Supplementary Material: Scenario Studies for Agentic UE-CoMIMO for 6G

Chao-Kai Wen, Yen-Cheng Chan, Lung-Sheng Tsai, Pei-Kai Liao, and Geoffrey Ye Li

***Abstract*—The main article presents Agentic UE-CoMIMO as a vision for AI-native virtualized terminals. It illustrates the vision with a system-level simulation of an AI-glass uplink live-streaming session across an outdoor-to-indoor (O2I) boundary. The magazine format cannot include displayed mathematics or the full scenario, model, and parameter details. This supplement provides them. Section II defines the six functions of an agentic controller and shows how each scenario implements them. It also gives the generic intent-dependent utility that the two scenario scores instantiate. The main article foregrounds planning, intent interpretation, and feedback-driven replanning. Its capability-matched baselines omit these three functions. This supplement documents the other three and explains why they are fixed across the comparisons.**

In *Scenario 1*, the glasses stream across the O2I boundary and then dwell indoors for an extended session. A smartphone hub agent performs scene-aware band pre-switching. It recruits a *shared* indoor customer-premises equipment (CPE) and chooses between full delegation (phone cellular radio off) and packet-split cooperation as the CPE load changes. It also runs a predictive energy-aware controller over the long dwell. Beyond the thermal/energy bitrate governor, the controller uses the mapped environment and known trajectory to predict scene complexity. It holds the most quality-efficient target rung as an energy budget, lowers the encoded bitrate for a predicted-static scene, and adapts the capture frame rate to reduce camera power. These actions keep the *perceived* quality unchanged. The objective is quality of experience (QoE), including glass temperature, glass battery, and phone battery, rather than peak throughput. We score it per step and through longevity-aware useful quality-seconds, which reward longer target-quality streaming. Across four trajectory families and twelve seeds under non-ideal prediction, the proposed policy has the most useful quality-seconds and the best mean QoE in every family. A capability-matched myopic controller observes the same variables and commands the same actions, but maximizes only current-step QoE. It matches the proposed policy's instantaneous QoE in the canonical episode. However, it runs hotter, consumes more battery, dies about eight minutes earlier, and provides 22–47% fewer useful quality-seconds in the long episodes. This comparison isolates the value of prediction and intent from mere adaptivity.

We then study *Scenario 2*, a wearable blind-spot sensing episode. A rear-approaching cyclist causes smartwatch body blockage and, later, a transient scooter-CPE failure. The watch sensing receiver draws non-trivial power from a battery of about one watt-hour. Therefore, every adaptive controller idles it while the scooter-powered CPE covers the rear. The policies differ in how they recover when the delegated device disappears. Watch-only sensing fails during blockage. A static CPE commitment fails during the CPE outage. A capability-matched reactive controller detects the outage only after a report timeout and remains blind for over a second in the hazard zone. The agentic hub instead identifies the failure immediately from a micro-agent token. Its intent-dependent redundancy keeps the watch active inside the warning zone. It therefore maintains the delivered rear-hazard warning through both events, with 98.7% in-zone coverage. The other policies reach 82.9%, 65.8%, and 56.6%. The agent also uses less wearable energy than an always-on receiver. For each scenario, we give the narrative, the complete system model with all symbols defined, a parameter table that separates standards-based from illustrative values, and the simulation results. An interactive implementation accompanies this document, and every equation and number below matches that code.

This document is supplementary material for the article "Agentic UE-CoMIMO for 6G Terminals: From Virtual Antenna Augmentation to AI-Native Virtualization." It provides the definitional apparatus, scenario narrative, system model, parameter settings, and full simulation discussion that the magazine format cannot accommodate. Each scenario is self-contained, and additional scenarios will be appended as separate sections. The reference simulation code and interactive browser implementations of both scenarios are available at https://chaokaiwen.github.io/agentic-ue-comimo/.

I. Purpose and Scope

This supplement complements the main article in two ways. First, it records the engineering detail (geometry, channel and device models, control logic, and parameters) behind the agentic-control evaluations, at a level of rigor appropriate for reproduction but inappropriate for a magazine. Second, it serves as a growing library of *scenarios*, where each section defines one self-contained case study that exercises a particular aspect of Agentic UE-CoMIMO. The condensed illustration in the main article and the full scenario here both target the agentic control plane (UE-CoMIMO 2.0), which sits on top of the physical-layer virtual-antenna mechanisms (UE-CoMIMO 1.0): the main article reports the headline result, and this supplement gives the reproducible detail.

Throughout, we separate two classes of parameters. **Standards-based** values are taken from 3GPP, IEEE 802.11, IEC, or ITU-T references and are marked accordingly. **Illustrative** values are engineering choices selected to expose a qualitative trade-off. Their absolute magnitudes are plausible but are not measured device figures, and the conclusions are stated to depend only on their ordering, not on their exact values.

The magazine format admits no displayed mathematics and a strict figure and table budget. Section II therefore carries two definitional items that the main article states in prose and delegates here: the full six-function definition of an agentic controller, with its instantiation in the two scenarios, and the generic intent-dependent utility that the per-scenario scores of Sections III-F4 and IV-H instantiate.

II. Agentic Functions and the Intent-Dependent Utility

A. The Six Functions

We use the term *agentic* in a specific, testable sense rather than as a synonym for adaptive. An agentic controller is

defined by six functions.

- **Intent interpretation** maps the service objective onto the utility weights and constraint set of (1).
- **Task decomposition** translates the active intent into communication, sensing, and computation subgoals with explicit targets.
- **Planning** arranges actions over a forward horizon rather than one step at a time.
- **Tool selection** chooses and combines heterogeneous mechanisms, such as relay modes, transmitter and receiver roles, and compute placement, instead of tuning the parameters of one fixed mechanism.
- **Memory** maintains state that outlives a single decision, including radio maps, device-reliability weights, and thermal models.
- **Feedback-driven replanning** lets downstream task quality and device-state tokens revise earlier decisions, including events outside the original formulation, such as the failure of a recruited device.

A conventional context-aware cross-layer controller adapts within a fixed problem formulation. The agentic layer instead selects and combines mechanisms. It can also revise the formulation when the intent or observed task quality changes. Closed-form or learned optimization remains useful inside each selected mechanism.

Two clarifications apply. First, none of the six functions requires a large AI model. The controllers in Sections III and IV implement all six functions with interpretable rules and state machines. They stand in for future LLM-based hub agents. Thus, the reported gains arise from the functional structure, not from a specific model. Second, richer observations or a larger action space alone do not make a controller agentic. Both scenarios therefore include a capability-matched baseline that observes the same variables and commands the same actions.

The main article foregrounds three of the six, namely planning (there called prediction), intent interpretation, and feedback-driven replanning, because those are precisely the functions the capability-matched baselines are stripped of. Task decomposition and tool selection are instead held *fixed* across each comparison by construction: baseline D of Section III-G commands the identical mode set and pursues the identical rate subgoal, and the Reactive controller of Section IV-I commands the identical recruit, release, idle, and wake actions. Memory enters the comparison only through the predictive state that planning consumes, such as the position-indexed channel map and the thermal steady-state model, which is why it is not isolated separately. Table I records how each function is realized in the two scenarios.

B. The Intent-Dependent Utility

At each decision instant the hub agent selects the action that maximizes a generic intent-dependent utility,

$$U(t) = \lambda_{\mathrm{c}} Q_{\mathrm{comm}}(t) + \lambda_{\mathrm{s}} Q_{\mathrm{sense}}(t) + \lambda_{\mathrm{p}} Q_{\mathrm{perc}}(t) - \lambda_{\mathrm{e}} E(t) - \lambda_{\mathrm{o}} O_{\mathrm{coord}}(t) - \lambda_{\mathrm{th}} P_{\mathrm{thermal}}(t), \quad (1)$$

Here, $Q_{\mathrm{comm}}(t)$, $Q_{\mathrm{sense}}(t)$, and $Q_{\mathrm{perc}}(t)$ denote communication, sensing, and perception quality. The last term measures application-level performance. The terms $E(t)$, $O_{\mathrm{coord}}(t)$, and $P_{\mathrm{thermal}}(t)$ denote energy consumption, coordination overhead, and thermal stress. The weights λ_{c}, λ_{s}, λ_{p}, λ_{e}, λ_{o}, and λ_{th} depend on the current intent. They let the cluster balance communication, sensing, computation, and device constraints for the active objective instead of spectral efficiency alone.

Equation (1) is a template, not a separately implemented objective: each scenario instantiates it with concrete, fully specified terms, and it is those concrete scores that the simulations optimize and report.

Scenario 1 is communication-only, so $\lambda_{\mathrm{s}} = \lambda_{\mathrm{p}} = 0$. The per-step score (27) implements the remaining terms. The communication term becomes the perceptual utility $w_r\, u(r_v, \phi, C)$, gated on a non-empty buffer, minus the stall penalty w_s. The energy term becomes the glass-radio and phone-power penalties w_e and w_{ph}. The coordination term becomes the switching-delay and mode-switch penalties w_d and w_{sw}. The thermal term becomes the excess-temperature penalty w_{th}. Intent changes both the weights and the constraint set. A long-stream intent increases the cost of energy and heat relative to instantaneous quality. It also arms the governor, which lowers the permitted bitrate rung b_{top} from 25 to 15 Mbps. This constraint change is the main effect. Thus, a weight-only interpretation of intent is incomplete.

Scenario 2 is a joint instance in which the sensing term is active. In the joint score of Section IV-H, $\lambda_{\mathrm{c}} Q_{\mathrm{comm}}$ is the streaming utility w_{stream}; $\lambda_{\mathrm{s}} Q_{\mathrm{sense}}$ is the delivered-warning reward w_{warn} less the missed-warning, false-alarm, and warning-latency penalties w_{miss}, w_{false}, and w_{lat}; $\lambda_{\mathrm{e}} E$ is w_{energy}; and $\lambda_{\mathrm{o}} O_{\mathrm{coord}}$ is the sensing-token overhead and mode-switch penalties w_{token} and w_{switch}. No thermal model is included, so $\lambda_{\mathrm{th}} = 0$. Here the intent switch is explicit and time-varying: outside the warning zone the efficiency intent dominates and the watch receive chain idles, whereas inside the zone the safety intent raises the effective weight on Q_{sense} enough to justify three actions that all cost utility elsewhere, namely holding the watch warm as redundancy, capping the video enhancement layer, and duplicating the warning over both uplink paths.

In neither scenario is a separate perception term instantiated, so $\lambda_{\mathrm{p}} = 0$ throughout; application-level perception quality is folded into the communication and sensing terms. We keep Q_{perc} in (1) because the compute-placement action space of the architecture requires it, and we note its absence here rather than implying it has been evaluated.

III. Scenario 1: AI-Glass Uplink Live Streaming Across an O2I Boundary

A. Narrative and Motivation

A creator wearing AI glasses live-streams first-person video while walking from an outdoor street into an indoor venue such as a mall, studio, or exhibition hall, then settles at a fixed spot and keeps streaming for an extended session. The glasses generate an *uplink* video flow, so this is not a downlink consumption problem. The glasses have limited RF and thermal headroom, so they offload video over 5 GHz Wi-Fi to a smartphone and never transmit at 700 MHz, because

TABLE I
THE SIX AGENTIC FUNCTIONS, THEIR OPERATIONAL DEFINITIONS, AND THEIR INSTANTIATION IN THE TWO SCENARIO STUDIES.

Function	Operational definition	Scenario 1 (live streaming)	Scenario 2 (blind-spot sensing)
Intent interpretation	Service objective → utility weights and constraints in (1)	Long-stream intent arms the thermal/energy governor	Streaming → safety-critical intent switch inside the warning zone
Task decomposition	Intent → comm/sensing/compute subgoals with targets	Rate requirement derived from the permitted quality rung	Video layer capped; resources reserved for sensing tokens
Planning	Actions arranged over a forward horizon	Predicted scene complexity over a 4-s horizon; make-before-break pre-switch	Confirmed recruitment with setup scheduling
Tool selection	Choose and combine heterogeneous mechanisms	Direct, relay, packet-split, or full CPE delegation by phone-power order	CPE recruitment; bistatic transmitter/receiver role assignment
Memory	State that outlives one decision	Position-indexed channel map; thermal steady-state model	Device-reliability weights ρ in the fusion rule
Feedback-driven replanning	Task quality and device tokens revise earlier decisions	Thermal trend triggers pre-emptive bitrate reduction	Failure token triggers immediate fallback to phone–watch sensing

a low-band antenna is hard to integrate into an eyeglass form factor. The smartphone is the hub agent. It carries the cellular uplink, runs lightweight scene understanding, and decides how the cluster cooperates.

Two difficulties drive the scenario. The first is the O2I transition. Outdoors, the 3.5 GHz uplink offers ample capacity. As the user approaches and crosses the entrance, the 3.5 GHz link suffers heavy building penetration loss and a transient doorway blockage, and can collapse within a second or two. A purely reactive controller only switches bands *after* the link has degraded, paying a reconnection gap on an already-failing link and stalling the live stream. Indoors, a 700 MHz fallback cannot sustain the target bitrate through a high-loss building, so the hub agent engages a nearby indoor CPE. Worse, parts of the venue are outright RF dead-zones (a shielded lift lobby or a metal-clad hall) where the macro link to the gNB vanishes entirely, so a reactive band-switch can only stall, and the stream survives only by delegating to the in-venue CPE. The second difficulty is the long indoor dwell itself. The CPE is *shared* infrastructure whose spare uplink capacity fluctuates with other users, the smartphone battery drains fastest when its cellular transmitter is active, and the glasses heat toward their skin-contact limit if the encoder keeps running at the top bitrate. The control objective is the creator's QoE, namely visual quality, stall avoidance, latency, glass battery and temperature, and phone battery, rather than peak throughput. This scenario therefore exercises five agentic behaviors at once: intent-aware band selection, scene-aware proactive (make-before-break) switching, task-driven recruitment of an assisting device, resource-rational delegation to shared infrastructure (turning the phone's cellular radio off when the CPE alone suffices), and thermal foresight (relaxing the bitrate before the temperature limit is reached, rather than after). Over the long dwell, the controller adds three predictive energy-aware mechanisms. First, it holds the most quality-efficient bitrate as an energy budget instead of chasing the top rung. Second, it lowers the encoded bitrate when the predicted scene is static. Third, it lowers the capture frame rate during a predicted-static dwell to reduce camera power that bitrate control cannot change.

TABLE II
TRANSMISSION MODES IN SCENARIO 1.

Mode	Path	Description
M0	glass → gNB	AI-glass direct 3.5 GHz uplink (fallback only, high glass power and heat)
M1	glass → phone → gNB	Wi-Fi offload, smartphone relays on 3.5 GHz (outdoor workhorse)
M2	glass → phone → gNB	Wi-Fi offload, smartphone relays on 700 MHz (indoor robustness)
M3	glass → phone + CPE → gNB	Non-coherent packet-split cooperation on 700 MHz across smartphone and CPE (Option C)
M4	glass → phone → CPE → gNB	Full delegation: phone relays over Wi-Fi to the CPE and turns its cellular radio *off* (phone battery saver)

B. Geometry and Mobility

Time is discretized with step $\Delta t = 0.1$ s over an episode of $T_{\mathrm{ep}} = 900$ s, long enough to expose both the glass thermal dynamics (time constant ≈150 s, Section III-E) and the divergence in battery drain across policies. A single gNB sits at horizontal coordinate $x_{\mathrm{g}} = -100$ m at height $h = 25$ m. The user walks along a straight line at speed $v = 1$ m/s and stops at an indoor spot $x_{\mathrm{stop}} = 10$ m, so the user position is $x(t) = \min(x_0 + vt,\ x_{\mathrm{stop}})$ with $x_0 = -30$ m. The entrance is at $x = 0$ (crossed at $t = 30$ s), the fully indoor region begins at $x = 5$ m (reached at $t = 35$ s), and the user dwells at x_{stop} from $t = 40$ s onward. The 3D user-to-gNB distance is $d(t) = \sqrt{(x(t) - x_{\mathrm{g}})^2 + h^2}$. An indoor CPE near a window is placed at $x_{\mathrm{c}} = 15$ m and becomes discoverable once the user reaches the entrance; its gNB distance $d_{\mathrm{c}} = \sqrt{(x_{\mathrm{c}} - x_{\mathrm{g}})^2 + h^2}$ is constant. The indoor crossing fraction and indoor depth are

$$p_{\mathrm{f}}(t) = \mathrm{clip}\big(x(t)/5,\ 0,\ 1\big), \tag{2}$$

$$\delta(t) = \max\big(0,\ x(t) - 5\big), \tag{3}$$

where $\mathrm{clip}(z, 0, 1) = \min(\max(z, 0), 1)$. The AI glasses always reach the smartphone over a short 5 GHz Wi-Fi link that is provisioned not to be the bottleneck. The cluster uses one of five transmission modes, summarized in Table II.

C. Channel Model

All distances are in meters and all power quantities in dB or dBm. The path loss follows the 3GPP TR 38.901 urban-macro (UMa) NLOS model

$$\mathrm{PL}(d, f) = 13.54 + 39.08 \log_{10}\big(\max(d, 1)\big) + 20 \log_{10}(f), \tag{4}$$

with carrier frequency f in GHz. The receiver noise power on a band of width B (in Hz) is

$$N(B, I) = -174 + 10 \log_{10}(B) + \mathrm{NF} + I, \tag{5}$$

where $\mathrm{NF} = 7$ dB is the noise figure and I is an interference margin, set to $I_{700} = 6$ dB on the busier 700 MHz band and $I = 0$ otherwise. The three bandwidths are $B_{\mathrm{w}} = 80$ MHz (Wi-Fi) and the per-user cellular shares $B_{3.5} = 50$ MHz and $B_{0.7} = 10$ MHz, each half of the underlying 100 and 20 MHz carriers (Table III).

The O2I penetration loss ramps with $p_{\mathrm{f}}(t)$ using a 3GPP TR 38.901 *high-loss* building model that is consistent across bands:

$$L^{3.5}_{\mathrm{pen}}(t) = L^{3.5}_{\mathrm{w}}\, p_{\mathrm{f}}(t) + L_{\mathrm{door}}\left[4\, p_{\mathrm{f}}(t)\big(1 - p_{\mathrm{f}}(t)\big)\right], \tag{6}$$

$$L^{0.7}_{\mathrm{pen}}(t) = L^{0.7}_{\mathrm{w}}\, p_{\mathrm{f}}(t), \tag{7}$$

with steady wall losses $L^{3.5}_{\mathrm{w}} = 28$ dB and $L^{0.7}_{\mathrm{w}} = 18$ dB. The factor $4p_{\mathrm{f}}(1 - p_{\mathrm{f}})$ equals 1 at mid-crossing ($p_{\mathrm{f}} = 0.5$) and 0 at both ends. Thus, $L_{\mathrm{door}} = 8$ dB is the peak transient doorway blockage on the 3.5 GHz link. It represents body and door-frame diffraction. Indoor clutter adds $L_{\mathrm{cl}}(t) = \min\big(0.2\,\delta(t),\ 4\big)$ dB. Each mobile link also has a random term $X_b(t)$ for $b \in \{3.5, 0.7\}$. This term has two components. The first is spatially correlated log-normal shadowing with standard deviations $\sigma_{3.5} = 4$ dB and $\sigma_{0.7} = 3$ dB and a 10 m decorrelation distance. The second is first-order autoregressive temporal fading with standard deviation $\sigma_{\mathrm{t}} = 1.5$ dB and a 2 s correlation time. It represents people moving around the user, so the channel continues to fluctuate after the user stops. Body loss is $L^{\mathrm{g}}_{\mathrm{b}} = 2$ dB for the glasses and $L^{\mathrm{p}}_{\mathrm{b}} = 3$ dB for the smartphone.

Finally, a shielded interior pocket (for example a lift lobby or a metal-clad hall) over the position interval $x \in [x_{\mathrm{d0}}, x_{\mathrm{d1}}] = [8, 14]$ m adds a *dead-zone* loss

$$L_{\mathrm{dz}}(t) = L_{\mathrm{d}}\, \mathbf{1}\big\{x(t) \in [x_{\mathrm{d0}}, x_{\mathrm{d1}}]\big\}, \quad L_{\mathrm{d}} = 60 \text{ dB}, \tag{8}$$

to the *outdoor macro* uplink only (both the 3.5 GHz and the 700 MHz links to the gNB). The in-venue CPE, reached over the local Wi-Fi hop and backhauled on its own fixed link, is unaffected. The dead-zone thus blocks the cellular-only policies while leaving CPE delegation intact, isolating a coverage-hole regime from the thermal and energy regime that governs the rest of the episode.

With transmit power $P_{\mathrm{g}} = 20$ dBm and antenna gain $G_{\mathrm{g}} = -2$ dBi for the glasses, and $P_{\mathrm{p}} = 23$ dBm, $G_{\mathrm{p}} = 0$ dBi for the smartphone, the three mobile-link SINRs (in dB) are

$$\gamma_{\mathrm{g}}(t) = P_{\mathrm{g}} + G_{\mathrm{g}} - \mathrm{PL}(d, 3.5) - X_{3.5} - L^{3.5}_{\mathrm{pen}} - L_{\mathrm{cl}} - L^{\mathrm{g}}_{\mathrm{b}} - L_{\mathrm{dz}} - N(B_{3.5}, 0), \tag{9}$$

$$\gamma^{3.5}_{\mathrm{p}}(t) = P_{\mathrm{p}} + G_{\mathrm{p}} - \mathrm{PL}(d, 3.5) - X_{3.5} - L^{3.5}_{\mathrm{pen}} - L_{\mathrm{cl}} - L^{\mathrm{p}}_{\mathrm{b}} - L_{\mathrm{dz}} - N(B_{3.5}, 0), \tag{10}$$

$$\gamma^{0.7}_{\mathrm{p}}(t) = P_{\mathrm{p}} + G_{\mathrm{p}} - \mathrm{PL}(d, 0.7) - X_{0.7} - L^{0.7}_{\mathrm{pen}} - L_{\mathrm{cl}} - L^{\mathrm{p}}_{\mathrm{b}} - L_{\mathrm{dz}} - N(B_{0.7}, I_{700}). \tag{11}$$

The CPE link is treated as a *static* link with its own fixed wall loss $L_{\mathrm{c}} = 12$ dB and no shadowing, clutter, or body loss, since the CPE is fixed near a window:

$$\gamma_{\mathrm{c}} = P_{\mathrm{c}} + G_{\mathrm{c}} - \mathrm{PL}(d_{\mathrm{c}}, 0.7) - L_{\mathrm{c}} - N(B_{0.7}, I_{700}), \tag{12}$$

with $P_{\mathrm{c}} = 23$ dBm and $G_{\mathrm{c}} = 10$ dBi (a fixed directional panel antenna). The efficiency-discounted achievable rate of a link with SINR γ (dB), efficiency η, and bandwidth B is

$$R(\gamma; \eta, B) = \eta\, B\, \log_2\big(1 + 10^{\gamma/10}\big), \tag{13}$$

with implementation efficiencies $\eta_{\mathrm{c}} = 0.55$ (cellular) and $\eta_{\mathrm{w}} = 0.50$ (Wi-Fi). The individual link rates are $R_{\mathrm{g}} = R(\gamma_{\mathrm{g}}; \eta_{\mathrm{c}}, B_{3.5})$, $R^{3.5}_{\mathrm{p}} = R(\gamma^{3.5}_{\mathrm{p}}; \eta_{\mathrm{c}}, B_{3.5})$, and $R^{0.7}_{\mathrm{p}} = R(\gamma^{0.7}_{\mathrm{p}}; \eta_{\mathrm{c}}, B_{0.7})$. The Wi-Fi rate uses $\gamma_{\mathrm{w}} = 25$ dB outdoors and 22 dB indoors, $R_{\mathrm{w}} = R(\gamma_{\mathrm{w}}; \eta_{\mathrm{w}}, B_{\mathrm{w}})$.

The CPE is *shared* infrastructure. Its total uplink grant is $C_{\mathrm{tot}} = 40$ Mbps, and a two-state Markov background load $\beta(t)$, representing other users of the venue, occupies 5 Mbps (mean dwell 30 s) or 25 Mbps (mean dwell 15 s). The capacity allocatable to our cluster is therefore

$$R_{\mathrm{c}}(t) = \min\big(R(\gamma_{\mathrm{c}}; \eta_{\mathrm{c}}, B_{0.7}),\ C_{\mathrm{tot}} - \beta(t)\big) \in \{35,\ 15\} \text{ Mbps}, \tag{14}$$

The physical CPE link rate is approximately 38 Mbps at $\gamma_{\mathrm{c}} \approx 21$ dB, which exceeds the grant. Therefore, the CPE backhaul can carry the target rung by itself. The changing CPE load forces the agent to rebalance between delegation and cooperation indoors. Inside the dead-zone of (8), the phone's 700 MHz leg is blocked. The cooperative split (M3) then loses its second link and reduces to pure CPE delegation (M4). The CPE backhaul is the only path that keeps the stream active there.

D. Cooperative and Delegated Uplink Modes

Packet-split cooperation (M3) uses the non-coherent hybrid defined by Option C. A fraction $\alpha = 0.8$ of the traffic is split across the smartphone and CPE 700 MHz links. The remainder is duplicated on the stronger link. A coordination overhead $O = 2$ Mbps is then subtracted:

$$R_3 = \max\Big(0,\ \alpha\big(R^{0.7}_{\mathrm{p}} + R_{\mathrm{c}}\big) + (1 - \alpha)\max\big(R^{0.7}_{\mathrm{p}}, R_{\mathrm{c}}\big) - O\Big). \tag{15}$$

Because the split operates at the packet level (multi-path scheduling, in the spirit of multi-path transport, not PHY combining), no additional spectral-efficiency penalty is applied; the entire coordination cost is the explicit overhead O.

Full delegation (M4) routes everything through the CPE: the glasses offload to the phone over Wi-Fi as usual, the phone

forwards over a second Wi-Fi hop to the CPE, and the CPE alone transmits at 700 MHz, so the phone's cellular radio is switched off entirely. The two Wi-Fi hops are assumed to operate on separate channels and do not contend for airtime, so

$$R_4 = R_{\mathrm{c}}. \tag{16}$$

Because the glasses offload to the smartphone, every relayed mode is also bounded by the Wi-Fi rate. The delivered rate of each mode is therefore

$$R_{\mathrm{M0}} = R_{\mathrm{g}}, \quad R_{\mathrm{M}k} = \min\big(R_{\mathrm{w}},\ R_k\big),\ k \in \{1,2,3,4\}, \tag{17}$$

where $R_1 = R_{\mathrm{p}}^{3.5}$ and $R_2 = R_{\mathrm{p}}^{0.7}$. Modes M0, M1, and M2 are always available; M3 and M4 require the CPE, i.e. $x(t) \geq 0$.

E. Traffic, Energy, and Thermal Models

The video source uses adaptive bitrate (ABR) streaming over the ladder $\mathcal{B} = \{3, 8, 15, 25\}$ Mbps, corresponding respectively to a low-quality fallback, HD live video, high-quality HD/2K live video, and premium 2K/near-4K live video. The target rung is set to $r_{\mathrm{tgt}} = 15$ Mbps, which represents a high-quality live-streaming operating point that avoids the disproportionate thermal cost of the top rung. A rate margin $\Delta r = 3$ Mbps is imposed to reduce buffer depletion. Given the delivered rate $R_{\mathrm{del}}(t)$ of the active mode, the selected bitrate $r_v(t)$ is the highest $b \in \mathcal{B}$ satisfying $R_{\mathrm{del}}(t) \geq b + \Delta r$. The governor in Section III-F may cap the available ladder at r_{tgt}. In addition, the selected bitrate is dropped by one rung when the buffer falls below 1 s and is forced to the lowest rung when $T > T_{\max}$.

The live buffer has a capacity $B_{\mathrm{cap}} = 2.5$ s, reflecting that a live sender cannot pre-deliver an arbitrarily long future segment, and is initialized at 2.0 s. Its evolution is modeled as

$$\begin{aligned} &B_{\mathrm{uf}}(t{+}\Delta t) \\ &= \min\Big(B_{\mathrm{cap}}, \max\big(0,\ B_{\mathrm{uf}}(t) + \Delta t\big(R_{\mathrm{del}}(t)/r_v(t) - 1\big)\big)\Big), \end{aligned} \tag{18}$$

where $R_{\mathrm{del}}(t)/r_v(t) > 1$ increases the buffer and $R_{\mathrm{del}}(t)/r_v(t) < 1$ drains it. A stall occurs when $B_{\mathrm{uf}}(t) = 0$.

1) Scene complexity and content-adaptive coding: A static, low-motion first-person view reaches a given *perceived* quality at a far lower encoded bitrate, and a far lower frame rate, than a fast-moving one. We model the instantaneous scene complexity $C(t) \in [0, 1]$ as a smoothed sum of ego-motion, the O2I doorway transition, and nearby crowd motion,

$$C_{\mathrm{raw}}(t) = \mathrm{clip}\big(C_0 + w_m m(t) + w_{\mathrm{d}} \mathbf{1}\{\text{doorway}\} + w_c \rho(t),\ 0,\ 1\big), \tag{19}$$

where $m(t) = \mathbf{1}\{|x_{\mathrm{stop}} - x(t)| > 0\}$ indicates motion toward the intended stop x_{stop}. The crowd level is $\rho(t) \in \{0, 0.5, 1\}$ for an idle scene, passers-by, and a busy CPE spell. Complexity is low-pass filtered as $C(t) = a_c C(t - \Delta t) + (1 - a_c) C_{\mathrm{raw}}(t)$, where $a_c = e^{-\Delta t/\tau_c}$. We use $C_0 = 0.10$, $w_m = 0.60$, $w_{\mathrm{d}} = 0.50$, $w_c = 0.40$, and $\tau_c = 1$ s. A quiet indoor dwell gives $C \approx 0.1$. Walking, doorway crossing, or a passing crowd can drive C toward 1.

The bitrate and frame rate needed to reach the target perceived quality scale with complexity,

$$r_{\mathrm{need}}(C) = r_{\mathrm{tgt}}\big(\rho_r + (1 - \rho_r) C\big), \tag{20}$$

$$\phi_{\mathrm{need}}(C) = \phi_0 + (1 - \phi_0) C, \tag{21}$$

with floors $\rho_r = 0.40$ and $\phi_0 = 0.60$. Thus, a zero-complexity scene needs 6 Mbps at 18 fps, whereas full motion needs 15 Mbps at 30 fps. The *perceptual* rate utility in the QoE score (Section III-F) uses the weaker of the bitrate and frame-rate adequacies,

$$\begin{aligned} u\big(r_v, \phi, C\big) = \min\Big(&\min\big(1, \tfrac{r_v}{r_{\mathrm{need}}(C)}\big),\ \min\big(1, \tfrac{\phi}{\phi_{\mathrm{need}}(C)}\big)\Big) \\ &+ w_o \frac{\max(0,\ r_v - r_{\mathrm{need}}(C))}{r_{\max} - r_{\mathrm{need}}(C)}, \end{aligned} \tag{22}$$

where $\phi \in [\phi_0, 1]$ is the capture frame-rate fraction (relative to 30 fps). Under-provisioning either the bitrate or the frame rate for the current motion lowers u (visible as blocking or judder), whereas a static scene streamed at 8 Mbps and 20 fps and a moving scene at 15 Mbps and 30 fps both reach $u = 1$. For all baseline policies C is irrelevant ($\phi \equiv 1$ and the utility reduces to the saturating bitrate ratio at r_{tgt}); only the proposed policy exploits C through Mechanisms 2 and 3 below.

2) AI-glass power and temperature: With the normalized bitrate $\hat{r} = r_v / r_{\max}$ and $r_{\max} = 25$ Mbps (the top ladder rung), the glass power draw combines the camera, a superlinear encoder term, and the active radio,

$$P(t) = \phi(t)\big(P_{\mathrm{cam}} + P_{\mathrm{e0}}\big) + P_{\mathrm{e1}}\, \hat{r}^{\,1.5} + P_{\mathrm{link}}, \tag{23}$$

with $P_{\mathrm{cam}} = 0.7$ W, encoder base $P_{\mathrm{e0}} = 0.1$ W and slope $P_{\mathrm{e1}} = 1.6$ W. The exponent 1.5 reflects that encoder power scales with pixel rate, so the top rung is disproportionately expensive. The frame-rate fraction $\phi(t)$ scales the camera and per-frame encoder base (sensor and ISP capture cost), whereas the bitrate-driven encoder term and the radio are data-volume and link costs that ϕ does not change; baselines run $\phi \equiv 1$. The radio term is $P_{\mathrm{link}} = P_{\mathrm{cell}} = 2.0$ W in the direct mode M0, and $P_{\mathrm{link}} = P_{\mathrm{w0}} + P_{\mathrm{w1}}\, \hat{r}$ with $P_{\mathrm{w0}} = 0.3$ W and $P_{\mathrm{w1}} = 0.3$ W in every offloaded mode. A crucial consequence for longevity is that the *fixed* overhead $P_{\mathrm{fix}} = P_{\mathrm{cam}} + P_{\mathrm{e0}} + P_{\mathrm{w0}} = 1.1$ W is independent of bitrate; at the target rung it is 54% of the offloaded draw and at 8 Mbps it rises to 74%. The bitrate lever therefore saturates, and only reducing the capture frame rate (Mechanism 3) can attack the camera share of this fixed floor. A first-order thermal model gives the glass temperature,

$$T(t{+}\Delta t) = T_{\mathrm{amb}} + a\big(T(t) - T_{\mathrm{amb}}\big) + b\, P(t), \tag{24}$$

with ambient $T_{\mathrm{amb}} = 28$ °C, initial temperature $T_{\mathrm{amb}} + 6 = 34$ °C, and coefficients $a = 0.99933$ and $b = 0.004$ per 0.1 s step, which correspond to a thermal time constant $\tau = -\Delta t / \ln a \approx 150$ s and a steady-state thermal resistance $b/(1 - a) = 6$ °C/W. The 150 s time constant gives the glasses a realistic multi-minute warm-up rather than an instantaneous jump. The safe and emergency skin-contact limits are $T_{\mathrm{safe}} = 42$ °C and $T_{\max} = 48$ °C. Two steady states anchor the thermal story: streaming the top rung over Wi-Fi draws 3.0 W, whose steady-state temperature 46.0 °C *violates* the

safe limit, whereas the target rung draws 2.02 W and settles at 40.1 °C, safely below it. The thermal feasibility boundary therefore lies between the top two ladder rungs, which is exactly the headroom the governor exploits.

3) Smartphone power: The smartphone power is composed of additive states: a baseline $P_{\mathrm{id}} = 0.3$ W, Wi-Fi reception from the glasses $P_{\mathrm{rx}} = 0.4$ W (active in every relayed mode), the cellular transmitter $P_{\mathrm{t35}} = 2.0$ W at 3.5 GHz or $P_{\mathrm{t07}} = 1.2$ W at 700 MHz, and Wi-Fi forwarding to the CPE $P_{\mathrm{tx,w}} = 0.6$ W. The per-mode totals are

$$P_{\mathrm{ph}} = \begin{cases} 0.3\ \mathrm{W} & \mathrm{M0\ (phone\ not\ in\ the\ path)}, \\ 2.7\ \mathrm{W} & \mathrm{M1}, \\ 1.9\ \mathrm{W} & \mathrm{M2}, \\ 2.5\ \mathrm{W} & \mathrm{M3}, \\ 1.3\ \mathrm{W} & \mathrm{M4\ (cellular\ radio\ off)}. \end{cases} \tag{25}$$

M4 is the key entry: by delegating the whole uplink to the CPE, the phone drops to roughly half the draw of any mode that keeps its cellular transmitter alive. Phone energy is tracked as a KPI and enters the QoE score below.

4) Battery and longevity: Both devices carry finite batteries that drain at their instantaneous power. The AI glasses hold $E_{\mathrm{g}} = 0.8$ Wh, the order of magnitude of a current streaming-class smart-glass (a 210 mAh cell sustains about 45 min of continuous capture in measured reviews), and the smartphone holds $E_{\mathrm{ph}} = 15$ Wh (a typical 4000 mAh handset). Each device's remaining charge evolves as $E(t{+}\Delta t) = E(t) - P(t)\,\Delta t/3600$ (Wh). If the glass battery reaches zero the stream ends: that creator is removed and its QoE average is *frozen* at the moment of death, so the headline metric measures QoE while actually streaming, and battery longevity is reported separately as remaining charge and uptime. Because the glass battery is small relative to a 0.8 to 4.4 W draw, longevity becomes a first-class objective over a long dwell. A policy that streams a higher bitrate buys instantaneous quality at the cost of a shorter session, so an agent optimizing the whole session must ration energy rather than chase rate.

F. Scene-Aware Prediction and Agentic Control

The hub agent predicts indoor entry from on-device scene understanding. With the logistic function $\sigma(z) = 1/(1+e^{-z})$, the indoor probability is a sigmoid of position triggered 6 m before the entrance, passed through a 1 s inference delay τ_{inf} and additive Gaussian noise $\nu(t) \sim \mathcal{N}(0, 0.05^2)$:

$$p_{\mathrm{in}}(t) = \mathrm{clip}\Big(\sigma\big(k\,(x(t-\tau_{\mathrm{inf}}) - x_{\mathrm{trig}})\big) + \nu(t),\ 0,\ 1\Big), \tag{26}$$

with sharpness $k = 0.8$ and trigger position $x_{\mathrm{trig}} = -6$ m.

1) Thermal- and battery-aware bitrate governor: The proposed policy alone runs a governor that caps the ABR ladder with three controls. The other policies leave the ladder uncapped, except that the hard protection $T > T_{\max}$ forces the lowest rung for all policies. *(i) Predictive intent trigger.* From (24), the steady-state temperature at glass power P is $T_{\mathrm{ss}}(P) = T_{\mathrm{amb}} + \frac{b}{1-a}P$. The agent uses this equation to evaluate the future thermal state. The governor engages when four conditions hold: the user has been fully indoors for at least 10 s; the buffer is at least 80% full; the session intent is a long stream; and T_{ss} at the top rung exceeds $T_{\mathrm{safe}} - 0.5$ °C. It releases when the scene module signals departure ($p_{\mathrm{in}} < 0.4$). A myopic QoE maximizer does not make this trade. The thermal penalty is zero until T_{safe} is crossed, when heat is already stored. Baseline D tests this point and stays near its reactive throttling threshold of 41 °C during the dwell (claim 10 of Section III-I). *(ii) Reactive thermal cap.* This control also protects a long outdoor dwell. It caps the ladder at r_{tgt} when $T > T_{\mathrm{safe}} - 1$ °C and at 8 Mbps when $T > T_{\mathrm{safe}}$. *(iii) Energy-budget tier (Mechanism 1).* The fixed overhead in (23) makes the target rung the most quality-efficient operating point. Because $P_{\mathrm{fix}} = 1.1$ W is independent of bitrate, quality per joule is highest at $r_{\mathrm{tgt}} = 15$ Mbps. A lower rung loses rate utility faster than it saves power. A higher rung uses more power for the small w_o bonus. The proposed policy therefore caps the ladder at r_{tgt} and uses it until battery depletion. It does not use the reactive 8 and 3 Mbps step-downs of a naive battery saver. Operating at 3 Mbps near depletion extends runtime but provides little useful quality, which lowers the integrated useful quality of Section III-F4. At r_{tgt}, the glasses draw approximately 2.0 W. This gives an approximately 24 min ceiling without reducing target quality. Together, these controls keep the proposed glasses within the thermal limit and extend useful streaming life. Ablation C lacks the governor and does neither.

2) Predictive complexity-aware streaming (Mechanisms 2 and 3): The mapped static environment and intended trajectory let the agent predict the position-dependent components of scene complexity. The proposed policy predicts the worst-case complexity over a horizon $H = 4$ s. It follows the path $\{x(t{+}k\Delta t)\}_{k=1}^{H/\Delta t}$ and reads the deterministic ego-motion and doorway terms. The current crowd level is held as a conservative estimate. The prediction is $\hat{C}(t) = \max_k C\big(x(t{+}k\Delta t)\big)$. The agent is not clairvoyant. Gaussian estimation error gives $\tilde{C} = \mathrm{clip}(\hat{C} + \epsilon, 0, 1)$, where $\epsilon \sim \mathcal{N}(0, \sigma_p^2)$ and $\sigma_p = 0.15$. A 3 s exponential moving average suppresses per-step noise and keeps the encoder on a stable rung.

Mechanism 2 (perceptual-rate bitrate). Rather than always streaming r_{tgt}, the policy encodes only what the predicted scene needs: the ladder is additionally capped at the lowest rung $b \in \mathcal{B}$ with $b \geq r_{\mathrm{need}}(\tilde{C})$ from (20). A predicted static dwell ($\tilde{C} \approx 0.1$, $r_{\mathrm{need}} \approx 7$ Mbps) thus streams at 8 Mbps, perceptually identical to 15 Mbps by (22) but at lower encoder power, while predicted motion keeps the full 15 Mbps.

Mechanism 3 (adaptive frame rate). The capture frame-rate fraction is matched to the predicted motion, $\phi(t) = \phi_{\mathrm{need}}(\mathrm{clip}(\tilde{C} + \beta_\phi, 0, 1))$ from (21), with a small conservative bias $\beta_\phi = 0.08$ so that a misprediction errs toward smoother video rather than judder. On a static dwell this drops the camera to ≈20 fps, attacking the fixed P_{fix} overhead that the bitrate lever cannot reach; on motion it returns to 30 fps. The β_ϕ margin is deliberately decoupled from the bitrate choice, so the encoder still captures the full $8 \to 15$ Mbps energy saving while the frame rate stays on the safe side of judder.

3) Mode selection: The controller is a priority policy that implements the intent-dependent utility (1). The rate requirement is $r_{\mathrm{req}} = b_{\mathrm{top}} + \Delta r$. Here, b_{top} is the intent/thermal

cap before the additional complexity-aware encoder cap. It is 25 Mbps normally and 15 Mbps under the governor, so r_{req} is 28 or 18 Mbps:

1) indoors with the CPE available, choose among {M4, M2, M3}, ordered by *increasing phone power* (25), the first mode whose delivered rate meets r_{req}; switching to a lower-power mode additionally requires a 5 Mbps hysteresis margin to prevent load-driven ping-pong; if no mode qualifies, pick the highest-rate one (best effort);
2) else if $p_{\mathrm{in}} \geq 0.65$ or the 3.5 GHz SINR is no longer usable ($\gamma_{\mathrm{p}}^{3.5} \leq 0$), *pre-switch* to 700 MHz (M2) before the link collapses;
3) else if outdoors with a usable 3.5 GHz link, use the smartphone 3.5 GHz relay (M1).

Rule 1 is the resource-rational core: when the shared CPE has headroom, M4 wins and the phone's cellular radio sleeps; when the CPE gets busy, M4 can no longer meet r_{req} and the controller escalates to M3, whose packet split adds the phone's own 700 MHz capacity on top of the loaded CPE. All band switches use *make-before-break*: during the 1.5 s reconnection gap (plus a one-time 1.0 s setup gap the first time the CPE is engaged), the old link keeps serving, so a proactive switch costs little while a reactive switch pays the gap on an already-collapsed link. A minimum mode dwell of 2 s further damps oscillation.

4) QoE score and useful-quality longevity: The per-step QoE score balances perceptual quality against stalls, switching delay, glass radio energy, phone energy, and thermal stress:

$$\begin{aligned} q(t) = {} & w_r\, u\big(r_v, \phi, C\big)\, \mathbf{1}\{B_{\mathrm{uf}}(t) > 0\} \\ & - w_s\, \mathbf{1}\{\mathrm{stall}\} - w_d\, \frac{\tau_m}{0.2} \\ & - w_e\, \frac{P_{\mathrm{rad}}}{P_{\mathrm{cell}}} - w_{\mathrm{ph}}\, \frac{P_{\mathrm{ph}}}{P_{\mathrm{ph}}^{\max}} \\ & - w_{\mathrm{th}}\, \frac{\max(0,\, T - T_{\mathrm{safe}})}{T_{\max} - T_{\mathrm{safe}}} - w_{\mathrm{sw}}\, \mathbf{1}\{\mathrm{switch}\}, \end{aligned} \tag{27}$$

Here, $\mathbf{1}\{\cdot\}$ is the indicator function. The multiplier $\mathbf{1}\{B_{\mathrm{uf}}(t) > 0\}$ sets the quality reward to zero when the playback buffer is empty. The quality term is the *perceptual* utility $u(r_v, \phi, C)$ of (22). It saturates at 1 when the bitrate and frame rate meet $r_{\mathrm{need}}(C)$ and $\phi_{\mathrm{need}}(C)$. Exceeding these needs up to $r_{\max} = 25$ Mbps adds only the small bonus $w_o = 0.1$. For a baseline policy with $\phi \equiv 1$ and unused C, this becomes the saturating ratio at the 15 Mbps target. Perceived-quality scoring lets the proposed policy use 8 Mbps and 20 fps on a static scene without a QoE loss. It still penalizes a policy that drops to 8 Mbps during motion. The radio power is $P_{\mathrm{rad}} = P_{\mathrm{link}}$ from (23). Encoder power is already represented by the thermal term and is not counted again. The normalization is $P_{\mathrm{ph}}^{\max} = 2.7$ W, the largest per-mode phone power (M1). The per-mode switching delays are $\tau_m \in \{0.02, 0.03, 0.03, 0.05, 0.05\}$ s for M0 through M4. The weights are $w_r = 1$, $w_o = 0.1$, $w_s = 4$, $w_d = 1.5$, $w_e = 0.8$, $w_{\mathrm{ph}} = 0.6$, $w_{\mathrm{th}} = 3$, and $w_{\mathrm{sw}} = 0.5$. Their order is $w_s > w_{\mathrm{th}} > w_d > w_r > w_e > w_{\mathrm{ph}} > w_{\mathrm{sw}}$. A stall has the largest penalty, consistent with ITU-T P.1203 and streaming-QoE measurements. Exceeding the on-face skin-safety limit has the next-largest penalty because it is a hard wearable constraint. The score accumulates only while the device is alive and freezes after battery depletion, as stated in Section III-E4.

Because average QoE alone is blind to how *long* a creator keeps delivering, and can even reward dying early to stop accumulating a slightly negative score, we adopt a longevity-aware figure of merit, the *useful quality-seconds*, that integrates the *saturating* perceived quality over the watchable session,

$$Q_{\mathrm{use}} = \int_0^{T_{\mathrm{ep}}} \tilde{u}\big(r_v, \phi, C\big)\, \mathbf{1}\{\mathrm{alive}\}\, \mathbf{1}\{B_{\mathrm{uf}} > 0\}\, dt, \tag{28}$$

where $\tilde{u} = \min\big(\min(1, r_v/r_{\mathrm{need}}(C)),\, \min(1, \phi/\phi_{\mathrm{need}}(C))\big) \in [0, 1]$ is the min-of-adequacies term of (22) *without* the over-target bonus, so the quality credit caps at 1 per second and Q_{use} can never exceed the alive, unstalled time. This single number captures both of the stated goals: it rewards higher perceived quality and a longer stall-free life, it saturates the quality credit at the target (so it does not over-reward chasing 4K), and it is never increased by dying sooner. It directly encodes the value the proposed policy adds, namely sustaining target-quality video for longer, and is used as the headline verdict alongside average QoE in Section III-I.

G. Baseline Policies

The proposed policy is compared against two fixed-or-reactive baselines, one ablation, and one capability-matched myopic controller, all evaluated on the *same* precomputed channel and CPE-load traces for a paired comparison. (Earlier revisions also included a fixed-3.5 GHz relay and a fixed-700 MHz relay; these were transitional variants of the cases below, and since a real handset already band-switches they are dropped.)

- **A**: AI-glass direct 3.5 GHz (M0 fixed), the worst thermal case;
- **B**: reactive 700 MHz switch on a SINR threshold, where the smartphone starts on the 3.5 GHz relay (M1) and moves to 700 MHz (M2) only after an exponentially smoothed $\gamma_{\mathrm{p}}^{3.5}$ falls below -5 dB (smoothing time constant 3.5 s);
- **C (ablation)**: the identical scene-aware mode orchestration as the proposed policy, including M4 delegation, but *without* the predictive energy-aware extensions, that is, no thermal/energy governor and no complexity-aware bitrate or frame-rate adaptation, so its ABR keeps chasing the top rung at full frame rate indoors. C isolates the value of the predictive control plane from that of the mode orchestration it shares with the proposed policy.
- **D (capability-matched myopic)**: a non-agentic controller with the *same observation and action space* as the proposed policy. It reads every current measurement the proposed policy reads (per-mode link rates, CPE availability and load, glass temperature, buffer, batteries, and the current scene complexity) and commands the same actions (all modes M0–M4, the bitrate-ladder cap, and the capture frame rate), but it selects, each step, the mode that maximizes the *current-step* QoE estimate, with

no horizon prediction, no intent trigger, and no thermal-trend reasoning. Its protection is reactive only: the same thermal caps as the proposed policy's reactive tier (r_{tgt} above 41 °C, 8 Mbps above 42 °C) and the naive battery step-downs (8 Mbps below 25%, 3 Mbps below 10%); its frame rate follows the *current* complexity (a one-step maximizer is indifferent above the perceptual need, so it breaks the tie toward lower power), with no lookahead and no conservative bias. Predictive observations such as p_{in} carry no value for a one-step objective, so a rational myopic controller ignores them. D answers the question of whether the proposed policy's gains come merely from richer observations and a larger action space: any gap between D and the proposed policy is attributable to prediction, intent, and planning alone.

H. Parameter Settings

Table III lists all parameters with their provenance. The path-loss model, noise floor, O2I penetration, UE power class, Wi-Fi and NR bandwidths, and skin-contact thermal limits are standards-based, whereas device power, thermal dynamics, cooperation overhead, control thresholds, and QoE weights are illustrative.

I. Results and Discussion

We report results at two levels. Table IV and Figs. 1–5 document the full proposed controller (all predictive mechanisms of Section III-F2 enabled, non-ideal prediction $\sigma_p = 0.15$) against all four baselines, including the capability-matched myopic controller D, on the canonical 900 s walk-in-and-dwell episode at seed 7, across the O2I transition and the indoor dwell. Table V then evaluates the same controller against all four baselines, including D, over four trajectory families and 12 seeds each, where the small glass battery is exercised to depletion, scored by perceived quality (22) and useful quality-seconds (28).

Table IV reports the key metrics over the 900 s episode (random seed 7). Figure 1 shows the delivered bitrate and smoothed QoE across the O2I transition and the long indoor dwell, and Figs. 2 to 5 show the mode timeline, glass temperature, battery drain, and the aggregate comparison. The proposed policy spends 26 s in M1, 9 s in M2, and 865 s in M4; because the dwell lies inside the dead-zone, where the phone's 700 MHz leg is blocked and the packet split (M3) collapses onto pure delegation, it delegates fully to the CPE throughout the indoor dwell.

The results support eleven claims, all consistent with the physical setup and robust to the illustrative parameters.

1) **The capability constraint is enforced.** The AI glasses never transmit at 700 MHz, and only the smartphone and CPE do. Mode M0 exists solely as a thermal- and battery-worst fallback.
2) **Direct cellular on the glasses overheats and drains fastest.** Policy A reaches 45.1 °C, above the 42 °C safe limit, and is down to 8.8% battery by $t = 900$ s, the lowest of all, because the glasses both encode and transmit at high power. Offloading the uplink off the glasses is therefore a thermal and battery necessity, not just a capacity choice.
3) **Direct 3.5 GHz on the glasses stalls through the boundary.** Policy A loses the link on entry and stalls for about 864 s of the 900 s episode, collapsing to the lowest bitrate and the worst QoE (-5.84); only offloading the uplink and adapting the band avoids it.
4) **A reactive radio cannot escape the macro dead-zone.** Two limits compound here. First, even where its macro link survives, 700 MHz alone cannot meet the target: with 18 dB penetration and only a 10 MHz per-user share, the reactive policy (B) is held to the lowest rungs (average delivered ≈ 3.8 Mbps), never reaching the 15 Mbps target. Second, inside the shielded pocket the macro link is blocked outright (8), and because B's band-switch only ever reaches the smartphone 700 MHz and never the CPE, it stalls for 860 s of the dwell (QoE -4.57), almost as badly as the direct-3.5 GHz baseline A. A reactive band-switch is helpless in a coverage hole; escaping it needs a proactive delegation to alternative infrastructure (the CPE), which the scene-aware policies C and P make and which carries them through the hole with zero stalls.
5) **Proactive beats reactive because of make-before-break.** The proposed policy leaves 3.5 GHz 4.5 s before entry (the scene prediction and an early shadowing dip both push it off the band), whereas the reactive policy (B) switches 5.5 s after entry, paying the reconnection gap on a collapsing link. The proposed transition minimum QoE (-1.23) is a brief dip, whereas the direct-3.5 GHz policy A falls to -5 and stays there.
6) **The CPE makes the target bitrate available, and the encoder uses only what the scene needs.** The proposed policy spends 865 s on the CPE (here almost entirely on M4, since the dead-zone blocks the phone's 700 MHz and collapses the M3 split onto pure delegation), whose boosted backhaul can sustain the 15 Mbps target throughout the dwell, headroom the form-factor-limited cluster cannot reach on the phone's 700 MHz link alone. The encoder then streams at the target during motion and the busy spells and content-adaptively at 8 Mbps on the static dwell (Mechanism 2), with zero stalls and only brief dips at mode-switch instants, so the average delivered bitrate (8.8 Mbps) sits below the target while the perceived quality stays at it.
7) **Delegation puts the phone's cellular radio to sleep.** Whenever the shared CPE has headroom the controller picks M4 (1.3 W phone draw) over M3 (2.5 W) and M2 (1.9 W). Inside the dead-zone, where the phone's own 700 MHz is unusable anyway, full delegation is also the only mode that keeps the stream alive. The proposed policy spends its whole indoor dwell in M4 and uses 1214 J of phone energy, against 1740 J for the reactive baseline B, which keeps its 700 MHz transmitter running uselessly into the dead-zone. The ablation C, likewise forced onto the CPE, draws the same 1214 J, so among the streaming policies delegation roughly halves the phone radio cost relative to keeping a cellular transmitter

TABLE III
SCENARIO 1 PARAMETERS AND THEIR PROVENANCE. ILLUSTRATIVE VALUES ARE ENGINEERING CHOICES; CONCLUSIONS DEPEND ON THEIR ORDERING, NOT THEIR EXACT MAGNITUDES.

Group / parameter	Value	Basis
Geometry and mobility		
gNB position x_{g} / height h	-100 m / 25 m	3GPP TR 38.901 UMa geometry (10 m to 5 km valid range)
User speed v / start x_0 / stop x_{stop}	1 m/s / -30 m / 10 m	Pedestrian; entrance at $t = 30$ s, dwell from $t = 40$ s
Episode length T_{ep} / step Δt	900 s / 0.1 s	Long dwell exposes thermal ($\tau \approx 150$ s) and battery drain
CPE position x_{c} / discovery	15 m / from entrance	Illustrative (near-window placement)
Radio and channel		
Bandwidths $B_{\mathrm{w}}/B_{3.5}/B_{0.7}$	80 / 50 / 10 MHz	IEEE 802.11ax; 3GPP NR n78 / n28 (cellular = per-user share, half the carrier)
Noise figure NF	7 dB	Typical UE receiver
Efficiency $\eta_{\mathrm{c}}/\eta_{\mathrm{w}}$	0.55 / 0.50	Illustrative (Shannon-to-implementation gap); M3 split is packet-level, no extra PHY penalty
Wall loss $L_{\mathrm{w}}^{3.5}/L_{\mathrm{w}}^{0.7}$	28 / 18 dB	3GPP TR 38.901 §7.4.3 high-loss building (consistent)
Doorway peak L_{door} (3.5 GHz)	8 dB transient	Illustrative (body and door-frame diffraction)
Indoor clutter	0.2 dB/m, cap 4 dB	Illustrative (cf. 3GPP O2I indoor loss)
Shadowing $\sigma_{3.5}/\sigma_{0.7}$	4 / 3 dB, corr. 10 m	3GPP TR 38.901 UMa NLOS (conservative)
Temporal fading σ_{t} / corr. time	1.5 dB / 2 s	Illustrative (people moving near the user)
Interference margin I_{700}	6 dB	Illustrative (busier low band)
Body loss $L_{\mathrm{b}}^{\mathrm{p}}/L_{\mathrm{b}}^{\mathrm{g}}$	3 / 2 dB	Illustrative
CPE wall loss L_{c}	12 dB, static link	Illustrative (fixed near-window CPE)
RF dead-zone $[x_{\mathrm{d0}}, x_{\mathrm{d1}}]$ / loss L_{d}	$[8, 14]$ m / 60 dB	Illustrative (shielded interior pocket; macro only, CPE exempt)
Wi-Fi SINR γ_{w} (out / in)	25 / 22 dB	Illustrative (short link, not the bottleneck)
Devices		
Tx power $P_{\mathrm{p}}/P_{\mathrm{c}}/P_{\mathrm{g}}$	23 / 23 / 20 dBm	3GPP TS 38.101 Power Class 3 (glass reduced)
Antenna gain $G_{\mathrm{g}}/G_{\mathrm{p}}/G_{\mathrm{c}}$	-2 / 0 / 10 dBi	Illustrative (form-factor dependent; CPE directional panel)
Glass power $P_{\mathrm{cam}}/P_{\mathrm{e0}}/P_{\mathrm{e1}}/P_{\mathrm{w0}}/P_{\mathrm{w1}}/P_{\mathrm{cell}}$	0.7/0.1/1.6/0.3/0.3/2.0 W	Illustrative; superlinear encoder (exponent 1.5), Eq. (23)
Phone power $P_{\mathrm{id}}/P_{\mathrm{rx}}/P_{\mathrm{t35}}/P_{\mathrm{t07}}/P_{\mathrm{tx,w}}$	0.3/0.4/2.0/1.2/0.6 W	Illustrative; per-mode totals 0.3 to 2.7 W, Eq. (25)
Thermal a / b	0.99933 / 0.004	Illustrative ($\tau \approx 150$ s, 6 °C/W)
Thermal limits $T_{\mathrm{amb}}/T_{\mathrm{safe}}/T_{\max}$	28 / 42 / 48 °C	IEC 62368-1 / EN 563 skin-contact limits
Battery E_{g} / E_{ph}	0.8 / 15 Wh	Streaming-class glasses ($\sim$210 mAh); 4000 mAh phone
CPE grant C_{tot} / bg. load / mean dwell	40 / $\{5, 25\}$ Mbps / 30 / 15 s	Illustrative (shared venue CPE, two-state Markov load)
Coop overhead O / setup / α	2 Mbps / 1 s / 0.8	Illustrative (Option C packet split)
Traffic and control		
Bitrate ladder $\mathcal{B}$ / target r_{tgt}	$\{3, 8, 15, 25\}$ / 15 Mbps	Illustrative (720p to 4K-live ladder)
Buffer init / cap / low	2.0 / 2.5 / 1.0 s	Low-latency live streaming
p_{in} threshold	0.65	Illustrative
Reactive threshold / smoothing	-5 dB / 3.5 s	Illustrative (baseline B)
Min dwell / switch gap / power hysteresis	2.0 s / 1.5 s / 5 Mbps	Illustrative (gap $\gg$ real NR sub-100 ms, conservative)
Scene trigger x_{trig} / k / τ_{inf}	-6 m / 0.8 / 1 s	Illustrative (on-device inference latency)
Governor intent trigger: dwell / buffer / margin / exit	10 s / 80% / 0.5 °C / p_{in}<0.4	Illustrative (predictive thermal governor)
Governor reactive thermal cap	r_{tgt}@41°C, 8 Mbps@42°C	Illustrative (location-agnostic thermal tier)
Energy-budget tier (Mech. 1)	cap at r_{tgt}, ride to death	Illustrative (max quality-per-joule under fixed P_{fix})
Scene complexity $C_0/w_m/w_{\mathrm{d}}/w_c/\tau_c$	0.10/0.60/0.50/0.40/1 s	Illustrative (ego-motion, doorway, crowd)
Perceptual floors ρ_r/ϕ_0	0.40/0.60	Illustrative (content-adaptive coding, 15 fps min)
Prediction horizon / noise / smoothing / fps bias	H=4 s / σ_p=0.15 / 3 s EMA / β_ϕ=0.08	Illustrative (non-ideal trajectory/complexity prediction)
QoE weights $w_{r,s,d,e,\mathrm{ph,th,sw}}$ / w_o	$1, 4, 1.5, 0.8, 0.6, 3, 0.5$ / 0.1	Illustrative; ordering per ITU-T P.1203, thermal a hard wearable constraint

alive. The instantaneous M4 saving is 48% relative to M3.

8) **The agent adapts to shared infrastructure.** Where the phone's 700 MHz leg is available (outside the dead-zone, as in the indoor-pacing family of Table V), the two-state CPE load drives an M3/M4 alternation: under light load (35 Mbps allocatable) the cluster delegates fully (M4), and when other users occupy the CPE (15 Mbps) the controller escalates to the packet split (M3), adding the phone's own 700 MHz capacity through the busy bursts. In the canonical dwell this lever is unavailable, because the dead-zone blocks the 700 MHz leg, so the cluster relies on the CPE backhaul alone, which is precisely why a CPE strong enough to carry the target on its own (the 10 dBi directional panel) matters.
9) **The governor wins on QoE, safety, and longevity at once.** Without it (ablation C), chasing the top rung on the strong CPE backhaul pushes the glasses past the safe limit during the dwell (peaking at 44.2 °C, Fig. 3) and drains the glass battery to 21.2% by $t = 900$ s (Fig. 4). With the full predictive controller, the proposed policy holds 36.1 °C (safe), retains 58.2% battery (about

TABLE IV

SCENARIO 1 RESULTS FOR THE *full* PROPOSED CONTROLLER (ALL PREDICTIVE MECHANISMS OF SECTION III-F2 ENABLED; QOE IS THE PERCEPTUAL UTILITY (22)), SEED 7, 900 S EPISODE, ENTRANCE AT $t = 30$ S, INDOOR DWELL FROM $t = 40$ S. "AVG BITRATE" IS THE AVERAGE *delivered video* BITRATE (THE PROPOSED POLICY ENCODES AT 8 MBPS ON THE STATIC DWELL, SO ITS AVERAGE SITS BELOW THE 15 MBPS TARGET WHILE ITS PERCEIVED QUALITY STAYS AT THE TARGET). "GLASS" AND "PHONE" ARE THE BATTERY CHARGE REMAINING AT $t = 900$ S. THE BEST PRACTICAL POLICY IS IN BOLD. "SWITCH LEAD" IS THE TIME BY WHICH THE CELLULAR BAND CHANGES BEFORE ENTRY, SO POSITIVE IS PROACTIVE AND NEGATIVE IS REACTIVE; A NEVER LEAVES 3.5 GHZ (M0), B SWITCHES REACTIVELY AFTER THE LINK HAS ALREADY DEGRADED, AND THE MYOPIC CONTROLLER D SHOWS +30 S ONLY BECAUSE ITS PER-STEP POWER ACCOUNTING PREFERS THE (OUTDOORS, AMPLY SUFFICIENT) 700 MHZ RELAY FROM THE START, NOT BECAUSE IT PREDICTS THE ENTRY. THE INDOOR DWELL FALLS INSIDE THE RF DEAD-ZONE OF (8), WHICH BLOCKS THE MACRO 3.5/700 MHZ BANDS, SO THE CELLULAR-ONLY BASELINES A AND B STALL THERE, WHEREAS THE THREE CPE-CAPABLE POLICIES (C, D, AND P) DELEGATE TO THE IN-VENUE CPE AND STREAM THROUGH. D NOMINALLY EDGES THE PROPOSED POLICY ON AVERAGE QOE (+0.12 VERSUS +0.10, A MARGIN TRACEABLE TO THE SMALL OVER-TARGET BONUS w_o); THE PROPOSED POLICY IS BOLDED ON THE RESOURCE COLUMNS BECAUSE THAT IS WHERE THE POLICIES ACTUALLY SEPARATE (CLAIM 10).

Policy	Avg bitrate (Mbps)	Stall (s)	Avg QoE	Min QoE (transition)	Max temp (°C)	Glass (%)	Phone (%)	Switch lead (s)
A: glass 3.5 direct	3.75	864	−5.84	−5.02	45.1	8.8	99.5	n/a
B: reactive 700 switch	3.78	860	−4.57	−4.78	36.3	60.4	96.8	−5.5
C: scene-aware w/o governor	19.65	0	−0.08	−1.23	44.2	21.2	97.8	+4.5
D: myopic full-observation	18.60	0	+0.12	−1.23	41.0	31.8	97.8	+30.0
Proposed: scene-aware UE-CoMIMO	**8.78**	**0**	+0.10	−1.23	**36.1**	**58.2**	97.8	+4.5

2.8 times C's remaining battery charge), *and* achieves a higher average QoE than C (+0.10 versus −0.08). Beyond the safe limit the thermal penalty outweighs C's bitrate advantage, and the complexity-aware coding further cools and powers down the glasses at no perceptual cost, so the controller yields a net QoE win while respecting the wearable's hard thermal and battery constraints. A myopic QoE maximizer cannot make this trade, because the thermal penalty activates only after the limit is crossed and the battery cost is paid only later; baseline D verifies this mechanism empirically in the next claim.

10) **Adaptivity with full information matches the instantaneous score; prediction and intent buy the resource trajectory.** The capability-matched myopic controller D is the direct test of whether the proposed gains come merely from richer observations and actions. Holding the same toolbox, D eliminates every stall and even discovers low-band relaying and full delegation on its own (it spends 35 s in M2 and 865 s in M4, never touching M1, because its per-step power accounting prefers the lower-power modes whenever they meet the demand). On average QoE it is statistically indistinguishable from the proposed policy (+0.12 versus +0.10; the nominal edge traces to the small over-target bonus w_o marginally outweighing the radio-energy penalty) and its useful quality-seconds tie (836 both). The separation is entirely in the resource trajectory: D streams 18.6 Mbps where 8.8 Mbps yields the same perceived quality, holds 41.0 °C for the whole dwell by riding its *reactive* thermal cap one degree below the skin-contact limit, and ends the episode at 31.8% glass battery against the proposed 58.2%. On the 2400 s multi-trajectory batch of Table V, D's glasses die at 1662 s (≈28 min) in the walk-in family, its reactive battery saver stretching the tail at near-useless rates, whereas the proposed policy streams to 2179 s (≈36 min) and delivers 2012 against 1369 useful quality-seconds; across all four families the quality-seconds gap is 22–47%. The comparison isolates the agentic contribution cleanly: cooperation eliminates the stalls, full-information adaptivity matches the instantaneous score, and prediction and intent are what make the same score affordable over the session.

11) **Quality trades against session longevity on a fixed energy budget.** Glass battery at $t = 900$ s orders the streaming-quality policies by their encoder load: the ablation C (≈ 20 Mbps) ends at 21.2%, while the proposed policy (content-adaptive, 8.8 Mbps average) ends at 58.2% (Fig. 4). The reactive baseline B coasts to ≈ 60% only because it is stalling at near-zero delivered rate in the dead-zone, not because it is streaming efficiently. Streaming more bits than the scene needs buys nothing in perceived quality but shortens the session, which is precisely why an agent optimizing the whole session rations energy through the energy budget and the complexity-aware bitrate and frame-rate, spending it only where the picture benefits.

These claims are qualitative. The reported numbers correspond to the fixed seed 7; the conclusions follow from designed regime boundaries (the penetration budget, the two CPE load states, the thermal steady states straddling T_{safe}, and the small glass battery relative to the streaming power) rather than from a particular noise realization. The accompanying interactive implementation reproduces the same ordering for any walk the user draws with the mouse.

J. Multi-Trajectory Evaluation of the Predictive Extensions

The canonical episode above uses one walk-in trajectory. We also evaluate all five policies over four trajectory families. Each family uses seeds $\{1, \ldots, 12\}$ and an episode length of $T_{\text{ep}} = 2400$ s so that the small glass battery can deplete. The *walk-in & dwell* family moves from $x_0 = -30$ m to $x = 10$ m and then stays there. The *indoor↔outdoor shuttle* family moves to $x = 10$ m, returns to $x = -20$ m, and repeats. It crosses the O2I boundary and dead-zone on each pass. The *outdoor pacing* family moves between $x = -30$ and $x = -6$ m and never reaches the CPE. The *indoor pacing*

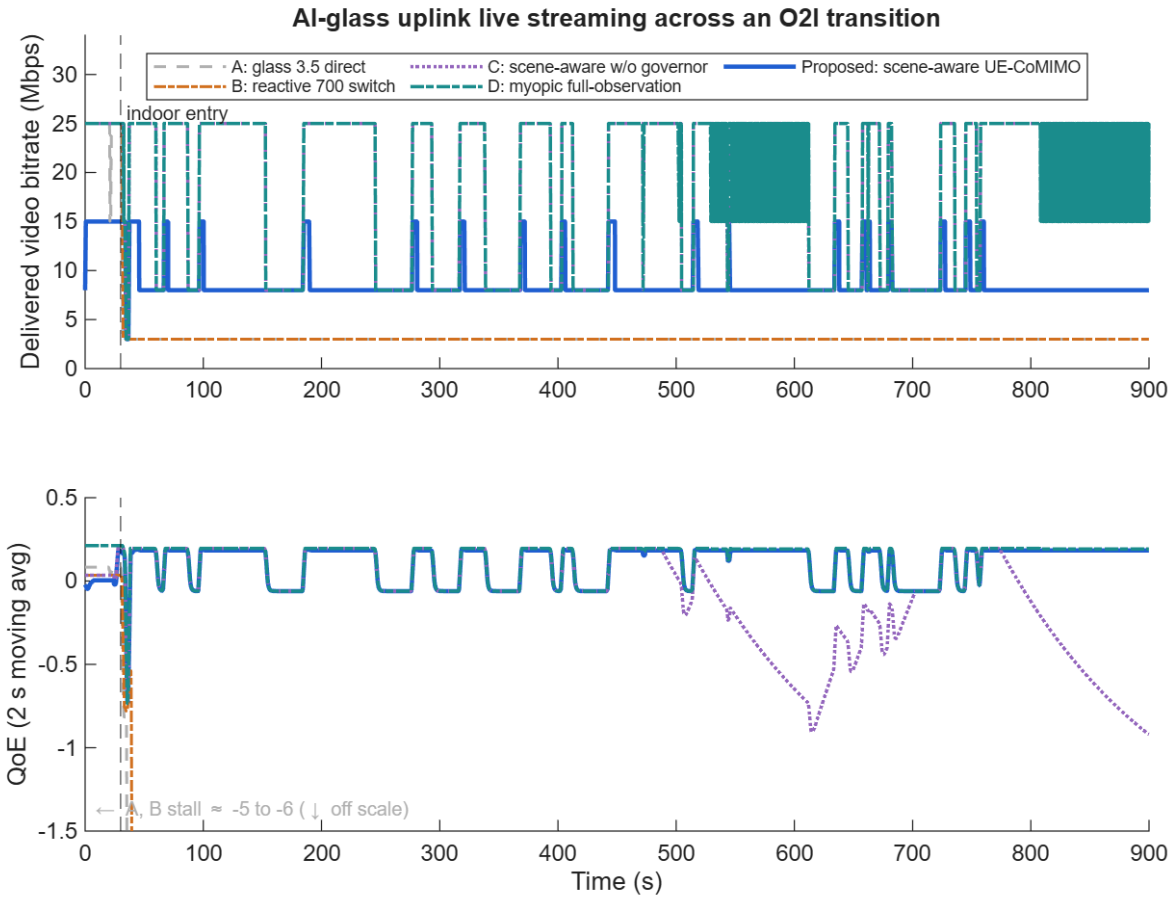


Fig. 1. Scenario 1: delivered video bitrate (top) and smoothed QoE (bottom) across the O2I transition (entrance at $t = 30$ s) and the long indoor dwell. The proposed policy pre-switches before entry, engages the shared CPE, and then encodes content-adaptively: it streams at the 15 Mbps target during motion and the busy CPE spells but drops to 8 Mbps on the static dwell (perceptually equivalent, Section III-F2), so its delivered-bitrate curve alternates between the two levels rather than sitting at the top rung, and its QoE (lower panel) stays at the top throughout. The ablation C keeps chasing the top rung on the CPE; as its glasses heat past the safe limit, its QoE drifts downward over the dwell and falls below the proposed policy, while A and the reactive baseline B both sit at the bottom, stalling in the RF dead-zone that blocks their macro link. The capability-matched myopic controller D also streams stall-free through the dead-zone and tracks the proposed policy's QoE closely, but it does so at the top permitted rung, spending more than twice the encoder power and riding its reactive thermal cap at 41 °C (claim 10).

family starts at $x = 10$ m and moves between $x = 2$ and $x = 16$ m, repeatedly crossing the dead-zone of (8). Prediction is non-ideal throughout ($\sigma_p = 0.15$). Table V reports arithmetic means over the 12 seeds. QoE is rounded to two decimals. Uptime and useful quality-seconds are rounded to the nearest integer.

Four observations follow.

(i) The proposed policy leads on quality-seconds everywhere, including against the capability-matched controller. P attains the highest mean quality-seconds in all four trajectory families and the highest mean QoE in all four (tying D only in outdoor pacing, where prediction has nothing to exploit; its one negative mean QoE is a marginal -0.03 in the relentlessly indoor-pacing case). Its nearest quality-seconds competitor everywhere is now the myopic controller D, not the ablation, and P still delivers 22 to 47% more useful quality-seconds (2012 versus 1369 in the walk-in family alone).

(ii) Raw uptime can be misleading; useful quality-seconds measures delivered service. The reactive baseline B often survives *longest* (up to 2349 s) yet collapses on quality-seconds, because in the venue-entering families it stalls in the macro dead-zone it cannot escape: 33 quality-seconds in the walk-in case and 638 in indoor pacing, against P's 2012 and 1334, while policy A stalls even harder (as few as 1 quality-second). D also outlives P in the shuttle and outdoor families (1524 versus 1500 s and 1516 versus 1476 s), but only by limping at its reactive battery saver's lowest rungs near end-of-life, and it still trails P by 22–28% quality-seconds there. Surviving longer on a dead or near-useless link is worthless; the energy-budget tier (Mechanism 1) instead rides the target rung to death on a link that is actually delivering.

(iii) The predictive extensions convert static time into longevity. In the static-heavy walk-in family, the ablation C, which shares P's mode orchestration but chases the top rung at full frame rate, dies at 1164 s with 1073 quality-seconds; the myopic controller D, protected by reactive caps only, reaches

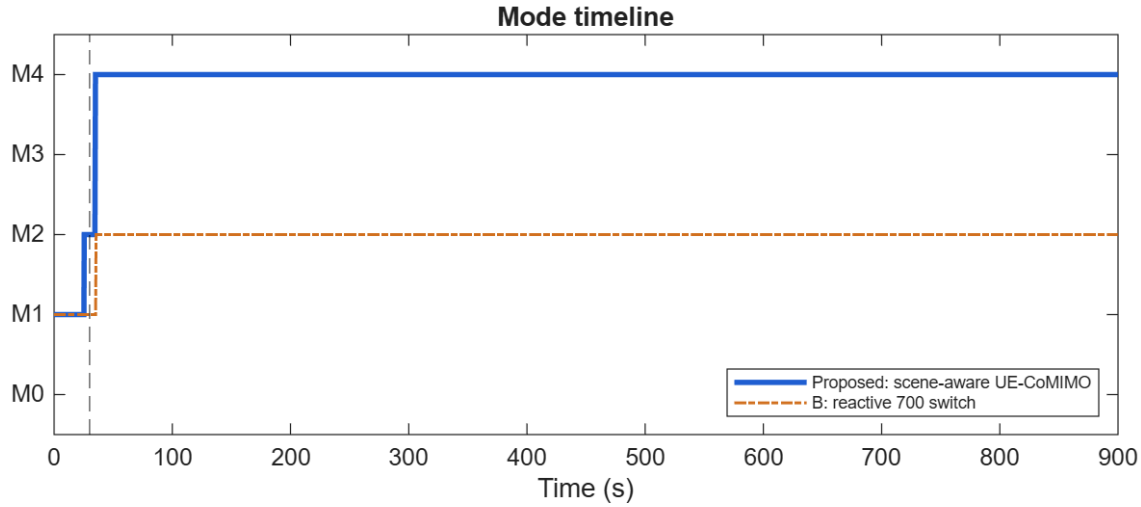


Fig. 2. Scenario 1 mode timeline. The proposed policy pre-switches 3.5 → 700 MHz before the entrance, then settles into full delegation (M4, phone cellular off) for the indoor dwell: the dwell lies inside the RF dead-zone, which blocks the phone's 700 MHz leg and so collapses the packet-split cooperation (M3) onto pure CPE delegation. The reactive baseline switches only after the link has degraded, stays on M2, and then stalls when the dead-zone blocks 700 MHz. (Where the macro 700 MHz survives, e.g. the indoor-pacing family of Table V, the proposed policy instead alternates M3/M4 with the CPE load.)

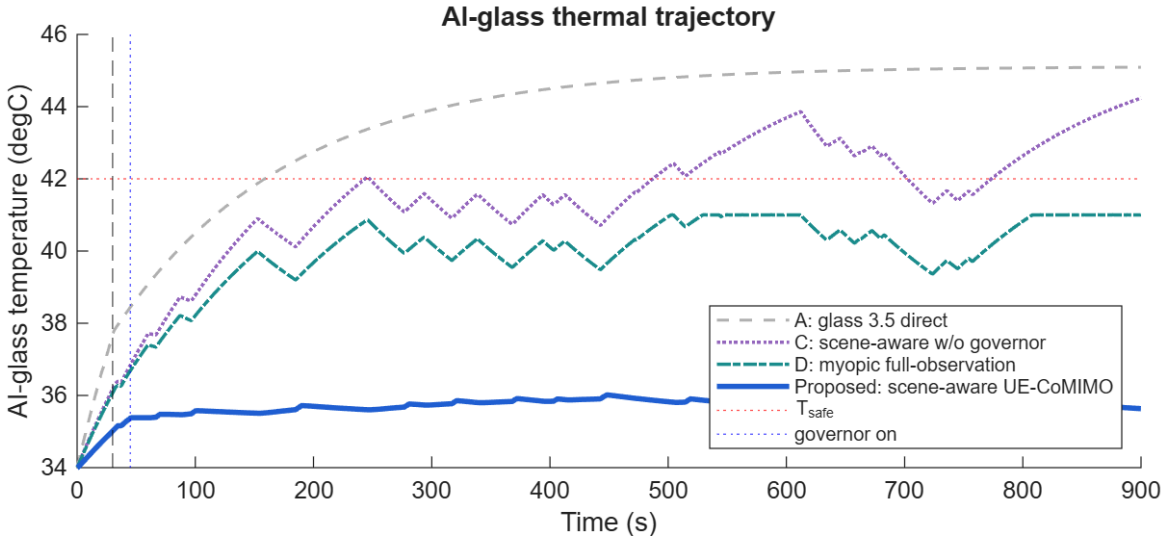


Fig. 3. Scenario 1 AI-glass temperature. Direct cellular (A) climbs toward 45 °C, well above the 42 °C safe limit. The ablation C (no governor) crosses the limit partway through the dwell while chasing the top bitrate rung on the CPE and then oscillates above it as the shared CPE load forces intermittent rate drops, peaking at 44.2 °C. The myopic controller D neither predicts nor pre-empts: it heats until its reactive cap engages at 41 °C and then oscillates just below the safe limit for the rest of the dwell. The proposed policy activates the governor (vertical dotted line) once a long stable indoor dwell is recognized, and the complexity-aware coding then drops the encoder to 8 Mbps and 20 fps on the static dwell, so the temperature settles around 36 °C, comfortably below the limit.

1662 s and 1369 quality-seconds; P, encoding 8 Mbps at ≈20 fps on the predicted-static dwell, reaches 2179 s and 2012 quality-seconds. Reactive protection thus recovers part of the gap and prediction accounts for the rest, attacking the fixed P_{fix} overhead that bitrate alone cannot. Even in the motion-dominated outdoor-pacing family, P's quality-seconds (1475) nearly equal its uptime, that is, it streams at target quality essentially until battery death.

(iv) The advantage comes from environment structure, not clairvoyance. Glass power depends only on bitrate and frame rate (23), and only the proposed policy adapts both to *predicted* scene complexity. D adapts the frame rate to current complexity but chases the top rung. C chases high bitrate rungs at full frame rate on static scenes, whereas A and B are capacity-limited. The ordering between P and A–C is unchanged as the prediction noise is raised to $\sigma_p = 0.3$ (verified on the three-baseline batch), and D is prediction-free by construction, so the comparison does not hinge on clairvoyance.

K. Limitations

This is a system-level study built on standard link budgets and channel models. It does not re-run a physical-layer MIMO simulation, and the cooperative-rate expressions are tractable abstractions rather than detailed waveform results. The glass and phone power and battery figures are illustrative additive

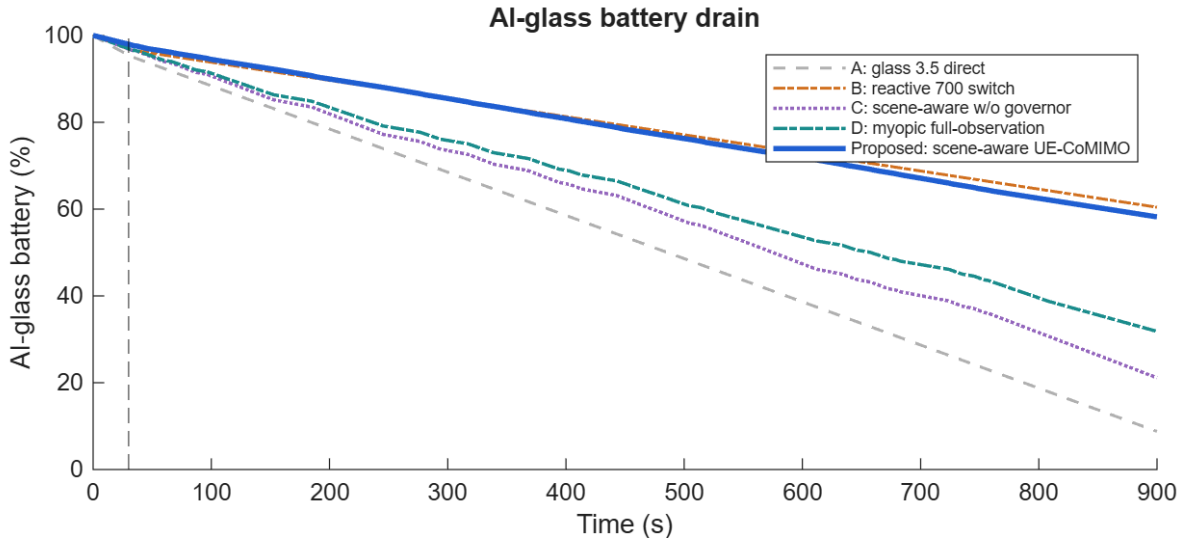


Fig. 4. Scenario 1 AI-glass battery drain. All glasses start full; the slope is set by the streaming power. Direct cellular (A) and the no-governor ablation (C) drain fastest, ending at 8.8% and 21.2%; the myopic controller D, throttled only by its reactive caps, ends at 31.8%. The proposed policy ends at 58.2%, retaining about 2.8 times C's remaining battery charge at $t = 900$ s, because it caps the bitrate at the target and then, on the static dwell, encodes content-adaptively at 8 Mbps and 20 fps (Mechanisms 2 and 3), cutting the encoder and camera power at no perceptual cost. The reactive baseline B coasts longest ($\approx 60\%$) only because it is stalling at near-zero rate in the macro dead-zone rather than streaming, illustrating that raw battery life on a dead link is no substitute for sustained useful quality.

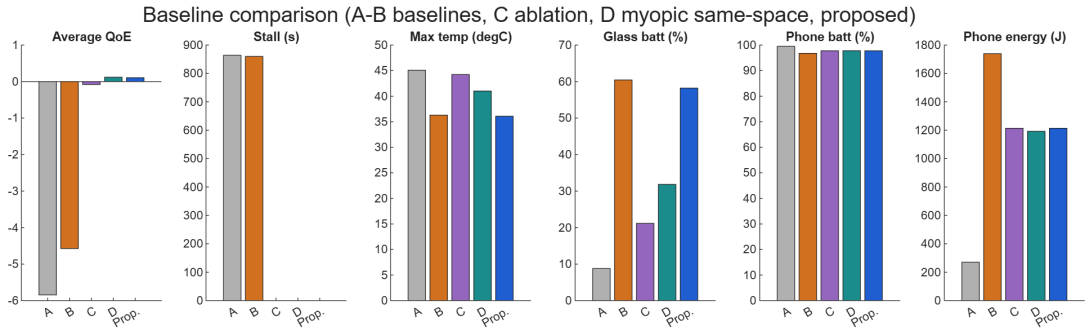


Fig. 5. Scenario 1 aggregate comparison across policies (average QoE, stall time, maximum glass temperature, glass and phone battery remaining, and phone energy). The proposed policy and myopic controller D are both stall-free and remain within the thermal limit, but the proposed policy retains substantially more glass battery.

device models, not measured numbers, and the QoE weights encode a design preference rather than a calibrated user study. In particular, the thermal weight makes the skin-safety limit a hard constraint, and the battery-death rule freezes rather than penalizes the QoE average. The scene-complexity model, the linear bitrate- and frame-rate-versus-complexity needs of (20)–(21), and the min-of-adequacies perceptual utility of (22) are deliberately simple stand-ins for a perceptual encoder model and a measured rate-distortion-versus-motion curve. Their role is to expose the qualitative lever (a static scene is cheap to encode and to capture), and the frame-rate power scaling in (23) likewise charges the camera and ISP per captured frame without modeling the full sensor pipeline. The trajectory and complexity prediction is idealized as the true near-future path plus Gaussian error. A deployed system would estimate it from motion sensors and on-device scene understanding, and the robustness to σ_p is the relevant guard. The shared-CPE background load is a two-state Markov abstraction of a much richer multi-user process, and the two Wi-Fi hops of M4 are assumed not to contend for airtime, which holds only when separate channels are provisioned. The make-before-break gap is deliberately conservative relative to standardized NR handover interruption. The battery-longevity findings assume the glasses run on their internal cell. In practice, an external accessory (for example a magnetic power capsule of several Wh) would decouple session length from the internal battery and let the agentic QoE advantage persist indefinitely, and the interactive implementation includes such an accessory as a toggle. These choices illustrate how an agentic controller converts physical-layer headroom into sustained user-perceived performance. The absolute values are not device specifications.

TABLE V
MULTI-TRAJECTORY EVALUATION OF ALL FIVE POLICIES, 12 SEEDS PER CELL, EPISODE $T_{\mathrm{ep}} = 2400$ S, NON-IDEAL PREDICTION ($\sigma_p = 0.15$), OVER THE FOUR EXPLICITLY SCRIPTED TRAJECTORY FAMILIES DEFINED IN THE TEXT. METRICS ARE THE MEANS OF AVERAGE QOE, GLASS UPTIME (TIME TO BATTERY DEATH, CAPPED AT T_{ep}), AND USEFUL QUALITY-SECONDS (28); THE PROPOSED POLICY (P) IS HIGHLIGHTED. P ATTAINS THE HIGHEST MEAN QUALITY-SECONDS IN EVERY FAMILY AND THE HIGHEST MEAN QOE IN EVERY FAMILY (TYING THE CAPABILITY-MATCHED MYOPIC CONTROLLER D ONLY IN OUTDOOR PACING, WHERE PREDICTION HAS NOTHING TO EXPLOIT).

Trajectory	Metric	A	B	C	D	P
Walk-in & dwell	Avg QoE	-5.91	-4.70	-0.15	$+0.01$	$\mathbf{+0.10}$
	Uptime (s)	989	2349	1164	1662	**2179**
	Quality-sec	31	33	1073	1369	**2012**
Indoor↔outdoor	Avg QoE	-3.76	-0.50	-0.84	-0.11	$\mathbf{+0.03}$
	Uptime (s)	823	1112	1030	1524	**1500**
	Quality-sec	426	947	978	1110	**1424**
Outdoor pacing	Avg QoE	-2.90	-1.28	-1.28	$+0.00$	$\mathbf{+0.00}$
	Uptime (s)	806	960	960	1516	**1476**
	Quality-sec	492	960	960	1210	**1475**
Indoor pacing	Avg QoE	-6.04	-1.58	-0.16	-0.17	$\mathbf{-0.03}$
	Uptime (s)	1005	2102	1200	1538	**1594**
	Quality-sec	1	638	1007	1047	**1334**

IV. SCENARIO 2: WEARABLE-COLLABORATIVE BLIND-SPOT SENSING AND WARNING

A. Narrative and Motivation

Scenario 2 keeps the same creator on the same electric scooter, but stresses the *sensing* path rather than the uplink. The rider wears AR glasses and a smartwatch, carries a smartphone that acts as the hub agent, and the scooter carries a CPE node. While the primary service is still first-person streaming, a cyclist approaches from the rear blind spot. The smartphone transmits a probing signal toward the rear. The smartwatch is the default bistatic receiver, and the scooter CPE can be recruited as an auxiliary bistatic receiver. Because the transmitter and receiver are physically separated, this bistatic geometry suppresses self-interference relative to a monostatic small-form-factor radar and forms a wider virtual aperture.

The episode contains two successive disturbances. First, the rider's arm and body block the watch-to-target echo path during part of the approach, so the watch becomes nearly blind. Second, the scooter CPE becomes unavailable while the target is still in the rear warning zone. Possible causes include a link drop, power glitch, or movement out of range. No fixed cooperation choice covers both events. A watch-only policy fails during blockage, while static CPE delegation fails during the CPE outage. The watch's wideband sensing receiver draws non-trivial power from a battery of about one watt-hour. Thus, every adaptive controller idles it after the scooter-powered CPE covers the rear. The policies differ in how they recover when that device disappears. A capability-matched reactive controller has no micro-agent failure token. It detects the loss only after a report timeout and remains blind during that delay. The agentic policy identifies the failure immediately from the token. Its safety intent keeps the watch active as redundancy inside the warning zone, so it switches to phone–watch bistatic sensing without a coverage gap. The joint QoE combines sensing reliability, warning delivery, uplink streaming, and device energy. It rewards timely delivered warnings and penalizes misses, energy use, and mode switching.

B. Geometry and Mobility

The rider body origin is at the scooter. Device positions (in meters) are as follows: smartphone transmitter $(-0.1, 0.58, 1.2)$; scooter CPE $(0.55, -0.62, 0.75)$; and watch receiver $(0.0, 0.58, 1.15)$ in the visible pose or $(0.48, 0.0, 1.15)$ in the body-blocked pose. The target follows a two-phase trajectory. During the approach $0 \le t \le T_{\mathrm{a}}$ with $T_{\mathrm{a}} = 10$ s,

$$\mathbf{p}_{\mathrm{tar}}(t) = \mathbf{p}_0 + \tfrac{t}{T_{\mathrm{a}}}(\mathbf{p}_1 - \mathbf{p}_0), \tag{29}$$

with $\mathbf{p}_0 = (-24, 0, 1.1)$ and $\mathbf{p}_1 = (-8, 0, 1.1)$; afterwards it passes the rider at $v_{\mathrm{p}} = 6$ m/s along $+x$. The episode is 15 s at a 0.1 s step. The rear warning-zone event is

$$\mathcal{Z}_{\mathrm{rear}}(t) = \{x_{\mathrm{tar}}(t) < 0\} \cap \{\|\mathbf{p}_{\mathrm{tar},xy}(t)\| \le R_{\mathrm{warn}}\}, \tag{30}$$

with $R_{\mathrm{warn}} = 18$ m. The target enters the zone at about 3.8 s and leaves it as it passes the rider near 11.3 s.

C. Bistatic Sensing Link Budget

For receiver r the target echo power is

$$P_{\mathrm{echo},r}(t) = \frac{P_{\mathrm{tx}} G_r \lambda^2 \sigma_{\mathrm{rcs}}}{(4\pi)^3 d_{\mathrm{tx}}^2(t)\, d_{\mathrm{rx},r}^2(t)}\, 10^{-A_r(t)/10}, \tag{31}$$

where d_{tx} is the phone-to-target distance, $d_{\mathrm{rx},r}$ the target-to-receiver distance, G_r the receiver antenna gain, and $A_r(t)$ the blockage loss on the receiver's echo path.

Why the CPE outperforms the watch. The two bistatic receivers differ in two ways. Both favor the CPE. First, the wrist-worn watch has antenna gain $G_{\mathrm{w}} = 2.0$, whereas the scooter-mounted CPE has $G_{\mathrm{c}} = 4.0$. Equation (31) therefore gives the CPE twice the echo power, or a 3 dB advantage, at the same geometry. Second, body-blockage loss $A_r(t)$ affects only the watch echo path. The CPE remains away from the rider's torso and has a clear line of sight. The two effects compound. Even without blockage, the CPE detection-probability curve is above the watch curve. During blockage, the watch loses up to 20 dB and the CPE does not (Fig. 7). This gain and line-of-sight advantage explains why CPE recruitment restores sensing when watch confidence drops.

Self-interference and direct-path leakage. The probing transmitter and the receivers are never perfectly isolated, and the leakage must be driven below the weak target echo. The monostatic phone receiver, whose transmitter and receiver are co-located, sees a strong self-interference (SI) term suppressed in two stages,

$$P_{\mathrm{SI}}^{\mathrm{pre}} = P_{\mathrm{tx}}\, 10^{-A_{\mathrm{SI}}/10}, \tag{32a}$$

$$P_{\mathrm{SI}}^{\mathrm{res}} = P_{\mathrm{tx}}\, 10^{-(A_{\mathrm{SI}}+D_{\mathrm{SI}})/10}, \tag{32b}$$

where A_{SI} is the *pre-ADC analog* suppression (isolation, analog cancellation) and D_{SI} the *post-ADC digital* cancellation. The bistatic watch and CPE receivers instead see a direct-path (DP) leakage from the phone transmitter, set by free-space loss over the transmitter-to-receiver distance $d_{\mathrm{DP},r}$ and the receiver gain,

$$P_{\mathrm{DP},r} = P_{\mathrm{tx}} G_r \left(\frac{\lambda}{4\pi d_{\mathrm{DP},r}}\right)^2, \tag{33a}$$

$$P_{\mathrm{DP},r}^{\mathrm{res}} = P_{\mathrm{DP},r}\, 10^{-(A_{\mathrm{DP}}+D_{\mathrm{DP}})/10}, \tag{33b}$$

with A_{DP} and D_{DP} the analog (pre-ADC) and digital (post-ADC) direct-path suppression. The aligned suppression budget is

Interference path	Pre-ADC analog	Post-ADC digital	Total
Monostatic ($A_{\mathrm{SI}}, D_{\mathrm{SI}}$)	65 dB	35 dB	100 dB
Bistatic ($A_{\mathrm{DP}}, D_{\mathrm{DP}}$)	40 dB	60 dB	100 dB

Both paths reach the same 100 dB total, but the *split* is what matters, because only the pre-ADC (analog) stage protects the converter. A b-bit ADC with back-off B_{ADC} has a quantization-noise floor

$$P_{\mathrm{q}} = P_{\mathrm{pre}}\, 10^{-(6.02b+1.76-B_{\mathrm{ADC}})/10}, \tag{34}$$

with $b = 10$ and $B_{\mathrm{ADC}} = 6$ dB, so a large residual still present at the ADC input raises the quantization floor by an amount that no post-ADC digital cancellation can later remove. The co-located monostatic SI is by far the strongest interferer, so it is given the most analog headroom (65 dB before the ADC) to keep the converter out of compression; the spatially separated bistatic direct path is already weaker, so 40 dB of analog suppression suffices before the ADC and more of the budget is left to digital cancellation (60 dB).

Collecting the target echo against this disturbance, the mean effective sensing SINR at receiver r is

$$\gamma_r = \frac{\overline{P}_{\mathrm{echo},r}}{P_{\mathrm{th}} + P_{\mathrm{cl}} + P_{\mathrm{res},r} + P_{\mathrm{q},r}}, \tag{35}$$

where $P_{\mathrm{th}} + P_{\mathrm{cl}}$ is the thermal-plus-clutter floor, $P_{\mathrm{res},r}$ is the residual interference (monostatic SI for the phone, residual direct path for the watch and CPE), and $P_{\mathrm{q},r}$ the ADC quantization-noise power of (34). This SINR, not the raw range, is what drives detection: the sensing model maps γ_r through a receiver-specific waveform-CFAR curve (Section IV-E). The remaining sensing parameters are listed in Table VI.

D. Body Blockage

Body blockage affects only the watch echo, not the CPE link, the monostatic self-interference, or the bistatic direct path. The blockage fraction $\beta(t) \in [0,1]$ is a smooth ramp-in, dwell, and ramp-out built from the smoothstep $s(u) = u^2(3-2u)$: zero before 2.7 s, ramping over $[2.7, 3.1)$, unity over $[3.1, 5.9)$, ramping out over $[5.9, 6.3)$, and zero thereafter. The full blockage loss is $A_{\mathrm{watch}}(t) = 20\,\beta(t)$ dB, applied to the watch echo power, so it enters the watch SINR of (35) directly,

$$\gamma_{\mathrm{w}}(t) = \frac{\overline{P}^{0}_{\mathrm{echo,w}}(t)\, 10^{-20\beta(t)/10}}{P_{\mathrm{th}} + P_{\mathrm{cl}} + P_{\mathrm{res,w}} + P_{\mathrm{q,w}}}, \tag{36}$$

and the watch detection probability is read from its single SINR curve, $P_{d,\mathrm{w}}(t) = \mathcal{L}_{\mathrm{w}}(\gamma_{\mathrm{w,dB}}(t))$. Folding the blockage into the SINR, rather than blending a separate blocked detection curve, avoids counting the loss twice.

E. Waveform-CFAR Detection-Probability Lookup

Detection probabilities are estimated offline with a waveform-level CA-CFAR Monte Carlo. They are stored as functions of the mean effective SINR in (35), because the link budget provides SINR rather than range. The sweep is $-35{:}2{:}20$ dB. At each SINR point and for each receiver, the simulation generates QPSK-OFDM symbols with an 18 m target delay. It scales the target amplitude to obtain the requested SINR. The waveform then passes through clutter and receiver-specific residual interference: monostatic self-interference or the bistatic direct path. The next stages are ADC clipping and quantization, digital cancellation, known-symbol channel estimation, a range IFFT, and a range-1D CA-CFAR. Each trial uses a shared Swerling-I exponential RCS factor. The empirical probability is $\widehat{P}_{d,r}(\gamma) = \frac{1}{N_1}\sum_i \mathbf{1}\{Z_{r,i} > \eta_{r,i}\}$ with $N_1 = 2000$ target-present trials. The $N_0 = 5000$ target-absent trials are split into 3500 calibration trials and 1500 validation trials. The calibration targets a false-alarm rate of 10^{-4}. The controller applies the cumulative maximum of each curve to remove finite-sample dips.

Crucially, the curves are *receiver-specific*: equal analytical SINR does not imply equal CFAR statistics, because the phone carries structured residual self-interference while the watch and CPE carry a geometry-dependent direct path, giving different pre-ADC power, clipping, quantization, and calibrated thresholds. As Fig. 7 shows, the CPE reaches $P_d = 0.8$ near -13 dB, the phone near $+1$ dB, and the watch only near $+5$ dB, so at equal SINR the CPE is the most reliable receiver; together with its higher gain and its immunity to body blockage, this is why recruiting it rescues the sensing task. During the episode each receiver token is $P_{d,r}(t) = \mathcal{L}_r\big(\gamma_{r,\mathrm{dB}}(t) + X_r(t)\big)$, where $\mathcal{L}_r$ is the receiver-r lookup and X_r an AR(1) dB fluctuation. SINR values outside the grid use the nearest endpoint.

F. Agentic Sensing Controller

Let $a_{\mathrm{CPE}}(t) \in \{0,1\}$ denote whether the CPE is engaged for sensing and $a_{\mathrm{w}}(t) \in \{0,1\}$ whether the watch receive

chain is on. Recruitment is confidence-triggered: define the approach-and-degrade event

$$\mathcal{S}(t) = \mathcal{A}(t) \cap \{P_{d,\mathrm{w}}(t) < 0.8\}, \tag{37}$$

where $\mathcal{A}(t)$ is the event that the target is behind and approaching. If $\mathcal{S}(t)$ holds for three consecutive steps (with the watch on and the CPE's micro-agent heartbeat present), the CPE is recruited after a 0.4 s setup delay. The engagement is then *episode-scoped* and released when the target is no longer behind the rider, since with the watch idled during delegation, the trigger confidence is no longer observable and holding the scooter-powered CPE costs the wearables nothing. Communication need is logged but does not by itself recruit or retain the CPE.

While the CPE is engaged and healthy, the watch duty follows the *intent*. Outside the warning zone the dominant intent is efficiency, so the watch chain idles ($a_{\mathrm{w}} = 0$) and sensing is fully delegated to the CPE. Inside the zone, the intent switches to safety-critical sensing: the watch is woken (a 0.2 s wake-up) and kept warm as redundancy, and its output is fused with the CPE's (Section IV-G). The same intent switch caps the video enhancement layer and duplicates the warning over both uplink paths (Section IV-H).

The second agentic ingredient is closed-loop *resilience through cause attribution*. The CPE is physically available only outside a failure window $[t_{\mathrm{f},0}, t_{\mathrm{f},1}) = [8.0, 10.5)$ s. When the micro-agent failure token arrives, the hub drops the CPE immediately, $a_{\mathrm{CPE}}(t) = 0$, and re-plans to phone–watch bistatic sensing: if the watch was warm (an in-zone failure, as in the episode) the fallback is seamless, and if it was idled the fallback costs one wake-up delay. In the canonical timeline the failure begins at 8.0 s with the target inside the zone, so the watch is already warm and healthy (the body blockage cleared by 6.3 s) and coverage continues without a gap. When the heartbeat resumes, the engagement re-forms for as long as the target remains behind.

G. Detection Fusion and Warning Decision

The agent's effective detection probability follows its duty state. With no CPE engaged it uses the watch directly, $P_{d,\mathrm{agent}}(t) = P_{d,\mathrm{w}}(t)$. While delegated with the watch idled (CPE healthy, target outside the zone), it is the reliability-limited CPE-only value, $P_{d,\mathrm{agent}}(t) = \rho_{\mathrm{c}} P_{d,\mathrm{c}}(t)$. When the safety intent holds the watch warm inside the zone, the watch and CPE probabilities are fused by a reliability-limited soft OR,

$$P_{d,\mathrm{agent}}(t) = \min\Big[P_{\max},\, 1 - \big(1 - \rho_{\mathrm{w}} P_{d,\mathrm{w}}\big)\big(1 - \rho_{\mathrm{c}} P_{d,\mathrm{c}}\big)\Big], \tag{38}$$

with $\rho_{\mathrm{w}} = 0.98$, $\rho_{\mathrm{c}} = 0.85$, and $P_{\max} = 0.97$. Warnings are produced from noisy token-level range estimates $\widehat{r}(t) = r(t) + \epsilon(t)$ with $\epsilon \sim \mathcal{N}(0, \sigma_r^2)$ and $\sigma_r(t) = 0.45 + 2.0(1 - P_{d,\mathrm{agent}}(t))$; a warning is issued when the target is in the rear zone, $\widehat{r}(t) \le R_{\mathrm{warn}}$, and $P_{d,\mathrm{agent}}(t) \ge 0.6$. Warnings inside the rear zone count as successes; warnings outside it count as false alarms.

H. Communication, Energy, Warning Delivery, and Joint QoE

The phone and CPE uplinks use a UMa-style path loss $L(d) = 13.54 + 39.08 \log_{10} d + 20 \log_{10} f_c$ and a discounted Shannon rate $R = \eta B \log_2(1 + \Gamma_{\mathrm{lin}})$ with $\eta = 0.55$. The communication mode is C2 (phone only), C3 (phone+CPE when the phone link cannot meet the video target with margin), or C4 (safety-priority multi-link with the video bitrate capped) while a rear target is in the zone, the CPE is active, and the policy performs intent switching; the reactive baseline never enters C4. Energy accounts for phone sensing and communication, watch receive (charged only while the watch chain is on, per the duty masks of each policy), CPE receive and uplink, and glass capture and encoding, integrated over the episode (Table VI). Because the CPE runs off the scooter, we report wearable energy (glass, phone, watch) separately from the total.

A generated warning must also *reach* the far end over the uplink. Each transmission attempt on a single moving uplink is lost with probability $p_\ell = 0.08$ (outage plus contention, illustrative); the safety-priority mode C4 duplicates the warning over the phone and CPE paths, whose independent losses square to $p_\ell^2 \approx 0.006$. The delivered warning probability per step is the sensing-level warning probability multiplied by the corresponding delivery factor, so duplication carries a measurable benefit alongside its bit cost, and all coverage figures below refer to *delivered* warnings.

The per-step joint QoE adds streaming utility and a warning reward and subtracts penalties for misses, false alarms, latency, sensing overhead, energy, and mode switching, with weights $(w_{\mathrm{stream}}, w_{\mathrm{warn}}, w_{\mathrm{miss}}, w_{\mathrm{false}}, w_{\mathrm{lat}}, w_{\mathrm{token}}, w_{\mathrm{energy}}, w_{\mathrm{switch}}) = (1, 2, 3, 0.5, 1, 0.3, 0.5, 0.2)$. The missed-warning weight exceeds the warning reward because the scenario is safety-sensitive. As the headline safety metric we use *warning coverage*, the fraction of the in-zone interval over which the warning probability stays reliable (≥ 0.5). Unlike an "ever-warned" success rate, which saturates at one for every policy once the target is close, coverage exposes the mid-hazard blind windows that the two baselines suffer.

I. Baseline Policies

Four policies share the geometry, link budget, and SINR-to-P_d lookup and differ only in cooperation logic.

No-Coop senses with the always-on watch and never recruits the CPE. It is blind during the body blockage and pays the watch receive power for the whole episode.

Static-CPE is pre-configured cooperation with no closed loop: it delegates rear sensing to the higher-gain CPE for the whole encounter (whenever the target is behind) and idles the watch chain, so its detection probability is the reliability-limited CPE-only value $\rho_{\mathrm{c}} P_{d,\mathrm{c}}(t)$. It is strictly better than watch-only from the start, but it keeps no fallback of any kind: it is blind for the entire CPE failure and resumes the watch only after the target has passed.

Reactive is the capability-matched same-space baseline. It observes the same watch-confidence stream and uses the same confidence threshold. It can recruit and release the CPE with the same 0.4 s setup delay. It can also idle and wake the

watch with the same 0.2 s wake-up delay. Its decision logic is non-agentic. After delegation, it idles the watch whenever the CPE appears healthy, whether the target is in the warning zone or not. It receives no failure token. It infers a CPE loss after 1.0 s of missing reports. During this timeout, it pays the coordination overhead and performs no sensing. It also has no communication-intent switch, so it remains in $\mathrm{C2}/\mathrm{C3}$ without a video cap or warning duplication. Its blind window is the timeout plus watch wake-up, or about 1.2 s inside the hazard zone.

Proposed is the agentic controller of Section IV-F. Against Reactive, exactly three ingredients differ, and each is one of the agentic functions of Section II-A: intent-dependent redundancy (the watch is kept warm and fused inside the warning zone), semantic failure attribution (the token replaces the timeout), and communication-intent switching ($\mathrm{C4}$ with warning duplication). It is the only policy covered through both events, and its wearable energy remains below No-Coop's always-on watch.

J. Parameter Settings

TABLE VI
SCENARIO 2 PARAMETERS. STANDARDS-BASED VALUES FOLLOW COMMON 5G NR FR1 ISAC ASSUMPTIONS; ILLUSTRATIVE VALUES ARE DESIGN CHOICES DOCUMENTED HERE FOR REPRODUCIBILITY.

Parameter	Value
Carrier f_c / bandwidth B	3.65 GHz / 100 MHz
Phone sensing Tx power	23 dBm
Receiver noise figure	7 dB
Target RCS σ_{rcs}	0.1 m^2
Watch / CPE Rx antenna gain $G_{\mathrm{w}}/G_{\mathrm{c}}$	2.0 / 4.0
Monostatic SI suppr. (pre+post-ADC)	65+35 dB
Bistatic direct-path suppr. (pre+post-ADC)	40+60 dB
ADC resolution / back-off	10 bit / 6 dB
Full body-blockage loss	20 dB
Blockage window (ramp/dwell)	$[2.7, 6.3)$ s
CPE failure window	$[8.0, 10.5)$ s
Warning range R_{warn}	18 m
Recruit trigger / setup	$P_{d,\mathrm{w}}$<0.8 (3-step confirm) / 0.4 s
Engagement release	target no longer behind (episode-scoped)
Watch Rx-chain power / wake-up	0.4 W / 0.2 s
CPE report timeout (Reactive)	1.0 s
Warning delivery loss (single / dup.)	0.08 / 0.08^2
Fusion $\rho_{\mathrm{w}}, \rho_{\mathrm{c}}, P_{\max}$	$0.98, 0.85, 0.97$
gNB distance / height	200 / 25 m
Phone / CPE uplink Tx	2.0 W / 23 dBm
Video ladder / target	$3, 8, 15, 25$ / 15 Mbps
Lookup SINR grid / ref. range	$-35{:}2{:}20$ dB / 18 m
CFAR N_1 / N_0 / P_{fa}	2000 / 5000 / 10^{-4}

K. Results and Discussion

Figure 6 reports the 1000-trial mean for each policy. The outer trials use seeds 31 through 1030. Figure 8 gives the seven-panel diagnostic. Three observations follow.

(i) Cooperation rescues the blockage. During the body blockage the mean watch SINR collapses and its detection probability falls to 0.003, so No-Coop loses the rear-hazard warning and its in-zone coverage is capped at 56.6%. Every cooperating policy bridges the blockage through the scooter CPE, whose reliability-limited detection probability holds near 0.835, and issues the first delivered warning about 0.10 s after the target enters the zone. This gain requires no closed loop at all. Even static delegation achieves it.

(ii) The failure separates the policies exactly along their control loops. When the CPE dies inside the warning zone, Static-CPE, with no loop, is blind for the entire 2.5 s outage and only resumes the watch after the target has passed: 65.8% coverage. The capability-matched Reactive controller does recover, but only after its 1.0 s report timeout plus the 0.2 s watch wake-up, a 1.2 s blind window inside the hazard zone: 82.9% coverage. The agent, warned by the failure token and holding the watch warm under its in-zone safety intent, falls back to phone–watch bistatic sensing seamlessly: 98.7% coverage. The residual gap between Reactive and Proposed is attributable to three agentic ingredients: semantic failure attribution, intent-dependent redundancy, and warning duplication in $\mathrm{C4}$. Duplication reduces delivered-warning loss from 8% to below 1% per step.

(iii) The safety gain is not bought with wearable energy or trigger-happy warnings. The mean per-step joint QoE is $+0.86$ for the agent against $+0.54$ for Reactive, $+0.08$ for Static-CPE, and $+0.10$ for No-Coop. Reactive streams the highest raw bitrate (18.6 Mbps, since it never caps video in $\mathrm{C4}$), yet its missed-warning penalty during the blind window dominates. Static-CPE additionally pays a standing energy penalty for holding the CPE transmit radio the whole encounter. On the wearable side the agent consumes 81.0 J against 84.3 J for No-Coop's always-on watch: intent-scoped watch duty (4.7 J of watch energy versus 6.0 J always-on) more than pays for the redundancy it buys, while Reactive (2.1 J) and Static (1.5 J) idle the watch harder but pay in coverage. The false-alarm rate stays at 0.022, so the coverage does not come from lowering the warning bar. The headline claim mirrors Scenario 1 in the sensing regime: cooperation restores the capability lost to blockage, any closed loop eventually recovers from a device failure, but only semantic failure attribution plus intent-dependent redundancy holds the safety function through the failure without a gap, at wearable energy below an always-on receiver.

REPRODUCIBILITY

Source code and browser demos are available at https://chaokaiwen.github.io/agentic-ue-comimo/. The scripts generate all traces internally. No external dataset is required.

For Scenario 1, download `scene_aware_ue_comimo_uplink.m` from the site (`matlab/`), start MATLAB in its containing folder, and run `scene_aware_ue_comimo_uplink`. The default seed is 7. The script runs all five policies and writes figures to a `results_matlab/` folder created next to the script. Compare the printed metrics with Table IV and the generated figures with Figs. 1–5. The MATLAB controls are `cfg.ebudget.enable`, `cfg.cx.enable`, `cfg.fps.enable`, `cfg.pred.cx_noise`, and `cfg.deadzone.enable`.

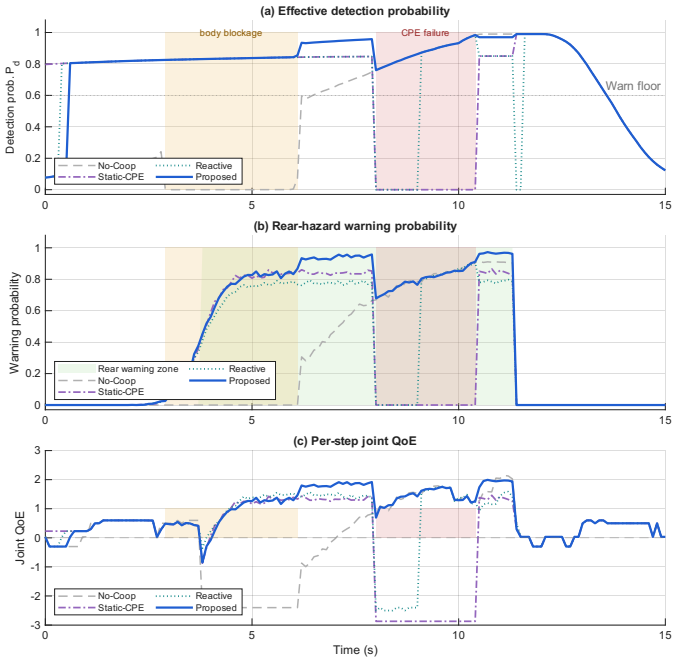


Fig. 6. Scenario 2 policy comparison over the episode (1000-trial mean). The rear target triggers a body-blockage interval (left band) and a CPE-failure interval (right band). (a) effective detection probability, (b) rear-hazard warning probability inside the warning zone, and (c) per-step joint QoE, for No-Coop, Static-CPE, the capability-matched Reactive controller, and the proposed agentic policy. No-Coop loses the warning during the blockage; Static-CPE is blind for the entire CPE failure; Reactive recovers from the failure only after its report timeout and watch wake-up; the agent, token-informed and with the watch kept warm in-zone, maintains the warning through both events.

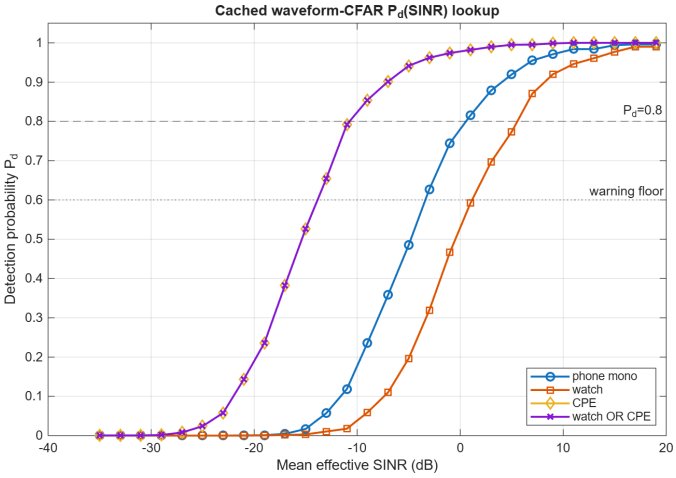


Fig. 7. Offline waveform-CFAR lookup used by the Scenario 2 episode: detection probability $P_d(\gamma)$ versus mean effective SINR γ, estimated separately for the monostatic phone, the watch, and the CPE (the watch-or-CPE fusion overlaps the CPE curve). At equal SINR the CPE reaches $P_d = 0.8$ near -13 dB, the phone near $+1$ dB, and the watch only near $+5$ dB, because equal analytical SINR does not imply equal CFAR statistics.

The Scenario 1 browser implementation is the Case 1 interactive demo on the site (`case1/`). Open its browser console and run `batchMulti(2400)`. This command runs the four trajectory families with seeds 1–12 and reports arithmetic means. The browser controls are

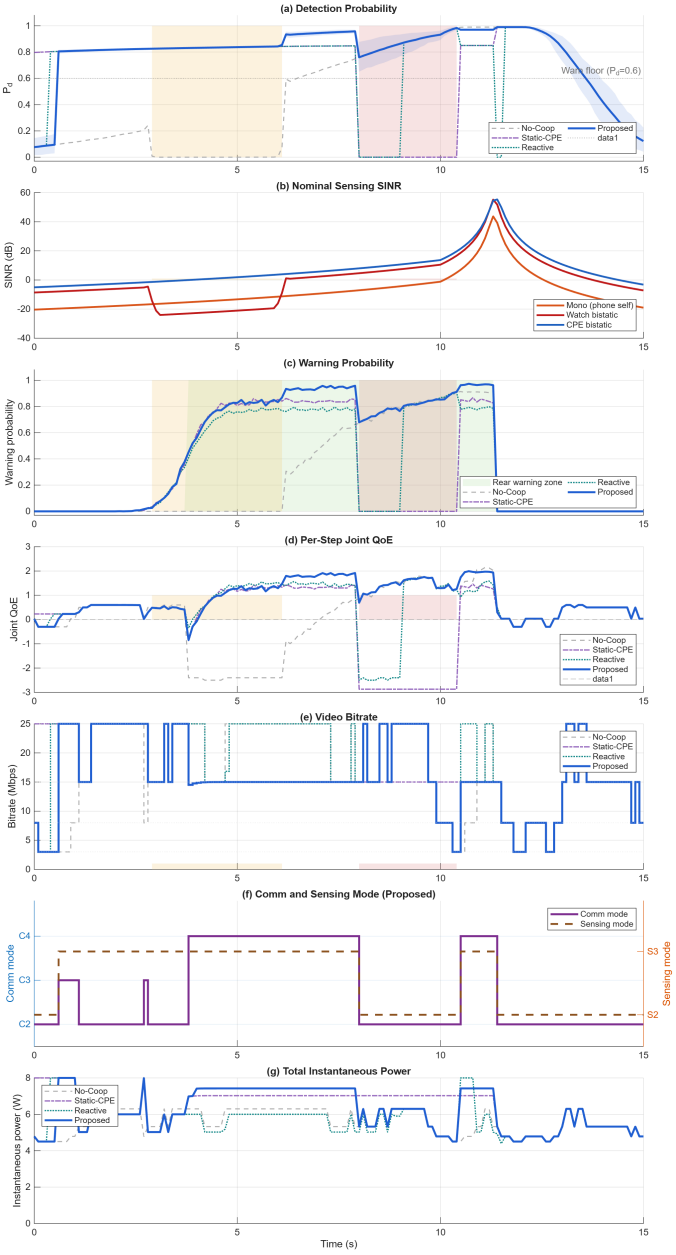


Fig. 8. Scenario 2 seven-panel diagnostic: detection probability, nominal sensing SINR, warning probability, per-step joint QoE, video bitrate, communication/sensing mode (proposed), and total instantaneous power, for the four policies. Shaded bands mark the body blockage (left) and the CPE failure (right).

`CFG.m3.enable`, `CFG.cx.enable`, `CFG.m6.enable`, `CFG.pred.cx_noise`, and `CFG.deadzone.enable`. The MATLAB and browser implementations use different random-number generators. Do not compare individual shadowing traces across them. Both reproduce the policy ordering. On the canonical trajectory, compare aggregate QoE within a few hundredths and battery within a few percent.

For Scenario 2, download `run_scenario2_sinr_pd_lookup.m` from the same page and start MATLAB in its containing folder. The entry function is `run_scenario2_sinr_pd_lookup`, and all parameters are in `defaultConfig()`. The

default waveform-CFAR settings are 2048 subcarriers, one OFDM symbol, 60 kHz spacing, two guard cells, 16 training cells, five search cells, and a ± 2-bin half-width (five-bin total) detection window. The SINR-to-P_d lookup is cached in `cfar_cache/` by a setup hash. The CPE-failure controls are `cfg.cpeFailStartS` and `cfg.cpeFailEndS`; set the start to `Inf` to disable the event. The other controls are `cfg.pWatchRxW`, `cfg.watchWakeS`, `cfg.cpeReportTimeoutS`, and `cfg.warnDeliverLossSingle`. Use the following command to run all four policies, including Reactive:

```
run_scenario2_sinr_pd_lookup('seed',31,'numTrials',1000,
'cfarNumMonteCarlo',2000,'cfarNumH0Trials',5000,
'outputTag','scenario2_matlab_range1d_cfar');
```

The outer trials use seeds 31 through 1030. The command writes to `artifacts/scenario2_matlab_range1d_cfar/` next to the script. Its outputs include `blindspot_paper_combined.pdf`, the SINR-to-P_d curves, the seven-panel overview, and JSON and MAT summaries. The four-policy metrics are printed in the MATLAB command window. The interactive browser version is the Case 2 demo on the site (`case2/`). It uses the same SINR link budget, lookup curves, and four policies. It also includes the intent-scoped watch duty, report-timeout inference, and warning-delivery model. Its expectation-based warning model reproduces the Reactive and Proposed coverage values in Section IV-K to three decimal places.