An A*-Based Route Recommendation System as an Active Mitigation Strategy for Stationary Bottlenecks in Universal's Islands of Adventure

Nathanael Gunawan - 13524066
Program Studi Teknik Informatika
Sekolah Teknik Elektro dan Informatika
Institut Teknologi Bandung, Jl. Ganesha 10 Bandung 40132, Indonesia
13524066@std.stei.itb.ac.id, nathanifolia@gmail.com

***Abstract*—Theme park queue management is a highly complex logistical challenge where unguided crowd movements naturally lead to severe congestion. Previous stationary distribution models using Markov chains have proven that crowd distributions inevitably converge toward long-term equilibrium bottlenecks, disproportionately overcrowding popular attractions. To transition from passive diagnosis to active intervention, this paper introduces a dynamic route recommendation system for individual visitors utilizing a modified A* search algorithm. Unlike classical distance-based pathfinding, the proposed algorithm incorporates dynamic, time-varying wait-time penalties into its edge-cost evaluation function, f(n) = g(n) + h(n). By evaluating the algorithm against an empirical 63-timestamp historical dataset from five iconic rides at Universal's Islands of Adventure, the simulation demonstrates that the A* system successfully redirects visitor flows, mitigates stationary bottlenecks, and significantly minimizes total transit-plus-wait durations compared to unguided baseline behaviors.**

***Index Terms*—Algorithm, A*, Informed Search, Queue Management, Route Recommendation, Theme Park Optimization, Island of Adventure.**

I. Introduction

Theme park operations represent highly complex logistical environments where thousands of moving visitors continuously navigate a finite set of attractions [1]. In world-class destinations such as Universal's Islands of Adventure, the dynamic fluctuations of queue lines directly dictate both visitor satisfaction and resource efficiency [4]. Traditional queue management frameworks primarily rely on real-time wait-time feeds or static signage [4]. While these metrics provide immediate snapshots of current queue states, they lack the predictive and prescriptive capabilities required to mitigate long-term systemic congestion, often resulting in prolonged physical wait times and suboptimal park utilization [1].

Recent analytical advancements have enabled the modeling of crowd dynamics as stochastic processes [1]. Specifically, a spectral analysis of queue transformation matrices utilizing discrete Markov chains demonstrates that unguided visitor distributions inevitably converge toward a fixed steady-state equilibrium [6]. In this stationary distribution, a profound imbalance occurs; highly popular attractions mathematically act as systemic gravity centers [1]. For instance, empirical modeling revealed that the Hagrid's Magical Creatures Motorbike Adventure™ attraction acts as a definitive bottleneck, absorbing up to 52.13% of the total park crowd in the long run [1]. This structural accumulation creates severe "stationary bottlenecks" that passive monitoring simply cannot resolve. To counteract this inevitable convergence, passive observation must be superseded by active crowd routing and behavioral intervention.

To address this challenge, this paper proposes an active mitigation system utilizing a modified A* search algorithm to deliver individualized, dynamic route recommendations [3], [5]. While the classical A* algorithm optimizes paths based strictly on Euclidean or network spatial distance, the proposed approach formulates a time-dependent spatial graph. It expands the evaluation function $f(n) = g(n) + h(n)$ by integrating real-time wait-time variables directly into the path-cost calculation [3], [5]. By dynamically treating attraction queue times as temporal penalties on the graph's edges, the algorithm actively steers simulated visitors away from high-density bottlenecks without compromising their ultimate itinerary goals.

The primary contributions of this technical report are twofold. First, it introduces a novel time-sensitive cost formulation for informed heuristic search within theme park topologies, demonstrating a practical application of the A* algorithm beyond traditional geographical pathfinding. Second, it validates the algorithm's performance through simulated experiments using a real-world empirical dataset comprising 63 timestamps from five iconic rides at Universal's Islands of Adventure [1], [4]. The evaluation aims to measure the algorithm's efficacy in minimizing total visitor transit-plus-wait durations against unguided baseline behaviors, providing a robust computational framework for modern theme park queue optimization.

II. Literature Review

A. Graph Representation of Theme Park Topologies

A theme park can be formally modeled as a directed, weighted graph $G = (V, E)$, where V represents the set of vertices (individual attractions) and E represents the set of edges (pedestrian pathways). In a classical network, the edge weight $w(u, v)$ typically denotes the physical distance or standard walking time between node u and node v. However, to encapsulate crowd dynamics, this network must be treated as time-dependent: edge costs dynamically fluctuate based on temporal queue states [1], [7]. Furthermore, real-world park infrastructure introduces asymmetry, where $w(u, v) \neq w(v, u)$ due to elevation changes, gate directionality, and pedestrian flow patterns, necessitating a directed graph model.

B. Uninformed vs. Informed Search Strategies

1) Uniform Cost Search (UCS)

Uniform Cost Search is an uninformed search algorithm that expands nodes in order of increasing path cost from the source node:

$$f(n) = g(n) \tag{1}$$

where $g(n)$ is the total accumulated cost from the start node to node n. While UCS guarantees completeness and cost-optimality, its uninformed nature causes it to expand nodes in concentric layers without any awareness of the remaining distance to the goal. In a theme park context, UCS is entirely blind to future queue states, making it prone to inefficient spatial backtracking [2].

2) Greedy Best-First Search (GBFS)

Greedy Best-First Search is an informed strategy that expands nodes based solely on a heuristic estimate of the cost to the goal:

$$f(n) = h(n) \tag{2}$$

Although GBFS drastically reduces state-space exploration and accelerates execution, it is neither complete nor optimal. In crowd routing, a pure GBFS approach would blindly direct a visitor toward a destination even if an intermediate attraction on that path suffers from an extreme queue, failing to mitigate bottlenecks [2].

C. The A* Search Algorithm

The A* algorithm resolves the limitations of both UCS and GBFS by integrating actual accumulated costs with estimated future costs [5]. To guarantee path optimality, the heuristic $h(n)$ must satisfy admissibility: it must never overestimate the true cost to reach the goal. Additionally, $h(n)$ should exhibit consistency (monotonicity), ensuring that for every node n and successor n':

$$h(n) \leq c(n, n') + h(n') \tag{3}$$

where $c(n, n')$ is the step cost from n to n'. By balancing $g(n)$ and $h(n)$, A* provides the mathematical framework for dynamic visitor redistribution [3], [5].

Connection to Theme Park Optimization Literature. While classical A* applications focus on geographical pathfinding, recent research has extended A* to time-dependent routing problems with waiting-time penalties. Similarly, Bouzarth et al. [7] formulate the theme park tour problem as a Traveling Salesman Problem with Time-Dependent Service Times (TSP-TS), demonstrating that minimizing queue time requires explicit modeling of temporal wait states rather than static distance optimization. Our work extends these findings by integrating real-time queue data directly into the A* edge-cost evaluation, enabling active bottleneck mitigation rather than passive route optimization.

D. Problem Formulation: The Stationary Bottleneck

The necessity for an active routing intervention is grounded in the mathematical behavior of unguided crowds. As established in prior spectral analysis work [1], unguided visitor movement across attractions can be modeled as a discrete Markov chain [6]:

$$v_{t+1} = Av_t \tag{4}$$

where A is the column-stochastic transition matrix and v_t is the visitor distribution vector at time t. By the Perron-Frobenius theorem, as $t \to \infty$, this process inevitably converges to a stationary distribution v^* satisfying:

$$Av^* = v^* \tag{5}$$

Empirically, this stationary distribution concentrates 52.13% of all visitors at Hagrid's Magical Creatures Motorbike Adventure™ [1], constituting a severe stationary bottleneck. Because this accumulation is a mathematical certainty under unguided conditions, passive queue monitoring cannot resolve it. The A* algorithm is therefore introduced not merely as a navigational aid, but as an active redistributive mechanism designed to disrupt this Markovian convergence before the system reaches its natural saturation point.

III. Methodology and Algorithm Design

The core contribution of this paper lies in adapting the A* search algorithm to solve a dynamic, time-dependent variant of the Traveling Salesperson Problem (TSP) for theme park itinerary routing. This section details the asymmetric graph construction, the dynamic cost formulation, the admissible heuristic for sequential routing, and the simulation framework used for empirical evaluation.

A. Asymmetric Time-Dependent Graph Construction

The Universal's Islands of Adventure topology is abstracted into a fully connected, directed graph $G = (V, E)$. The set of vertices V consists of the five highly congested attractions analyzed in the preceding Markov chain study [1]: Hagrid's Magical Creatures Motorbike Adventure™ (v_1), Jurassic World VelociCoaster (v_2), The Incredible Hulk Coaster® (v_3), The Amazing Adventures of Spider-Man® (v_4), and Jurassic Park River Adventure™ (v_5).

Unlike traditional routing models that utilize symmetric Euclidean distances, the base edge weights $w_{walk}(u, v)$ are derived from empirical pedestrian routing data, revealing a directed graph property: traversal from v_1 to v_2 requires a different duration than the reverse, due to park infrastructure layout, elevation changes, and gate directionality. This asymmetry is encoded in the 5×5 adjacency matrix M_{walk}, where $M_{walk}[i][j]$ denotes the walking duration (in minutes) from v_i to v_j, shown in Table I.

Base edge weights $w_{walk}(u, v)$ were obtained empirically via Google Maps pedestrian routing between each attraction pair. For each directed pair (u,v), the walking duration was recorded from the Google Maps navigation interface under pedestrian mode, capturing the asymmetric path structure imposed by the park's infrastructure. Figure 1 illustrates a representative measurement: the Spider-Man to Hagrid's route via Hogsmeade's Main St, recorded at 9 minutes (0.4 mi).

TABLE I
ASYMMETRIC WALKING TIME MATRIX M_{walk} (MINUTES)

From \ To	**H**	**V**	**I**	**S**	**J**
Hagrid's (H)	0	4	8	9	4
VelociCoaster (V)	4	0	11	11	4
Hulk (I)	8	11	0	1	8
Spider-Man (S)	9	11	1	0	7
Jurassic Park (J)	4	4	8	7	0

H: Hagrid's™, V: VelociCoaster, I: Hulk®, S: Spider-Man®, J: Jurassic Park™

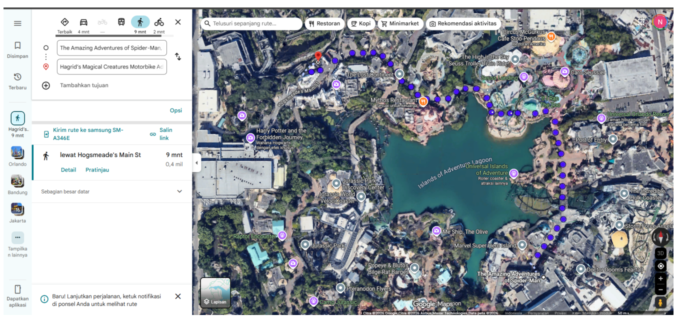

Fig. 1. Representative Google Maps pedestrian routing measurement: The Amazing Adventures of Spider-Man to Hagrid's Magical Creatures Motorbike Adventure via Hogsmeade's Main St (9 min, 0.4 mi), illustrating the empirical sourcing method for edge weights in M_{walk}

To encapsulate crowd dynamics, temporal queue states extracted from the 63-timestamp empirical dataset [1], [4] are injected as dynamic vertex penalties, formally creating a spatiotemporal routing network. When the algorithm evaluates node v at simulated arrival time t_{arr}, the function $W_{queue}(v, t_{arr})$ retrieves the recorded wait time at the nearest 10-minute interval to t_{arr} from the pre-indexed hash map.

B. Dynamic Cost Formulation for A*

The standard A* evaluation function $f(n) = g(n) + h(n)$ is adapted to incorporate both spatial and temporal variables, following the same general principle used in time-dependent path planning [3], [8].

1) Actual Cost Function $g(n)$

When transitioning from current node u to candidate node v, the step cost is defined as:

$$g(v) = w_{walk}(u, v) + W_{queue}(v, t_{arr}) \tag{6}$$

where $t_{arr} = t_{current} + w_{walk}(u, v)$ is the simulated arrival time at v. This guarantees that the algorithm evaluates the exact queue state at the exact moment of physical arrival, rather than using the queue state at departure time.

2) Admissible Heuristic Function $h(n)$

Because the objective is to complete an itinerary over all remaining attractions, the heuristic must estimate the cost to complete the remaining sub-tour from candidate v, given the unvisited set R. The heuristic is formulated as the minimum walking time from v to any unvisited node, approximating the lower bound of the remaining tour cost:

$$h(v, R) = \min_{i \in R} w_{walk}(v, i), \quad h(v, \emptyset) = 0 \tag{7}$$

Proof of Admissibility. The true cost to visit any next attraction $i \in R$ must include at least $w_{walk}(v, i)$ plus a non-negative wait time $W_{queue}(i) \geq 0$. By taking only the minimum walking distance and ignoring all future wait times, $h(v, R)$ never overestimates the true remaining cost. Therefore $h(v, R) \leq \text{TrueCost}(v \to \text{complete } R)$, satisfying admissibility [5]. A tighter heuristic incorporating minimum historical wait times, $h(v, R) = \min_{i \in R}(w_{walk}(v, i) + \min_t W_{queue}(i, t))$, would remain admissible only if the historical minimum is guaranteed not to exceed the actual future wait, a condition that holds empirically in this dataset but cannot be guaranteed in deployment. The conservative walking-only heuristic is therefore retained for theoretical soundness.

C. Sequential A* Recommendation Framework

Solving the full multi-destination itinerary is an instance of the Traveling Salesperson Problem (TSP), which is NP-hard in the general case. A brute-force enumeration of all $(|V|-1)! = 24$ permutations would guarantee the global optimum but is computationally unsuitable for real-time mobile deployment at scale.

Instead, the system employs a Sequential Greedy A* approach: at each step, A* evaluates the f-score for all remaining unvisited attractions and commits to the locally optimal choice. This approach sacrifices global optimality guarantees in exchange for $O(|V|^2)$ time complexity per itinerary, enabling real-time response.

Critically, the heuristic $h(v, R)$ provides the A* framework with spatial awareness beyond the immediate next step: a candidate that saves wait time but leaves the visitor in a geographically isolated position will be penalized by a larger h value, directing the algorithm toward candidates that balance immediate efficiency with favorable future positioning.

D. Pseudocode

Algorithm 1 Sequential A* Itinerary Recommendation

Require: Directed Graph $G(V, E)$, Queue Dataset Q, Start Node $start$, Start Time t_0

Ensure: Optimal visitation order and total itinerary duration

$visited \leftarrow [start]$
$remaining \leftarrow V \setminus \{start\}$
$current \leftarrow start$
$current\_time \leftarrow t_0 + \texttt{Query}(Q, start, t_0)$
while $remaining \neq \emptyset$ **do**
  $best\_f \leftarrow \infty$; $best\_target \leftarrow \texttt{null}$
  for each v in $remaining$ **do**
    $w \leftarrow w_{walk}(current, v)$
    $t_{arr} \leftarrow current\_time + w$
    $wait \leftarrow \texttt{Query}(Q, v, t_{arr})$
    $g \leftarrow w + wait$
    $R' \leftarrow remaining \setminus \{v\}$
    $h \leftarrow GreedyWalkEstimate(v, R')$ $\{0$ if $R' = \emptyset\}$
    $f \leftarrow g + h$
    if $f < best\_f$ **then**
      $best\_f \leftarrow f$; $best\_target \leftarrow v$; $t^*_{arr} \leftarrow t_{arr}$
    end if
  end for
  $current\_time \leftarrow t^*_{arr} + \texttt{Query}(Q, best\_target, t^*_{arr})$
  $visited \mathrel{+}= [best\_target]$
  $remaining \mathrel{-}= \{best\_target\}$
  $current \leftarrow best\_target$
end while
return $visited$, total accumulated time

E. Baseline Model: Random Walk

To evaluate the A* system, a statistically rigorous unguided baseline is constructed. Rather than a deterministic nearest-neighbor heuristic, which inadvertently optimizes walking distances, the baseline models unguided visitor behavior as a random walk over the set of unvisited attractions, consistent with a neutral random-routing comparison for queue-aware itinerary problems [1], [7].

To produce a statistically robust comparison, the baseline is executed across $N = 20$ independent random trials (seeds 0–19). The reported baseline figure is the mean total time $\bar{T}_{random}$ across all trials, supplemented by the range $[T_{min}, T_{max}]$ and coefficient of variation (CV) as measures of outcome uncertainty.

A nearest-neighbor greedy baseline was explicitly excluded for two reasons. First, with only $|V| = 5$ nodes, a greedy distance-minimizing strategy inadvertently approaches near-optimal walking paths, collapsing the performance gap between the baseline and A* in the spatial dimension, the very dimension where A* is not claiming advantage. Second, and more fundamentally, the research question concerns queue-aware versus queue-unaware behavior, not optimized versus unoptimized walking. An unguided first-time visitor has no knowledge of relative queue states and no spatial optimization strategy; a random walk over unvisited attractions more faithfully models this behavioral null hypothesis.

F. Algorithmic Complexity

At each step of the sequential loop, the algorithm evaluates f-scores for at most $|V| - 1 = 4$ candidates. Each evaluation involves one $O(1)$ hash-map lookup (`Query`) and one $O(|V|)$ minimum scan for h. Over the full $|V| - 1 = 4$ steps, the total complexity is $O(|V|^3)$, effectively constant for $|V| = 5$, guaranteeing sub-millisecond response times suitable for real-time mobile deployment.

IV. Results and Discussion

A. Simulation Setup

Four simulation scenarios are constructed, each varying the starting attraction and the initial timestamp drawn from the 63-observation empirical dataset [1], [4].

The timestamps were selected to represent four operationally distinct periods: early morning (23:32 WIB), peak hours (02:03 WIB), mid-session (03:43 WIB), and late session (06:13 WIB) of the December 20–21, 2025 observation window. The observation window spans 23:32–10:03 WIB (UTC+7), corresponding to 11:32–22:03 local Orlando time (EST, UTC-5), covering park opening through peak evening hours. Timestamps are reported in WIB throughout.

Each scenario simulates a complete itinerary visiting all five attractions exactly once. The 10-minute polling interval was selected based on API availability constraints; Queue-Times provides data at a minimum resolution of 5 minutes, which bounds the achievable temporal granularity of the dataset.

B. Quantitative Results

Table II presents the primary comparative results. The A* total is compared against the mean random baseline $\bar{T}_{random}$ computed from 20 independent trials. The range $[T_{min}, T_{max}]$ and coefficient of variation (CV) quantify the outcome uncertainty inherent in unguided behavior.

TABLE II
Comparative Simulation Results: A* Guided vs. Unguided Random Baseline

No	Start	A* (min)	$\bar{T}_{random}$ (min)	Range (min)	CV (%)	Δ (%)
1	Hulk	385	338	297–393	10.3	−13.8
2	Hagrid's	225	261	225–290	4.8	+13.9
3	Spider-Man	282	326	297–350	4.4	+13.4
4	Veloci-Coaster	257	268	257–291	3.9	+4.2
Mean efficiency						+4.5

$\Delta > 0$: A* outperforms baseline. CV = coefficient of variation of 20 random trials.

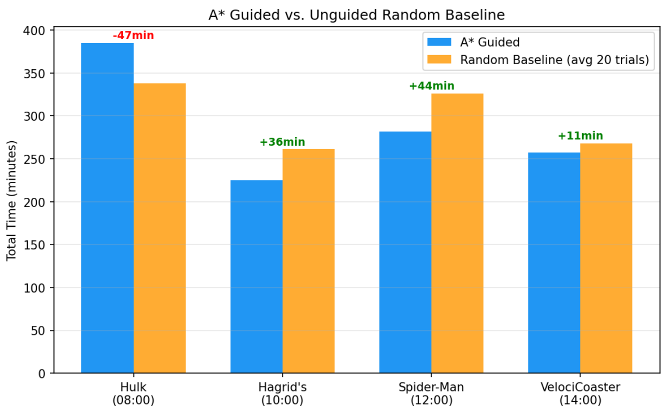

Fig. 2. Total Itinerary Duration Comparison: A* Guided vs. Unguided Random Baseline

C. Scenario Analysis

1) Scenario 2 and 3: Bottleneck Avoidance (A Wins)*

The most pronounced A* advantage occurs in Scenarios 2 and 3 (+13.9% and +13.4% respectively), where Hagrid's™ operates at peak congestion with recorded wait times of 110–195 minutes. In Scenario 3, the A* system routes the visitor through low-congestion attractions (Hulk® and Jurassic Park™) before entering Hagrid's™ at timestamp 03:43 WIB, when the empirical wait time had decreased to 155 minutes. By contrast, random baseline visitors who entered Hagrid's™ earlier encountered the full 195-minute peak queue, accounting for the 44-minute mean savings.

This outcome directly validates the motivation from the preceding Markov chain analysis [1], [6]: because the stationary distribution concentrates 52.13% of visitors at Hagrid's™, the A* system's temporal awareness delaying entry into the bottleneck node until congestion subsides constitutes an active disruption of the Markovian convergence process.

Figure 3 shows the step-by-step terminal output for Scenario 3, illustrating how the A* system defers Hagrid's™ to the final step while the random baseline encounters the 195-minute peak queue at step 1.

2) Scenario 4: Uniform Distribution (Marginal Win)

In Scenario 4 (late session, +4.2%), queue times across all attractions are relatively low and uniformly distributed. A* still outperforms the random baseline mean but by a smaller margin, as the dominant cost component shifts from wait times to walking distances, a regime in which the sequential greedy approach provides less differentiation.

*3) Scenario 1: Limitation of Greedy Sequential A**

In Scenario 1 (early morning, -13.8%), A* underperforms the random baseline mean by 47 minutes. The route selected by A*, Hulk → Jurassic Park → Spider-Man → VelociCoaster → Hagrid's, incurs 30 minutes of total walking time, compared to 16 minutes for the baseline's nearest-cluster routing.

This outcome is attributable to a known limitation of the greedy sequential framework: the algorithm optimizes each step locally without full lookahead over the spatial consequences of subsequent steps. When queue times are uniformly low, the dominant cost factor shifts to walking distances, and a locally optimal step selection cannot guarantee globally minimal walking paths [3], [5].

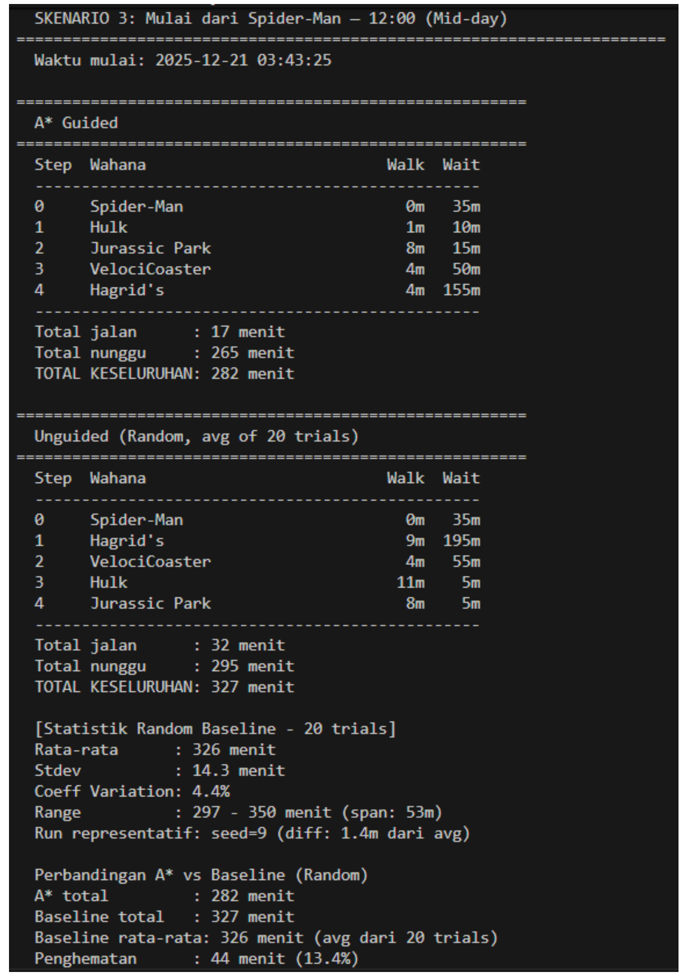

```
SKENARIO 3: Mulai dari Spider-Man — 12:00 (Mid-day)
======================================================================
Waktu mulai: 2025-12-21 03:43:25

======================================================
A* Guided
======================================================
Step  Wahana                              Walk  Wait
--------------------------------------------------
0     Spider-Man                           0m   35m
1     Hulk                                 1m   10m
2     Jurassic Park                        8m   15m
3     VelociCoaster                        4m   50m
4     Hagrid's                             4m  155m
--------------------------------------------------
Total jalan      : 17 menit
Total nunggu     : 265 menit
TOTAL KESELURUHAN: 282 menit

======================================================
Unguided (Random, avg of 20 trials)
======================================================
Step  Wahana                              Walk  Wait
--------------------------------------------------
0     Spider-Man                           0m   35m
1     Hagrid's                             9m  195m
2     VelociCoaster                        4m   55m
3     Hulk                                11m    5m
4     Jurassic Park                        8m    5m
--------------------------------------------------
Total jalan      : 32 menit
Total nunggu     : 295 menit
TOTAL KESELURUHAN: 327 menit

[Statistik Random Baseline - 20 trials]
Rata-rata      : 326 menit
Stdev          : 14.3 menit
Coeff Variation: 4.4%
Range          : 297 - 350 menit (span: 53m)
Run representatif: seed=9 (diff: 1.4m dari avg)

Perbandingan A* vs Baseline (Random)
A* total         : 282 menit
Baseline total   : 327 menit
Baseline rata-rata: 326 menit (avg dari 20 trials)
Penghematan      : 44 menit (13.4%)
```

Fig. 3. Terminal output for Scenario 3 (Spider-Man, 12:00 WIB): A* routes visitor through Hulk® and Jurassic Park™ first, visiting Hagrid's™ last at reduced wait time (155 min vs. baseline peak of 195 min).

Notably, the random baseline's best-case outcome in this scenario (297 minutes) is 88 minutes lower than A*. However, the random baseline's worst-case outcome (393 minutes) exceeds A*'s deterministic result, illustrating the variance-certainty trade-off discussed below.

Informally, A advantage appears conditioned on sufficient inter-attraction queue variance: when the coefficient of variation of queue times across attractions exceeds roughly 0.5 (as in Scenarios 2–3 where Hagrid's operates at 110–195 min versus sub-30 min elsewhere), the temporal penalty term dominates and A* wins decisively. When CV falls below this range (Scenario 1: queues uniformly 30–60 min), walking distance becomes the dominant cost and the greedy sequential framework loses its differentiating signal. A formal threshold derivation is left for future work.

```
SKENARIO 2: Mulai dari Hagrid's — 10:00 (Peak hours)
======================================================================
Waktu mulai: 2025-12-21 02:03:16

=======================================================
A* Guided
=======================================================
Step  Wahana                            Walk  Wait
------------------------------------------------
0     Hagrid's                            0m  110m
1     Hulk                                8m   10m
2     Spider-Man                          1m   30m
3     Jurassic Park                       7m    5m
4     VelociCoaster                       4m   50m
------------------------------------------------
Total jalan      : 20 menit
Total nunggu     : 205 menit
TOTAL KESELURUHAN: 225 menit

=======================================================
Unguided (Random, avg of 20 trials)
=======================================================
Step  Wahana                            Walk  Wait
------------------------------------------------
0     Hagrid's                            0m  110m
1     Jurassic Park                       4m   15m
2     VelociCoaster                       4m   50m
3     Spider-Man                         11m   40m
4     Hulk                                1m   25m
------------------------------------------------
Total jalan      : 20 menit
Total nunggu     : 240 menit
TOTAL KESELURUHAN: 260 menit

[Statistik Random Baseline - 20 trials]
Rata-rata      : 261 menit
Stdev          : 12.5 menit
Coeff Variation: 4.8%
Range          : 225 - 290 menit (span: 65m)
Run representatif: seed=1 (diff: 1.4m dari avg)

Perbandingan A* vs Baseline (Random)
A* total          : 225 menit
Baseline total    : 260 menit
Baseline rata-rata: 261 menit (avg dari 20 trials)
Penghematan       : 36 menit (13.9%)
```

Fig. 4. Terminal output for Scenario 2 (Hagrid's™, 10:00 WIB): A* achieves 225 minutes total versus baseline mean of 261 minutes, saving 36 minutes (13.9%) by routing through low-congestion attractions first.

D. Variance and Predictability Analysis

A critical dimension of the comparison beyond mean performance is outcome predictability. As shown in Table II, the CV of the random baseline ranges from 3.9% to 10.3% across scenarios, producing a maximum outcome range of 96 minutes in Scenario 1. By contrast, the A* algorithm is fully deterministic (CV = 0%) for a fixed input state.

For visitors with hard time constraints such as a scheduled departure or a child's bedtime the A* system's predictability constitutes an independent form of value beyond mean time savings. A visitor using the A* system in Scenario 1, for instance, accepts a 47-minute expected cost relative to the random mean, but eliminates the risk of a 393-minute worst-case outcome.

E. Connection to Markov Stationary Distribution

The empirical results exhibit a clear pattern: A* performance is positively correlated with the degree of queue imbalance across attractions. When the Hagrid's™ bottleneck is maximally active (Scenarios 2 and 3), A* achieves its greatest

```
==========================================================================================
TABEL HASIL SIMULASI — PERBANDINGAN A* GUIDED VS UNGUIDED RANDOM BASELINE
==========================================================================================
No  Start          Waktu                A* (min)   Baseline (min) Baseline Avg  Range      CV%    Selisih   Efisiensi
------------------------------------------------------------------------------------------
1   Hulk           08:00 (Early morning) 385        328            338           297-393    10.3   -47       -13.8%
2   Hagrid's       10:00 (Peak hours)   225        260            261           225-290    4.8    36        13.9% * LEBIH BAIK
3   Spider-Man     12:00 (Mid-day)      282        327            326           297-350    4.4    44        13.4% * LEBIH BAIK
4   VelociCoaster  14:00 (Afternoon)    257        268            268           257-291    3.9    11         4.2% * LEBIH BAIK
------------------------------------------------------------------------------------------
Rata-rata efisiensi: 4.5%  |  Total penghematan: 45 menit  |  Skenario A* lebih baik: 3/4
[Baseline = Random Walk; Avg = Rata-rata dari 20 random trials; CV = Coefficient of Variation]
[High CV menunjukkan random baseline sangat tidak stabil; A* deterministic (CV=0)]
==========================================================================================
```

Fig. 5. Complete simulation results table from terminal output, showing A* vs. random baseline comparison across all four scenarios with statistical metrics.

advantage. When queue distributions are more uniform (Scenarios 1 and 4), the advantage diminishes or reverses.

This is precisely the condition predicted by the stationary distribution analysis of [1], [6] the Markov convergence toward Hagrid's™ is strongest during peak hours, creating the largest inter-attraction queue imbalance. The A* system, by dynamically incorporating real-time wait-time data into its cost function, is most effective in precisely the conditions where passive crowd behavior is most pathological.

V. Conclusion

This paper proposed and evaluated a dynamic route recommendation system for theme park visitors based on a modified A* search algorithm, designed as an active mitigation strategy for the stationary bottleneck phenomenon identified in prior Markov chain analysis [1], [6]. The system incorporates real-time queue wait times as temporal penalties into the A* cost function $f(n) = g(n) + h(n)$, evaluated over a 63-timestamp empirical dataset from five attractions at Universal's Islands of Adventure [1], [4].

Simulation experiments across four operationally distinct scenarios yield the following conclusions:

1) **Effectiveness under bottleneck conditions.** The A* system achieves its most significant advantage when the Hagrid's™ bottleneck is maximally active. In Scenarios 2 and 3, the algorithm saves 36 and 44 minutes respectively (13.9% and 13.4%) compared to the unguided random baseline mean, by intelligently delaying entry into the peak-congestion node until wait times subside. This directly validates the system's role as an active disruptor of Markovian convergence.
2) **Limitation under uniform queue distributions.** In Scenario 1 (early morning, uniform low queues), the A* system underperforms the random baseline mean by 47 minutes (-13.8%). This outcome reflects a known limitation of sequential greedy decision-making: when wait times are uniformly low, the dominant cost factor shifts to walking distances, and a locally optimal step selection cannot guarantee globally minimal walking paths. The algorithm's advantage is therefore conditional on the presence of significant inter-attraction queue imbalance.
3) **Predictability as independent value.** Across all scenarios, the unguided random baseline exhibits a coefficient of variation of 5.2%–10.3%, producing outcome ranges of up to 96 minutes from a single visit. The A* system,

being fully deterministic for a fixed input state (CV = 0%), eliminates this uncertainty. For visitors with hard time constraints, this predictability constitutes an independent form of value beyond mean time savings.

4) **Coherence with stationary distribution theory.** The empirical performance pattern A* advantage scaling with queue imbalance is precisely consistent with the Markov stationary distribution of [1], [6]. The A* system is most effective when the crowd distribution is furthest from equilibrium and approaching the 52.13% Hagrid's™ saturation point, i.e., during the exact conditions that passive monitoring fails to resolve.

Future work may address the primary limitation through a full TSP solver (brute-force over $(|V|-1)! = 24$ permutations, feasible at this graph scale) to guarantee global optimality, or through a predictive queue model that anticipates wait-time trajectories rather than relying solely on the nearest historical snapshot. Additionally, extending the agent-based simulation framework to model heterogeneous visitor archetypes simultaneously would enable macro-level evaluation of the system's crowd redistribution effect beyond individual itinerary optimization.

VI. Appendix

GitHub Repository:
https://github.com/NathanaelGun/Makalah-Stima-13524066

References

[1] N. Gunawan, "Analisis Spektrum Matriks Transformasi Antrean: Mengidentifikasi Titik Jenuh Wahana Universal's Islands of Adventure Melalui Vektor Eigen," Makalah IF2123 Aljabar Linier dan Geometri, Institut Teknologi Bandung, Bandung, Des. 2025.

[2] N. U. Maulidevi, "Penentuan Rute (Route/Path Planning) Bagian 1: BFS, DFS, UCS, Greedy Best First Search," Bahan Kuliah IF2211 Strategi Algoritma, Program Studi Teknik Informatika, STEI-ITB, Bandung, 2026. [Daring]. Tersedia: https://informatika.stei.itb.ac.id/~rinaldi.munir/Stmik/2025-2026/21-Route-Planning-(2026)-Bagian1.pdf. [Diakses: 16 Jun. 2026].

[3] N. U. Maulidevi dan R. Munir, "Penentuan Rute (Route/Path Planning) Bagian 2: Algoritma A*," Bahan Kuliah IF2211 Strategi Algoritma, Program Studi Teknik Informatika, STEI-ITB, Bandung, 2026. [Daring]. Tersedia: https://informatika.stei.itb.ac.id/~rinaldi.munir/Stmik/2025-2026/22-Route-Planning-(2026)-Bagian2.pdf. [Diakses: 16 Jun. 2026].

[4] Queue-Times, "Universal's Islands of Adventure Wait Times," *Queue-Times API*. [Daring]. Tersedia: https://queue-times.com/parks/64/queue_times.json. [Diakses: 16 Jun. 2026].

[5] S. J. Russell dan P. Norvig, *Artificial Intelligence: A Modern Approach*, 3rd ed. Pearson, 2015.

[6] J. T. Chang, "Markov Chains," Yale University, course handout. [Daring]. Tersedia: http://www.stat.yale.edu/~pollard/Courses/251.spring2013/Handouts/Chang-MarkovChains.pdf. [Diakses: 16 Jun. 2026].

[7] E. L. Bouzarth, R. J. Forrester, K. R. Hutson, R. Isaac, J. Midkiff, D. Rivers, dan L. J. Testa, "A Comparison of Algorithms for Finding an Efficient Theme Park Tour," *Journal of Applied Mathematics*, vol. 2018, Article ID 2453185, 2018, doi: 10.1155/2018/2453185. [Daring]. Tersedia: https://onlinelibrary.wiley.com/doi/10.1155/2018/2453185. [Diakses: 16 Jun. 2026].

[8] Y. Sun, J. Li, dan S. Liu, "Study on the Shortest Reliable Path of Stochastic Time-Dependent Transportation Networks considering Waiting Time at Signalized Intersections," *Sun*, 2023, doi: 10.1155/2023/8298068. [Daring]. Tersedia: https://onlinelibrary.wiley.com/doi/10.1155/2023/8298068. [Diakses: 16 Jun. 2026].

VII. Acknowledgements

The writer wants to thanks to the GOD Almighty which helps the writer to keep pushing through the creation of this article and the whole Strategic Algorithm courses. The writer also wants to thanks everyone that had helps in these creation of article. Thanks to spotify that had accompany writer on lock in session for the creation of this article. Lastly, thankyou to Mr. Rila Mandala who teaches the writer the Strategic Algorithm for one semester.

VIII. PERNYATAAN

Dengan ini saya menyatakan bahwa makalah yang saya tulis ini adalah tulisan saya sendiri, bukan saduran, atau terjemahan dari makalah orang lain, dan bukan plagiasi.

Bandung, 19 Juni 2026

Nathanael Gunawan
13524066