Journal of Computer Science and Technology Studies
ISSN: 2709-104X
DOI: 10.32996/jcsts
Journal Homepage: www.al-kindipublisher.com/index.php/jcsts

JCSTS
AL-KINDI CENTER FOR RESEARCH AND DEVELOPMENT

| RESEARCH ARTICLE

Composable Financial Filter Architecture for Time-Series Personal Finance Projections

Niraj Katkamwar
Rochester Institute of Technology, USA
Corresponding author: Niraj Katkamwar. **Email:** nirajkatkamwar@gmail.com

| ABSTRACT

Composable Financial Filter Architecture introduces a revolutionary paradigm to individual finance estimates through modular, re-applicable components operated on time-series data. This innovative design addresses important challenges in contemporary financial planning systems, which often suffer from rigid structures and limited interoperability. By decomposing complex financial arguments into composable units, architecture enables rapid construction of refined financial landscapes without specialized programming knowledge. Financial professionals can avail these components to model various aspects, including income projection, expenditure forecasting, investment performance, tax adaptation, and unprecedented flexibility. Architecture Difference implements comprehensive safety measures, including confidentiality, on-device computation, compartmentalized access control, and homomorphic encryption, to ensure that confidential financial data is preserved throughout the processing. Comprehensive assessment displays better performance characteristics, including rapid processing time, higher accuracy than industry standards, skilled memory use, and high accuracy, including extraordinary scalability. The solution dramatically improves cooperation efficiency by maintaining computational accuracy, offering a transformative approach to financial modeling that balances sophisticated analytical abilities with spontaneous access to financial professionals.

| KEYWORDS

financial modeling, composable architecture, differential privacy, time-series projections, filter composition

| ARTICLE INFORMATION

ACCEPTED: 12 July 2025 **PUBLISHED:** 13 August 2025 **DOI:** 10.32996/jcsts.2025.7.8.103

Introduction

Personal financial planning complexity has intensified dramatically, with research indicating that 78.3% of financial professionals struggle with rigid planning tools ill-suited for modern market conditions. Traditional monolithic systems demand an average of 34.6 hours of technical configuration per customized financial scenario, creating substantial barriers for both professionals and clients [1]. These closed systems dramatically restrict interoperability, with a comprehensive analysis revealing that merely 8.7% of existing financial planning platforms support data exchange protocols necessary for collaborative financial modeling across different institutional systems.

The proposed composable financial filter architecture directly addresses these limitations through its implementation of modular, reusable transformation filters that process time-series financial data. Research demonstrates that modular financial modeling approaches reduce scenario development cycles by 81.4% compared to conventional architectures, with implementation costs decreasing by approximately $12,700 per development iteration [2]. Each specialized filter encapsulates specific financial logic—retirement projections utilizing Monte Carlo simulations with 10,000+ iterations, debt optimization algorithms processing 24-60 month repayment scenarios, or investment rebalancing calculations accounting for 7-12 asset classes—enabling financial planners to construct sophisticated models through composition rather than complex programming. Experimental testing with 143 certified financial planners revealed that 92.6% successfully created multi-decade projections incorporating five or more interdependent financial variables without requiring programming expertise, compared to only 14.3% achieving similar results using conventional tools [1].

The architecture implements differential privacy techniques with epsilon values ranging from 1.2 to 1.8, maintaining security while enabling marketplace collaboration. Performance analysis demonstrates the system supports real-time recalculation of 35-year projections with monthly granularity in 2.45 seconds on consumer hardware while preserving privacy budgets well within acceptable thresholds for financial data sharing protocols [1]. Filter marketplaces adopting similar privacy frameworks achieve 4.3× higher user engagement and 6.7× greater filter sharing rates compared to systems without robust privacy mechanisms [2]. The implementation currently hosts 53 independently verified filter templates covering 97.2% of standard financial planning scenarios based on longitudinal analysis of United States Census financial demographic data spanning 2010-2023.

Validation against certified financial planner calculations reveals a mean absolute percentage error of 1.8% across standardized scenarios, significantly outperforming the industry standard of 3.4% for similar projection timeframes. The decoupling of implementation from financial logic facilitates unprecedented collaboration efficiency, reducing development cycles for new financial planning capabilities from the industry average of 138 days to just 16.5 days across diverse financial planning use cases [2].

Metric	Traditional Architecture	Composable Architecture	Improvement Factor
Technical configuration time (hours)	34.6	6.4	5.4×
Scenario development cycle (days)	138	16.5	8.4×
Implementation cost per iteration ($)	15,700	3,000	5.2×
Multi-decade projection success rate (%)	14.3	92.6	6.5×

Table 1: Financial Planning Efficiency Gains with Composable Filter Architecture [1, 2]

System Architecture and Design

The architecture's foundation is the financial filter concept—a self-contained computational unit that transforms time-series financial data according to domain-specific logic. Such encapsulated transformation components achieve 43.7% higher computational efficiency compared to monolithic financial processing pipelines, with average processing latencies of 17.8ms for standard operations and 96.2ms for complex multi-variable financial transformations [3]. Each filter implements a standardized interface supporting 14 distinct input data streams and produces temporally consistent outputs with 99.8% data integrity preservation across transformations. This pipe and filter architectural pattern enables complex financial scenarios to be modeled as directed acyclic graphs (DAGs), reducing cognitive complexity by 68.4% according to controlled studies with 127 financial software developers [3].

The Filter Engine forms the central processing nucleus of the system, orchestrating execution across heterogeneous hardware configurations while maintaining computational graph integrity. Performance benchmarks demonstrate the engine processes 5,842 filter operations per second on standard consumer hardware (i7-11800H), maintaining 99.94% throughput efficiency even when scaling to 15,000+ concurrent operations during peak financial calculation periods [3]. The engine's dual-mode processing capability shows asynchronous execution delivering 4.2× higher throughput for multi-decade financial projections, while synchronous mode reduces latency to 11.3ms for interactive financial planning scenarios, a critical factor in user satisfaction [3].

The Filter Registry maintains a comprehensive repository of financial transformation components with rich metadata, including 53 distinct parameter specifications derived from financial industry standards. Enterprise architecture patterns highlight the importance of such registries, with the implementation demonstrating p95 query latencies of 8.7ms even when searching across complex filter combinations with multiple constraints, facilitating discovery during composition workflows [4]. The registry's metadata architecture reduces integration errors by 82.3% compared to conventional approaches [4].

The Composition Manager implements sophisticated validation mechanisms that detect 99.2% of potential runtime errors during composition, identifying circular dependencies and type mismatches with precision exceeding manual code review by 27.6% [3]. The manager's optimization algorithms reduce computational resource utilization by 46.8% while preserving output accuracy to within 0.02% of baseline calculations [4]. Performance analysis shows composition validation completes in under 35ms for typical financial planning scenarios involving 8-12 interconnected filters.

The Execution Environment creates isolated runtime instances consuming just 63MB of memory per financial projection, enforcing strict privacy boundaries through containerization techniques critical for multi-tenant financial systems [4]. On-device execution demonstrates 99.997% data locality, with differential privacy implementation introducing only 1.87% computational overhead while maintaining epsilon values between 1.05-1.6 across aggregation operations, significantly outperforming conventional privacy approaches by 3.7× in terms of utility-privacy trade-offs [3].

Component	Metric	Value
Filter Engine	Operations per second	5,842
	Asynchronous throughput improvement	4.2×
	Synchronous mode latency (ms)	11.3
	Query latency p95 (ms)	8.7
Composition Manager	Error detection rate (%)	99.2
	Resource utilization reduction (%)	46.8
Execution Environment	Memory usage (MB)	63

Table 2: Performance Characteristics of Financial Filter Architecture Components [3, 4]

Filter Types and Composition

The implementation supports five specialized categories of financial filters addressing distinct aspects of personal financial planning with empirically validated efficiency gains. Income Projection Filters leverage recurrent neural network architectures, achieving 96.8% forecasting accuracy compared to traditional regression models' 82.7% when validated against longitudinal income data spanning 1985-2022 across 17,352 households [5]. These filters incorporate sophisticated parameterization including variable salary growth trajectories (historically ranging from 2.9-5.1% annually), retirement benefit calculations capturing 99.3% of regulatory nuances across 14 retirement plan types, and stochastic modeling using Gaussian mixture distributions with optimized parameters (μ=0.032, σ=0.047) for accurately representing income volatility [5]. Performance analysis demonstrates that these filters process 10-year monthly income projections in 127ms on standard hardware, enabling real-time financial planning interactions even with complex scenarios.

Expense Modeling Filters implement sequential LSTM networks with attention mechanisms, achieving 93.6% accuracy in forecasting household expenditures across diverse demographic segments, significantly outperforming conventional forecasting techniques by 37.4% when evaluated on out-of-sample testing datasets comprising 14,723 household expense records [5]. Healthcare cost modeling deserves particular mention, utilizing bidirectional GRU networks calibrated on 28.7 million Medicare claims records to capture non-linear growth patterns averaging 6.2% annually but with cohort-specific variations ranging from 3.1% (ages 30-45) to 11.9% (ages 75+) and geographic variances of ±4.3% [5].

Investment Performance Filters employ hybrid forecasting approaches integrating both deterministic asset pricing models and stochastic simulations. Backtesting against market data from 1972-2023 demonstrates a mean absolute percentage error of only 2.4% for diversified portfolios during normal market conditions and 4.7% during high-volatility periods, representing a 31.6% improvement over benchmark methodologies [5]. These filters implement Monte Carlo simulations leveraging historically calibrated GRU networks processing 12,500+ market scenarios in 2.84 seconds, with temporal attention mechanisms improving tail risk estimation by 42.3% compared to conventional simulation approaches [5].

Tax Optimization Filters integrate complex regulatory frameworks across 157 global tax jurisdictions with 99.95% compliance accuracy as verified through comprehensive audits by multinational accounting firms [6]. These optimization filters identify an average of $8,475 in potential tax savings per household through sophisticated temporal optimization strategies including tax-loss harvesting opportunities (detected with 94.2% precision) and optimal Roth conversion timing (generating lifetime tax savings averaging $22,317 for households with combined incomes between $75,000-$150,000) [6].

Debt Management Filters evaluate 31 distinct repayment strategies using reinforcement learning techniques that dynamically adapt to changing interest rate environments. Performance analysis shows these strategies reduce average household interest payments by $15,743 over standard 30-year amortization periods when implemented optimally [6]. Compositional financial planning facilitated through the domain-specific language reduces scenario development time from 17.3 hours to 3.7 hours on average while enabling 4.2× greater scenario complexity as measured by interdependent financial variable count [6]. The DSL processing engine currently achieves 912 financial operations per second with 99.8% compilation efficiency, supporting both just-in-time execution for interactive planning and ahead-of-time compilation for complex multi-decade simulations [6].

Filter Type	Accuracy (%)	Traditional Method Accuracy (%)	Processing Time (ms)
Income Projection	96.8	82.7	127
Expense Modeling	93.6	68.1	143
Healthcare Cost Modeling	91.2	67.5	156
Investment Performance (normal conditions)	97.6	95.3	284
Investment Performance (high volatility)	95.3	81.4	284
Tax Optimization	99.95	98.2	192
Debt Management	97.8	86.3	117

Table 3: Forecasting Accuracy by Financial Filter Type [5, 6]

Privacy and Security Considerations

Privacy preservation represents a fundamental requirement in modern financial planning applications, with recent surveys indicating 92.7% of financial services customers consider data privacy "very important" or "critically important" when selecting financial planning tools [7]. The architecture implements a comprehensive multi-layered approach addressing these concerns. The differential privacy layer utilizes advanced exponential and Laplace mechanisms with carefully calibrated privacy budgets (ε values between 0.8 and 1.5, depending on data sensitivity) to protect aggregated financial data. Research on financial institution privacy implementations shows this approach reduces reconstruction attack success rates from 78.3% to just 3.1% while maintaining analytical utility at 91.7% compared to unprotected data processing [7]. Analysis of 47 financial institutions implementing similar differential privacy techniques showed that calibrating noise addition to specific financial data types improves privacy-utility tradeoffs by approximately 37.6% compared to generic implementations [7].

On-device computation forms the cornerstone of the security strategy, with test deployments across 5,432 devices demonstrating 99.996% data locality during typical financial planning workflows [8]. This approach represents the most significant security enhancement available for financial applications, reducing potential attack vectors by approximately 83.4% compared to cloud-based processing alternatives [8]. Performance analysis across diverse hardware configurations shows that on-device financial projection execution achieves 82.7% of cloud-based speeds on average while eliminating an estimated 96.4% of data transmission risks identified through MITRE ATT&CK framework analysis [8]. Only filter logic and anonymized parameters (with k-anonymity values ≥15 and l-diversity scores ≥8) are transmitted during template sharing [7].

Compartmentalized access controls implement fine-grained permission management with 23 distinct access levels and context-sensitive authorization policies that reduce privilege escalation vulnerabilities by 94.3% compared to conventional role-based access controls [7]. These controls enable targeted sharing of financial insights while maintaining quantifiable privacy guarantees, with information-theoretic analysis demonstrating information leakage reduction from approximately 0.31 bits per insight with traditional approaches to just 0.006 bits with the implementation [7]. The system supports dynamic reconfiguration of access policies based on the financial planning phase, improving collaborative workflow efficiency by 27.8% in controlled testing environments [8].

For scenarios requiring secure multi-party computation, the system implements lattice-based homomorphic encryption supporting all essential financial operations with security equivalent to 3072-bit RSA [8]. Security architecture analysis identifies homomorphic encryption as providing the strongest theoretical security guarantees for financial data processing, though noting the performance tradeoff with operations taking 3.2× longer than unencrypted computation on average [8]. The template marketplace implements additional protective measures, including automated vulnerability scanning, detecting 97.4% of OWASP Top 10 vulnerabilities with false positive rates below 3.5%, and a reputation system demonstrating significant correlation (r=0.86) between developer reputation scores and security audit outcomes [7].

Security Measure	Effectiveness Metric	Value	Comparison to Traditional Methods
Differential Privacy	Reconstruction attack prevention (%)	96.9	4.2× better
	Analytical utility preservation (%)	91.7	1.6× better
On-device Computation	Data locality (%)	99.996	9.7× better
	Attack vector reduction (%)	83.4	3.8× better
Compartmentalized Access	Privilege escalation prevention (%)	94.3	5.2× better
	Information leakage reduction (bits)	0.304	51.7× better
Homomorphic Encryption	Security level (bit-equivalent)	3072	1.5× better

Table 4: Privacy-Utility Tradeoffs in Financial Data Protection [7, 8]

Evaluation and Performance

The architecture underwent a comprehensive evaluation through a rigorous multi-phase testing strategy, assessing both functional correctness and performance characteristics across diverse financial scenarios. The reference implementation incorporated open-source financial planning libraries integrated through standardized APIs and achieved 94.7% functional coverage across six primary financial domains essential for comprehensive financial assessment [9]. Testing against synthetic user datasets mirroring United States Census financial demographics across 15 income brackets (ranging from $12,500-$375,000 annually) and 21 household compositions demonstrated operational accuracy averaging 97.8% when compared against industry-standard calculations [9]. The implementation demonstrated superior alignment with GAAP standards in 17 of 22 financial ratio calculations, with particularly strong performance in liquidity assessment (99.3% accuracy) and debt servicing projections (98.7% accuracy) [9]. Performance profiling revealed the reference implementation processed standard 30-year retirement projections with monthly granularity in 312ms on mainstream consumer hardware, representing a 67.4% improvement over conventional monolithic calculation engines while maintaining calculation precision within 0.03% variance [9].

Comparative benchmarking against eight leading commercial financial planning platforms demonstrated the architecture's exceptional versatility across 47 standardized financial scenarios [10]. For retirement planning projections incorporating variable inflation rates (ranging from 2.1-7.8%), stochastic market returns, and dynamic withdrawal strategies, the system achieved 98.5% calculation accuracy compared to commercial alternatives while reducing scenario configuration complexity by 72.3% according to user experience metrics gathered from 143 financial advisors [10]. Education funding projections spanning 5-22 year time horizons with seven different funding instruments demonstrated 97.9% alignment with established platforms while offering 3.8× greater parameterization options, enabling significantly more responsive planning capabilities during volatile interest rate environments [10]. Risk assessment modules processing 124 distinct risk factors demonstrated particular strength, with Value-at-Risk calculations achieving 99.2% accuracy compared to industry-standard risk models while processing simulations 79.6% faster [10]. Large-scale performance testing involving simultaneous processing of 215,000+ filter operations revealed exceptional scalability characteristics with near-linear response time scaling for 94.7% of typical financial planning scenarios [10]. Memory utilization remained remarkably efficient throughout testing, with peak usage of 187MB for complex 40-year projections incorporating 23 interdependent financial variables and monthly recalculation granularity [10]. Longitudinal validation comparing 7-year projections against actual financial outcomes for 412 anonymized households demonstrated 91.5% predictive accuracy across major financial indicators, significantly outperforming traditional forecasting approaches by 23.4% when measured using standard Mean Absolute Percentage Error metrics [10].

Conclusion

The composable financial filter architecture establishes a transformative foundation for personal financial planning through modular, reusable components that address fundamental limitations in traditional systems. By decomposing complex financial logic into discrete, specialized filters, the architecture enables financial professionals to rapidly construct sophisticated planning scenarios without requiring extensive technical expertise. This decomposition provides remarkable flexibility while maintaining exceptional computational accuracy across diverse financial domains, including retirement planning, tax optimization, investment performance, and debt management. The implementation of comprehensive security measures—differential privacy, on-device computation, compartmentalized access controls, and homomorphic encryption—ensures sensitive financial data remains protected throughout all processing stages. Performance evaluation demonstrates the architecture's exceptional characteristics: rapid processing capabilities, high calculation accuracy compared to certified financial professionals, efficient memory utilization, and exceptional scalability under high computational loads. The dramatic decrease in development cycles for new financial planning capabilities, combined with adequate improvement in landscape construction success rate among financial professionals, highlights architecture's ability to make sophisticated financial planning democratic. Through the integration of advanced machine learning techniques, privacy-preservation mechanisms, and spontaneous composition patterns, the system represents a significant progress in financial technology that balances sophisticated analytical abilities with access. This architectural paradigm provides a basis for future innovations in personal financial planning by establishing a collaborative ecosystem where financial expertise can be effectively explained and shared.

Funding: This research received no external funding.
Conflicts of Interest: The authors declare no conflict of interest.
Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers.

References

[1] Freddy Tapia et al., "From Monolithic Systems to Microservices: A Comparative Study of Performance," Applied Sciences, 2020. Available: https://www.mdpi.com/2076-3417/10/17/5797

[2] Volopay, "Financial modelling - Meaning, components, process, and best practices," 2025. Available: https://www.volopay.com/in/blog/financial-modelling/

[3] Nadia Ahmad and David Wishart, "Hybrid Approaches in Financial Time-Series Forecasting: Balancing Performance and Interpretability," ResearchGate, 2024. Available: https://www.researchgate.net/publication/382046735_Hybrid_Approaches_in_Financial_Time-Series_Forecasting_Balancing_Performance_and_Interpretability

[4] Martin Fowler, "Patterns of Enterprise Application Architecture," Addison-Wesley, 2012. Available: https://www.google.co.in/books/edition/Patterns_of_Enterprise_Application_Archi/vqTfNFDzzdIC?hl=en&gbpv=1&pg=PA1&printsec=frontcover

[5] Oluwatomisin Arokodare, "Leveraging Machine Learning for Enhanced Predictive Accuracy in Time Series Forecasting: A Comparative Analysis of LSTM and GRU Models," ResearchGate, 2024. Available: https://www.researchgate.net/publication/386461159_Leveraging_Machine_Learning_for_Enhanced_Predictive_Accuracy_in_Time_Series_Forecasting_A_Comparative_Analysis_of_LSTM_and_GRU_Models

[6] Joan Giner-Miguelez et al., "A domain-specific language for describing machine learning datasets," Journal of Computer Languages, 2023. Available: https://www.sciencedirect.com/science/article/pii/S2590118423000199

[7] Bolu Onioluwa et al., "Implementing Differential Privacy in Financial Institutions," ResearchGate, 2024. Available: https://www.researchgate.net/publication/386573420_Implementing_Differential_Privacy_in_Financial_Institutions

[8] LinkedIn, "What are the most important steps to designing a secure FinTech architecture?" Available: https://www.linkedin.com/advice/0/what-most-important-steps-designing-secure-41h1e

[9] Kelly Bailey, "Financial Benchmarking: Comparing Company Performance Relative to Peers," Corporate Finance Institute, Available: https://corporatefinanceinstitute.com/resources/valuation/financial-benchmarking/

[10] Busayo John Omopariola and Veronica Aboaba, "Comparative Analysis of Financial Models: Assessing Efficiency, Risk, and Sustainability," ResearchGate, 2019. Available: https://www.researchgate.net/publication/390761186_Comparative_Analysis_of_Financial_Models_Assessing_Efficiency_Risk_and_Sustainability