

Event-Driven Safety-Critical Control for Distributed Electric Vehicles with Energy-Efficient Lateral Stability Management

Gmash Rhandra

Department of Computer Science, Colorado State University, Fort Collins, CO, USA.
umesh.chandra937@colostate.edu

Pavel Love

Department of Computer Science, University of North Texas, Denton, TX, USA.
pavellove455@unt.edu

Kartik M. Srivastava

Department of Computer Science and Engineering, University at Buffalo, Buffalo, NY, USA.
kartik1972@buffalo.edu

Abstract

The electrification of vehicle platforms has accelerated the shift from centralized internal combustion drivelines to distributed electric propulsion architectures, in which independently controlled wheel actuators offer new opportunities and new safety obligations. This paper presents a system-level examination of event-driven safety-critical control for distributed electric vehicles, focusing on the coordination of lateral stability management and energy-efficient actuation. Rather than proposing a single control law, the work analyzes the architectural trade-offs that arise when event-triggered stability supervision is coupled with constrained torque allocation across multiple actuators. The discussion addresses the role of safety regions, event conditions, and the sequencing of stability interventions in reducing unnecessary actuation while preserving yaw stability margins. It further examines how such event-driven architectures interact with in-vehicle communication, software certification, and cross-domain data analytics. The paper argues that energy-efficient lateral stability management is not merely a control allocation problem but a socio-technical system design challenge involving robustness, fairness in actuator utilization, lifecycle energy impacts, and regulatory compliance. By integrating perspectives from vehicle dynamics, real-time systems, machine learning governance, and safety engineering, the paper develops a structured research agenda for future distributed electric vehicle platforms. The analysis highlights the importance of treating event-driven safety control as an infrastructural capability rather than a standalone algorithm, with implications for certification, deployment, and policy.

Keywords

event-driven control; distributed electric vehicles; lateral stability; energy efficiency; safety-critical systems; system architecture.

1. Introduction

The transition toward electrified mobility has fundamentally altered the control architecture of modern road vehicles. Traditional stability control systems, developed around hydraulic braking and internal combustion drivelines, were designed to modulate yaw moments through

differential braking with limited bandwidth and significant energy penalties [1], [2]. Distributed electric vehicles replace this centralized actuation model with multiple independently controlled motor drives, enabling direct yaw moment generation, finer torque resolution, and new possibilities for regenerative energy management [3]. However, this architectural shift also introduces a more complex safety-critical control problem. The availability of multiple actuators increases the dimensionality of the control allocation problem, while the need to preserve lateral stability under rapidly changing operating conditions places stringent demands on real-time supervision. Within this setting, lateral stability management must reconcile two potentially conflicting objectives: maintaining the vehicle within safe yaw rate and sideslip bounds, and minimizing the electrical energy consumed by stability interventions [4]. The system-level challenge is therefore not only to design a stabilizing controller but also to determine when and how strongly such control should be engaged.

The emergence of data-driven methods in large-scale cyber-physical systems provides additional context for distributed electric vehicle control. Cross-domain learning approaches have demonstrated that operational patterns can be transferred across heterogeneous sources to improve prediction and resource allocation [5]. Although this class of methods originated in network traffic analysis, its conceptual value extends to vehicle fleets, where high-dimensional operational data can inform stability monitoring, actuator health estimation, and energy-aware control strategies. Nevertheless, safety-critical vehicle control cannot rely solely on learning-based inference because of latency, explainability, and certification constraints. This tension motivates a hybrid architecture in which event-driven mechanisms select when complex computations or control updates are necessary, while well-understood feedback principles enforce stability.

This paper examines the structural, algorithmic, and policy dimensions of event-driven safety-critical control for distributed electric vehicles. The analysis focuses on lateral stability management as a representative safety-critical function and on energy efficiency as a first-order system constraint. The remainder of the paper is organized as follows. Section 2 reviews related concepts in vehicle stability control and event-triggered systems. Section 3 presents an architectural view of event-driven safety-critical control. Section 4 discusses distributed electric vehicle dynamics and lateral stability. Section 5 examines energy-efficient stability management. Subsequent sections address system integration, robustness, fairness, governance, deployment, and policy implications.

2. Related Work and System Context

Classical vehicle stability control has been extensively studied in the context of yaw moment generation and sideslip regulation. Early electronic stability programs used differential braking to generate corrective yaw moments, relying on estimates of yaw rate and lateral acceleration to detect incipient instability [2]. These systems were effective but limited by actuator bandwidth, hydraulic pressure dynamics, and the need to avoid excessive brake wear. More recent work has shown that constraints on path following can be explicitly incorporated into vehicle control design, allowing stability and trajectory tracking objectives to be handled within a unified framework [6]. For distributed electric vehicles, torque vectoring techniques have been compared under varying road friction conditions, demonstrating that independently controlled motors can achieve more precise yaw moment distribution than brake-based systems [7]. These developments illustrate a broader trend from centralized braking-based stability control toward distributed actuation and optimization-based allocation.

The event-driven control paradigm offers a complementary approach to constrained real-time systems. In time-triggered systems, control updates occur at fixed intervals determined by worst-case assumptions, which leads to over-utilization of computation and communication resources under benign operating conditions. Event-triggered control instead initiates updates only when a specified condition on system state or performance crosses a threshold [8]. Subsequent surveys formalized event-triggered and self-triggered strategies, clarifying stability-preserving event conditions, inter-event time properties, and implementation requirements [9]. Early comparative studies established that event-based sampling can outperform periodic sampling for certain stochastic systems by reducing the number of samples while preserving performance [10]. When applied to vehicle stability, event-driven strategies can significantly reduce the frequency of high-level control recomputation without sacrificing safety if the triggering conditions are designed using safety regions [11]. This observation connects vehicle stability control to the broader literature on event-based sampling and networked control, where stability preservation under reduced communication and computation is a central concern [9], [10]. In distributed electric vehicles, however, the event condition must reflect not only state error thresholds but also actuator energy budgets, road friction uncertainty, and the possibility of actuator degradation. Consequently, the design of an event-triggered stability manager becomes a systems problem involving state estimation, prediction of future vehicle states, and coordination among multiple control layers.

3. Architectural Foundations of Event-Driven Safety-Critical Control

A layered architecture is essential for reconciling the requirements of safety supervision and energy optimization in distributed electric vehicles. The lowest layer comprises per-wheel actuator controllers that enforce torque commands within physical and thermal limits. The middle layer performs control allocation, mapping the desired longitudinal force and yaw moment onto the available wheel actuators while respecting tire force constraints. The upper layer is responsible for model predictive trajectory and stability supervision, generating nominal motion commands and monitoring safety margin. In an event-driven formulation, the upper layer does not recompute at a fixed period. Instead, the system maintains a nominal command trajectory and an associated safety region, and a new optimization is triggered only when the measured or estimated vehicle state approaches the boundary of that region or when a disturbance invalidates the predicted margin [11]. This structure reduces the computational load on high-performance processors and lowers the average communication traffic between distributed controllers.

The safety region concept is central to this architecture. A safety region defines the admissible set of vehicle states, such as yaw rate, sideslip angle, lateral acceleration, and their rates, within which the vehicle remains controllable under the available actuator authority. The boundary of this region depends on road friction, tire characteristics, vehicle speed, and actuator health. Event conditions therefore become boundary-proximity tests rather than simple threshold comparisons. When the vehicle state is deep inside the safety region, energy efficiency can be prioritized, allowing the control allocator to favor actuators with higher efficiency or regenerative potential. When the state approaches the boundary, safety becomes dominant, and the allocator must sacrifice energy optimality to generate corrective yaw moments. The event-driven mechanism thus acts as a governance layer that schedules the transition between efficiency-oriented and safety-oriented control modes [12]. This

scheduling is not merely a software detail; it defines the system-level behavior under transient maneuvers and degraded conditions.

The design trade-offs in such architectures are significant. A narrow safety region triggers stability interventions more frequently, reducing energy efficiency and increasing actuator wear. A wide safety region may delay necessary corrections, especially under abrupt friction transitions or tire saturation. Moreover, the safety region itself is not static; it evolves with vehicle speed, load transfer, steering input, and environmental conditions. Modern control strategies can compute robust invariant sets or reachable sets to characterize admissible states under bounded disturbances [13]. These methods provide formal guarantees but often require high-dimensional optimization that is difficult to execute in real time. Event-driven computation addresses this difficulty by performing expensive set computations only when the vehicle approaches the boundary, while using lightweight propagation models betweenwhile using lightweight propagation models between safety-critical updates. This separation of concerns mirrors layered architectures in large-scale infrastructure systems, where high-frequency local control operates autonomously and infrequent global coordination occurs only when operational envelopes are threatened. The event-driven stability manager therefore functions as a supervisory governance layer that enforces the transition from nominal energy-prioritized allocation to safety-prioritized intervention without requiring continual global optimization. In practice, this requires an explicit interface contract among the sensing, estimation, allocation, and actuation layers. The contract must define the data freshness bounds, the authority of each layer to override nominal commands, and the preconditions under which a safety-critical event may interrupt energy-optimal operation. Such contracts are not merely technical specifications; they encode assumptions about actuator health, communication latency, and environmental uncertainty that must be verified during system integration.

A significant architectural trade-off concerns the granularity of event conditions. If events are defined at the level of individual wheel torque commands, the system may generate a large number of low-amplitude adjustments that lead to excessive coordination overhead and actuator chattering. If events are defined only at the vehicle yaw rate level, the system may miss local tire saturation or actuator saturation that precedes global instability. The appropriate granularity depends on the dynamic coupling among wheels, the predictive horizon of the stability monitor, and the available sensing precision. In distributed electric vehicles, local wheel controllers can often detect early signs of tire force saturation through changes in slip slope or motor current signatures. Allowing these local controllers to trigger localized events, while a higher-level monitor filters and prioritizes them, can improve robustness without sacrificing centralized safety oversight. However, this dual-level event structure increases the state space of the event scheduler and complicates the formal certification of timing behavior.

Another architectural consideration is the choice between event-triggered and self-triggered mechanisms. In event-triggered approaches, control updates occur when a state-dependent condition crosses a threshold, which reduces average resource consumption but introduces variable and potentially unbounded inter-event times. In self-triggered approaches, the next update time is computed from the current state, allowing more predictable scheduling but requiring a conservative bound on future disturbances. Distributed electric vehicles operating in mixed traffic may benefit from a hybrid strategy in which self-triggered predictions are used under nominal conditions, while event-triggered overrides occur when unexpected

disturbances or sensor anomalies are detected. The resulting timing behavior is neither fully periodic nor fully stochastic, and its certification requires compositional reasoning across vehicle layers, a topic that remains open in the safety-critical embedded systems literature.

4. Distributed Electric Vehicle Dynamics and Lateral Stability

Lateral stability in distributed electric vehicles is governed by the coupling between tire lateral forces, longitudinal force distribution, vehicle speed, and road surface conditions. Unlike conventional vehicles with a single powertrain, distributed architectures can generate direct yaw moments by commanding asymmetric wheel torques. This capability allows more rapid yaw rate correction and finer resolution than brake-based stability systems, but it also introduces sensitivity to individual actuator faults and tire force saturation. The allocation of longitudinal forces across wheels changes the available lateral tire force, because tire friction limits are shared between longitudinal and lateral directions. When a wheel is used heavily for propulsion or braking, its lateral force capacity decreases, which can reduce the vehicle's ability to resist side slip. Consequently, lateral stability management in distributed electric vehicles cannot be treated independently of longitudinal torque allocation. The two domains are physically coupled through the tire force ellipse and through load transfer effects during acceleration, braking, and steering [14].

The safety region formulation provides a systematic way to represent admissible combinations of yaw rate and sideslip angle under current friction and actuator constraints. In contrast to fixed threshold-based stability programs, safety regions adapt to vehicle speed and estimated road friction, thereby avoiding unnecessary interventions when the vehicle is crossing moderate sideslip transients that remain fully recoverable [11]. Predictive approaches extend this idea by projecting the vehicle state forward over a short horizon and checking whether the predicted trajectory stays within the safety region. Model predictive control has been widely applied to vehicle stability and path tracking because it can naturally incorporate actuator limits, tire force constraints, and multiple objectives [15]. In distributed electric vehicles, the optimization problem includes the allocation of forces among wheels, which increases the number of decision variables compared to conventional steering or braking-only systems. Event-driven variants reduce the frequency of solving this optimization by invoking it only when the predicted state approaches the safety boundary or when a change in measured road friction invalidates the current plan.

Motion planning and control frameworks for autonomous vehicle systems provide useful architectural parallels, even though the present paper focuses on human-driven or semi-autonomous distributed electric vehicles. In these frameworks, global planning and local feedback control are separated by an interface that enforces feasibility and safety [16]. Similarly, a distributed electric vehicle stability system can be organized into a nominal trajectory generation layer, a safety monitor, and an actuator allocation layer. The safety monitor evaluates whether current commands, if continued, would keep the vehicle within a robust safe set. In cases where the nominal commands are energy-optimal but approach the safe set boundary, the monitor must decide whether to adjust the commands proactively or wait for a triggering event. The trade-off between proactive and reactive intervention is a central topic in safety-critical control and has direct implications for energy efficiency, because proactive adjustments may sacrifice some energy savings to avoid a later aggressive and more energy-intensive correction [17].

5. Energy-Efficient Stability Management

Energy efficiency in lateral stability management arises from the fact that corrective yaw moments consume electrical energy, reduce regenerative braking opportunities, and may increase tire wear and thermal losses. In centralized hydraulic stability systems, the energy cost of stability intervention is primarily expressed through brake energy dissipation. In distributed electric vehicles, the cost structure is more nuanced. Some stability corrections can be performed by redistributing torque from one wheel to another without increasing total tractive effort, thereby preserving kinetic energy. Other corrections require a global decrease in propulsion torque or a regenerative braking action, which may still recover some energy. The optimal allocation depends on the instantaneous vehicle state, motor efficiency maps, battery state of charge, and the available tire force margins [14]. Energy-efficient control allocation for over-actuated vehicles therefore aims to minimize a weighted sum of energy loss and stability violation, where the weights change as the vehicle approaches the safety region boundary [18].

The event-driven perspective reframes this optimization as a dual-mode problem. When the vehicle state is far from the safety boundary, the allocator can select the most energy-efficient distribution of longitudinal forces among wheels, possibly favoring motors with higher efficiency under the current torque and speed conditions. When the state approaches the boundary, the allocator must prioritize yaw moment generation, and the energy optimality criterion becomes secondary. The event scheduler determines when the transition occurs. Poorly designed event thresholds can cause frequent oscillation between modes, leading to energy losses through actuator transients and failing to achieve either strong energy savings or consistent safety margins. A well-designed scheduler incorporates hysteresis or dwell times to prevent rapid mode switching, and it may use predicted time-to-boundary rather than instantaneous distance as the triggering variable. This prediction reduces unnecessary interventions caused by brief, self-correcting disturbances.

One often overlooked aspect is the energy impact of the control computation itself. Although motor actuation dominates vehicle energy use, high-frequency optimization on embedded processors also consumes power and generates heat. Event-driven strategies can reduce the computational duty cycle of the stability supervisor, allowing lower average processor utilization and potentially enabling the use of less power-hungry control units. This is particularly relevant for distributed electric vehicles, where multiple electronic control units are distributed around the vehicle and communicate over shared networks. Reducing the number of high-priority messages and optimization invocations lowers the energy footprint of the embedded computing platform and reduces the thermal load in the control modules. These effects are smaller than propulsion energy but nontrivial in fleet-scale or long-duration operation contexts.

The robustness of energy-efficient stability management under parameter uncertainty is also a systems concern. Road friction is difficult to estimate accurately in real time, and friction estimators may respond slowly to sudden surface changes such as encountering ice or gravel. If the friction estimate is optimistic, the safety region expands beyond the true safe set, and the event scheduler may delay intervention. If the friction estimate is conservative, the scheduler triggers stability control too often, sacrificing energy efficiency and causing unnecessary torque modulation. Robust and gain-scheduled approaches attempt to address this by introducing formal robustness margins into the control design [18]. In distributed architectures, the uncertainty can be mitigated by combining friction estimates with wheel-level observations of slip and torque response, thereby reducing the reliance on a single global

friction estimate. Nevertheless, the certification of such fused and event-gated systems remains an active area of research.

6. System Integration, Communication, and Real-Time Constraints

The deployment of event-driven stability control in a production distributed electric vehicle requires integration with in-vehicle communication networks, real-time operating systems, and functional safety processes. Time-triggered architectures dominate current automotive control networks because they provide deterministic latency and simplified certification. Event-triggered control introduces asynchronous message flows that may conflict with fixed-priority scheduling assumptions. To preserve safety, event-driven stability messages must be assigned appropriate priority and bounded latency. The system designer must verify that the worst-case delay between the detection of a boundary-crossing event and the application of corrective torque does not exceed the allowable response time for the vehicle state. This worst-case analysis is complicated by the fact that the event itself depends on state estimates that may be delayed or corrupted. In this context, the interplay between event-triggered control and networked control theory is central, because stability preservation under communication constraints has been formally studied but often with simplified vehicle models [9]. Transferring these guarantees to high-dimensional, nonlinear vehicle dynamics remains challenging.

The use of cross-domain data analytics and learning-based inference adds further integration complexity. While deep transfer learning has shown strong performance in other large-scale cyber-physical systems by leveraging cross-domain data [5], its application to safety-critical vehicle control is limited by the need for explainability and guaranteed behavior. Nevertheless, learned models can serve an advisory role, for example by predicting upcoming road conditions or estimating the likelihood of future stability events based on fleet data. These predictions can pre-tune the event scheduler and the safety region parameters without directly commanding low-level actuators. Such a hybrid architecture separates safety-critical feedback from data-driven adaptation, allowing the latter to influence scheduling policies while preserving deterministic safety enforcement. This separation also aligns with recent trends in automotive safety standards, which distinguish between nominal performance functions and safety mechanisms.

Fault tolerance is another system-level requirement. Distributed electric vehicles expose multiple actuators and local controllers, which can be exploited for redundancy. If one motor inverter fails or one wheel speed sensor degrades, the stability manager must reconfigure the remaining actuators to preserve yaw moment capability. Event-driven architectures can support this by treating fault detection as a triggering event that forces immediate recomputation of the safety region and allocation policy. The event conditions must therefore include not only state proximity but also actuator health indicators, communication timeouts, and sensor plausibility checks. The resulting event scheduler is a multimodal monitor that integrates physical state, computation timing, and subsystem health. This integration is necessary to meet the safety requirements of highly distributed platforms.

7. Robustness, Fairness, and Actuator Utilization

Robustness in distributed electric vehicle stability control extends beyond classical disturbance rejection. The system must tolerate uneven tire wear, motor thermal derating, battery voltage sag, and asymmetrical road adhesion. Set-theoretic methods provide a foundation for verifying that a given control policy keeps the vehicle inside an invariant safe

set under bounded disturbances [13]. However, the computational burden of high-dimensional set computations often limits their direct use online. Event-driven strategies can reduce this burden by precomputing safe sets for a discrete set of operating modes and using online interpolation or conservative approximations between modes. The accuracy of these approximations directly affects the trade-off between safety and energy efficiency. Overly conservative approximations shrink the safety region, causing more frequent interventions and higher energy consumption. Under-approximations may violate safety guarantees. The engineering challenge is to find decompositions of the vehicle dynamics that admit efficient set computations without excessive conservatism.

Fairness in actuator utilization is a distinct dimension that emerges in multi-actuator distributed systems. A stability manager that repeatedly uses the same two wheels for yaw moment generation may cause uneven thermal loading, tire wear, and motor degradation. Over time, this uneven utilization can increase energy consumption because degraded components operate less efficiently. Fairness-aware allocation policies can distribute corrective effort across actuators according to their thermal margins, wear states, and efficiency maps, while still meeting the instantaneous yaw moment requirement. The event-driven scheduler can enforce fairness by maintaining a utilization history and including fairness penalties in the allocation objective when the vehicle is in a noncritical state. When safety becomes critical, fairness is necessarily subordinate to the requirement to generate the necessary corrective moment. This hierarchical prioritization mirrors broader debates in resource-constrained infrastructure systems, where fairness and efficiency are balanced differently under normal and emergency operating conditions.

The interplay between fairness and robustness also raises questions about certification. A fair allocation policy may be more complex and time-varying than a fixed allocation map, making it harder to verify worst-case behavior. Conversely, a purely efficiency-oriented allocation may produce predictable but uneven wear, creating long-term robustness risks that are not captured by short-horizon stability analysis. A system-level governance framework is therefore needed to manage the trade-off between short-term energy efficiency, long-term actuator health, and immediate safety. Such a framework must specify not only control algorithms but also monitoring and maintenance policies that operate on longer time scales.

8. Governance, Certification, and Policy Implications

The introduction of event-driven safety-critical control in distributed electric vehicles has implications for vehicle certification and regulatory governance. Current functional safety standards for road vehicles are built around assumptions of periodic task execution and deterministic communication. Event-driven control introduces variable task activation that must be justified through new kinds of timing analyses and safety cases. Regulators and standards bodies are increasingly interested in the safety of automated and software-intensive vehicle functions, but the certification of adaptive, event-gated control remains an open issue. A key governance challenge is to establish acceptable evidence for proving that an event scheduler does not systematically delay safety interventions under rare but critical conditions. This requires not only simulation and testing but also formal methods for timing verification and fault injection campaigns.

From a policy perspective, the energy savings enabled by event-driven stability management can contribute to fleet-level efficiency targets. The adoption of distributed electric vehicles is closely tied to lifecycle emissions and energy consumption. Stability systems that reduce unnecessary actuation and regenerative losses can improve real-world energy economy, but

these benefits must be demonstrated across diverse driving conditions and over long operational lifetimes. Policymakers may need to consider incentives for vehicles that incorporate energy-aware safety control, provided that safety is not compromised. The verification challenge is significant, because energy efficiency improvements must not be achieved by relaxing stability margins in ways that increase crash risk. A robust regulatory framework should require independent validation of safety region calibration and event threshold settings across representative driving scenarios.

Software governance also matters. Event-driven systems often rely on multiple interacting software components, including state estimators, friction estimators, event schedulers, control allocators, and low-level motor controllers. The integration of these components must be managed through configuration and version control processes that preserve traceability between field data, simulation results, and deployed parameters. In fleets of connected vehicles, over-the-air updates may adjust event schedules or safety region calibration based on aggregated data. While such updates can improve performance, they also introduce new risks if the update process is not governed by the same rigor as the original certification. Policy frameworks for software-defined vehicles must therefore address the lifecycle management of safety-critical parameters, including post-deployment updates and their interaction with event-driven control.

9. Forward-Looking Perspectives

The future of event-driven safety-critical control for distributed electric vehicles will likely involve tighter integration with predictive perception and cooperative infrastructure. If vehicles can receive reliable short-term forecasts of road surface conditions from connected infrastructure or preceding vehicles, the event scheduler can pre-activate stability resources before entering a low-friction zone. This anticipatory mode could reduce the need for reactive high-torque corrections and improve both safety and energy efficiency. At the same time, the reliance on external data introduces new attack surfaces and data trust issues. The safety region boundaries would become dependent on the credibility of incoming information, requiring robust fusion and fallback mechanisms. Cross-domain learning methods developed for large-scale traffic prediction can inform the design of such cooperative perception pipelines [5], but their translation to safety-critical control must be conducted with caution.

Another forward-looking direction is the use of high-fidelity digital twins for the certification and continuous validation of event-driven control. A digital twin of the vehicle dynamics and control stack can be used to stress-test event conditions under rare scenarios that are difficult to reproduce on proving grounds. By combining field data with simulation-based counterfactual analysis, manufacturers can build stronger safety cases for event-driven behavior. However, digital twin fidelity and simulation-to-real transfer remain nontrivial barriers. The governance of such validation pipelines must ensure that the digital twin is not used to justify unsafe parameter choices through selective evidence.

Event-driven stability management may also evolve toward a more formal hierarchical control theory in which safety, energy, and actuator health are treated as interacting constraints at different time scales. In such a formulation, a slow outer loop updates energy and health policies, a medium loop manages stability event scheduling, and fast inner loops enforce local torque limits. The integration of these loops requires not only new algorithms but also new architectural patterns for in-vehicle software, communication scheduling, and processor allocation. These patterns will shape the next generation of distributed electric

vehicle control platforms and will require interdisciplinary collaboration across control theory, embedded systems, software engineering, and regulatory science.

10. Conclusion

This paper has examined event-driven safety-critical control for distributed electric vehicles from a system-level perspective, with particular attention to the management of lateral stability and energy efficiency. The analysis has shown that event-driven approaches can reduce unnecessary actuation and computational load while preserving safety margins, provided that safety regions and event conditions are designed with care. The architectural trade-offs involve the granularity of event triggers, the interaction between local and global monitoring, and the certification of asynchronous real-time behavior. In distributed electric vehicles, the coupling between longitudinal torque allocation and lateral tire force capacity makes energy-efficient stability management a cross-domain problem rather than a simple add-on to conventional stability control.

The discussion has further argued that energy efficiency, robustness, fairness in actuator utilization, and regulatory compliance must be treated as co-designed properties of the control architecture. Event-driven scheduling is not merely a software optimization; it is a governance mechanism that determines when the system shifts from efficiency-oriented operation to safety-critical intervention. This shift has implications for vehicle certification, software lifecycle management, and policy frameworks. The use of safety regions as dynamic envelopes rather than fixed thresholds offers a principled basis for balancing these competing objectives, but it also requires rigorous validation.

Looking forward, the integration of predictive perception, cross-domain data analytics, and digital twin validation may enable more anticipatory and less energy-intensive stability management. However, these advances must be accompanied by formal methods, robust safety cases, and clear regulatory guidance. The challenge is not only to design better event conditions but to build the institutional and technical infrastructure necessary to trust adaptive safety-critical control in distributed electric vehicles across their operational lifetimes.

References

1. Ehsani, M., Gao, Y., Longo, S., & Ebrahimi, K. (2018). *Modern electric, hybrid electric, and fuel cell vehicles* (3rd ed.). CRC Press.
2. van Zanten, A. T. (2000). Bosch ESP systems: 5 years of experience. SAE Technical Paper 2000-01-1633. <https://doi.org/10.4271/2000-01-1633>
3. Hori, Y. (2004). Future vehicle driven by electricity and control—Research on four-wheel-motored "UOT Electric March II." *IEEE Transactions on Industrial Electronics*, 51(5), 954–962.
4. Chen, Y., & Wang, J. (2014). Adaptive energy-efficient control allocation for planar motion control of over-actuated electric ground vehicles. *IEEE Transactions on Control Systems Technology*, 22(4), 1362–1373.
5. 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.

6. Erlien, S. M., Fujita, S., & Gerdes, J. C. (2016). Shared steering control using safe envelopes for obstacle avoidance and vehicle stability. *IEEE Transactions on Intelligent Transportation Systems*, 17(2), 441–451.
7. De Novellis, L., Sorniotti, A., & Gruber, P. (2014). Wheel torque distribution criteria for electric vehicles with torque-vectoring differentials. *IEEE Transactions on Vehicular Technology*, 63(4), 1593–1602.
8. Tabuada, P. (2007). Event-triggered real-time scheduling of stabilizing control tasks. *IEEE Transactions on Automatic Control*, 52(9), 1680–1685.
9. Heemels, W. P. M. H., Johansson, K. H., & Tabuada, P. (2012). An introduction to event-triggered and self-triggered control. *Proceedings of the IEEE*, 100(12), 3130–3140.
10. Åström, K. J., & Bernhardsson, B. (2002). Comparison of periodic and event based sampling for first-order stochastic systems. *Proceedings of the 41st IEEE Conference on Decision and Control*, 2, 1640–1645.
11. Cai, G., Yin, G., Pi, D., Zhuang, W., Feng, J., Ren, Y., & Ding, H. (2025). Safety region based event-driven lateral stability control for DDEVs with energy conservation. *IEEE Transactions on Transportation Electrification*.
12. Hrovat, D., Di Cairano, S., Tseng, H. E., & Borrelli, F. (2012). The hybrid model predictive control approach to vehicle stability control. *IEEE Control Systems Magazine*, 32(4), 37–54.
13. Blanchini, F. (1999). Set invariance in control. *Automatica*, 35(11), 1747–1767.
14. Chen, Y., & Wang, J. (2013). Fast and global optimal energy-efficient control allocation with applications to over-actuated electric ground vehicles. *IEEE Transactions on Control Systems Technology*, 21(1), 267–275.
15. Falcone, P., Borrelli, F., Asgari, J., Tseng, H. E., & Hrovat, D. (2007). Predictive active steering control for autonomous vehicle systems. *IEEE Transactions on Control Systems Technology*, 15(3), 566–580.
16. Paden, B., Čáp, M., Yong, S. Z., Yershov, D., & Frazzoli, E. (2016). A survey of motion planning and control techniques for self-driving urban vehicles. *IEEE Transactions on Intelligent Vehicles*, 1(1), 33–55.
17. Yin, G., Jin, X., & Chen, N. (2016). Dynamic stability control of four-wheel independent drive electric vehicles based on integrated yaw moment and lateral force control. *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering*, 230(5), 643–655.
18. Zhang, H., & Wang, J. (2016). Vehicle lateral dynamics control through AFS/DYC and robust gain-scheduling approach. *IEEE Transactions on Vehicular Technology*, 65(1), 489–494.
19. Rawlings, J. B., Mayne, D. Q., & Diehl, M. (2017). *Model predictive control: Theory, design, and applications* (2nd ed.). Nob Hill Publishing.
20. Lian, Y., Zhao, Y., Hu, X., & Zhang, H. (2018). Lateral stability control of distributed drive electric vehicle based on model predictive controller. *Proceedings of the 2018 IEEE Intelligent Vehicles Symposium*, 1672–1677.