

ARGUS AUDIT REPORT

SESSION ID: ARGUS-S-2025-12-28-6D3A

DATE: 2025-12-28T20:33:31.164Z

This document constitutes a rigorous adversarial critique.

All claims have been subjected to multi-agent simulation.

CLAIM 1: ACCEPTED

"Carbon-aware scheduling of flexible AI workloads, which can tolerate delays of 2-24 hours, can reduce carbon emissions by 6-18% with a cost increase of less than 5% in regions with volatile carbon intensity like the UK and California."

[THESIS_DESTROYER]

```
{  
  "attack": "The claim asserts a cost increase of less than 5%. What specific costs are being considered? Does this include the cost of implementing the carbon-aware scheduling system, the cost of monitoring carbon intensity, and potential costs associated with delayed AI workload completion if delays impact downstream processes or SLAs? A more comprehensive accounting might reveal the cost increase exceeds 5% in some scenarios."  
}
```

[JOURNAL_REVIEWER_SIMULATOR]

```
{  
  "verdict": "ACCEPTED",  
  "fatal": false,  
  "noveltyClassification": [],  
  "reasons": [  
    "The original claim includes specific caveats (regions, delay tolerance, cost increase under 5%). The adversary questions the scope of cost accounting, which is an appropriate challenge but does not invalidate the claim given its hedged nature. The claim acknowledges a cost increase; the adversary seeks clarification on cost components, not demonstrating the claim is false."  
  ]  
}
```

CLAIM 2: ACCEPTED

"Optimizing AI workload scheduling for cost leads to a 15-25% increase in carbon emissions, while optimizing for carbon minimization results in an 8-12% increase in cost due to regional energy pricing disparities."

[THESIS_DESTROYER]

```
{  
  "attack": "Can you provide data showing a scenario where optimizing for cost  
*decreases* carbon emissions, or optimizing for carbon *decreases* cost? For  
example, if a 'green' region has unexpectedly low prices due to oversupply of  
renewables during off-peak hours, wouldn't carbon minimization and cost  
optimization align?"  
}
```

[JOURNAL_REVIEWER_SIMULATOR]

```
{  
  "verdict": "ACCEPTED",  
  "FAILURE_TAGS": [],  
  "REVISION_PATH": "N/A",  
  "JUSTIFICATION": "The claim includes specific ranges (15-25% and 8-12%),  
indicating a level of caution. The adversary's request for scenarios where the  
trend reverses doesn't invalidate the core claim; it highlights edge cases which  
are acceptable given the bounded claim."  
}
```

CLAIM 3: REVISE

"The effectiveness of carbon-aware scheduling for AI workloads depends on the volatility of the grid's carbon intensity and the flexibility of the workloads themselves."

[THESIS_DESTROYER]

```
{  
  "attack": "If carbon-aware scheduling depends on BOTH volatility and flexibility,  
does that imply that high flexibility can compensate for low volatility, or vice  
versa? Provide empirical evidence showing the degree to which one can substitute  
for the other. Otherwise, the claim is unfalsifiable because any result can be  
rationalized after the fact."  
}
```

```
[JOURNAL_REVIEWER_SIMULATOR]
{
  "verdict": "REVISE",
  "fatal": false,
  "noveltyClassification": [
    "VAGUE CLAIM"
  ],
  "reasons": [
    "The adversary correctly identifies the claim's ambiguity regarding the interplay between volatility and flexibility, rendering it currently unfalsifiable. The degree to which one can compensate for the other is unclear."
  ]
}
```

CLAIM 4: REVISE

"Integer linear programming can be used to model and optimize the trade-offs between minimizing CO2 emissions and minimizing cost when scheduling flexible AI workloads subject to SLA constraints."

```
[THESIS_DESTROYER]
{
  "attack": "Does 'optimizing trade-offs' guarantee a Pareto-optimal solution, or just *an* improved solution? If it's the latter, can you provide a counterexample where the ILP approach converges to a sub-optimal solution compared to other scheduling heuristics, even within the SLA constraints? Specifically, can an ILP model get stuck in a local optimum, failing to identify a superior schedule that simultaneously reduces both CO2 emissions and cost?"
}
```

```
[JOURNAL_REVIEWER_SIMULATOR]
{
  "verdict": "REVISE",
  "FAILURE_TAGS": [
    "VAGUE CLAIM",
    "LACK_OF_PRECISION"
  ],
  "REVISION_PATH": [
    "Clarify the meaning of 'optimizing trade-offs'.",
    "Specify whether the approach guarantees a Pareto-optimal solution.",
    "Address the possibility of the ILP model converging to a sub-optimal solution.",
    "Provide evidence or justification for the optimality of the solution."
  ],
  "JUSTIFICATION": "The claim lacks precision regarding the optimality of the solution obtained using Integer Linear Programming. The adversary raises a valid concern about potential sub-optimality, which needs to be addressed through clarification or further analysis."
}
```

CLAIM 5: ACCEPTED

"Governance mechanisms such as KPIs, SLA windows, and carbon budgets are crucial for the successful enterprise adoption of carbon-aware AI workload scheduling, but their efficacy hinges on the organization's specific carbon reduction priorities."

[THESIS_DESTROYER]

```
{  
  "thesis_attack": "If the organization's carbon reduction priorities are poorly defined or internally contradictory, wouldn't these governance mechanisms actually *hinder* the adoption of carbon-aware AI workload scheduling, by providing a false sense of control and accountability while masking actual increases in carbon emissions? For instance, a KPI focused solely on energy consumption might incentivize shifting workloads to regions with dirtier grids, thereby reducing local energy use but increasing overall carbon footprint. This contradicts the implied goal of *reducing* carbon emissions."  
}
```

[JOURNAL_REVIEWER_SIMULATOR]

```
{  
  "verdict": "ACCEPTED",  
  "fatal": false,  
  "noveltyClassification": [],  
  "reasons": [  
    "The claim is appropriately hedged, acknowledging the dependency on specific carbon reduction priorities. The adversarial argument, while valid, highlights a potential failure mode *within* that dependency, not a fundamental flaw in the claim itself. The caveat protects the claim from this specific attack."  
  ]  
}
```