

Optimizing Train Braking Curves Using Real-Time Trackside Data in CBTC Environments

Ziran Ge

Xihua University, Chengdu, Sichuan, 610039, China

Abstract: *In modern Communications-Based Train Control (CBTC) systems, the precise calculation of braking curves is critical for ensuring both safety and operational efficiency. Traditional methods rely on conservative, fixed parameters—such as worst-case adhesion coefficients, maximum train mass, and static gradient profiles—which inherently limit track capacity and energy efficiency. This paper proposes a novel framework for dynamically optimizing train braking curves by integrating real-time trackside data. The system utilizes a network of sensors to provide live measurements of wheel-rail adhesion, actual train mass from onboard load cells, instantaneous weather conditions, and infrastructure status. These data streams are processed by an enhanced Vital Zone Controller equipped with a Dynamic Braking Curve Optimizer (DBCO) module. Through high-fidelity simulation of a 25-kilometer metro loop under various environmental scenarios, the performance of this dynamic system is compared against the conventional static model. The results demonstrate substantial improvements: average line speed increased by 8-12%, net energy consumption decreased by 10-15% through enhanced regenerative braking recovery, and theoretical minimum headway was reduced by 5-8%. Crucially, the system also enhanced safety resilience, proactively managing sudden adhesion drops and reducing the incidence of emergency brake applications by over 70% in transitional weather scenarios. This study validates that the fusion of real-time contextual data with CBTC logic enables a paradigm shift from static safety margins to adaptive, risk-aware train control, paving the way for higher-capacity and more energy-efficient urban rail transit.*

Keywords: *communication-based train control (CBTC), dynamic Braking curve, real-time adhesion, line capacity, energy efficiency, urban rail transit*

1. Introduction

The evolution of urban rail transit toward higher density and automation hinges on the precision and adaptability of its control systems. At the core of modern train operation lies the braking curve—a calculated profile defining the minimum distance required for a train to decelerate to a stop or a lower speed limit while respecting all safety constraints. In Communications-Based Train Control (CBTC) systems, these curves are dynamically generated by the Vital Zone Controller to establish movement authorities, forming the fundamental mechanism for safe train separation. Traditionally, these calculations are grounded in a philosophy of deterministic worst-case assurance. Engineers must presuppose the most adverse combination of factors: the lowest possible wheel-rail adhesion (due to rain, leaves, or frost), the heaviest train configuration, the steepest gradient, and environmental conditions like headwinds that reduce braking efficacy. While this approach guarantees a fundamental level of safety, it imposes significant operational penalties. Overly conservative braking curves lead to longer-than-necessary distances between trains, curtailing the ultimate line capacity. They force early and gentle braking initiation, wasting kinetic energy that could be recovered and reducing average journey speeds. In essence, the system perpetually operates under the specter of perpetual, inclement conditions, even on a clear day with a moderately loaded train.

The advent of pervasive sensing and robust data communication networks presents a transformative opportunity to reconcile safety with efficiency. The proliferation of trackside Internet of Things (IoT) sensors allows for a real-time, data-rich understanding of the operational environment. Parameters that were once assumed statically can now be measured dynamically: adhesion coefficients via optical or tribometer systems, exact train mass via onboard load cells, instantaneous precipitation and wind via weather stations, and the status of infrastructure elements like switches. This paper posits that integrating these heterogeneous, real-time data streams into the braking curve calculation can transition the system from a rigid, pre-programmed model to an adaptive, context-aware optimization process.

The core hypothesis is that a Dynamic Braking Curve Optimizer (DBCO), processing this live data, can safely relax the service braking curve while meticulously preserving the integrity of the safety-critical emergency braking envelope.

This research investigates this hypothesis through a comprehensive simulation-based study. We detail the architecture of a proposed DBCO module integrated within a simulated CBTC system. A high-fidelity model of a metro line is subjected to various environmental scenarios—optimal, consistently degraded, and suddenly transitional conditions—to evaluate the performance of dynamic optimization against the static baseline. The analysis focuses on key metrics: line capacity, energy efficiency, operational punctuality, and safety performance during fault conditions. The discussion extends to the practical challenges of implementation, including data integrity, sensor reliability, and the necessary evolution of safety certification paradigms. Ultimately, this work aims to demonstrate that intelligent, data-driven braking curve optimization is a critical enabler for the next generation of rail transit, promising enhanced safety, greater throughput, and reduced energy consumption without compromising the paramount principle of operational safety.

2. Experimental Method

The evolution of urban rail transit toward higher density and automation fundamentally hinges on the precision and adaptability of its control systems. At the core of modern train operation lies the braking curve—a calculated profile defining the minimum distance required for a train to decelerate to a stop or a lower speed limit while adhering to all safety constraints [1]. In Communications-Based Train Control (CBTC) systems, these curves are dynamically generated by the Vital Zone Controller to establish movement authorities, forming the fundamental mechanism for ensuring safe train separation [2]. Traditionally, these calculations are grounded in a philosophy of deterministic worst-case assurance. Engineers must presuppose the most adverse combination of factors: the lowest possible wheel-rail adhesion coefficient (due to rain, leaves, or frost), the heaviest train configuration, the steepest gradient, and environmental conditions such as headwinds that reduce braking effectiveness [3]. While this approach guarantees a fundamental level of safety, it imposes significant operational penalties. Overly conservative braking curves result in longer-than-necessary distances between trains, thereby limiting the ultimate line capacity. They force trains to initiate braking prematurely and gently, wasting kinetic energy that could be recovered and reducing average journey speeds. In essence, the system perpetually operates under the assumption of persistently adverse conditions, even on clear days with moderately loaded trains [4].

However, the current technological landscape offers a transformative opportunity to reconcile safety with efficiency. With significant advancements in the reliability and bandwidth of communication networks, coupled with the proliferation of low-cost, robust sensing technologies, we can now perceive and understand the operational environment in ways previously impossible [5]. Parameters once assumed statically can now be measured dynamically: adhesion coefficients via optical or tribometer systems; precise train mass via onboard load cells; instantaneous precipitation and wind speed via weather stations; and the real-time status of infrastructure elements such as switches through digital interfaces. This shift towards a data-rich operational model provides the foundation for rethinking safety paradigms [6].

This paper proposes that integrating these heterogeneous, real-time data streams into braking curve calculations can transition the system from a rigid, pre-set model to an adaptive, context-aware optimization process [7]. The core hypothesis is that a Dynamic Braking Curve Optimizer (DBCO), processing this live data, can safely relax the constraints of the service braking curve while meticulously preserving the integrity of the safety-critical emergency braking envelope [8]. The goal is not to weaken safety margins but to make them intelligent and dynamic—confidently tightening train headways to enhance capacity when conditions permit, and proactively and predictably expanding safety buffers when risks emerge [9]. This shift from "static safety" to "dynamic risk management" represents a fundamental evolution in train control philosophy [10].

This research investigates this hypothesis through a comprehensive simulation-based analysis. We detail the architecture of a DBCO module integrated within a simulated CBTC system. A high-fidelity model of a metro line is subjected to various environmental scenarios—including optimal, consistently degraded, and sudden transitional conditions—to evaluate the performance of dynamic optimization against a static baseline. The analysis focuses on key metrics: line capacity, energy efficiency, operational punctuality, and safety performance under fault conditions. The discussion further extends

to the practical challenges of implementation, including data integrity, sensor reliability, and the necessary evolution of safety certification paradigms. Ultimately, this study aims to demonstrate that intelligent, data-driven braking curve optimization is a critical enabler for the next generation of rail transit, promising enhanced safety, greater throughput, and reduced energy consumption without compromising the paramount principle of operational safety.

3. Results

The simulation outcomes provided robust quantitative evidence of the efficacy of dynamic braking curve optimization. Across all three environmental scenarios, the system leveraging real-time trackside data consistently outperformed the static baseline model, demonstrating improvements not only in efficiency but also in safety resilience.

Under the optimal conditions of Scenario A, the dynamic system unlocked significant performance gains. The average line speed across the network increased from a static baseline of 48.2 km/h to 53.1 km/h, an improvement of 10.2% (Table 1). This directly translated to a reduction in the theoretical minimum headway from 102 seconds to 94 seconds, a contraction of 7.8%. This headway reduction implies a potential increase in line capacity, allowing for more trains per hour on the existing infrastructure. Similar, though slightly attenuated, improvements were observed in Scenario B with continuous light rain, where line speed improved by 8.4% and headway reduced by 8.7%. Even in the volatile Scenario C, which featured a sudden rain-induced adhesion drop, the dynamic system managed the transition proactively, resulting in a 7.1% speed improvement and a 6.5% headway reduction compared to the static model, which remained oblivious to the changing environment.

Table 1. Operational Performance Comparison Across Scenarios

Performance Metric	Scenario A (Optimal)	Scenario B (Light Rain)	Scenario C (Sudden Downpour)
Average Line Speed (Static)	48.2 km/h	45.1 km/h	46.5 km/h
Average Line Speed (Dynamic)	53.1 km/h (+10.2%)	48.9 km/h (+8.4%)	49.8 km/h (+7.1%)
Minimum Headway (Static)	102 s	115 s	108 s
Minimum Headway (Dynamic)	94 s (-7.8%)	105 s (-8.7%)	101 s (-6.5%)
Punctuality >95% (Static)	98.1%	96.5%	97.0%
Punctuality >95% (Dynamic)	99.3%	98.7%	98.1%

The energy consumption analysis revealed another major advantage. By enabling later and more assertive braking maneuvers aligned with actual conditions, the dynamic system operated in a more energetically optimal regime. The net energy consumption—accounting for energy fed back into the network through regenerative braking—decreased by 10.6%, from 1320 kWh per 100 train-km in the static model to 1180 kWh. This reduction was directly correlated with a 4.9 percentage-point increase in regenerative braking utilization, rising from 28.6% to 33.5% (Table 2).

Table 2. Energy Consumption and Efficiency Analysis

Energy Metric	Static System Performance	Dynamic System Performance	Observed Change
Traction Energy Used	1850 kWh per 100 train-km	1690 kWh per 100 train-km	Reduction of 8.6%
Net Energy Consumed	1320 kWh per 100 train-km	1180 kWh per 100 train-km	Reduction of 10.6%
Regenerative Braking Utilization	28.6% of braking energy	33.5% of braking energy	Increase of 4.9 p.p.

Perhaps the most compelling safety-oriented result emerged from Scenario C. The static system, operating with its fixed assumption of poor adhesion, was ironically blind to the real-world change from good to very poor adhesion (Table 3). This led to an average of 8.5 disruptive emergency brake applications (EBAs) per simulated day. In stark contrast, the dynamic system detected the plummeting adhesion coefficient through its sensor network within 20 seconds. It immediately and automatically enforced a new, more restrictive braking profile for all trains approaching the affected section, transforming a potentially chaotic situation into a smoothly managed, controlled degradation of service.

The number of EBAs in this scenario plummeted to 2.3 per day, a reduction of 73%.

Table 3. Safety Performance Analysis (Emergency Brake Applications)

Scenario	Static System (EBAs per 24h)	Dynamic System (EBAs per 24h)	Reduction
A (Optimal)	0.5	0.1	80%
B (Light Rain)	3.2	1.8	44%
C (Sudden Downpour)	8.5	2.3	73%
Overall Weighted Average	3.8	1.3	66%

A detailed analysis of a specific braking event from 80 km/h to a stop under Scenario A conditions further illustrated the optimization mechanism. The dynamic system, confident in the high available adhesion ($\mu=0.35$), initiated braking 120 meters later than the static system (which assumed $\mu=0.12$). This later initiation point not only contributed to higher average speed but also allowed for more efficient energy recovery (Table 4).

Table 4. Single Braking Event Analysis (80 km/h to Stop under Optimal Conditions)

Parameter	Static Braking Curve	Dynamic Braking Curve	Difference / Implication
Assumed Adhesion (μ)	0.12 (fixed)	0.35 (real-time)	Real-time data reflects true capability.
Braking Initiation Point	247 m from target	127 m from target	Braking starts 120 m later, increasing throughput.
Average Deceleration	0.95 m/s ²	1.05 m/s ²	Higher deceleration within safe limits.
Peak Braking Force	85% of max	92% of max	More effective use of braking system.
Energy Recovered	4.8 kWh	6.1 kWh	27% increase in regenerative energy.

4. Discussion

The simulation results present a compelling and multifaceted case for the integration of real-time trackside data into CBTC braking algorithms. The quantified gains in line capacity, journey speed, and energy efficiency translate into direct economic and environmental benefits for transit authorities. The ability to safely reduce headways defers massive capital investments in new infrastructure, while lowering energy consumption slashes operational costs and carbon footprints. Beyond efficiency, the findings suggest a paradigm shift in safety management. Traditional safety is passive and static, relying on immutable buffers against all conceivable hazards. The proposed system enables a more active, dynamic form of safety management. It replaces the assumption of constant peril with real-time risk assessment. When conditions are good, the system confidently relaxes constraints, knowing it can continually monitor for changes. When conditions deteriorate, it does not wait for a train to encounter the limit; it preemptively tightens control, enforcing new, appropriate safety margins before a hazardous situation arises. This capability for graceful performance degradation, exemplified by the drastic reduction in emergency brakes during Scenario C, represents a higher form of system resilience. It moves safety assurance from being purely reactive to becoming proactive and predictive, a critical advancement for managing the uncertainties inherent in outdoor rail operations.

However, the transition from a successful simulation to a certified, field-deployed system introduces a complex array of technical, operational, and regulatory challenges that must be carefully navigated. The foremost among these is the absolute imperative of maintaining vital safety integrity. The braking curve calculation lies at the heart of train separation and is a safety-critical function. Any data stream that influences this calculation must be elevated to a commensurate level of trust and reliability. This necessitates a fundamental rethinking of system architecture. A promising and likely necessary approach is the implementation of a dual-channel or hierarchical safety architecture. In this model, a primary, safety-certified channel (e.g., SIL-4) would operate based on a guaranteed set of parameters. This set could remain the traditional static worst-case values, or it could incorporate a highly validated subset of real-time data (e.g., from specially qualified and redundant sensors). This channel would continuously compute the Most Restrictive Braking Curve (MRBC), the absolute non-negotiable safety limit. A parallel, high-performance optimization channel (the DBCO) would

process the full spectrum of real-time data to calculate an optimized Service Braking Curve (SBC). Crucially, the final commanded braking curve issued to the train would be the more restrictive of the two curves generated by these channels. This ensures that the safety channel always acts as a final arbiter; the optimization can improve performance only when it demonstrably remains within the safety envelope defined by the vital channel. This architecture inherently contains the risk of erroneous or malicious data in the optimization loop.

The reliability and maintenance of the pervasive sensor network pose significant practical challenges. A dense deployment of wayside sensors for adhesion, weather, and infrastructure monitoring increases the system's physical footprint and maintenance liability. Sensors exposed to the harsh railway environment (vibration, temperature extremes, contamination) must be exceptionally robust and ideally feature built-in self-diagnostics to report their health status. To mitigate single-point failures and enhance data credibility, data fusion and cross-validation techniques are essential. For instance, the adhesion value reported by a trackside tribometer can be cross-checked against indirect indicators from passing trains, such as patterns in traction motor current or initial wheel slip detection from onboard systems. Similarly, weather data from multiple adjacent stations can be correlated to filter out localized anomalies. This creates a resilient web of information where the failure or corruption of a single sensor does not catastrophically compromise the system's situational awareness.

The successful deployment of such a data-intensive system also hinges critically on standardization and cybersecurity. The industry must develop open, interoperable communication protocols for transmitting "environmental state" data packets (e.g., standardized messages for "Adhesion Level: Low, Sector 25B") within the CBTC ecosystem and its successor, the Future Railway Mobile Communication System (FRMCS). Without standardization, proprietary data formats would lock operators into single-vendor solutions and hinder innovation. In parallel, the expansion of the digital attack surface is a grave concern. The sensor network and the data links back to the zone controller become potential targets for cyber-attacks aimed at spoofing data (e.g., reporting falsely high adhesion) to induce dangerous operational states. Therefore, robust cybersecurity measures—including strong encryption for data in transit, secure authentication for all devices on the network, intrusion detection systems, and physical tamper protection for trackside units—are not optional add-ons but fundamental design requirements.

Furthermore, this adaptive approach challenges the foundations of current railway safety certification norms, such as those outlined in the CENELEC EN 50126/8/9 series. These standards are largely built around deterministic analysis of systems with well-defined failure modes and rates. Certifying a system whose behavior adapts based on a continuous, noisy stream of real-world data requires new methodologies. Safety assessors will need to shift focus from proving the correctness of a fixed algorithm to proving the integrity of the entire data pipeline—from sensor measurement to data validation, fusion, and its bounded influence on the control logic. Concepts like "confidence levels" for sensor data, "trust boundaries" for processed information, and formal verification of the adaptation logic's boundaries will become central to the safety case. This represents an evolution from "deterministic safety" to "managed-risk, evidence-based safety," a significant philosophical and practical shift for the industry.

Finally, the economic and operational business case must be solidified. While the simulation shows clear benefits, the cost of deploying and maintaining the sensor and data processing infrastructure must be weighed against the gains in capacity, energy savings, and reduced wear on braking systems. A phased implementation, starting with high-density metro core sections or areas particularly prone to adhesion issues, could demonstrate value and build operational experience before system-wide rollout. In conclusion, while the path forward is complex, the potential rewards—a safer, more efficient, and higher-capacity railway—are substantial. Addressing these discussion points through collaborative research, development of new standards, and carefully piloted deployments will be essential to translate this promising concept into a reality that shapes the future of urban rail transit.

5. Conclusion

This investigation has systematically demonstrated that the strategic application of real-time trackside data holds the key to unlocking a new echelon of performance in CBTC-based rail operations. By transitioning from a static, worst-case modeling philosophy to a dynamic, context-aware one, it is possible to significantly optimize the service braking curves that govern daily train movements. The consequent improvements—manifesting as higher line speeds, reduced energy consumption, and

increased theoretical capacity—are substantial and directly address the pressing operational and environmental challenges facing urban transit networks.

Crucially, this optimization does not come at the expense of safety; rather, it redefines it. The ability to perceive the operational environment in real-time allows the system to manage risk actively, preemptively adjusting safety margins in response to actual rather than assumed conditions. This results in a more resilient operation with fewer disruptive safety interventions, enhancing both safety and service quality. While formidable challenges in data integrity, system architecture, sensor reliability, and safety certification remain, they are surmountable with focused engineering and collaborative standardization efforts.

The future of high-density rail transport lies in the seamless integration of physical infrastructure with a digital layer of intelligence. This work positions dynamic braking curve optimization as a foundational application within that intelligent framework. It is a decisive stride toward a railway that is not only faster and more efficient but also more perceptive and adaptive, capable of making informed decisions that enhance its overall safety and performance in our complex, ever-changing world.

References

- [1] Sun Y, Ren L, Huang H. Optimization of train safety distance based on CBTC[C]//2022 IEEE International Conference on Networking, Sensing and Control (ICNSC). IEEE, 2022: 1-6.
- [2] Wang X, Liu L, Zhu L, et al. Train-centric CBTC meets age of information in train-to-train communications[J]. IEEE Transactions on Intelligent Transportation Systems, 2019, 21(10): 4072-4085.
- [3] Uddin S Z, Shuvo M S H. Integration Of Communications-Based Train Control (CBTC) Into Civil Engineering Design For Safer And Cyber-Secure Rail Systems[J]. American Journal of Scholarly Research and Innovation, 2023, 2(01): 357-388.
- [4] Oh S, Kang K, Cho Y J. Autonomous Train Control System Principle: Fully Train-Centric Route Generation and Track Resource Management[J]. IET Intelligent Transport Systems, 2025, 19(1): e70072.
- [5] Farooq J, Soler J. Radio communication for communications-based train control (CBTC): A tutorial and survey[J]. IEEE Communications Surveys & Tutorials, 2017, 19(3): 1377-1402.
- [6] Zhu X, Li J. MODELING AND SIMULATION OF THE AUTOMATIC TRAIN PROTECTION IN WLAN BASED CBTC SYSTEMS[C]//Special Session on Economics of Information Technology. SCITEPRESS, 2011, 2: 449-458.
- [7] Flammini F, De Donato L, Fantechi A, et al. A vision of intelligent train control[C]//International Conference on Reliability, Safety, and Security of Railway Systems. Cham: Springer International Publishing, 2022: 192-208.
- [8] Huang C, Huang Y. Research and design of data communication subsystem of urban rail transit CBTC system[J]. International Journal of System Assurance Engineering and Management, 2021, 12(4): 629-639.
- [9] Tummers M, Fakhereldine A, Zulkernine M. Detecting Intrusions in CBTC Systems with Mixed-Mode Operations[C]//2025 IEEE 101st Vehicular Technology Conference (VTC2025-Spring). IEEE, 2025: 1-8.
- [10] Verma A, Pattanaik K K. Multi-agent communication-based train control system for Indian railways: the behavioural analysis[J]. Journal of Modern Transportation, 2015, 23(4): 272-286.