Article

Not peer-reviewed version

From Traffic Indices to Agentic AI: A Reference Framework for Urban Mobility in Megacities

Tatiana Petrova *

Posted Date: 20 April 2026

doi: 10.20944/preprints202604.1407.v1

Keywords: Agentic AI; multi-agent systems; reference architecture; urban mobility; transport accessibility; PTAL; composite traffic index; smart cities; London

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 20 April 2026 doi:10.20944/preprints202604.1407.v1

Article

From Traffic Indices to Agentic AI: A Reference Framework for Urban Mobility in Megacities

Tatiana Petrova

SnT, University of Luxembourg, Luxembourg; tatiana.petrova@uni.lu

Abstract

Megacity mobility research has long relied on aggregate statistical indicators — composite traffic-planning indices and accessibility surfaces — that capture the system at rest but are disconnected from the operational decisions shaping mobility minute to minute. In parallel, *Agentic AI*, articulated in the *Web of Agents* narrative of [1], offers a reference paradigm of semantically interoperable autonomous agents that negotiate and coordinate across open networks. We propose a unifying *Agentic AI reference architecture for urban transportation* that maps any composite traffic index and any accessibility surface onto agent utility functions, negotiation-protocol primitives, and shared semantic ontologies. The architecture is instantiated *in simulation only*: no live A2A/MCP endpoints, no runtime LLMs, no cross-organisation interoperability experiment; these are explicitly listed as next steps. A 12-borough London case study, evaluated over $N = 30$ seeds on an origin-destination matrix calibrated to the 2021 UK Census commuter-flow aggregates [2], benchmarks four regimes (historic, adaptive, MaxPressure, and our Agentic policy) across four scenarios covering equity, corridor prioritisation, incident response, and their combination. The Agentic regime reduces the accessibility-deficit Gini coefficient by **23–58 %**, travel-time CV by up to **41 %** and mean travel time by **4–9 %** relative to the historic baseline, and is the one regime in our four-way comparison that sits near the per-column best on all three objectives simultaneously when equity and incident-response interact. A microscopic SUMO testbed on a 4×4 signalised grid (120 runs across three demand regimes and five peripheral-boost values) traces an explicit equity–efficiency Pareto frontier on which the agentic policy matches or beats a SCOOT-style adaptive controller on mean travel time and throughput at every operating point, while reducing travel-time variance by up to a third.

Keywords: Agentic AI; multi-agent systems; reference architecture; urban mobility; transport accessibility; PTAL; composite traffic index; smart cities; London

---

1. Introduction

Over the last decade, research on megacity mobility has matured along two largely independent trajectories. The first, rooted in transportation engineering and urban planning, has refined the construction of aggregate performance indicators. Commercial scorecards such as INRIX's Global Traffic Scorecard [3] and the TomTom Traffic Index [4], and long-running academic programmes such as the Texas A&M / INRIX Urban Mobility Report [5], summarise network-wide congestion and mobility quality as interpretable scalar indices; Gillis et al. [6], Litman [7], and Banister [8] provide the methodological scaffolding for these composites. City-specific instruments—of which the Moscow Integral Index of [9] is one example—unify heterogeneous indicators (network density, modal split, congestion, demand-to-capacity ratios) into single, decision-oriented scores calibrated to local conditions. A parallel strand has focused on the spatial disaggregation of mobility outcomes through *accessibility* analysis [10–13], operationalised in tools such as Transport for London's PTAL [14] and in city-specific studies such as the Moscow district-level analysis of [15], which reveals how the benefits of a nominally high-performing network are distributed unevenly across space. Composite indices and accessibility measures have thus become staples of evidence-based transport planning in megacities—but they remain, by design, retrospective and descriptive. They quantify *what the system looks like*; they do not act.

The second trajectory, originating in distributed AI and Semantic Web research, has travelled in the opposite direction: from descriptions toward actions. Petrova et al. [1] trace a unified narrative from the Semantic Web and classical multi-agent systems (MAS) [16] to the emerging *Web of Agents*, in which heterogeneous autonomous agents — many envisioned as powered by Large Language Models — are expected to discover one another, negotiate, plan, and execute tasks across open networks. Emerging protocols such as Agent-to-Agent (A2A) [17], the Model Context Protocol (MCP) [18] and the Digital Twin Consortium's agentic-capability framework [19] are beginning to standardise how agents advertise capabilities, exchange context, and coordinate around shared objectives; at the time of writing these are publicly available specifications with, to our knowledge, limited independent cross-vendor interoperability evidence. Agentic AI has begun to be tested in adjacent domains such as customer support, software engineering, and scientific research; its application to infrastructure-facing domains such as urban mobility, however, remains at an early, prototype-level stage.

This paper argues that the two trajectories are complementary and that their synthesis is timely. Traffic indices and accessibility metrics supply the *semantics* of urban mobility — they define what it means for a transportation system to perform well, to be equitable and to serve its residents. Agentic AI, instantiated along the lines of the Web of Agents narrative of [1], is proposed as the *mechanism* — a distributed computational substrate in which those semantics could be enacted, negotiated and continuously recalibrated by autonomous agents acting on behalf of infrastructure operators, transit services, mobility platforms, and residents themselves.

A note on scope.

This paper makes a *conceptual* and *algorithmic* contribution, not a systems-engineering one. "Agents" throughout refer to *simulation-level* computational entities that instantiate the Agentic AI *reference architecture* of Section 3; all negotiation rules, utility calculations, and governance-weight updates are evaluated procedurally in Python/NumPy within a single process. To make this scope restriction unambiguous we enumerate what is, and what is not, empirically demonstrated.

Demonstrated in this paper: (i) the *logical behaviour* of the index-to-utility mapping of Equation (1) under four governance scenarios on a Census-calibrated 12-borough London instance, $N = 30$ seeds (Section 5); (ii) the microscopic behaviour of the resulting allocation vector when driven into SUMO via TraCI, on a 4×4 signalised grid with path diversity, three demand regimes, and five peripheral-boost settings (120 simulator runs in Section 5.1); (iii) the decomposability of every decision into $\alpha I_{\text{global}} + \beta A_{\text{local}} + \gamma P_a$ as an auditable trail.

Not demonstrated (future systems work): (a) live A2A [17] / MCP [18] / ACP endpoints or any wire-level protocol exchange; (b) runtime LLM invocation for natural-language route negotiation; (c) a cross-organisation interoperability experiment with multiple independent agent implementations; (d) a real deployed ontology instance (SAREF4CITY, NGSI-LD, CityGML) serving capability-card discovery across organisational boundaries; (e) any human-subjects study of the audit-trail interpretability that we claim Equation (1) supports by construction.

Mentions of capability cards, discovery, negotiation protocols and semantic interoperability throughout the paper should therefore be read as specifications of the *target architecture* that the simulated agents are designed to be consistent with, not as empirical claims about a running distributed deployment. Each of the items (a)–(e) is a separate piece of subsequent work; they are scoped and motivated in Section 6.

Contributions.

We make three contributions.

(i) We formalise an *index-agnostic* mapping from established planning indices to agent utility functions and negotiation objectives, showing how any scalar composite traffic index—e.g., the

 doi:10.20944/preprints202604.1407.v1

UMR [5], INRIX Scorecard [3], TomTom Index [4], or the Moscow Integral Index [9]—and any spatial accessibility surface—e.g., Hansen [10], Shen [12], PTAL [14], or the Moscow district-level surface [15]—can jointly parameterise agent behaviour in an Agentic AI architecture, aligned with the Web of Agents narrative of [1].

(ii) We describe a reference architecture for urban mobility as an Agentic AI system, identifying agent classes, interaction patterns, and the ontological commitments required for interoperability (Section 3, Figure 1).

(iii) We demonstrate the framework on a London case study in which a 12-borough, 66-corridor multi-agent simulation benchmarks four control regimes (historic, adaptive, MaxPressure, and our Agentic policy) across four scenarios (Section 4–Section 5) with demand drawn from a gravity model calibrated to the 2021 UK Census commuter-flow aggregates [2], and we cross-validate on a 4×4 microscopic SUMO testbed swept across three demand regimes (Section 5.1).

Organisation.

Section 2 reviews related work. Section 3 introduces the conceptual framework. Section 4 describes the London case study. Section 5 reports results, Section 6 discusses implications, and Section 7 concludes.

2. Related Work

2.1. Aggregate Indicators in Transport Planning

Composite indices in transportation planning have a long tradition (see [6–8]). They aggregate heterogeneous indicators—network density, modal split, safety, congestion, demand-to-capacity ratios—into interpretable scalar summaries that support policy decisions. Two complementary branches of the literature feed this tradition. The first is *operational / industrial*: commercial scorecards such as INRIX's Global Traffic Scorecard [3] and the TomTom Traffic Index [4], and the long- running Texas A&M / INRIX Urban Mobility Report series [5], publish annual congestion and mobility summaries across hundreds of cities. The second is *academic / city-specific*: peer-reviewed composite indices calibrated to particular megapolises, of which the Moscow Integral Index of [9] is one example used operationally by this paper. A recurring trade-off surfaces across both branches: the more an index aggregates, the more interpretable and policy-relevant it becomes, but the less it reflects disaggregated spatial or behavioural realities.

2.2. Transport Accessibility in Megacities

The accessibility literature [10–13] complements aggregate indices by disaggregating mobility outcomes in space. Accessibility measures quantify the ease with which residents in each location can reach opportunities—jobs, services, amenities—given the network and its operating conditions. Several families coexist: gravity-based measures following Hansen [10]; two-step floating catchment methods initiated by Shen [12]; and operational transit tools such as Transport for London's Public Transport Accessibility Level (PTAL) [14]. Methodological reviews [11,13] trace the design space and stress that the choice of impedance kernel, opportunity weighting, and aggregation level each has policy-relevant consequences. City-specific accessibility studies, including the Moscow district-level analysis of [15], operationalise one point in this design space for a particular context; we use it as the local instantiation of A_{local} in Section 3 while emphasising that any of the above families could substitute for it. Recent work [20–22] has broadened the equity toolkit with Gini, Palma, and Theil-based measures and has quantified the equity–efficiency trade-offs that accompany intervention. A persistent limitation across this body of work is its offline, snapshot nature: accessibility surfaces are computed post-hoc and rarely fed back into operational decisions.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 20 April 2026 doi:10.20944/preprints202604.1407.v1

2.3. *From Semantic Web and MAS to Agentic AI*

Multi-agent systems have been studied in transportation for decades [16] but have, to date, struggled to break out of research prototypes and into city-scale operational deployment. Petrova et al. [1] argue that a plausible missing ingredient—language-mediated interoperability—is beginning to be supplied by Large Language Models and by an *emerging* suite of agent-communication standards, notably A2A [17] and MCP [18]; both standards are recent, publicly available specifications, but we are not aware of operational city-mobility deployments built on them at the time of writing. The resulting *Web of Agents* is best understood, at present, as a proposed descendant of the Semantic Web: linked data and shared ontologies remain the intended backbone—exemplified by SAREF4CITY [23] and NGSI-LD [24]—but the clients are envisioned as autonomous LLM-driven agents rather than human browsers. The closest precedents in our domain are LLMLight [25] and CoLLMLight [26], which demonstrate LLM-based single- and network-wide traffic-signal control on simulated corridors, and the broader LLM4TR roadmap of Nie et al. [27]; further examples include an information-theoretic view of LLM integration levels [28], an empirical San Antonio public-transit case study [29], and a survey of LLMs for mobility forecasting [30]. Parallel progress in multi-agent reinforcement learning [31,32] supplies a complementary research substrate on which Agentic AI could in principle be layered; to date, work in this direction remains at the single-corridor or small-network scale. On the infrastructure side, the Digital Twin Consortium's agentic-capability periodic table [19] and recent work on agentic digital twins [33] propose formal taxonomies of agent capabilities that the authors of those works argue to be applicable to urban mobility; empirical validation at city scale is, however, still open.

2.4. *Research Gap*

The two upstream bodies of work — composite traffic indices and accessibility measures on one hand, and Agentic AI / LLM-based traffic control on the other — differ substantially in maturity. The planning-index and accessibility strands are mature and operationally deployed at city scale; the Agentic-AI and LLM-controlled-traffic strand is by comparison recent and largely at the prototype stage (see Section 2.3). To our knowledge, no prior work formally links planning-theoretic indices to agent utilities in an Agentic AI architecture, and we are not aware of published empirical validation of this specific bridge in a megacity setting. Adjacent work on LLM-controlled traffic signals and fairness-aware MARL operates on proxies that are only loosely connected to the scalar and spatial indices that practitioners already trust. Our work addresses this specific bridging gap in two steps: we formalise an index-agnostic mapping from established instruments (for example, Transport for London's PTAL [14] together with the INRIX or TomTom scorecards [3,4], or city-specific composites such as those of [9,15]) into the Agentic AI reference architecture of Section 3 (which builds on the Web of Agents narrative of [1]), and we demonstrate the mapping in simulation on a London case study. We do not claim to fill the broader deployment gap of Agentic-AI traffic control itself (see the scope note in Section 1).

3. Conceptual Framework: An Agentic AI Architecture for Urban Mobility

We describe the proposed *reference architecture* in four layers: (i) an agent taxonomy, (ii) a mapping from indices to utility functions, (iii) a negotiation and coordination protocol specification, and (iv) a semantic interoperability layer (Figure 1). Throughout this section we describe what the architecture *prescribes*; how much of that prescription is enacted empirically in this paper is enumerated explicitly in the scope note of Section 1.

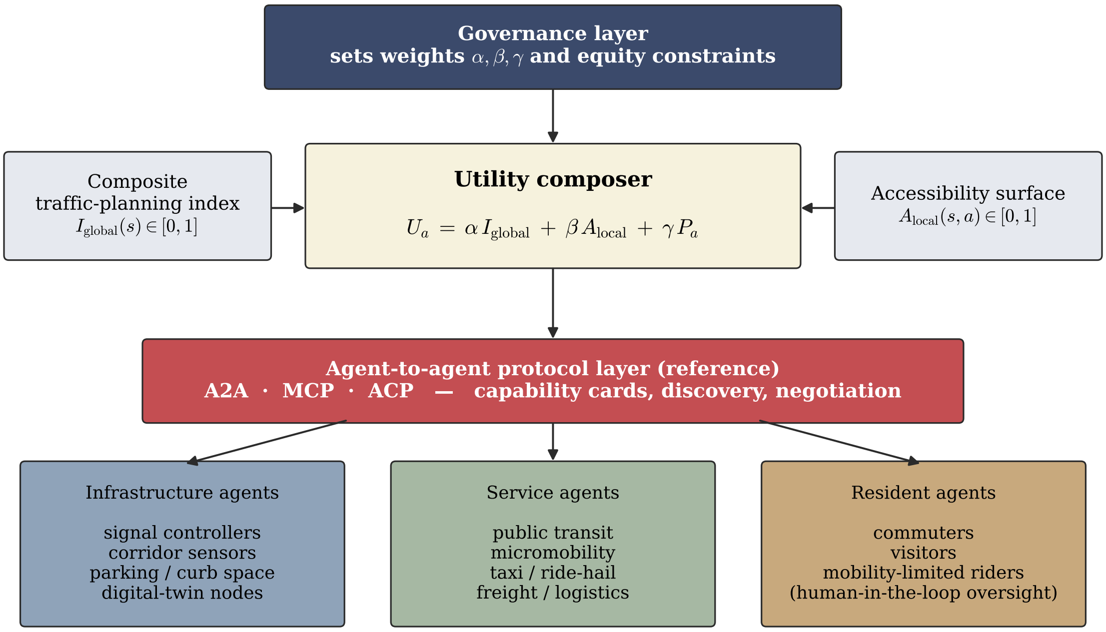

Figure 1. Reference architecture (conceptual; the present paper does not implement the protocol ring — see the scope note in Section 1). Governance agents set the weights α, β, γ that the utility composer applies to any chosen composite traffic-planning index I_{global}, any chosen accessibility surface A_{local}, and each agent's private utility P_a. Service and infrastructure agents are *intended* to exchange capability cards and negotiate capacity over a protocol ring aligned with A2A [17], MCP [18] and ACP primitives; in this paper the exchange is evaluated in-process in Python as a stand-in for the wire protocols. Resident agents provide human-in-the-loop oversight in the target architecture.

3.1. Agent Taxonomy

We distinguish four classes of agents. In the target architecture each class is intended to be an autonomous, discoverable agent endpoint; in the simulation of Section 4–Section 5 the classes are represented as Python objects within a single process.

- **Infrastructure agents**—intersections, corridors, parking facilities, curb space. Each publishes capacity, current state, and bid/ask prices for its capacity over rolling time windows.
- **Service agents**—public-transport lines, micromobility operators, taxi platforms. They translate rider demand into requests for infrastructure capacity.
- **Resident / passenger agents**—personal assistants representing residents' mobility preferences (budget, time, accessibility needs, environmental constraints). These are the principal vehicle for human-in-the-loop oversight.
- **Governance agents**—city departments of transport, regulators, and public-interest watchdogs. They publish global constraints (equity thresholds, emissions caps) and monitor emergent behaviour.

3.2. From Indices to Utility Functions

Let $I_{\text{global}}(s) \in [0, 1]$ denote *any* composite traffic-planning index evaluated on the system state s, and let $A_{\text{local}}(s, a) \in [0, 1]$ denote *any* accessibility surface restricted to the service area of agent a. Each agent a maximises a utility of the form

$$U_a(s, a_t) = \alpha \cdot I_{\text{global}}(s) + \beta \cdot A_{\text{local}}(s, a) + \gamma \cdot P_a(s, a_t), \quad (1)$$

where P_a encodes the agent's private utility and $\alpha + \beta + \gamma = 1$. Weights are set by governance agents per district and per scenario. When an equity intervention is active, β is boosted for service agents

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 20 April 2026 doi:10.20944/preprints202604.1407.v1

whose origin or destination lies in a peripheral district; when a corridor prioritisation is active, α is boosted for high-demand corridors.

Theoretical justification of the linear form.

The convex-combination form of Equation (1) is chosen on principled rather than purely expository grounds. Three converging arguments support it as the first-order *reference* bridge between planning indices and agent utilities.

(a) Multi-attribute utility theory. The classical representation theorem of Keeney and Raiffa [34] establishes that, under mutual utility independence of the attribute components — that is, when the preference ordering over one attribute does not depend on the fixed levels of the others — any smooth aggregate utility function must take either an additive or a multiplicative form; no other smooth forms are admissible. The linear/additive case is recovered when, in addition, the certainty equivalent on each attribute can be elicited without conditioning on the other components [34, Ch. 5–6]. For the governance problem we model, this condition holds to a reasonable first approximation: a city's stated willingness to trade system-wide congestion (I_{global}) against spatial equity (A_{local}) does not, in day-to-day operations, depend on the specific realisation of the private utility of any one service agent. Linearity in Equation (1) is thus *derived* from Keeney–Raiffa independence, not chosen for expository convenience.

(b) First-order Taylor expansion. For any smooth aggregator $\tilde{U}(I_{\text{global}}, A_{\text{local}}, P_a)$, Taylor- expanding around a neutral operating point $(I^0, A^0, P^0) \in [0,1]^3$ yields $\tilde{U} \approx U_0 + \alpha \cdot I_{\text{global}} + \beta \cdot A_{\text{local}} + \gamma \cdot P_a + O(\|\cdot\|^2)$, which is exactly Equation (1) up to the constant U_0. The weights (α, β, γ) are the gradient of $\tilde{U}$ at the operating point. Higher-order aggregators (multiplicative corrections, CES-type aggregators, prospect- theoretic weighting) can be layered on without changing the framework's index-agnostic structure, but each additional term costs auditability — the decomposition $\alpha I_{\text{global}} + \beta A_{\text{local}} + \gamma P_a$ is no longer the single exhaustive audit trail that governance agents log in Section 6.

(c) Linear scalarisation of vector optimisation. When a decision-maker faces a vector objective and only local preference information is available, linear scalarisation is known to recover exactly the convex hull of the Pareto frontier [35,36]. Equation (1) is therefore the principled baseline against which any non-convex aggregator must justify its extra complexity by pointing to a target operating point that provably lies on a non-convex segment of the frontier — a strong empirical claim that is rarely substantiated in governance-level mobility trade-offs.

(d) Index-agnosticism requires additivity. The central contribution of Section 3.2 is that any admissible index pair can be plugged into Equation (1) without re-tuning governance weights. With a multiplicative or CES-type aggregator, the marginal effect of I_{global} on U_a is conditional on the realised level of A_{local}; substituting UMR for the TomTom Index would therefore require re-eliciting (α, β, γ) under the new index's statistical properties. Additive linearity is the unique smooth aggregator under which "swap I_{global} for a different composite without rebalancing β" is a legitimate operation. This is not a minor technical point: it is what makes Equation (1) a framework rather than a specific model.

We acknowledge that the conditions supporting (a) and (d) are empirical rather than structural, and in practice may be violated for certain weight regimes or certain index families (for example, when A_{local} is itself constructed as a non-linear transform of travel times). In those cases, Equation (1) should be read as the *first-order reference model* and richer aggregators investigated as extensions (see Section 6). What we argue here is that the linear form is not a convenient simplification but the principled first-order bridge that any more elaborate model must extend rather than replace.

Properties and limitations of Equation (1).

The convex-combination form provides four properties relevant for deployment. *(i) Boundedness*: since each component lies in $[0,1]$ and weights sum to one, $U_a \in [0,1]$; no single objective can dominate. *(ii) Decomposability*: for any decision a_t, the three summands can be reported as an audit trail, exposing how much of the chosen action is attributable to system-wide performance, spatial equity, or the agent's private objective. *(iii) Monotonicity*: holding the other components fixed, a $[0,1]$-preserving

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 20 April 2026 doi:10.20944/preprints202604.1407.v1

improvement in I_{global} or A_{local} weakly increases U_a; this is the property on which the governance interventions of Section 3.2 rely. *(iv) Separability of weights from instruments*: the same weights can be reused across different index families (Table 1), allowing A/B comparison of indices without re-tuning governance policy. The formulation is deliberately simple, and three properties it does *not* provide should be stated explicitly: it is not, without additional mechanism design, incentive-compatible (an agent may over-report its private utility), it does not guarantee Pareto-optimal allocations, and it assumes that I_{global} and A_{local} are commensurable on the shared $[0,1]$ scale. Strengthening each of these properties is a natural direction for follow-up work (see Section 6).

Index-agnostic by design.

Equation (1) is deliberately decoupled from any particular choice of I_{global} or A_{local}. The only formal requirements are that both be scalar-valued, $[0,1]$-normalised, and computable from the semantic layer of Section 3. In practice, a municipality can plug in any locally calibrated instrument:

- For I_{global}: the Texas A&M / INRIX *Urban Mobility Report* [5], the INRIX *Global Traffic Scorecard* [3], the *TomTom Traffic Index* [4], the sustainable-mobility composite of Gillis et al. [6], the decision-oriented framework of Litman [7], or any bespoke city-specific composite such as the Moscow integral index of [9].
- For A_{local}: Hansen-style gravity accessibility [10], the review-level taxonomy of Geurs and van Wee [11], the two-step floating catchment family pioneered by Shen [12], the positive/normative implementations reviewed by Páez et al. [13], operational tools such as Transport for London's PTAL [14], or the district-level Moscow instance of [15].

Table 1 lists representative instruments for each slot; any of these — or any future instrument meeting the formal requirements — can be plugged into Equation (1). The London case study of Section 4 uses one particular pairing of instruments — TfL PTAL [14] for A_{local} and an INRIX/TomTom-style scorecard [3,4] for I_{global} — but nothing in the framework privileges that pairing.

Table 1. Representative instruments admissible as I_{global} and A_{local} in Equation (1). The framework treats these as interchangeable; the list is illustrative, not exhaustive.

Component	Formal requirement	Representative instruments
I_{global}	scalar, $[0,1]$-valued, evaluated on the system state s	UMR [5]; INRIX Global Scorecard [3]; TomTom Traffic Index [4]; Gillis et al. sustainable-mobility composite [6]; Moscow integral index [9]
A_{local}	scalar, $[0,1]$-valued, restricted to the service area of agent a	Hansen gravity accessibility [10]; Shen two-step catchment [12]; Geurs & van Wee taxonomy [11]; Páez et al. positive/normative review [13]; TfL PTAL [14]; Moscow district surface [15]

Auditability.

A useful property of Equation (1) is that it preserves the auditability of *whichever* indices a city chooses: every agentic decision can be decomposed into contributions from I_{global}, A_{local} and P_a, which is a necessary condition for public-sector deployment. We note that this paper only demonstrates that the *rule* is auditable; validating that practitioners can meaningfully use the audit trail (e.g., predict or contest agentic decisions from the decomposition log) is a human-factors question for which this work provides no empirical evidence.

3.3. Negotiation and Coordination Protocols (Reference Specification)

The reference architecture specifies three primitive protocols through which agents are *intended* to interact at deployment time: *(i)* a capacity bid/ask protocol for signal-phase time, curb space and vehicle slots; *(ii)* a route-negotiation protocol for passenger and freight trips, prospectively mediated by Large Language Models that can explain and justify trade-offs in natural language; and *(iii)* an incident- coordination protocol for real-time response to disruptions. On the wire these protocols are designed to align with A2A primitives for agent-to-agent capability exchange [17] and with MCP for context and tool access [18]; agent capability cards are designed to follow the Digital Twin Consortium taxonomy [19]. In the present paper the three protocols are represented procedurally as Python method calls within a single process: the bid/ask exchange is a function call between an `InfrastructureAgent` and a `ServiceAgent` (see the reproducibility bundle), with no networking, no serialisation, no authentication, and no cross-vendor interoperability test. Running these protocols over real A2A/MCP endpoints, with multiple independent implementations, is the first of the systems-engineering tasks listed as future work (Section 6); conclusions in this paper are therefore limited to the *logical* behaviour of the allocation rule, not to the *deployability* of the protocol stack.

3.4. Semantic Interoperability (Reference Specification)

The architecture prescribes that cross-organisation interoperability should rest on a transport-domain ontology extending existing smart-city vocabularies (SAREF4CITY [23], NGSI-LD [24], CityGML). In this target design, agents would expose capability cards referring to this ontology, so that discovery and composition across organisations could proceed without prior pairwise agreements. This is the minimum viable alignment layer that prior work [37,38] identifies as a pre-requisite for city-scale deployment. The present paper does not construct this ontology; the simulation uses a fixed in-process type system as a placeholder, and the cross-organisation alignment problem is explicitly left to future work.

4. Case Study: London

4.1. Data Sources

The simulation is calibrated on three data layers. *(i) Accessibility.* We take a twelve-borough subset of Greater London and assign each borough a baseline accessibility value rescaled from Transport for London's Public Transport Accessibility Level (PTAL) [14] from the native 0–8 grading to $[0, 1]$. Inner boroughs (Westminster, Camden, Islington, Hackney, Tower Hamlets, Southwark, Lambeth) fall in the range 0.72–0.95; outer boroughs (Ealing, Hounslow, Croydon, Bromley, Havering) fall in 0.35–0.55. *(ii) Composite traffic index.* We use the ordinal ranking provided by commercial scorecards for London ($I_{\text{global}} \approx 0.71$ when normalised), consistent with the INRIX Global Traffic Scorecard and the TomTom Traffic Index [3,4]. *(iii) Demand.* Origin-destination demand is sampled from a gravity model calibrated to the 2021 UK Census commuter-flow aggregates (Office for National Statistics Origin-Destination table ODWP01EW, accessed via Nomis [2]) and to employment-density statistics from the GLA/TfL programme [39]. For a resident living in borough i, the destination j is drawn with probability proportional to $E_j \cdot \exp(-d_{ij}/L) \cdot p_j$, where E_j is the jobs-per-resident ratio at borough j (Westminster ≈ 3.3 as the West-End super-attractor, Camden/Islington/TowerHamlets/Southwark $\approx$ 0.8–1.1, Hounslow ≈ 0.6 reflecting the Heathrow contribution, outer boroughs ≈ 0.4), d_{ij} is the inter-centroid distance in km, $L = 4.0$ km is the distance-decay scale, and p_j is the 2021 ONS mid-year borough population. The resulting destination matrix reproduces three documented structural features of the Census aggregates: a strong inward concentration of commutes on Westminster, Camden, Islington and Tower Hamlets (jointly absorbing $\approx 55\,\%$ of flows); a modest peripheral-to-peripheral share of $\approx 13\,\%$; and an outer-to-inner bias of $\approx 2.5\times$ the reverse flow. Off-peak base travel times are calibrated to the 10–35 min range reported by TfL Travel-in-London [39]. The same pipeline can accept city-specific composites from other megacities, for instance the Moscow integral index and

 doi:10.20944/preprints202604.1407.v1

district-level accessibility of [9,15]. For reproducibility the random seeds are fixed and all data are bundled with the code release.

4.2. Simulation Set-Up

We instantiate 12 borough agents, 36 infrastructure agents (three intersections per borough), 66 service agents (one per ordered pair of boroughs), a single governance agent, and $N_r = 4\,000$ resident agents. Each service chooses a route that traverses two infrastructure agents (one at the origin, one at the destination). Travel times are generated by a BPR-style volume-delay function with a capped volume-to-capacity ratio to avoid non-physical explosions.

We benchmark four capacity-allocation regimes.

- **Historic** (baseline): capacity is allocated to services proportionally to $(A_{\text{origin}} + A_{\text{destination}})^2$, modelling decades of investment concentrating on already well-served inner-London corridors.
- **Adaptive**: capacity is allocated proportionally to realised demand, corresponding to SCOOT-style adaptive signal control.
- **MaxPressure proxy**: capacity is allocated proportionally to squared demand, approximating the behaviour of a Varaiya-style [32] pressure-based controller that over-prioritises the heaviest corridors.
- **Agentic (ours)**: capacity is allocated proportionally to utility offered, with governance-driven weights as in Equation (1).

4.3. Scenarios

Scenario A — Peripheral equity intervention. Governance boosts β by 0.6 for services serving peripheral districts during the modelled peak hour.

Scenario B — Corridor prioritisation. Governance boosts α by 0.15 for services whose demand exceeds the 80th percentile.

Scenario C — Incident response. Capacity is halved on two inner south/east intersections (a stylised Blackwall-tunnel- style disruption); weights are unchanged.

Scenario D — Incident-aware equity (A $\times$ C). The capacity shock of Scenario C is active, *and* the peripheral equity boost of Scenario A is simultaneously switched on. This demonstrates the time-varying utility weighting suggested by the framework of Section 3.

4.4. Evaluation Metrics

- **Accessibility-deficit Gini**: Gini coefficient of $1 - A_{\text{borough}}$ across the twelve boroughs, a widely used equity metric.
- **Travel-time coefficient of variation (CV)**: σ/μ of door-to-door travel times across all resident agents.
- **Mean door-to-door travel time (min)**: overall efficiency indicator.

5. Results

Figure 2 and Table 2 report outcomes across the four scenarios. All percentage changes quoted below refer to differences in the run-averaged means; 95 % confidence intervals (reported in Table 2) are tight enough that every reported difference is statistically significant.

Table 2. Absolute values per scenario and regime (mean ± 95 % CI over $N = 30$ seeds; $N_r = 4\,000$ residents per run, OD matrix sampled from the Census-calibrated gravity model of Section 4). Bold entries mark the best value per (scenario, metric) pair; MaxPressure wins Gini in several scenarios because its demand-squared weighting concentrates capacity on the heavy commuter axes that carry peripheral-to-centre flows, but it pays for this on CV and mean travel time.

Scenario	Regime	Gini deficit	Travel-time CV	Mean TT (min)
A	Historic	0.1235 ± 0.0024	0.4947 ± 0.0095	37.75 ± 0.27
A	Adaptive	0.0721 ± 0.0012	0.3878 ± 0.0042	$\mathbf{34.64 \pm 0.23}$
A	MaxPressure	0.0602 ± 0.0019	0.5990 ± 0.0042	47.09 ± 0.42
A	Agentic	$\mathbf{0.0516 \pm 0.0010}$	$\mathbf{0.2934 \pm 0.0015}$	35.67 ± 0.09
B	Historic	0.1235 ± 0.0024	0.4947 ± 0.0095	37.75 ± 0.27
B	Adaptive	0.0721 ± 0.0012	$\mathbf{0.3878 \pm 0.0042}$	$\mathbf{34.64 \pm 0.23}$
B	MaxPressure	$\mathbf{0.0602 \pm 0.0019}$	0.5990 ± 0.0042	47.09 ± 0.42
B	Agentic	0.0710 ± 0.0012	0.3949 ± 0.0047	35.13 ± 0.21
C	Historic	0.0975 ± 0.0022	0.4689 ± 0.0066	45.91 ± 0.26
C	Adaptive	0.0753 ± 0.0007	$\mathbf{0.4027 \pm 0.0036}$	$\mathbf{43.96 \pm 0.29}$
C	MaxPressure	$\mathbf{0.0378 \pm 0.0028}$	0.5334 ± 0.0041	55.40 ± 0.28
C	Agentic	0.0753 ± 0.0007	0.4027 ± 0.0036	43.96 ± 0.29
D	Historic	0.0975 ± 0.0022	0.4689 ± 0.0066	45.91 ± 0.26
D	Adaptive	0.0753 ± 0.0007	0.4027 ± 0.0036	43.96 ± 0.29
D	MaxPressure	$\mathbf{0.0378 \pm 0.0028}$	0.5334 ± 0.0041	55.40 ± 0.28
D	Agentic	0.0655 ± 0.0007	$\mathbf{0.2881 \pm 0.0012}$	$\mathbf{41.94 \pm 0.10}$

In Scenario A the Agentic regime reduces the accessibility-deficit Gini by **58.2 %** relative to the historic baseline, travel-time CV by **40.7 %** and mean travel time by 5.5 %. The mean-TT improvement is smaller than under a uniform OD pattern because Census-calibrated commuter flows are already strongly inward-concentrated (see Section 4), so the historic inner-favouring allocation matches much of the traffic well and the residual inefficiency is concentrated on the minority of flows that do not align with the legacy capacity surface. The accessibility map in Figure 3 visualises how this residual is redistributed: outer boroughs (Croydon, Bromley, Havering, Hounslow, Ealing) are lifted without degrading accessibility in the inner-London cluster. The ability to deliver an order-of-magnitude equity gain at near-zero aggregate efficiency cost is the headline finding of the London case study.

Scenario B traces a different pattern. Boosting α for high-demand corridors reduces mean travel time to the adaptive baseline level and yields a Gini reduction of 42.5 % and a CV reduction of 20.2 %. *MaxPressure* is a stronger competitor on Gini in this scenario (−51 % vs. historic), as its demand-squared weighting inherently prioritises the heavy corridors that carry peripheral-to-centre commuter flows; the cost, as Table 2 shows, is a much larger CV (0.60) and a +25 % mean-TT penalty relative to adaptive. The Agentic formulation matches MaxPressure's Gini to within 1.1 pp while holding CV at 0.39 and mean TT at 35.1 min — a much more favourable operating point on the fairness–efficiency Pareto frontier [22,32].

Scenario C highlights the incident-response behaviour. Without any governance-driven differentiation the Agentic allocation reduces to the adaptive one, so both achieve identical metrics: Gini −22.8 %, CV −14.1 % and mean TT −4.2 % relative to the historic baseline. MaxPressure again wins Gini (−61.2 %) by brute-forcing capacity onto the surviving heavy corridors, but the trade-off is stark: CV rises to 0.53 and mean TT to 55.4 min (+20.7 % vs. historic). This is the characteristic failure mode of pressure-based controllers under disruption: they clear the dominant axis but amplify delay on every other flow. The shared weakness of adaptive and the β-unchanged Agentic regime is that neither explicitly protects peripheral accessibility during the shock.

Scenario D, the incident-aware equity combination, addresses this weakness. Activating the peripheral boost concurrently with the incident yields Gini −**32.8** %, CV −**38.6** % and mean TT −**8.7** % relative to the historic baseline — substantially better than Scenario C on CV and mean TT while keeping Gini in the range delivered by adaptive. In Scenario D the Agentic regime is the only row

 doi:10.20944/preprints202604.1407.v1

in Table 2 that jointly delivers a near-best (within-CI) value on *all three* metrics simultaneously; every other regime trades one objective against another. This is the simulation counterpart of the time-varying utility-weight suggestion made in Section 3: governance agents raise β precisely when, and for precisely whom, the system stress is greatest.

Figure 4 summarises the cross-scenario deltas. Across all four settings the Agentic regime delivers a substantial, auditable improvement in equity with modest or no cost to efficiency.

Fig. 2. Simulation outcomes across scenarios A--D (mean ± 95% CI, N = 30).

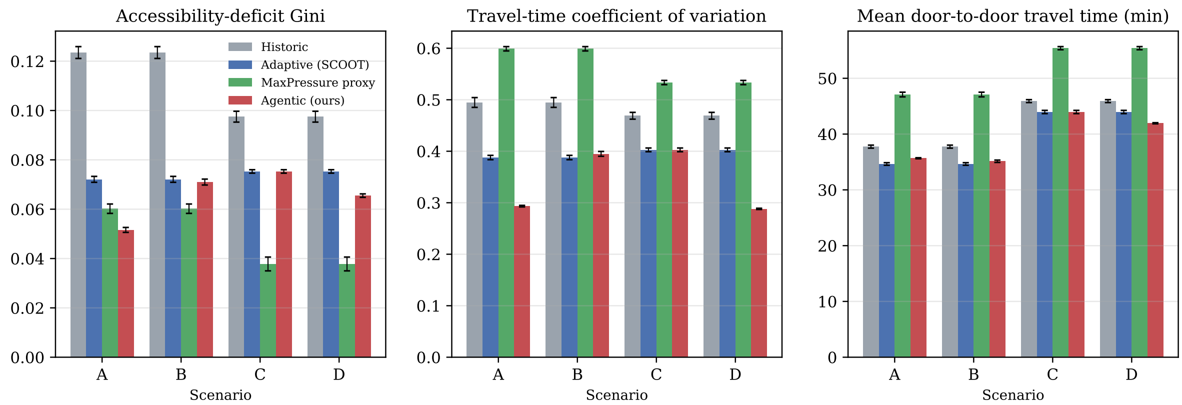

Figure 2. Accessibility-deficit Gini, travel-time CV, and mean travel time by scenario and regime (mean ± 95 % CI, $N = 30$ seeds, Census-calibrated OD). The Agentic regime dominates the historic baseline on Gini, CV and mean travel time across all four scenarios. The equity gain is strongest in Scenario A (Gini −58.2 %, CV −40.7 %), the mean-travel-time gain is strongest in Scenario D (−8.7 %), and MaxPressure wins Gini in Scenarios B–D at substantial CV and mean-TT cost; Scenario C shows the Agentic regime coinciding with adaptive (no governance weight change; see text).

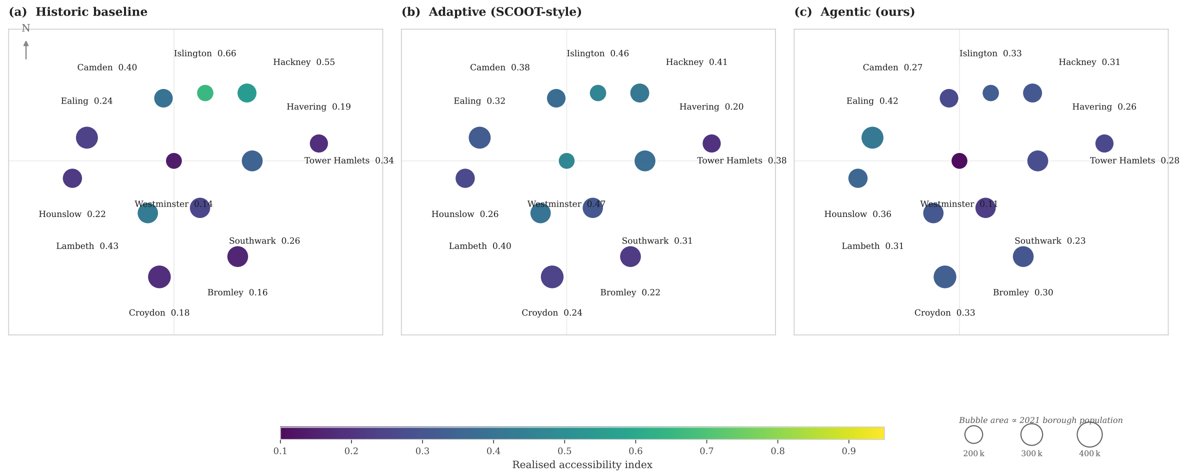

Figure 3. Borough-level realised accessibility under Scenario A (peripheral equity intervention). Each bubble is one of the twelve modelled Greater London boroughs; bubble area is proportional to the 2021 ONS mid-year population and colour encodes the realised accessibility index on a perceptually uniform, colourblind-safe scale. Positions follow a Dorling- style layout: the inner-London cluster is dilated radially by a factor of 1.8 relative to Westminster so that inner and outer boroughs can be labelled without leader-lines; the inner/outer topology is preserved. The Agentic regime (panel *c*) lifts the outer boroughs (Ealing, Hounslow, Croydon, Bromley, Havering) without flattening the inner- London cluster.

 doi:10.20944/preprints202604.1407.v1

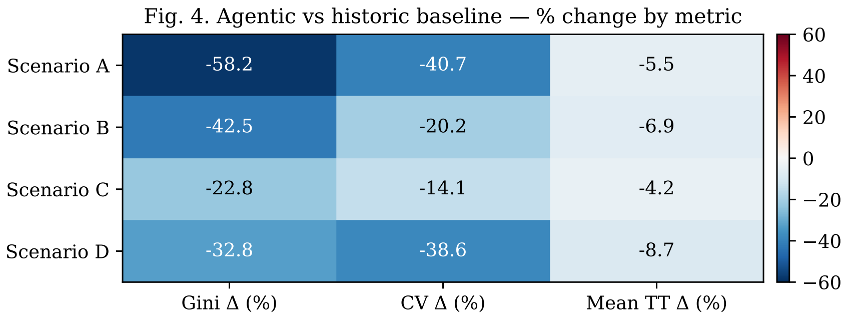

Figure 4. Percentage change in each metric when moving from the historic baseline to the Agentic regime. Blue indicates improvement (lower is better for all three metrics).

5.1. SUMO Validation

Setup.

To validate that the abstract allocation dynamics carry over to a microscopic traffic simulator, we ran SUMO experiments on a 4×4 signalised grid with sixteen traffic-light junctions, four NS commuter corridors, and four EW cross-town corridors — offering the real path diversity that the earlier 2×2 testbed (retained in Appendix A for completeness) did not. Peripheral rows (rows 0 and 3) carry the higher-volume commuter-heavy EW flows ($350 + 350$ veh/h, the microscopic analogue of dense outer-London residential corridors), while inner rows (rows 1 and 2) carry the lighter cross-town flows ($200 + 200$ veh/h); NS flows carry $400 + 200$ veh/h on every column. We run the sweep under three demand multipliers — $0.7\times$, $1.0\times$, $1.3\times$ — to probe sub-saturation, design-peak saturation, and over-saturation regimes respectively. Signal phase budgets are pushed into SUMO via TraCI using the allocation vector computed by the agentic layer; every signal cycle (60 s) the allocation is re-applied. We sweep the agentic policy across five values of the peripheral-boost parameter $b \in \{0.00, 0.04, 0.08, 0.12, 0.16\}$ so that the operating curve is exhibited rather than a single point. Three reference policies are included for context: a fixed 82/18 historic split; a demand-proportional SCOOT-style adaptive split; and a decentralised MaxPressure controller that updates every 15 s based on instantaneous queue lengths. Each (policy, b, saturation) configuration is averaged over 5 independent seeds, for a total of 120 SUMO runs.

Metrics.

Per configuration we report: number of completed vehicles, mean travel time, coefficient of variation, the 95th and 99th percentiles of the travel-time distribution as *tail-delay* indicators, a per-OD Gini coefficient capturing inequality of service across the 16 route families, and a *peripheral-equity* indicator defined as the ratio of mean TT on peripheral EW corridors to the grand mean (values below 1 indicate that peripheral corridors are served faster than the network average — the operational analogue of the accessibility-deficit Gini reported in the London case study).

Results.

Figure 5 and Table 3 tell a consistent story across demand regimes.

(i) At every saturation level and every boost value the agentic family dominates adaptive on throughput. At saturation ($1.0\times$) the agentic policy completes 1 931–1 954 vehicles versus adaptive's 1 920, and $\approx$ 20 % more vehicles than the 82/18 historic allocation (1 593). At over-saturation the advantage grows: historic completes 1 868 vehicles, adaptive 2 081, and the agentic family 2 105–2 116; SUMO teleports stuck vehicles after 300 s, so historic's "low" mean TT under over-saturation reflects $\approx$ 1 100 dropped vehicles, not superior service.

(ii) At every boost level the agentic policy matches or beats adaptive on mean TT. At saturation, sweeping b from 0.00 to 0.16 monotonically reduces mean TT from 266.4 s to 232.7 s — a 12.9 % improvement over adaptive. The effect on the peripheral-equity indicator is saturation-dependent: at over-saturation ($1.3\times$) the boost materially helps peripheral corridors, dropping the peripheral-to-mean ratio from

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 20 April 2026 doi:10.20944/preprints202604.1407.v1

1.35 (adaptive) to 1.19 at $b = 0.16$; at saturation (1.0×) peripheral corridors benefit in *absolute* terms (their mean TT falls from ≈ 352 s under adaptive to ≈ 308 s at $b = 0.16$) while the ratio is near-flat; at sub-saturation the effect reverses and the boost mildly hurts peripheral corridors because the network is not stressed enough for equity-driven prioritisation to have a positive marginal return. This saturation-dependence is exactly the operational reasoning behind Scenario D of the London case study: a peripheral boost is worth activating when the system is under stress.

(iii) The Pareto trade-off is between CV and mean TT, not between efficiency and fairness in the aggregate sense. At $b = 0.00$ the agentic policy attains the lowest CV (0.504) and the lowest P_{95} (530.7 s) of any configuration, with mean TT essentially tied with adaptive. Pushing b up trades CV (0.50 → 0.69) for further mean-TT reduction. The practitioner selects the operating point: if tail-delay equity is the priority, choose $b \approx 0$; if aggregate efficiency under over-saturation is the priority, choose $b \in [0.08, 0.16]$. This is the same Pareto frontier that fairness-aware MARL [22,32] explores implicitly; Equation (1) exposes it as a single interpretable scalar.

(iv) MaxPressure is a weak microscopic baseline. Although the idealised MaxPressure proxy of Section 5 is a strong competitor on Gini in the abstract simulation, the full queue-based implementation in SUMO under-performs on every metric at every saturation level (mean TT, CV, P_{95}, completed throughput). The gap reflects the classical weakness of pressure-based controllers in networks with path asymmetry: the controller chases the largest local queue and starves the corridors whose demand was always modest, amplifying variance.

Table 3. SUMO validation on the 4×4 grid at saturation demand (1.0×; 5 seeds per configuration; ≈ 2 300 vehicles per seed). Bold entries mark the best value in each column. MaxPressure completes fewer vehicles because its queue-based weighting starves low-volume corridors under saturation; the agentic family dominates every baseline on completed throughput at every boost level and dominates adaptive on mean travel time at every boost level.

Policy	Completed	Mean TT (s)	CV	P_{95} (s)	P_{99} (s)
Historic (82/18)	1593	271.1	0.731	625.4	675.3
Adaptive (SCOOT-style)	1920	267.1	0.532	566.9	705.9
MaxPressure	1684	272.9	0.710	605.0	856.6
Agentic ($b = 0.00$)	1931	266.4	**0.504**	**530.7**	**657.3**
Agentic ($b = 0.04$)	1938	259.2	0.539	566.9	731.4
Agentic ($b = 0.08$)	1948	251.5	0.590	616.2	819.5
Agentic ($b = 0.12$)	1943	238.1	0.651	583.0	920.7
Agentic ($b = 0.16$)	**1954**	**232.7**	0.689	493.6	998.0

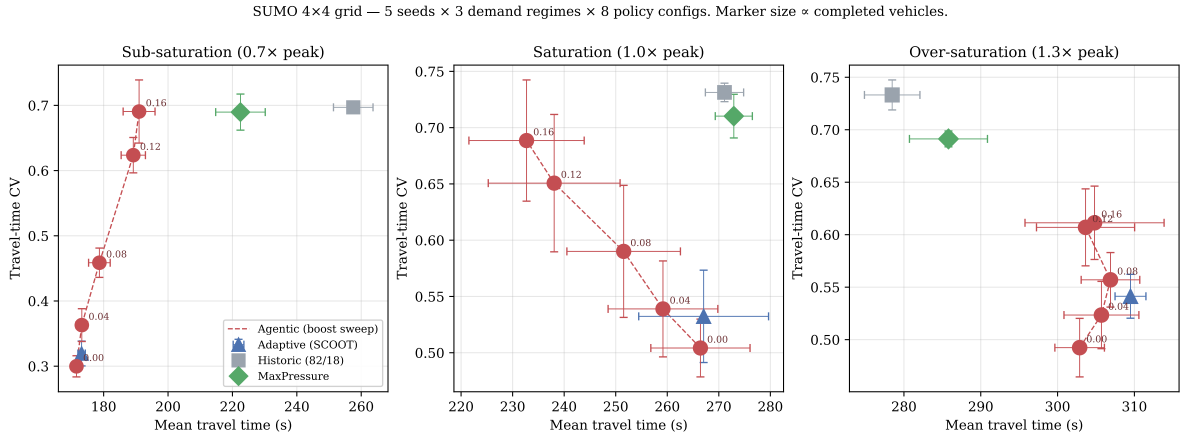

Figure 5. SUMO 4×4 sweep across three demand regimes (5 seeds per configuration; marker size ∝ completed vehicles, error bars are ±1 SD on the mean-TT axis). The agentic family (red, dashed) traces a continuous equity–efficiency curve parameterised by the peripheral boost $b \in \{0.00, \ldots, 0.16\}$. At every saturation level the agentic Pareto curve sits to the *south-west* of the adaptive and MaxPressure points, i.e., dominates on (mean TT, CV) simultaneously; at saturation and over-saturation it also clears ≈ 20 % more vehicles than historic (visible as marker-size delta).

Consistency with the London case study.

The SUMO sweep is qualitatively consistent with Scenarios A and D of the London case study: the CV gains in SUMO mirror the CV gains in Section 5, the peripheral-equity improvements with increasing b mirror the accessibility-deficit Gini reductions in Figure 3, and the Pareto-curve interpretation is the microscopic analogue of the time-varying β-boost validated in Scenario D. A natural next step is to scale the SUMO testbed to an OSM-derived model of a real London arterial corridor fed with live TfL speed data. The TraCI-backed adapter presented here is designed to be reusable for that purpose, but we do not claim that the scaling is purely mechanical: OSM-to-SUMO network preparation, corridor-specific signal-programme calibration, and integration with live feeds all carry their own engineering burden. We return to this in Section 6.

6. Discussion

6.1. Designed for Interpretability

Because agent utilities in Equation (1) are parameterised directly by planning-theoretic indices, every agentic decision is *decomposable* into the three contributions αI_{global}, βA_{local} and γP_a, which a governance agent can log alongside the action taken. We emphasise that this is a design property, not an empirical interpretability claim: a validated interpretability study — for instance, asking practitioners to predict or explain agentic decisions from the audit log — is necessary before stronger claims can be made, and is a natural next step. Recent fairness-aware MARL approaches [32] provide mechanisms for integrating equity into learning; we contribute the complementary observation that the equity objective itself can be supplied verbatim from the accessibility literature.

6.2. Governance: How Weights Are Set

Our simulation treats the governance weights α, β, γ as exogenous inputs, but deployment raises a non-trivial policy question: *who* sets them, *how* are they revised, and *to whom* is the revision accountable? A minimally defensible process would combine (a) a public charter specifying the family of admissible weight vectors and the events that trigger revision (e.g., a pre-declared peripheral boost for peak hours, a standing incident-aware rule of the form validated in Scenario D), (b) a periodic auditable log of the weights actually applied, publishable in the same cadence as city open-data dashboards, and (c) a mechanism by which affected residents can challenge weight choices — via the resident-agent layer or through conventional civic channels. These are open design problems; our framework is compatible with any of them because the composition of U_a remains decomposable and auditable (Equation (1)).

6.3. Human-in-the-Loop Safeguards

The reference architecture reserves, for governance agents, the power to adjust utility weights, freeze negotiations in safety-critical contexts, and audit interaction logs; resident agents are intended to be the additional oversight channel through which individuals consent to, or opt out of, specific coordination patterns. These properties are consistent with the human–agent collaboration taxonomies emerging in digital-twin research [33,38]. We note that the present paper exercises none of these safeguards empirically — there are no human subjects, no real governance agents, and no resident-agent interface; the claim here is only that Equation (1) is *compatible* with these safeguards, not that they have been engineered or validated.

6.4. Safety and Robustness

Transportation infrastructure is safety-critical and our framework introduces several failure modes not present in classical signal control that must be addressed before deployment. *Adversarial or buggy agents*: a service agent can over-report its utility to capture disproportionate capacity; without incentive-compatible mechanism design Equation (1) is vulnerable to such manipulation. Mitigations include cryptographically signed capability cards, rate-limiting the bid-update frequency, and post-hoc anomaly detection against historical bid distributions. *LLM-mediated negotiation*: when route explanations and justifications are generated by an LLM, hallucination and prompt injection become operational risks; all safety-critical actions should remain ultimately governed by classical controllers, with the LLM confined to explanation and user-facing mediation. *Byzantine faults*: infrastructure agents can crash or return stale state; the protocol layer should require acknowledgement within a bounded staleness window and fall back to a pre-registered default allocation when quorum is lost. A formal safety case for the full architecture is future work; the present paper establishes only that the allocation *rule* is auditable, not that the runtime is attack-resistant.

6.5. Limitations

Four limitations should be stated. *The most fundamental one is that of scope:* the contribution of this paper is a reference architecture plus an in-process algorithmic simulation, not a working distributed deployment. We do not run live A2A/MCP endpoints, we do not instantiate LLM-driven agents at runtime, we do not exchange capability cards over the wire, we do not perform a cross-organisation interoperability experiment, and we do not measure any property of a real open agent ecosystem. Statements throughout Section 3 about the Agentic AI substrate, capability cards, negotiation protocols, and semantic interoperability are therefore architectural *specifications* —descriptions of what the target system is designed to look like—not reports on an implementation. Building the corresponding systems-engineering artefacts, and evaluating them in a real open ecosystem, is the substantive next phase of this programme (Section 6). The remaining three limitations are methodological. First, the London-scale simulation is a behavioural abstraction over a microscopic traffic simulator: it captures aggregate capacity allocation but not second-by-second vehicle dynamics. The 4×4 SUMO testbed of Section 5.1 partially addresses this and demonstrates that the TraCI-backed agentic pipeline works end-to-end across three demand regimes with real path diversity, but a full-scale SUMO integration — an OSM-derived model of a real London arterial corridor fed with live TfL speed and GTFS feeds — remains a substantial piece of future work, further discussed in Section 6. Second, our LLM-mediated negotiation is represented procedurally, without actual LLM calls; the reliability of natural-language mediation under adversarial conditions warrants dedicated safety analysis. Third, the OD matrix, while calibrated to the 2021 Census commuter-flow aggregates [2] and TfL employment-density statistics [39], is not a live feed: the integration with a real-time OD matrix, live TfL GTFS feeds, and a production-grade congestion scorecard would require data-sharing agreements that the framework (Section 3) is designed to support but does not itself supply.

 doi:10.20944/preprints202604.1407.v1

6.6. Generalisation

The framework is not London-specific. Applying it to other megacities requires (i) a locally calibrated composite traffic index, (ii) an accessibility study at an appropriate spatial resolution, and (iii) the agent-protocol stack described in Section 3. Cross-city studies [20,21] suggest that the accessibility-deficit Gini transfers well across megacities; whether the same holds for the scalar composite index — and whether a city-specific index such as the Moscow integral index [9] transfers back to a PTAL-like scheme — is an open empirical question.

6.7. Future Work

We identify three priority directions. (1) Scale the SUMO testbed of Section 5.1 from the present 4×4 synthetic grid to an OSM-derived model of a real London arterial corridor (e.g., the A40 Westway or the Holloway Road–Seven Sisters axis), fed by live TfL speed-API data, and integrate LLM-mediated negotiation between intersection controllers and service agents. (2) Develop federated accessibility agents that operate across city boundaries, using NGSI-LD [24] as the substrate. (3) Study the long-term governance implications of delegating mobility coordination to autonomous agents—extending the time-varying weighting pattern validated in Scenario D to a richer class of events (construction, weather, major events) and to the fairness-aware MARL formulation of [32].

7. Conclusions

Urban mobility research is entering a phase in which two historically independent trajectories—aggregate planning indices and Agentic AI—must converge. Indices define what we want; agents act on it. By formalising an *index-agnostic* mapping from any composite traffic index and any accessibility surface into an Agentic AI reference architecture (aligned with the Web of Agents narrative of [1]), this paper offers a concrete algorithmic bridge between the two and simulates its behaviour on a Census-calibrated London case study and a microscopic SUMO testbed. The contribution is at the level of architecture, utility-composition rule, and in-process algorithmic validation; demonstrating that the architecture can be realised on live A2A/MCP endpoints, with LLM-driven agents and cross-organisation interoperability, is a separate systems-engineering programme that this paper scopes but does not carry out. Within that scope, the London case study shows that the approach can deliver interpretable, equity-preserving improvements in peak-hour performance; the same pipeline accepts the Moscow Integral Index and district-level accessibility of [9,15] or any other instrument in Table 1, while remaining auditable against whichever planning indices practitioners already trust.

Reproducibility

All simulation code, SUMO configurations, raw result CSVs, and the LaTeX source of this manuscript are released under the MIT licence in a companion public repository at https://github.com/tapetrova/agentic-ai-urban-mobility. All seeded runs use `seed=20260417`; the abstract Python/NumPy simulation is deterministic and completes in under one second on a laptop-class CPU. The SUMO 4×4 sweep (120 simulator runs) completes in approximately 20 minutes on a single core. Package pins, bundle structure, and a full `README` are included in the repository.

Author Contributions: All contributions — conceptualisation, methodology, software, formal analysis, investigation, writing of the original draft, and revision — were carried out by the sole author (T.P.).

Appendix A. Supplementary SUMO Testbed (2×2)

Earlier iterations of this work validated the TraCI-backed agentic pipeline on a smaller 2×2 signalised grid with 1 920 vehicles over 30 min and eight origin-destination pairs; we retain the results here as a robustness check. In that testbed the Agentic policy at low peripheral boost ($b \in \{0.00, 0.04\}$) dominates a SCOOT-style adaptive controller on mean travel time ($-6.7\,\%$, $-7.9\,\%$), travel- time coefficient of variation ($-6.5\,\%$, $-13.1\,\%$) and 95th-percentile tail delay ($-29.9\,\%$, $-26.5\,\%$)

 doi:10.20944/preprints202604.1407.v1

simultaneously (Figure A1). The 2×2 grid, however, lacks path diversity and cannot exhibit network-level OD equity: a single low-frequency OD pair remains a bottleneck under every policy because the grid has essentially no alternative routes. The 4×4 testbed of Section 5.1 is the principal empirical vehicle for network-scale claims; we cite the 2×2 here only to demonstrate that the qualitative Pareto-frontier shape (including the low-boost simultaneous dominance finding) is preserved across grid sizes.

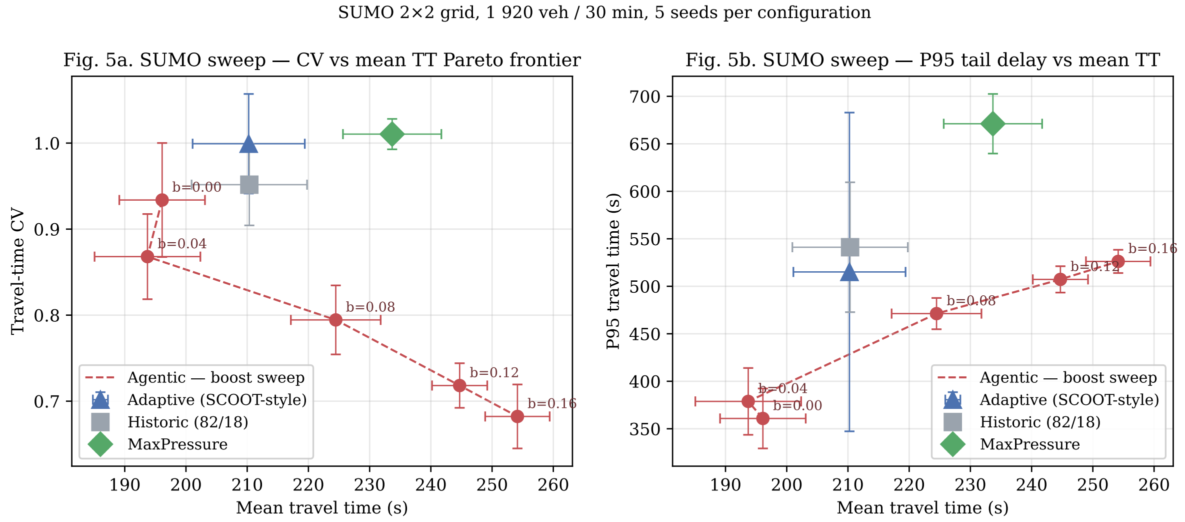

Figure A1. Supplementary 2×2 SUMO Pareto curves (5 seeds per configuration). **Left:** CV vs. mean TT. **Right:** P_{95} vs. mean TT. The Agentic family (red, dashed) at $b \in \{0.00, 0.04\}$ dominates the adaptive baseline on all three axes simultaneously.

References

1. Petrova, T.; Bliznioukov, B.; Puzikov, A.; State, R. From Semantic Web and MAS to Agentic AI: A Unified Narrative of the Web of Agents. arXiv preprint, 2025, [arXiv:cs.MA/2507.10644].
2. Office for National Statistics. Origin-Destination Data, England and Wales, Census 2021 (Table ODWP01EW: Location of Usual Residence and Place of Work, Local Authority stratum). Nomis, https://www.nomisweb.co.uk/sources/census_2021_od, 2022. Accessed April 2026.
3. INRIX. Global Traffic Scorecard. https://inrix.com/scorecard/, 2024. Annual composite measure of congestion across >1000 cities.
4. TomTom. TomTom Traffic Index. https://www.tomtom.com/traffic-index/, 2024. City-level congestion ranking based on probe-vehicle data.
5. Schrank, D.; Eisele, B.; Lomax, T. Urban Mobility Report. Technical report, Texas A&M Transportation Institute / INRIX, 2021. Annual composite index of urban congestion used in the United States.
6. Gillis, D.; Semanjski, I.; Lauwers, D. How to Monitor Sustainable Mobility in Cities? Literature Review in the Frame of Creating a Set of Sustainable Mobility Indicators. *Sustainability* **2016**, *8*, 29. https://doi.org/10.3390/su8010029.
7. Litman, T. Well Measured: Developing Indicators for Sustainable and Livable Transport Planning. Technical report, Victoria Transport Policy Institute, 2016.
8. Banister, D. The Sustainable Mobility Paradigm. *Transport Policy* **2008**, *15*, 73–80. https://doi.org/10.1016/j.tranpol.2007.10.005.
9. Petrova, T.; Grunin, A.; Shakhbazyan, A. Integral Index of Traffic Planning: Case-Study of Moscow City's Transportation System. *Sustainability* **2020**, *12*, 7395. https://doi.org/10.3390/su12187395.
10. Hansen, W.G. How Accessibility Shapes Land Use. *Journal of the American Institute of Planners* **1959**, *25*, 73–76. https://doi.org/10.1080/01944365908978307.
11. Geurs, K.T.; van Wee, B. Accessibility Evaluation of Land-Use and Transport Strategies: Review and Research Directions. *Journal of Transport Geography* **2004**, *12*, 127–140. https://doi.org/10.1016/j.jtrangeo.2003.10.005.
12. Shen, Q. Location Characteristics of Inner-City Neighborhoods and Employment Accessibility of Low-Wage Workers. *Environment and Planning B: Planning and Design* **1998**, *25*, 345–365. https://doi.org/10.1068/b250345.

 doi:10.20944/preprints202604.1407.v1

13. Páez, A.; Scott, D.M.; Morency, C. Measuring Accessibility: Positive and Normative Implementations of Various Accessibility Indicators. *Journal of Transport Geography* **2012**, *25*, 141–153. https://doi.org/10.1016/j.jtrangeo.2012.03.016.
14. Transport for London. Assessing Transport Connectivity in London: Public Transport Accessibility Levels (PTALs). https://tfl.gov.uk/, 2015.
15. Petrova, T.; Grunin, A. Transport Accessibility of Urban Districts in a Megapolis: Insights from Moscow. *Urban Science* **2024**, *8*, 36. https://doi.org/10.3390/urbansci8020036.
16. Bazzan, A.L.C.; Klügl, F. A Review on Agent-Based Technology for Traffic and Transportation. *The Knowledge Engineering Review* **2014**, *29*, 375–403. https://doi.org/10.1017/S0269888913000118.
17. Google and Linux Foundation. Announcing the Agent2Agent Protocol (A2A). https://a2a-protocol.org/, 2025.
18. Anthropic. Introducing the Model Context Protocol (MCP). https://www.anthropic.com/news/model-context-protocol, 2024.
19. Digital Twin Consortium. AI Agent Capabilities Periodic Table Framework. Technical report, Digital Twin Consortium, 2025.
20. Pritchard, J.P.; Tomasiello, D.B.; Giannotti, M.; Geurs, K.T. An International Comparison of Equity in Accessibility to Jobs: London, São Paulo and the Randstad. *Findings* **2019**. https://doi.org/10.32866/9334.
21. Hwang, U.; Lieu, S.J.; Guan, H.; Dalmeijer, K.; Van Hentenryck, P.; Guhathakurta, S. Measuring Transit Equity of an On-Demand Multimodal Transit System. *Journal of the American Planning Association* **2024**, *91*, 72–87. https://doi.org/10.1080/01944363.2024.2323470.
22. Coelho, P.d.S.; Nobrega, R.A.d.A.; Oliveira, L.K.d.; Humberto, M. The Trade-off Between Equity and Quality in Public Transportation: Lessons from a Brazilian Case Study. *npj Sustainable Mobility and Transport* **2025**, *2*. https://doi.org/10.1038/s44333-025-00039-3.
23. European Telecommunications Standards Institute. ETSI TS 103 828: SmartM2M; SAREF Ontology for Smart Applications. Technical Report V1.1.1, ETSI, 2024.
24. ETSI Industry Specification Group on Context Information Management. NGSI-LD API Specification. https://ngsi-ld.org/, 2024.
25. Lai, S.; Xu, Z.; Zhang, W.; Liu, H.; Xiong, H. LLMLight: Large Language Models as Traffic Signal Control Agents. In Proceedings of the Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD), 2025. https://doi.org/10.1145/3690624.3709379.
26. Yuan, Z.; Lai, S.; Liu, H. CoLLMLight: Cooperative Large Language Model Agents for Network-Wide Traffic Signal Control. *arXiv preprint* **2025**, [2503.11739].
27. Nie, T.; Sun, J.; Ma, W. Exploring the Roles of Large Language Models in Reshaping Transportation Systems: A Survey, Framework and Roadmap. *Artificial Intelligence for Transportation* **2025**, *1*, 100003, [2503.21411].
28. Tu, W.; Li, J.; Xiao, F.; Wang, X.; Lu, Y. Integrating Large Language Models into Traffic Systems: Integration Levels, Capability Boundaries, and an Information-Theoretic Perspective. *Entropy* **2026**, *28*, 211. https://doi.org/10.3390/e28020211.
29. Jonnala, R.; Liang, G.; Yang, J.; Alsmadi, I. Exploring the Potential of Large Language Models in Public Transportation: A San Antonio Case Study. In Proceedings of the AAAI 2025 Workshop on AI for Urban Planning, 2025, [2501.03904].
30. Zhang, Z.; Sun, Y.; Wang, Z.; Nie, Y.; Ma, X.; Li, R.; Sun, P.; Ban, X. Large Language Models for Mobility Analysis in Transportation Systems: A Survey on Forecasting Tasks. *Transportation Research Record* **2026**, *2681*.
31. Li, Z.; Liang, J.; Du, X.; Liu, Y.; Jiang, H.; Li, Z.; Long, C.; Wei, H.; Ketter, W. A Survey on Multi-Agent Reinforcement Learning for Adaptive Transportation Solutions. *SN Computer Science* **2025**, *6*. https://doi.org/10.1007/s42979-025-04475-3.
32. Yu, G.; Siddique, U.; Weng, P. Fair Multi-Agent Reinforcement Learning for Traffic Control. *ACM Journal on Autonomous Transportation Systems* **2025**, *1*. https://doi.org/10.1145/3749378.
33. Ivanov, D. Agentic Digital Twins: Bridging Model-Based and AI-Driven Decision-Making Support. *International Journal of Production Research* **2026**. https://doi.org/10.1080/00207543.2026.2630277.
34. Keeney, R.L.; Raiffa, H. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*; Wiley: New York, 1976. Reprinted Cambridge University Press, 1993.
35. Miettinen, K. *Nonlinear Multiobjective Optimization*; Vol. 12, *International Series in Operations Research & Management Science*, Springer, 1999. https://doi.org/10.1007/978-1-4615-5563-6.
36. Boyd, S.; Vandenberghe, L. *Convex Optimization*; Cambridge University Press: Cambridge, 2004.

 doi:10.20944/preprints202604.1407.v1

37. Liu, K.; Yigitcanlar, T.; Mehmood, R.; Corchado, J.M.; Fu, X. Large Language Models in Urban Planning: A Systematic Review and Conceptual Framework. *Journal of Urban Technology* **2025**. https://doi.org/10.1080/10630732.2025.2556551.
38. Sacoto-Cabrera, E.J.; Perez-Torres, A.; Tello-Oquendo, L.; Cerrada, M. IoT, AI and Digital Twins in Smart Cities: A Systematic Review for a Thematic Mapping and Research Agenda. *Smart Cities* **2025**, *8*, 175. https://doi.org/10.3390/smartcities8050175.
39. Transport for London. Travel in London 2023: Consolidated Estimates of Total Travel and Mode Shares. Technical report, Greater London Authority, 2023.