International Journal of Marine Engineering Innovation and Research, Vol. 11(3), Sept. 2026. 915-924
(pISSN: 2541-5972, eISSN: 2548-1479)

Linear Programming-Based Maritime Speed Allocation and Fuel-Cost Efficiency: An Intelligent Decision Support System with an LLM Explanation Layer

Moejiono Moejiono[1*], Achmad D. Fachrudin[2], Diana Alia[3], Novitasari Novitasari[4], Abadi Abadi[5], Damoyanto Purba[6]

(Received: 19 June 2026 / Revised: 21 June 2026 / Accepted: 28 June 2026 / Available Online: 1 September 2026)

***Abstract*—** This applied mathematical study develops a linear-programming model for maritime speed allocation to minimize estimated voyage operating cost under route-distance and sailing-deadline constraints. The decision variables represent the distance assigned to each predefined speed mode, enabling the model to determine a continuous mixed-speed allocation that remains feasible within the permitted sailing time. The model was implemented in a web-based Intelligent Decision Support System (IDSS) using a HiGHS-based linear-programming solver and demonstrated on the Tanjung Priok–Tanjung Perak route using an application-generated distance of 469.84 nautical miles. Under a baseline VLSFO price of USD 755/MT and a 35-hour sailing deadline, the model allocated 390.69 NM to the 13-knot mode and 79.15 NM to the 16-knot mode, producing an estimated minimum operating cost of USD 34,748.16 and total fuel consumption of 38.30 MT. Sensitivity analysis shows that tighter deadlines shift the optimal allocation toward higher-speed modes, increase estimated voyage cost, and eventually produce infeasibility when the required travel time falls below the capability of the fastest available speed mode. At the fixed 35-hour deadline, fuel-price variation changes the objective-function value without altering the selected speed allocation in the tested scenarios. As a supplementary interpretive component, an LLM-based layer transforms structured optimization outputs into a natural-language Voyage Master Plan while preserving the ship master's final authority. The study therefore contributes a traceable mathematical optimization model and a functionally integrated explanation interface. Because the analysis uses model-input parameters rather than empirically calibrated vessel records and does not include user-based evaluation, further validation is required before operational deployment.

Keywords— linear programming; maritime speed allocation; fuel-cost efficiency; sensitivity analysis; intelligent decision support system; LLM explanation layer.

Corresponding Author: moejiono@poltekpel-sby.ac.id

I. INTRODUCTION

Maritime transportation is required to maintain service reliability while facing continuing pressure to improve energy efficiency and control operating costs. Within voyage operations, fuel consumption is particularly sensitive because decisions concerning sailing speed directly influence voyage expenditure and environmental performance. This issue is especially relevant in archipelagic shipping networks, where domestic routes must accommodate sailing distance, arrival-time commitments, and operational uncertainty simultaneously.

A central operational decision in voyage planning is the selection of an appropriate speed profile. When a vessel experiences delay or is assigned a restrictive estimated time of arrival (ETA), increasing speed may appear to be an immediate response. However, the mathematical relationship between vessel speed and fuel consumption is not proportional. Fagerholt et al. [1] represented fuel consumption and emissions on shipping routes through an approximately cubic function of sailing speed. Psaraftis and Kontovas [2] likewise identified speed selection as a key mathematical decision variable in energy-efficient maritime transportation, particularly when schedule constraints and fuel-price conditions are considered simultaneously. This relationship indicates that voyage planning requires an optimization framework rather than an intuitive preference for higher speed.

From an applied mathematics perspective, this problem can be formulated as a constrained optimization problem. For a route with known total distance and a set of available sailing-speed modes, the distance assigned to each speed mode can be treated as a decision variable. The objective is to minimize estimated voyage operating cost, while the total allocated distance must equal the required voyage distance and the accumulated sailing time must remain within the permitted deadline. This formulation yields a linear-programming model when the cost and time contributions of each predefined speed mode are expressed per nautical mile. Linear programming is therefore not merely a computational tool in this study, but the principal analytical framework for characterizing feasible solutions, identifying the

Moejiono Moejiono, Politeknik Pelayaran Surabaya, Surabaya, 60155, E-mail: moejiono@poltekpel-sby.ac.id
Achmad D. Fachrudin, Politeknik Pelayaran Surabaya, Surabaya, 60155, E-mail: achmad.dhany@poltekpel-sby.ac.id
Diana Alia, Politeknik Pelayaran Surabaya, Surabaya, 60155, E-mail: diana.alia@poltekpel-sby.ac.id
Novitasari Novitasari, Politeknik Pelayaran Surabaya, Surabaya, 60155, E-mail: novitasari@poltekpel-sby.ac.id
Abadi Abadi, Universitas Negeri Surabaya, Surabaya, 60231, E-mail: abadi@unesa.ac.id
Damoyanto Purba, Politeknik Pelayaran Surabaya, Surabaya, 60155, E-mail: damoyanto.purba@poltekpel-sby.ac.id

minimum-cost allocation, and examining changes in the solution under altered parameter values.

Previous maritime optimization studies provide a strong foundation for this modelling approach. Fagerholt et al. [1] demonstrated how speed optimization can reduce fuel consumption and emissions along shipping routes, while Psaraftis and Kontovas [2] provided a systematic account of speed models for energy-efficient maritime transportation. In the Indonesian maritime-logistics context, Sirait et al. [3] demonstrated the applicability of mixed-integer linear programming for ship allocation and routing decisions involving cost minimization in multi-product oil distribution. Collectively, these studies confirm the importance of mathematical optimization in maritime planning. Nevertheless, most such studies emphasize the numerical optimum and operational implications of the model, whereas fewer studies address how an optimization result can be communicated as a structured and readily reviewable recommendation for operational users.

This interpretive issue does not diminish the centrality of the mathematical model; rather, it concerns the communication of its outputs. A cost-minimizing allocation may be mathematically valid and computationally traceable, yet the resulting numerical values still require interpretation when presented in a decision-support environment. Recent work on user-centered explainable artificial intelligence (xAI) for maritime decision support emphasizes transparency and interpretability in human–machine interaction within safety-critical maritime environments [4]. In parallel, Miller et al. [5] identify potential roles for Large Language Models (LLMs) in maritime safety and decision assistance. These developments provide a rationale for investigating an LLM as a supporting explanation interface, provided that the mathematical optimization remains the sole mechanism for deriving the solution and that professional navigational authority is not delegated to generative AI.

Against this background, the present study is positioned primarily within applied mathematical modelling and operations research, with an additional decision-support interface component. The study addresses an operational interpretation gap by developing an Intelligent Decision Support System (IDSS) in which a HiGHS-based linear-programming engine generates a traceable minimum-cost distance allocation across predefined speed modes under route-distance and sailing-time constraints. The resulting allocation is then interpreted through sensitivity analysis to reveal how the optimal solution responds to changes in deadline tightness and fuel-price parameters. Only after the mathematical solution has been obtained is an LLM employed as a supplementary explanation layer to translate structured optimization outputs into a natural-language Voyage Master Plan. Accordingly, the LLM does not determine the optimum, alter the feasible region, or replace the decision authority of the ship master.

The novelty of this study lies in two connected contributions. The primary mathematical contribution is the formulation and implementation of a transparent linear-programming model for continuous mixed-speed distance allocation, accompanied by sensitivity analysis of operational and economic parameters in a domestic maritime-route case. The supplementary implementation contribution is the integration of the optimized numerical output with an LLM-based explanation layer, allowing the mathematically generated recommendation to be communicated in an operationally readable form without transferring the optimization function to AI. In this sense, the study complements existing maritime optimization literature by retaining the mathematical model as the analytical core while extending the presentation of its results within a human-readable decision-support prototype.

Accordingly, this study has three objectives. First, it formulates a linear-programming model for allocating voyage distance across predefined speed modes while minimizing estimated operating cost under total-distance and sailing-deadline constraints. Second, it analyzes changes in the optimal solution and objective-function value under different deadline and fuel-price scenarios through sensitivity analysis. Third, it demonstrates how the resulting numerical solution can be communicated through an LLM-based explanation layer in a web-based IDSS prototype using a domestic Indonesian shipping-route case. Because the present study is limited to mathematical computation, sensitivity analysis, and functional prototype demonstration, empirical evaluation of officer workload, trust, usability, and safety outcomes is positioned as necessary future work.

II. METHOD

This study employed a computational applied mathematical modelling approach grounded in Operations Research. Its principal analytical component was a continuous linear-programming model developed to determine the minimum-cost allocation of sailing distance across predefined vessel-speed modes under route-distance and sailing-deadline constraints. The mathematical model was subsequently implemented in a web-based Intelligent Decision Support System (IDSS). The system implementation served to demonstrate the computational solution and to provide an LLM-based explanation of the optimized output; it did not replace or modify the mathematical optimization process. The research procedure comprised five stages: (1) specification of route, vessel, cost, and deadline parameters; (2) formulation of the linear-programming model; (3) numerical solution of the model using a HiGHS-based solver; (4) sensitivity analysis of sailing-deadline and fuel-price parameters; and (5) functional demonstration of an LLM explanation layer that communicates the mathematically derived solution in a natural-language voyage plan.

A. Mathematical Model Formulation

Let $i = 1, 2, \dots, n$ denote the available sailing-speed modes. For each mode i, the parameter v_i represents sailing speed in knots, and q_i represents fuel consumption in metric tons per day (MT/day). The parameter Pdenotes fuel price in USD/MT, while Fdenotes the vessel's daily fixed operating cost in

USD/day. The total required route distance is denoted by D in nautical miles (NM). The total voyage deadline is denoted by H in hours, and Sdenotes the specified total port-stay duration in days. Accordingly, the maximum time available for sailing is defined as:

$$T_{max} = H - (S \times 24) \quad (1)$$

The continuous decision variable x_i represents the sailing distance, in NM, allocated to speed mode i. For each speed mode, the estimated sailing operating cost per nautical mile is derived from its fuel-cost contribution and fixed daily operating-cost contribution as follows:

$$c_i = \frac{(q_i \cdot P) + F}{24 \cdot v_i} \quad (2)$$

The fixed operating cost associated with port stay is defined as:

$$C_{port} = S \cdot F \quad (3)$$

The linear-programming objective function minimizes estimated total voyage operating cost:

$$\text{Minimize } Z = \sum_{i=1}^{n} (c_i \cdot x_i) + C_{port} \quad (4)$$

The model is subject to three conditions. First, the allocated distances across all speed modes must equal the required voyage distance:

$$\sum_{i=1}^{n} x_i = D \quad (5)$$

Second, the accumulated sailing time must not exceed the available sailing time after accounting for port stay:

$$\sum_{i=1}^{n} \frac{x_i}{v_i} \leq T_{\max} \quad (6)$$

Third, each distance allocation is continuous and non-negative:

$$x_i \geq 0 \forall i \in \{1,2, \dots, n\} \quad (7)$$

The resulting solution is therefore interpreted as a continuous mixed-speed distance allocation, rather than a mixed-integer solution. In addition to the optimal objective-function value, total fuel consumption is calculated from the optimized allocation as:

$$\text{Total Fuel (MT)} = \sum_{i=1}^{n} \left(\frac{x_i}{v_i} \cdot \frac{q_i}{24} \right) \quad (8)$$

while the optimized total voyage duration is calculated as:

$$\text{Total Duration (hours)} = \sum_{i=1}^{n} \frac{x_i}{v_i} + (S \times 24) \quad (9)$$

This formulation positions the study within applied mathematical modelling: the route, speed-mode, cost, and time parameters define the feasible region, while the objective function determines the minimum-cost allocation within that region.

B. *IDSS Computational Implementation*

The linear-programming formulation was implemented in a web-based IDSS developed in Python using the FastAPI framework [6]. Consistent with the analytical focus of this study, the application architecture was organized into three functional components: route-distance pre-processing, mathematical optimization, and natural-language explanation of the computed solution.

C. *Route-Distance Pre-processing*

The route-distance component received two or more port names or geographic coordinate points as input. Port coordinates were retrieved from an embedded Indonesian port database. For selected ports, predefined offshore-coordinate adjustments were applied before sea-route computation in order to avoid routing directly from inland or terminal coordinate points. Maritime-route distance was then estimated using the Python searoute routing module [7]. If a sea-route solution could not be generated for a segment, the application applied a fallback estimate based on the Haversine great-circle distance multiplied by a factor of 1.3. Where an offshore adjustment was applied, the corresponding distance between the original port point and the adjusted offshore point was additionally included through the Haversine calculation. Accordingly, the computed distance is treated as a route-distance input generated by the prototype for optimization, rather than as a certified navigational passage plan.

D. *Mathematical Optimization Engine*

The optimization model was solved using the linprog function in the SciPy optimization library with method='highs' [8]. In the implemented configuration, HiGHS automatically selects an appropriate linear-programming solver; accordingly, the numerical procedure is referred to in this study as HiGHS-based linear-programming optimization rather than as a specific simplex implementation. The coefficient vector supplied to the solver contained the calculated operating cost per NM for each available speed mode. The equality-constraint matrix represented the total-distance condition, while the inequality-constraint matrix represented the maximum sailing-time condition. The lower bound of every decision variable was set to zero, with no integer or binary restriction imposed. A scenario was classified as feasible when the solver returned an optimal solution satisfying the specified constraints. When no combination of available speed modes could satisfy the sailing deadline, the scenario was classified as infeasible.

E. LLM Explanation Layer

After completion of the mathematical optimization, the system constructed a structured output containing the route description, total computed distance, optimized distance assigned to each speed mode, sailing duration, estimated voyage cost, total fuel consumption, and port-stay information. This structured optimization output was transmitted to the Google Gemini generative-AI service [9] through a predefined prompt requesting a natural-language Voyage Master Plan. The generated narrative was structured to include an executive summary, speed and bunkering strategy, cost information, and nautical cautions relevant to Indonesian waters, together with a statement preserving the ship master's final authority.

The LLM layer was not involved in defining the feasible region, computing the objective-function minimum, or modifying the optimal allocation. Its role was limited to communicating the numerical solution in a readable operational format. Consequently, the IDSS demonstration evaluates functional integration between an optimization model and an explanation interface; it does not constitute empirical validation of explanation accuracy, user trust, cognitive-workload reduction, or navigational safety improvement.

F. Case Demonstration and Sensitivity-Analysis Design

The computational case examined the domestic maritime route between Tanjung Priok (Jakarta) and Tanjung Perak (Surabaya). To maintain consistency between the mathematical experiment and the functional system demonstration, all reported numerical analyses used the route distance generated by the finalized IDSS route-distance component. For the specified route, the application produced a total sailing distance of 469.84 NM. This distance was treated as the input distance for the optimization model and for all sensitivity analyses reported in this study. As described previously, the route-distance output is a prototype-generated computational input and is not intended to constitute a certified navigational passage plan.

The vessel profile consisted of three predefined sailing-speed modes representing a medium-sized feeder-vessel scenario: Slow Steaming at 10 knots with fuel consumption of 15 MT/day, Normal Speed at 13 knots with fuel consumption of 24 MT/day, and Service Speed at 16 knots with fuel consumption of 40 MT/day. The daily fixed operating cost was specified as USD 4,000/day. A baseline VLSFO price of USD 755/MT, based on the bunker-price source selected for the computational experiment, was used for the principal analysis [10]. The final manuscript should state the precise observation or retrieval date associated with this market-price input.

For the reported numerical experiments, the time constraint was defined as an available sailing deadline, and port-stay duration was set to zero days. This setting isolates the mathematical effect of sailing-time restrictions on optimal speed allocation. Although the developed application retains a port-stay input field for broader functional use, port-stay cost and time were not activated in the reported sensitivity analyses or in the aligned case-demonstration output.

Two sensitivity analyses were undertaken. First, in the sailing-deadline analysis, the fuel-price parameter and vessel profile were held constant while the maximum permitted sailing time was varied from loose to restrictive conditions. The observed outputs were the optimal objective-function value, the optimized speed-mode distance allocation, the time-constraint status, and the feasibility status of the model. This procedure identifies the deadline levels at which the time constraint becomes binding and the condition under which no available speed-mode combination can complete the voyage within the permitted time.

Second, in the fuel-price sensitivity analysis, the sailing deadline was fixed at 35 hours while the VLSFO price parameter was varied across selected price scenarios. The observed outputs were the optimal objective-function value and any change in the selected distance allocation. This procedure evaluates whether fuel-price variation changes only the cost magnitude or also changes the optimal speed-mode composition within the examined range.

G. Scope and Methodological Boundaries

The study evaluates a mathematical optimization model and its implementation as a functional IDSS prototype. Several methodological boundaries therefore apply. First, the fuel-consumption rates associated with the available speed modes are treated as predefined model inputs rather than parameters calibrated from vessel-specific operational records. Second, the 469.84-NM route distance is generated through the prototype route-distance component and serves as an optimization input; it is not a certified navigational route or an approved passage plan. Third, the numerical experiments isolate sailing-time and fuel-price effects by setting port-stay duration to zero, although the application architecture permits port-stay inputs for future operational scenarios. Fourth, the LLM-generated Voyage Master Plan is an explanatory output derived only after the LP solution has been obtained; it does not alter the feasible region, calculate the optimum, or provide autonomous navigational instruction. Finally, no bridge-officer usability test or expert safety-validation procedure was performed. Claims concerning workload reduction, trust, operational appropriateness, and navigational safety therefore remain outside the scope of the reported findings.

III. RESULTS AND DISCUSSION

This section presents the numerical outcomes of the linear-programming model, followed by their mathematical interpretation and the functional demonstration of the IDSS prototype. Consistent with the applied-mathematics orientation of this study, the analysis centers on the optimal objective-function value, the resulting continuous mixed-speed distance allocation, changes in the solution under varied parameter values, and the feasibility status of the model. The LLM-based output is subsequently discussed as an explanation layer applied only after the mathematical optimum has been computed.

All numerical results reported in this section use a single aligned case setting: the Tanjung Priok–Tanjung Perak route with an application-generated total distance of 469.84 NM. Accordingly, the sensitivity analyses, the dashboard demonstration, and the LLM-generated Voyage Master Plan refer to the same route-distance basis. This unified configuration improves traceability between the mathematical formulation, numerical results, and functional prototype output.

A. Baseline Parameters and Cost Coefficients

The aligned case-study scenario used the application-generated Tanjung Priok–Tanjung Perak route distance of 469.84 NM. The vessel profile consisted of three predefined speed modes: Slow Steaming at 10 knots with fuel consumption of 15 MT/day, Normal Speed at 13 knots with fuel consumption of 24 MT/day, and Service Speed at 16 knots with fuel consumption of 40 MT/day. The daily fixed operating cost was set at USD 4,000/day, and the baseline VLSFO price was specified as USD 755/MT based on the bunker-price source selected for the computational experiment. Because fuel prices vary over time, the final manuscript should specify the date on which this price observation was retrieved.

Using the cost-coefficient formulation presented in the method section, the operating-cost coefficients under the baseline fuel-price scenario were obtained as shown in Table 1. The coefficient values remain independent of the total route distance because each coefficient represents estimated operating cost per nautical mile for a specific speed mode.

TABLE 1.
BASELINE PARAMETERS AND DERIVED SAILING-COST COEFFICIENTS

Speed Mode	Speed (knots)	Fuel Consumption (MT/day)	Derived Cost Coefficient (USD/NM)
Slow Steaming	10	15	63.8542
Normal Speed	13	24	70.8974
Service Speed	16	40	89.0625

The coefficient values show that Slow Steaming is the least-cost mode per nautical mile under the baseline parameters. However, completing the 469.84-NM route entirely at 10 knots requires approximately 46.984 hours. Thus, Slow Steaming can remain the sole selected mode only when the available sailing deadline is at least 46.984 hours. For tighter deadlines, the model must allocate part or all of the route to higher-speed modes despite their higher cost coefficients.

B. Sensitivity of the Optimal Solution to Sailing-Deadline Constraints

The first sensitivity analysis examined changes in the optimal solution when the maximum permitted sailing time was progressively tightened while maintaining the baseline VLSFO price at USD 755/MT and retaining all vessel parameters. The deadline scenarios were selected to represent a loose deadline, the all-Slow-Steaming feasibility boundary, intermediate binding conditions, the fastest-mode feasibility boundary, and an infeasible condition. The results are presented in Table 2.

TABLE 2.
DEADLINE SENSITIVITY ANALYSIS FOR THE ALIGNED 469.84-NM ROUTE SCENARIO

Scenario ID	**Sailing Deadline (hours)**	**Slow Steaming, 10 kn (NM)**	**Normal Speed, 13 kn (NM)**	**Service Speed, 16 kn (NM)**	**Optimal Cost, USD**	**Time Slack (hours)**	**Status**
S-Time-01	55.000	469.84	0.00	0.00	30,001.24	8.016	Optimal
S-Time-02	46.984	469.84	0.00	0.00	30,001.24	0.000	Optimal; binding
S-Time-03	41.000	210.53	259.31	0.00	31,827.61	0.000	Optimal; binding
S-Time-04	35.000	0.00	390.69	79.15	34,748.16	0.000	Optimal; binding
S-Time-05	29.365	0.00	0.00	469.84	41,845.13	0.000	Optimal; binding
S-Time-06	28.000	–	–	–	–	–	Infeasible

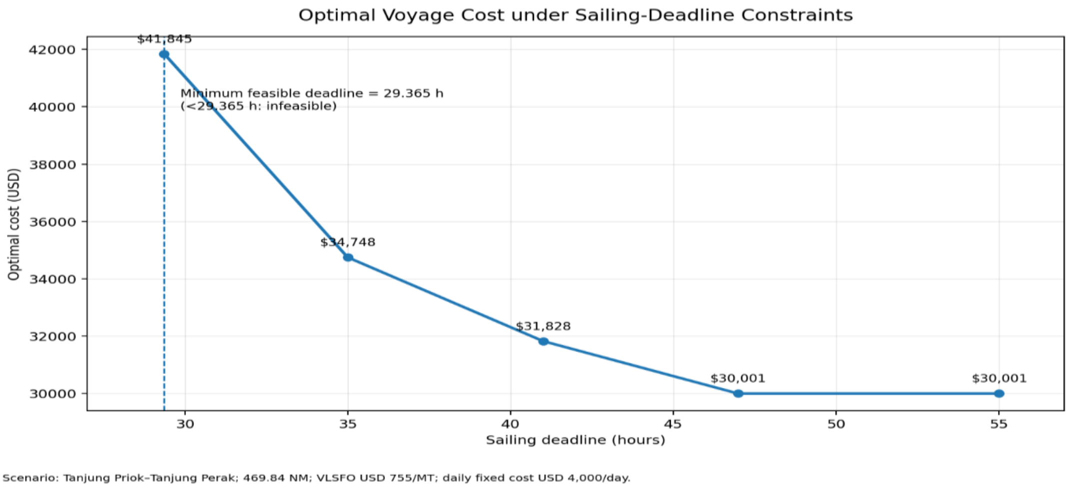

Figure 1. Piecewise-linear response of the optimal voyage cost to the sailing-deadline constraint for the aligned Tanjung Priok–Tanjung Perak route scenario of 469.84 NM.

The results demonstrate that the sailing-deadline constraint directly governs the composition of the optimal distance allocation. Under a 55-hour sailing deadline, the model assigns the entire route to Slow Steaming, with 8.016 hours of unused time. At a deadline of 46.984 hours, the same least-cost allocation remains optimal, but the time constraint becomes binding because the full 469.84-NM route completed at 10 knots requires exactly 46.984 hours.

When the permitted sailing time is reduced to 41 hours, an all-Slow-Steaming allocation becomes infeasible. The minimum-cost feasible solution therefore shifts to a continuous combination of 210.53 NM at 10 knots and 259.31 NM at 13 knots, producing an estimated operating cost of USD 31,827.61. At the focal 35-hour deadline, an all-Normal-Speed solution is also infeasible because completing 469.84 NM at 13 knots would require approximately 36.142 hours. Consequently, the optimal allocation consists of 390.69 NM at 13 knots and 79.15 NM at 16 knots, resulting in a minimum estimated operating cost of USD 34,748.16.

At the limiting feasible sailing deadline of 29.365 hours, the entire route must be completed at Service Speed. Any deadline shorter than 29.365 hours is infeasible because even the fastest available speed mode, operating at 16 knots, cannot complete the 469.84-NM route within the required time.

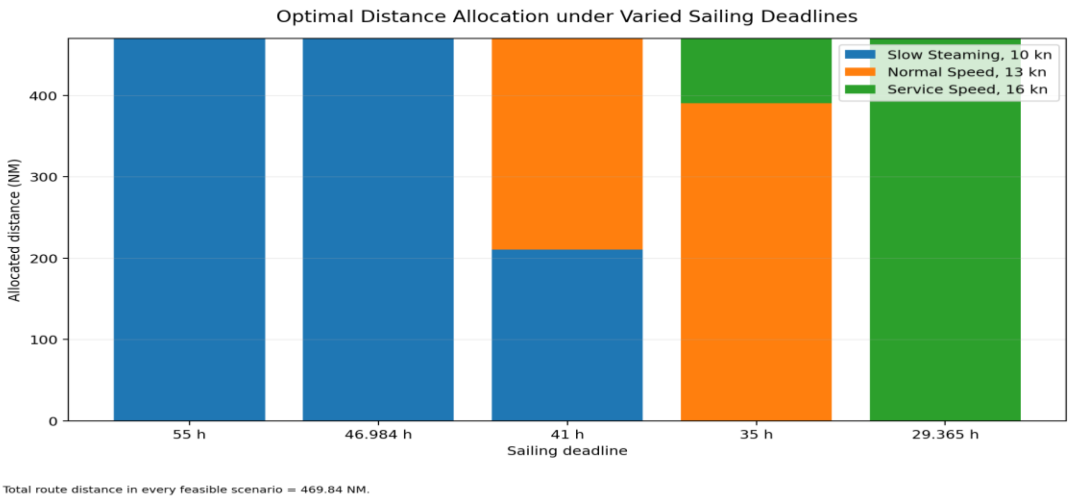

Figure 2. Continuous mixed-speed distance allocations under varied sailing-deadline scenarios for the aligned 469.84-NM route case.

From a mathematical perspective, Figure 1 should not be interpreted as evidence of an exponential cost relationship. Within this linear-programming formulation, tightening the right-hand-side value of the sailing-time constraint produces a piecewise-linear change in the optimal objective-function value. Changes in slope occur because the selected optimal basis changes across deadline intervals: Slow Steaming alone is optimal for deadlines of 46.984 hours or more; combinations of Slow Steaming and Normal Speed are selected for deadlines between approximately 36.142 and 46.984 hours; combinations of Normal Speed and Service Speed are selected for deadlines between approximately 29.365 and 36.142 hours; and no feasible solution exists below 29.365 hours.

The observed increase in operating cost under tighter deadline conditions is consistent with maritime speed-optimization literature showing that restrictive schedules require the use of higher-speed, higher-cost operational modes. The present result additionally provides a transparent model-based explanation: the cost increase occurs because the binding time constraint forces movement away from the least-cost per-NM mode toward progressively more expensive speed-mode combinations.

C. Sensitivity of the Optimal Solution to Fuel-Price Variation

The second sensitivity analysis examined changes in the objective-function value and distance allocation when the VLSFO price was varied while the sailing deadline remained fixed at 35 hours. For the aligned 469.84-NM route, this deadline is binding and requires a combination of Normal Speed and Service Speed because the route cannot be completed within 35 hours using Normal Speed alone.

TABLE 3.
FUEL-PRICE SENSITIVITY ANALYSIS AT A FIXED 35-HOUR SAILING DEADLINE FOR THE 469.84-NM ROUTE

Fuel-Price Scenario	VLSFO Price (USD/MT)	Slow Steaming, 10 kn (NM)	Normal Speed, 13 kn (NM)	Service Speed, 16 kn (NM)	Optimal Cost, USD	Status
Low-Price Assumption	400.00	0.00	390.69	79.15	21,152.44	Optimal
Baseline Scenario	755.00	0.00	390.69	79.15	34,748.16	Optimal
Higher-Price Scenario I	864.00	0.00	390.69	79.15	38,922.61	Optimal
Higher-Price Scenario II	993.50	0.00	390.69	79.15	43,882.18	Optimal

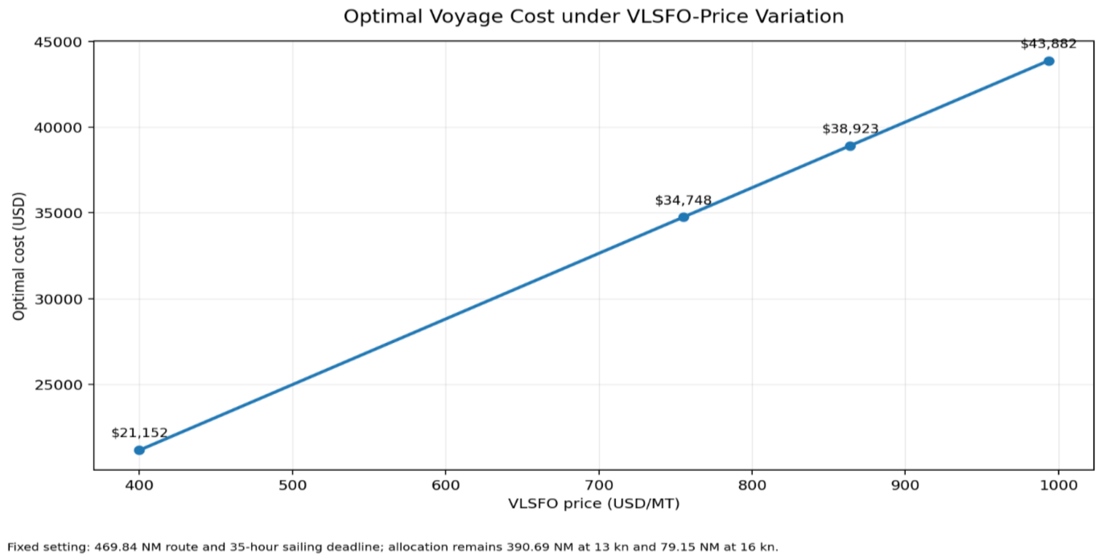

Figure 3. Linear change in optimal voyage cost under varied VLSFO-price inputs at a fixed 35-hour sailing deadline for the aligned 469.84-NM route case.

The results show that, within the examined fuel-price range, the optimal allocation remains unchanged at 390.69 NM under Normal Speed and 79.15 NM under Service Speed. Only the value of the objective function changes. This behavior is mathematically interpretable: at the fixed 35-hour deadline, the active distance and sailing-time constraints determine the selected two-mode allocation. Within the tested price range, changes in the cost vector do not cause an alternative feasible extreme point to become less costly. The optimal allocation therefore remains stable while the estimated total cost increases linearly with the fuel-price parameter.

This result should be interpreted as allocation stability within the examined parameter range, not as evidence that the model is universally unaffected by economic or operational changes. Changes in fuel-consumption assumptions, daily fixed cost, available speed modes, route distance, or uncertainty conditions may alter the ordering of the cost coefficients and consequently the selected optimal basis.

D. Functional Demonstration of the IDSS Prototype and LLM Explanation Layer

The web-based IDSS prototype was demonstrated using the same Tanjung Priok–Tanjung Perak case setting employed in the numerical analysis. The route-distance component produced a total sailing distance of 469.84 NM, which was submitted to the linear-programming engine together with the three predefined speed modes, a daily fixed operating cost of USD 4,000/day, a baseline VLSFO price of USD 755/MT, zero port-stay duration for the reported experiment, and a 35-hour sailing deadline.

Under these input parameters, the HiGHS-based optimization engine generated an optimal continuous mixed-speed allocation of 390.69 NM at the 13-knot Normal Speed mode and 79.15 NM at the 16-knot Service Speed mode. The computed sailing duration was

35.00 hours, the estimated operating cost was USD 34,748.16, and the estimated fuel consumption was 38.30 MT. Minor differences in displayed decimals may arise from interface-level rounding.

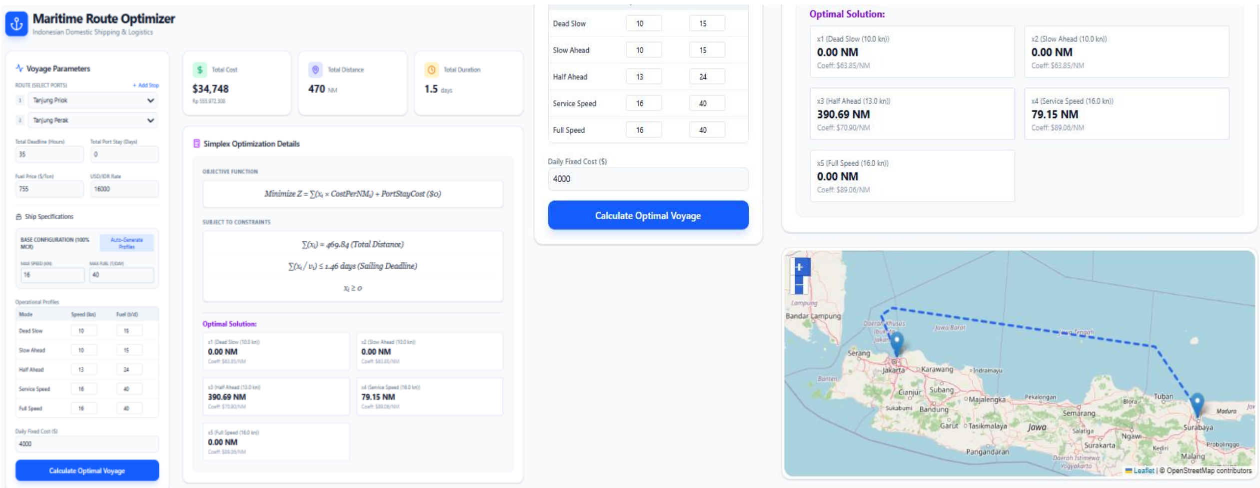

Figure 4. IDSS dashboard presenting the route, aligned input parameters, estimated operating cost, fuel consumption, and optimized continuous mixed-speed distance allocation for the 469.84-NM Tanjung Priok–Tanjung Perak case.

Figure 4 documents that the mathematical solution can be presented in a traceable numerical form within the application interface. The selected speed modes, allocated distances, sailing-time outcome, estimated cost, and fuel-consumption output can be reviewed directly against the case-study input parameters. Because the same configuration is used in the sensitivity analysis and in the dashboard demonstration, the prototype output is directly traceable to the numerical findings reported above.

Following computation of the optimum, the application transmitted the structured optimization output to the LLM explanation layer. The transmitted data included the route description, total route distance, selected speed modes and allocated distances, estimated sailing duration, voyage cost, fuel consumption, and relevant parameter information. The LLM then generated a natural-language Voyage Master Plan presenting the computed speed phases and supplementary nautical cautions.

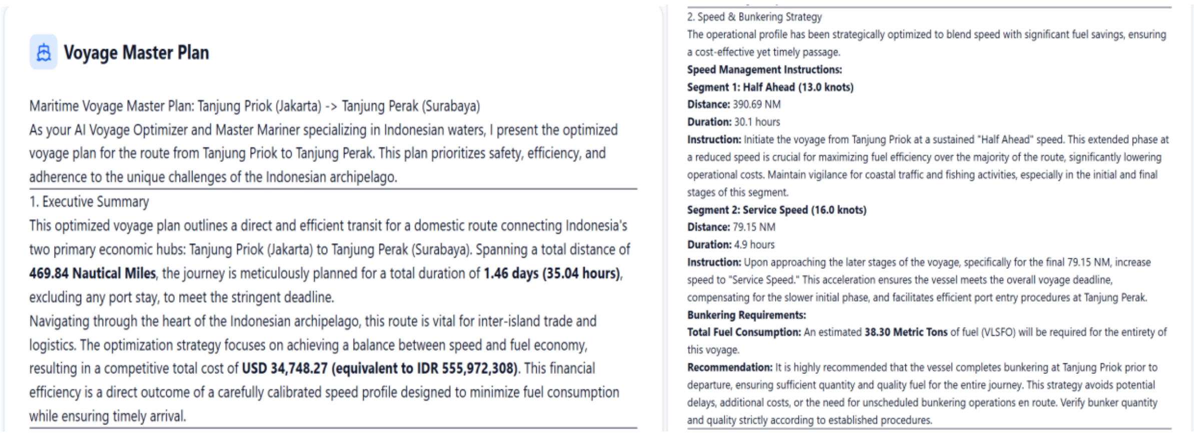

Figure 5. Natural-language Voyage Master Plan generated by the LLM explanation layer from the optimized output of the aligned 469.84-NM route scenario.

The output illustrated in Figure 5 provides evidence of functional integration between the mathematical optimization engine and the LLM-based explanation interface. The LLM reformulates structured optimization outputs into an operationally readable narrative; it does not define the constraints, determine the minimum-cost solution, modify the selected distance allocation, or validate navigational safety. Final decision authority remains with the ship master.

Accordingly, the prototype demonstration should be interpreted conservatively. It establishes that structured output from a mathematically traceable maritime speed-allocation model can be communicated through an LLM-supported narrative within the developed IDSS. It does not demonstrate reduced cognitive workload, improved trust, prevention of automation bias, or improved navigational safety. Such claims require subsequent empirical evaluation involving maritime experts and intended users.

E. Relation to Previous Studies and Contribution of the Present Study

The numerical findings reinforce the established relevance of mathematical speed optimization in maritime transportation. Consistent with Fagerholt et al. [1] and Psaraftis and Kontovas [2], the aligned case analysis demonstrates that tighter schedule requirements lead to the selection of higher-cost speed operations. The present study makes this relationship explicit through a continuous linear-programming formulation: as the sailing deadline is reduced, the time constraint becomes binding and progressively shifts the optimum from Slow Steaming toward Normal Speed and Service Speed. In the aligned Tanjung Priok–Tanjung Perak case, the 35-hour deadline requires a Normal-Speed and Service-Speed combination, while any deadline shorter than 29.365 hours becomes infeasible under the available speed-mode set.

At the same fixed 35-hour deadline, the examined fuel-price scenarios increase the estimated minimum cost without changing the selected distance allocation. This result clarifies the distinct roles of the model parameters in the reported case: deadline tightness determines the required speed-mode composition, whereas fuel-price variation affects the magnitude of the resulting operating cost within the tested interval.

In relation to maritime optimization studies in the Indonesian context, including Sirait et al. [3], this study addresses a focused voyage-level decision: the minimum-cost allocation of sailing distance across available speed modes for a specified domestic route. The model does not seek to replace broader fleet-allocation, inventory, routing, or scheduling models. Rather, it provides an analytically transparent component that may support evaluation of speed and arrival-time trade-offs within a wider decision-support architecture.

The LLM component occupies a deliberately bounded position within this contribution. In line with emerging interest in explainable and user-centered maritime decision support, the developed IDSS illustrates how a mathematically generated solution may be communicated in a structured operational narrative. The contribution is therefore not an AI-derived optimization method, but a model-centered architecture in which the mathematical solution remains traceable and the LLM supports interpretation only after the optimum has been obtained.

Research findings must be presented clearly and address the research questions. Results may be supplemented with figures and tables presenting research data. The discussion is written descriptively and narratively, avoiding excessive tables, figures, and graphs, while including relevant references (citations from sources) related to the research findings. The discussion differs from the results section and does not repeat what has already been described in the results. The discussion may explain the interpretation of research findings, compare results with similar studies, and describe the limitations of the study and how further research can complement the findings.

IV. Conclusion

This study developed a continuous linear-programming model for maritime speed allocation and fuel-cost efficiency within a web-based Intelligent Decision Support System (IDSS). The mathematical model allocates sailing distance across predefined vessel-speed modes while minimizing estimated voyage operating cost under route-distance and sailing-deadline constraints. In the aligned Tanjung Priok–Tanjung Perak case, the application-generated route distance was 469.84 NM. Under a 35-hour sailing deadline and a baseline VLSFO price of USD 755/MT, the optimal solution allocated 390.69 NM to the 13-knot Normal Speed mode and 79.15 NM to the 16-knot Service Speed mode. This allocation produced an estimated minimum operating cost of USD 34,748.16 and estimated fuel consumption of 38.30 MT.

The sensitivity analyses clarify the mathematical roles of the principal model parameters. Tightening the sailing deadline progressively shifts the optimal allocation from Slow Steaming toward higher-speed modes with larger cost coefficients. The resulting objective-function response is piecewise linear, reflecting changes in the active optimal basis as the time constraint becomes increasingly restrictive. For the examined route, deadlines shorter than 29.365 hours are infeasible because the fastest available mode cannot complete the voyage within the required time. In contrast, changes in the VLSFO-price parameter at a fixed 35-hour deadline alter the estimated minimum cost without changing the selected speed allocation within the tested range. This finding indicates allocation stability under the examined fuel-price scenarios, rather than universal insensitivity to economic or operational changes.

The developed IDSS also demonstrates a bounded integration of mathematical optimization and generative AI. The HiGHS-based solver remains solely responsible for determining the optimal allocation, while the LLM-based explanation layer translates the structured numerical output into a natural-language *Voyage Master Plan*. The contribution of the LLM is therefore interpretive rather than computational. The present prototype should not be regarded as an autonomous navigation system or a certified passage-planning instrument. Further research should incorporate vessel-specific operational data, professionally validated route information, combined uncertainty analysis, and user-based evaluation involving maritime experts to assess explanation accuracy, usability, trust, and operational appropriateness.

Acknowledgements

The authors would like to express their sincere gratitude to Politeknik Pelayaran Surabaya for providing institutional support and research facilities throughout this study. Appreciation is also extended to colleagues and reviewers who provided valuable feedback during the development of the mathematical model and the Intelligent Decision Support System (IDSS) prototype. Their insights contributed significantly to improving the quality and clarity of this research. Any remaining errors or limitations are solely the responsibility of the authors.

References

[1] K. Fagerholt, G. Laporte, and I. Norstad, "Reducing Fuel Emissions by Optimizing Speed on Shipping Routes," *J. Oper. Res. Soc.*, vol. 61, no. 3, pp. 523–529, 2010, doi: 10.1057/jors.2009.77.

[2] H. N. Psaraftis and C. A. Kontovas, "Speed Models for Energy-Efficient Maritime Transportation: A Taxonomy and Survey,"

Transp. Res. Part C Emerg. Technol., vol. 26, pp. 331–351, 2013, doi: 10.1016/j.trc.2012.09.012.

[3] M. Sirait, P. Charnsethikul, and N. Paoprasert, "A Multi-Type Ship Allocation and Routing Model for Multi-Product Oil Distribution in Indonesia with Inventory and Cost Minimization Considerations: A Mixed-Integer Linear Programming Approach," *Logistics*, vol. 9, no. 1, p. 35, 2025, doi: 10.3390/logistics9010035.

[4] D. Jirak, P. Maes, A. Saroukanoff, and D. van Rooy, "From Sea to System: Exploring User-Centered Explainable AI for Maritime Decision Support," *arXiv Prepr.*, 2025, [Online]. Available: https://arxiv.org/abs/2509.15084

[5] T. Miller, I. Durlik, E. Kostecka, A. Łobodzińska, K. Łazuga, and P. Kozlovska, "Leveraging Large Language Models for Enhancing Safety in Maritime Operations," *Appl. Sci.*, vol. 15, no. 3, p. 1666, 2025, doi: 10.3390/app15031666.

[6] FastAPI, "FastAPI Documentation." 2026. [Online]. Available: https://fastapi.tiangolo.com/

[7] G. Halili, "searoute (Version 1.6.0) [Python Package]." PyPI, 2026. [Online]. Available: https://pypi.org/project/searoute/

[8] P. Virtanen *et al.*, "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python," *Nat. Methods*, vol. 17, no. 3, pp. 261–272, 2020, doi: 10.1038/s41592-019-0686-2.

[9] Google AI for Developers, "Gemini API Documentation." 2026. [Online]. Available: https://ai.google.dev/gemini-api/docs

[10] Ship and Bunker, "Singapore Bunker Prices." 2026. [Online]. Available: https://shipandbunker.com/prices/apac/sea/sg-sin-singapore