

ZERO-WAIT STS

A predictive coordination layer that eliminates dead time between STS crane moves

Praharsh Pulla, Thanu Reddy Pathipati, Kamaleswar Urranki, Sai Naga Krishna Kaushik Nerusu

Academic Advisor: Dr. Lawrence Henesey

Blekinge Institute of Technology, Sweden

2026 Ottonel Popesco Student Challenge by PEMA

Abstract. Every container terminal in the world measures gross crane productivity in moves per hour, yet the most consistent source of lost productivity is not how fast the STS Cycles but how long it sits idle between moves waiting for the next ITV (Internal Transfer Vehicle), the next instruction, or the next decision. This proposal introduces **ZERO-WAIT STS**, a predictive coordination layer that prepares each move before the current one ends, assigns vehicles to precise arrival slots, and always holds a ready fallback. The concept requires no new crane actuation hardware and no major terminal reconstruction. It is designed to run primarily on existing operational systems and data streams, with limited integration effort where telemetry, dispatch, or edge-computing components are not already available. It is prototypable on a student-scale budget using open tools, and pilot-deployable in a terminal environment where the required operational data streams and interfaces are available. Research on coordinated quay crane and vehicle scheduling confirms that synchronisation alone can raise effective throughput by 8-15% without any mechanical upgrade [3, 5, 10, 11]. Our conservative pilot targets are a 10-20% reduction in crane idle time between moves, a 6-10% reduction in total cycle time, and a 5-8% gain in effective moves per hour. At a representative crane operating cost baseline of approximately EUR 280-300 per hour [15, 16, 17], recovering around 180 idle Operating Hours (GOH) annually per STS unit translates to roughly EUR 50,000 in direct OPEX(Operating Expenditure) savings per crane per year, with a pilot investment of EUR 11,000-25,000 and an indicative payback period of under one year.

Analysis of current STS crane operations

Most people looking at an STS crane for the first time assume that the speed of the hoist or trolley is what primarily limits productivity. This assumption is understandable, but it is not accurate. As summarised in Table 1, an STS crane cycle consists of multiple phases, including hoisting, trolley travel, spreader centering, landing, trolley spotting, and inter-move waiting [9, 12].

The major mechanical phases such as hoisting and trolley travel are already highly optimised in modern cranes through advanced drive systems and anti-sway control. These phases account for a significant portion of the cycle duration, yet they are not the dominant source of inefficiency. The real inefficiencies arise in the transition and coordination phases of the cycle, particularly during spreader centering, landing, and trolley spotting.

More importantly, a substantial portion of lost time occurs outside the physical movement phases, in what the literature refers to as **inter-cycle dwell time**. This is the idle period between two consecutive crane moves that occurs when the crane completes one operation and must wait for the next task to become executable. In real-world operations this gap arises from specific coordination friction: **delayed synchronisation of trucks for twin-lift moves**, manual **pinning station exceptions** at the pinning station, and unexpected **restows** where a container is buried under unplanned cargo.

Although each delay is relatively small, typically in the range of 5 to 15 seconds, its cumulative impact is significant. For a crane performing 20 to 30 moves per hour, even a 10-second average delay per cycle results in over 30 minutes of lost productive time during a 10-hour vessel call [13, 14]. Existing research confirms

that actual throughput deviates from theoretical capacity not due to mechanical limitations but due to coordination gaps between the crane, transport systems, and yard operations [6, 9].

Based on this analysis, the primary focus of this work is not to further optimise crane motion but to **eliminate inter-cycle dwell time** by improving real-time coordination and decision readiness. The objective is to ensure that the next crane move and its required vehicle are always ready before the current move is completed, thereby minimising idle time and enabling continuous crane operation.

Phase	Description	Duration
Hoist (loaded)	Container lifted clear of vessel bay	6-9 s
Trolley travel (outbound)	Trolley traverses ship bay to landside	15-28 s
Fine positioning	Spreader aligned over transport target	5-15 s
Hoist (lower)	Container lowered onto truck or AGV	6-9 s
Twistlock engagement	Lock state confirmed; spreader released	3-8 s
Hoist return (unloaded)	Empty spreader raised	4-7 s
Trolley return (inbound)	Trolley traverses landside back to ship	15-28 s
Re-alignment over bay	Final spreader positioning above next container	4-12 s
Inter-cycle Dwell	Idle gap (Wait for Truck/Task/Confirmation)	5-15 s
Mechanical and handling cycle subtotal (excluding extended operational waiting)	All phases, live operational conditions	63-131 s

Table 1 Typical STS crane cycle phases and indicative durations showing the Coordination Gap

In practice, total observed move duration at the operational level may be higher once queueing, vehicle unavailability, yard-side delay, and other coordination losses are included.

Bottleneck and constraints

A review of terminal operations literature and the operational structure described in the PEMA challenge brief reveals four specific situations that create avoidable inter-cycle dwell time. These are not edge cases; they occur in normal operations many times per vessel call [9, 10].

1. The next vehicle arrives late or in the wrong lane. When a truck or AGV is not positioned at the handoff point by the time the crane finishes its current move, the crane waits. Research on coordinated crane-AGV scheduling demonstrates that even moderate improvements in vehicle timing, arriving 15-20 seconds earlier on average, can raise crane utilisation by 8-12% because the crane never enters an idle state between moves [3, 13]. The problem is not that vehicles are slow; it is that they are called too late or directed to the wrong position.

2. The originally planned next move is no longer the best choice. A Baplie is written before the ship arrives. Once loading or discharge begins, conditions change: a container is not where it was expected, a bay becomes temporarily inaccessible, or a priority shift is called from the yard. If the crane operator is locked into the original sequence, any deviation requires a manual decision under time pressure. If the system had already identified the next best available move, that decision would take zero seconds [10, 11].

3. The operator has to decide without enough information. Studies of crane operator interfaces show that real-time situational awareness on modern cranes is often incomplete; operators know what the current task is, but they do not have a clear, updated view of what comes next and whether everything is ready [4]. This creates a micro-pause at the end of every cycle while the operator checks, confirms, and then acts. Multiplied across hundreds of moves per vessel, these micro-pauses represent a significant source of lost time.

4. When anything goes wrong, recovery is improvised. In current operations, if the planned next move fails due to a truck delay, a blocked task, or a system error, the crane either waits or the operator improvises a workaround. The workaround is usually correct, but it takes time to formulate and creates ripple effects on subsequent moves [14]. A pre-computed fallback move, ready before it is ever needed, eliminates this recovery time entirely.

Each of these four failures is reducible with better information, delivered earlier. None of them requires a new crane or major terminal reconstruction, although some terminals may require additional integration with existing telemetry, dispatch, or operator-support systems.

Existing Technologies

The port equipment industry has invested heavily in crane productivity over the past two decades, and that investment has produced real gains. It is worth being clear about what already works well and where the remaining gap lies, because ZERO-WAIT STS is designed to fill that gap, not to replace what already works.

Mechanical and motion improvements have significantly raised the ceiling on crane speed. Modern anti-sway control systems reduce positioning time, and high-speed trolley and hoist drives have shortened travel durations. These improvements are mature and well-documented [2, 6, 13]. They have made

the crane faster without making the crane more continuously productive, because the idle gaps described above occur outside the motion phases and remain invisible to motion-control systems.

Full terminal automation addresses the coordination problem more radically, replacing human decisions with automated dispatching and using AGVs instead of manually driven trucks. The results at fully automated terminals such as ECT Rotterdam, HHLA Hamburg, and Qingdao are measurably better: throughput per crane-hour is higher and variability is lower [7, 18]. However, full automation requires total terminal reconstruction, capital investment in the hundreds of millions of euros, and years of transition. It is not available to most operating terminals today.

Scheduling algorithms and TOS optimisation sit between these two and are the area most directly relevant to our concept. Comprehensive literature surveys establish that optimised task sequencing can recover 10-15% of throughput currently lost to poor move ordering [9, 10, 11, 14]. Coordinated AGV and crane dispatching research demonstrates 8-12% crane utilisation improvement purely from synchronised timing [3]. These results are primarily achieved through improved software logic and operational integration, without requiring changes to crane actuation hardware. A key gap in current practice is that much of this scheduling logic is generated before execution begins and is not always refreshed dynamically at the move-by-move operational level. When the plan meets reality and deviations occur, the crane is left without real-time guidance [19, 20].

ZERO-WAIT STS closes that gap. It applies the coordination logic that scheduling research has validated, but runs it continuously and in real time, move by move, throughout the entire vessel call.

ZERO-WAIT STS: The Proposed Solution

The core idea is simple: **while the crane is executing the current move, the system is already preparing the next one.** In current operations, pre-staging happens reactively; it begins after the crane finishes. ZERO-WAIT STS shifts that preparation to run in parallel with the active move. By the time the crane is ready, the next task is confirmed, the vehicle is positioned, and a fallback is ready. The system is designed to reduce or eliminate waiting for decisions that could have been made earlier in the cycle.

The system operates through five functions that work together to mitigate "Quayside Chaos":

1. Look-ahead Move Preparation

Look-ahead move preparation is the foundation. As soon as a move begins, the system evaluates the next 2-3 tasks from the vessel plan and checks each for readiness. Crucially, this module is **Load Moment aware**: it recognises that a 40ft High-Cube or a heavy Tank container requires different hoist speeds and "Damping Profile" than a standard 20ft box. It specifically manages **Twin-lift synchronisation**; if the next move involves lifting two 20ft containers, the system triggers a dual-vehicle request immediately. The result is a ranked shortlist of moves that accounts for the physical reality of the cargo, ensuring that the moment the current move ends, the next one is already selected.

How we achieve this: We utilise a real-time data pipeline that scrapes the TOS "Baplie" (Stowage) file and move lists. By pulling read-only data from the Crane PLC or equivalent crane-state interface, the system estimates when the active move has entered its final execution phase.

The System Logic: The “Trigger & Evaluate” Cycle

The system does not merely wait; it constantly compares the Digital Plan against the Physical Reality using a three-step internal loop:

- **The Real-Time Trigger (PLC Polling):** The system listens to the Crane’s PLC via a standard industrial protocol (OPC-UA or Modbus). It monitors the **Z-axis (hoist height)** and Spreader Status. Using a predefined trigger condition, the system triggers the next-move evaluation while the trolley is still in motion.
- **The sub-second Evaluation (“Readiness Check”):** Once triggered, the Decision Engine runs a sanity check. It asks: “Is the target truck’s GPS within 50 metres?” If Item A is not ready, the engine instantly promotes Item B.
- **The Synchronisation (Micro-slotting):** The system sends a dispatch recommendation to the Vehicle Mount Terminal (VMT). The truck is guided into position just as the crane begins the final lowering phase, reducing inter-cycle dwell time.

2. Micro-slot Vehicle Positioning

Micro-slot vehicle positioning replaces the current bulk-call approach with precision timing. Instead of the Terminal Operating System (TOS) issuing a generic command, ZERO-WAIT STS issues a **Targeted Arrival Window**. The system communicates a targeted arrival instruction, for example: “Truck A, proceed to Lane 2, Slot B; target arrival window aligned with the next handoff.” This ensures vehicles do not arrive too early, causing congestion, or too late, causing crane idle time.

How we achieve this: We replace broad dispatching with **Stochastic ETA calculations**. Using GPS/telematics from the truck fleet, the driver receives a countdown or target-arrival cue on the existing VMT, aiming to bring the vehicle into position shortly before the spreader is ready.



Figure 1 Overview of the ZERO-WAIT STS Operational Environment.

3. Dynamic Resequencing

Dynamic resequencing handles the inevitable deviations of port life. If the top-ranked move becomes unavailable due to a **Vessel Dig** or a jammed **Twistlock** at the The Sill Beam area, the system does not pause for manual re-planning. It immediately selects the next best move from the pre-evaluated shortlist. By identifying such ground-blocks in real time, it presents a confirmed recommendation to the operator, reducing delay.

How we achieve this: We implement a **Rolling Horizon Optimiser**. Every 5 seconds, the engine re-checks the “Readiness Score” of candidate moves. If operational status indicates an obstruction, the engine re-ranks the shortlist, promoting a verified fallback move before the crane trolley stops.

4. The Operator Interface

The operator interface ties everything together. The operator sees a single, clean display showing current status, confirmed next move, and vehicle readiness. The display factors in complex variables such as wind speed and weight so the operator does not have to. Research on interface design confirms that giving the operator a clear recommendation rather than raw data to interpret measurably reduces hesitation time and decision errors.

How we achieve this: We deploy a **WebSockets-based dashboard** running on existing hardware. Using “Traffic Light” logic (Green for locked, Yellow for transition, Red for fallback), we filter raw telemetry into actionable recommendations.



Figure 2 The ZERO-WAIT STS Predictive Operator Interface

Interface Functional Overview

As shown in Figure 2, the display is divided into three functional zones:

1. **Current Execution (Left):** Monitors the hoist height and horizontal trolley speed via PLC integration. A dynamic countdown timer allows the operator to pace the move while the system prepares the transition.
2. **Predictive Staging (Center):** Displays a pre-staged task (Move 43) complete with load weight and twin-lift synchronization requirements, removing the cognitive burden of task checking.
3. **Vehicle/Safety Streams (Right):** Identifies the assigned vehicle (ITV-108) and confirms its arrival in the correct lane and slot. A **System Fallback** panel presents an alternative (Move 44) if the primary staged move becomes ground-blocked.

5. Fallback and Safety Design

Fallback and safety design ensures that no software failure can stop the crane. The system operates in three modes:

1. **Coordinated mode:** Full predictive synchronisation and micro-slot assignment.
2. **Recommendation mode:** Guidance provided even when some vehicle or yard inputs are uncertain.
3. **Manual mode:** The system steps back; the terminal operates exactly as it does today.

This degradation path is essential, ensuring that the terminal retains full control and that operational risk remains bounded during deployment.

Technology Stack

The proposed pilot is designed to be student-budget friendly

and scalable, avoiding new infrastructure by utilising existing terminal data streams:

1. **Data integration layer:** Python-based API connectors that aggregate TOS "Move Lists," GPS truck telemetry, and read-only feeds from Crane PLCs.
2. **Optional context feeds:** Yard-slot readiness, local obstruction status, and weather/wind data where available.
3. **Decision engine:** A rules-based engine built on open-source optimisation libraries (e.g., Google OR-Tools) for sub-second re-ranking and fallback selection.
4. **Prediction module:** A lightweight **Kalman Filter** algorithm for short-horizon ETA estimation, calculating high-probability arrival windows for the vehicle fleet.
5. **Operator interface:** A browser-based dashboard (React/Vue.js) that can run on existing cabin workstations or equivalent operator terminals, without requiring specialised display hardware.
6. **Event logging and audit trail:** A local SQLite database providing a "Black Box" log of recovered idle time, operator overrides, and performance metrics for ROI analysis.
7. **Deployment layer:** A containerised (Docker) deployment on an industrial edge computer or equivalent local processing node, deployed on-crane or nearby depending on terminal architecture, to support low-latency processing and high system availability.




Figure 3 Detailed Architectural Diagram of the ZERO-WAIT STS.

Operational flow of ZERO-WAIT STS

Figure 3 illustrates how the coordination layer operates continuously during crane execution. The system receives multiple parallel input streams, including the vessel plan from the TOS, real-time vehicle positions from the dispatch system, and, where available, yard-slot readiness and related operational state information. These inputs are not processed sequentially but combined into a single coordination view.

At the centre, the *coordination engine* evaluates this combined state in real time. While the crane is still executing the current move, the engine performs three parallel actions. First, a short *look-ahead* evaluates the next two to three possible moves and checks their readiness. Second, *micro-slotting* assigns a specific vehicle to a precise lane and time window, ensuring arrival exactly when needed. Third, *dynamic resequencing* selects the best available next move based on current conditions rather than relying strictly on the original plan.

If any uncertainty or disruption occurs, a pre-computed *fallback*

option is always available, allowing immediate recovery without delay. The result of these processes is presented to the operator as a single, clear instruction: the confirmed next move, vehicle readiness status, and fallback if required.

The Feedback Loop and Cycle Repeat: As shown in the architectural diagram, the flow concludes with a critical **feedback loop**. Once the crane executes the next move, the real-time outcome is fed back into the coordination engine. This updates the look-ahead horizon, allowing the system to verify that the executed move matched the plan and immediately begin evaluating the subsequent 2-3 tasks. This continuous loop ensures that decision-making, vehicle positioning, and task validation all happen ahead of time, eliminating the idle gap between crane moves.

Validation Plan and Pilot Evaluation

To verify that ZERO-WAIT STS produces measurable operational value without introducing disruption, the system will be evaluated in three progressive stages.

Stage 1: Offline replay validation. Historical crane-task, vehicle-dispatch, and yard-readiness records, with a focus on **high-complexity vessel calls with frequent twin-lifts and vessel digs**, will be replayed through the coordination engine to test whether the system would have selected earlier-ready moves and generated valid fallback recommendations.

Stage 2: Shadow-mode validation. The system will run in parallel with live operations on one selected STS crane without influencing crane execution or mandatory safety decisions. ZERO-WAIT STS will generate next-move and vehicle-slot recommendations in real time while actual terminal decisions continue as normal. This stage specifically monitors the engine's ability to predict **twistlock-related delays** without operational risk.

Stage 3: Advisory pilot validation. After shadow-mode verification, the system will be deployed in advisory-only mode on a single crane. Operators and, where relevant, dispatch personnel will see the recommended next move, vehicle readiness status, and fallback option, but will retain full authority to accept, override, or ignore the recommendation. No direct crane control will be introduced at any stage.

Assumptions for Pilot Success. The projected performance gains rest on the following operational assumptions:

1. **Data Latency:** Terminal IT infrastructure provides TOS and GPS telemetry updates with a latency of less than 1 second.
2. **Connectivity:** Reliable Wi-Fi or LTE coverage is available under the crane and across the ITV/AGV travel lanes for VMT synchronisation.
3. **Vehicle Availability:** The terminal maintains a vehicle-to-crane ratio sufficient to allow for the "Micro-slot" buffer (typically 3-5 vehicles per crane).
4. **PLC Accessibility:** Read-only access to the Crane PLC (via OPC-UA or Modbus) is granted for Z-axis and Spreader status monitoring.
5. **Operator Buy-in:** Crane operators participate in the Advisory Pilot with a baseline willingness to follow "Green" recommendations when physical conditions allow.

Evaluation metrics. Pilot performance will be assessed using the following measurable indicators:

1. Average inter-move idle time (seconds per move),
2. Synchronisation accuracy for Twin-lift vehicle arrivals,
3. System response time to unplanned Vessel Digs,
4. Vehicle on-time arrival rate to assigned micro-slots,

5. Operator override rate,
6. Percentage of recommendation events handled in recommendation-only and manual fallback modes during the pilot, and coordinated mode in later deployment phases if enabled,
7. Number of fallback activations,
8. Estimated vessel-call time recovered.

Success criteria. The pilot will be considered successful if it demonstrates a measurable reduction in inter-move idle time and cycle time, an increase in effective moves per hour, stable operator acceptance of recommendations, and safe continuity of operation under fallback conditions. In line with the conservative targets of this proposal, the primary success threshold is a 6-10% reduction in average cycle time and a 5-8% increase in effective moves per hour during the advisory pilot.

Operational safeguards. All validation stages will preserve existing terminal operating authority. If live data feeds become unavailable or uncertain, the system will degrade from coordinated mode to recommendation-only mode or manual mode, ensuring that validation cannot interrupt crane operations.

Expected Outcomes

Performance targets for a first deployment are deliberately conservative, grounded in validated literature, and account for the practical realities of human-operated terminals:

1. **Vehicle–crane synchronisation:** Optimised vehicle timing alone raises crane utilisation by 8–12% [10, 13], with recent simulations showing baseline gains from 68% to 78% [3]. ZERO-WAIT STS achieves this via micro-slotting, replacing bulk calls with targeted VMT instructions so vehicles arrive exactly at handoff.
2. **Real-time resequencing:** Adapting sequences to live ground conditions reduces average cycle time by 8–14% compared to static scheduling [2, 11, 20]. Our engine evaluates candidate moves in under 200 ms, instantly promoting a fallback during trolley transit if a vessel dig or twistlock jam occurs.
3. **Operator interface improvements:** Contextualised next-step recommendations reduce average decision latency by 4–7 seconds per cycle [4]. PLC-triggered monitoring locks in the confirmed next move early, removing the cognitive burden of manual readiness checks.

Combined Impact: Taking these effects together, we target a **10–20% reduction in inter-move idle time** and a **10–15% reduction in vehicle wait time** at the micro-slots. This produces a 6–10% reduction in average cycle time and a 5–8% increase in effective moves per hour. For an STS crane averaging 25 moves per hour, this equates to 13–20 additional moves, recovering **45 to 70 minutes of productive crane time per 10-hour vessel call**.

Measure	Target	Basis
Avg. cycle time	6–10% Reduction	[2, 3, 11]
Moves per hour	5–8% Increase	[3, 10]
Inter-move idle	10–20% Reduction	[2, 4, 13]
Vehicle wait (Micro-slot)	10–15% Reduction	[3, 13]

Table 2 System Performance Targets and Empirical Basis

The projected operational improvements, including cycle time and idle reduction, are summarized in **Table 2**.

Business Impact & Economic Value

ZERO-WAIT STS converts incremental coordination improvements into quantifiable financial gains. By eliminating 5-15 seconds of idle time per move, the system recovers crane capacity that is currently lost to avoidable waiting, without any new capital equipment.

The Math of Recovery (Per Crane)

The business case is built on conservative, referenced industry baselines.

Using representative peer-reviewed cost baselines for terminal operations and container-handling economics, this proposal assumes an STS operating cost of approximately **EUR 280-300 per crane-hour** for business-case estimation [15, 16, 17]. This figure is used here as a conservative planning baseline and should be recalibrated against local terminal labour, energy, and utilisation conditions during pilot implementation [15, 17].

Applying the lower bound of our pilot target to a representative operating profile of 250 vessel calls per year, the model suggests recovery on the order of **165-200 Operating Hours (GOH) annually per STS unit**. This estimate is directly derived from our cycle-time analysis: 10 seconds saved per move at 25 moves per hour yields approximately 4 minutes recovered per crane-hour; across a 10-hour vessel call and 250 vessel calls per year, this accumulates to roughly 167 hours at minimum [2, 3, 13]. Taking EUR 280 as the lower-bound operating cost, annual OPEX savings per crane are approximately **EUR 47,000-56,000**. We use **EUR 50,000** as the central conservative estimate.

A new STS ship-to-shore crane represents a multi-million-euro capital investment, so even modest throughput recovery from software coordination may defer or better utilise future capital expenditure [15, 16]. A 5% Intensity Gain across a fleet of ten cranes would translate into a substantial increase in annual move capacity, effectively providing the equivalent of additional crane time without the same capital outlay [3, 10].

ROI and Scalability

At a conservative annual gain of around EUR 50,000 and a maximum pilot cost of EUR 25,000, the indicative payback period is within the first year and may approach six months under favourable operating conditions. Once the integration blueprint is established on the first crane, subsequent cranes require only configuration, substantially reducing marginal deployment cost.

Metric (Per Crane)	Description	Annual Impact
Maint. & Support	Annual software/integration upkeep	(EUR 4k–10k)
Recovered Hours	165–200 GOH @ EUR 280/hr	EUR 50,000
OPEX Efficiency	Direct cost recovery per unit	EUR 47k–56k
Net Annual Gain	Post-maintenance recurring value	+ EUR 40k–46k

Table 3 Annual Recurring Costs vs. Expected Gains (Conservative Lower Bound)

As shown in **Table 3**, the ZERO-WAIT STS system generates

significant value by recovering lost operational capacity. Based on the conservative cost baselines and recovered hours the project achieves a net positive impact exceeding EUR 40,000 per crane annually.

Conclusion: ZERO-WAIT STS is not merely a technical upgrade; it is a productivity-focused operational investment. It allows terminals to extract more value from existing equipment, improving Vessel Turnaround Time (VTT) and deferring the need for multi-million euro hardware expansions [16]. Importantly, these figures represent a deliberately conservative lower bound: terminals with higher vessel call frequency, larger crane fleets, or more complex cargo mixes should expect proportionally greater returns [9, 14].

Safety and Sustainability

Safety. ZERO-WAIT STS is designed not to command the crane directly and to preserve operator authority over execution. It makes recommendations that the operator accepts, modifies, or ignores. Every move still requires execution under the authority of certified operational personnel and within existing terminal safety protocols. The three-mode fallback ensures the system can be overridden or switched off at any moment without disrupting crane operation. In practice, better preparation should improve safety: when the operator knows what comes next before it happens, there are fewer rushed decisions, fewer unexpected vehicle positions under the crane, and fewer moments of mismatch between plan and reality that lead to fast, high-pressure improvisation.

Sustainability. Container terminals are among the most energy-intensive logistics facilities in the world, and their transition to electrified equipment is accelerating under pressure from both regulation and rising fossil fuel costs. A systematic review of port energy efficiency strategies identifies operational scheduling, including the timing and sequencing of crane and vehicle activity, as one of the highest-leverage levers for reducing terminal energy consumption, because it lowers peak demand without requiring new hardware investment [7]. ZERO-WAIT STS contributes to this directly: when vehicles arrive at precisely the right moment and cranes operate with fewer idle intervals, the sharp simultaneous power spikes that occur when multiple cranes hoist concurrently are naturally smoothed. This matters because peak electricity tariffs in large industrial ports are often structured around instantaneous demand peaks rather than total energy consumption, meaning that even modest improvements in scheduling can produce disproportionately large reductions in electricity costs [8]. As terminals add electric AGVs, shore power connections, and renewable generation to their energy mix, the coordination logic underpinning ZERO-WAIT STS becomes directly relevant to energy management as well as crane productivity [15]. A terminal that already operates with low inter-move idle time and precise vehicle timing is also a terminal whose power demand is smoother, more predictable, and easier to manage against a grid capacity constraint.

Operational Impact

The port industry is moving toward greater automation across all operational tiers, with automated and semi-automated terminals proliferating at an accelerating pace [18, 20]. Successful transitions consistently depend on three prerequisites: advance knowledge of the next task before the current one ends, a defined response when the plan fails, and clean switching between advisory, semi-automated, and manual control states. ZERO-WAIT STS establishes all three from day one. A terminal deploying it today recovers idle time immediately while building

the data infrastructure and operator familiarity that any subsequent automation upgrade will require, making it one of the few port technology investments that delivers value at the point of deployment and remains relevant as the terminal evolves [19, 14].

Conclusion

Somewhere on every quay, right now, a crane is waiting. Not because it is broken, not because the operator is inattentive, but because the next truck has not arrived, or the next task has not been confirmed, or no one has yet decided what the best next move is. That wait happens dozens of times every hour, on every crane, at every terminal in the world. It is so routine that it has become invisible.

ZERO-WAIT STS makes it visible and then eliminates it. By preparing the next move while the current one is still running, the system removes the reactive gap that sits at the heart of lost crane productivity. It does this without changing crane actuation hardware and without requiring full terminal reconstruction, while fitting within existing operational safety governance. It does this with software that a student team can build and demonstrate using open tools, while real terminal pilot validation would require access to operational data and interfaces.

The performance numbers are grounded in published research. The pilot cost is modest. The path to full deployment is short and the return is swift. Once proven at one crane, the concept becomes the coordination foundation that every future automation upgrade will depend on.

The crane should never have to wait for the next move. ZERO-WAIT STS ensures it does not.

References

- [1] R. Tan et al., "The bay-based quay crane scheduling problem considering time-varying handling capacities in automated container terminals," *Transportation Research Part E: Logistics and Transportation Review*, vol. 201, p. 104254, 2025, doi: <https://doi.org/10.1016/j.tre.2025.104254>.
- [2] Y. Hu et al., "Coordinated scheduling optimization of quay cranes and AGVs in automated container terminals," *Computers & Operations Research*, vol. 182, p. 107147, 2025, doi: <https://doi.org/10.1016/j.cor.2025.107147>.
- [3] J. Goh, Y. Fang, and B. Ens, "Embedded visualizations in crane operation user interfaces for real-time assistance," *Automation in Construction*, vol. 173, p. 106078, 2025, doi: <https://doi.org/10.1016/j.autcon.2025.106078>.
- [4] X. Li, L. Wu, H. Lang, Z. Jin, and P. He, "A critical literature review on layout designs and handling technology in traditional and automated container terminals," *Transport Economics and Management*, vol. 3, pp. 269–280, 2025, doi: <https://doi.org/10.1016/j.team.2025.06.002>.
- [5] N. Makhado, T. Paepae, M. Sejeso, and C. Harley, "Berth Allocation and Quay Crane Scheduling in Port Operations: A Systematic Review," *Journal of Marine Science and Engineering*, vol. 13, no. 7, 2025, doi: <https://doi.org/10.3390/jmse13071339>.
- [6] H. J. Carlo, I. F. A. Vis, and K. J. Roodbergen, "Transport operations in container terminals: Literature overview, trends, research directions and classification scheme," *European Journal of Operational Research*, vol. 236, no. 1, pp. 1–13, 2014, doi: <https://doi.org/10.1016/j.ejor.2013.11.023>.
- [7] Ç. Iris and J. S. L. Lam, "A review of energy efficiency in ports: Operational strategies, technologies and energy management systems," *Renewable and Sustainable Energy Reviews*, vol. 112, pp. 170–182, 2019, doi: <https://doi.org/10.1016/j.rser.2019.04.069>.
- [8] F. D. Kanellos, E.-S. M. Volanis, and N. D. Hatzigiorgiou, "Power Management Method for Large Ports With Multi-Agent Systems," *IEEE Transactions on Smart Grid*, vol. 10, no. 2, pp. 1259–1268, 2019, doi: <https://doi.org/10.1109/TSG.2017.2762001>.

- [9] D. Kizilay and D. T. Eliyi, "A comprehensive review of quay crane scheduling, yard operations and integrations thereof in container terminals," *Flexible Services and Manufacturing Journal*, vol. 33, no. 1, pp. 1–42, Mar. 2021, doi: <https://doi.org/10.1007/s10696-020-09385-5>.
- [10] K. H. Kim and J. W. Bae, "A Look-Ahead Dispatching Method for Automated Guided Vehicles in Automated Port Container Terminals," *Transportation Science*, vol. 38, no. 2, pp. 224–234, May 2004, doi: <https://doi.org/10.1287/trsc.1030.0082>.
- [11] L. Zhen, D. Zhuge, S. Wang, and K. Wang, "Integrated berth and yard space allocation under uncertainty," *Transportation Research Part B: Methodological*, vol. 162, pp. 1–27, 2022, doi: <https://doi.org/10.1016/j.trb.2022.05.011>.
- [12] S.-L. Chao and Y.-J. Lin, "Evaluating advanced quay cranes in container terminals," *Transportation Research Part E: Logistics and Transportation Review*, vol. 47, no. 4, pp. 432–445, Jul. 2011, doi: <https://doi.org/10.1016/j.tre.2010.12.002>.
- [13] L. Chen, A. Langevin, and Z. Lu, "Integrated scheduling of crane handling and truck transportation in a maritime container terminal," *European Journal of Operational Research*, vol. 225, no. 1, pp. 142–152, Feb. 2013, doi: <https://doi.org/10.1016/j.ejor.2012.09.019>.
- [14] C. Bierwirth and F. Meisel, "A follow-up survey of berth allocation and quay crane scheduling problems in container terminals," *European Journal of Operational Research*, vol. 244, no. 3, pp. 675–689, Aug. 2015, doi: <https://doi.org/10.1016/j.ejor.2014.12.030>.
- [15] N. K. Park and Y. An, "Financial Analysis of Automated Container Terminal Capacity from the Perspective of Terminal Operating Company," *JMSE*, vol. 8, no. 11, p. 954, Nov. 2020, doi: <https://doi.org/10.3390/jmse8110954>.
- [16] B. Dragović, A. Dragović, and M. A. Dulebenets, "The quay crane operation problem at marine container terminals: bibliometric analysis, emerging trends, and future research opportunities," *Transportation Research Part E: Logistics and Transportation Review*, vol. 201, p. 104266, Sep. 2025, doi: <https://doi.org/10.1016/j.tre.2025.104266>.
- [17] Y. Ding, K. Chen, X. Wei, and Y. Yang, "A novel cost-management system for container terminals using a time-driven Activity-Based Costing approach," *Ocean & Coastal Management*, vol. 217, p. 106011, Feb. 2022, doi: <https://doi.org/10.1016/j.ocecoaman.2021.106011>.
- [18] B. Kim, G. Kim, and M. Kang, "Study on Comparing the Performance of Fully Automated Container Terminals during the COVID-19 Pandemic," *Sustainability*, vol. 14, no. 15, p. 9415, Aug. 2022, doi: <https://doi.org/10.3390/su14159415>.
- [19] R. Choe, H. Cho, T. Park, and K. R. Ryu, "Queue-based local scheduling and global coordination for real-time operation control in a container terminal," *J Intell Manuf*, vol. 23, no. 6, pp. 2179–2192, Dec. 2012, doi: <https://doi.org/10.1007/s10845-011-0564-y>.
- [20] W. K. Kon, N. S. F. Abdul Rahman, R. Md Hanafiah, and S. Abdul Hamid, "The global trends of automated container terminal: a systematic literature review," *MABR*, vol. 6, no. 3, pp. 206–233, Sep. 2021, doi: <https://doi.org/10.1108/MABR-03-2020-0016>.