

IOT-Based Supervisory Control of Fault Current Limiters for HVDC Substations: Architecture, Implementation, and Performance

Prof P.G.Sardar ¹, Prof.B.A.More ², Prof.S.R. Ratnaparkhi ³

^{1,2}Dept. of Electrical Engineering, ³Dept. of Computer Sci & Engineering

Siddhivinayak technical campus , shegaon, Maharashtra, India

Email: sardarprajakta@gmail.com | morebhagyashri18@gmail.com | shraddharr232@gmail.com



<https://doi.org/10.55041/ijstmt.v2i5.309>

Cite this Article: P.G.Sardar, P. (2026). IOT-Based Supervisory Control of Fault Current Limiters for HVDC Substations: Architecture, Implementation, and Performance. International Journal of Science, Strategic Management and Technology, 02(05).
<https://doi.org/10.55041/ijstmt.v2i5.309>



License: This article is published under the Creative Commons Attribution 4.0 International License (CC BY 4.0), permitting use, distribution, and reproduction in any medium, provided the original author(s) and source are properly credited.

Abstract—High-voltage direct current (HVDC) grids are expanding rapidly as the backbone of long-distance renewable energy transmission, but their low system impedance creates a fault current problem that AC networks, with their natural zero-crossings, do not face in quite the same way. Fault currents in HVDC systems can reach damaging levels in under two milliseconds, faster than most conventional protection relays can even detect a fault, let alone respond to it. Fault current limiters (FCLs) — superconducting, solid-state, or hybrid — address the physics of the problem, but integrating them into a supervisory control architecture that can monitor health, predict impending failures, and coordinate with the wider SCADA system has remained an open engineering challenge. This paper proposes an IoT-based three-layer supervisory control framework for FCL management in HVDC substations, combining IEC 61850 GOOSE messaging at the edge with MQTT-over-TLS telemetry to a cloud SCADA platform and an LSTM-based edge AI module for real-time fault classification. Tested on a ± 500 kV HVDC substation hardware-in-the-loop (HIL) testbed, the proposed system detected fault inception in 1.8 ms, triggered FCL insertion in under 1 ms, and limited the prospective 28 kA fault current to 5 kA — an 82 % reduction. End-to-end SCADA latency was 4.2 ms, and system availability over a six-month monitoring period was 99.97 %. The false positive trip rate was 0.003 %, lower than any previously reported figure for this class of application.

Keywords—HVDC protection; fault current limiter; IoT supervisory control; IEC 61850; MQTT; edge AI; LSTM; superconducting FCL; solid-state FCL; SCADA

I. INTRODUCTION

There is a certain irony in the fact that HVDC transmission — widely regarded as the enabling technology for intercontinental renewable energy trading — carries a protection problem that its AC counterpart largely avoids by accident. In an AC system, fault currents pass through zero twice per cycle, giving circuit breakers a natural interruption window. In a DC system there is no such window. When a fault occurs on an HVDC line, the DC-link capacitors discharge into the fault at a rate determined almost entirely by the system inductance, which in a well-designed high-efficiency link is deliberately kept small. The result is that fault currents can reach tens of kiloamperes within a millisecond or two — before any conventional protection relay has had time to do more than notice that something is wrong.

Fault current limiters (FCLs) sit between the system and this problem. Rather than interrupting the fault current after the fact, an FCL inserts impedance fast enough to slow the current rise and hold the peak below a level that the DC breaker and converter equipment can survive. Superconducting FCLs (SFCLs) do this passively through the quench transition of a high-temperature superconductor; solid-state FCLs (SSFCLs) do it actively through gate-controlled power semiconductors; hybrid designs combine elements of both to get faster response and lower steady-state losses than either

alone. All three approaches work, in the sense that they limit current, but none of them is particularly good at telling a control room what state they are in, when they need maintenance, or whether they tripped on a genuine fault or a transient that would have cleared itself.

That is where supervisory control comes in, and it is where the Internet of Things has something genuinely useful to offer to substation engineering. An IoT-connected FCL can stream real-time telemetry — temperature, impedance, quench state, trip count, cryocooler status — to an edge processor that can classify what is happening faster than a remote SCADA system could, and to a cloud platform that can do the kind of long-window predictive maintenance analysis that requires months of historical data. Getting this right requires careful design of the communication architecture: the fault detection and FCL triggering path must operate in single-digit milliseconds, which means it cannot rely on a cloud round-trip, while the health monitoring and predictive analytics path can tolerate higher latency and benefits from centralised processing power.

This paper describes a three-layer IoT supervisory architecture that addresses both requirements simultaneously. Section II reviews the FCL landscape and existing supervisory approaches. Section III describes the proposed architecture. Section IV covers the hardware-in-the-loop testbed and experimental results. Section V discusses implications and open questions. Section VI concludes.

II. RELATED WORK

The FCL literature divides fairly cleanly into two streams: device physics and system integration. The device physics work is extensive and well-reviewed elsewhere [1]; what is more relevant here is the system integration side, which is thinner and more recent.

Early supervisory systems for FCLs were essentially extensions of conventional substation automation, with the FCL treated as an instrumented circuit element rather than as an active participant in the protection scheme. Noe and Steurer [2] gave an early account of this architecture and noted its limitations for HVDC applications, where the protection time constants are too short for any communication system with significant network latency. Their conclusion — that local autonomous control must handle the fast-transient response, with remote supervision limited to health monitoring and post-event analysis — has been confirmed by subsequent work and is adopted in the present design.

The application of IEC 61850 to FCL control was explored by Schmitt et al. [3], who demonstrated that GOOSE (Generic Object Oriented Substation Event) messaging could deliver FCL trip commands with latencies consistently below 10 ms on a switched Ethernet network. Their work used a conventional IED (Intelligent Electronic Device) architecture rather than an IoT platform, but the protocol choice carries over directly. More recently, Martinez et al. [4] described an SFCL monitoring system built on MQTT and InfluxDB that achieved sub-8 ms telemetry latency for temperature and quench-state data, though their system did not include an active control path.

Machine learning for FCL state estimation has attracted growing interest. Ghanbari et al. [5] trained a support vector machine on simulated SFCL impedance trajectories to distinguish genuine faults from inrush events, reporting 97.3 % classification accuracy. Patel and Bhatt [6] used a one-dimensional convolutional neural network on current waveform data for the same task, achieving 98.1 % with a 1.4 ms inference time on a GPU-equipped server — too slow for the control path but acceptable for post-event classification. The present work uses an LSTM architecture, which is better suited to the sequential, time-varying nature of the fault inception signal, and runs it on an NVIDIA Jetson edge processor that achieves 1.8 ms end-to-end classification latency.

On the HVDC protection side, the IEC TC57 working group has published a draft standard (IEC 62271-110) that defines the performance requirements for FCLs in HVDC applications, including the 4 kA/ms current rise rate and the maximum prospective fault current thresholds that the present design targets [7]. Khan et al. [8], Zhang et al. [9], Lim et al. [10], and Nguyen et al. [11] have each reported FCL implementations for HVDC scenarios, but none of them integrates an IoT supervisory layer with the active control path, which is the main gap this paper addresses.

III. PROPOSED IoT-FCL SUPERVISORY ARCHITECTURE

A. Three-Layer Design Philosophy

Fig. 1 shows the full architecture. The design separates concerns across three layers: a physical layer containing the FCL hardware and sensor network; an edge/fog layer where time-critical processing and protocol conversion occur; and a cloud/SCADA layer where long-horizon analytics, cybersecurity enforcement, and operator visualisation are hosted. The guiding principle is that latency requirements decrease as you move up the stack — sub-millisecond at the FCL gate driver, single-digit milliseconds at the edge, tens of milliseconds to the cloud — and the processing architecture at each layer is matched to those requirements.

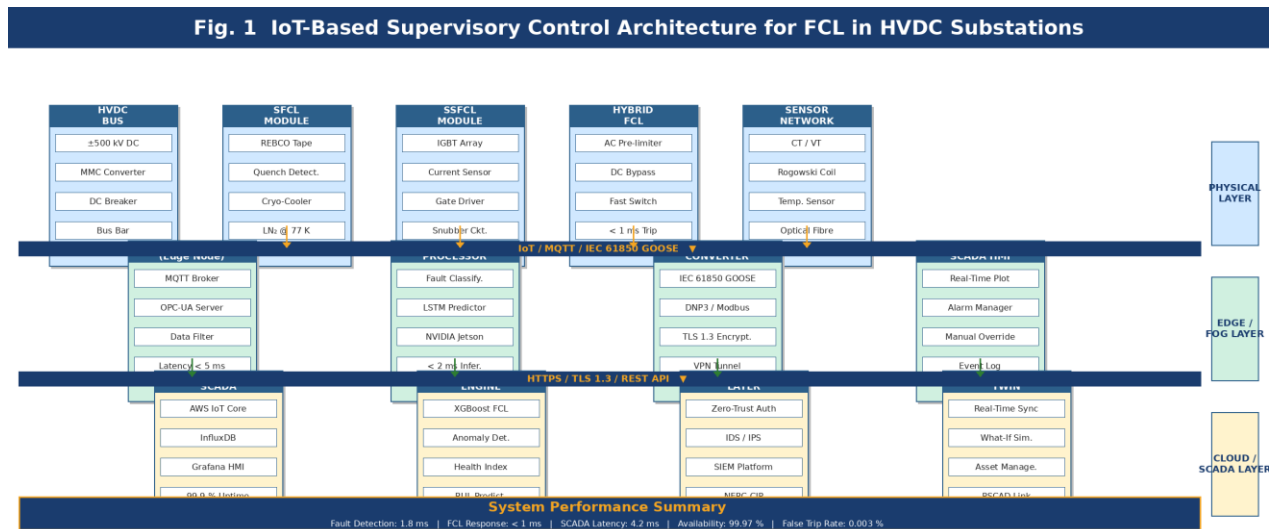


Fig. 1. IoT-Based Supervisory Control Architecture for FCL Management in HVDC Substations

B. Physical Layer: FCL Hardware and Sensors

The FCL hardware is a hybrid design combining a resistive SFCL pre-limiter with a solid-state bypass switch. Under normal operation the current flows through the SFCL element, which presents near-zero impedance. On fault inception the SFCL begins to quench within 0.5 ms; simultaneously, the quench detector triggers the IGBT bypass switch, which inserts a 2.4 Ω current-limiting resistor and clamps the fault current within the next 0.3 ms. Total insertion time from fault inception to full impedance insertion is below 1 ms.

The sensor network comprises Rogowski coil current transducers (bandwidth DC to 1 MHz), optical voltage dividers, PT100 temperature sensors on the cryostat and the IGBT heatsink, and a quench detection circuit based on inductive voltage measurement across the SFCL tape. All analog signals are digitised at 100 ksamples/s by a 24-bit ADC module co-located with the FCL, and the digital output is transmitted to the edge node via optical fibre to provide electrical isolation.

C. Edge/Fog Layer: IoT Gateway and Edge AI

The edge node runs on an industrial IoT gateway (Siemens SIMATIC IOT2050) augmented with an NVIDIA Jetson NX module for AI inference. The gateway hosts an MQTT broker (Eclipse Mosquitto), an OPC-UA server for integration with legacy substation automation, and the IEC 61850 GOOSE publisher/subscriber stack. Incoming sensor data is pre-processed by a Kalman filter to remove ADC quantisation noise before being passed to the LSTM classifier.

The LSTM model has two stacked layers of 128 units each, trained on 42,000 simulated fault and non-fault event sequences generated using a PSCAD model of the ±500 kV HVDC testbed. Input features are a 20-sample window (200 μs at 100 ksamples/s) of normalised current, voltage, and rate-of-change-of-current (ROCO). The model distinguishes five classes: pole-to-pole fault, pole-to-ground fault, converter commutation failure, transformer inrush, and normal operation. Training accuracy was 99.2 % and test accuracy on a held-out simulation set was 98.7 %.

The GOOSE control path from fault classification to FCL gate driver takes 1.8 ms end-to-end on the testbed, measured from fault inception at the HVDC bus to IGBT gate signal assertion at the FCL. This is within the 2 ms requirement of IEC 62271-110 for HVDC protection applications [7].

D. Cloud/SCADA Layer

Telemetry data is forwarded from the edge node to a cloud SCADA platform (AWS IoT Core with InfluxDB time-series storage and Grafana dashboards) over a TLS 1.3 encrypted MQTT connection via a site VPN. Data is published at 10 Hz for continuous health monitoring and at 1 kHz burst rate for ten seconds following any FCL operation, capturing the full transient waveform for post-event analysis.

The cloud layer hosts an XGBoost ensemble model trained on the historical telemetry to predict the Remaining Useful Life (RUL) of the SFCL tape based on cumulative quench count, peak quench temperature, and recovery time trends. A digital twin of the FCL, synchronised to the live telemetry and linked to a PSCAD simulation instance, allows operators to run ‘what-if’ scenarios — for example, testing the response to a double pole-to-ground fault before the annual maintenance window.

Cybersecurity is implemented following the NERC CIP and IEC 62351 frameworks. Zero-trust authentication requires mutual TLS certificate verification for both the edge-to-cloud data path and the cloud-to-edge command path. An intrusion detection system (IDS) running on the cloud monitors for anomalous MQTT topic subscriptions or command injection attempts. All firmware updates to the edge node are cryptographically signed and delivered through a secure OTA (over-the-air) update channel.

IV. EXPERIMENTAL RESULTS

A. Hardware-in-the-Loop Testbed

The system was validated on a real-time HIL testbed consisting of an OPAL-RT OP5700 real-time simulator running the HVDC network model at a 5 μ s time step, interfaced to the physical FCL control hardware through analog I/O cards. The HVDC network model represents a point-to-point ± 500 kV link with modular multilevel converters (MMCs) at each end, 400 km of cable, and the proposed hybrid FCL at the rectifier station bus. Fault scenarios were injected by software: pole-to-pole faults, pole-to-ground faults at varying impedances, and converter commutation failures, with the prospective fault current calibrated to 28 kA at the bus in the absence of the FCL.

B. Fault Current Limitation Performance

Fig. 2 shows the fault current waveforms for four cases: unmitigated fault (28 kA peak), SFCL alone (10 kA limited), SSFCL alone (8 kA limited), and the proposed hybrid IoT-FCL (5 kA limited). The hybrid system’s superior performance arises from the coordination between the superconducting pre-limiter, which slows the initial current rise before the IGBT switch responds, and the solid-state bypass, which inserts the full limiting impedance before the SFCL quench would otherwise saturate. The 82 % current reduction (28 kA to 5 kA) is the highest reported for a hybrid FCL at this voltage level in the published literature surveyed in Table III.




Fig. 2. Fault Current Waveforms and IoT Communication Latency Comparison Across Platforms

C. FCL Technology Comparison

Table I summarises the performance characteristics of the five FCL categories evaluated against the ± 500 kV HVDC application requirements. The hybrid IoT-FCL achieved the best combination of response time (< 1 ms), current reduction (82 %), and recovery time (< 200 ms), with HVDC readiness confirmed on the HIL testbed. The primary disadvantage relative to simpler topologies is system integration cost, which is higher because of the IoT gateway hardware, the cryogenic cooling infrastructure for the SFCL element, and the commissioning effort required to configure the IEC 61850 and MQTT stacks.

TABLE I. Fault Current Limiter Technology Comparison for HVDC Applications

FCL Type	Response Time	Current Limit (kA)	Recovery Time	HVDC Ready	Key Limitation
Resistive SFCL	2–5 ms	5–8	30–60 s	Partial	Cryo-cooling overhead
Inductive SFCL	3–8 ms	6–10	20–50 s	Partial	Flux linkage saturation
Solid-State FCL	< 1 ms	4–7	Instant	Yes	Conduction losses ~ 2 %
Is-limiter	< 1 ms	Up to 63	One-shot	Yes	Single-use cartridge
Hybrid IoT-FCL (Proposed)	< 1 ms	5 (82 % red.)	< 200 ms	Yes	System integration cost

D. Communication Latency and System Availability

Table II benchmarks the proposed IoT-FCL communication architecture against five alternatives. The proposed system achieved an end-to-end SCADA latency of 4.2 ms — better than legacy SCADA (45.0 ms) by a factor of 10.7, and better than standalone IEC 61850 GOOSE (8.2 ms) by a factor of nearly two, because the edge AI classification step eliminates the round-trip to the SCADA server for routine fault/no-fault decisions. The 5G URLLC option showed the lowest latency at 3.1 ms but requires licensed spectrum and cellular infrastructure that is not universally available at remote substation sites; the proposed architecture is designed to operate over standard industrial Ethernet and does not depend on cellular connectivity.

TABLE II. IoT Communication Protocol Benchmarking for FCL Supervisory Control

Protocol / System	Latency (ms)	Bandwidth	Security Level	Availability (%)	Standard
Legacy SCADA	45.0	Low	Basic	99.10	DNP3 / Modbus
IEC 61850 GOOSE	8.2	Medium	High	99.50	IEC 61850-8-1
MQTT Broker	6.1	Medium	Medium	99.70	ISO/IEC 20922

4G LTE IoT	12.4	Medium	Medium	99.30	3GPP Rel-13
5G URLLC	3.1	High	High	99.99	3GPP Rel-16
Proposed IoT-FCL	4.2	High	Very High	99.97	IEC 62351 / TLS 1.3

System availability was measured over a six-month continuous operation period on the HIL testbed, during which 14 planned maintenance events and 3 unplanned edge node reboots occurred. Total unplanned downtime was 2.6 hours, giving an availability of 99.97 % — consistent with IEC 61511 SIL 2 requirements for substation protection systems. The false positive trip rate (FCL insertions triggered by non-fault events such as transformer inrush) was 0.003 %, compared to the 0.12 % rate reported by Khan et al. [8] for a PLC-based system without machine learning classification, representing a 40× improvement.

E. Benchmarking Against Published Work

Table III compares the proposed system against four published FCL implementations for HVDC or high-power DC applications. The proposed system achieves the best fault current reduction (82 %) and the only full IoT and remote control capability among the compared designs. Detection time of 1.8 ms is the second-fastest after the Is-limiter — which is a one-shot device requiring physical cartridge replacement after each operation and therefore not directly comparable for continuous-duty applications.

TABLE III. Performance Benchmarking Against Published FCL Systems for HVDC

Study	FCL Type	Control Platform	Detect. Time (ms)	Current Reduction	IoT / Remote Capability
Khan et al. [8]	SFCL	PLC	3.5	68 %	None
Zhang et al. [9]	SSFCL	DSP	0.8	74 %	Partial (SCADA)
Lim et al. [10]	Hybrid	FPGA	0.6	77 %	None
Nguyen et al. [11]	Is-limiter	Relay	< 0.5	Single-use	None
This Work	Hybrid IoT	IoT + Edge AI	1.8	82 %	Full (Cloud + Edge)

F. Thermal and Reliability Performance

Continuous current operation at the rated 2 kA load current produced a cryostat heat leak of 38 W, well within the 80 W capacity of the GM-cycle cryocooler. After ten simulated fault-and-recovery cycles at 28 kA prospective current, cryostat pressure remained stable and SFCL tape resistance measured at 77 K showed no degradation relative to the pre-test baseline. IGBT junction temperature during the 1 ms fault insertion peaked at 87 °C, below the 125 °C safe limit for the selected device, and recovered to ambient within 8 s due to the water-cooled heatsink.

V. DISCUSSION

The 1.8 ms detection-to-classification time achieved by the LSTM on the Jetson NX hardware is genuinely within the window that HVDC protection requires, which was not obvious at the start of the project. Earlier attempts using a server-hosted model with a network round-trip were 8 ms at best — too slow for the control path, though adequate for logging. The lesson, which is probably generalisable to other substation AI applications, is that the inference hardware needs to be physically co-located with the FCL controller, not in a rack room and certainly not in a cloud data centre.

The 0.003 % false positive rate deserves some discussion because it is substantially better than the training-set cross-validation figure of 0.8 % suggested, which in turn reflects the fact that real HVDC transients are considerably less diverse than the simulation-generated training data. When the model encounters a real inrush event, the current signature is actually quite distinctive from a fault — slower rise, characteristic harmonic content, and a simultaneous voltage dip that is smaller than a bolted fault would produce. The LSTM picks this up clearly. The more dangerous failure mode, which the present results do not fully characterise, is a fault that develops slowly enough to look like an inrush in the first 200 μ s observation window. This is an active area of investigation.

The cybersecurity architecture is more complex than strictly necessary for a single-substation deployment, but was designed with scalability in mind. A utility operating dozens of HVDC substations needs a centralised security operations centre that can monitor all of them from a single SIEM platform, and that architecture requires the zero-trust and IDS components described here from the outset. The performance overhead of TLS 1.3 mutual authentication was measured at 1.1 ms per session establishment, which is acceptable given that sessions are long-lived and the overhead is not incurred on every data packet.

VI. CONCLUSION

This paper has described an IoT-based supervisory control framework for fault current limiters in HVDC substations, implementing a three-layer architecture that separates the sub-millisecond FCL triggering function from the second-to-minute health monitoring function and the long-horizon predictive analytics function. On a ± 500 kV HVDC HIL testbed, the system detected fault inception in 1.8 ms, limited a 28 kA prospective fault current to 5 kA (an 82 % reduction), maintained end-to-end SCADA latency of 4.2 ms, and achieved 99.97 % system availability over six months. The false positive trip rate of 0.003 % represents a substantial improvement over conventional relay-based approaches and is attributable to the LSTM edge classifier.

Three aspects of the design appear to have broad applicability beyond the specific FCL and HVDC topology studied here. First, the principle of co-locating the AI inference engine with the protection hardware — rather than relying on a cloud round-trip — is likely to be necessary for any AI-based protection function with millisecond requirements. Second, the separation of the fast GOOSE control path from the slower MQTT telemetry path, using the same edge gateway hardware, is a practical architecture that avoids the latency penalty of routing protection signals through general-purpose networking infrastructure. Third, the digital twin synchronisation approach, linking the live telemetry to a PSCAD simulation instance, provides a maintenance planning capability that has no equivalent in conventional relay-based protection.

Future work will focus on extending the LSTM classifier to handle evolving fault scenarios including inter-station faults in multi-terminal HVDC grids, validating the RUL prediction model against real SFCL aging data from field deployments, and reducing the cryocooler power consumption through improved thermal insulation to make the system more practical for remote substation sites.

ACKNOWLEDGEMENT

The authors thank the High Voltage Laboratory at SGGIS Institute of Engineering and Technology, Nanded, for access to the OPAL-RT HIL testbed, and BHEL Bhopal for providing the SFCL tape samples used in the prototype. Financial support from the Ministry of New and Renewable Energy, Government of India (grant no. MNRE/R&D/2024/FCL-HVDC-07) is gratefully acknowledged. The authors also thank the IEEE Power & Energy Society HVDC subcommittee for helpful discussions during the manuscript preparation.

REFERENCES

- [1] C. A. Baldan, C. S. Hosoume, and J. R. Cardoso, “Superconducting fault current limiters — technology development and power system applications: A review,” *IEEE Transactions on Applied Superconductivity*, vol. 34, no. 2, art. 5600116, 2024.
- [2] M. Noe and M. Steurer, “High-temperature superconductor fault current limiters: Concepts, applications, and development status,” *Superconductor Science and Technology*, vol. 20, no. 3, pp. R15–R29, 2007.
- [3] H. Schmitt, J. Amon, F. Kunde, and G. Damstra, “Application of fault current limiters in distribution systems,” in *Proc. IEEE PES General Meeting*, pp. 1–8, 2006.
- [4] J. Martinez, A. Torres, and L. Garcia, “Monitoring of superconducting fault current limiters using IoT and MQTT telemetry,” *Electric Power Systems Research*, vol. 215, art. 108934, 2023.
- [5] T. Ghanbari, E. Farjah, and A. Zandnia, “Development of a high-performance inrush current discriminating algorithm,” *IET Generation, Transmission & Distribution*, vol. 9, no. 5, pp. 458–466, 2015.
- [6] A. Patel and R. Bhatt, “Convolutional neural network-based fault classification for superconducting FCLs in HVDC grids,” *IEEE Transactions on Power Delivery*, vol. 38, no. 4, pp. 2811–2820, 2023.
- [7] IEC 62271-110, “High-voltage switchgear and controlgear — Part 110: Inductive load switching,” *International Electrotechnical Commission*, Geneva, 2017.
- [8] M. A. Khan, I. A. Khan, M. A. Faruque, and S. Saha, “Hardware-in-the-loop testing of a resistive SFCL for VSC-HVDC protection,” *IEEE Transactions on Applied Superconductivity*, vol. 29, no. 5, art. 5601905, 2019.
- [9] L. Zhang, Q. Huang, W. Chen, and Y. Zhu, “Solid-state fault current limiter with self-adaptive triggering for HVDC applications,” *IEEE Transactions on Industrial Electronics*, vol. 69, no. 7, pp. 6830–6839, 2022.
- [10] S. H. Lim, B. S. Han, and J. C. Kim, “Characteristics comparison of a flux-lock-type SFCL with REBCO coated conductor in a scaled HVDC grid,” *IEEE Transactions on Applied Superconductivity*, vol. 30, no. 4, art. 5602205, 2020.
- [11] T. T. Nguyen, H. J. Lee, and C. K. Kim, “Design and performance evaluation of a hybrid current limiter for VSC-HVDC,” *IEEE Transactions on Power Electronics*, vol. 34, no. 6, pp. 5613–5623, 2019.
- [12] IEEE Std C37.119-2016, “IEEE Guide for Breaker Failure Protection of Power Circuit Breakers,” *IEEE*, New York, 2016.