

# HiMAP-Travel

Hierarchical Multi-Agent Planning for Long-Horizon Travel

Accepted at **ICML 2026**

Authors: **Bui The Viet** [1], Li Wenjun [1], Liu Yong [2]

[1] School of Computing and Information Systems, Singapore Management University

[2] Singapore Search & Recommendation Lab, Huawei

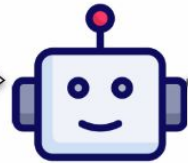
# Problem Statement



User

I'm going from *Seattle to California* from *November 6 to 10, 2023*. I have a *budget of \$6,000*. For lodging, I prefer an *entire room* and the accommodations must be *pet-friendly*.

Let me help! To solve this problem, I need to (1) **analyze certain constraints**, (2) **collect information** through reasonable use of necessary tools.



Agent

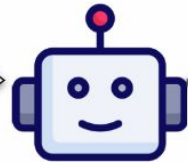
# Problem Statement



User

I'm going from *Seattle to California* from *November 6 to 10, 2023*. I have a *budget of \$6,000*. For lodging, I prefer an *entire room* and the accommodations must be *pet-friendly*.

Let me help! To solve this problem, I need to (1) **analyze certain constraints**, (2) **collect information** through reasonable use of necessary tools.



Agent

## Planning

The plan must adhere to certain **constraints**, e.g., **user needs** and **commonsense**. It's also vital to ...



### User Needs (Hard Constraints)

1. Budget: \$6000
2. Room Type: Entire Room
3. Room Rule: Pet-friendly



### Commonsense Constraints

1. Reasonable City Route
2. Diverse Restaurants
3. Diverse Attractions
4. Non-conflicting Transportation
5. Accommodation meets Minimum Night



## Delivery Plan

2023-11-06	2023-11-07	2023-11-08	2023-11-09	2023-11-10
Seattle -> Los Angeles	Los Angeles	Los Angeles -> San Diego	San Diego	San Diego -> Seattle
<ul style="list-style-type: none"> <li>Flight: F123 (13:40-16:12), Cost: \$120</li> <li>Accommodation: Luxury building studio</li> <li>Dinner: The Attraction</li> </ul>	<ul style="list-style-type: none"> <li>Breakfast: Chicken Minar</li> <li>Lunch: Rajdhani Restaurant</li> <li>Dinner: Domino's Pizza</li> <li>Attractions: Santa Monica Pier, Griffith Park</li> <li>Accommodation: Luxury building studio</li> </ul>	<ul style="list-style-type: none"> <li>Take taxi to San Diego</li> <li>Breakfast: Open Yard</li> <li>Lunch: The Lost Mughal</li> <li>Dinner: Burger King</li> <li>Attractions: Cabrillo Monument</li> <li>Accommodation: East Side Apartment</li> </ul>	<ul style="list-style-type: none"> <li>Breakfast: Baskin Robbins</li> <li>Lunch: Harry's Bar</li> <li>Dinner: Dragon Way</li> <li>Attractions: La Jolla Shores Park, California Tower</li> <li>Accommodation: East Side Apartment</li> </ul>	<ul style="list-style-type: none"> <li>Flight: F789 (7:30-10:56), Cost: \$300</li> </ul>

2023-11-06

### Seattle -> Los Angeles

- Flight: F123, (13:40-16:12), Cost: \$120
- Accommodation: Luxury building studio
- Dinner: The Attraction

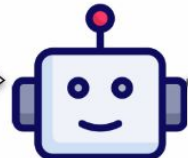
# Problem Statement



User

I'm going from *Seattle to California* from *November 6 to 10, 2023*. I have a *budget of \$6,000*. For lodging, I prefer an *entire room* and the accommodations must be *pet-friendly*.

Let me help! To solve this problem, I need to (1) **analyze certain constraints**, (2) **collect information** through reasonable use of necessary tools.



Agent

## Toolbox

1. CitySearch(·)
2. AttractionSearch(·)
3. FlightSearch(·)
4. DistanceMatrix(·)
5. RestaurantSearch(·)
6. AccommodationSearch(·)



## Interaction with Environment



## Planning

The plan must adhere to certain **constraints**, e.g., **user needs** and **commonsense**. It's also vital to ...



### User Needs (Hard Constraints)

1. Budget: \$6000
2. Room Type: Entire Room
3. Room Rule: Pet-friendly



### Commonsense Constraints

1. Reasonable City Route
2. Diverse Restaurants
3. Diverse Attractions
4. Non-conflicting Transportation
5. Accommodation meets Minimum Night



## Delivery Plan

2023-11-06	2023-11-07	2023-11-08	2023-11-09	2023-11-10
Seattle -> Los Angeles	Los Angeles	Los Angeles -> San Diego	San Diego	San Diego -> Seattle
<ul style="list-style-type: none"> <li>Flight: F123 (13:40-16:12), Cost: \$120</li> <li>Accommodation: Luxury building studio</li> <li>Dinner: The Attraction</li> </ul>	<ul style="list-style-type: none"> <li>Breakfast: Chicken Minar</li> <li>Lunch: Rajdhani Restaurant</li> <li>Dinner: Domino's Pizza</li> <li>Attractions: Santa Monica Pier, Griffith Park</li> <li>Accommodation: Luxury building studio</li> </ul>	<ul style="list-style-type: none"> <li>Take taxi to San Diego</li> <li>Breakfast: Open Yard</li> <li>Lunch: The Lost Mughal</li> <li>Dinner: Burger King</li> <li>Attractions: Cabrillo Monument</li> <li>Accommodation: East Side Apartment</li> </ul>	<ul style="list-style-type: none"> <li>Breakfast: Baskin Robbins</li> <li>Lunch: Harry's Bar</li> <li>Dinner: Dragon Way</li> <li>Attractions: La Jolla Shores Park, California Tower</li> <li>Accommodation: East Side Apartment</li> </ul>	<ul style="list-style-type: none"> <li>Flight: F789 (7:59-10:56), Cost: \$300</li> </ul>

2023-11-06

### Seattle -> Los Angeles

- Flight: F123, (13:40-16:12), Cost: \$120
- Accommodation: Luxury building studio
- Dinner: The Attraction

## Information Collection

[Tool] CitySearch[California]

[Result] San Francisco, Los Angeles, ..., San Diego



[Tool] FlightSearch[Seattle, San Francisco, 2023-11-06]

[Result] No Flights.

[Tool] FlightSearch[Seattle, Los Angeles, 2023-11-06]

[Result] Flight Number: F123, 13:40-16:12, Cost: \$120

[Tool] DistanceMatrix[Los Angeles, San Diego, taxi]

[Result] Duration: 1 hour 57 mins, Distance: 193 km, Cost: \$200

[Tool] TransportationSearch[San Diego, Seattle, 2023-11-10]

[Result] Flight Number: F789, (7:59-10:56), Cost: \$300



[Tool] AccommodationSearch[Los Angeles]

[Result] 'Cozy Room for U', \$130/night, Minimum night: 8, Entire Room, Pets allowed  
'Luxury building studio', \$150/night, Minimum night: 1, Entire Room, Pets allowed



[Tool] RestaurantSearch[Los Angeles]

[Result] The Attraction, Cuisine: French, ...

# Long-Horizon Travel Planning is Hard

”

”State-of-the-art LLMs cannot handle complex planning tasks... even GPT-4 successfully produces a plan that meets all the constraints for a few tasks (0.6%).”

”

— TravelPlanner Benchmark (Jian Xie et al., AAAI 2024)

# Why Hard?

**Query:** Please plan a travel itinerary for me. I'm departing from Cincinnati and heading to Norfolk for three days. The dates of travel are from March 10th to **March 12th, 2022**. I have a budget of \$1,400 for this trip.

## Trajectory:

Action 1: FlightSearch[Cincinnati, Norfolk, **2023-03-10**]

Observation 1: There is no flight from Cincinnati to Norfolk on **2023-03-10**.

Action 2: FlightSearch[Cincinnati, Norfolk, **2023-03-09**]

Action 3: FlightSearch[Cincinnati, Norfolk, **2023-03-11**]

Action 4: FlightSearch[Cincinnati, Richmond, **2023-03-10**]

...

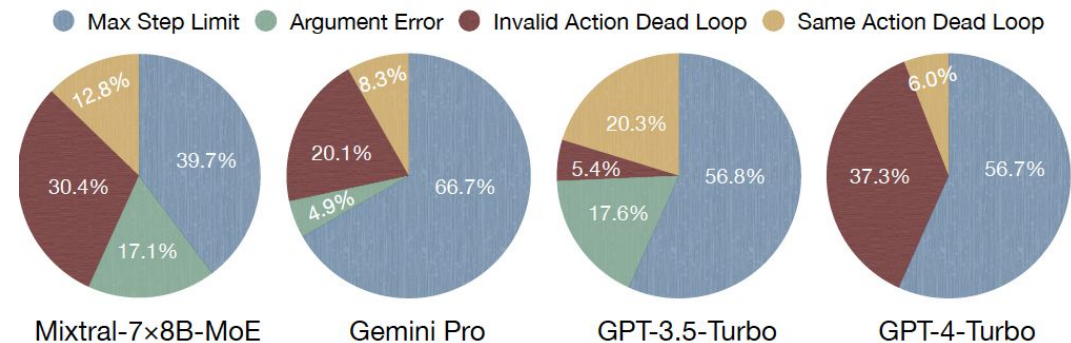
Action 8: FlightSearch[Cincinnati, Charlotte, 2023-03-10]

Action 9: **I'm sorry, but as an AI language model, I am unable to execute actions or observe their outcomes.**

## Analysis:

In this case, even though all actions are successful, the final step still fails. The query requires the date to be in 2022, but language agent is fixed in 2023. Therefore, even though it tries cities near Norfolk, the returned information remains null. It can also be observed that GPT-4 ultimately chooses not to continue, rather than trying other tools, after 8 failed attempts.

	Validation (#180)						Test (#1,000)					
	Delivery Rate	Commonsense Pass Rate		Hard Constraint Pass Rate		Final Pass Rate	Delivery Rate	Commonsense Pass Rate		Hard Constraint Pass Rate		Final Pass Rate
		Micro	Macro	Micro	Macro			Micro	Macro	Micro	Macro	
Greedy Search	100	74.4	0	60.8	37.8	0	100	72.0	0	52.4	31.8	0
<b>Two-stage</b>												
Mistral-7B-32K (Jiang et al., 2023)	8.9	5.9	0	0	0	0	7.0	4.8	0	0	0	0
Mixtral-8x7B-MoE (Jiang et al., 2024)	49.4	30.0	0	1.2	0.6	0	51.2	32.2	0.2	0.7	0.4	0
Gemini Pro (G Team et al., 2023)	28.9	18.9	0	0.5	0.6	0	39.1	24.9	0	0.6	0.1	0
GPT-3.5-Turbo (OpenAI, 2022)	86.7	54.0	0	0	0	0	91.8	57.9	0	0.5	0.6	0
GPT-4-Turbo (OpenAI, 2023)	<b>89.4</b>	<b>61.1</b>	<b>2.8</b>	<b>15.2</b>	<b>10.6</b>	<b>0.6</b>	<b>93.1</b>	<b>63.3</b>	<b>2.0</b>	<b>10.5</b>	<b>5.5</b>	<b>0.6</b>
<b>Sole-planning</b>												
Direct <sub>GPT-3.5-Turbo</sub>	100	60.2	4.4	11.0	2.8	0	100	59.5	2.7	9.5	4.4	0.6
CoT <sub>GPT-3.5-Turbo</sub>	100	66.3	3.3	11.9	5.0	0	100	64.4	2.3	9.8	3.8	0.4
ReAct <sub>GPT-3.5-Turbo</sub>	82.2	47.6	3.9	11.4	6.7	0.6	81.6	45.9	2.5	10.7	3.1	0.7
Reflexion <sub>GPT-3.5-Turbo</sub>	93.9	53.8	2.8	11.0	2.8	0	92.1	52.1	2.2	9.9	3.8	0.6
Direct <sub>Mixtral-8x7B-MoE</sub>	100	68.1	5.0	3.3	1.1	0	99.3	67.0	3.7	3.9	1.6	0.7
Direct <sub>Gemini Pro</sub>	93.9	65.0	8.3	9.3	4.4	0.6	93.7	64.7	7.9	10.6	4.7	2.1
Direct <sub>GPT-4-Turbo</sub>	<b>100</b>	<b>80.4</b>	<b>17.2</b>	<b>47.1</b>	<b>22.2</b>	<b>4.4</b>	<b>100</b>	<b>80.6</b>	<b>15.2</b>	<b>44.3</b>	<b>23.1</b>	<b>4.4</b>



# Case Study

**Query:** Can you create a 5-day travel plan for 2 people departing from Syracuse to visit 2 cities in Georgia? We are planning to travel from March 16th to March 20th, 2022. Our budget is approximately \$2,000. We are interested in trying both American and Mediterranean cuisines during our trip.

## Trajectory:

...

Action 13: FlightSearch[Syracuse, Atlanta, **2023-03-20**]

Thought 14: Since I have **exhausted the options for finding flights** from Syracuse to any city in Georgia on the desired dates, and I am unable to search for alternative transportation methods due to the nested use of tools restriction, the only remaining step is to use the Planner tool with the information that has been gathered so far.

Action 14: Planner[Create a 5-day travel plan from Syracuse to Georgia]

## Plan:

Day 1:  
transportation: **Flight Number: F1234567, from Syracuse to Atlanta**

...

## Analysis:

In this example, the language agent's persistent date errors and fixation on airplane travel lead it to bypass valid information search. Instead, the agent choose to directly plan, resulting in the fabrication of details like the fictitious flight number "F1234567." This approach highlights the agent's tendency to create erroneous information when unable to access accurate data.

**Query:** Please create a travel itinerary for a solo traveler departing from Jacksonville and heading to Los Angeles for a period of 3 days, specifically from March 25th to March 27th, 2022. The budget for this trip is now set at \$2,400.

## Plan:

Day 1:  
transportation: **Flight Number: F3765485**, from Jacksonville to Los Angeles

...

accommodation: **Midtown East 2 Bed/2 Bath Elevator Building Apt (7-day minimum night)**

...

Day 2:  
accommodation: **Midtown East 2 Bed/2 Bath Elevator Building Apt (7-day minimum night)**

...

Day 3:  
transportation: **Flight Number: F3765485, from Los Angeles to Jacksonville**

## Analysis:

The first issue is the chosen accommodation does not meet the 7-day minimum stay requirement, and we find language agent just directly select the first option from search results. The second issue is that language agent incorrectly copies the flight number from day 1 for the day 3 flight from Los Angeles to Jacksonville, despite being provided with the correct flight number.

**Query:** Could you help develop a week-long travel itinerary suitable for a group of 6 people, departing from Baton Rouge and planning to visit 3 different cities in Texas? The travel dates are set from March 17th to March 23rd, 2022. Our travel budget has been adjusted to \$14,600. Bearing in mind that we have children under ten years old, our accommodations need to allow young children and we prefer to occupy entire rooms. We also prefer not to self-drive during this trip.

## Plan:

```
{  
  "days": 1,  
  "current_city": "from Baton Rouge to Texarkana",  
  "breakfast": "-",  
  "lunch": "Big City Bread Cafe, Texarkana",  
  "dinner": "Poets Cafe, Texarkana",  
},  
{  
  "days": 2,  
  "current_city": "Texarkana",  
  "transportation": "-",  
  "breakfast": "Big City Bread Cafe, Texarkana",  
  ...  
}
```

## Analysis:

In this scenario, the language agents select the same restaurant for both lunch on the first day and breakfast on the second day, a choice that seems to defy commonsense.

# The Problem: Pathology of Monolithic Planning

## The Challenge

### Semantic vs. Combinatorial Reasoning

- LLMs excel at semantic reasoning
- **But fail at Long-Horizon Combinatorial Optimization**

### Hard Constraints

- Rigid logical invariants (e.g., budget caps, temporal consistency)

### Soft Constraints

- User preferences (e.g., maximize attraction quality)

## Current Paradigm

Monolithic Agents (ReAct/CoT)

- Uses a single policy to generate the entire trajectory sequentially

### Failure Mode: Constraint Drift

- A single error early in the sequence (e.g., overspending on Day 1) cascades
- This reduces the valid solution space to zero

# Solution 1: Reinforcement Learning



---

## DEEPTRAVEL: AN END-TO-END AGENTIC REINFORCEMENT LEARNING FRAMEWORK FOR AUTONOMOUS TRAVEL PLANNING AGENTS

**Yansong Ning<sup>1\*</sup>, Rui Liu<sup>2</sup>, Jun Wang<sup>2</sup>, Kai Chen<sup>2</sup>, Wei Li<sup>2</sup>, Jun Fang<sup>2</sup>  
Kan Zheng<sup>2</sup>, Naiqiang Tan<sup>2</sup>, Hao Liu<sup>1†</sup>**

<sup>1</sup> The Hong Kong University of Science and Technology (Guangzhou)

<sup>2</sup> Didichuxing Co. Ltd

yning092connect.hkust-gz.edu.cn, liuh@ust.hk  
{invincibleliu, tannaiqiang}@didiglobal.com

# Solution 2: Multi-agent Approach



---

## **ATLAS: Constraints-Aware Multi-Agent Collaboration for Real-World Travel Planning**

**Jihye Choi<sup>1 2 \*</sup>, Jinsung Yoon<sup>1</sup>, Jiefeng Chen<sup>1</sup>, Somesh Jha<sup>2</sup> and Tomas Pfister<sup>1</sup>**

<sup>1</sup>Google Cloud, <sup>2</sup>University of Wisconsin-Madison

# Solution 2: ATLAS



Can you please give me a 7-day itinerary for me and a labmate traveling from San Francisco to Rio de Janeiro from April 21-27 2026? Our combined budget is \$6000, and we would like to try Brazilian, Mexican, and Korean cuisines. We want to explore two nearby cities if possible.

Nothing should be hallucinated. The itinerary should cover all the essentials. Travel route should make sense. I want to have diverse experience.

## ✦ Gemini-2.5-pro

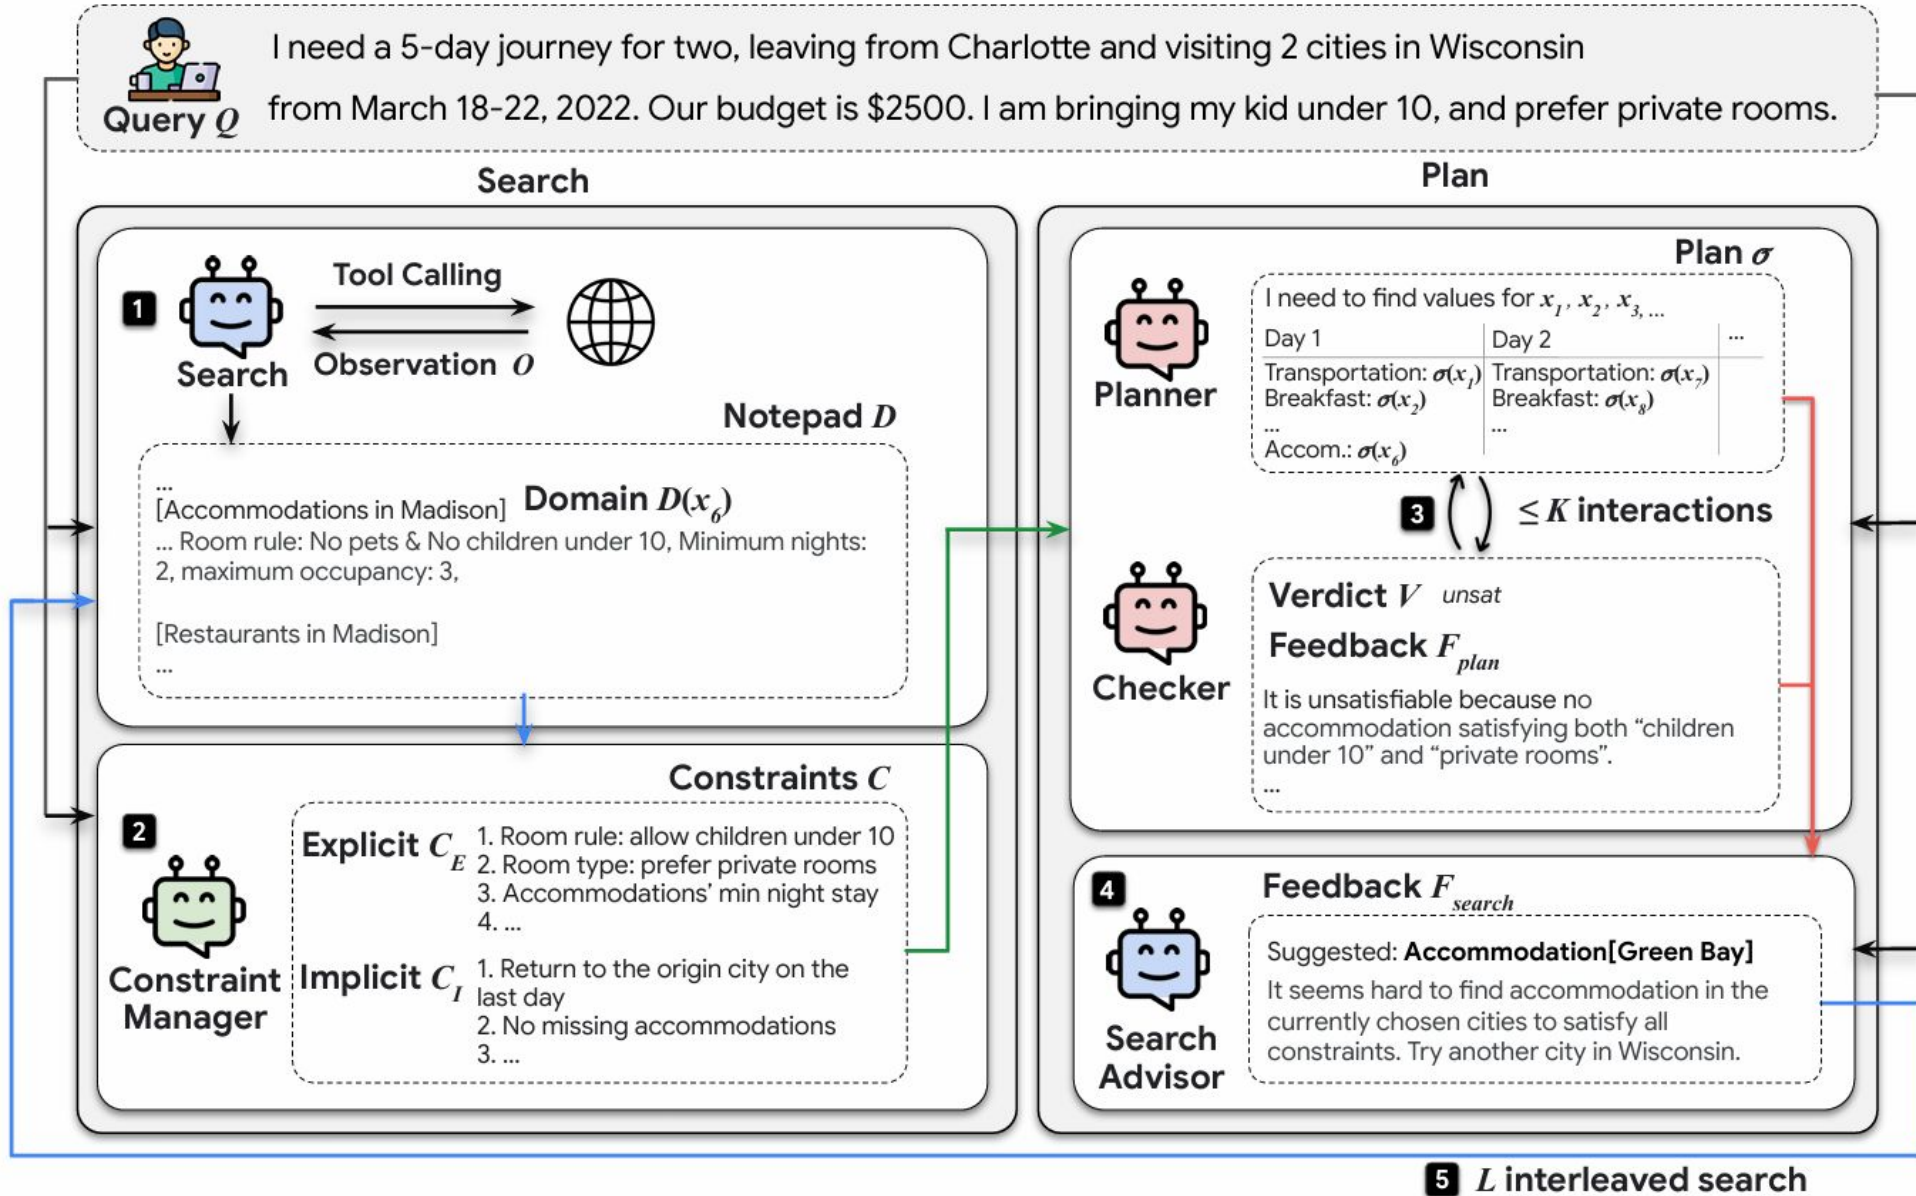
Day 1	<p><b>X</b> <b>Transportation:</b> AV805, Departure: 14:00, Arrival: 09:00, 04-21-2026 <i>Hallucinated</i></p> <p><b>Breakfast:</b> Cultivar</p> <p><b>Attraction:</b> Christ the Redeemer, Rio de Janeiro</p> <p><b>X</b> <b>Lunch:</b> - <i>Missing necessary details</i></p> <p><b>Dinner:</b> Churrascaria Palace</p> <p><b>Accommodation:</b> Hotel Atlantico Prime</p>
Day 2	<p><b>X</b> <b>Transportation:</b> Self-driving, Rio de Janeiro to Petrópolis</p> <p><b>X</b> <b>Breakfast:</b> Cultivar <i>Repeated suggestions</i></p> <p><b>Attraction:</b> Quitandinha Palace; Imperial Museum, Petrópolis</p> <p><b>X</b> <b>Lunch:</b> Puebla Cafe, Rio de Janeiro <i>Invalid city sequence</i></p>
...	...

VS

## 🌐 ATLAS (ours)

Day 1	<p><b>✓</b> <b>Transportation:</b> AA2802, Departure: 13:27, Arrival: 08:20, 04-21-2026</p> <p><b>Breakfast:</b> Cultivar</p> <p><b>Attraction:</b> Christ the Redeemer, Rio de Janeiro</p> <p><b>✓</b> <b>Lunch:</b> Restaurante Guacamole Taqueria</p> <p><b>Dinner:</b> Churrascaria Palace</p> <p><b>Accommodation:</b> Santa Vista Rio</p>
Day 2	<p><b>✓</b> <b>Transportation:</b> Taxi, Rio de Janeiro to Petrópolis</p> <p><b>Breakfast:</b> Galetto 183</p> <p><b>Attraction:</b> Quitandinha Palace; Imperial Museum, Petrópolis</p> <p><b>✓</b> <b>Lunch:</b> Restaurante Alcântara Petrópolis, Petrópolis</p>
...	...

# Solution 2: ATLAS



# Baseline 1: DeepTravel (Yansong Ning et al., 2025)

## Methodology

- Train the agent via **Reinforcement Learning (RL)**
  - Use GRPO method
- **Autonomously plan, execute tools, and refine actions**
- **Maximize rewards over the entire episode**
- **Global Constraints**
  - Trajectory-Level Verifier
- **Local Constraints**
  - Turn-Level Verifier

## Limitations

- **Single agent + Monolithic policy**
- **Sequential Inference**
  - Day 1 → Day 2 → ... → Day 7
- **Large Context Length**
  - Higher training/inference cost & time
  - Requires more GPU memory

# Baseline 2: ATLAS (Ali Behrouz et al., 2025)

## Methodology

- **Multi-agent framework**
- **Search Agent**
  - Interacts with tools
- **Constraint Manager**
  - Handles constraints
- **Planner**
  - Proposes a candidate schedule
- **Checker**
  - Verifies the plan against all constraints
- **Search Advisor**
  - Diagnoses failures and suggests search actions

## Limitations

- **High Resource Cost**
  - 3-day plan = 6 mins inference
  - Excessive loops and context usage
- **Post-hoc Checking Loop**
  - Relies heavily on LLM quality
  - Not trainable/fine-tunable
  - More trial-and-error context length
- **Sequential Bottlenecks**
  - Sequential Planner-Checker loop

# Baseline Limitations Comparison

## DeepTravel

Ning et al., 2025

### Single Agent Bottleneck

- Relies on a monolithic policy for the entire multi-day trajectory.

### Sequential Inference

- Generates Day 1 to Day 7 in order; early errors cascade downstream.

### Constraint Drift

- As context grows, attention to global constraints (like budget) dilutes.

### High Memory Usage

- Long tool traces saturate the context window, requiring significant GPU resources.

## ATLAS

Behrouz et al., 2025

### High Computational Cost

- Verification loops are resource-heavy (e.g., 6 mins for a 3-day plan).

### Post-hoc Checking Loop

- Repairs plans after generation rather than enforcing constraints by construction.

### Sequential Bottlenecks

- Iterative Planner-Checker loops increase latency significantly.

### Fixed Architecture

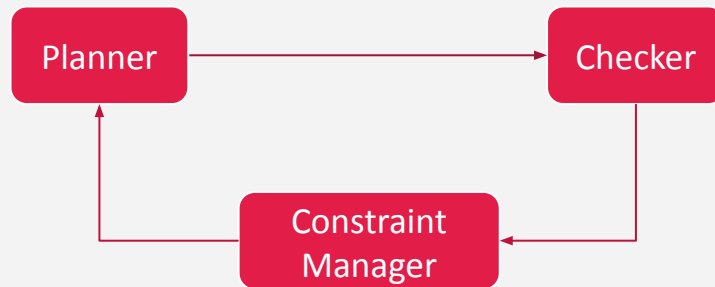
- Typically not fine-tunable, relying heavily on the base LLM quality.

# Baseline Limitations

## ATLAS: Post-hoc Latency

Behrouz et al., 2025

RECURSIVE



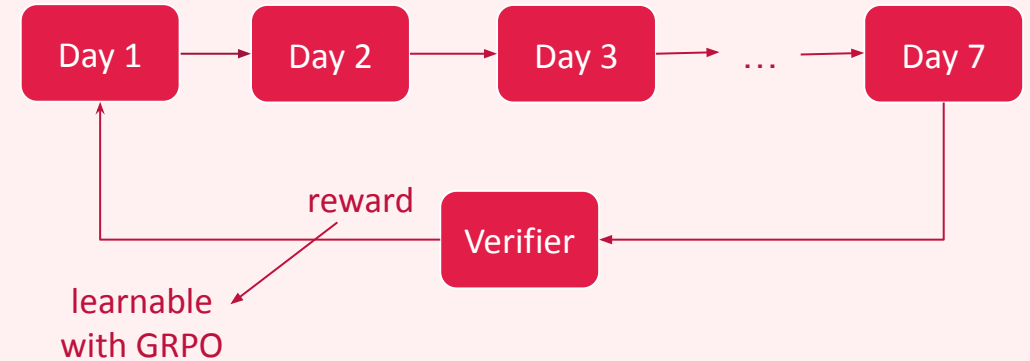
### Post-hoc Checking Loop

- Relies heavily on LLM quality
- Not trainable/fine-tunable
- More trial-and-error context length

## DeepTravel: Monolithic Drift

Ning et al., 2025

SEQUENTIAL



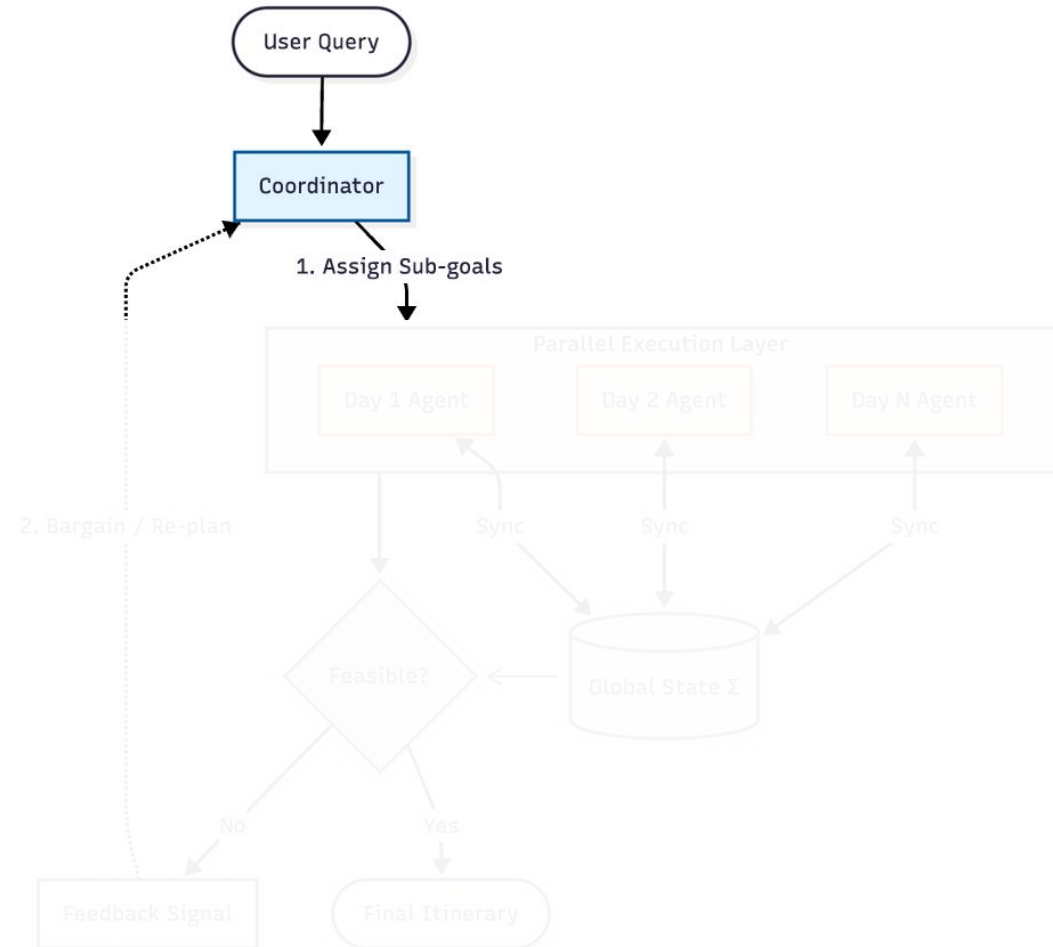
### Large Context Length

- Higher training/inference cost & time
- Requires more GPU memory

# Our Solution: HiMAP-Travel

## Methodology

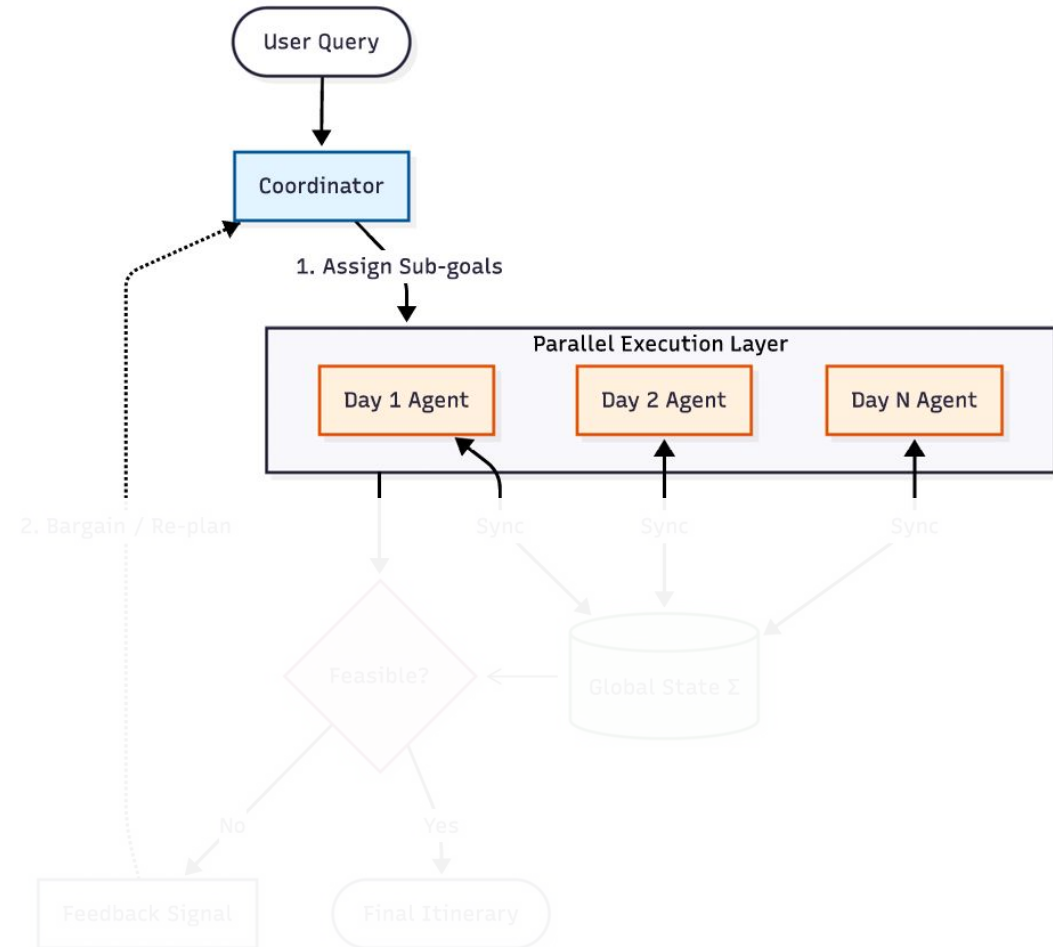
- **MARL + GRPO**
- **Coordinator**
  - Extract query information
  - Manage constraints
  - Distribute tasks into Executors



# Our Solution: HiMAP-Travel

## Methodology

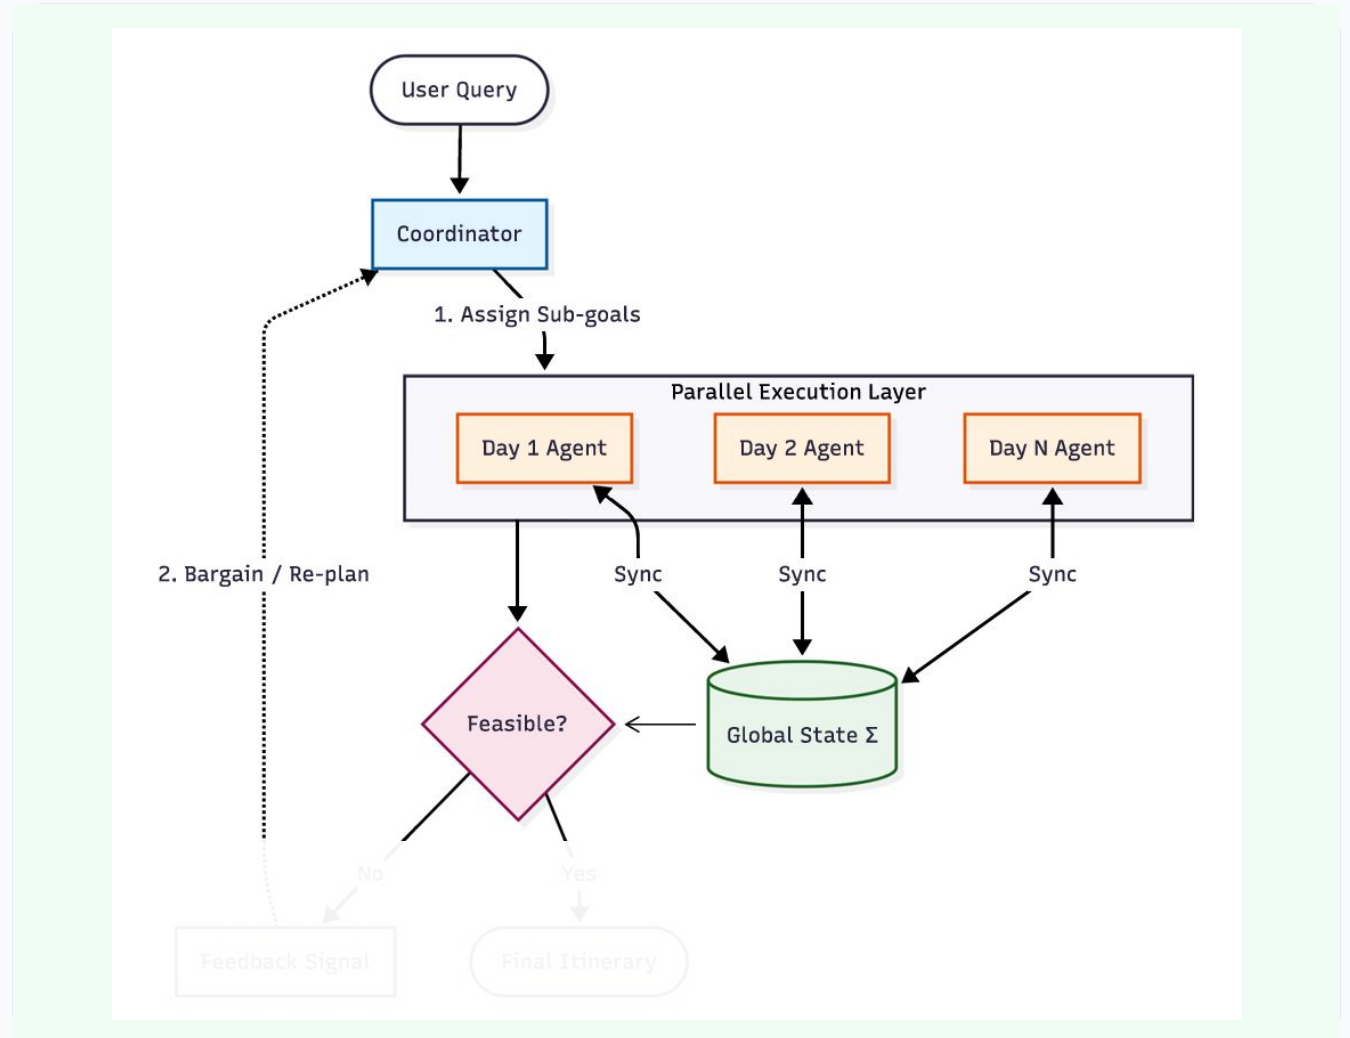
- **MARL + GRPO**
- **Coordinator**
- **Executors**
  - Departure Agent
  - Stay Agent
  - Travel Agent
  - Return Agent



# Our Solution: HiMAP-Travel

## Methodology

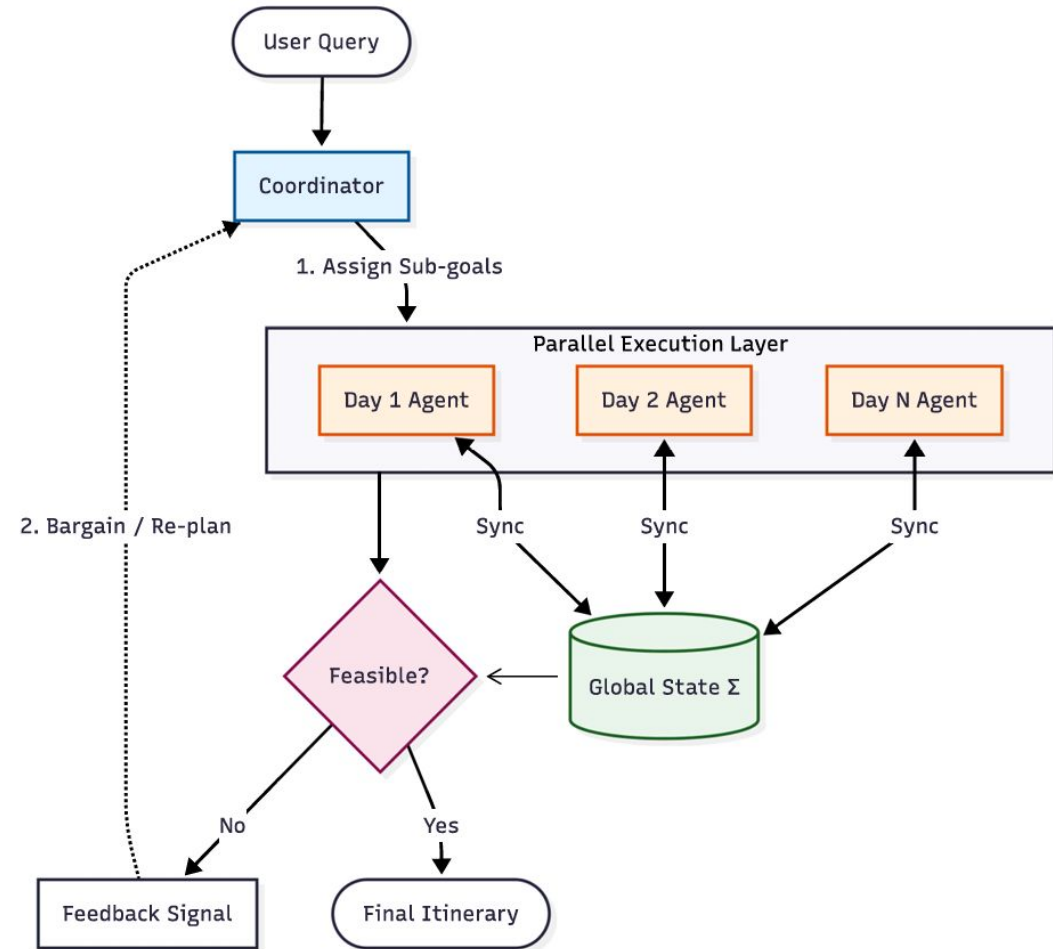
- MARL + GRPO
- Coordinator
- Executors
- Synchronized Global State
  - `check(action)`
  - `commit(action)`
  - `checkpoint()`
  - `rollback()`



# Our Solution: HiMAP-Travel

## Methodology

- **MARL + GRPO**
- **Coordinator**
- **Executors**
- **Synchronized Global State**
- **Cooperative Bargaining Protocol**
  - Feedback to Coordinator
  - Re-plan if infeasible
  - Reward calculation



We are a group of 5 looking for a 7-day travel itinerary, departing from Sault Ste. Marie towards Michigan. We plan to visit 3 different cities from March 15th to March 21st, 2022. Our budget allows us to spend up to \$12,700. We require entire rooms for our stay. Could you help us with this plan?

Coordinator Agent

Step 1: `information_extractor()`

- Origin: Sault Ste. Marie
- Destination: Michigan
- Duration: 7 days
- Start date: 2022-03-15
- Cities to visit: 3
- Number of people: 5
- Budget: \$12700

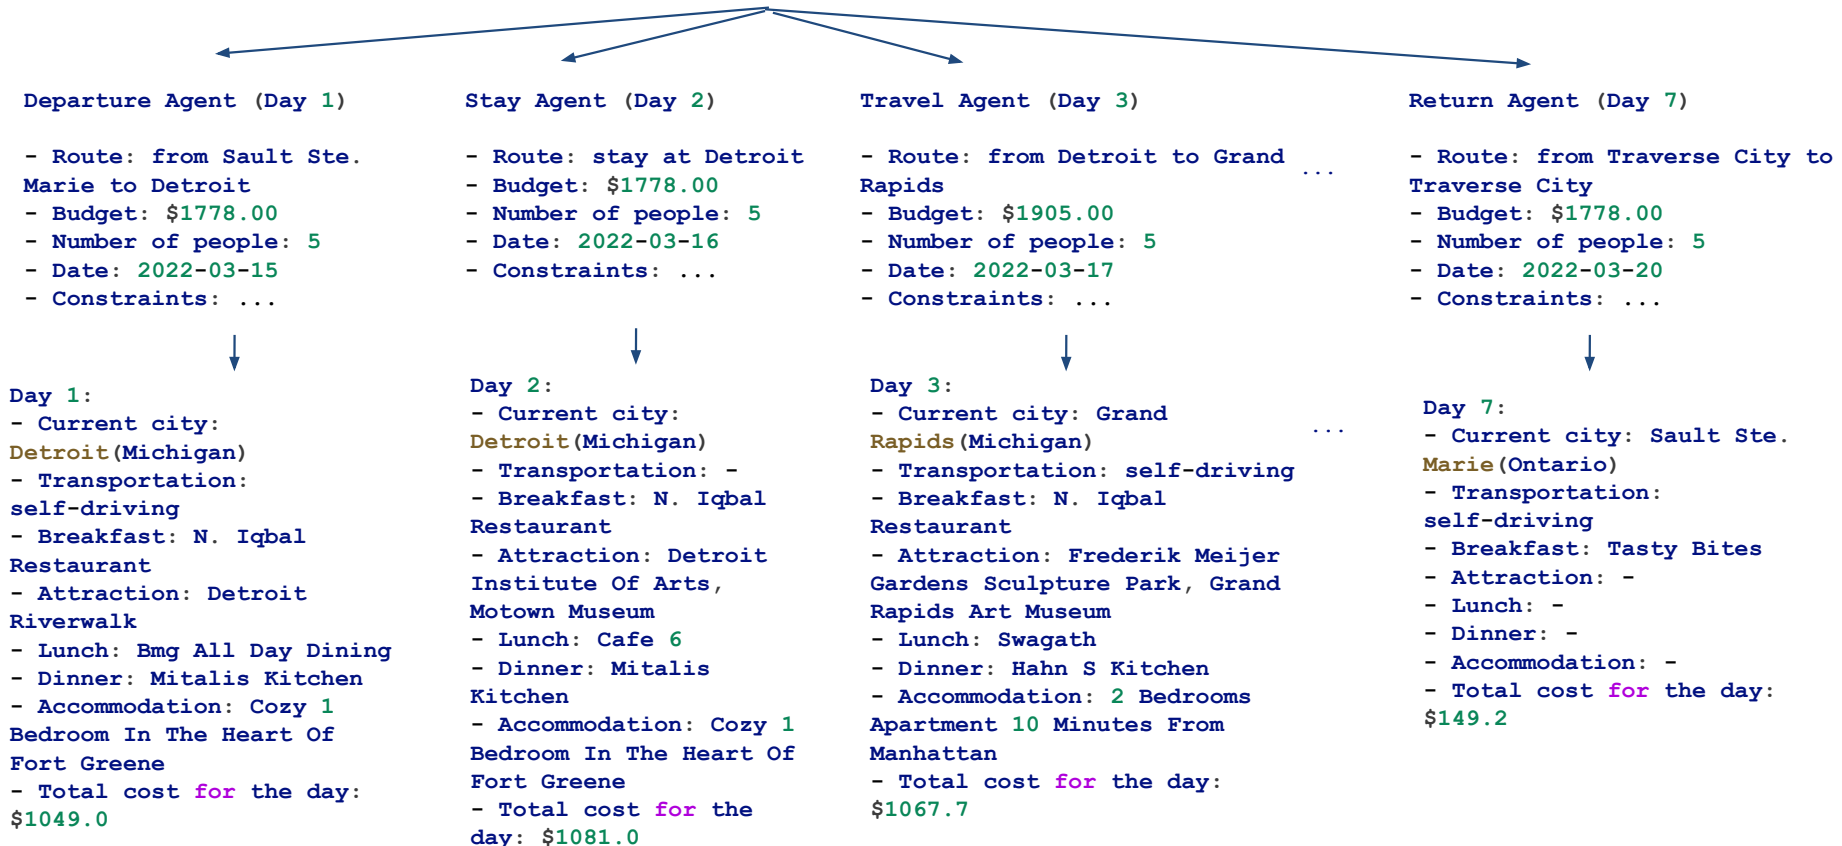
- Constraints:

- \* House rules: entire room
- \* Cuisine preferences: -
- \* Room type: entire room
- \* Transportation: -

Step 2: `city_search(Michigan)`: Alpena, Detroit, Escanaba, ...

Step 3: `distribute_task()`

- `visiting_cities`: Detroit, Grand Rapids, Traverse City
- `budget_rate_per_day`: \$1778, \$1778, \$1905, ...



# Experimental Results

Method	Backbone	Test FPR (Final Pass Rate)
MTP	GPT-4	0.60%
ATLAS	Gemini-2.5-Pro	35.00%
DeepTravel	Qwen3-8B	43.98%

# Experimental Results

Method	Backbone	Test FPR (Final Pass Rate)
MTP	GPT-4	0.60%
ATLAS	Gemini-2.5-Pro	35.00%
DeepTravel	Qwen3-8B	43.98%
<b>HiMAP-Travel</b>	Qwen3-8B	<b>52.65%</b>

# Experimental Results

Method	Backbone	Test FPR (Final Pass Rate)
ATLAS	Gemini-2.5-Pro	35.00%
DeepTravel	Qwen3-8B	43.98%
<b>HiMAP-Travel</b>	Qwen3-8B	<b>52.65%</b>

Test Final Pass Rate (FPR)

**52.65%**

Outperforms ATLAS by +17.65%

Reliability

**93%**

Variance reduction vs. DeepTravel

Efficiency

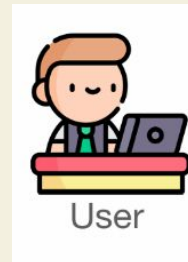
**2.5x**

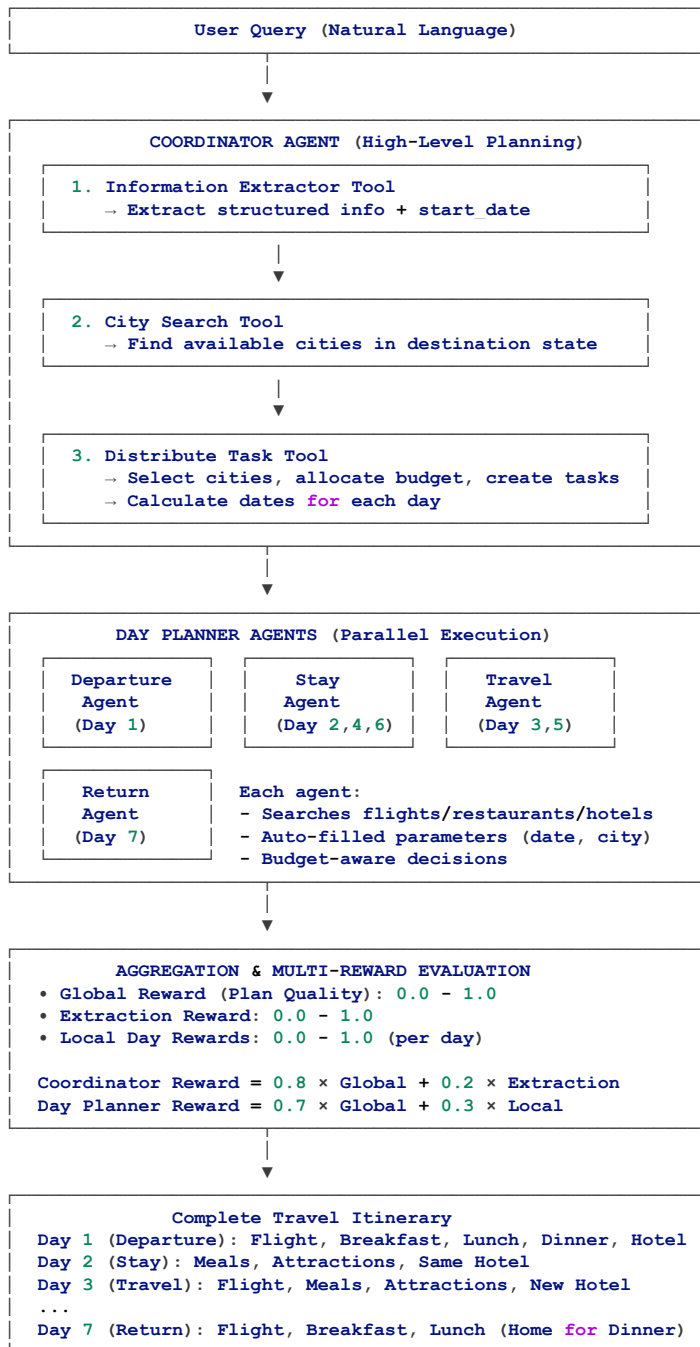
Latency reduction and GPU usage

# Conclusion

## Key Takeaways

- HiMAP-Travel successfully mitigates **Constraint Drift** in long-horizon planning.
- The **Synchronized Global State** ensures deterministic safety for shared resources like budgets and venues.
- Achieved state-of-the-art **52.65% Test FPR** on TravelPlanner, outperforming sequential baselines.
- Parallel execution enables a **2.5x latency reduction**, making real-world deployment feasible.





We are a group of 5 looking for a 7-day travel itinerary, departing from Sault Ste. Marie towards Michigan. We plan to visit 3 different cities from March 15th to March 21st, 2022. Our budget allows us to spend up to \$12,700. We require entire rooms for our stay. Could you help us with this plan?

Coordinator Agent

Step 1: `information_extractor()`

- Origin: Sault Ste. Marie
  - Destination: Michigan
  - Duration: 7 days
  - Start date: 2022-03-15
  - Cities to visit: 3
  - Number of people: 5
  - Budget: \$12700
- Constraints:
    - \* House rules: entire room
    - \* Cuisine preferences: -
    - \* Room type: entire room
    - \* Transportation: -

Step 2: `city_search(Michigan)`: Alpena, Detroit, Escanaba, ...

Step 3: `distribute_task()`

- `visiting_cities`: Detroit, Grand Rapids, Traverse City
- `budget_rate_per_day`: \$1778, \$1778, \$1905, ...

Departure Agent (Day 1)

- Route: from Sault Ste. Marie to Detroit
- Budget: \$1778.00
- Number of people: 5
- Date: 2022-03-15
- Constraints: ...

Stay Agent (Day 2)

- Route: stay at Detroit
- Budget: \$1778.00
- Number of people: 5
- Date: 2022-03-16
- Constraints: ...

Travel Agent (Day 3)

- Route: from Detroit to Grand Rapids
- Budget: \$1905.00
- Number of people: 5
- Date: 2022-03-17
- Constraints: ...

Return Agent (Day 7)

- Route: from Traverse City to Traverse City
- Budget: \$1778.00
- Number of people: 5
- Date: 2022-03-20
- Constraints: ...

Day 1:

- Current city: Detroit (Michigan)
- Transportation: self-driving
- Breakfast: N. Iqbal Restaurant
- Attraction: Detroit Riverwalk
- Lunch: Bmg All Day Dining
- Dinner: Mitalis Kitchen
- Accommodation: Cozy 1 Bedroom In The Heart Of Fort Greene
- Total cost for the day: \$1049.0

Day 2:

- Current city: Detroit (Michigan)
- Transportation: -
- Breakfast: N. Iqbal Restaurant
- Attraction: Detroit Institute Of Arts, Motown Museum
- Lunch: Cafe 6
- Dinner: Mitalis Kitchen
- Accommodation: Cozy 1 Bedroom In The Heart Of Fort Greene
- Total cost for the day: \$1081.0

Day 3:

- Current city: Grand Rapids (Michigan)
- Transportation: self-driving
- Breakfast: N. Iqbal Restaurant
- Attraction: Frederik Meijer Gardens Sculpture Park, Grand Rapids Art Museum
- Lunch: Swagath
- Dinner: Hahn S Kitchen
- Accommodation: 2 Bedrooms Apartment 10 Minutes From Manhattan
- Total cost for the day: \$1067.7

Day 7:

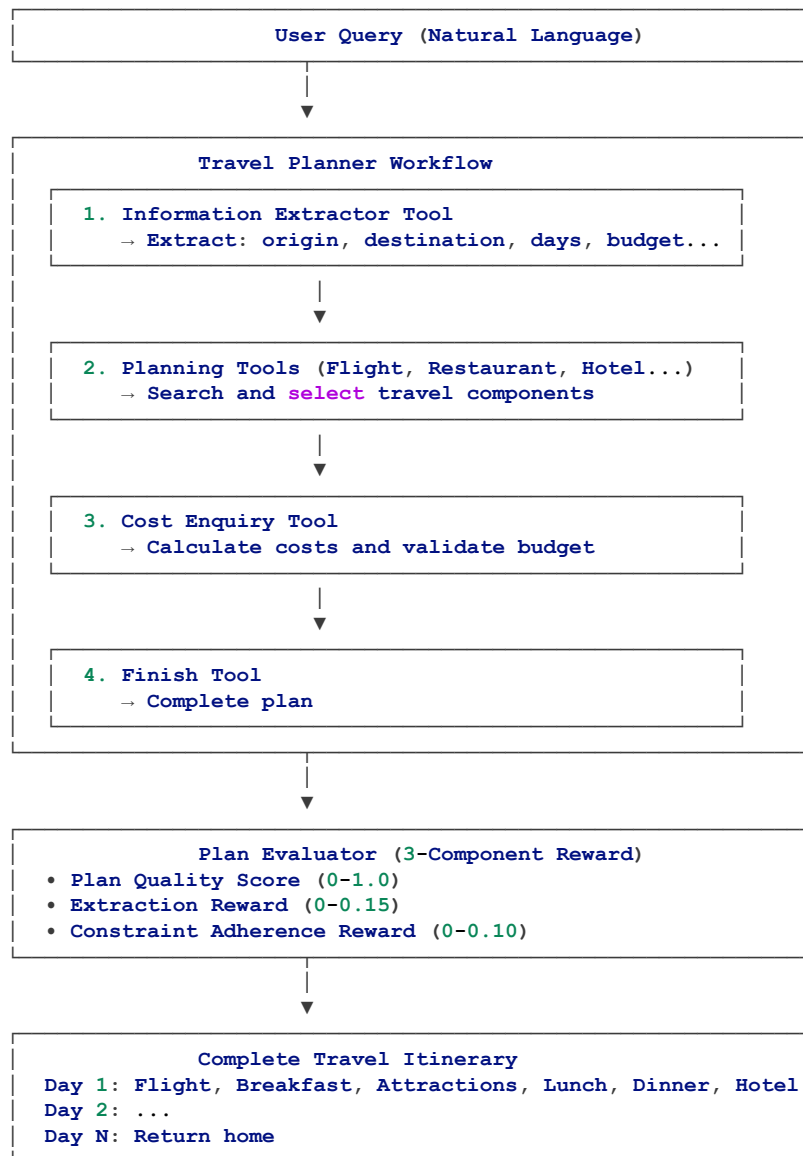
- Current city: Sault Ste. Marie (Ontario)
- Transportation: self-driving
- Breakfast: Tasty Bites
- Attraction: -
- Lunch: -
- Dinner: -
- Accommodation: -
- Total cost for the day: \$149.2

Complete Travel Itinerary

- Day 1 (Departure): Flight, Breakfast, Lunch, Dinner, Hotel
- Day 2 (Stay): Meals, Attractions, Same Hotel
- Day 3 (Travel): Flight, Meals, Attractions, New Hotel
- ...
- Day 7 (Return): Flight, Breakfast, Lunch (Home for Dinner)

**V1 vs. V2**

# V1 Architecture



Can you generate a 5-day travel itinerary for a group of 4, departing from Myrtle Beach and planning to visit 2 cities in Tennessee? The trip is scheduled from March 14th to March 18th, 2022, and we have allocated a budget of \$5,500. We would prefer to stay in private rooms during our accommodations.

Step 1: Information Extractor:

- Origin: Myrtle Beach
- Destination: Tennessee
- Duration: 5 days
- Cities to visit: 2
- Number of people: 4
- Budget: \$5500
- Travel dates: 2022-03-14 to 2022-03-18
- Room requirement: private room
- ...

Step 2.1: `city_search(Tennessee): Bristol, Chattanooga, Knoxville, Nashville, ...`

Step 2.2: `accommodation_search(Chattanooga, Tennessee)`  
`accommodation_search(Nashville, Tennessee)`  
...

Step 2.3: `restaurant_search(Chattanooga, Tennessee)`  
`restaurant_search(Nashville, Tennessee)`  
...

Step 2.4: `flight_search(Myrtle Beach, South Carolina, Chattanooga, Tennessee)`  
`distance_search(Chattanooga, Tennessee, Nashville, Tennessee)`  
...

Step 3.1: `cost_enquiry(...)`

Day 1:

- Current city: Chattanooga(Tennessee)
- Transportation: self-driving
- Breakfast: Ichiban
- Attraction: Ruby Falls, Tennessee Aquarium
- Lunch: Ichiban
- Dinner: Ichiban
- Accommodation: Sunny Room Pvrte Office In Huge Loft
- Total cost for the day: \$2538.6

Step 3.2: `cost_enquiry(...)`

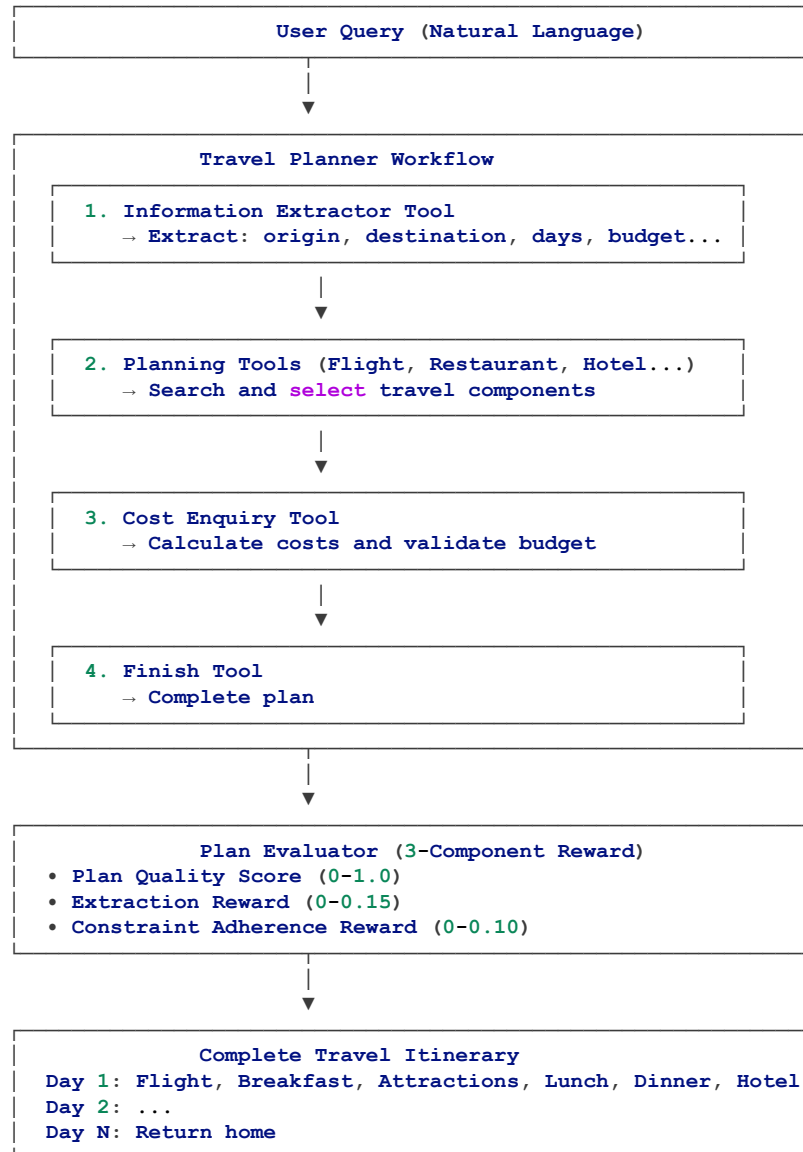
Day 2:

- Current city: Chattanooga(Tennessee)
- Transportation: -
- Breakfast: Ichiban
- Attraction: Ruby Falls, Tennessee Aquarium
- Lunch: Ichiban
- Dinner: Ichiban
- Accommodation: Sunny Room Pvrte Office In Huge Loft
- Total cost for the day: \$2500.0

...

Step 4: `finish()`

# Plan Evaluation



## Hard Constraints:

Raw score: 60.00/100

Issues (4):

- ✗ Budget exceeded: \$9358.9 > \$5500 (+\$3858.9, +70.2%)
- ✓ Days: 5 = 5
- ✓ People: 4 consistent across all days
- ✓ Cities: 2 = 2

## Field Completeness:

Raw score: 100.00/100

Issues (4):

- Day 2: ✓ Stay day - no transport needed (staying in same city)
- Day 4: ✓ Stay day - no transport needed (staying in same city)
- Day 5: ✓ Last day - returning home (100% valid in training)
- Day 5: ✓ Last day - attraction dash acceptable (62.2% in training)

## Transportation Logic:

Raw score: 100.00/100

✓ No issues found - Enhanced evaluation passed

## Location Consistency:

Raw score: 100.00/100

Issues (3):

- Day 5: Breakfast in departure city Nashville(Tennessee) - travel day pattern (acceptable)
- Day 5: Lunch in departure city Nashville(Tennessee) - travel day pattern (acceptable)
- Day 5: Dinner in departure city Nashville(Tennessee) - travel day pattern (acceptable)

## Meal Coverage:

Raw score: 100.00/100

✓ No issues found - Enhanced evaluation passed

## Attraction Quality:

Raw score: 100.00/100

✓ No issues found - Enhanced evaluation passed

## Accommodation Logic:

Raw score: 90.00/100

Issues (2):

- Day 1: Minor accommodation capacity shortage (2 vs 4 people)
- Day 2: Minor accommodation capacity shortage (2 vs 4 people)

## Cost Accuracy:

Raw score: 60.00/100

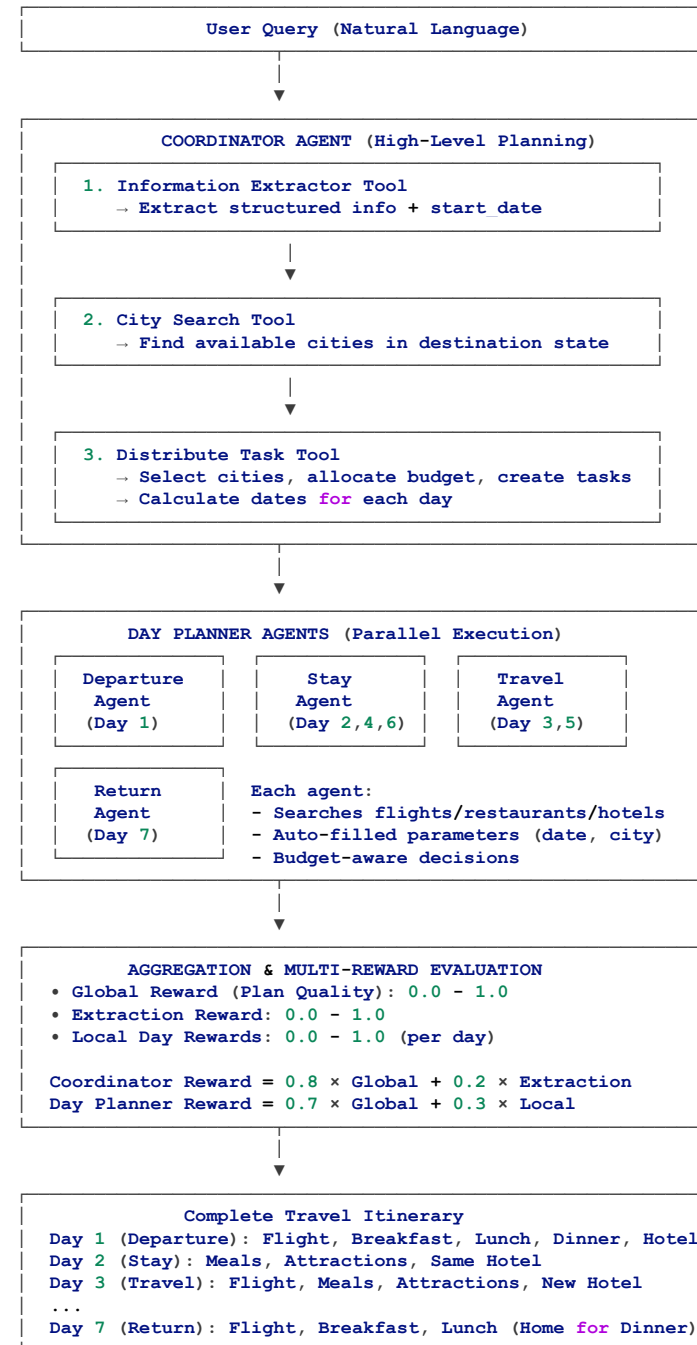
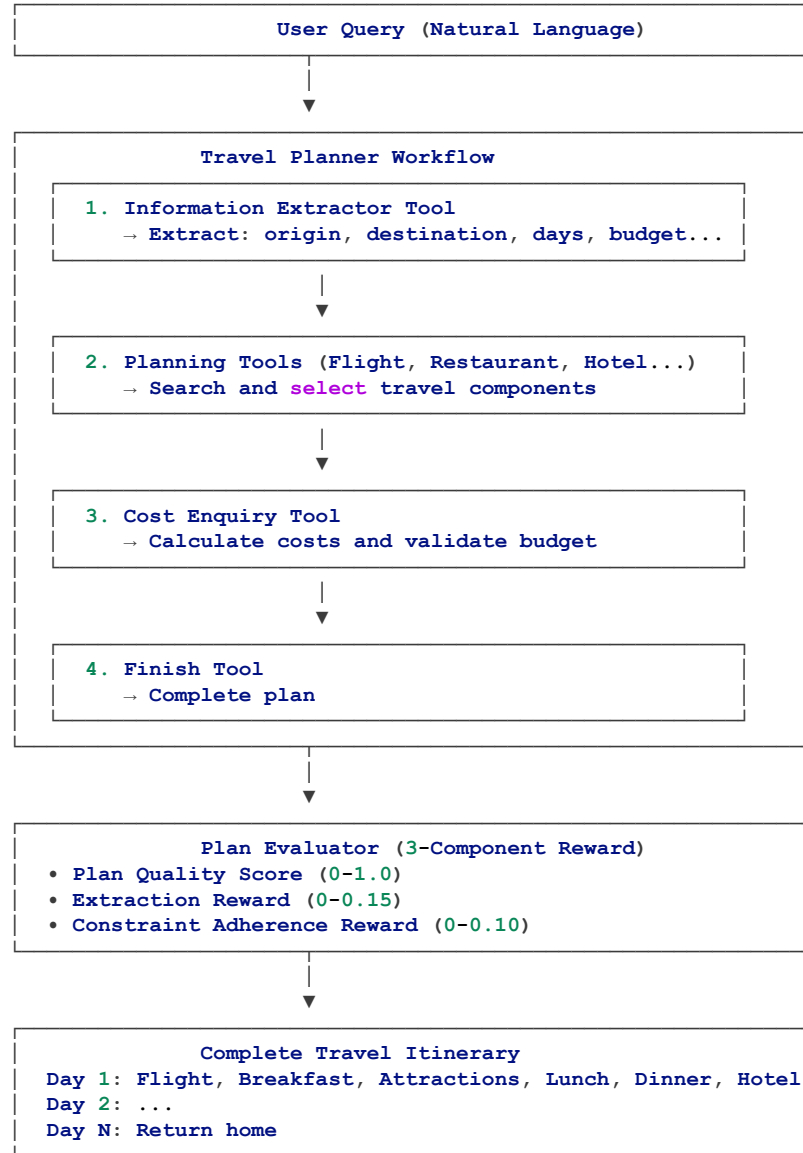
Issues (1):

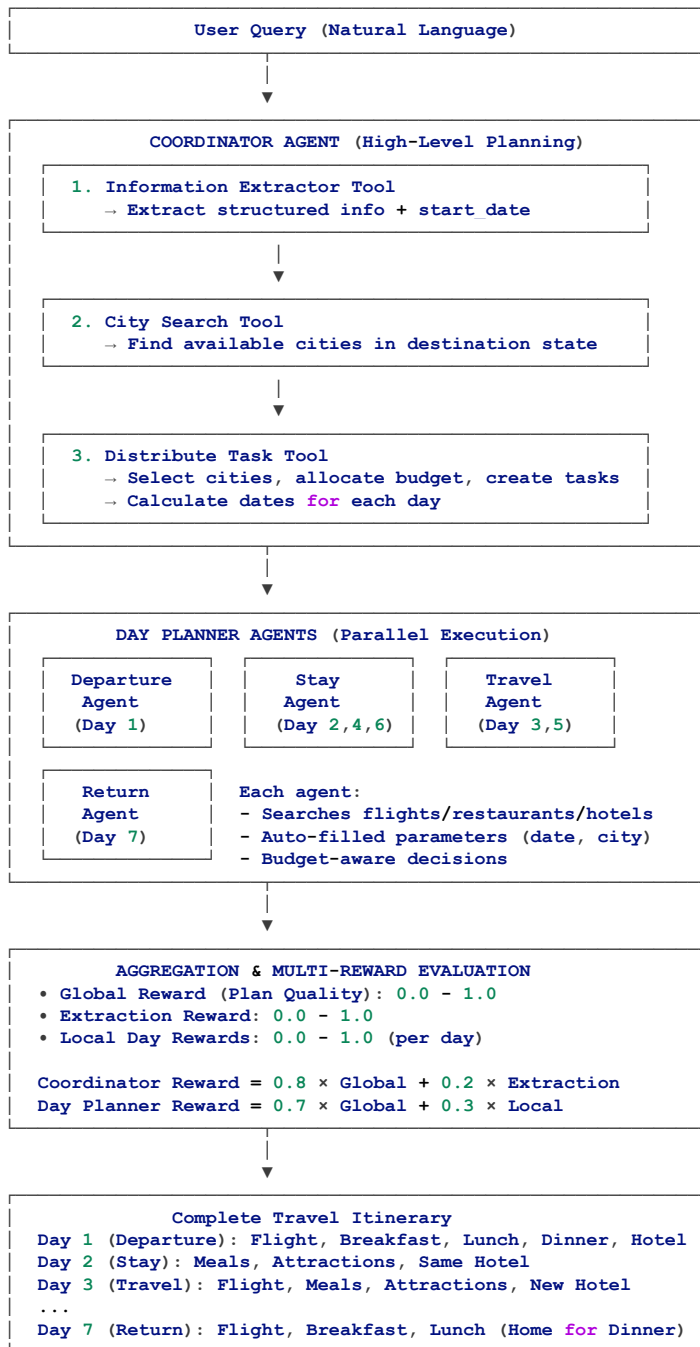
- Budget exceeded: \$9358.90 / \$5500.00 = 170.16% > 100% (penalty: -40.0)

## Extraction Evaluation:

✓ No issues found - Enhanced evaluation passed

# v1 vs v2 Architecture





We are a group of 5 looking for a 7-day travel itinerary, departing from Sault Ste. Marie towards Michigan. We plan to visit 3 different cities from March 15th to March 21st, 2022. Our budget allows us to spend up to \$12,700. We require entire rooms for our stay. Could you help us with this plan?

Coordinator Agent

Step 1: `information_extractor()`

- Origin: Sault Ste. Marie
  - Destination: Michigan
  - Duration: 7 days
  - Start date: 2022-03-15
  - Cities to visit: 3
  - Number of people: 5
  - Budget: \$12700
- Constraints:
    - \* House rules: entire room
    - \* Cuisine preferences: -
    - \* Room type: entire room
    - \* Transportation: -

Step 2: `city_search(Michigan)`: Alpena, Detroit, Escanaba, ...

Step 3: `distribute_task()`

- `visiting_cities`: Detroit, Grand Rapids, Traverse City
- `budget_rate_per_day`: \$1778, \$1778, \$1905, ...

Departure Agent (Day 1)

- Route: from Sault Ste. Marie to Detroit
- Budget: \$1778.00
- Number of people: 5
- Date: 2022-03-15
- Constraints: ...

Stay Agent (Day 2)

- Route: stay at Detroit
- Budget: \$1778.00
- Number of people: 5
- Date: 2022-03-16
- Constraints: ...

Travel Agent (Day 3)

- Route: from Detroit to Grand Rapids
- Budget: \$1905.00
- Number of people: 5
- Date: 2022-03-17
- Constraints: ...

Return Agent (Day 7)

- Route: from Traverse City to Traverse City
- Budget: \$1778.00
- Number of people: 5
- Date: 2022-03-20
- Constraints: ...

Day 1:

- Current city: Detroit (Michigan)
- Transportation: self-driving
- Breakfast: N. Iqbal Restaurant
- Attraction: Detroit Riverwalk
- Lunch: Bmg All Day Dining
- Dinner: Mitalis Kitchen
- Accommodation: Cozy 1 Bedroom In The Heart Of Fort Greene
- Total cost for the day: \$1049.0

Day 2:

- Current city: Detroit (Michigan)
- Transportation: -
- Breakfast: N. Iqbal Restaurant
- Attraction: Detroit Institute Of Arts, Motown Museum
- Lunch: Cafe 6
- Dinner: Mitalis Kitchen
- Accommodation: Cozy 1 Bedroom In The Heart Of Fort Greene
- Total cost for the day: \$1081.0

Day 3:

- Current city: Grand Rapids (Michigan)
- Transportation: self-driving
- Breakfast: N. Iqbal Restaurant
- Attraction: Frederik Meijer Gardens Sculpture Park, Grand Rapids Art Museum
- Lunch: Swagath
- Dinner: Hahn S Kitchen
- Accommodation: 2 Bedrooms Apartment 10 Minutes From Manhattan
- Total cost for the day: \$1067.7

Day 7:

- Current city: Sault Ste. Marie (Ontario)
- Transportation: self-driving
- Breakfast: Tasty Bites
- Attraction: -
- Lunch: -
- Dinner: -
- Accommodation: -
- Total cost for the day: \$149.2

User Query (Natural Language)

COORDINATOR AGENT (High-Level Planning)

1. Information Extractor Tool
  - Extract structured info + start date

2. City Search Tool
  - Find available cities in destination state

3. Distribute Task Tool
  - Select cities, allocate budget, create tasks
  - Calculate dates for each day

DAY PLANNER AGENTS (Parallel Execution)

- Departure Agent (Day 1)
- Stay Agent (Day 2,4,6)
- Travel Agent (Day 3,5)

- Return Agent (Day 7)
- Each agent:
  - Searches flights/restaurants/hotels
  - Auto-filled parameters (date, city)
  - Budget-aware decisions

AGGREGATION & MULTI-REWARD EVALUATION

- Global Reward (Plan Quality): 0.0 - 1.0
  - Extraction Reward: 0.0 - 1.0
  - Local Day Rewards: 0.0 - 1.0 (per day)
- Coordinator Reward = 0.8 × Global + 0.2 × Extraction
- Day Planner Reward = 0.7 × Global + 0.3 × Local

Complete Travel Itinerary

- Day 1 (Departure): Flight, Breakfast, Lunch, Dinner, Hotel
- Day 2 (Stay): Meals, Attractions, Same Hotel
- Day 3 (Travel): Flight, Meals, Attractions, New Hotel
- ...
- Day 7 (Return): Flight, Breakfast, Lunch (Home for Dinner)

**Questions?**