{LA}\right)^{\delta_3}, \tag{15}$$

here, η_j^{LC}, η_j^{UR} and η_j^{LA} represent the local cost, task urgency, and look-ahead factors, respectively. The exponents δ_1, δ_2, δ_3 are weights that calibrate the relative influence of each factor.

Local cost factor provides a fundamental, greedy measurement of local efficiency. It favors tasks that can be started and completed quickly, minimizing both travel time and processing time. For an idle crew c currently at the location of its last task i, the local cost factor for a candidate task j is defined as the inverse of the sum of travel and processing time:

$$\eta_j^{LC}=\frac{1}{T_{travel}(i,j)+T_{proc}(j)}, \tag{16}$$

where $T_{travel}(i,j)$ is the travel time for the crew from location i to j, and $T_{proc}(j)$ is the processing time of task j. This factor ensures that the algorithm maintains a baseline of efficient resource utilization.

Task urgency factor is designed to mitigate workflow bottlenecks, a common issue in multi-stage projects. In the PFCSP, a significant risk is the accumulation of the completed piling tasks awaiting grouting, which can lead to prolonged project duration. The urgency factor addresses this by prioritizing grouting tasks that have been waiting the longest. For a grouting p_j, its urgency is proportional to time elapsed since its preceding piling task d_j was completed. For a piling task, this factor is neutral.

$$\eta_j^{UR}=\begin{cases}1+\dfrac{t_{current}-End(d_j)}{C^{avg}_{\max}}, & \text{if } j \text{ is a grouting task } p_j,\\ 1, & \text{if } j \text{ is a piling task}\end{cases} \tag{17}$$

where $t_{current}$ is the current decision time, $End(d_j)$ is the completion time of the predecessor piling task, and $C^{avg}_{\max}$

is the running average of the makespan from previously constructed solutions, serving as a dynamic normalization constant. By prioritizing older waiting tasks, this factor actively promotes a smoother flow through the two-stage construction process.

The third and the most critical component of the heuristic function is the look-ahead factor, designed to overcome the myopic nature of standard constructive heuristics. Unlike the local cost or urgency factors which evaluate the desirability of a task based solely on the current state S_t, this factor evaluates a candidate task by recursively simulating the optimal trajectory of the crew for a depth of $K = 3$ steps to assess its long-term impact on the spatio-temporal topology.

Specifically, at each recursive step, we employ a beam search strategy with a width of B. Instead of expanding all feasible tasks, the algorithm only selects the top-B most promising tasks, ranked by a lightweight base heuristic, to propagate into the next simulation level. This truncation effectively prunes the search tree, reducing the complexity from $O\left(\left|\Omega\right|^K\right)$ to $O\left(B\cdot K\right)$, ensuring the algorithm remains efficient even in high-density instances.

To facilitate this calculation, we explicitly define the global scheduling state at any decision step t as a tuple $S_t = \langle \mathbf{M}_{done}, \mathbf{L}_{crews}, \mathbf{T}_{avail}, \mathbf{Z}_{active} \rangle$, representing the task completion status, crew locations, availability times, and the set of active vibration exclusion zones. A deterministic transition function $\mathcal{T}(S, j, c)$ is employed to return the new state resulting from crew c executing task j, specifically capturing the endogenous activation of new exclusion zones based on the task type.

For a piling crew c, the look-ahead factor $\left(\eta_{d_j}^{LA}\right)$ focuses on avoiding future spatial deadlocks. The key idea is to temporarily simulate the execution of a candidate piling task and then examine whether sufficient feasible workfronts remain in the following virtual steps. A candidate that appears efficient in the current state may receive a low look-ahead score if it is likely to isolate the remaining piles or force the crew into idle waiting. Accordingly, the factor is defined as the inverse of the minimum cumulative cost over the 3-step horizon:

$$\eta_{d_j}^{LA} = \frac{1}{Cost(S_0, j, c) + \epsilon}. \tag{18}$$

The cumulative cost function is calculated via nested recursion:

$$\widehat{Cost} = \underbrace{C_{trans}(S_0, j)}_{\text{Step 1 (Real)}} + \min_{k\in\Omega_B(S_1,c)} \left[\underbrace{C_{trans}(S_1, k) + \Psi(S_1)}_{\text{Step 2 (Virtual)}} + \min_{l\in\Omega_B(S_2,c)} \left(\underbrace{C_{trans}(S_2, l) + \Psi(S_2)}_{\text{Step 3 (Virtual)}} \right) \right], \tag{19}$$

here, S_1 and S_2 represent the sequential virtual states generated by the transition function, and $\Omega_B(S, c)$ denotes the set of beam subset. C_{trans} represents the transition cost (travel plus processing time). Crucially, a Deadlock Penalty Function $\psi(S)$ is integrated; if the set of feasible tasks $\Omega_B(S, c)$ becomes empty in any virtual state due to spatial locking, $\psi(S)$ assigns a prohibitive cost M, forcing the ant to reject paths leading to unavoidable idleness.

Conversely, for a grouting crew g, the factor $\left(\eta_{p_j}^{LA}\right)$ aims to maximize the schedule's unlocking potential. The heuristic value is derived from the cumulative number of piling tasks released from potential blockage constraints over the next three steps:

$$\eta_{p_j}^{LA} = 1 + Unlock(S_0, j, g). \tag{20}$$

The recursive gain function is defined as:

$$\widehat{Unlock} = N_{blocked}\left(j \mid S_0\right) + \gamma \cdot \max_{k\in\Omega_B(S_1,g)} \left[N_{blocked}\left(k \mid S_1\right) + \gamma \cdot \max_{l\in\Omega_B(S_2,g)} \left(N_{blocked}(l \mid S_2) \right) \right], \tag{21}$$

where $N_{blocked}(task \mid S)$ calculates the number of feasible piling tasks that would be covered by the exclusion zone generated by the candidate task, and $\gamma \in (0, 1)$ is a discount factor that weighs immediate unlocking higher than distant future gains. By integrating this predictive factor, the algorithm effectively shifts from reactive selection to proactive planning, directly addressing the challenge of endogenous constraints.

To clarify the role of the recursive look-ahead factor, consider a decision point at which a piling crew can choose from two candidate piles, denoted as A and B. Candidate pile A may have the shortest travel distance and would therefore be preferred by a purely local-cost rule. However, if the virtual execution of A generates a vibration exclusion zone that covers several nearby unscheduled piles, the candidate set in the next simulated state may become very small or even empty. In this case, the deadlock penalty in the look-ahead evaluation assigns a large cost to this branch. Candidate pile B, although slightly farther away, may leave the surrounding workfront more spatially open and preserve feasible piles for the next two virtual steps. The look-ahead factor therefore gives B a higher heuristic value because it maintains future constructability rather than only minimizing the immediate travel cost. The beam search strategy makes this evaluation computationally tractable: instead of expanding all possible future sequences, only the top-B candidate branches ranked by the base heuristic are propagated at each virtual level. For grouting crews, the same principle is used in the opposite direction: a grouting task receives a higher look-ahead value if its completion is expected to release more potentially blocked piling tasks in the subsequent states.

3.2.3. Adaptive dynamic normalization mechanism

The three heuristic factors proposed, η^{LC}, η^{UR}, and η^{LA} possess distinct physical dimensions and varying orders of magnitude. For instance, the value of η^{LA} for grouting crews is typically in the range of [1, 9], whereas η^{LC} is often in the order of 10^{-2}.

Direct multiplication of these heterogeneous factors in Eqn (15) would result in the factor with the largest magnitude dominating the decision probability, rendering the weight parameters δ ineffective. To mitigate this scale bias and ensure that δ strictly reflects the relative sensitivity of the algorithm to each factor, we implement an adaptive normalization strategy within the local candidate set.

At each decision step t, for an idle crew c with a feasible candidate task set $\mathcal{N}_c(t)$, the raw heuristic values are normalized relative to the maximum value present in the current candidate set. The normalized heuristic factor $\bar{\eta}_{ij}^k$ is calculated as follows:

$$\bar{\eta}_{ij}^k = \frac{\eta_{ij}^k}{\max_{m\in\mathcal{N}_c(t)}(\eta_{im}^k)} + \epsilon, \tag{22}$$

where $k \in \{LC, UR, LA\}$ represents the type of heuristic factor, and ϵ is a small constant to prevent zero probabilities. Consequently, the composite heuristic function in Eqn (15) is reformulated as:

$$\eta_{ij}(t) = [\bar{\eta}_{ij}^{LC}]^{\delta_1} \cdot [\bar{\eta}_{ij}^{UR}]^{\delta_2} \cdot [\bar{\eta}_{ij}^{LA}]^{\delta_3}. \tag{23}$$

3.3. Pheromone update mechanism

The pheromone update mechanism is the core learning component of the TA-ACO algorithm, enabling it to iteratively refine its search strategy based on past performance. The pheromone update mechanism allows the algorithm to reinforce task transitions that appear in high-quality schedules. A distinctive feature of our TA-ACO is the maintenance of two separate and independent pheromone matrices: τ^P for the piling colony and τ^G for the grouting colony. This partitioned learning structure mirrors the problem's heterogeneous resource division and allows for specialized knowledge acquisition for each task type. The update process, executed at the end of each iteration, comprises two complementary phases: pheromone evaporation and pheromone deposition.

To prevent premature convergence to a locally optimal solution and to encourage exploration of the search space, the pheromone values on all trails are first uniformly reduced. This process simulates the natural decay of old information, diminishing the influence of previously found solutions and allowing the colonies to adapt to new, potentially more promising search directions. The evaporation is applied to both matrices as follows:

$$\tau_{ij} \leftarrow (1-\rho)\cdot\tau_{ij}, \ \forall(i,j), \tag{24}$$

where τ_{ij} represents the pheromone value on the trail connecting task i to task j, and $\rho \in (0,1]$ is the pheromone evaporation rate.

Following evaporation, pheromone is deposited to reinforce the pathways that constitute high-quality solutions. We use a global-best elitist strategy. In this strategy, only the best schedule found so far is used for pheromone deposition. In this strategy, only the single best solution found so far across all iterations, denoted as the global-best schedule S_{gbest} with a makespan of $C_{\max}^{gbest}$, is used to deposit pheromone.

The amount of pheromone to be added, $\Delta\tau^{gbest}$, is inversely proportional to the quality of the global-best solution, ensuring that better solutions provide stronger reinforcement:

$$\Delta\tau^{gbest} = \frac{Q}{C_{\max}^{gbest}}. \tag{25}$$

Critically, the deposition process respects the dual-colony structure. The complete schedule S_{gbest} is decomposed, and pheromone is added only to the corresponding matrices:

For the piling pheromone matrix: we iterate through the task sequence of each individual piling crew within S_{gbest}. For every pair of consecutively executed piling tasks (d_i, d_j) by the same crew, the pheromone on the corresponding trail in τ^P is reinforced.

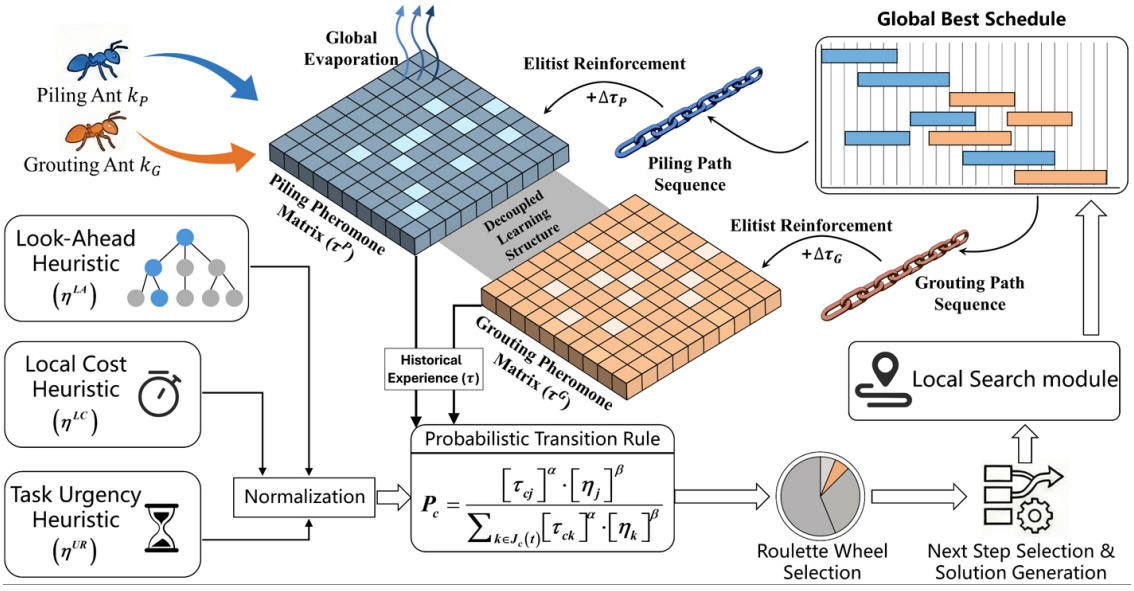


Figure 3. General framework of the proposed TA-ACO algorithm

For the grouting pheromone matrix: similarly, for every pair of consecutive grouting task $\left(p_i, p_j\right)$ executed by the same grouting crew, the trail in τ^G is reinforced.

The update rule for deposition is therefore:

$$\tau_{ij} \leftarrow \tau_{ij} + \Delta\tau^{gbest}, \ \forall \text{consecutive pairs } (i,j) \in S_{gbest}. \quad (26)$$

As comprehensively illustrated in Figure 3, the general framework of the proposed TA-ACO algorithm seamlessly integrates the dual-colony state-dependent construction process, multi-factor heuristics, partitioned pheromone updates, and problem-specific local search into a unified iterative optimization loop.

3.4. Problem-specific local search enhancement

Metaheuristic algorithms like ACO are powerful at exploring a vast search space to identify promising regions, but they may lack the fine-tuning capability to pinpoint the local optimum within those regions. To address this, we embed a local search module within the TA-ACO framework, which is applied to the best solution generated in each iteration. This module does not employ generic neighborhood operators; instead, it consists of two highly specialized operators, each designed to mitigate a specific source of the inefficiency inherent in the PFCSP.

The first operator focuses on the critical path. In makespan minimization, reducing the project duration usually requires shortening the sequence of activities that determines the final completion time. We first identify this path by a backward pass from the last-finishing task and by tracing zero-slack predecessors.

After the critical path is identified, two neighborhood moves are tested. The first move swaps two adjacent critical tasks performed by the same crew. The second move relocates a critical task to another position in the same crew sequence, especially when non-critical tasks provide available slack. Each move is followed by a schedule re-evaluation because small sequence changes may affect subsequent task timings. A first-improvement rule is used: once a move reduces the makespan, it is accepted and the search restarts from the improved schedule.

Another operator is tailored to address the most unique and challenging constraint of the PFCSP: the forced idle time imposed upon piling crews by the dynamic vibration exclusion zones. Such idle periods are a primary source of inefficiency. The operator's procedure begins by systematically scanning the schedule of each piling crew to detect idle windows, periods of forced inactivity where a crew must wait for a vibration zone to dissipate at its next task's location. Upon identifying such a window, the operator searches for a suitable opportunistic task from later in the same crew's schedule to insert into this gap. A task is deemed a viable candidate if it satisfies both temporal and spatial feasibility: its entire operation, including travel, must fit within the idle window, and its location must be clear of any vibration zones during the proposed execution time. Should multiple candidates exist, the operator employs a best improvement criterion, selecting the relocation that yields the maximum reduction in the overall project makespan. By executing this logic, the operator effectively transforms periods of constraint-induced inefficiency into productive work, significantly improving the quality of the final schedule.

3.5. Exact algorithm for performance benchmarking

To systematically evaluate the solution quality of the proposed TA-ACO algorithm and to establish an absolute performance benchmark, we employed an exact optimization method for solving problem instances. This approach provides a theoretical performance ceiling against which the heuristic solutions can be rigorously compared. This comprehensive MILP model detailed in Section 2 was then implemented and solved using the state-of-the-art commercial optimization solver, Gurobi Optimizer. Gurobi is widely recognized as one of the most powerful tools for tackling large-scale MILP problems, integrating a sophisticated suite of algorithms that includes advanced presolving techniques, cutting-plane methods, and an efficient branch-and-bound engine. For the computational experiments, we imposed a uniform computational time of 3600 CPU seconds for Gurobi on each test instance. The objective function value of a proven optimal solution found within this timeframe, or the best feasible solution found if the time limit was exhausted, was recorded. This value serves as the definitive optimal or near-optimal benchmark for that specific instance, providing a solid foundation for quantifying the optimality gap of the TA-ACO algorithm.

4. Computational experiments

4.1. Case study

The empirical data for this study is derived from a mega-integrated transportation hub project in Southwest China. As shown in Figure 4, integrating an airport terminal, a high-speed railway station, and a global traffic center (GTC), the project is characterized by the profound coupling of its underground spaces. Specifically, high-speed rail and subway lines traverse longitudinally beneath the terminal's foundation, creating an intricate three-dimensional intersection. This high-density spatial integration channels all superstructure development into a singular spatio-temporal bottleneck: the massive pile foundation construction phase.

Due to complex geological conditions, the project employs impact bored piling. The intense vibrations generated by this process introduce rigorous ESTC. Construction must adhere to strict skip-piling protocols: the initial curing period of a newly cast pile endogenously triggers a dynamic vibration exclusion zone. These zones emerge as a direct consequence of scheduling decisions and dissipate over time, resulting in a highly volatile feasible region for the scheduling problem.

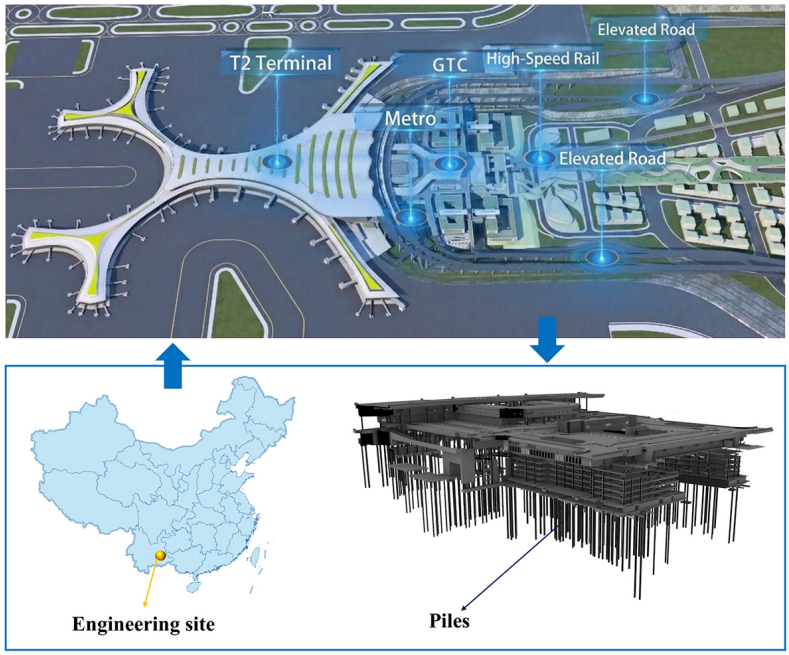


Figure 4. Overview of the mega-integrated transportation hub case study: geographical location, project layout, and 3D structural model of the pile foundation

To validate algorithm performance across different dimensions, we constructed a high-fidelity instance set with three distinct gradients based on real engineering data. The spatial distribution of piles and constraint characteristics are directly mapped from the actual construction scenarios:

- Small-scale instances (S-01–S-04, 12–27 piles): selected from spatially isolated single concourse areas. These smaller-scale instances are designed to verify the logical correctness and optimization precision of the TA-ACO algorithm by comparing its results against exact solutions obtained from the commercial solver, Gurobi.
- Medium-scale instances (M-01–M-04, 37–62 piles): derived from typical intersection zones between the terminal and the GTC, characterized by extremely high pile density and complex structural interference. These instances aim to test the computational efficiency and scalability of the algorithm under strong coupling constraints compared to the exact solver.
- Large-scale instances (L-01–L-02, 143–197piles): encompassing the complete pile foundation works of the main terminal's core area. Representing the full magnitude of the real-world engineering challenge, this set evaluates the practical utility of the algorithm in coordinating massive task volumes with multiple construction crews, focusing on the magnitude of makespan reduction relative to traditional heuristic strategies.

4.2. Parameter tuning and sensitivity analysis

The performance of the proposed TA-ACO algorithm is significantly influenced by a series of key parameters, which collectively govern the algorithm's search behavior, particularly the balance between exploration and exploitation. To ensure the algorithm exhibits robust and efficient performance across problem instances with diverse characteristics, and to establish a fair and rational parameter configuration for the subsequent computational experiments, we selected representative medium-scale instances to conduct a systematic parameter tuning and sensitivity analysis. The parameters identified for analysis are the core components critical to the TA-ACO algorithm: the pheromone importance factor α, the heuristic information importance factor β, the pheromone evaporation rate ρ, and the weights for the heuristic factors δ_1, δ_2, δ_3.

The relative magnitude of α and β directly determines the extent to which ants rely on historical experience versus current heuristic information during path selection. Figure 5 illustrates the algorithm's convergence behavior under different parameter combinations. The experimental results indicate that when the value of α is excessively large, the algorithm relies too heavily on paths generated in the early stages. This leads to a rapid depletion of

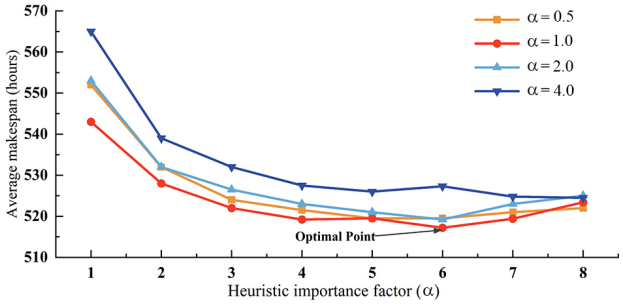


Figure 5. Sensitivity analysis: impact of α and β on solution quality

population diversity, causing the algorithm to prematurely converge to local optima. The results demonstrate that the combination of α = 1.0, β = 6.0 achieves the optimal balance between solution quality and convergence speed. This observation underscores a critical insight: in problems characterized by strong dynamic constraints, such as the PFCSP, heuristic information is more critical than historical information.

Building upon the adaptive dynamic normalization mechanism introduced in Section 3.2.3, the sensitivity analysis presented here reflects the algorithm's intrinsic behavioral response to different heuristic components, rather than mere magnitude corrections. Since the heuristic values are normalized to a unified scale of (0, 1], the hyperparameters δ_1, δ_2, δ_3 strictly represent the relative importance of local efficiency, task urgency, and look-ahead topology, respectively. This structural decoupling ensures the robustness of the following parameter tuning process.

Preliminary experiments indicated that the system's performance is relatively robust to variations in the Urgency Factor weight (δ_2) within the range [0.5, 2.0]. While δ_2 shows lower sensitivity regarding the makespan in the parameter sweep, it plays a necessary role in preventing the starvation of grouting tasks, a function that will be further verified in the subsequent ablation study. Therefore, we fixed δ_2 = 1.0 to conduct a focused bivariate sensitivity analysis on the interaction between the local cost weight (δ_1) and the look-ahead weight (δ_3).

Figure 6 illustrates the interplay between the local cost weight (δ_1) and the look-ahead weight (δ_3) on the project makespan. A distinct saturation behavior is observed along the vertical axis: as δ_3 increases from 0.0 to 5.0, the makespan declines precipitously, confirming that the cooperative look-ahead mechanism is the dominant driver for resolving endogenous spatial constraints; however, this performance gain diminishes and stabilizes beyond δ_3 = 4.0, indicating that the algorithm effectively reaches its topological optimization capacity. Concurrently, the system exhibits a convex, basin-shaped response to the Local Cost weight. While a moderate δ_1 (range [0.0, 5.0]) is essential for maintaining basic travel efficiency, excessive emphasis on this factor (δ_1 > 3.5) induces myopic behavior, where ants prioritize immediate proximity over long-term spatial connectivity, ironically leading to severe spatial fragmentation and inferior global outcomes. The optimal performance, represented by the deep blue valley in the heatmap, is achieved where the strong cooperative guidance ($\delta_3 \approx 4.0$) is balanced by a moderate sensitivity to local costs ($\delta_1 \approx 2.0$). Based on this empirical equilibrium, and fixing the robust Urgency Factor at δ_2 = 1.0, we determine the final parameter configuration for the TA-ACO as δ_1 = 2.0, δ_2 = 1.0, and δ_3 = 4.0. The final parameter configuration is shown in Table 4.

Table 4. TA-ACO algorithm parameters

Parameter symbol	Description	Value
N_{ants}	Population Size	20
α	Pheromone Factor	1.0
β	Heuristic Factor	6.0
B	Beam Width	15
ρ	Evaporation Rate	0.1
Q	Pheromone Constant	100
δ_1	Weight for Local Cost	2.0
δ_2	Weight for Urgency	1.0
δ_3	Weight for Look-ahead	4.0
K_{iter}	Max Iterations	500

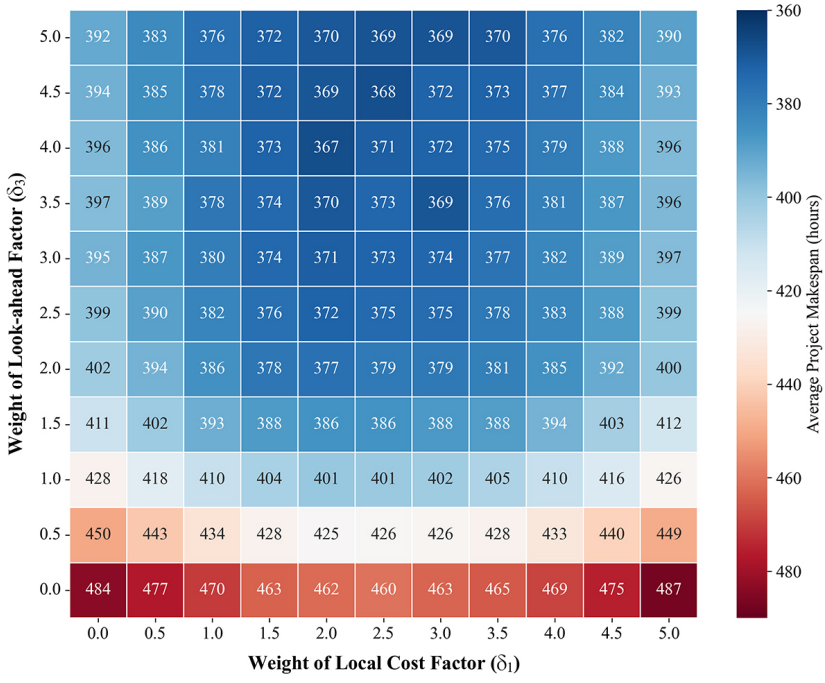


Figure 6. Sensitivity analysis of heuristic components weight

4.3. Comparison of results

4.3.1. Validation against exact solver

Prior to presenting the large-scale experimental results, we analyze a simplified 12-pile instance to elucidate the fundamental complexity introduced by ESTC. In this pilot scenario, parameterized with a uniform processing time of 18 hours and a mandatory curing period of 36 hours, the impact of scheduling logic on project duration is isolated and visibly amplified. Figures 7 and 8 illustrate the construction trajectories of a Greedy Strategy versus the Optimal Strategy, respectively. As depicted in Figure 7, the Greedy algorithm, driven solely by distance minimization, adopts a path that appears locally efficient in the initial stages. However, this myopic sequencing leads to a severe spatial blockage: the consecutive execution of piles No. 9 and No. 10 triggers overlapping vibration exclusion zones that simultaneously cover the locations of the remaining pending piles. This creates a spatial deadlock where the crew is forced into idleness, unable to operate until the curing periods expire. In sharp contrast, the Optimal Strategy shown in Figure 8 demonstrates a topologically aware sequence. By strategically staggering the spatial execution order, it ensures that active exclusion zones never completely blockade the remaining workfront. The quantitative disparity is significant: the improper sequencing by the greedy approach extends the makespan to 247.3 hours, whereas the optimal strategy completes the project in 224.5 hours. This illustrative example underscores a critical premise: in the presence of endogenous constraints, myopic decisions are not merely inefficient – they are fatal to the continuous flow of construction, inducing disproportionate delays that cannot be recovered in later stages.

To evaluate the global optimization capability of the TA-ACO algorithm, we first benchmarked it against the commercial solver Gurobi (v10) using the small-scale instance set. As presented in Table 5, TA-ACO successfully identified makespans identical to those obtained by Gurobi across all small-scale test cases. This result robustly validates the logical correctness of the TA-ACO algorithm, demonstrating that its dual-population cooperative mechanism is capable of effectively pinpointing the global optimal solution within the solution space.

However, as the problem size expands, the computational complexity of the PFCSP exhibits the exponential escalation characteristic of NP-Hard problems. For medium-scale instances, Gurobi failed to prove optimality within the imposed 3600-second time limit, yielding only the best feasible solution found so far. In contrast, TA-ACO demonstrates exceptional scalability, converging to high-quality feasible solutions within approximately 500 seconds. On some test sets, it achieved results superior to the solutions reported by Gurobi.

To validate the algorithm's performance in a complex, real-world engineering context, we utilized the large-scale instance L-02, which was constructed based on actual data from the core area of the T2 Terminal. We benchmarked the proposed TA-ACO against the industry-standard Greedy strategy and three mainstream metaheuristic algorithms: Simulated Annealing (SA), Hybrid Genetic Algorithm (HGA), and Adaptive Large Neighborhood Search (ALNS). As illustrated in Figure 9, TA-ACO exhibits significant superiority in both convergence speed and final solution quality, achieving the lowest makespan of 608.4 hours. This represents a distinct advantage over the next best algorithm, ALNS (638.6 hours), and a substantial improvement over HGA and SA.

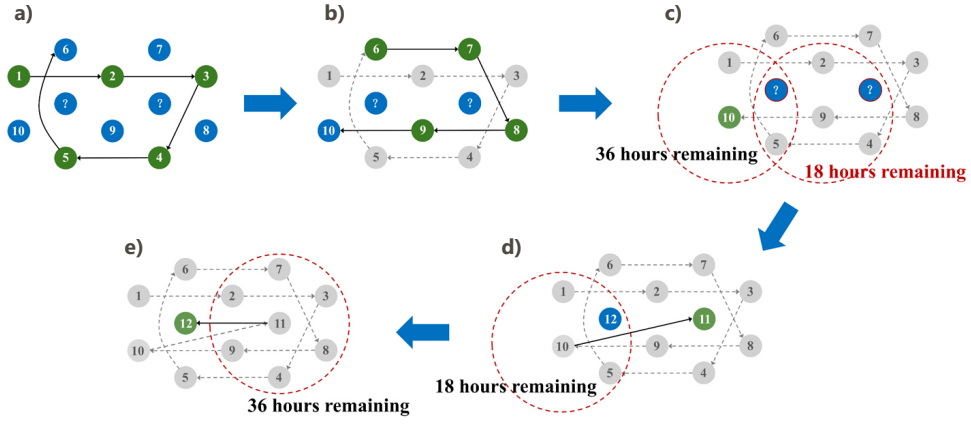


Figure 7. Greedy construction route diagram for 12 piles

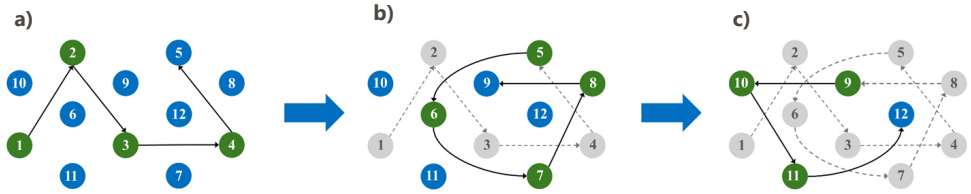


Figure 8. Optimal construction route diagram for 12 piles

Table 5. Comparison of results between the Gurobi solver and TA-ACO

ID	Piles	Crews	Gurobi		TA–ACO		Gap
			Obj(h)	CPU(s)	Obj(h)	CPU(s)	
S–01	12	1	224.5	4.8	224.5	6.4	0.00%
S–02	15	1	280.0	6.7	280.0	9.8	0.00%
S–03	23	2	212.2	45.2	212.2	45.5	0.00%
S–04	27	2	254.3	245.6	255.5	62.3	0.47%
M–01	37	3	266.4	1880.3	268.5	117.9	0.78%
M–02	48	3	310.0*	3600	309.2	249.5	–0.25%
M–03	55	3	398.5*	3600	382.6	328.6	–3.99%
M–04	62	4	445.0*	3600	427.5	491.7	–3.93%
L–01	143	5	–	N/A	720.5	880.9	–
L–02	197	8	–	N/A	608.4	1036.2	–

Note: 1. The computational upper bound for CPU time is set at 3600 seconds. 2. * indicates that Gurobi was unable to prove optimality; in such cases, the comparison is based on the current best solution reported by the solver.

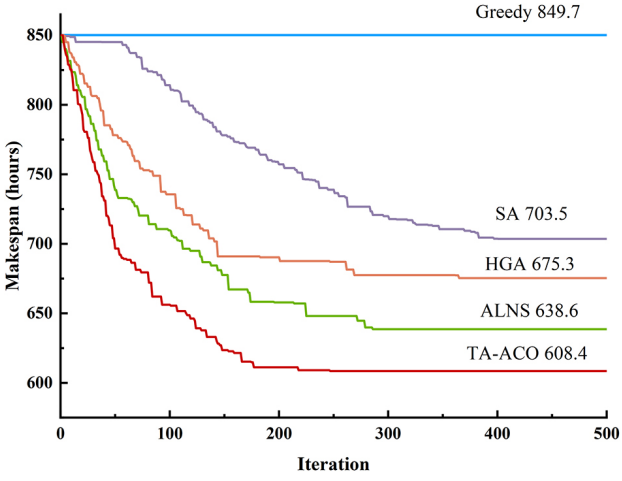


Figure 9. Convergence curve diagram of five algorithms

In contrast, the Greedy strategy, which is widely adopted in current on-site engineering practices, yields the poorest performance with a makespan of 849.7 hours. Its fundamental limitation lies in the myopic logic of always prioritizing the nearest available pile. While computationally simple, this approach frequently traps the machinery in local optima: after clearing immediate tasks, the rig often faces a situation where no reachable piles remain nearby due to exclusion constraints, forcing unproductive long-distance transfers. Furthermore, this lack of global planning tends to cause an accumulation of uncompleted piles in the final stages, resulting in deadlocks where remaining tasks are clustered or mutually exclusive, severely hindering parallel operations.

Concurrently, we observed that the Genetic Algorithm and its variants typically rely on permutation-based encoding. This mechanism, which decouples sequencing from scheduling, renders the algorithm unable to perceive the dynamic vibration exclusion zones, which are endogenously triggered by decisions – at the precise moment of decision-making. This separation inevitably leads to significant forced waiting times during the decoding process. In contrast, the TA-ACO algorithm, incorporating the look-ahead factor, demonstrates superior spatio-temporal planning capabilities. Piling ants shift their focus from merely minimizing travel distance to adopting a sparsity maximization strategy. Subject to safety distance constraints, the algorithm guides crews to prioritize pile locations that impose minimal interference on surrounding potential tasks. This strategy is essentially analogous to dynamically solving the Maximum Independent Set problem in graph theory, ensuring that the site can accommodate the maximum number of concurrent crews within any given time window. By strategically sacrificing local travel efficiency, TA-ACO secures the maximization of global parallelism, thereby significantly shortening the overall project duration.

4.3.2. Comparative evaluation against state-of-the-art metaheuristics

To intuitively visualize the divergent behaviors of different algorithms in managing complex endogenous constraints, we tracked the construction snapshots of the large-scale instance L-01 (comprising 143 piles, 5 piling crews, and 3 grouting crews) across its lifecycle. As illustrated in Figure 10, the distinction between the baseline ALNS (Left Column) and the proposed TA-ACO (Right Column) is initially subtle but progressively stark. In the early phase (Figure 10a), the abundance of available pile positions allows both algorithms to operate without significant interference; the spatial distribution of completed tasks and active crews appears uniform, as the relatively low constraint density masks the strategic differences between the two approaches. However, as the project advances to the intermediate stage (Figure 10b), a fundamental disparity in spatial planning emerges. The ALNS snapshot reveals the formation of several isolated islands (highlighted in red dashed circles), where clusters of uncompleted piles are surrounded by completed ones. This pattern indicates that the ALNS strategy greedily exhausts peripheral or easy-to-access nodes, leaving behind fragmented clusters that are spatially disconnected. In sharp contrast,

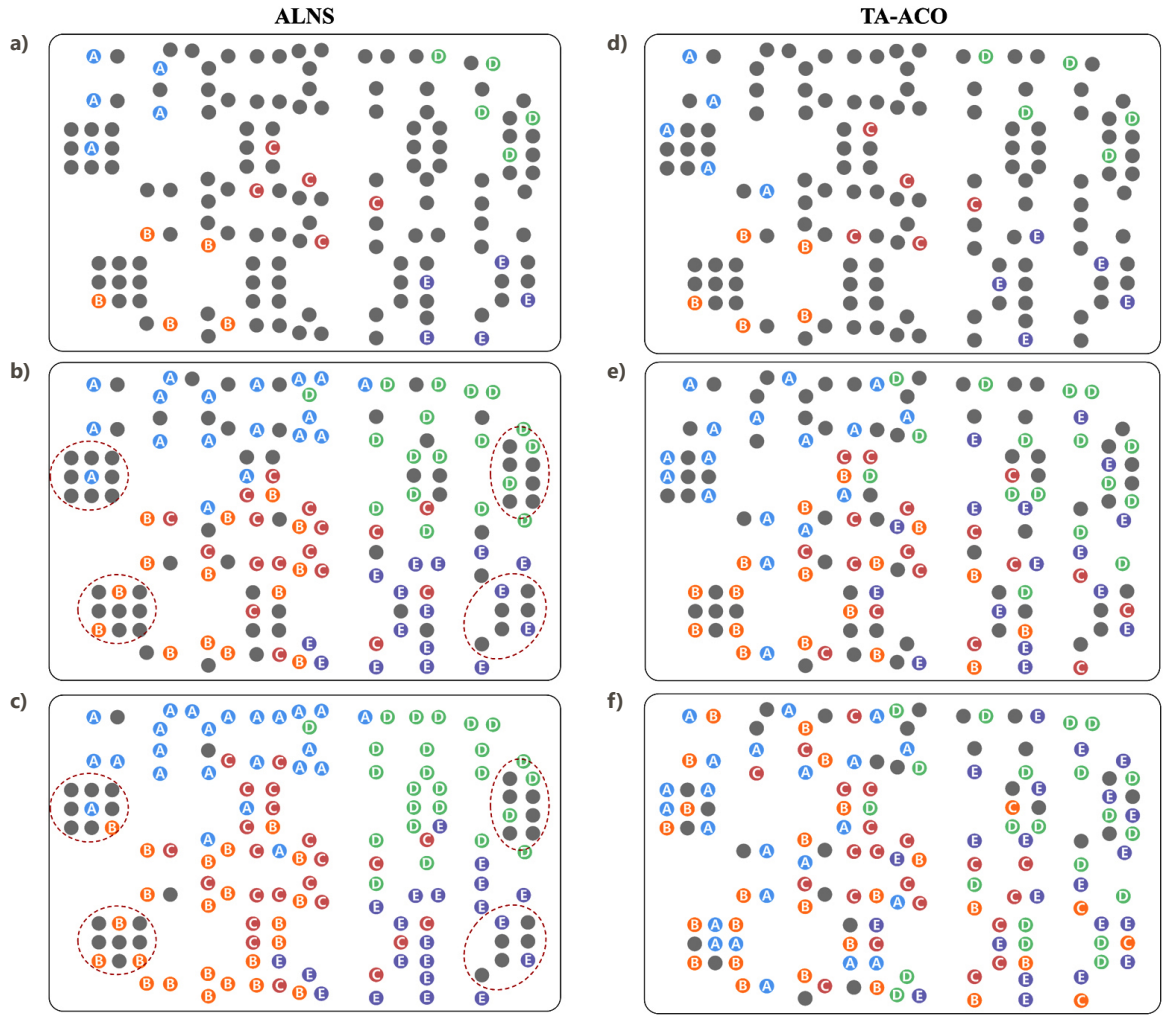


Figure 10. A microscopic comparison between ALNS and TA-ACO: a–c – Snapshots of the construction process using the ALNS algorithm at early, middle, and late stages, respectively; d–f – Snapshots of the construction process using the TA-ACO algorithm at early, middle, and late stages, respectively

the topology maintained by TA-ACO remains significantly more connected; the uncompleted piles are fewer and effectively dispersed, ensuring that crews consistently maintain a high workfront density with feasible adjacent nodes for transfer, thereby avoiding non-productive relocation.

The consequences of these intermediate spatial patterns fully manifest in the final stage (Figure 10c), where the scheduling efficiency of ALNS collapses due to severe spatial deadlocks. In the ALNS solution, the remaining piles are concentrated in dense, island-like clusters. Governed by the strict skip-piling constraint, executing any single pile within these clusters instantaneously generates a vibration exclusion zone that renders all immediate neighbors infeasible for the duration of the curing period. This physical bottleneck forces multiple crews into involuntary idleness, as they must wait for constraints to dissipate before accessing the adjacent tasks. Conversely, the TA-ACO schedule avoids this pitfall entirely; thanks to the look-ahead mechanism applied in earlier stages, the remaining tasks are spatially sparse and satisfy safety distance requirements. This strategic sparsity allows multiple crews to work in parallel on the final set of piles without mutual interference, facilitating a smooth and efficient project close-out.

The performance gap observed in the microscopic evolution can be attributed to the fundamental difference in how the two algorithms perceive and handle ESTC. ALNS relies on destroy and repair operators, which typically manipulate the solution in the permutation space rather than the physical space. When a heuristic operator inserts a task into a sequence, it is often blind to the dynamic vibration zones generated by that insertion, leading to the construction of sequences that unknowingly violate spatial constraints and necessitate costly delays. On the other hand, TA-ACO operates as a constructive method embedded with a forward-looking spatial awareness. By explicitly evaluating the future topological impact of potential moves via the look-ahead factor, the algorithm simulates the generation of vibration zones before finalizing a decision. This allows TA-ACO to treat physical space as a depletable resource, prioritizing the preservation of connectivity and sparsity over immediate local optimality, thus proactively preventing the formation of spatial deadlocks.

4.4. Ablation study

To thoroughly evaluate the effectiveness of each core component within the TA-ACO algorithm and quantify the marginal contributions and synergistic effects of different heuristic strategies in addressing key constraints of the PFCSP, we designed and implemented an ablation study comprising seven distinct configurations. The experiments were conducted on the large-scale instance L-01, characterized by high spatio-temporal coupling, with each variant run independently 20 times. The variants are defined as shown in Table 6.

To quantify the marginal contribution of each module, we conducted an ablation study on the large-scale instance L-01, as summarized in Table 6 and Figure 11. The results unequivocally identify the Dual-colony Structure (Variant 4) and the look-ahead Factor (δ_3, Variant 3) as the two absolute cornerstones of the TA-ACO's performance; their removal resulted in the most precipitous deterioration in solution quality. Specifically, the degradation observed in Variant 4 (Single Colony), where the makespan spiked to 825.9 hours, highlights the necessity of decoupling heterogeneous resources. When piling and grouting crews are forced to share a single pheromone matrix, the distinct topological requirements of space-filling and constraint-unlocking become conflated. This leads to pheromone confusion, where conflicting guidance signals prevent the algorithm from learning specialized strategies for different crew types. Similarly, Variant 3 exhibited a significant performance drop. Without the predictive guidance of δ_3, the crews reverted to myopic behavior. This regression reproduced the specific failure mode observed in the ALNS algorithm: the construction of isolated islands of piles. The absence of the look-ahead mechanism meant crews failed to perceive the risk of self-entrapment, leading to severe spatial fragmentation and unavoidable delays in the later stages of the schedule.

The contribution of the remaining heuristic factors, while less dominant than the core structural components, proves essential for fine-tuning the solution to global optimality. As evidenced by Variant 1 and Variant 2, removing these factors resulted in moderate makespan increases to 738 h and 742.6 h, respectively. This confirms that while topological planning is the primary driver, mechanisms for balancing task aging and minimizing travel distance remain requisite for eliminating local inefficiencies. Additionally, the local search module reduced the project duration by approximately 1.62%, while significantly decreasing the standard deviation of the schedule. Finally, regarding computational efficiency, we report a transparent trade-off: variants incorporating the beam-search-based look-ahead mechanism incur a higher computational overhead, requiring approximately twice the CPU time compared to the simplified Variant 3. However, even with this increased complexity, the full TA-ACO algorithm consistently converges to high-quality solutions within 1200 seconds. Given the scale of the problem, where a 21.46% reduction in makespan translates to weeks of saved construction time, this computational investment is highly cost-effective and falls well within the acceptable time window for static project scheduling.

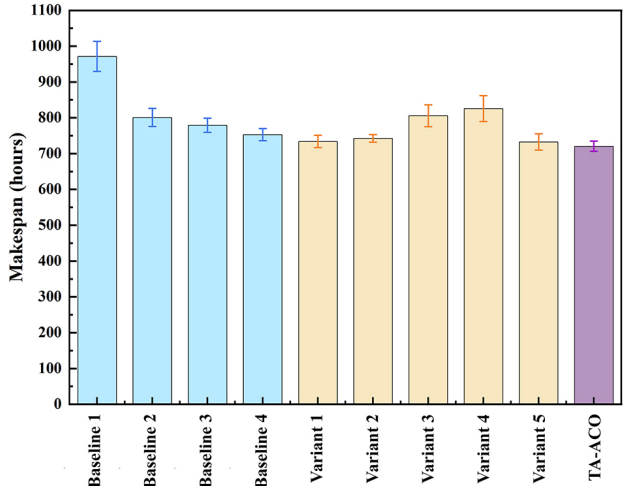


Figure 11. Bar chart of project duration distribution

4.5. Field implementation

To bridge the gap between the theoretical TA-ACO algorithm and on-site engineering management, we developed a web-based decision support system, named Alpha-Pile. This system encapsulates the complex optimization engine

Table 6. Design of the ablation study

Variant name	Description	Average makespan	Standard deviation	Gap
Baseline 1	Greedy algorithm	917.4	42.2	21.46%
Baseline 2	Simulated Annealing	800.9	25.2	10.04%
Baseline 3	Hybrid genetic algorithm	779.6	19.7	7.58%
Baseline 4	Adaptive Large Neighborhood Search	753.5	22.8	4.38%
Variant 1	No urgency factor	738.1	19.2	2.38%
Variant 2	No cost factor	742.6	10.5	2.97%
Variant 3	No look-ahead factor	805.7	30.5	10.57%
Variant 4	Single colony	825.9	36.1	12.76%
Variant 5	No local search enhancement	732.4	26.9	1.62%
TA-ACO	*The proposed algorithm*	*720.5*	*14.4*	–

into a user-friendly interface, enabling project managers to utilize advanced scheduling strategies without requiring expertise in metaheuristics.

As illustrated in Figure 12, the system features a centralized human-machine interface tailored for the specific workflow of the aviation hub project, which is deployed in the on-site management center. Bridging the optimized theoretical schedule with actual field execution, the interface provides project managers with a real-time dynamic visualization of the physical construction process. This live feedback module displays vertical operational profiles, critical depth indicators, and instantaneous mix status, allowing engineers to intuitively verify that the ongoing machinery operations strictly adhere to the spatio-temporal constraints dictated by the TA-ACO algorithm. Furthermore, the integrated control panel serves as the operational command center, enabling management personnel to seamlessly configure project parameters, integrate sensor data, and coordinate construction activities without disrupting the on-site workflow.

The system was deployed in a critical pilot zone of the T2 Terminal foundation works, comprising 97 pile foundations and characterized by high-density structural interference. According to the original baseline schedule generated using the traditional nearest-neighbor strategy, the projected construction duration for this zone was 47 days. By strictly adhering to the optimized schedule, the construction team successfully navigated the ESTC and completed the piling works in 34 days. This achievement represents a significant lead time of 13 days ahead of schedule. This empirical efficiency gain, which closely aligns with the simulation results presented in Section 4.3, provided a crucial time buffer for the subsequent superstructure assembly, thereby validating the practical robustness and value of the TA-ACO algorithm in a real-world engineering environment.

5. Discussion

The experimental results, particularly the microscopic comparison in Figure 10, reveal a counter-intuitive phenomenon inherent to the PFCSP: strategies that maximize local travel efficiency often yield inferior global outcomes. This study identifies this as the paradox of local efficiency under endogenous constraints. Empirical evidence from the large-scale instance L-01 demonstrates this starkly: while the Greedy strategy prioritized immediate proximity, it resulted in a prolonged makespan of 917.4 hours. In comparison, the proposed TA-ACO algorithm achieved a makespan of 720.5 hours on instance L-01, representing a substantial 21.4% reduction.

Traditional heuristics operate on the premise that minimizing non-productive travel time directly contributes to makespan reduction. However, in the presence of ESTC, myopic distance minimization creates a spatial overdraft. As observed in the ALNS simulation (Figure 10b), by rapidly consuming accessible clusters of piles, heuristic crews inadvertently construct spatial islands—isolated enclaves of uncompleted tasks surrounded by active vibration exclusion zones. This structural flaw forces crews into unavoidable idleness in the later stages, resulting in a suboptimal makespan of 753.5 hours compared to TA-ACO.

In contrast, the TA-ACO algorithm, driven by the look-ahead factor, functions similarly to a dynamic solver for

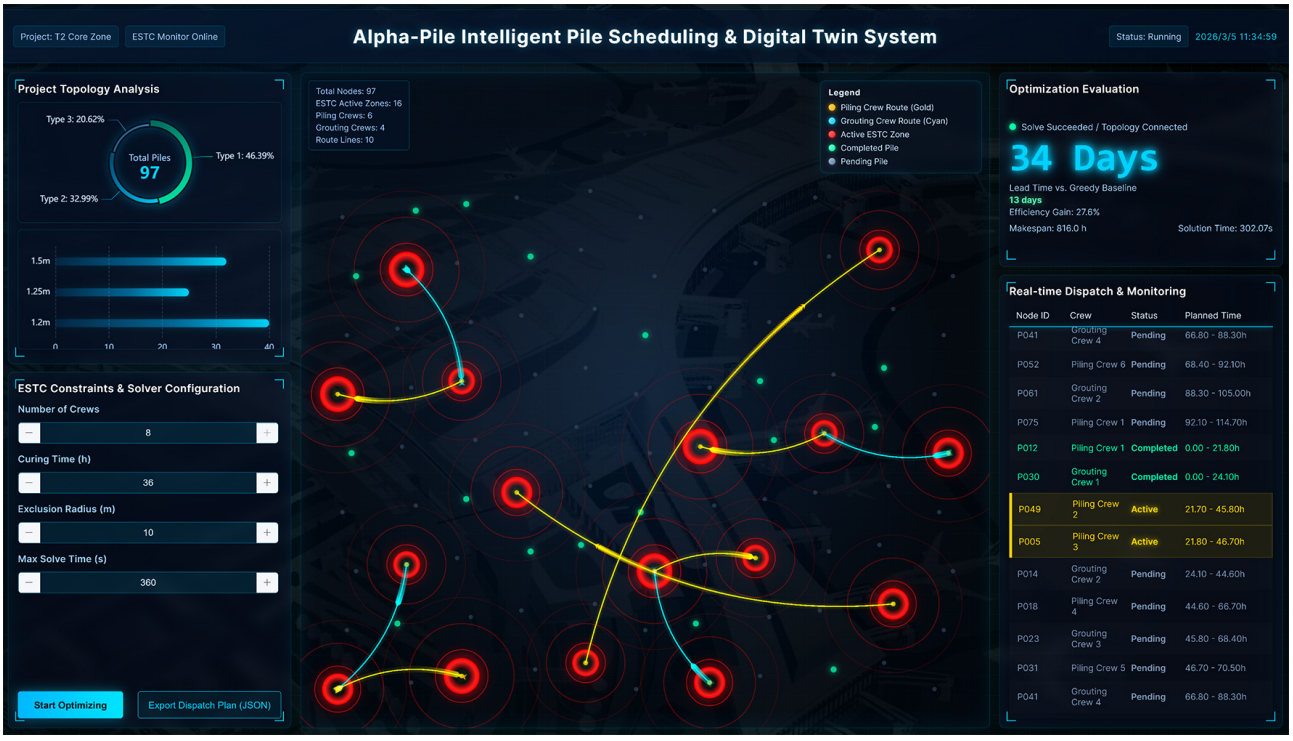


Figure 12. User interface of the system implemented at the construction site

the Maximum Independent Set problem. It strategically sacrifices immediate travel efficiency to preserve topological connectivity. The ablation study confirms the pivotal role of this mechanism: removing the look-ahead factor (Variant 3) caused the makespan to deteriorate significantly to 805.7 hours, mimicking the fragmentation failure mode. By maintaining a sparse distribution of active tasks, the algorithm ensures that the active workfront density remains high throughout the project lifecycle, preventing the late-stage bottlenecks that cripple traditional schedules.

The ablation study further underscores the critical necessity of the dual-colony structure. This necessity arises from the adversarial nature of the heterogeneous resources involved. Piling crews act as constraint generators, creating vibration zones that lock the search space, whereas grouting crews act as constraint relievers, whose completion of tasks eventually unlocks these zones. Attempting to guide these opposing roles with a unified learning mechanism leads to pheromone confusion, as evidenced by Variant 4, which yielded a poor makespan of 825.9 hours. In this single-matrix configuration, the drive to maximize sparsity (for piling) conflicts with the drive to maximize urgency (for grouting). The TA-ACO successfully decouples these objectives. The piling colony focuses on topological preservation, while the grouting colony focuses on bottleneck elimination. This specialized division of labor allows the system to navigate the non-monotonic search landscape where a good decision for one resource type might be a bad decision for the other if not properly coordinated.

6. Conclusions

This study addressed the decision-dependent uncertainty in subsurface engineering by formulating the PFCSP with ESTC. By integrating a MILP model with the TA-ACO, we successfully bridged the gap between theoretical scheduling and the dynamic physical realities of construction sites. The core findings and contributions are summarized as follows:

(1) The proposed TA-ACO algorithm shows stable performance across varying problem scales. In small-scale instances, it matches the global optimality of the exact solver. Crucially, in large-scale scenarios where exact methods fail to converge, TA-ACO establishes a new performance benchmark. In the rigorous test of Instance L-02, the algorithm achieved a makespan of 608.4 hours, significantly outperforming the industry-standard Greedy strategy and the ALNS. This corresponds to a project duration reduction of approximately 21.5% to 28.3% compared to baseline heuristics.

(2) The results indicate that the root cause of scheduling inefficiency under ESTC is spatial fragmentation rather than resource scarcity. Microscopic analysis reveals that traditional nearest-neighbor strategies inadvertently construct isolated work islands, leading to unavoidable deadlocks. The TA-ACO's success stems from its look-ahead factor, which acts as a dynamic sparsity maximizer. By treating physical space as a depletable resource and prioritizing topological connectivity, the algorithm maintains a high active workfront density, ensuring parallel operations can be sustained throughout the project lifecycle.

(3) This research proves that heterogeneous resources in PFCSP – piling and grouting – possess adversarial coupling that creates a non-monotonic search landscape. The ablation study confirms that a unified learning mechanism leads to pheromone confusion and suboptimal results. The dual-colony structure of TA-ACO effectively resolves this by decoupling the optimization objectives: the Piling Colony focuses on preventing spatial blockage, while the Grouting Colony targets the relaxation of temporal bottlenecks, achieving a synergistic global balance.

(4) The theoretical contributions have been translated into practice through the Alpha-Pile decision support system. Field implementation at a mega-aviation hub demonstrated that the optimized schedule reduced the construction period of a critical zone from a projected 47 days to 34 days, securing a 13-day lead time for subsequent superstructure assembly. These results suggest that, in complex subsurface projects, topologically aware algorithmic scheduling can provide useful support for reducing myopic decisions and improving schedule reliability.

Despite these results, several limitations remain. First, the current model assumes deterministic task durations and fixed curing requirements, whereas actual construction may be affected by weather, equipment breakdowns, material supply, crew availability, and geological variability. Future work could incorporate stochastic or robust scheduling formulations to improve schedule reliability under execution uncertainty. Second, the beam-search-based look-ahead mechanism improves spatial planning but increases computational overhead. Although the current runtime is acceptable for static pre-construction scheduling, faster strategies such as parallel computing, rolling-horizon optimization, or incremental schedule repair will be required for real-time adaptive scheduling. Third, the present objective focuses on makespan minimization. Future studies could extend the framework to multi-objective scheduling by considering crew idle time, equipment travel distance, safety buffers, and schedule robustness. Finally, the Alpha-Pile system could be further developed into a digital-twin-based platform, where real-time sensor and progress data are used to trigger dynamic rescheduling in response to on-site deviations.

Acknowledgements

This work was supported by the Yunnan Provincial S&T Program Project under Grant 202303AA080012; the National Natural Science Foundation of China under Grant 52309166; the National Natural Science Foundation of China under Grant 52509174; and the Tianjin Science and Technology Program Project, China (Grant No. 26JRZDSS00320).

Funding

This work was supported by the Yunnan Provincial S&T Program Project under Grant 202303AA080012; the National Natural Science Foundation of China under Grant 52309166; the National Natural Science Foundation of China under Grant 52509174; and the Tianjin Science and Technology Program Project, China (Grant No. 26JRZDSS00320).

Author contributions

JH and MZ conceived the study and were responsible for the design and development of the data analysis. QM, RW, and XX were responsible for data collection and analysis. LH, ML, and MZ were responsible for data interpretation. JH wrote the first draft of the article. All authors reviewed and approved the final manuscript.

Disclosure statement

The authors declare that they have no known competing financial, professional, or personal interests that could have appeared to influence the work reported in this paper.

References

Abotaleb, I., Nassar, K., & Hosny, O. (2016). Layout optimization of construction site facilities with dynamic freeform geometric representations. *Automation in Construction*, *66*, 15–28. https://doi.org/10.1016/j.autcon.2016.02.007

Chen, L., Jiang, C., Huang, Y., & Wu, W. (2026). Uncertainty-aware cloud resource management via fuzzy-adaptive hierarchical reinforcement learning. *Applied Soft Computing*, *186*, Article 114157. https://doi.org/10.1016/j.asoc.2025.114157

Cheng, C., Shan, R., Wu, X., & Lv, S. (2024). An integrated distributionally robust model for two-echelon patient appointment scheduling. *Computers & Industrial Engineering*, *198*, Article 110593. https://doi.org/10.1016/j.cie.2024.110593

Ding, H., Zhuang, C., & Liu, J. (2023). Extensions of the resource-constrained project scheduling problem. *Automation in Construction*, *153*, Article 104958. https://doi.org/10.1016/j.autcon.2023.104958

Ge, Y., & Wang, A. (2021). Spatial scheduling for irregularly shaped blocks in shipbuilding. *Computers & Industrial Engineering*, *152*, Article 106985. https://doi.org/10.1016/j.cie.2020.106985

Han, R., Li, X., Shen, Z., & Jia, D. (2024). A framework of robust project portfolio selection problem under strategic objectives considering the risk propagation. *Engineering, Construction and Architectural Management*, *31*(12), 4872–4896. https://doi.org/10.1108/ECAM-08-2022-0801

Hartmann, S., & Briskorn, D. (2022). An updated survey of variants and extensions of the resource-constrained project scheduling problem. *European Journal of Operational Research*, *297*(1), 1–14. https://doi.org/10.1016/j.ejor.2021.05.004

He, S., Zhou, Y., Yang, Y., Li, X., Liu, Y., & Liu, T. (2025). Vulnerable sequence identification for sequential cascading failure analysis in power grid. *IEEE Transactions on Industrial Informatics*, *21*(6), 4830–4840. https://doi.org/10.1109/TII.2025.3547022

Homem-de-Mello, T., Kong, Q., & Godoy-Barba, R. (2022). A simulation optimization approach for the appointment scheduling problem with decision-dependent uncertainties. *INFORMS Journal on Computing*, *34*(5), 2845–2865. https://doi.org/10.1287/ijoc.2022.1212

Hong, S., & Park, S.-K. (2015). Effect of vehicle-induced vibrations on early-age concrete during bridge widening. *Construction and Building Materials*, *77*, 179–186. https://doi.org/10.1016/j.conbuildmat.2014.12.043

Hu, M., Zhou, M., Zhang, Z., Zhang, L., & Li, Y. (2025). A novel disjunctive-graph-based meta-heuristic approach for multi-objective resource-constrained project scheduling problem with multi-skilled staff. *Swarm and Evolutionary Computation*, *95*, Article 101939. https://doi.org/10.1016/j.swevo.2025.101939

Jing, S., & Fu, W. (2026). Integrated optimization approach for skill-aware and collaborative IC design portfolios decision to enhance R&D project effectiveness and resilience. *Computers & Industrial Engineering*, *213*, Article 111765. https://doi.org/10.1016/j.cie.2025.111765

Kerkhove, L.-P., & Vanhoucke, M. (2017). Optimised scheduling for weather sensitive offshore construction projects. *Omega*, *66*, 58–78. https://doi.org/10.1016/j.omega.2016.01.011

Kwon, B., & Lee, G. M. (2015). Spatial scheduling for large assembly blocks in shipbuilding. *Computers & Industrial Engineering*, *89*, 203–212. https://doi.org/10.1016/j.cie.2015.04.036

Lei, J., Zhang, Y., Chen, Y., Zhao, J., & Ren, W. (2026). A multi-objective construction site layout and safety planning model based on evolutionary algorithms. *Engineering Applications of Artificial Intelligence*, *164*, Article 113177. https://doi.org/10.1016/j.engappai.2025.113177

Li, L., Zhang, J., Zhang, H., & Leus, R. (2025). A novel hyper-heuristic based on surrogate genetic programming for the three-dimensional spatial resource-constrained project scheduling problem under uncertain environments. *Computers & Operations Research*, *179*, Article 107013. https://doi.org/10.1016/j.cor.2025.107013

Ma, D., Cao, X., Zeng, B., Jia, Q.-S., Chen, C., Zhai, Q., & Guan, X. (2025). Proactive robust hardening of resilient power distribution network: Decision-dependent uncertainty modeling and fast solution strategy. *IEEE Transactions on Automation Science and Engineering*, *22*, 14953–14966. https://doi.org/10.1109/TASE.2025.3564483

Mozhdehi, S., Baradaran, V., & Hosseinian, A. H. (2024). Multi-skilled resource-constrained multi-project scheduling problem with dexterity improvement of workforce. *Automation in Construction*, *162*, Article 105360. https://doi.org/10.1016/j.autcon.2024.105360

Ostapska, K., Hingorani, R., Tecław, W., Kaufmann, F., Gradeci, K., & Rüther, P. (2025). Automated robotic deconstruction sequence planning from scan-to-BIM data for reinforced concrete structure reuse. *Automation in Construction*, *180*, Article 106548. https://doi.org/10.1016/j.autcon.2025.106548

Pan, Y., & Zhang, L. (2021). Roles of artificial intelligence in construction engineering and management: A critical review and future trends. *Automation in Construction*, *122*, Article 103517. https://doi.org/10.1016/j.autcon.2020.103517

Pfitzner, F., Braun, A., Borrmann, A., & Bosché, F. (2025). Spatial analysis and complexity evaluation for predicting rebar installation duration. *Automation in Construction*, *179*, Article 106462. https://doi.org/10.1016/j.autcon.2025.106462

Salhab, D., Sabek, M., Pourrahimian, E., Gonzalez, V. A., & Hamzeh, F. (2026). Integrated framework of computer vision and fuzzy systems for lean workspace management in con-

struction. *Automation in Construction*, *181*, Article 106618. https://doi.org/10.1016/j.autcon.2025.106618

Tao, S., Wu, C., Sheng, Z., & Wang, X. (2018). Space-time repetitive project scheduling considering location and congestion. *Journal of Computing in Civil Engineering*, *32*(3), Article 04018017. https://doi.org/10.1061/(ASCE)CP.1943-5487.0000745

Vermeire, G., & Vanhoucke, M. (2025). A priority rule heuristic for the multi-skilled resource-constrained project scheduling problem. *Applied Soft Computing*, *171*, Article 112776. https://doi.org/10.1016/j.asoc.2025.112776

Xu, Z., Lu, X., Guan, H., & Ren, A. (2014). High-speed visualization of time-varying data in large-scale structural dynamic analyses with a GPU. *Automation in Construction*, *42*, 90–99. https://doi.org/10.1016/j.autcon.2014.02.020

Yang, Y., Chen, C., & Li, T. (2024). Automated construction site layout design system for prefabricated buildings using transformer based conditional GAN. *Advanced Engineering Informatics*, *62*(Part D), Article 102885. https://doi.org/10.1016/j.aei.2024.102885

Zhang, X., Ma, S., & Chen, S. (2020). Healthcare service configuration based on project scheduling. *Advanced Engineering Informatics*, *43*, Article 101039. https://doi.org/10.1016/j.aei.2020.101039

Zhang, J., Li, L., Demeulemeester, E., & Zhang, H. (2024). A three-dimensional spatial resource-constrained project scheduling problem: Model and heuristic. *European Journal of Operational Research*, *319*(3), 943–966. https://doi.org/10.1016/j.ejor.2024.07.018

Zhong, H., Zhang, Y., & Gu, Y. (2022). A bi-objective green tugboat scheduling problem with the tidal port time windows. *Transportation Research Part D*: *Transport and Environment*, *110*, Article 103409. https://doi.org/10.1016/j.trd.2022.103409

Zhu, B., Zheng, Y., Luo, C., Zhang, H., Gong, L., Yang, J., & Yang, B. (2025). Experimental investigation and micro-mechanisms of early-age vibration effects on the strength and permeability of fiber reinforced concrete. *Construction and Building Materials*, *504*, Article 144634. https://doi.org/10.1016/j.conbuildmat.2025.144634