



A Review of Stability-Oriented FCS-MPC for High-Frequency SiC Power Converters

Yifei Gao

School of Energy Systems, Lappeenranta-Lahti University of Technology LUT, Lappeenranta,
53850, Finland

Yifei.Gao@student.lut.fi

Abstract. Wide-bandgap silicon-carbide (SiC) devices enable switching frequencies above 100 kHz. They keep high power density and fast dynamics. High-frequency operation also magnifies sampling-to-computation delay, parameter drift, and parasitic effects. These factors challenge the stability of Finite-Control-Set Model Predictive Control (FCS-MPC). This paper reviews three routes. The first is fast vector pre-selection. The second is explicit delay compensation. The third is online parameter/disturbance adaptation. We compare representative methods and reveal the trade-offs among computational complexity, stability margins, and robustness to thermal drift. We also note two engineering bottlenecks. One is model-error-delay coupling and the resource limit of FPGA-SoC platforms. The other is coordinated control in multi-module systems. Finally, we discuss future combinations of FCS-MPC with long prediction horizons, machine learning, and hardware acceleration. This research provides crucial theoretical support and practical guidance for the stable control of high-frequency power electronic systems. It holds significant value for promoting the industrialization of the application field of wide bandgap devices. By systematically analyzing the design bottlenecks of the integration of control algorithms and hardware platforms, this research has pointed out the innovative direction for the development of the next-generation highly reliable power converters.

Keywords: Delay Compensation, Extended State Observer, Vector Pre-selection.

1 Introduction

The increasing need for high power density and high efficiency continues to grow. Some of the big drivers are renewable generation, on-board fast charging, and more-electric aircraft. Traditional silicon-insulated-gate bipolar transistors (IGBTs) typically do not meet targets above 50 kHz. Silicon carbide metal-oxide-semiconductor field-effect transistors (SiC MOSFETs) have a higher critical field and low gate charge. They can drive three-phase inverters and Vienna rectifiers to ≥ 100 kHz, with system efficiencies $\geq 97\%$ [1][2]. However, high dv/dt and package parasitics raise common-mode voltage. They may also unbalance the DC-link neutral point. Temperature-driven

© The Author(s) 2026

S. Zhang (ed.), *Proceedings of the 2025 International Conference on Electronics, Electrical and Grid Technology (ICEEGT 2025)*, Advances in Engineering Research 292,

https://doi.org/10.2991/978-94-6463-986-5_86

parameter drift further reduces the phase margin of proportional – integral (PI) and proportional – resonant (PR) control at high frequency [3][4].

FCS-MPC is attractive for high-frequency SiC systems. It uses a discrete-time model. It handles multiple constraints. It also drives the optimal switching state without a modulator [5]. Microsecond sampling periods bring a triple challenge. A full search over 8–44 vectors is heavy for a microcontroller unit (MCU). A 1–2 T_s delay shifts the system roots and reduces phase margin. The on-state drain-to-source resistance $R_{ds(on)}$ grows with temperature. That growth causes model mismatch and degrades prediction accuracy [6][7]. These issues have driven new algorithms around “High-Frequency SiC + FCS-MPC.”

Existing studies follow three routes. The first is fast vector screening. The second is explicit delay compensation. The third is online parameter adaptation. However, a systematic view of the trade-offs is still missing. The trade-offs involve computational load, stability margins, and robustness against thermal drift. Multi-module paralleling also needs deeper study. The same is true for FPGA–SoC power consumption. The interpretability of machine learning (ML) remains an open issue.

This work focusses on real-time control and stability in excess of 100 kHz, for SiC converters. The predominant problem areas are computation burden, loop delay, and drift due to temperature. We present three different strategies for handling these issues: fast vector screening with a selected FCS-MPC model, an explicit delay compensation, and a trial based online compensation for disturbance and/or parameter drift. We present how each of these issues affect switching frequency, total harmonic distortion (THD), and power use of the controller. The performance of any controller has now become the locus of limitation to the performance of systems, with respect to real-time control, and has also become a bottleneck in commercial applications for high-power density electric transportation, more-electric aircraft, new energy storage, scouting drones, etc. This essentially limits the advantages of SiC devices. By systematically evaluating the three routes to FCS-MPC we can assist the cross-layer co-design requirements. The relevant layers are algorithm, hardware, and thermal management. By evaluating these criteria for operation in excess of 100 kHz we can also support open test standards.

2 Control challenges in high-frequency sic converters

High-frequency SiC converters face three linked challenges. First, electromagnetic coupling. dv/dt above 50 kV/ μ s and layout parasitics can excite MHz resonances. In a Vienna rectifier, they also aggravate neutral-point voltage oscillation [3]. Second, delay dominance. At 100 kHz, the sampling period T_s is about 10 μ s. ADC, computation, and PWM update can occupy 30–60% of one cycle. This cut alone may remove 20–30 degrees of phase margin [6]. Third, thermal-induced mismatch. Temperature rise changes $R_{ds(on)}$ and leakage inductance. The junction temperature can rise notably as the switching frequency increases. The associated model mismatch can push prediction errors beyond acceptable levels [7]. A minimum common-mode-voltage constraint improves EMC under high dv/dt . A Total Frequency Spread (TFS) metric quantifies

switching-frequency dispersion under FCS-MPC [4][8]. These effects act together. They define the stability boundary above 100 kHz. Any single fix must be co-designed with the other two compensations. Only then can we get a robust and practical solution.

3 Progress on Fast Algorithms and Stability Enhancements

3.1 Fast vector screening

There are two typical ideas. The first is vector reduction. Sun et al. compressed 27 vectors to 8 in a three-level Vienna rectifier. A dynamic hysteresis band controlled the effective switching frequency. The search cost dropped by about 75%. The method fits MCU platforms [3]. The second is vector expansion with precise localization. Yan et al. expanded the valid set to 44 vectors in a two-level voltage-source inverter (VSI). Sector identification enabled one-shot optimal selection. FPGA parallelism kept the 100 kHz decision delay near 1 μ s [9]. Each idea has limits. Vector reduction sacrifices output freedom. Vector expansion relies on hardware resources.

3.2 Explicit delay compensation

For delays of 1–3 T_s , the literature is split between two approaches: the first method uses a window of reasonable length with a forward lookahead to compensate for the delay. The second method uses a long window and rolling optimization approach. The predicted state can be written in at least two ways. The first uses a fixed time delay which is known. In that case, the compute cost is similar to that of basic FCS-MPC, and only the effect will decrease once the load or temperature change causes the delay to drift and therefore used as a compensator. The second approach uses a long window and a rolling optimization approach, and applies sphere decoding, or a branch-and-bound which searches and optimizes the switch sequence over a state space of length 4–6. It will tolerate a delay of $\geq 2 T_s$ and also be robust to parameter drift. The compute load rises by 30 – 50% [6].

3.3 Online parameter/disturbance adaptation

Real-world operation changes $R_{ds(on)}$ and leakage inductance with temperature and current. Babayomi et al. introduced a structurally-adaptive extended state observer (SA-ESO). The observer works with MPC in a bidirectional DC–DC converter. It estimates modeling errors and external disturbances in real time. It also improves steady-state accuracy in both charge and discharge modes [7]. Wan et al. added a loss-equalization term to the cost function in a three-level ANPC inverter. The term balances device conduction and commutation. It spreads heat and also helps neutral-point-voltage control [10]. Such methods often pair with vector screening or candidate compression. The observer and loss estimation add compute load. Reducing the search set helps keep real-time performance.

3.4 Representative results comparison

Table 1. Representative Results.

Route	Method / Platform	f _{sw} / Rating	Highlight	Limitation
Vector reduction	“Longest Extension Step” FCS-MPDPC / dSPACE	20 kHz / Vienna 3 kW	−75% search; MCU-friendly	Non-SiC; low frequency
Vector expansion	44-vector sector optimization / FPGA	100 kHz / 2L-VSI 5 kW	≈1 μs decision; THD < 3%	High resources; no thermal adaptation
Parallel convolution	LabVIEW-FPGA / cRIO	80 kHz / SiC-VSI 4 kW	Graphical workflow; 35% logic	T _s = 12.5 μs still long
Long window	Sphere decoding / ZCU104	100 kHz / SiC-VSI 7.5 kW	Tolerates 2 T _s ; better stability	+40% compute; needs cooling
SA-ESO + MPC	Bidirectional DC–DC / HIL	150 kHz / 11 kW	≤2% steady-state error	No physical prototype yet
Loss equalization	Loss-equalization MPC / Prototype	20–50 kHz / 3L-ANPC	Hot-spot mitigation; NP control	>100 kHz not shown

Table 1 illustrates six representative implementations. It provides a way to visually compare three additional objectives: minimizing complexity, robustness to delay, and mitigating thermal drift. In vector reduction, Sun reduced the Vienna candidates by 75% on an MCU. The run time was under 100 μs. The hardware threshold was the least, but the switching frequency was just 20 kHz [3]. For graphical hardware acceleration, Iranian used LabVIEW-FPGA to run parallel convolution search. At 80 kHz the logic usage was 35%. The sampling period of 12.5 μs was still long [5]. In vector expansion with delay compensation, Yan used an Artix-7 FPGA and reduced the 44-vector decision time to 1 μs. Villalón then applied a six-step sphere-decoding window. The controller tolerated a 2 T_s loop delay and kept a −7 dB phase margin. The extra dynamic power was about 4 W [6][9]. For online disturbance/parameter compensation, Babayomi used a structurally-adaptive extended state observer. The power error stayed within 2% from 25–125 °C. Wan added a loss-equalization term in a three-level ANPC prototype and reduced the hot-spot temperature by ~18 K [7][10].

In short, vector reduction is hardware-friendly but frequency-limited. Vector expansion with a long window delivers low THD and delay robustness above 100 kHz, but it has the highest controller power. Adaptive methods improve robustness against thermal drift. They add about 30% compute and still lack large-power physical validation. Future work should pair these methods with low-power hardware acceleration.

Beyond these cases, Ma ported high-frequency SiC control to a 10 kW back-to-back three-phase inverter in a microgrid. This shows good scalability for grid-connected and islanded modes [11]. Marsal reviewed fast computation methods such as lookup tables, pipelining, and near-optimal approximations. These methods can cut execution time by

20–60% on current FPGA platforms [12]. Determinedly, vector reduction is hardware-friendly but frequency-constrained. Vector expansion requires a long window before maintaining low THD and delay robustness for frequencies above 100 kHz. Vector expansion is characterized by the highest required controller power. Adaptive control is possible with some additional complexity (about 30% more compute), supporting aspects of thermal-drift robustness, although large-power prototypes have not followed. As previously discussed in the understanding of system-level application and further compute reduction from static modeling, all versions require ecosystem layer co-optimization.

4 Comparative case studies

To verify the three routes on real platforms, we select three cases. The first is a Vienna rectifier for vector reduction. The second is a two-level SiC inverter for vector expansion plus delay compensation. The third is a bidirectional SiC dual-active-bridge (DAB) for parameter adaptation.

4.1 Vienna rectifier — vector reduction with hysteresis

Sun compressed 27 space vectors to 8 and added a dynamic hysteresis band. The prototype rating was 3 kW. The nominal switching frequency was 20 kHz. The effective switching frequency was limited to 15 kHz. The computation dropped by about 75%. A dSPACE DS1103 finished prediction within 100 μ s. The input-current THD fell from 3.4% to 2.1%. The NP-voltage ripple stayed below 1 V. The devices were not SiC and the frequency was below 100 kHz. Even so, the idea is portable to mid-power platforms [3].

4.2 Two-level SiC inverter — vector expansion + long-window delay compensation

Yan expanded the candidate set to 44. Sector identification enabled one-shot selection. An Artix-7 FPGA performed parallel convolution. The worst-case decision delay at 100 kHz was about 1 μ s. The line-to-line voltage THD was 2.8% [9]. Villalón embedded a six-step sphere-decoding window. The design compensated a 2 T_s loop delay. Stability margins improved; the gain margin was about -7 dB. Low-frequency damping also improved. Logic utilization was about 42%. Dynamic power was about 4.3 W [6]. The scheme achieves microsecond-level latency and $<3\%$ THD. The cost is higher controller power and tighter thermal design.

4.3 Bidirectional SiC DC–DC (fast charging) — SA-ESO + MPC

Babayomi linked SA-ESO with MPC in hardware-in-the-loop (HIL) at 150 kHz. The observer estimated model errors, disturbances, and thermal drift accurately and in real-time. The power error was maintained within 2% at temperatures ranging from 25 to

125 °C. The DC-bus ripple was reduced by 35%. The compute load increased by ~30%. The overall latencies remained less than 13 μ s[7]. The method shows robust potential at high frequency and high temperature. A large-power physical prototype is still needed.

4.4 Cross-case performance comparison

Table 2. Performance Comparison [7]

	Vienna (3 kW)	2L-VSI (5 kW)	DAB (11 kW, HIL)
Device / f_{sw}	IGBT / 20 kHz	SiC / 100 kHz	SiC / 150 kHz
Route	Vector reduction + hysteresis	Vector expansion + sphere decoding	SA-ESO + MPC
Decision latency	100 μ s (MCU)	1 μ s (FPGA)	13 μ s (RT-CPU)
THD / stability margin	2.1% / —	2.8% / GM -7 dB	— / —
Extra controller power	+0.3 W (MCU)	+4.3 W (FPGA)	+30% CPU load
Limitation	20 kHz only	Resource and thermal limits	No physical prototype

Table 2 shows clear trends. Vector reduction suits mid-frequency MCU control. Vector expansion with a long window keeps low THD at high frequency and resists delay. It is sensitive to FPGA resources and cooling. SA-ESO with MPC improves steady-state accuracy under thermal and model drift by online estimation. So far it is validated only in HIL and adds about 30% compute load. None of the three has been evaluated with a unified metric on multi-module parallel systems or on >10 kW physical platforms. Cross-layer validation remains a priority.

5 Challenges and future directions

There are still challenges at cross-layers of the stack when operating >100 kHz with long-term stability. First, delay compensation algorithms generally assume a fixed delay. Due to queuing and processing jitter in distributed and parallel systems, this delay is random. The next effort is to connect delay estimation with online identification of thermal parameters, preferably in a predictive model that can be updated in each loop.

Second, the challenge of impedance coupling in parallel modules is unsolved. For example, while there are a couple of SiC units in parallel, there may be a high and positive dv/dt on a switching device inducing non-negligible circulating current and potentially subsynchronous oscillations on a DC bus and an AC side. While local FCS-MPC may not stabilize impedance coupling across modules, we cannot guarantee safety

beyond that level. We need impedance-aware MPC co-design or hierarchical predictive control.

Third, I believe limits on controller power and thermal limits could become a bottleneck. If we move search and prediction to an FPGA/SoC, we may achieve an approximate 1 μ s latency at the expense of approximately 4-6 W of added controller power. It could be with a low-power RISC-V vector unit, or dedicated AI accelerators, or chiplet-based, or simply SoCs.

Machine learning (ML) is now enabling real-time state estimations, fast vector predictions, and online weight tuning. However, while deep learning CNN-style black-box models can assist with systems where noise is low, they lack robustness in high-noise regimes and provide no reliability in extreme design scenarios. Physics-constrained neural networks or explainable AI (XAI) approaches are promising for improving safety and auditability.

There are myriad test benches used for experimentation today, each with highly varied power levels, types of converters, combinations of topologies, and various metrics. The diversity of all the test benches make an equitable comparison and ML transfer impossible! A major, and utmost foundational task remains an open test standard and dataset for >100 kHz SiC converters.

In conclusion, going forward, an algorithm, hardware, and thermal design need to work together. Use fast vector screening as the foundation of the algorithm design. Use long prediction windows to absorb random delays. Use adaptive estimation of (thermal) parameters to cancel thermal drifts. Then fuse the entire stack with low-power accelerator as many have suggested.

When all is said and done; safe, fast, frugal, explainable (>100 kHz) SiC converters have 'arrived', meaning they are ready for vehicles, aviation (the sky), and large-scale energy storage.

6 Conclusion

We compared three routes for FCS-MPC in >100 kHz SiC converters. Vector reduction on MCUs cuts search by ~75%. It is limited in frequency. Vector expansion with a long window achieves $\approx 1 \mu$ s decisions and ~2.8% THD at 100 kHz. It adds ≈ 4 W of FPGA power. SA-ESO-based adaptation holds power error within 2% from 25 to 125 $^{\circ}$ C. It increases compute by ~30% and still awaits large-power prototypes. These findings define a three-dimensional trade-off: switching-frequency capability, real-time performance, and robustness to thermal drift. The framework supports algorithm-hardware co-design for high-density applications.

Future work should fuse multi-objective adaptation with long prediction horizons. It should also exploit low-power RISC-V and heterogeneous FPGA-SoC acceleration. Finally, open datasets and unified tests will speed adoption in vehicles, aviation, and storage.

References

1. Y. Fan, L. Tong, Y. Lu et al., High-switching-frequency SiC power conversion systems with improved finite control set model predictive control. *Energies* **17**(18), 4601 (2024).
2. R. Gregor, S. Toledo, E. Maqueda et al., Advancements in power converter technologies: A focus on SiC-MOSFET-based voltage source converters. *Energies* **16**(16), 5994 (2023).
3. Y. Sun, X. Feng, Z. Dai et al., A novel optimal voltage vector selection control strategy for Vienna rectifier. *IEEE Access* **11**, 49422–49430 (2023).
4. P. Qing, J. Xiong, E. Ma, A model predictive control scheme with minimum common-mode voltage for PMSM drive system fed by VSI. *Machines* **12**(5), 292 (2024).
5. M. E. Iranian, E. Zamiri, A. de Castro, Optimal implementation of d–q frame finite control set model predictive control with LabVIEW FPGA. *Electronics* **14**(1), 100 (2025).
6. A. Villalón, C. Burgos-Mellado, M. Rivera et al., Real-time model predictive control for two-level voltage source inverters with optimized switching frequency. *Applied Sciences* **15**(13), 7365 (2025).
7. O. O. Babayomi, Z. Zhang, Z. Li et al., Bidirectional DC–DC converters for distributed energy resources: Robust predictive control with structurally-adaptive extended state observers. *International Journal of Electrical Power & Energy Systems* **158**, 109913 (2024).
8. T. Tricarico, J. A. Costa, D. Herrera et al., Total frequency spread: A new metric to assess the switching frequency spread of FCS-MPC. *Energies* **15**(14), 5273 (2022).
9. S. Yan, C. Li, Y. Cui et al., An improved FCS-MPC based on virtual vector expansion and sector optimization for 2L-VSCs. *IEEE Access* **10**, 127450–127460 (2022).
10. S. Wan, B. Wang, J. Chen et al., Model predictive control strategy based on loss equalization for three-level ANPC inverters. *Actuators* **13**(3), 111 (2024).
11. C.-T. Ma, Z.-Y. Zheng, Design and implementation of a SiC-based multifunctional back-to-back three-phase inverter for advanced microgrid operation. *Micromachines* **14**(1), 134 (2023).
12. E. Marsal, M. R. Arahal, M. G. Satué et al., Review of fast computation methods for finite-state predictive control of multi-phase drives. *Applied Sciences* **14**(24), 11774 (2024)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (<http://creativecommons.org/licenses/by-nc/4.0/>), which permits any noncommercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

