

Specifics of Software Quality Assurance in High-Frequency Trading (HFT) Systems

Khapankou Anton*

Senior quality assurance engineer, Belarus, Minsk

Email: khapabnov.anton@gmail.com

Abstract

The article is devoted to the analysis and reconsideration of specialized methods and approaches to quality assurance (QA) in high-frequency trading (HFT) systems. Within the framework of the study, key architectural vulnerabilities of HFT systems are detailed, contemporary approaches to regression testing are examined, and practices for incorporating hardware components into the development loop are considered. Special emphasis is placed on the problem of non-determinism in parallel computational pipelines and on methods for reducing the variability of temporal characteristics, including minimizing jitter. The aim of the work is to formulate an integral QA strategy that ensures an optimal balance between extremely high execution speed and guaranteed reliability of the software code. To achieve this goal, methods of systems analysis, comparative examination of architectural patterns, and the integration of best practices accumulated in ultra-low-latency engineering are employed. In the final part, the author presents the conceptual approach of Continuous Latency Integration, which is considered a logical extension of the ideas of continuous integration as applied to time-efficiency metrics. The material is addressed to software system architects, QA engineers in the field of fintech, and developers of trading algorithms.

Keywords: high-frequency trading; quality assurance; FPGA verification; low latency; algorithmic trading; performance testing; jitter.

1. Introduction

The relevance of the study is determined by the fact that the modern infrastructure of financial markets operates in regimes that a priori lie beyond the limits of human reaction. Well-known incidents of the Flash Crash class convincingly demonstrate that traditional approaches to software testing prove to be ineffective when applied to pipelines operating in nano- and microsecond intervals. Under these conditions, quality assurance in an HFT environment ceases to be reducible to the trivial identification and elimination of bugs and is transformed into the task of ensuring determinism of system behavior and predictability of response time.

Received: 10/16/2025

Accepted: 12/16/2025

Published: 12/26/2025

** Corresponding author.*

The aim of the work is to develop and provide theoretical and practical justification for a comprehensive approach to quality assurance of HFT systems, which simultaneously takes into account the logical correctness of the implemented algorithms and the strict timing constraints imposed on their execution.

To achieve this aim, the following research **tasks** are to be solved.

- to analyze the dominant architectural patterns for building HFT systems (including configurations using FPGA and GPU) and to identify specific failure vectors arising precisely in these technological contexts;
- to systematize approaches to the measurement and control of microstructural quality metrics such as tick-to-trade latency, jitter, and throughput under conditions of high market volatility;
- to formulate practice-oriented recommendations for incorporating methods of formal verification and stochastic testing into the CI/CD pipeline for trading platforms.

The novelty of the work lies in the fundamental shift of focus from classical functional testing to the verification of timing characteristics as a fully fledged functional requirement for the system.

The author's hypothesis is formulated as follows: in the context of HFT systems, it is impossible to strictly separate the concepts of code quality and infrastructure quality; consequently, an effective QA strategy must initially provide for elements of hardware verification (including FPGA co-simulation) at the earliest possible stages of the development life cycle, in line with the Shift-Left Testing approach.

2. Materials and Methods

In the preparation of this article, the method of systematic literature review was used, aimed at identifying advanced engineering practices and theoretical developments in the field of high-frequency trading (HFT). The methodological basis of the study was a detailed analysis of technical reports, publications in peer-reviewed scientific journals, as well as materials from leading technological conferences.

The search work was carried out using international abstract and full-text databases Scopus, Web of Science, IEEE Xplore, as well as the preprint repositories of arXiv.org (sections Quantitative Finance and Computer Science). The search strategy was based on the use of combinations of English-language keywords: High-frequency trading software quality, FPGA verification trading, Low latency testing methodologies, Algorithmic trading reliability.

The primary stage of source selection was based on the analysis of titles and abstracts. Of more than 50 identified publications, only those works were retained that contained specific information on architectural solutions, performance metrics, or testing methodologies. Studies focusing exclusively on the economic analysis of the profitability of trading strategies without consideration of the technical aspects of software implementation were excluded.

Special emphasis was placed on sources describing hybrid hardware–software systems, since the modern HFT

infrastructure increasingly relies on the use of FPGAs. To identify current trends in the field of verification, industrial reports of key vendors of EDA (Electronic Design Automation) systems were additionally analyzed.

In addition, the method of comparative analysis was applied, which made it possible to compare approaches to testing in the classical enterprise segment and in the sphere of ultra-low-latency systems. The synthesis of the obtained results provided the formation of a holistic understanding of quality requirements in the subject area under consideration.

3. Results

As a result of the analysis of contemporary scientific and applied literature, it has been established that quality assurance practices in high-frequency trading (HFT) systems fundamentally diverge from classical approaches to software development and testing. Current HFT strategies increasingly rely on machine learning (ML) models, which introduce a pronounced element of non-determinism into the system and radically complicate the predictability of its behavior under extreme market regimes [1]. In such a configuration, standard unit tests prove to be methodologically insufficient: they are unable to adequately cover scenarios in which the algorithm's output decisions are determined by a complex nonlinear response to microstructural features of the market and depend on the dynamic context rather than on a fixed set of input data.

One of the central challenges is the verification of FPGA-based solutions. The growing architectural and logical complexity of FPGA designs leads to projects requiring repeated iterations due to functional defects that were not detected at early stages of verification [2]. In HFT conditions, even a minimal error in the logic of an FPGA module (for example, in a market data parsing block) can trigger the generation and transmission of incorrect orders on a nanosecond scale, before the software layer has time to react and correct the behavior of the system. Consequently, quality assurance inevitably shifts toward the use of formal verification methods that make it possible to mathematically prove the correctness of the operation of critical hardware modules and thereby reduce the likelihood of catastrophic failures [2, 5].

The analysis of architectural patterns shows that the drive to minimize latency systematically pushes developers to abandon standard operating system protection mechanisms, including kernel bypass and extensive use of shared memory regions without locks. Such optimizations do indeed make it possible to significantly reduce latency, but at the same time, they make the system extremely sensitive to synchronization errors and race conditions [4]. Testing such low-level configurations requires specialized tools capable of injecting failures or anomalous behavior at the level of individual CPU cycles or single network packets in order to reproduce rare but potentially critical scenarios.

Within this paradigm, QA tasks go beyond testing isolated modules and require the use of agent-based modeling, in which the behavior of a trading robot is validated in an artificial environment saturated with other aggressive, competing algorithms. For the correct assessment of FPGA design timing, it is necessary to generate and replay representative time series that simulate actual market dynamics [6].

Significant attention in recent works is also paid to the problem of infrastructure testing. A substantial share of

latencies and errors arises not in the logic of the trading strategy itself, but in the network infrastructure and at the level of switching equipment [7,8]. In this regard, the area of responsibility of an HFT QA engineer inevitably includes quality control of the network configuration: the correctness of TCP offload settings, the consistency and accuracy of time synchronization via PTP, and other parameters of the low-level stack. Ignoring rare but significant outliers in response time is unacceptable, since it is precisely during periods of heightened volatility, when such outliers are most likely, that the largest number of both highly profitable and most loss-making trades is executed [9].

Taken together, the results of the analysis indicate a shift from a reductionist approach focused on testing only software code to a comprehensive Verification and Validation (V&V) concept, within which the object of verification is no longer an individual component but an integrated system that includes FPGA modules, the network stack, ML models, and business logic. This integrated system is considered a single real-time loop with strict timing constraints and high reliability requirements, which radically raises the bar for quality assurance methods and tools [10].

4. Discussion

This section is devoted to the integration of the results presented above and to the formulation of the author's proposed quality assurance model for high-frequency trading systems. In contrast to typical practices applied in banking information systems, HFT requires an actual inversion of the classical testing pyramid and the implementation of the Continuous Latency Integrity concept, in which the control and validation of temporal characteristics become the dominant aspect of the QA process rather than an ancillary non-functional requirement.

Drawing on the analyzed works, it can be asserted that the key methodological challenge is the Observability Effect: any attempt to increase the observability of the system through direct measurements of its parameters (for example, by means of detailed logging or embedded profiling) in itself introduces additional delays, often unacceptable for HFT, and thereby distorts the measured picture.

In this regard, the proposed architecture of the QA process should be organized around the principle of passive monitoring and the use of hardware simulation, which makes it possible to maximally separate the observation loop from the execution loop. The diagram below presents the conceptual structure of the test loop that implements this approach and demonstrates the flow of artifacts and quality control signals (Fig. 1).

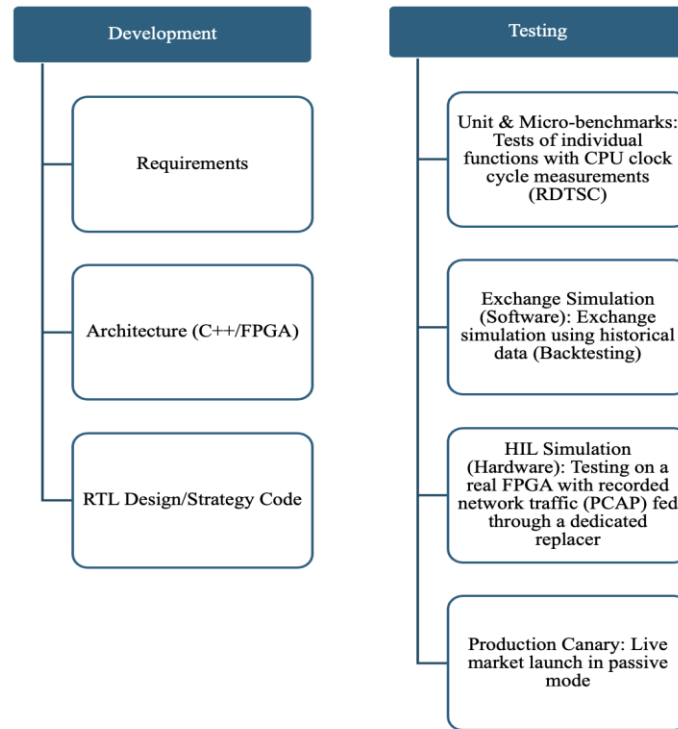


Figure 1: Conceptual structure of the test loop [2, 5, 7]

The key element of the proposed model is the HIL Simulation stage. It is precisely the execution of the configuration on real hardware that makes it possible to reveal subtle timing errors that, in principle, do not manifest themselves under conditions of purely software emulation. In other words, without the inclusion of physical hardware in the testing loop, a significant portion of temporal anomalies and violations of determinism remain latent. At the next step, it is necessary to perform a systematization of the tools used at various levels of the test loop.

Table 1 provides a comparative mapping of the corresponding approaches, demonstrating their area of applicability, limitations, and contribution to ensuring latency integrity.

Table 1: Comparative analysis of testing methods in classical software and HFT systems [3, 4, 8]

Characteristic	Classical software (Web/Enterprise)	HFT systems (Low Latency)
Success criterion	Functional correctness	Correctness + determinism
Networking	Abstraction (HTTP/REST)	Direct access (kernel bypass, Solarflare)
Test data	Synthetic data	Replay of real market traffic (L3 data)
Tools	Selenium, JUnit, JMeter	C++ microbenchmarks, PCAP replay, Verilator
Monitoring	Logs, APM (New Relic)	TAP aggregators, hardware timestamps

Continuing the exposition of the author's model, it is necessary to place special emphasis on the aspects of risk management. The works [3, 10] highlight the threat that arises when a trading robot, due to a failure, incorrect configuration, or unpredictable interaction with the external environment, begins to generate destructive trading activities with a potentially systemic effect on the market. To prevent such scenarios, the risk management mechanism should be regarded not as a superstructural component but as an integral part of the quality assurance architecture.

In this context, the QA loop must initially incorporate a procedure for verifying the correct operation of the Kill Switch as the last line of defense. This loop should be classified as critical: it is subject to testing with maximum priority and an increased depth of verification, since its failure or incorrect operation creates conditions for cascading risks in the financial system [3, 10].

In addition to the architecture, the choice of metrics is important. The average latency alone is insufficient. Table 2 proposes a set of metrics that are mandatory for HFT QA.

Table 2: Key quality metrics of an HFT system [3, 10]

Metric	Description	Measurement method
Tick-to-Trade (T2T)	Time from receiving the market data packet to sending the order packet.	Hardware (FPGA timestamping)
Wire-to-Wire	Total time of signal propagation through the system, including network interface cards.	TAP splitters + oscilloscope/analyzer
Jitter	Variability of latency.	Latency distribution histogram
Order-to-Ack	Time from sending an order to receiving an acknowledgment from the exchange (depends on the exchange, but is important for the strategy).	Log analysis (PCAP)

The approach proposed in this work, which combines FPGA simulation, rigidly formalized risk gates, and a closed analytics loop, in fact transforms the quality assurance process: from delayed, reactive defect elimination to predictive management of system reliability and performance. In the context of HFT, code quality is treated as synonymous with its determinism: a system is considered to be of high quality not when it exhibits low latency in the average case, but when it is guaranteed not to exhibit pathologically high latencies.

5. Conclusion

In the course of the study, all the objectives set were consistently addressed. First, a detailed analysis of the architectural features of modern HFT systems was performed, on the basis of which a stable trend towards the transition to hybrid solutions using an FPGA was identified. It has been shown that such a shift in architectural

focus requires a fundamental revision of traditional approaches to testing and verification. Second, methods for measuring specialized performance metrics were systematized, which made it possible to substantiate the key importance of accurate measurement of jitter and latency as the main indicators of operational and financial risk. Third, an integral approach to quality assurance was proposed, including HIL simulation and the use of formal verification methods for critically important components.

The key conclusion of the work is that quality assurance in high-frequency trading is, in essence, a multidisciplinary task at the intersection of software development, hardware design, and big data analytics. Ignoring the hardware component in testing leads to a loss of control over the behavioral characteristics of the system under production conditions and, consequently, to the emergence of uncontrollable financial risks.

References

- [1]. G.Ibikunle, B. Moews, D.Muravyev, K. Rzaev. "Data-driven measures of high-frequency trading." *arXiv preprint arXiv:2405.08101*. 2024. <https://doi.org/10.48550/arXiv.2405.08101>.
- [2]. E. D. Sozzo, D. Conficconi, A.Zeni, M. Salaris, D. Sciuto, M.D. Santambrogio. "Pushing the level of abstraction of digital system design: A survey on how to program FPGAs." *ACM Computing Surveys*, vol.55(5), pp.1-48, 2022. <https://doi.org/10.1145/3532989>.
- [3]. F.Afshan, K.Y. Leong, A. Najmi, U. Razi, B. Lelchumanan, C.W.H. Cheong. "Fintech advancements for financial resilience: Analysing exchange rates and digital currencies during oil and financial risk." *Resources Policy*, vol. 88, 104432, 2024. <https://doi.org/10.1016/j.resourpol.2023.104432>.
- [4]. P. Bilokon, B. Gunduz. "C++ design patterns for low-latency applications including high-frequency trading." *arXiv preprint arXiv:2309.04259*, 2023. <https://doi.org/10.48550/arXiv.2309.04259>.
- [5]. A.A. Abdulsamad, S.R. Répás. "Application of FPGA Devices in Network Security: A Survey." *Electronics*, vol.14(19), 2025. <https://doi.org/10.3390/electronics14193894>
- [6]. S. Ahmed, N.Ahmad, N.A. Shah, G.E. M. Abro, A. Wijayanto, A. Hirsi, A.R. Altaf. "Lightweight aes design for iot applications: Optimizations in fpga and asic with dfa countermeasure strategies." *IEEE Access*, 2025. <https://doi.org/10.1109/ACCESS.2025.3533611>.
- [7]. P.S. Yadav. "Optimizing Serverless Architectures for Ultra-Low Latency in Financial Applications". *European Journal of Advances in Engineering and Technology*, vol.11(3), pp.146-157, 2024.
- [8]. G. Dasari. "Tick Data Quality Control: Detecting and Correcting Inconsistencies in High-Frequency Trading." *Journal of Computer Science and Technology Studies*, vol.7(4), pp.814-821, 2025. <https://doi.org/10.32996/jcsts.2025.7.4.94>.
- [9]. Z. Cobandag Guloglu, C. Ekinici. "Liquidity measurement: A comparative review of the literature with a focus on high frequency." *Journal of Economic Surveys*, vol. 36(1), pp. 41-74, 2022. <https://doi.org/10.1111/joes.12440>.
- [10]. K.Z. Zaharudin, M.R. Young, W.H. Hsu. "High-frequency trading: Definition, implications, and controversies." *Journal of Economic Surveys*, vol.36(1), pp. 75-107, 2022. <https://doi.org/10.1111/joes.12434>.