

Adaptive Traffic Signal Control Using IoT and CCTV Integration in a Rapidly Urbanizing Sub-Saharan African City

Abdulkareem Olakunle Afolabi [1*], Mercy Osato Omoifo-Irefo [2]
[1] Department of Transportation and Highway Engineering, South Carolina State University, South Carolina, USA
[2] Department of Environmental and Public Health Sciences, University of Cincinnati. Ohio, USA

* Corresponding Author: **Abdulkareem Olakunle Afolabi**

Article Info

ISSN (Online): 3107-6580
Volume: 02
Issue: 06
Received: 06-09-2026
Accepted: 04-10-2026
Published: 02-11-2026
Page No: 24-37

Abstract
Nairobi requires traffic-management approaches that can respond to changing demand while accounting for heterogeneous traffic and imperfect observations. Official traffic surveys provide contextual evidence, but aggregate counts cannot establish whether integrating Internet of Things (IoT) sensing and CCTV improves intersection operations. This study develops a transparent fusion and queue-responsive control framework and evaluates its behavior through secondary-data analysis and a stylized simulation. Twenty station-year records covering ten Nairobi count stations in 2013 and 2023 were transcribed and arithmetically audited. A separate four-approach queue experiment compared equal fixed-time, demand-matched fixed-time, perfect-observation queue-responsive, CCTV-only, IoT-only, and fused controllers under six assumed demand and sensing conditions, with 30 paired random seeds per condition. Published station totals increased from 515,282 to 697,801 passenger-car units (PCU), or 35.42%, while station-specific changes varied substantially. In the synthetic unbalanced-demand scenario, fused control produced a mean workload-delay proxy of 12.91 s/PCU, compared with 22.03 for equal fixed-time and 16.41 for demand-matched fixed-time, reductions of 41.41% and 21.35%, respectively. Against CCTV-only control, the paired improvement was 0.38 s/PCU, with a 95% confidence interval of 0.25 to 0.52. The balanced-demand comparison with CCTV alone did not clearly establish an advantage. The findings support further testing of sensing integration, while demonstrating that adaptive-control gains and fusion gains must be evaluated separately. The experiment establishes conditional computational behavior, not field effectiveness, crash reduction, or deployment readiness in Nairobi.

DOI: https://doi.org/10.54660/.IJECA.2026.2.6.24-37

Keywords: Adaptive traffic signal control, Internet of Things, CCTV analytics, sensor fusion, queue-responsive control, Nairobi, simulation

1. Introduction

1.1. Urbanization and the Local Traffic-Management Problem

Nairobi is an appropriate setting for investigating demand-responsive traffic management because its development planning has explicitly addressed the relationship between urban growth, transport demand, and metropolitan connectivity. Earlier official surveys documented substantial changes in traffic activity, while later surveys provide a basis for examining changes across selected road-count locations [2]. These observations motivate a locally grounded investigation without assuming that every intersection has the same congestion mechanism.

Traffic operations depend on more than total demand. Approach imbalance, turning movements, pedestrian crossing requirements, public-transport stopping, and the spatial distribution of queues can change the timing decisions that an intersection requires. Signal timing practice therefore calls for a relationship between operational objectives, detection, phasing, and evaluation [3]. In this study, city-level evidence establishes context, whereas the executed controller experiment isolates a narrower computational question under declared assumptions.

1.2. Limitations of Existing Signal-Control Approaches

Fixed-time operation can allocate green time efficiently when its assumed demand pattern remains appropriate, but cannot directly respond to short-term queue fluctuations. Established adaptive systems already adjust signal settings in response to measured conditions. SCOOT, for example, adapts splits, cycle lengths, and offsets through repeated timing adjustments [4]. Consequently, the novelty of an IoT-CCTV proposal cannot rest on the general idea of changing green time dynamically.

Queue-based control offers an interpretable alternative to complex optimization. Max-pressure research connects stage selection with queues and service rates, under specific network and stochastic assumptions [5]. Reinforcement-learning approaches have also been investigated for adaptive signal operation [6]. More recent pressure-based learning has addressed arterial coordination [7]. These methods establish substantial prior art, but do not by themselves validate a particular sensing arrangement or its transfer to Nairobi.

The traffic state available to a controller may be incomplete. Roadside sensors differ in measurement coverage and installation needs [8]. Video-based systems introduce detection and tracking errors that must be distinguished from control errors [9]. Multi-object tracking can recover observations that would otherwise be discarded because of low detection confidence, but benchmark performance is not evidence of reliable operation at a new camera site [10]. Sensor integration must therefore be tested as an estimation problem as well as a control problem.

1.3. Research Gap and Problem Statement

A defensible research problem is whether combining independently informative observations produces operational benefits beyond those achievable with one sensing source and the same controller. Data-fusion literature identifies uncertainty, registration, inconsistent observations, and correlated errors as central difficulties [11]. Literature on intelligent transport in developing-country settings also emphasizes the relationship between technical design and infrastructure conditions [12]. The present study addresses this problem through a transparent computational benchmark, while identifying what remains untested locally.

The methodological gap is the need to isolate the incremental contribution of fusion from the larger contribution of adapting control. The validation gap is the absence, in the evidence assembled here, of synchronized local queue-reference, IoT, and CCTV measurements. The implementation gap concerns the hardware, maintenance, governance, and controller interfaces required for a field installation. These are the boundaries of this study, rather than a claim that no previous African adaptive-control research exists [11, 12].

1.4. Aim and Specific Objectives

The aim is to formulate an integrated IoT-CCTV adaptive-control framework and examine its conditional behavior in a reproducible experiment informed by Nairobi's transport context. The completed objectives are to audit published station-level traffic data, implement a declared fusion rule and queue controller, compare sensing ablations under identical assumed demand, and quantify uncertainty across paired replications. A further objective is to specify the field measurements needed to validate the framework. That final objective is a protocol contribution, not a completed field campaign.

1.5. Research Questions and Hypotheses

RQ1 asks whether fused queue-responsive control reduces the synthetic workload-delay proxy relative to equal fixed-time and demand-matched fixed-time control. RQ2 asks whether integration adds a measurable benefit relative to CCTV-only and IoT-only control when the decision policy is held constant. RQ3 asks how directional imbalance and a mid-run demand reversal affect those comparisons. RQ4 asks how a camera outage, camera bias, and shared sensing bias change performance.

For RQ1, the directional hypothesis is that the mean paired difference, fused minus demand-matched fixed-time delay proxy, is negative in the unbalanced-demand scenario. For RQ2, the analogous difference against CCTV-only control is examined separately in each scenario. These are exploratory computational hypotheses. They were not independently preregistered, and the confidence intervals are not confirmatory evidence of local field benefit. No hypothesis concerning crashes, emissions, or pedestrian outcomes is tested.

1.6. Original Contributions

The first contribution is a reproducible transcription and audit of an official, publicly accessible aggregate dataset. The second is an executable benchmark that holds the controller constant while varying the observation source. The third is a paired evaluation that includes unfavorable or uncertain comparisons, particularly when fusion offers little advantage over video. The fourth is a field-validation specification connecting perception, control, fallback, and governance. Neither inverse-variance fusion nor queue-responsive phase selection is claimed as a new algorithm [5, 11].

1.7. Paper Organization

Section 2 positions the framework within sensing and signal-control research. Section 3 defines the Nairobi evidence and the separate computational design. Section 4 presents the implementation and future field-validation procedure. Section 5 reports observed-data calculations and synthetic experiment results. Sections 6 and 7 discuss interpretation, transferability, and limitations. Section 8 states the conclusions supported by the completed analyses.

2. Related Work and Research Gap

2.1. Conventional and Adaptive Signal-Control Methods

Signal-control studies must distinguish timing adaptation from sensing improvements. A controller supplied with perfect queues answers a different question from one supplied with noisy measurements. Pressure-based methods and learning-based control provide useful methodological

reference points, but their assumptions, phase restrictions, and training requirements differ [5, 6, 7]. The present benchmark therefore includes a perfect-observation version of its own queue policy, rather than treating a different published algorithm as an interchangeable baseline.

SCOOT illustrates that coordinated adaptation is established engineering practice [4]. The benchmark here is deliberately simpler: an isolated two-phase intersection abstraction without offsets or downstream coordination. Its demand-matched fixed controller receives the initial directional demand ratio, avoiding a comparison exclusively against an obviously mismatched equal split. This is a demand-proportional heuristic, not a globally optimized timing plan, a reconstruction of Nairobi timings, or a substitute for a commercial adaptive-system evaluation.

2.2. IoT-based Traffic Monitoring

IoT describes connectivity and data exchange, rather than a unique traffic measurement technology. A connected loop, magnetometer, radar, or camera may measure different variables over different detection zones. Sensor selection must consider whether the intended output is presence, occupancy, count, speed, or queue state [8]. In the proposed field architecture, roadside sensing supplies flow and occupancy information that would need a separately validated transformation before becoming a queue estimate.

The synthetic IoT channel directly measures queue workload with assumed noise. This abstraction isolates fusion behavior but does not demonstrate that a single physical detector can observe an entire queue. A field implementation would require multiple detection zones, an arrival-departure accumulation model, radar coverage, or another justified queue estimator. Equipment models, mounting details, sampling rates, and installation costs are therefore left unclaimed until actual procurement and surveying occur [8].

2.3. CCTV-based Traffic Estimation

A practical video pipeline requires object detection, temporal association, geometric calibration, and a rule that converts trajectories into traffic quantities. The UA-DETRAC benchmark shows why detection and tracking should be evaluated together [9]. ByteTrack provides a relevant association method [10]. Neither publication establishes vehicle-count or queue-estimation accuracy for Nairobi, so no published benchmark score is imported as a local performance result.

Traffic datasets can differ substantially in viewpoint and coverage. The pNEUMA experiment used aerial observations in Athens and supplies a valuable research resource, but its geography and observation geometry differ from roadside CCTV [13]. A recent BMD-45 preprint provides developing-city CCTV imagery from Bengaluru and examines domain differences; it is supplementary evidence, explicitly not treated as a peer-reviewed Nairobi validation dataset [14]. These examples support testing models in the intended operating domain.

2.4. Integration of Heterogeneous Traffic Data

Fusion requires consistent units and aligned timestamps before observations can be combined. Independent unbiased measurements can justify inverse-variance weighting. When observations share calibration errors or common environmental bias, the same rule can produce misleading confidence [11]. The completed experiment tests independent random channels and an imposed shared bias, but does not estimate a full joint error covariance from real sensors.

The key design distinction is between statistical redundancy and operational redundancy. Two channels may reduce random error when both are available, while one surviving channel may preserve adaptive operation during an outage. Both effects are examined here. However, the assumed variance functions are known by the simulator. Their availability is an idealization requiring replacement with measured uncertainty models in a field study [11].

2.5. Evidence from Sub-Saharan African and Comparable Settings

Regional transfer should be based on observed operating conditions rather than a single label such as African traffic. Infrastructure compatibility, maintenance capacity, sensing coverage, and institutional arrangements can determine whether an apparently effective algorithm is useful [12]. The Nairobi contextual analysis therefore does not assign the same demand pattern or deployment constraints to all Sub-Saharan African cities.

The public Nairobi report is a transport-demand survey source, not an integrated-sensing experiment. It offers documented PCU-based observations and collection arrangements [1]. The international perception datasets and control papers offer different kinds of evidence, summarized in Table 1. This separation prevents a simulation result, a detector benchmark, and a metropolitan survey from being presented as equivalent validation.

Table 1: Comparative literature matrix. Limitations refer to applicability to this study, rather than a universal assessment of each source.

Study	Setting	Sensing	Method	Evaluation	Transfer limit	Relevance
[4]	UK	Road detectors	Coordinated adaptive	System development	Different setting	Established adaptation
[5]	Abstract network	Queue state	Max pressure	Theory	Idealized queues	Queue policy rationale
[7]	Arterial networks	Traffic state	Pressure learning	Simulation	Different sensing	Stronger future baseline
[9]	Traffic video benchmark	Camera	Detection and tracking	Annotated data	Not Nairobi	Perception validation
[11]	Multiple domains	Heterogeneous	Fusion methods	Review	No local control test	Error assumptions
[12]	Developing countries	ITS technologies	Methodological review	Literature	No present field trial	Implementation context
[13]	Athens	Drone trajectories	Observational	Field dataset	Different viewpoint	Data-domain boundary
[14]	Bengaluru	CCTV	Detection benchmark	Dataset preprint	Not peer reviewed	Domain testing

2.6. Synthesis and Positioning of the Present Study

The literature establishes adaptive control, video tracking, and heterogeneous fusion as existing fields [4, 5, 10, 11]. The proposed contribution is an auditable evaluation arrangement: real contextual data are distinguished from synthetic controller inputs, sensing ablations share a decision policy, and uncertainty is calculated from independent random seeds. Any future claim of algorithmic novelty would

require additional development and a broader comparison than this proof of concept provides.

3. Study Area and Research Design

3.1. Study Location and Urban Context

The selected city is Nairobi, Kenya. The contextual dataset consists of the ten stations reported in Table 2.7 of the June 2025 JICA-NaMATA technical report. Their identifiers are retained as DT01 to DT08, DT10, and DT11; DT09 is not added because it is absent from that comparison table [1]. These are road-count locations, not ten instrumented signalized intersections.

Table 2: Actual published station totals and original growth calculations. Units are daily PCU. Station identifiers are preserved; these are count stations, not instrumented experimental junctions.

Station	2013 PCU	2023 PCU	Change PCU	Change %
DT01	28,237	40,492	12,255	43.40
DT02	71,101	70,912	-189	-0.27
DT03	79,656	126,412	46,756	58.70
DT04	107,993	177,263	69,270	64.14
DT05	102,288	84,980	-17,308	-16.92
DT06	30,168	27,399	-2,769	-9.18
DT07	18,375	34,647	16,272	88.56
DT08	27,104	59,990	32,886	121.33
DT10	15,568	42,909	27,341	175.62
DT11	34,792	32,797	-1,995	-5.73

A geographic intersection map is not supplied because the executed controller experiment has no surveyed intersection coordinates or verified local geometry. Figure 1 instead depicts the proposed architecture. A field manuscript should add a map from verified coordinates with an appropriate base-map attribution, and should distinguish count stations, camera locations, detector zones, and actual signalized junctions. Inventing a map of experimental sites would misstate this study's design.

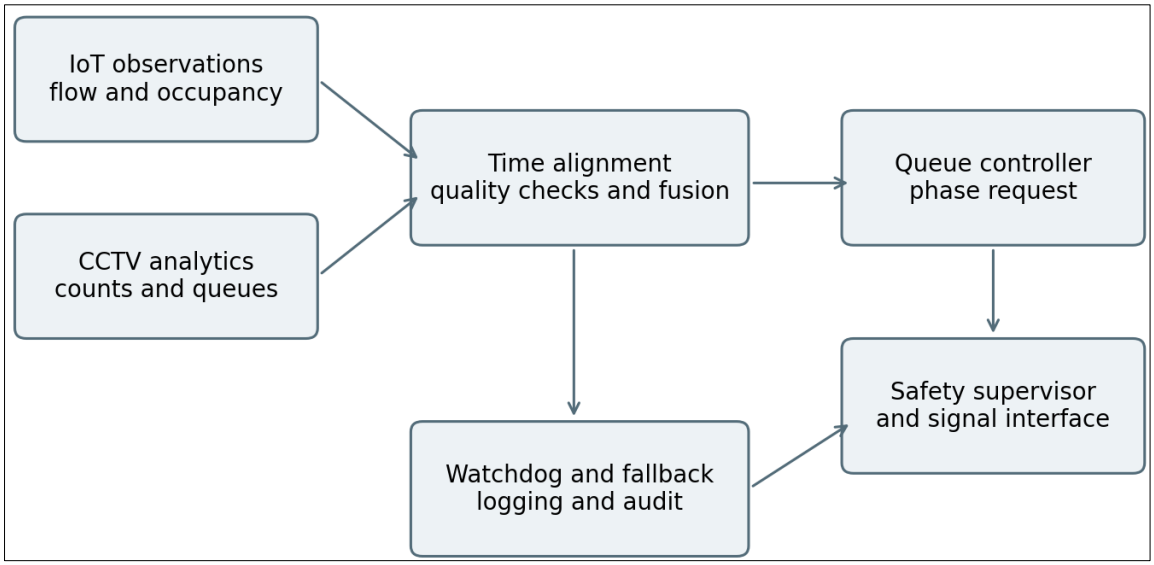


Fig 1: Proposed IoT-CCTV architecture. Original conceptual diagram. No installed hardware or field-tested safety interface is implied.

3.2. Intersection Selection and Characteristics

The computational site is a synthetic intersection with four queues, numbered 0 to 3. Phase 0 serves queues 0 and 2; phase 1 serves queues 1 and 3. Turning movements, lane widths, shared lanes, bus stops, and pedestrian crossings are not modeled. There is no assertion that this geometry corresponds to a named Nairobi junction. This abstraction allows the sensing comparison to be reproduced without manufacturing local geometry or timing observations.

Future field selection should identify junctions with accessible controller records, safe mounting opportunities, meaningful demand variation, and queue-reference coverage. Geometry, legal movements, pedestrian provisions, storage lengths, and existing timings should be recorded before building a calibrated model. Selection should avoid basing recruitment only on junctions where cameras work well, because this would hide relevant sensing failures [3, 8].

3.3. Overall Study Design

The study has two independent evidence streams. Stream A is a secondary descriptive analysis of published aggregate data. Stream B is an executed discrete-time stochastic queue experiment with assumed arrivals and sensing errors. Stream A motivates the need to consider heterogeneity, but its daily PCU totals do not calibrate Stream B's second-by-second arrivals. The study is neither a controlled before-and-after evaluation nor a hardware-in-the-loop pilot.

No current Nairobi signal timing plan, detector stream, CCTV video, or pedestrian count was supplied or collected. Statements about the proposed system describe its intended

field architecture. Numerical control results describe the executed software experiment. The distinction is preserved in the results tables, figure captions, and downloadable replication files.

3.4. Observation Schedule and Sampling Strategy

The official report distinguishes 24-hour Type A counts collected in 2023 from separate 16-hour Type B counts collected in 2025. The present transcription uses the published 2013-2023 Type A comparison only [1]. Sampling dates are not converted into a continuous annual series, and the two years are not treated as repeated daily observations. Weather coverage and within-day uncertainty cannot be reconstructed from the comparison table.

The synthetic experiment uses 30 independent arrival and noise seeds, numbered 1000 through 1029, for each of six scenarios and six controllers, totaling 1,080 runs. Each run has 3,600 seconds of arrivals and a subsequent clearance period. Thirty seeds provide a practical exploratory uncertainty estimate, not a prospective power justification. Common random numbers pair controllers within each scenario; seeds, rather than individual seconds, are the replication units.

3.5. Operational Definitions

Published traffic volume is expressed in PCU over the survey period. It is not a count of people, registered vehicles, or unique journeys. In the simulation, one arrival packet represents one assumed unit of PCU workload. Queue length is unfinished PCU workload, rather than measured spatial queue length. Throughput is workload discharged by second 3,600. The delay proxy is accumulated post-service queue workload divided by total arriving workload. It omits acceleration and deceleration losses and must not be called Highway Capacity Manual control delay.

A camera outage means the simulated camera observation is unavailable between seconds 1,200 and 1,799. Camera bias means a multiplicative reduction of 35% in its estimates on queues 0 and 2 during the arrival period. Common bias applies the same reduction to both channels. These are researcher-selected stressors, not estimated Nairobi failure frequencies. Availability, communications latency, and hardware failure rates are not measured.

4. Materials and Methods

4.1. System Architecture

The proposed deployment architecture contains roadside observations, video analytics, an edge processing unit, time alignment, quality checks, fusion, a phase-request policy, and an independent safety supervisor. The edge unit would retain a valid fallback plan and log observation freshness and accepted phase requests. Figure 1 is an original conceptual design. Only the queue, sensing-abstraction, fusion, and phase-selection components are exercised in the completed experiment.

A field implementation should transmit aggregate observations when feasible and should maintain local control during loss of a central connection. The controller interface must accept requests only through approved operational functions, with conflict checking and intergreen enforcement independent of the optimization process. The general relationship among detection, phasing, and timing constraints is grounded in signal-engineering guidance; the specific hardware interface remains to be selected and verified [3].

4.2. IoT Sensing Hardware and Measurements

No physical sensor model is reported as installed. The synthetic IoT observation for queue j at time t is yI = max(0, q + sigmaI zI), where sigmaI = 4 + 0.05q, zI follows a standard normal distribution, and q is queue workload in PCU. Noise is independently generated across observation times and approaches. The formula is an experimental assumption, not a manufacturer's accuracy specification.

For a field deployment, the equipment register should include model and firmware identifiers, mounting position, coverage polygon, output variables, sampling frequency, timestamp source, power design, and communication protocol. A radar or multi-zone detector may be appropriate for queue coverage, but sensor choice should follow the intended measurement and site conditions [8]. No bill of materials, installation price, or procurement recommendation is presented as measured evidence.

4.3. CCTV Acquisition and Processing

The completed experiment does not run an object detector. Its synthetic camera observation is yC = max(0, q + sigmaC zC), where sigmaC = 2 + 0.15q. The coefficients create an assumed channel with lower error at small queues and increasing uncertainty as queues grow. In biased scenarios, observations on queues 0 and 2 are subsequently multiplied by 0.65. This is a stress test of systematic underestimation, not a photometric simulation of rain or darkness.

A future video pipeline could combine a documented detector configuration with a documented tracker. YOLO provides a foundational example of single-stage detection [15], while the selected tracker should be evaluated for trajectory continuity as well as counting utility [10]. Camera resolution, focal length, frame rate, compression, mounting geometry, and regions of interest must be recorded from the actual installation. These parameters cannot be supplied credibly from the aggregate survey dataset.

Local annotations should include relevant vehicle groups, occlusion, crossing-line events, and queue-tail positions. A queue may be operationalized as contiguous low-speed traffic within a stated approach polygon, with sensitivity to the speed and gap thresholds. Training, tuning, and testing should be separated by recording session and preferably by location, rather than by adjacent frames. The choice of perception metrics should reflect downstream counting and queue-estimation needs [9].

4.4. Ground Truth and Sensing Validation

The simulator's internal queue is exact reference workload for the computational experiment. It is not an independent field annotation. The released output includes synthetic estimation error as a diagnostic, but it cannot establish CCTV or IoT accuracy. Since controllers generate different queue trajectories, their observation-error averages are also conditional on different states and should not be interpreted as a matched perception benchmark.

A field study should obtain independent manual or surveyed reference observations over stratified periods, including difficult conditions. Count mean absolute error, queue error in vehicles or metres, class-specific detection precision and recall, and tracking continuity should be evaluated on held-out data. Double annotation on a documented subset should quantify disagreements and adjudication. Detection scores alone do not validate an operational queue estimator [9, 10].

4.5. Data Preprocessing and Synchronization

The official comparison table was manually transcribed into a tidy file with station, year, seven traffic classes, published total, recomputed component sum, and arithmetic difference. Source values were retained even when they failed exact summation. Station growth uses published station totals, whereas class shares use recomputed class sums. The provenance file identifies the source URL, PDF page 17, printed page 11, and Table 2.7 [1].

Synthetic observations use a common one-second clock and are requested every five seconds when no phase clearance is active. There is no modeled transport latency or clock drift. For a field system, observations should be aligned by capture time, not merely server arrival time. Stale measurements should carry explicit age and quality fields. Duplicate, missing, and contradictory observations should be logged rather than silently replaced by zero [11].

4.6. IoT and CCTV Fusion Method

For two available observations, the executed estimate is qhat = (yC/sigmaC² + yI/sigmaI²)/(1/sigmaC² + 1/sigmaI²). Nonnegative channel observations make the weighted estimate nonnegative. This is an inverse-variance rule motivated by the independent, unbiased measurement case [11]. In the experiment, the simulator supplies its true state-dependent noise scales. That favorable assumption isolates control behavior but gives the fusion module information that a real installation must estimate.

When the simulated camera is unavailable, the fused controller uses the IoT observation alone. The CCTV-only controller alternates phases using a 40-second fallback target, subject to its update schedule and maximum-green rule. Shared bias is imposed without informing the fusion weights, deliberately testing variance weighting under misspecification. No learned fault classifier, covariance estimator, or disagreement-rejection mechanism was implemented.

For future deployment, bias estimates and uncertainty should come from held-out field references, with a minimum variance floor and stale-data exclusion. If residual dependence is estimated, the scalar extension is qhat = $(1^T R^{-1} y)/(1^T R^{-1} 1)$, where R is a positive-definite residual covariance matrix. This extension is proposed, not executed. If both channels are invalid, an approved local fallback should operate. Independence must not be assumed merely because the devices use different technologies [11].

4.7. Traffic-State Representation and Control Objective

Let qj(t) be unfinished workload and aj(t) the arriving workload during second t. With uj(t) equal to one for an actively served approach and zero otherwise, the executed update is qj(t+1) = max[0, qj(t) + aj(t) - 0.5uj(t)]. Each active approach can discharge 0.5 PCU/s. Arrivals occur before service, and queue workload is accumulated after service. The model permits fractional PCU workload and has unlimited storage.

The decision score for phase p is the sum of its two estimated queues. The operational objective is to reduce cumulative unfinished workload while limiting unnecessary phase changes. A score difference of 2 PCU creates hysteresis; minimum and maximum green rules constrain switching. This is a greedy queue-responsive heuristic inspired by queue-based control, not a full downstream-pressure model or a solution to a stated finite-horizon optimization problem [5].

4.8. Adaptive Signal-Control Algorithm

All controllers begin at phase 0 with empty queues. At eligible five-second ticks, the queue policy evaluates whether to switch. It retains the current phase until at least 15 seconds of green have elapsed. It then switches if the opposing score exceeds the active score by more than 2 PCU, if green has reached 60 seconds, or if the opposing phase's service age reaches 90 seconds. A requested switch inserts four service-free seconds before the next green. Green thresholds are checked on the five-second grid, so the nominal maximum may be exceeded by up to four seconds. During the clearance tail after arrivals stop, the last noise vectors are reused while state-dependent scales continue to update.

Pseudocode for each second is: add arrivals; if an eligible decision tick occurs, obtain the selected observation or fallback and evaluate the switching rules; enforce any active clearance; otherwise discharge the two active queues; update service age and accumulate queue workload. After 3,600 seconds, arrivals stop and the run continues until all workload exits. The executable script is authoritative for event ordering, tie handling, and the finite-horizon drain.

Green decisions depend on a fixed number of queues, so score computation is constant-size in this experiment. Processing an actual image is absent and would usually dominate the sensing workload. Python wall-clock execution time is recorded for reproducibility but is not a measurement of end-to-end system latency on an edge device. No inference accelerator throughput or communication performance is claimed.

4.9. Safety Constraints and Fallback Operation

The simulated minimum green, maximum green, clearance, and service-age rules are experimental parameters. They do not establish safe values for a real road. Yellow and all-red intervals are represented only by one aggregate service-free interval; pedestrian clearance, turning conflicts, emergency priority, and controller hardware interlocks are absent. Engineering design must determine those requirements from verified local geometry and applicable operating rules [3].

The experiment tests one camera outage with a surviving IoT channel. It does not test power loss, complete sensing loss, stale packets, corrupted commands, or recovery synchronization with physical signals. Field commissioning should include watchdog behavior, fallback entry, protected re-entry into adaptive operation, and independent acceptance of permitted phase transitions. Such checks are necessary additions to the proposed architecture, not completed outcomes of the simulation [3].

4.10. Simulation Environment and Calibration

The custom model is implemented in Python using NumPy and pandas, with Matplotlib for figures and SciPy for confidence intervals. Package versions and random seeds are included in the supplementary materials. There is no warm-up because each run represents a finite synthetic demand episode beginning with an empty intersection. The post-arrival drain prevents a controller from appearing efficient merely because some workload remains unserved at the reporting horizon.

This model is not calibrated. Its assumed total arrival

intensity is 0.5 PCU/s, or 1,800 PCU/h, across four approaches. Balanced rates are 0.125 PCU/s per approach; unbalanced rates are 0.175, 0.075, 0.175, and 0.075. Arrivals are independent Poisson packets. None of these rates is inferred from the Nairobi daily totals. A later calibrated microsimulation could use a documented package such as SUMO, but SUMO was not used to produce the present results [16].

A field-informed model would require turning counts, platoon arrivals, saturation flow, bus stopping, lane-use behavior, pedestrian demand, and measured queue or travel-time validation. Calibration and validation should use separate observation periods and declared acceptance criteria. A model that reproduces daily totals alone would remain insufficient to validate second-by-second adaptive operation. Evaluation guidance emphasizes matching operational objectives with suitable measurements and comparisons [17].

Table 3: Executed synthetic experiment parameters. Every numeric parameter below is an assumption, not an observed Nairobi junction value.

Parameter	Value or implementation
Queues and phases	Four workload queues; phase 0 serves 0 and 2; phase 1 serves 1 and 3
Arrival episode	3,600 s, empty initial state; full post-episode drain
Replications	30 seeds, 1000 through 1029; 6 controllers x 6 scenarios
Total intensity	0.5 PCU/s or 1,800 PCU/h; Poisson workload packets
Balanced rates	0.125 PCU/s on each approach
Unbalanced rates	0.175, 0.075, 0.175, 0.075 PCU/s
Service rate	0.5 PCU/s per active approach; unlimited queue storage
Adaptive timing	5 s decision grid; 15 s minimum and 60 s nominal maximum green
Switching	2 PCU hysteresis; 90 s opposing service-age trigger
Clearance	4 s aggregate service-free interval per switch
Fixed green times	40/40 s, or 56/24 s for initially unbalanced demand
Camera random error SD	2 + 0.15q PCU; clipped nonnegative observation
IoT random error SD	4 + 0.05q PCU; clipped nonnegative observation
Faults	35% underestimation on queues 0 and 2; or camera loss during 1200 <= t < 1800 s
Reversal	Directional rates swapped at second 1800; fixed plan unchanged
Tail observations	Final noise vectors reused after arrival period; state-dependent scales recomputed

4.11. Baseline Controllers and Experimental Scenarios

Equal fixed-time control uses 40 seconds of green for each phase. Demand-matched fixed-time control uses 56 and 24 seconds under initially unbalanced demand, preserving the 80-second total green budget while matching the 70:30 phase-demand ratio. It uses 40 and 40 seconds in the balanced scenario. The perfect-observation queue controller applies the same adaptive rules with exact queue state. It is an informational reference, not a mathematical optimum.

CCTV-only, IoT-only, and fused control share the adaptive decision rules. The demand-reversal scenario swaps the unbalanced approach rates at second 1,800; the demand-matched fixed plan retains its initial split. This deliberately tests a stale plan and must not be presented as a comparison with an optimally retimed plan. Camera bias, camera outage, and common bias otherwise share the unbalanced arrival pattern. High-demand saturation, mixed-vehicle dynamics, and full communication outages remain untested.

4.12. Performance Measures and Statistical Analysis

The principal reported control outcome is $D = \Sigma t \Sigma j\, qj(t)/\Sigma t \Sigma j\, aj(t)$, calculated over the arrival episode and full drain. Because service is fractional and the queues represent workload, D is a synthetic workload-delay proxy in seconds per PCU. Secondary outputs include mean and 95th-percentile network workload during the arrival period, discharged workload by second 3,600, maximum approach workload, drain duration, and switch count. No spatial spillback measure is available.

For each scenario-controller combination, the mean and Student t interval use 30 seed-level outcomes. For a comparator, paired differences are dr = Dfused,r - Dcomparator,r; their interval is mean(d) ± t0.975,29 sd(d)/√30. Relative reduction is 100(mean comparator - mean fused)/mean comparator. This is a ratio of means, not the mean of within-run percentages. Intervals are unadjusted exploratory 95% intervals; no familywise significance claim is made across the many comparisons.

There are 1,080 runs but only 30 independent random seeds per scenario. Seconds, queues, and observations inside a run are dependent and are not counted as additional replicates. Nairobi station changes are reported descriptively, without invented confidence intervals, because repeated-day sampling variation is unavailable. Field evaluation would require intersection- or day-level replication, a justified comparison design, and explicit control of temporal confounding [17].

4.13. Privacy Data Governance and Reproducibility

Only published aggregate survey values and synthetic workload are analyzed. No face image, registration plate, personal trajectory, or raw CCTV file is processed. The public-source documents remain hosted by their publishers; supplementary files provide the transcription, provenance, original code, generated tables, figure data, and software versions. The source report is cited rather than repackaged as an unrestricted dataset [1].

For a future camera study, the authors should document lawful access, institutional review requirements, purpose limitation, access roles, retention, masking, and rules for releasing annotations. These are proposed governance measures, not claims of permission or ethics approval. The experiment requires neither camera identities nor person re-identification. Actual author, institutional, funding, and ethics statements must be verified by the responsible researchers before journal submission.

5. Results

5.1. Observed Traffic Characteristics and Baseline Performance

Across the ten published station totals, traffic increased by 35.42% between the two survey years. Six stations increased and four decreased; the median station-specific change was 51.05%, with a range from -16.92% to 175.62%. These are descriptive changes across selected count locations, not a causal estimate of urbanization or a citywide congestion index. Table 2 and Figures 2 and 3 show the station-level variation calculated from the public data [1].

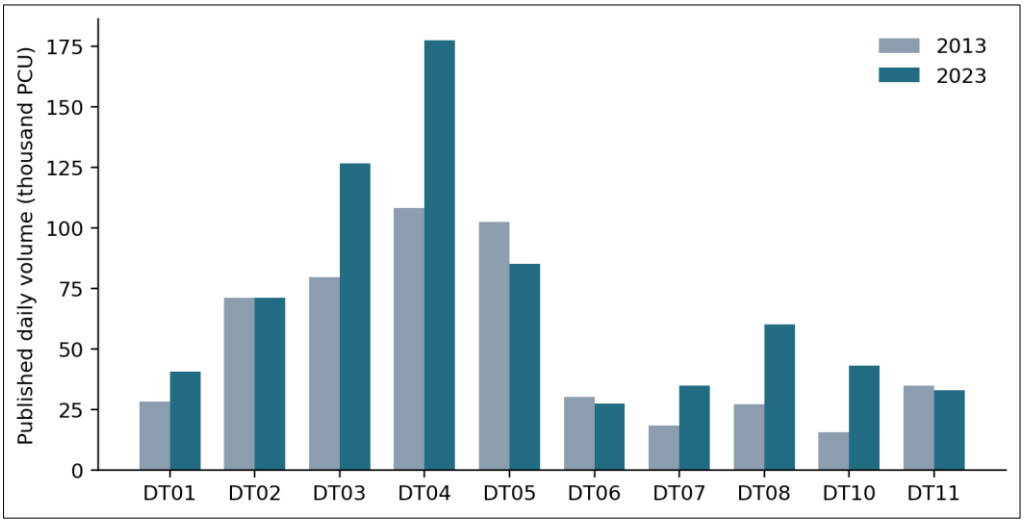


Fig 2: Published station totals for 2013 and 2023. Actual survey data in daily PCU, redrawn from JICA-NaMATA

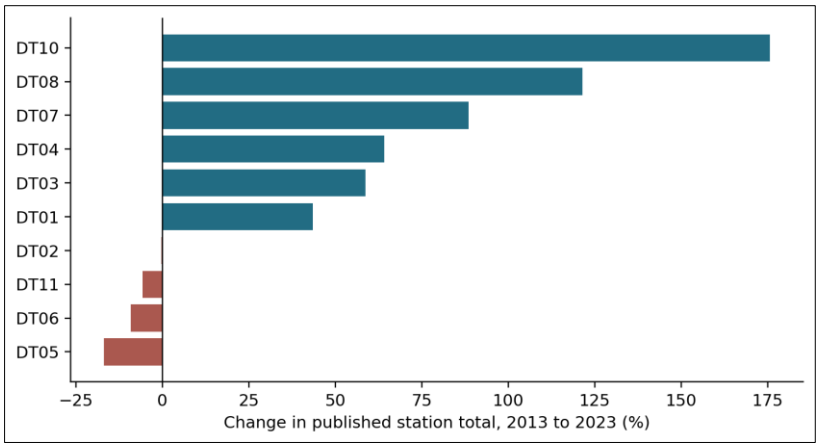


Fig 3: Percentage change in published station totals. Original calculations from actual survey data. The denominator is each station's 2013 published PCU total.

Recomputed class shares show substantial compositional change, but these are PCU shares rather than shares of vehicle numbers or passengers. Car/taxi accounts for 50.34% and 59.35% of recomputed class totals in the two years. The motorcycle component rises from 1.57% to 7.95%, while the heavy-truck/trailer component declines from 6.45% to 1.36%. Table 4 and Figure 4 use recomputed category sums consistently [1].

Table 4: Actual traffic-class data recomputed from detailed cells. Shares use 515,282 PCU in 2013 and 697,814 PCU in 2023. Source [1]. Published subtotals remain in the audit supplement.

Class	2013 PCU	2023 PCU	2013 share %	2023 share %
bicycle	1,833	1,231	0.36	0.18
motorcycle	8,090	55,486	1.57	7.95
car / taxi	259,381	414,156	50.34	59.35
light / truck	57,351	55,603	11.13	7.97
heavy / truck / trailer	33,242	9,494	6.45	1.36
matatu / minibus	87,402	42,567	16.96	6.10
bus	67,983	119,277	13.19	17.09

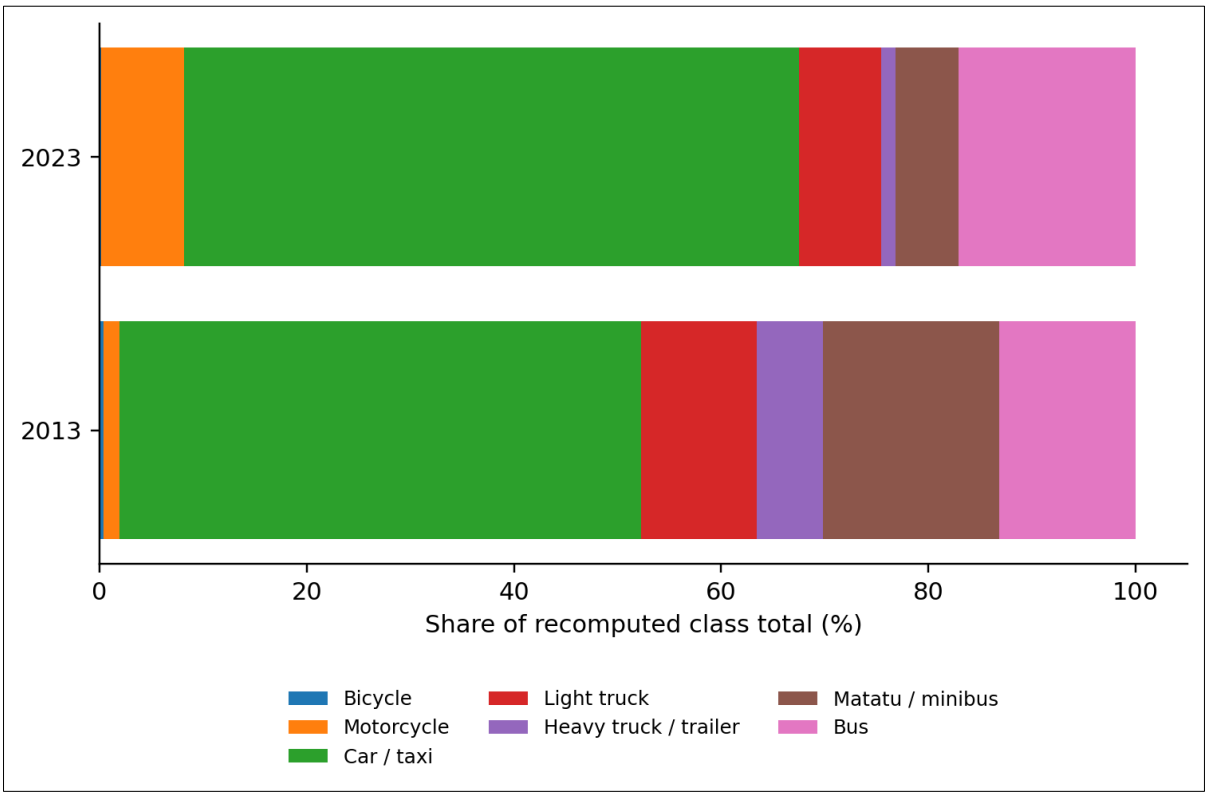


Fig 4: Traffic-class composition of actual survey PCU. Shares use recomputed class totals of 515,282 and 697,814 PCU. PCU shares must not be interpreted as vehicle-count or passenger shares.

The 2013 detailed entries sum exactly to their published total. In 2023, eight station rows differ by 1 or 2 PCU between detailed components and published row totals. Detailed cells sum to 697,814 PCU, whereas published station totals sum to 697,801 PCU, a difference of 13 PCU or approximately 0.0019%. Published 2023 class subtotals also sum to 697,802 PCU. These discrepancies are preserved in the audit file rather than silently corrected [1].

No observed junction-delay, spatial-queue, or existing-signal performance result can be derived from these aggregate records. The baseline-control figures in Table 5 are software experiment outputs. They are not baseline performance measured at the Nairobi count stations. This distinction also prevents daily traffic growth from being incorrectly used as a direct estimate of delay growth.

5.2. Accuracy of IoT CCTV and Fused Estimates

No field perception-validation result is available. The synthetic observation generator uses the declared Gaussian error functions and the simulator's exact workload as a reference. Inverse-variance fusion therefore has a favorable known-uncertainty setting in the unbiased scenarios. The resulting synthetic errors cannot validate a physical sensor, camera model, vehicle classifier, or local queue estimator.

The interpretation of perception accuracy is further limited by state dependence. Under the implemented model, uncertainty increases with the current queue, and each controller creates its own queue trajectory. Comparing each controller's average sensing error would combine estimation behavior with the controller's effect on the states being observed. A separate matched-state replay experiment is required before making a rigorous perception-only comparison. The current paper reports the observation mechanism and control ablations without claiming that such a perception benchmark was completed.

5.3. Comparative Traffic-Control Performance

In the unbalanced scenario, fused control gives a mean delay proxy of 12.91 s/PCU, with a 95% interval of 12.75 to 13.07. Equal fixed-time gives 22.03 s/PCU and demand-matched fixed-time gives 16.41 s/PCU. The paired fused-minus-demand-matched difference is -3.50 s/PCU, with a 95% interval of -3.74 to -3.27, corresponding to a 21.35% ratio-of-means reduction. These figures quantify performance within the assumed workload experiment.

The same scenario gives 13.29 s/PCU for CCTV-only control and 14.33 for IoT-only control. The incremental fusion benefit relative to CCTV is 2.87%, substantially smaller than its benefit relative to either fixed-time plan. Perfect-observation queue control gives 11.94 s/PCU, approximately 0.97 below fused control. Thus the results separate the value of adapting control, the incremental value of fusion, and the remaining penalty from imperfect observations.

Table 5: Synthetic mean workload-delay proxy in s/PCU, with unadjusted 95% confidence intervals. Each cell represents 30 seeds. No local field measurements are reported.

Scenario	Equal fixed	Demand matched	Perfect queue	CCTV only	IoT only	Fused
Balanced	19.34 (19.12, 19.55)	19.34 (19.12, 19.55)	11.81 (11.65, 11.97)	12.52 (12.37, 12.67)	13.17 (13.00, 13.35)	12.40 (12.23, 12.57)
Unbalanced	22.03 (21.49, 22.57)	16.41 (16.17, 16.66)	11.94 (11.82, 12.06)	13.29 (13.11, 13.48)	14.33 (14.10, 14.56)	12.91 (12.75, 13.07)
Demand reversal	21.66 (21.24, 22.08)	113.95 (107.55, 120.34)	11.93 (11.83, 12.03)	13.27 (13.10, 13.44)	14.50 (14.26, 14.73)	13.04 (12.87, 13.21)
Camera bias	22.03 (21.49, 22.57)	16.41 (16.17, 16.66)	11.94 (11.82, 12.06)	14.31 (14.11, 14.52)	14.33 (14.10, 14.56)	13.76 (13.53, 13.98)
Camera outage	22.03 (21.49, 22.57)	16.41 (16.17, 16.66)	11.94 (11.82, 12.06)	14.84 (14.66, 15.01)	14.33 (14.10, 14.56)	13.11 (12.95, 13.26)
Common bias	22.03 (21.49, 22.57)	16.41 (16.17, 16.66)	11.94 (11.82, 12.06)	14.31 (14.11, 14.52)	15.86 (15.59, 16.13)	14.25 (14.03, 14.46)

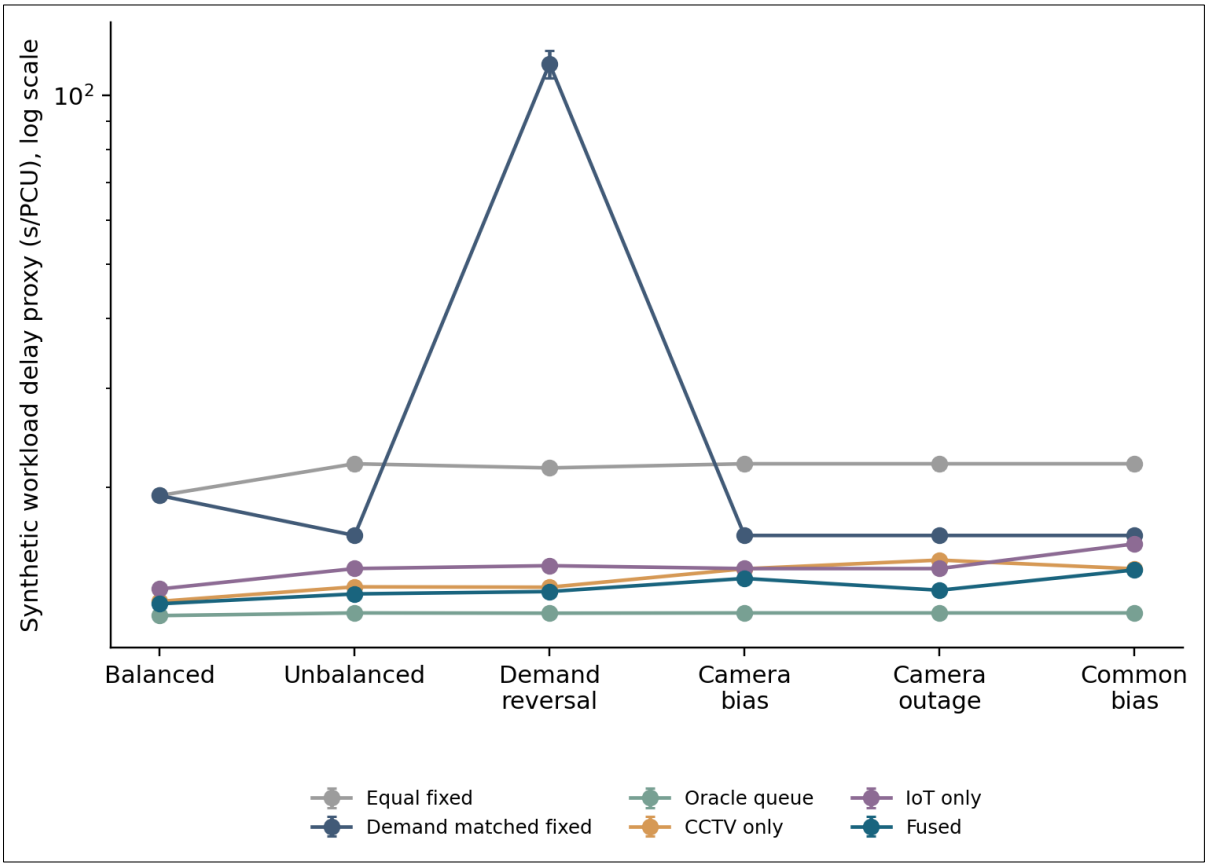


Fig 5: Executed synthetic workload experiment. Points are means over 30 paired seeds; bars are unadjusted 95% t intervals. The logarithmic vertical axis accommodates the stale demand-matched plan in the reversal scenario. Source: original simulation_replications.csv and simulate.py. No field traffic-control outcomes are plotted.

5.4. Performance Across Demand and Traffic Conditions

Balanced demand produces 12.40 s/PCU under fused control, compared with 19.34 for the identical equal and demand-matched fixed plans. Under demand reversal, fused control gives 13.04 s/PCU, whereas the initially matched but unchanged fixed plan gives 113.95. This large reversal difference is a consequence of the deliberately stale directional split, not evidence that fusion routinely improves real-world delay by a similar proportion.

The experiment does not vary total demand intensity or model different vehicle classes. Official PCU compositional changes provide contextual motivation for future mixed-traffic testing, but do not cause heterogeneous acceleration, motorcycle filtering, or bus dwell behavior in the simulator. Any interpretation concerning vehicle-group performance would therefore exceed the completed analysis.

5.5. Ablation Analysis

The balanced-demand paired difference between fused and CCTV-only control is -0.12 s/PCU, with a 95% interval of -0.26 to 0.01. Because the interval includes zero, this experiment does not clearly establish a balanced-demand advantage over CCTV alone. In unbalanced demand, the corresponding interval is -0.52 to -0.25. Table 6 and Figure 6 present these paired differences rather than relying on visual overlap of separate controller intervals.

Table 6: Synthetic paired ablation effects. Difference is fused minus CCTV-only delay proxy in s/PCU. Negative values favor fusion; intervals are unadjusted.

Scenario	Difference	95% CI lower	95% CI upper	Reduction %
Balanced	-0.124	-0.260	0.012	0.99
Unbalanced	-0.381	-0.516	-0.247	2.87
Demand reversal	-0.235	-0.376	-0.094	1.77
Camera bias	-0.556	-0.766	-0.346	3.88
Camera outage	-1.729	-1.902	-1.556	11.66
Common bias	-0.067	-0.273	0.138	0.47

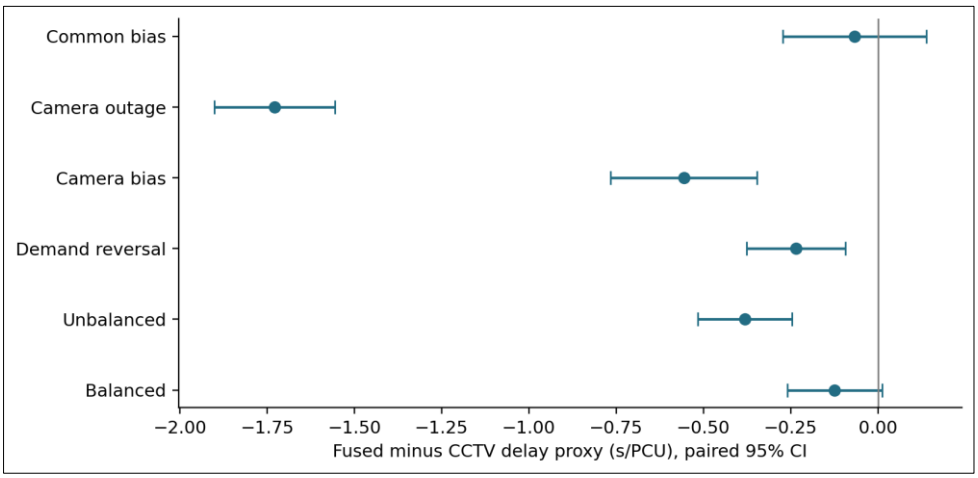


Fig 6: Incremental fusion effect relative to CCTV-only control under identical assumed arrivals and paired noise seeds. Negative values favor fusion. Intervals are exploratory unadjusted 95% t intervals across 30 seeds. Source: original table_paired_effects.csv.

With shared bias, the fused-minus-CCTV difference is -0.07 s/PCU, with an interval from -0.27 to 0.14. Fusion's reduction relative to CCTV is only 0.47% in that condition, compared with 11.66% during the camera outage. These results show that a surviving independent observation channel can be useful during an outage, while variance weighting alone does not establish a reliable improvement when both channels share systematic underestimation.

5.6. Robustness and Sensitivity Analysis

The camera-outage scenario yields 13.11 s/PCU for fused control and 14.84 for CCTV-only control with fallback. The paired difference is -1.73 s/PCU, with a 95% interval from -1.90 to -1.56. Camera bias gives 13.76 s/PCU for fused control, while common bias gives 14.25. These scenarios test particular fault mechanisms and do not represent a frequency-weighted estimate of real system reliability.

No latency sweep, saturation-demand sweep, covariance sweep, fusion-weight tuning study, or control-parameter sensitivity sweep was executed. The known variance functions, independent random errors, clearance length, green constraints, and hysteresis remain fixed. Reporting those omitted experiments as completed would be incorrect. They are prioritized in Section 7 because the present effects may change under other parameter choices.

5.7. Computational and Operational Feasibility

Every simulated workload episode was cleared completely, and the code asserts conservation of arrivals and total discharged workload. The released replication file retains queue, throughput, switch-count, drain, and software runtime diagnostics. Completion and conservation demonstrate internal computational consistency; they do not validate real-world traffic behavior, closed-loop safety, or equipment availability.

No hardware was purchased or benchmarked, and no installation or recurring cost was measured. Consequently, a cost-benefit ratio or claim of low-cost deployment is not reported. An implementation estimate would need camera reuse feasibility, mounting, power backup, communications, maintenance, annotation, software support, and controller integration. A field pilot should evaluate these alongside operational outcomes rather than assuming that software simulation establishes affordability [8, 17].

6. Discussion

6.1. Principal Findings and Answers to the Research Questions

RQ1 receives conditional support in the executed workload model: fused control reduces the reported delay proxy against both fixed plans under unbalanced demand. RQ2 has a more limited answer. Fusion improves on CCTV in some scenarios, but its balanced-demand and common-bias intervals include zero. This distinction matters because a large improvement over fixed-time operation can occur even when the additional sensing source contributes very little.

RQ3 shows that the initial split and subsequent demand changes materially affect the comparison. The reversal result highlights the vulnerability of a stale directional plan, but its size should not be generalized. RQ4 shows useful redundancy during the modeled camera outage and limited protection against shared bias. These answers support additional field validation, not an assertion that an integrated system is already effective at a Nairobi intersection.

6.2. Why the Integrated System Performed as Observed

The queue policy reallocates green in response to unfinished workload, while fixed plans cannot observe run-specific fluctuations. This mechanism explains much of the difference against fixed-time control. Queue-based research provides a

theoretical motivation for using congestion state, but the present heuristic lacks the full network structure and assumptions of max-pressure theory [5]. Its numerical performance should not be presented as a new stability theorem.

Fusion changes the information supplied to that policy. Under the declared independent random-error model, averaging observations by known variance can improve the queue estimate available at a decision tick. During camera loss, the IoT channel preserves state-responsive decisions. When both channels underestimate the same approaches, averaging cannot reconstruct the missing systematic component. The common-bias outcome is therefore consistent with the distinction between independent uncertainty and shared error discussed in fusion research [11].

The perfect-observation reference performs better on average in each reported scenario, but is not proven optimal. Its advantage identifies a potential sensing penalty under this policy, not the maximum possible value of improved perception under every controller. Furthermore, the uncertainty functions depend on simulator truth. A real fusion module would need to estimate its reliability from operationally available information, which could diminish the observed gains.

6.3. Comparison with Previous Studies

The findings are compatible with the broader motivation for responsive signal timing, but percentage improvements cannot be ranked directly against SCOOT, reinforcement-learning studies, or pressure-based arterial control. Those works address different networks, baselines, objectives, and observation assumptions [4, 6, 7]. The present study's most defensible comparison is internal: identical arrivals and decision rules, with only the sensing source changed.

Perception research likewise supplies methods rather than transferable accuracy guarantees. A tracking method validated on a recognized benchmark still requires a separate local counting and queue test [9, 10]. The aerial pNEUMA dataset and the Bengaluru CCTV preprint illustrate valuable but different observation domains [13, 14]. They are not substituted for missing Nairobi sensor streams. This prevents published detection scores from becoming invented local evidence.

6.4. Implications for Nairobi and Comparable Settings

The official data show that transport activity changes differently across selected locations. A useful municipal evaluation should therefore stratify sites by operating conditions and compare integration against well-maintained single-source alternatives. Historical PCU totals help define the context, but turning movements, pedestrian demand, storage, and actual phasing are the evidence needed to design a local controller evaluation [1, 3].

Operational adoption would require maintainable equipment, documented responsibilities, staff training, and a response process when observations degrade. Literature on intelligent transport in developing-country contexts supports considering infrastructure and implementation conditions alongside algorithms [12]. Transfer to another city should depend on comparable sensing coverage, control interfaces, demand structure, and maintenance arrangements. Geographic proximity alone does not establish transferability.

6.5. Service Distribution and Unintended Consequences

An average workload improvement can conceal delays on minor approaches or impose burdens on people who are not represented in the objective. The current model has phase-age constraints but no person-level weighting, bus occupancy, pedestrian queues, or vehicle-class outcomes. It cannot establish equitable service. A field evaluation should report approach-level delay, upper-tail waits, pedestrian service, and public-transport person-delay where reliable occupancy information exists [17].

Camera underestimation can also have distributional consequences if particular approaches or vehicle groups are consistently missed. A detector that performs well overall may still fail in crowded or poorly illuminated regions. Local annotation and error reporting should therefore be stratified by approach, relevant vehicle group, and operating condition. These are validation requirements motivated by perception research, not demographic or fairness findings established in this experiment [9].

6.6. Deployment and Scalability

Scaling from one synthetic junction to a corridor requires downstream storage, turning ratios, signal coordination, and communication dependencies. The present unlimited point queues cannot represent spillback or blocking. Network pressure-based research addresses interactions beyond this isolated model, and arterial learning methods consider coordination that is absent here [5, 7]. More sensors or a wider deployment will not automatically preserve the observed local effects.

A staged implementation should first verify independent sensing, then evaluate fused estimates through recorded replay, then use calibrated simulation, hardware testing, and an authorized pilot. Each stage should have explicit acceptance criteria and a tested fallback. Operational-objective validation should continue after initial deployment, because a short improvement does not establish long-term reliability or cost-effectiveness [17].

7. Limitations and Future Research

The geographical evidence is limited to a published comparison of ten Nairobi road-count stations. It does not represent all metropolitan roads or current intersection operations. Future work should collect synchronized approach-level observations at verified signalized sites and document site-selection bias. Survey-year changes also do not resolve seasonality, weather variation, or repeated-day measurement uncertainty [1].

The computational model is deliberately stylized. It assumes independent Poisson arrivals, constant service rates, two nonconflicting phases, unlimited storage, and fractional PCU service. It excludes turning interactions, heterogeneous vehicle dynamics, pedestrians, enforcement, downstream queues, and realistic incident behavior. A calibrated microscopic or mesoscopic model with held-out validation is needed before site-specific operational predictions are credible [16, 17].

The sensing model is unvalidated and gives the fusion rule known state-dependent variances. Independent Gaussian noise does not reproduce occlusion, false tracks, temporal drift, packet loss, or a real detector's accumulation error. Future work should first measure joint residual distributions and temporal correlation from synchronized references, then run a matched-state perception replay before closed-loop

testing. Full covariance fusion and explicit fault rejection should be evaluated against the simple weighting rule [11].

The experiment covers one total demand intensity, six scenarios, and fixed controller parameters. Future sensitivity work should vary demand to and beyond saturation, directional asymmetry, clearance time, maximum green, hysteresis, outage duration, error correlation, and uncertainty misspecification. The demand-matched baseline should also be supplemented by optimized time-of-day plans and a relevant actuated controller. Conclusions should be re-evaluated if those stronger comparators remove the apparent advantage.

The statistical analysis is exploratory. Thirty seeds estimate Monte Carlo variation within the assumed model, not variation across real days or junctions. Intervals are not adjusted for all comparisons, and the scenarios are not prevalence-weighted. Future confirmatory work should define a primary outcome, a minimum relevant benefit, an independent tuning set, and a site- or day-based sampling plan before testing. More simulation seeds alone cannot repair an invalid traffic or sensing model [17].

No operational device latency, energy use, maintenance cost, privacy implementation, or hardware safety behavior was validated. Future feasibility work should measure these directly and include unfavorable conditions. Claims concerning emissions, safety, or metropolitan economic benefit require their own valid outcomes and evaluation methods. This manuscript therefore confines its contribution to audited contextual analysis, executable computational comparisons, and a field-validation specification.

8. Conclusions

This study combines an audited official Nairobi traffic dataset with a separate, reproducible simulation of IoT-CCTV observation fusion and queue-responsive control. The contextual analysis identifies heterogeneous changes across ten road-count stations and preserves small source-table arithmetic inconsistencies. It does not convert daily survey totals into unobserved second-by-second junction measurements.

In the assumed unbalanced workload scenario, fused control gives 12.91 s/PCU compared with 22.03 for equal fixed-time and 16.41 for demand-matched fixed-time control. Its incremental reduction against CCTV alone is 2.87%, substantially smaller than its reduction against fixed-time control. Balanced demand and shared bias do not clearly establish a fusion advantage over CCTV alone. These results show why sensing ablations are necessary when evaluating integrated adaptive systems.

The established contribution is a transparent computational proof of concept and a reproducible data audit. Practical application in Nairobi requires independently validated sensing, stronger calibrated traffic models, approved safety integration, and field evaluation. The present findings do not establish real intersection delay reduction, improved pedestrian service, crash reduction, emissions savings, or network-wide benefit.

References

1. Japan International Cooperation Agency and Nairobi Metropolitan Area Transport Authority. The Project for Capacity Building for Bus Operation Policy and Management in Nairobi Metropolitan Area. Technical Report on Traffic Demand Forecast and Public Transport Strategy. June 2025. Attachment, Table 2.7, printed p. 11, PDF p. 17. Source link
2. Nairobi City County and Japan International Cooperation Agency. The Project on Integrated Urban Development Master Plan for the City of Nairobi in the Republic of Kenya. Final Report, Part III: Appendix. December 2014. Appendix 3, Traffic Survey Results. Source link
3. Koonce P, Rodegerdts L, Lee K, Quayle S, Beaird S, Braud C, Bonneson J, Tarnoff P, Urbanik T. Traffic Signal Timing Manual. Federal Highway Administration; 2008. Report FHWA-HOP-08-024. Source link
4. Hunt PB, Robertson DI, Bretherton RD, Winton RI. SCOOT: A Traffic Responsive Method of Coordinating Signals. Transport and Road Research Laboratory; 1981. Laboratory Report 1014. Source link
5. Varaiya P. Max pressure control of a network of signalized intersections. Transportation Research Part C: Emerging Technologies. 2013;36:177-195. doi:10.1016/j.trc.2013.08.014. Source link
6. Abdulhai B, Pringle R, Karakoulas GJ. Reinforcement learning for true adaptive traffic signal control. Journal of Transportation Engineering. 2003;129(3):278-285. doi:10.1061/(ASCE)0733-947X(2003)129:3(278). Source link
7. Wei H, Chen C, Zheng G, Wu K, Gayah V, Xu K, Li Z. PressLight: Learning Max Pressure Control to Coordinate Traffic Signals in Arterial Network. Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. 2019:1290-1298. doi:10.1145/3292500.3330949. Source link
8. Klein LA, Mills MK, Gibson DRP. Traffic Detector Handbook. Third Edition, Volume I. Federal Highway Administration; 2006. Report FHWA-HRT-06-108. Source link
9. Wen L, Du D, Cai Z, Lei Z, Chang MC, Qi H, Lim J, Yang MH, Lyu S. UA-DETRAC: A new benchmark and protocol for multi-object detection and tracking. Computer Vision and Image Understanding. 2020;193:102907. doi:10.1016/j.cviu.2020.102907. Source link
10. Zhang Y, Sun P, Jiang Y, Yu D, Weng F, Yuan Z, Luo P, Liu W, Wang X. ByteTrack: Multi-object tracking by associating every detection box. Computer Vision, ECCV 2022. Lecture Notes in Computer Science. 2022;13682:1-21. doi:10.1007/978-3-031-20047-2_1. Source link
11. Khaleghi B, Khamis A, Karray FO, Razavi SN. Multisensor data fusion: A review of the state-of-the-art. Information Fusion. 2013;14(1):28-44. doi:10.1016/j.inffus.2011.08.001. Source link
12. Mfenjou ML, Ari AAA, Abdou W, Spies F, Kolyang. Methodology and trends for an intelligent transport system in developing countries. Sustainable Computing: Informatics and Systems. 2018;19:96-111. doi:10.1016/j.suscom.2018.08.002. Source link
13. Barmpounakis E, Geroliminis N. On the new era of urban traffic monitoring with massive drone data: The pNEUMA large-scale field experiment. Transportation Research Part C: Emerging Technologies. 2020;111:50-71. doi:10.1016/j.trc.2019.11.023. Source link

14. Sharma A, Mhatre C, Gawali S, Bokkasam R, Sharma B, Pattanaik V, Rathore P, Krishnapuram R, Kovvali VG, Chakraborty A, Simmhan Y. BMD-45: A Large-Scale CCTV Vehicle Detection Dataset for Urban Traffic in Developing Cities. arXiv preprint. 2026. arXiv:2604.24419. Not peer reviewed at the cited version. Source link
15. Redmon J, Divvala S, Girshick R, Farhadi A. You only look once: Unified, real-time object detection. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 2016:779-788. doi:10.1109/CVPR.2016.91. Source link
16. Lopez PA, Behrisch M, Bieker-Walz L, Erdmann J, Flotterod YP, Hilbrich R, Lucken L, Rummel J, Wagner P, Wiessner E. Microscopic traffic simulation using SUMO. Proceedings of the IEEE International Conference on Intelligent Transportation Systems. 2018:2575-2582. doi:10.1109/ITSC.2018.8569938. Source link
17. Highway Administration F. Measures of Effectiveness and Validation Guidance for Adaptive Signal Control Technologies. July. Report FHWA-HOP-13-031. Source link; 2013.

How to Cite This Article

Afolabi AO, Omoifo-Irefo MO. Adaptive traffic signal control using IoT and CCTV integration in a rapidly urbanizing Sub-Saharan African city. Int J Eng Comput Appl. 2026;2(6):24-37. doi:10.54660/.IJECA.2026.2.6.24-37.

Creative Commons (CC) License

This is an open access journal, and articles are distributed under the terms of the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International (CC BY-NC-SA 4.0) License, which allows others to remix, tweak, and build upon the work non-commercially, as long as appropriate credit is given and the new creations are licensed under the identical terms.