

Resilient Cooperative Adaptive Cruise Control Under Intermittent V2V Connectivity and Cyber- Physical Disturbances

Maurice Larsen

Department of Computer Science, George Mason University, Fairfax, VA, USA.
contactmaurice@gmu.edu

Shengyong Yang

School of Information Technology, University of Cincinnati, Cincinnati, OH, USA.
yansiyuan@uc.edu

Abstract

Cooperative adaptive cruise control systems depend on real-time sensor data and vehicle-to-vehicle communication to maintain safe inter-vehicle gaps. However, operational deployments face stochastic distortions in sensor measurements and time-varying communication delays that can destabilize platoons, degrade string stability, and create fairness and safety risks. This paper provides a system-level examination of stochastic cooperative adaptive cruise control with emphasis on structural trade-offs rather than algorithmic derivation. It analyzes the layered sensing, communication, control, and actuation architecture, and discusses how data distortion enters through environmental noise, occlusion, calibration drift, and adversarial interference. It further examines how communication latency interacts with control loop timing, information flow topology, and feedback gains. The analysis shows that robustness cannot be achieved solely through improved filtering or more aggressive control; it requires architectural redundancy, delay-aware coordination, graceful degradation mechanisms, and context-sensitive trust between vehicles. The paper considers cross-domain lessons from networked control, cyber-physical systems, and intelligent transportation infrastructure. It also addresses governance challenges including safety certification, liability, algorithmic fairness, privacy, and long-term infrastructure investment. Policy implications are discussed in relation to heterogeneous fleets, equity in access to cooperative automation, and the need for standards that distinguish transient performance from bounded stability. The paper argues that sustainable deployment requires a shift from deterministic performance benchmarks toward risk-aware, adaptive, and institutionally embedded control architectures.

Keywords

cooperative adaptive cruise control; stochastic systems; communication delay; sensor distortion; string stability; intelligent transportation systems.

1. Introduction

Cooperative adaptive cruise control has emerged as a central mechanism for improving traffic stability, energy efficiency, and roadway capacity by enabling vehicles to coordinate longitudinal motion through vehicle-to-vehicle and vehicle-to-infrastructure communication. Early research and industry demonstrations treated the problem largely as deterministic: inter-vehicle spacing, relative velocity, and preceding acceleration were assumed to be available with bounded noise, while communication was assumed to be fast and reliable. A comprehensive review of communication, driver characteristics, and control dimensions

shows that these assumptions are often violated in realistic deployments [1]. In particular, sensor data can be distorted by weather, occlusion, surface reflections, and hardware aging, while communication links suffer from congestion, intermittent connectivity, and variable routing in vehicular environments. These uncertainties are not simple additive measurement errors; they interact with feedback dynamics, information topology, and controller gains in ways that can transform small perturbations into dangerous oscillations.

String stability has been a foundational concept in the analysis of vehicle platoons [2]. It describes whether spacing errors are attenuated or amplified as they propagate upstream through a platoon. The classical literature on disturbance propagation established that even stable individual vehicles can amplify disturbances in a string if their feedback policies are not designed with explicit coupling constraints [3]. Subsequent work on delayed acceleration feedback recognized that communication latency changes the phase of the feedback signal and can destroy string stability despite well-tuned local controllers [4]. These findings established a core system-level lesson: cooperative behavior depends not only on local control logic but also on the temporal and statistical properties of the signals exchanged between vehicles. A cyber-physical systems perspective places these challenges in a broader architecture of sensing, computation, communication, and actuation, where each component operates under resource constraints and uncertainty [5].

Although extensive literature exists on deterministic cooperative adaptive cruise control and on isolated delay compensation, there remains a need for a system-level examination of stochastic sensor data distortion and communication delay as coupled design constraints. This paper addresses that gap by discussing structural trade-offs across architecture, robustness, governance, deployment, and policy. Rather than proposing a specific estimator or control law, it focuses on how distortions and delays reshape system design priorities. The analysis draws on control theory, networked infrastructure studies, and policy scholarship to develop a perspective that is both technically informed and institutionally attentive.

2. System Architecture and Information Flow

Cooperative adaptive cruise control can be represented as a layered system spanning local sensing, data association, communication, decision making, and actuation. Within this structure, information flow topology is a primary determinant of stability and scalability. Research has shown that different predecessor-following, predecessor-leader following, and broadcast topologies produce distinct trade-offs between robustness and communication burden [6]. The choice of topology affects how delay and distortion are distributed: a local predecessor-following topology may reduce dependency on distant information but can propagate disturbances more readily, while broader topologies can improve preview but introduce more complex failure modes. The practical design of string-stable platoons therefore requires careful selection of feedback information, time headway, and control gains in relation to the expected communication environment [7].

Historical experimental comparisons of autonomous and cooperative follower control documented that cooperative strategies can substantially improve tracking performance and reduce inter-vehicle gaps, but they also revealed increased sensitivity to data quality and coordination failures [8]. This sensitivity is amplified in environments where security vulnerabilities can affect the integrity of connected vehicle data streams [9]. Malicious or accidental distortion may be indistinguishable from naturally stochastic sensor degradation, creating a difficult monitoring problem. From a system architecture perspective, cooperative driving thus requires layered defenses that include redundant sensing, plausibility checking,

source reputation management, and delay-aware trust assessment. Network-level solutions also matter, because connectivity, radio resource management, and handover behavior directly shape the delay distribution and packet loss characteristics experienced by control loops [10].

These architectural decisions are not only technical; they also determine certification boundaries and operational responsibilities. A design that relies on cloud-based trajectory planning will have different delay tolerances than one that uses only local predecessor state data. A design that assumes high-integrity roadside infrastructure may behave poorly when such infrastructure is absent or degraded. Therefore, system architecture must be evaluated under realistic failure conditions, including partial communication loss, sensor dropout, and conflicting information from multiple sources. The central challenge is to preserve an acceptable level of string stability without imposing excessive infrastructure and computational costs. This requires a deliberate allocation of authority among onboard sensing, vehicle-to-vehicle messages, roadside units, and central traffic management services.

3. Sensor Data Distortion and Stochastic Uncertainty

The characterization of sensor data distortion is complicated by the fact that errors are rarely stationary. Environmental obscuration, surface conditions, calibration drift, target occlusion, and sensor aging produce temporally correlated noise and occasional outlier events. In some cases, benign sensor degradation can mimic the statistical signatures of intentional attacks, making fault detection and security monitoring deeply intertwined. The broader institutional context matters as well: public adoption of automated and cooperative driving depends on societal trust in how such rare but high-consequence uncertainties are managed [11]. Consequently, the design of stochastic cooperative adaptive cruise control must consider not only average error but also tail risk, error correlation across vehicles, and the time scale over which sensors can recover from degraded states.

Field experiments with controlled autonomous vehicles have demonstrated that small perturbations in individual vehicle behavior can either dissipate or amplify into stop-and-go waves depending on the interaction between control logic, sensing, and vehicle response [12]. These empirical results underscore that stochastic distortion cannot be treated as a local filtering problem alone. When multiple vehicles share corrupted state estimates, their coordinated responses can synchronize into oscillatory patterns that are more severe than any single sensor error would suggest. The spatial and temporal correlation of distortion is therefore as important as its marginal distribution. Architects of cooperative systems need diagnostic strategies that can distinguish isolated sensor failure from correlated environmental degradation, because the two cases require different responses.

The formal analysis of string stability has expanded to include diverse information structures, heterogeneous vehicle dynamics, and uncertainty in feedback channels [13]. This literature shows that bounded measurement noise can produce unbounded spacing variation if the closed-loop sensitivity to high-frequency disturbances is not properly shaped. In a stochastic setting, the objective shifts from guaranteeing perfect attenuation of deterministic disturbances to bounding the probability with which spacing errors remain within safe envelopes. This requires a shift from deterministic pass or fail certification toward risk-aware performance indicators that account for rare events. Such indicators are more difficult to verify in real time, but they are necessary for meaningful safety assurance under uncertainty.

Data-driven methods from adjacent domains offer useful tools for understanding and predicting traffic and communication conditions. For example, cross-domain transfer learning has been used to predict cellular traffic patterns by leveraging large-scale data from multiple environments [14]. Similar approaches can support adaptive cooperative control by anticipating congestion in both the road traffic and the wireless communication domain. However, reliance on learned models introduces additional governance concerns, including model generalization, data provenance, and the possibility of silent performance degradation when deployment conditions differ from training distributions. These concerns are amplified in safety-critical tasks where sensor distortion and delay co-occur.

Recent work has addressed the integrated stochastic control problem in which sensor data distortion and communication delay are jointly modeled within cooperative adaptive cruise control [15]. That formulation recognizes that delay and distortion do not simply add noise; they change the information available for state reconstruction and the timing of corrective actions. From a system perspective, this integrated treatment is important because it forces designers to consider how controller parameters, estimate quality, and communication latency jointly determine safety margins. It also highlights that methods that are robust to delay but fragile under distortion, or vice versa, may be inadequate for real-world deployment. A coherent stochastic architecture must therefore balance estimator conservatism, controller responsiveness, and communication scheduling across heterogeneous vehicles.

4. Communication Delay and Cooperative Control

Communication latency in vehicular networks is shaped by radio channel conditions, medium access contention, network load, packet scheduling, and mobility-induced handover. Unlike a fixed computational delay, vehicular delay is strongly time-varying and often correlated with traffic density and roadside infrastructure availability. When vehicles respond to delayed state information, the effective feedback signal has a phase lag that can reduce damping and, under certain headway choices, produce oscillatory behavior. The system-level implication is that latency cannot be accepted as an external disturbance after controller design; it must be treated as part of the control architecture.

Several architectural responses to delay have been proposed. One strategy is to use predictive compensation, in which the controller estimates current states from delayed observations and a model of vehicle dynamics. Another strategy is to adjust communication scheduling so that safety-critical messages receive higher priority and lower queuing delay. A third strategy is to reduce dependence on time-sensitive data by designing control laws that tolerate bounded staleness. Each approach involves trade-offs. Predictive compensation improves responsiveness but increases sensitivity to model mismatch, especially when sensor distortion is also present. Priority scheduling can reduce average delay for critical links but may starve lower-priority traffic and create fairness concerns in mixed traffic. Robust delay-tolerant control can simplify real-time requirements but may sacrifice throughput by forcing larger gaps than necessary.

Delay effects also interact with platoon heterogeneity. Vehicles with different braking capabilities, communication stacks, or controller settings experience different effective delay sensitivities. A homogeneous platoon may be analyzable as a repetition of identical dynamics, but real deployments include mixed fleets and aftermarket devices. In such settings, the worst-performing node often determines the stability margin of the entire string. System architects therefore need to consider delay not only as a network-level average but as a distribution that varies across vehicles and time. This perspective supports decentralized

monitoring, adaptive headway policies, and the inclusion of communication quality as a first-class input to control decisions. It also suggests that certification should not assume a single delay bound for all vehicles and all roadways.

From an operational perspective, delay and sensor distortion are coupled because both reduce the quality of state estimates around the same decision point. A delayed but accurate measurement may be preferable to a timely but severely distorted one in some situations, while the opposite may hold in others. The appropriate response depends on the vehicle dynamics, traffic scene complexity, and available fallback sensing. This coupling makes it difficult to optimize delay compensation and signal processing independently. Instead, a cross-layer approach is needed, in which communication scheduling, sensor fusion, and controller adaptation are jointly configured to maintain safety under uncertain information.

5. Infrastructure Deployment and System Integration

Deployment of stochastic cooperative adaptive cruise control requires more than in-vehicle control. Roadside infrastructure, edge computing, and traffic management centers can provide complementary information that reduces individual uncertainty. Cooperative intersection management, for example, coordinates vehicles and signal systems to improve safety and capacity, but its benefits depend on reliable communication and appropriate fallback behavior [16]. When cooperative adaptive cruise control is integrated with intersection coordination, the same sensor distortion and delay problems reappear at a larger spatial scale. A vehicle approaching an intersection must fuse local sensing, infrastructure messages, and possibly data from other road users under tight timing constraints. System-level integration therefore demands standardized message sets, interoperability testing, and clear allocation of authority among vehicle, infrastructure, and cloud services.

Infrastructure can also serve as a stabilizing resource by providing reference trajectories, traffic state estimates, or communication relays. However, this creates dependency on public investment and maintenance. The reliability of public infrastructure may vary across jurisdictions, creating spatial inequities in the quality of cooperative automation. A vehicle that benefits from well-maintained roadside units in one region may lose that support in another, requiring graceful degradation to local sensing with larger headways and lower speeds. From a governance perspective, this variability argues for minimum service levels and open standards that prevent vendor-specific fragmentation. It also raises questions about who bears responsibility when an infrastructure message is delayed or distorted and contributes to a safety-critical event.

The integration of stochastic control with infrastructure also requires alignment between real-time control loops and slower planning cycles. Traffic management centers typically operate on longer timescales than vehicle controllers, and their advice may be useful even if it is not fresh enough for immediate control. One architectural option is to separate safety-critical real-time coordination from advisory services that influence route choice, speed smoothing, and platoon formation. Such separation can reduce coupling and make certification more tractable, but it also requires careful handoff procedures when advisory information becomes safety-relevant. These handoffs are especially challenging when sensor distortion or communication delay creates uncertainty about whether the vehicle should rely on its own estimate or an external reference.

6. Policy, Fairness, and Sustainability

The long-term societal impact of stochastic cooperative adaptive cruise control depends not only on technical performance but also on how rapidly and equitably the technologies are adopted. Forecasting studies indicate that consumer adoption of connected and autonomous vehicle technologies is shaped not only by perceived safety but also by cost, regulatory clarity, and the perceived reliability of cooperative systems [11]. However, the deployment of such systems also raises questions about equitable access to the benefits of automation. If the infrastructure and communication support required for stochastic cooperative adaptive cruise control are deployed unevenly, then early benefits will accrue to affluent corridors and well-resourced jurisdictions, while lower-income and rural communities may experience delayed or degraded functionality. A justice-oriented analysis of transportation technology suggests that this divergence can reinforce existing mobility inequities [19]. From a systems governance standpoint, policies must therefore address minimum service levels, open interfaces, and public investment in connected infrastructure, rather than assuming that private-sector adoption alone will produce equitable outcomes.

Sustainability is also a multi-dimensional concern. Cooperative adaptive cruise control can reduce fuel consumption through smoother acceleration and braking, but the energy burden of sensing, computing, and communication equipment, along with the energy demands of roadside and edge infrastructure, must be accounted for in lifecycle assessments. Studies on the travel, energy, and carbon impacts of highly automated vehicles indicate that the net effect depends heavily on induced demand, platooning penetration, and vehicle occupancy [20]. A stochastic control architecture that tolerates high sensor distortion and communication delay may require larger headways or additional infrastructure, which could erode throughput benefits and increase energy use. Therefore, the design objective should not be minimal headway under ideal conditions but bounded risk under realistic uncertainty with acceptable energy and equity outcomes.

Deployment of stochastic cooperative adaptive cruise control also depends on the availability of low-latency computation and reliable communication. Multi-access edge computing can offload some state estimation and platoon coordination tasks from vehicles to nearby infrastructure, reducing the computational burden on individual agents and enabling richer traffic-level coordination [17]. However, edge resources are themselves finite and must be scheduled under varying load. As the number of connected vehicles grows, contention for spectrum and edge processing can introduce new sources of delay and data distortion. System architects must therefore consider not only the nominal capacity of these resources but also their performance under peak load, failure, and adversarial conditions. A robust deployment strategy would allocate edge computing and communication priority based on safety-criticality and the current level of sensor uncertainty, rather than on a fixed service class.

Safety certification for stochastic cooperative adaptive cruise control is more complex than for deterministic systems because it must account for joint distributions of sensor distortion, communication delay, and environmental conditions. Formal verification methods that rely on worst-case bounds may be overly conservative, while purely statistical testing may fail to capture rare combinations of adverse events. Validation approaches must combine high-fidelity simulation, closed-track testing, and carefully designed real-world pilots to stress the system under correlated distortions and delays [18]. Furthermore, certification should be executable throughout the system lifecycle, including over-the-air updates and changes in the communication environment. A governance framework that requires continuous risk

assessment and incident reporting, rather than one-time type approval, is more consistent with the dynamic nature of stochastic cooperative control.

7. Conclusion

This paper has examined stochastic cooperative adaptive cruise control as a systems problem in which sensor data distortion and communication delay are not isolated disturbances but interacting constraints that shape architecture, robustness, governance, and deployment. The analysis has shown that traditional deterministic approaches, which assume reliable sensing and negligible latency, are insufficient for real-world operation. Instead, system designers must account for the statistical properties of distortion, the temporal structure of delay, and the correlation between uncertainty sources across multiple vehicles. Information flow topology, controller gains, and time headway choices must be co-designed with sensing redundancy and delay-aware communication strategies, because their interactions determine whether spacing errors remain bounded or propagate dangerously through a platoon.

A central finding is that robustness cannot be achieved by adding a single filter or estimator. It requires layered defenses that include plausibility checking, source reputation management, gracefully degrading control modes, and the ability to fall back to local sensing when cooperative information becomes unreliable. The coupling between sensor distortion and communication delay means that optimization of one without regard to the other may produce fragile designs. Similarly, edge computing and roadside infrastructure can reduce uncertainty but also introduce new failure modes and resource contention. Therefore, infrastructure investment must be accompanied by open standards and equitable deployment policies to avoid deepening mobility divides.

Policy and certification frameworks must move from deterministic pass or fail criteria toward risk-aware, continuous assessment methods that reflect the stochastic nature of cooperative control. This shift is necessary not only for safety but also for building public trust, since rare high-consequence events under uncertain information can undermine confidence even when average performance is excellent. Future research should continue to integrate control theory, communication network modeling, and institutional analysis to develop architectures that are safe, fair, and sustainable under the combined stress of sensor distortion and communication delay.

References

1. Dey, K. C., Yan, L., Wang, X., Wang, Y., Shen, H., Chowdhury, M., Yu, L., Qiu, C., & Soundararaj, V. (2016). A review of communication, driver characteristics, and controls aspects of cooperative adaptive cruise control. *IEEE Transactions on Intelligent Transportation Systems*, 17(2), 491-509.
2. Swaroop, D., & Hedrick, J. K. (1996). String stability of interconnected systems. *IEEE Transactions on Automatic Control*, 41(3), 349-357.
3. Seiler, P., Pant, A., & Hedrick, K. (2004). Disturbance propagation in vehicle strings. *IEEE Transactions on Automatic Control*, 49(10), 1835-1842.
4. Xiao, L., & Gao, F. (2011). Practical string stability of platoon of adaptive cruise control vehicles. *IEEE Transactions on Intelligent Transportation Systems*, 12(4), 1184-1194.

5. Lee, E. A. (2008). Cyber physical systems: Design challenges. In 2008 11th IEEE International Symposium on Object and Component-Oriented Real-Time Distributed Computing (ISORC) (pp. 363-369). IEEE.
6. Zheng, Y., Li, S. E., Wang, J., Cao, D., & Li, K. (2016). Stability and scalability of homogeneous vehicular platoon: Study on the influence of information flow topologies. *IEEE Transactions on Intelligent Transportation Systems*, 17(1), 14-26.
7. Milanés, V., Shladover, S. E., Spring, J., Nowakowski, C., Kawazoe, H., & Nakamura, M. (2014). Cooperative adaptive cruise control in real traffic situations. *IEEE Transactions on Intelligent Transportation Systems*, 15(1), 296-305.
8. Rajamani, R., & Shladover, S. E. (2001). An experimental comparative study of autonomous and cooperative vehicle-follower control systems. *Transportation Research Part C: Emerging Technologies*, 9(1), 15-31.
9. Petit, J., & Shladover, S. E. (2015). Potential cyberattacks on automated vehicles. *IEEE Transactions on Intelligent Transportation Systems*, 16(2), 546-556.
10. Zheng, K., Zheng, Q., Chatzimisios, P., Xiang, W., & Zhou, Y. (2015). Heterogeneous vehicular networking: A survey on architecture, challenges, and solutions. *IEEE Communications Surveys & Tutorials*, 17(4), 2377-2396.
11. Kyriakidis, M., Happee, R., & de Winter, J. C. F. (2015). Public opinion on automated driving: Results of an international questionnaire among 5000 respondents. *Transportation Research Part F: Traffic Psychology and Behaviour*, 32, 127-140.
12. Stern, R. E., Cui, S., Delle Monache, M. L., Bhadani, R., Bunting, M., Churchill, M., Hamilton, N., Haulcy, R., Pohlmann, H., Wu, F., Piccoli, B., Seibold, B., Sprinkle, J., & Work, D. B. (2018). Dissipation of stop-and-go waves via control of autonomous vehicles: Field experiments. *Transportation Research Part C: Emerging Technologies*, 89, 205-221.
13. Ploeg, J., Semsar-Kazerooni, E., Lijster, G., van de Wouw, N., & Nijmeijer, H. (2015). Graceful degradation of cooperative adaptive cruise control. *IEEE Transactions on Intelligent Transportation Systems*, 16(1), 488-497.
14. Zhang, C., Zhang, H., Qiao, J., Yuan, D., & Zhang, M. (2019). Deep transfer learning for intelligent cellular traffic prediction based on cross-domain big data. *IEEE Journal on Selected Areas in Communications*, 37(6), 1389-1401.
15. Cai, G., Yin, G., Liu, Y., Feng, J., Liang, J., Wang, F., & Liu, H. (2025). Stochastic cooperative adaptive cruise control with sensor data distortion and communication delay. *IEEE Transactions on Intelligent Transportation Systems*, 26(7), 9500-9515.
16. Lee, J., & Park, B. (2012). Development and evaluation of a cooperative vehicle intersection control algorithm under the connected vehicles environment. *IEEE Transactions on Intelligent Transportation Systems*, 13(1), 81-90.
17. Taleb, T., Samdanis, K., Mada, B., Flinck, H., Dutta, S., & Sabella, D. (2017). On multi-access edge computing: A survey of the emerging 5G network edge cloud architecture and orchestration. *IEEE Communications Surveys & Tutorials*, 19(3), 1657-1681.
18. Koopman, P., & Wagner, M. (2016). Challenges in autonomous vehicle testing and validation. *SAE International Journal of Transportation Safety*, 4(1), 15-24.

19. Golub, A., & Martens, K. (2014). Using principles of justice to assess the modal equity of regional transportation plans. *Journal of Transport Geography*, 41, 10-20.
20. Wadud, Z., MacKenzie, D., & Leiby, P. (2016). Help or hindrance? The travel, energy and carbon impacts of highly automated vehicles. *Transportation Research Part A: Policy and Practice*, 86, 1-18.