Street Art & Urban Creativity | Vol. 12, No. 6, 2026 | ISSN 2183-9956

Street Art & Urban Creativity. Scientific Journal

Infrastructure Project Delivery in Riyadh's Road Network

AI-Assisted Optimization and the Policy Trade-offs of Urban Coordination

Yousif Humeida[1], Abdulrazaq Alqarni[2]

[1] Jacobs Solutions, Riyadh, Saudi Arabia;

[2] Riyadh Infrastructure Projects Center (RIPC), Riyadh, Saudi Arabia

KEYWORDS

urban infrastructure coordination
Tabu Search
spatiotemporal optimization
artificial intelligence
roadworks
synthetic data
Riyadh

ABSTRACT

Coordinating cross-sector infrastructure works is a city-scale scheduling problem with direct implications for service continuity and neighborhood disruption. This study presents an AI-assisted spatiotemporal simulation using Tabu Search and a synthetic portfolio of 200 road-aligned projects on Riyadh's public road network. This approach eliminates all hard-overlaps (intersecting projects occurring simultaneously) while reducing proximity-overlaps (simultaneous projects within 1 km) by 76.4% relative to greedy coordination. However, reducing proximity overlaps increases delivery delays. Sensitivity analysis exposes a Pareto frontier between disruption reduction and delivery delay, showing that AI optimization alone is not sufficient and policymakers must determine acceptable trade-offs.

1. Introduction

Cities rarely deliver infrastructure as isolated projects. Road renewal, water, wastewater, drainage, electricity, telecommunications, streetscape, and public realm works compete for the same corridors and affect the same communities. Each project may be technically viable on its own while the combined portfolio creates avoidable excavation, access restrictions, traffic diversions, service risk, and repeated disruption. The coordination question is therefore not only whether individual projects can start as requested, but whether the city should allow multiple projects to occupy the same or nearby corridors at the same time.

This problem is particularly important in rapidly growing metropolitan regions, where infrastructure portfolios involve many asset owners, permitting authorities, utility providers, contractors, consultants, and public agencies. In Riyadh, the Riyadh Infrastructure Projects Center (RIPC) describes integrated planning, stakeholder coordination, data-driven delivery, cost efficiency, and the reduction of infrastructure-related disruption as core objectives (RIPC, n.d.). A public report stated that RIPC have analyzed and coordinated more than 200,000 current and future projects to avoid overlap with planned works, illustrating the scale at which portfolio-level coordination can influence urban expenditure and service quality (Saudi Press Agency, 2024).

To enable such coordination, RIPC have developed and launched the Integrated Infrastructure Coordination Plan (IICP) (Riyadh Daily, 2025). The IICP provides the motivating institutional context for this study. Its practical challenge is to convert a large collection of project locations, requested dates, durations, priorities, and interdependencies into a schedule that is feasible at corridor level and defensible at city level. Due to the sensitivity and confidentiality of the IICP projects data, this study reproduces the challenges the IICP has been developed to address through a synthetic, road-network-based simulation. This simulation does not reproduce actual IICP projects. Instead, it preserves selected structural features of the coordination problem: projects follow public road geometry; each project has a requested start, duration, and priority; some corridors intersect directly; and other projects are close enough that simultaneous execution could compound local disruption. Such synthetic data allow the scheduling logic, objective design, and trade-offs to be reviewed without exposing real project records.

Tabu Search has been adopted as the optimization component in this study's simulation. Tabu Search is a memory-based metaheuristic within the broader family of computational-intelligence methods. It searches a large combinatorial schedule space, remembers recently reversed moves, and accepts strategically selected changes that can escape local optima.

Three research questions guide this study. First, can a synthetic road-network representation reveal the scale and distribution of direct and nearby project interactions in a citywide portfolio? Second, can a Tabu Search schedule reduce simultaneous nearby works while preserving zero direct work-zone clashes? Third, how does the preferred schedule change when planners assign different importance to delivery delay, proximity disruption, completion time, and program makespan?

The contribution is fivefold. The paper develops a reproducible urban-infrastructure simulation using public road geometry; separates non-negotiable hard clashes from softer proximity-based disruption; compares several scheduling approaches on the same portfolio; quantifies delay distribution and priority effects; and exposes a policy frontier showing that different weight settings produce materially different urban outcomes. The result is a framework that can be adopted for testing coordination policies before applying them to real-world operational data.

2. Literature Review

Project scheduling research provides a mature foundation for planning and coordinating activities while accounting for task dependencies, limited resources, and project time

constraints. The Resource-Constrained Project Scheduling Problem (RCPSP) represents the core framework in this field, determining the start times of activities while satisfying precedence relationships and resource capacity constraints (Brucker et al., 1999). Over time, researchers have extended RCPSP to better reflect real-world projects by incorporating multiple concurrent projects, alternative execution methods, uncertainty, and multiple optimization objectives (Hartmann & Briskorn, 2010, 2022). More recently, research on Resource-Constrained Multi-Project Scheduling (RCMPS) has highlighted that optimizing projects independently does not necessarily produce the best overall system performance, as projects often compete for shared resources and interact across organizational boundaries (Gómez Sánchez et al., 2023).

Urban infrastructure coordination shares a similar system-level perspective but differs from the standard formulation. Rather than optimizing the sequence and timing of individual project activities, the unit of analysis in this study is the infrastructure project corridor, where multiple utility projects are coordinated within a shared spatial and temporal context. The objective therefore extends beyond efficient project execution to identifying coordination opportunities that reduce repeated excavations, minimize disruption, and maximize the benefits of integrated infrastructure delivery.

Construction planning also has a long spatial tradition. Riley and Sanvido (1995) demonstrated that construction activities use identifiable space patterns, and Thabet (1992) treated work space as a constraint that interacts with resources and schedule. Later 4D and workspace-interference studies linked temporal overlap with physical adjacency and showed that parallel activities can become infeasible or inefficient when their work areas intersect (Moon et al., 2014). Ardila and Francis (2020) reviewed spatiotemporal planning and concluded that location-based, takt, and space-planning approaches emerged partly because conventional activity lists do not represent production flow and spatial conflict adequately.

While most workspace scheduling research has focused on buildings, industrial facilities, or other bounded construction sites, urban infrastructure projects operate within a fundamentally different environment. Road corridors are shared public assets where construction activities can simultaneously affect traffic flow, access to properties, utility services, emergency response, and neighboring projects. As a result, spatial relationships become a critical component of project coordination. Geographic Information Systems (GIS) provide the analytical capabilities needed to capture these relationships by supporting operations such as buffering, proximity analysis, spatial indexing, intersection detection, and the transformation of network geometry into scheduling constraints. OpenStreetMap offers an openly accessible source of road network data (Haklay & Weber, 2008), while OSMnx enables the efficient retrieval, projection, modeling, and analysis of these networks within a reproducible Python-based workflow (Boeing, 2025).

Metaheuristics are widely used when the search space is too large for exhaustive enumeration or when the objective combines several competing considerations. Tabu Search uses adaptive memory to discourage immediate reversal, manage intensification and diversification, and permit aspiration when a tabu move improves the best-known solution (Glover, 1989; Glover & Laguna, 1997). Thomas and Salhi (1998) applied Tabu Search to resource-constrained project scheduling, demonstrating the relevance of structured neighborhoods, tabu tenure, and objective approximations. These features are useful for corridor coordination because moving one project may resolve one overlap while creating delay or new proximity interactions elsewhere.

The AI and smart-city literature broadens the evaluation criteria. AI can support urban management, planning, mobility, and infrastructure, but implementation benefits depend on governance, data quality, institutional capability, and attention to risk (Yigitcanlar et al., 2020; Wolniak & Stecuła, 2024). Cugurullo (2020) argues that urban AI should be understood within its political and institutional context rather than treated as an autonomous technical layer. For infrastructure coordination, this means that an optimization score cannot substitute for public policy. Objective weights encode choices about whose delays matter, which disruptions are unacceptable, and how much schedule flexibility is available.

Digital-twin research similarly emphasizes the value of linking a digital representation to urban decision making. City digital twins can combine geometry, state, scenarios, and

stakeholder interaction, but their usefulness depends on clear purposes and credible data rather than visual fidelity alone (Dembski et al., 2020; Ketzler et al., 2020). This study is not a full digital twin because it has no live data connection or bidirectional operational control. It is better characterized as a synthetic simulation component that could be integrated into a broader planning platform.

The literature therefore provides four complementary foundations for this study: multi-project scheduling, spatial construction planning, GIS-enabled urban modeling, and AI-assisted optimization. Building on these foundations, the study examines how tools such as the IICP can support the coordinated delivery of urban infrastructure projects, improve network-level efficiency, and reduce the adverse impacts of construction in a rapidly growing city such as Riyadh.

3. Methodology

A structured approach has been adopted in this study. A citywide infrastructure coordination problem is first defined, after which a synthetic portfolio of projects is created to represent the key characteristics of infrastructure delivery in Riyadh. The spatial and temporal relationships between projects are then translated into a scheduling model that distinguishes between hard-overlap and proximity-overlap. Hard-overlap occurs when the work zones of two projects physically intersect and are scheduled to be active during the same period, creating an infeasible conflict that must be eliminated. Proximity-overlap occurs when two projects are located within a predefined distance of one another and are active simultaneously. While proximity-overlap does not make a schedule infeasible, it is treated as undesirable because it may increase cumulative disruption to road users, nearby businesses, and surrounding communities.

The synthetic project portfolio is subsequently evaluated using three scheduling approaches: the requested schedule, the greedy coordinated schedule, and the Tabu Search optimized schedule. Their performance is assessed using a consistent set of evaluation metrics, including hard-overlap, proximity-overlap, project delay, and overall duration.

The framework developed in this study is designed to enable different scheduling strategies to be compared and the trade-offs between timely project delivery and reduced urban disruption to be systematically evaluated.

The scheduling problem is represented as a portfolio of n infrastructure projects distributed across a shared urban road network. Each project i is defined by a road-aligned geometry (g_i), a requested start day (r_i), a fixed duration (d_i), and a priority value (p_i). Priority represents the relative importance of a project compared with other projects in the portfolio. Priority assignments may be based on considerations such as service criticality, project type, strategic importance, and the population affected. For instance, maintenance works may be prioritized over upgrades, water and electricity infrastructure may be considered more critical than telecommunications or irrigation assets, and transmission facilities serving larger populations may be given greater priority than local distribution connections.

The decision variable is the actual start day (s_i) assigned to each project. Projects may be delayed to improve coordination but cannot start earlier than their requested start day. Once construction begins, each project is assumed to remain continuously active until its completion.

The scheduling objective is to determine project start days that eliminate hard-overlap while reducing proximity-overlap and limiting the delivery impacts associated with coordination. These impacts are represented by project delay, priority-weighted completion time, and overall programme makespan. For simplicity, the developed framework does not consider resource availability, contractor capacity, permitting dependencies, or activity-level sequencing.

To evaluate the proposed framework, three scheduling approaches are applied to the same project portfolio. The requested schedule preserves the original requested start dates and serves as the uncoordinated baseline. The greedy coordinated schedule applies a deterministic hard-overlap avoidance rule to generate the first feasible solution. The Tabu Search optimized schedule uses the greedy schedule as its initial solution and iteratively searches for improved

schedules that further reduce proximity-overlap while maintaining acceptable project delivery performance.

The effectiveness of the proposed framework is evaluated by applying each scheduling approach to the same synthetic project portfolio and comparing the resulting schedules using a consistent set of performance indicators. The evaluation considers the framework's ability to eliminate hard-overlap, reduce proximity-overlap, and quantify the associated delivery impacts through measures such as project delay and programme makespan. In addition, a sensitivity analysis is undertaken to examine how changes to the objective-function coefficients influence the balance between reducing urban disruption and maintaining timely project delivery. This provides insight into the trade-offs associated with different coordination policies and demonstrates the flexibility of the proposed framework for supporting alternative planning objectives. Figure 1 summarizes this sequence.

Figure 1. The synthetic Riyadh road-network simulation methodology.

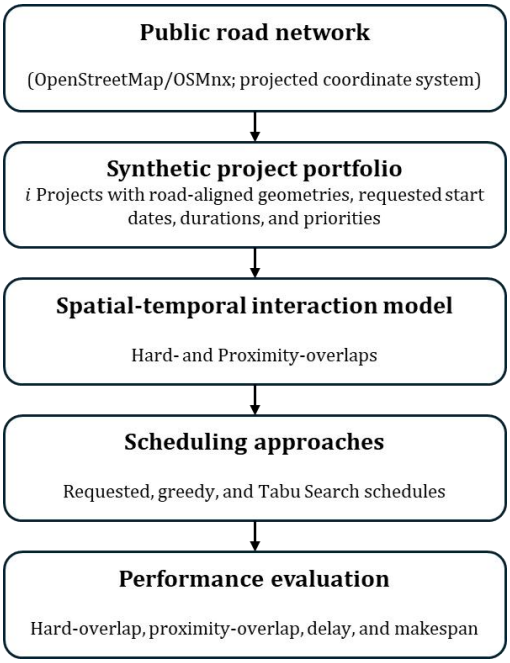

Source: Authors' model design.

3.1. Public Road Network

The study area was defined by a 20 km radius centered at latitude 24.7136 and longitude 46.6753 in Riyadh. A drivable road network was downloaded from OpenStreetMap through OSMnx and projected to a meter-based coordinate reference system.

3.2. Synthetic project portfolio

Synthetic infrastructure projects were generated by randomly selecting 200 road-aligned corridors approximately 3 km in length from the OpenStreetMap road network. A fixed random seed was used to ensure that the project portfolio could be reproduced.

Each project was assigned a requested start day (r_i) selected uniformly between Day 0 and Day 365. This interval represents the planning horizon for project commencement rather than programme completion, allowing projects that begin later in the year to finish beyond Day 365. Project durations (d_i) were sampled from a triangular distribution with a minimum of 30 days, a most likely value of 60 days, and a maximum of 90 days. Projects were also assigned one of three priority (p_i) levels, with higher-priority projects (i.e., $p_i = 3$) receiving greater weighting during schedule optimization.

Figure 2 illustrates the resulting synthetic project portfolio, showing the spatial distribution of the generated road-aligned corridors across the study area. The corridor colours represent the requested project start day.

Each project will have an actual start date (s_i) which will be updated based on the scheduling decisions and as a result a project finish date (f_i) where $f_i = s_i + d_i$.

Figure 2. The synthetic Riyadh road-network simulation methodology.

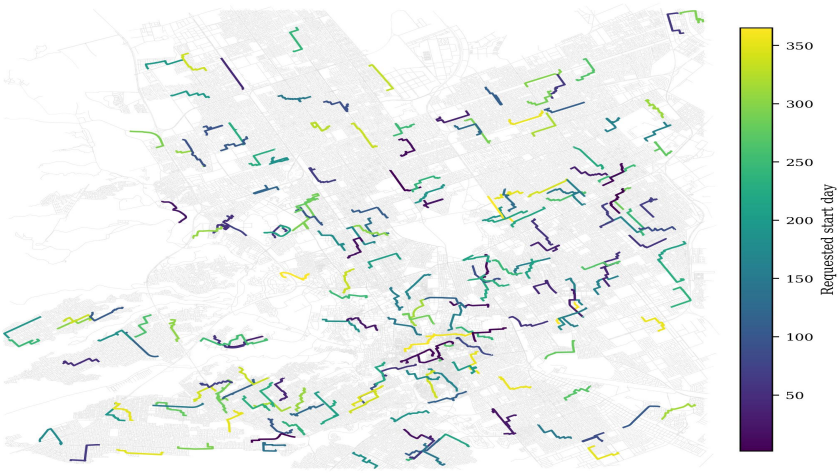

Source: Authors' simulation output.

The requested schedule will already have some hard-overlaps and proximity overlaps. Figure 3 below shows an example for each.

Figure 3. Hard- and Proximity-overlap examples: (a) geospatial location on Riyadh map; (b) geospatial location zoom-in; and (c) requested schedule.

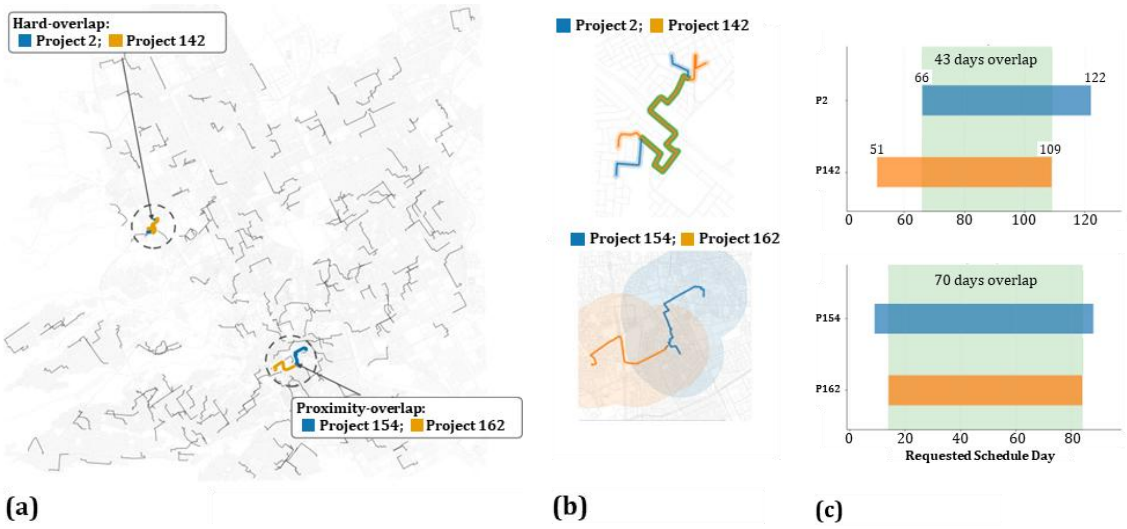

Source: Authors' simulation output.

3.3. Spatial-temporal interaction model

To evaluate any given project schedule, the occurrence of hard-overlap and proximity-overlap is determined through a two-stage assessment. First, the temporal relationship between every pair of projects is evaluated. Spatial relationships are then assessed only for projects that overlap in

time. A hard-overlap or proximity-overlap is therefore recorded only when both temporal and spatial overlap conditions are satisfied.

The temporal overlap between projects i and j is defined as:

$$O_{ij} = max \quad \{0,\ min(f_i, f_j) - max(s_i, s_j)\} \tag{1}$$

where O_{ij} represents the number of days during which both projects are simultaneously active. A value of zero indicates that no temporal overlap exists.

Hard-overlap is evaluated using buffered project work zones. For each project, a work-zone polygon W_{ij} is created by applying a buffer of width b to the project geometry g_i,

$$W_i = buffer(g_i, b) \tag{2}$$

For the purposes of this study, $b = 30\ m$ is considered a reasonable buffer assumption for hard-overlaps.

A project pair is considered a hard-overlap when their work zones intersect and they are active simultaneously:

$$H_{ij} = \begin{cases} 1, & W_i \cap W_j \neq \emptyset \ \ and \ \ O_{ij} > 0 \\ 0, & otherwise. \end{cases} \tag{3}$$

where H_{ij} is a hard-overlap indicator for the project pair i and j.

Projects that do not satisfy the hard-overlap condition are subsequently evaluated for proximity-overlap. Let $d(g_i, g_j)$ denote the shortest distance between the two project geometries. A proximity-overlap exists when the projects are within a predefined proximity threshold r and are simultaneously active,

$$P_{ij} = \begin{cases} 1, & d(g_i, g_j) \leq r \ \ and \ \ O_{ij} > 0 \ \ and \ \ H_{ij} = 0 \\ 0, & otherwise. \end{cases} \tag{4}$$

where $r = 1\ km$ has been adopted for this study.

For project pairs exhibiting proximity-overlap, a distance-based weighting is assigned as

$$w_{ij} = 1 - \frac{d(g_i, g_j)}{r} \tag{5}$$

so that closer projects incur larger penalties than projects located near the edge of the proximity threshold.

3.4. Scheduling approaches

The scheduling problem is formulated as the minimization of an objective function that balances the competing objectives of reducing hard-overlap, minimizing proximity-overlap, limiting project delay, prioritizing earlier completion of higher-priority projects, and reducing the overall programme duration. The objective function is expressed as

$$\min_{\mathbf{s}} J = \alpha_H H + \alpha_P P + \alpha_D D + \alpha_C C + \alpha_M M \tag{6}$$

where s is the vector of project start days.

The five objective-function components are defined in Equations (7a)–(7e). H denotes the total hard-overlap days which is the total number of days during which physically intersecting project work zones are simultaneously active, P is the weighted proximity-overlap score which represents The duration and severity of simultaneous nearby projects, with closer projects receiving greater weight, D is the priority-weighted delay, represnting the total delay relative to requested start dates, weighted by project priority, C the priority-weighted completion time which is the sum of project completion times, weighted by priority, and M the programme makespan which is the completion time of the latest-finishing project.

The coefficients α_H, α_P, α_D, α_C and α_M are non-negative weighting parameters that control the relative importance assigned to each objective component. Different combinations of these coefficients represent different coordination policies. For example, increasing α_P encourages the optimizer to reduce proximity-overlap, whereas increasing α_D favours earlier project delivery.

The generic formulation therefore provides the flexibility to evaluate alternative planning priorities.

The values of objective function coeffiients for the primary implementation adopted in this study are presented in table 1. The adopted coefficients define the relative importance assigned to the different objective-function components. They were selected to ensure that hard-overlaps were eliminated, prioritize the reduction of proximity-overlaps during optimization, and limit excessive schedule movement. While these coefficient alues represent one implementation, the influence of different coefficient values is further investigated through the sensitivity analysis.

Table 1. Objective function coefficient values for primary simulation run

α_H	α_P	α_D	α_C	α_M
10,000,000	500	25	1	5

Source: Authors' simulation output.

$$H = \sum_{i<j} H_{ij} O_{ij} \tag{7a}$$

$$P = \sum_{i<j} P_{ij} w_{ij}\, O_{ij} \tag{7b}$$

$$D = \sum_{i} p_i \max\{0, s_i - r_i\} \tag{7c}$$

$$C = \sum_{i} p_i f_i \tag{7d}$$

$$M = \max_i f_i \tag{7e}$$

To evaluate the proposed scheduling framework, three scheduling approaches of increasing sophistication were applied to the same synthetic project portfolio.

For the greedy coordinated schedule and the Tabu Search optimized schedule, hard-overlap is also enforced as an explicit feasibility condition. The hard-overlap term is nevertheless retained in the generic score so that the infeasible requested schedule can be evaluated consistently and to provide a defensive penalty if the feasibility screening is altered in later implementations.

Requested Schedule

The requested schedule assigns $s_i = r_i$ for every project. It represents the portfolio before coordination and can contain both hard-overlap and proximity-overlap. No optimization is applied. The schedule is retained as a reference for measuring the amount of intervention required by the coordinated approaches.

Greedy Coordinated Schedule

The greedy coordinated schedule processes projects in ascending requested-start order, with higher-priority projects processed first when requested dates are equal. For each project, the algorithm tests the requested start against projects already placed in the schedule. If one or more hard-overlap candidate projects would be active simultaneously, the project is moved to the latest finish day among the current blockers. The test is repeated until no hard-overlap remains. This method approximates a straightforward coordination rule in which an earlier or higher-priority project proceeds and the blocked project waits.

The greedy method is deterministic for a fixed portfolio and produces the initial hard-feasible solution for Tabu Search. It is not designed to reduce proximity-overlap unless that reduction

occurs incidentally while resolving hard-overlap. Its role is therefore to separate the benefit of simple feasibility enforcement from the additional benefit of metaheuristic search.

Tabu Search Optimized Schedule

In this study, a solution is the vector of 200 actual start days. The primary search begins from the greedy coordinated schedule, samples candidate single-project shifts, rejects any candidate that would create hard-overlap, and scores the remaining candidates using the objective function in Equation (6).

At each iteration, the best admissible candidate among the sampled moves becomes the next current solution, even when it does not improve the global best. This allows the search to move out of local minima. A tabu record prevents the moved project from immediately returning to its previous start day. The aspiration rule overrides tabu status when the candidate improves the best-known objective. The best solution observed over all iterations is returned as the Tabu Search optimized schedule.

3.5. Performance evaluation

The proposed framework was evaluated through two complementary experiments. The first experiment compared the performance of the three scheduling approaches using a common synthetic project portfolio to assess their ability to coordinate infrastructure projects and reduce spatial-temporal interactions. The second experiment investigated the sensitivity of the optimization model to different objective-function coefficients, allowing the influence of alternative planning priorities on the resulting schedules to be examined.

The performance of each schedule was assessed using a consistent set of evaluation metrics, including hard-overlap days, proximity-overlap score, project delay, and programme makespan. The reported Tabu Search implementation used the objective function presented in Equation (9), with a maximum of 3,000 iterations and 120 candidate moves evaluated during each iteration.

The coefficient sensitivity analysis was undertaken to investigate how different planning priorities influence the optimized schedule. Rather than identifying a single optimal set of coefficients, the experiment examined how alternative weight combinations affect the balance between minimizing urban disruption and maintaining timely project delivery.

A range of coefficient combinations representing different coordination policies was evaluated, from delivery-oriented configurations that emphasized minimizing project delay to disruption-oriented configurations that emphasized reducing proximity-overlap. Throughout the analysis, the hard-overlap coefficient remained fixed to ensure that feasible schedules were maintained. A summry of the coefficient combinations selected for the sensetivity assessment is shwon in table 2. The eleven coefficient configurations (W00–W10) represent a progression from delivery-oriented to proximity-oriented coordination policies. The delivery-oriented configurations assigned greater importance to minimizing project delay and programme duration, whereas the proximity-oriented configurations increasingly prioritized reducing simultaneous nearby works. Intermediate configurations provided a balanced compromise between these objectives, allowing the trade-offs associated with different coordination strategies to be systematically evaluated.

Table 2. Objective function coefficient values for sensitivity assessment runs

	α_H	α_P	α_D	α_C	α_M
$W00$	10,000,000	0	250	2.00	30.00
$W01$	10,000,000	25	200	2.00	25.00
$W02$	10,000,000	50	150	1.50	20.00
$W03$	10,000,000	100	100	1.25	15.00
$W04$	10,000,000	200	75	1.00	10.00
$W05$	10,000,000	500	50	1.00	5.00

W06	10,000,000	1,000	35	0.75	3.00
W07	10,000,000	2,500	20	0.50	2.00
W08	10,000,000	5,000	10	0.25	1.00
W09	10,000,000	10,000	5	0.10	0.50
W10	10,000,000	25,000	2	0.05	0.25

Source: Authors' simulation assumptions.

Each sensitivity scenario was initialized using the greedy coordinated schedule and optimized independently using the Tabu Search algorithm. The same optimization procedure was applied across all scenarios to provide a consistent basis for comparison. The resulting schedules were subsequently evaluated to examine the trade-offs between project delivery performance and urban disruption under different policy priorities.

4. Results

4.1. Primary simulation run

Table 3 compares the three schedules. As mentioned in section 3.2, the requested uncoordinated schedule contained 24 hard-overlaps and 917 hard-overlap days. It also contained 55 proximity-overlaps with proximity-overlap score of 901.9.

The greedy coordinated schedule eliminated all hard-overlaps and hard-overlap days. It did so with an average delay of 5.0 days and a maximum delay of 91 days. However, it retained all 55 proximity-overlaps with a proximity-overlap score of 845.5, showing that resolving direct clashes alone does not necessarily reduce local construction concentration.

The Tabu optimized schedule also maintained zero hard-overlaps and zero hard-overlap days. Relative to the greedy schedule, it reduced the proximity-overlap score from 845.5 to 29.3 and reduced proximity-overlaps from 55 to 13. This represents a 96.5% reduction in weighted proximity score and a 76.4% reduction in proximity-overlaps. The trade-off was an increase in average delay from 5.0 to 15.7 days and an increase in makespan from Day 493 to Day 544.

Table 3. Primary run performance of requested, greedy and Tabu optimized schedules

	Hard-overlap	Hard-overlap days	Proximity-overlap	Proximity-overlap score	Mean delay	Max delay	Makespan
Requested	24	917	55	901.9	0.0	0	439
Greedy	0	0	55	845.5	5.0	91	493
Tabu	0	0	13	29.3	15.7	255	544

Source: Authors' simulation output.

Figure 4. The synthetic Riyadh road-network simulation methodology.

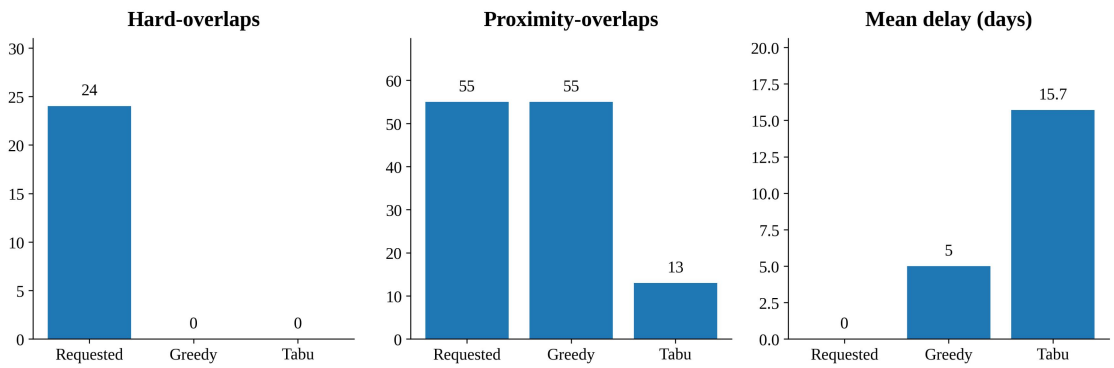

Source: Authors' simulation output.

Figure 5 shows the value of the objective function resulting from all of the 3000 Tabu Search iterations. The primary search improved rapidly and then refined the solution more gradually. The best recorded objective declined from 510,724 at the initial sampled state to 236,663 by iteration 50, 218,550 by iteration 100, and 205,285 at the end of the recorded run. This pattern suggests that the search found the major proximity reductions early, while later iterations made smaller adjustments among already feasible schedules.

Figure 5. Objective function values during the reported Tabu Search run.

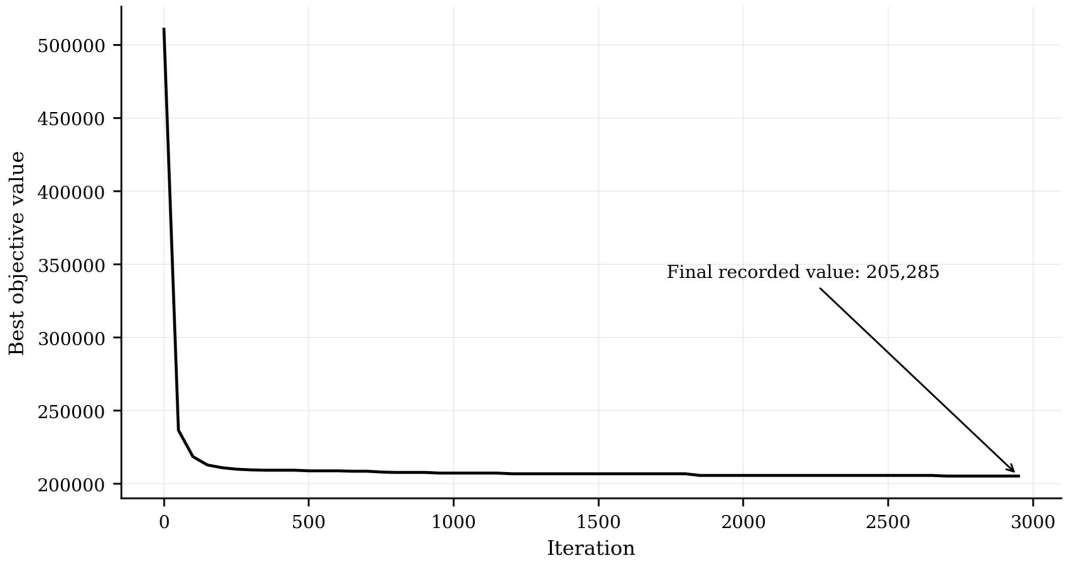

Source: Authors' simulation output.

4.2. Sensitivity assessment

The sensitivity analysis shows that the framework is not limited to producing one optimized schedule. Instead, it can generate a pareto frontier for making policy decisions. In delivery-first settings, the optimizer retains low delay but leaves many proximity overlaps. As proximity penalties increase, the optimizer progressively reduces nearby simultaneous works at the cost of additional delay. The trade-off is nonlinear. Moving from W00 to W04 increases mean delay from 5.0 to 9.8 days while reducing proximity-overlap pairs from 55 to 33. The balanced W05 configuration increases mean delay to 14.5 days and reduces proximity-overlap pairs to 15. More proximity-oriented configurations reduce the remaining pairs to five, two, and ultimately zero, but with mean delays exceeding approximately 18-26 days depending on the configuration. The endpoint in which proximity is treated as a hard constraint also reaches zero proximity overlaps but with a mean delay of 23.0 days and a makespan of 583 days. More details of the outputs of the sensitivity analysis are show in table 4 and figure 6

Table 4. Sensitivity assessment schedule runs performance

	Hard-overlap	Hard-overlap days	Proximity-overlap	Proximity-overlap score	Mean delay	Max delay	Makespan
W00	0	0	55	845.50	5.02	91	493
W01	0	0	55	845.50	5.02	91	493
W02	0	0	55	845.50	5.02	91	493
W03	0	0	53	808.95	5.16	91	493
W04	0	0	33	248.17	9.82	161	493

W05	0	0	15	48.16	14.53	324	508
W06	0	0	5	7.09	18.61	486	575
W07	0	0	2	1.75	22.53	484	572
W08	0	0	0	0.00	25.91	484	578
W09	0	0	0	0.00	22.71	484	572
W10	0	0	0	0.00	24.90	485	574

Source: Authors' simulation output.

Figure 6. Pareto frontier of mean delivery delay versus proximity-overlaps and weighted proximity-overlap scores. Point labels on the frontier indicate remaining proximity-overlap pair counts.

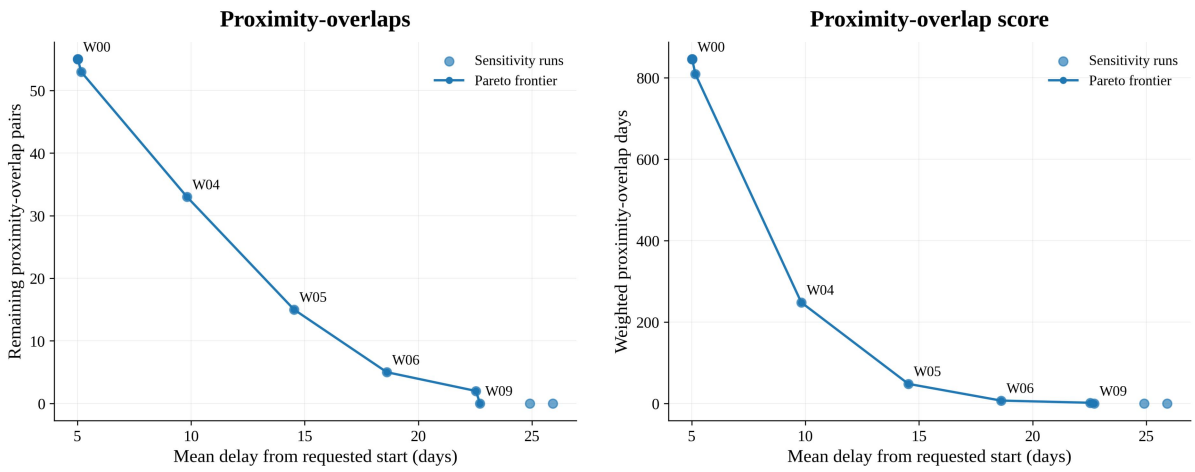

Source: Authors' simulation output.

5. Discussion and conclusion

The results demonstrate that citywide infrastructure coordination involves more than the removal of direct scheduling conflicts. While hard-overlap can be eliminated through relatively simple coordination rules, reducing the concentration of simultaneous nearby works requires a broader portfolio-level optimization approach. The comparison between the requested, greedy coordinated, and Tabu Search schedules therefore reveals two distinct but related planning objectives: maintaining a clash-free programme and minimizing the cumulative disruption created by concurrent works across the urban network.

The first finding is that eliminating hard-overlap and managing proximity-overlap are distinct planning objectives. The greedy coordinated schedule was effective in removing all hard-overlap with only a limited increase in average delay. However, it did not reduce the number of projects operating simultaneously within nearby areas. A portfolio may therefore be free from direct spatial-temporal conflicts while still concentrating roadworks within the same neighbourhoods. This distinction is important for public communication and performance reporting: achieving "zero hard-overlap" should not be interpreted as achieving minimum urban disruption.

The second finding is that Tabu Search can move beyond the immediate blocker logic applied by the greedy approach. While preserving the elimination of hard-overlap, the algorithm redistributed project start dates across the portfolio to reduce proximity-overlap. The 96.5% reduction in the weighted proximity-overlap score is substantial because the metric reflects both the duration of simultaneous works and the distance between projects.

The third finding is that developing a pareto frontier for policy decision making is more useful than any single optimized schedule. Infrastructure agencies operate under changing priorities and constraints. During time-critical delivery periods, configurations similar to W00–W03 may be appropriate because they protect project dates while accepting higher levels of nearby simultaneous work. In contrast, configurations such as W05–W07 may be more suitable

for sensitive districts, major events, school periods, emergency routes, or public-realm programmes where disruption reduction is given greater importance. Treating all proximity-overlaps as prohibited may be justified only in exceptional circumstances, as the model shows that doing so can result in substantial increases in delay and makespan.

For Riyadh, the framework is aligned with the institutional need to coordinate large numbers of projects across multiple service providers. However, it should be applied within a human-governed decision process. Planners should define the hard constraints, approve the proximity criteria, review the projects experiencing the greatest schedule changes, and assess the neighbourhoods most affected. Rather than returning a single recommendation, the optimization process should provide several alternatives together with clear explanations of which overlaps were removed, which projects were moved, why each change was made, and how the selected policy compares with other available options. Planners should also be able to lock dates, protect critical projects, define event windows, and reject algorithmic recommendations where necessary.

The framework developed in planning tools such as the IICP and simulated using synthetic data in this study is potentially transferable beyond Riyadh because its core elements (project geometries, requested schedules, hard-overlap relationships, proximity-overlap relationships, and policy weights) are generic. However, the model parameters and objectives should not be transferred without local calibration. Cities differ in road-network structure, permitting practices, construction seasons, public-transport dependence, governance arrangements, and tolerance for delay. Application in another context should therefore begin with a local definition of the coordination problem.

6. Acknowledgements

While this study used synthetic project data and publicly available road-network data, RIPC and Jacobs Solutions nonetheless provided valuable institutional support that helped shape the research. Therefore, the authors gratefully acknowledge both organizations.

References

Ardila, F., & Francis, A. (2020). Spatiotemporal planning of construction projects: A literature review and assessment of the state of the art. *Frontiers in built environment*, *6*, 550305.

Boeing, G. (2025). Modeling and analyzing urban networks and amenities with OSMnx. *Geographical Analysis*, *57*(4), 567-577.

Brucker, P., Drexl, A., Möhring, R., Neumann, K., & Pesch, E. (1999). Resource-constrained project scheduling: Notation, classification, models, and methods. *European journal of operational research*, *112*(1), 3-41.

Cugurullo, F. (2020). Urban artificial intelligence: From automation to autonomy in the smart city. *Frontiers in Sustainable Cities*, *2*, 38.

Dembski, F., Wössner, U., Letzgus, M., Ruddat, M., & Yamu, C. (2020). Urban digital twins for smart cities and citizens: The case study of Herrenberg, Germany. *Sustainability*, *12*(6), 2307.

Glover, F. (1989). Tabu search—part I. *ORSA Journal on computing*, *1*(3), 190-206.

Glover, F., & Laguna, M. (1997). *Tabu search* (Vol. 104). Boston: Kluwer academic publishers.

Sánchez, M. G., Lalla-Ruiz, E., Gil, A. F., Castro, C., & Voß, S. (2023). Resource-constrained multi-project scheduling problem: A survey. *European Journal of Operational Research*, *309*(3), 958-976.

Haklay, M., & Weber, P. (2008). Openstreetmap: User-generated street maps. *IEEE Pervasive computing*, *7*(4), 12-18.

Hartmann, S., & Briskorn, D. (2010). A survey of variants and extensions of the resource-constrained project scheduling problem. *European Journal of Operational Research*, 207(1), 1-14.

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.

Herroelen, W. (2005). Project scheduling—Theory and practice. *Production and operations management*, *14*(4), 413-432.

Ketzler, B., Naserentin, V., Latino, F., Zangelidis, C., Thuvander, L., & Logg, A. (2020). Digital twins for cities: A state of the art review. *Built environment*, *46*(4), 547-573.

Kolisch, R., & Hartmann, S. (2006). Experimental investigation of heuristics for resource-constrained project scheduling: An update. *European journal of operational research*, *174*(1), 23-37.

Moon, H., Kim, H., Kim, C., & Kang, L. (2014). Development of a schedule-workspace interference management system simultaneously considering the overlap level of parallel schedules and workspaces. *Automation in Construction*, *39*, 93-105.

Riley, D. R., & Sanvido, V. E. (1995). Patterns of construction-space use in multistory buildings. *Journal of Construction Engineering and management*, *121*(4), 464-473.

Riyadh Daily. (2025, February 3). *With local and international presence: RIPC at Riyadh Region showcases the ceremony of the Integrated Infrastructure Coordination Plan*. https://alriyadhdaily.com/index.php/article/c8310e02f5b143ff9c378f948926660e

Riyadh Infrastructure Projects Center. (n.d.). Our mission. Retrieved July 18, 2026, from https://ripc.gov.sa/en/about

Saudi Press Agency. (2024, October 9). Riyadh Infrastructure Project Center enhances road quality sustainability. https://www.spa.gov.sa/en/N2185688

Thabet, W. Y. (1992). *A space-constrained resource-constrained scheduling system for multi-story buildings* (Doctoral dissertation, Virginia Polytechnic Institute and State University). Virginia Polytechnic Institute and State University.

Thomas, P. R., & Salhi, S. (1998). A tabu search approach for the resource constrained project scheduling problem. *Journal of heuristics*, *4*(2), 123-139.

Wolniak, R., & Stecuła, K. (2024). Artificial intelligence in smart cities—applications, barriers, and future directions: a review. *Smart cities*, *7*(3), 1346-1389.

Yigitcanlar, T., Desouza, K. C., Butler, L., & Roozkhosh, F. (2020). Contributions and risks of artificial intelligence (AI) in building smarter cities: Insights from a systematic review of the literature. *Energies*, *13*(6), 1473.