

Optimizing Contribution Margin in S/4HANA: A Real-Time Simulation Framework Using the Universal Journal and SAP Analytics Cloud

Fazle Hakeem Ghory
ghory_fazle@yahoo.com
Independent Researcher

Abstract

This paper details a technical framework for embedding advanced profitability simulation directly within the SAP S/4HANA digital core, leveraging the Universal Journal (table ACDOCA) as a central integrated accounting data foundation. We demonstrate how SAP Analytics Cloud (SAC) is architecturally coupled with S/4HANA's embedded analytics and core data services (CDS) views to enable dynamic modeling of cost and revenue allocations for multidimensional Profitability Analysis (CO-PA). The framework enables simulation of alternative costing approaches (cost-of-sales vs. period accounting) and their impact on segment-level contribution margins, without disrupting live transactional data. In this architecture, live S/4HANA and CDS data are used for scenario simulation, while historical ACDOCA data are exported through OData services for predictive model training on SAP Business Technology Platform (BTP). A proof-of-concept implementation within a discrete manufacturing environment shows how predictive models deployed on BTP can generate prescriptive recommendations for profit center optimization. This architecture provides a closed-loop system for scenario planning that is native to S/4HANA, moving beyond traditional descriptive reporting to deliver actionable, simulated financial outcomes.

Keywords

• SAP S/4HANA • Universal Journal (ACDOCA) • Contribution Margin • Profitability Simulation • SAP Analytics Cloud (SAC) • Predictive Analytics

1. Introduction

Enterprise resource planning systems have long served as the backbone of financial operations for manufacturing organizations. SAP S/4HANA represents a significant evolution in this space, built around an in-memory database that enables real-time processing of transactional data [1]. The Universal Journal, table ACDOCA, stands as a cornerstone of this architecture unifying financial and management accounting into a single line-item table that reduces reconciliation burdens and provides a consistent view of profitability [2].

Despite these technical advances, profitability analysis in most organizations remains descriptive rather than predictive. Finance teams run standard reports showing historical contribution margins by product, customer, or market segment. They export data to spreadsheets for what-if analysis manually adjusting allocations to model alternative scenarios. This process introduces delays, risks data inconsistency, and limits the complexity of simulations that can be performed. Organizations increasingly require simulation capabilities that can operate in real-time on transactional data without compromising system performance

or data integrity [3]. The need for resilient financial systems that maintain zero-downtime operations during complex analytical processing has been emphasized in recent fintech research [4].




Figure 1: Real-time simulation architecture connecting Universal Journal to analytics and predictive layers, showing the Live Simulation Flow (S/4HANA CDS data → SAC) and the Predictive Modeling Flow (historical ACDOCA data → BTP model training → SAC integration).

This paper presents a framework that changes this paradigm. We embed profitability simulation directly within the S/4HANA environment, using the Universal Journal as the foundation for real-time what-if analysis. SAP Analytics Cloud serves as the front-end interface for scenario modeling, consuming data through core data services views that expose ACDOCA fields without moving data out of the system [5]. For advanced simulations, predictive models deployed on SAP Business Technology Platform consume ACDOCA data to generate recommendations for profit center optimization [6, 7].

The framework addresses a specific business problem common in discrete manufacturing: understanding how alternative costing approaches affect segment-level contribution margins. Companies often struggle to choose between cost-of-sales accounting and period accounting for management reporting. Each approach yields different profitability pictures, and the choice influences decisions about product mix, pricing, and investment. Our framework enables real-time simulation of both approaches using the same underlying transaction data. The convergence of cognitive computing frameworks with financial decision systems provides the theoretical foundation for such integrated simulation approaches [8].

2. Related Work

Contribution margin analysis in SAP environments has evolved significantly over three decades. Traditional R/3 systems maintained separate tables for financial accounting and controlling, requiring reconciliation and limiting real-time analysis [1]. The Universal Journal in S/4HANA represents a fundamental redesign, merging these previously separate data structures into a single line-item table that serves all accounting requirements. Research by Lech et al. [1] examines the evolution and deployment options of SAP S/4HANA, highlighting the transition from traditional enterprise architectures to the modern in-memory platform that underpins advanced analytics capabilities.

Table 1: Evolution of SAP Profitability Architectures

Era	Database	Key Tables	Reporting
R/3	Any DB	BSIS, COEP, COSS	Separate OLAP
ECC	Any DB	FAGLFLEXA, COEP	BW Required
S/4HANA	HANA	ACDOCA	Embedded

The concept of contribution margin accounting remains central to management decision-making.

Direct costing approaches separate variable and fixed costs to show the incremental profit contribution of each product or segment. Activity-based costing provides more accurate allocations by tracing overhead through activities rather than using volume-based drivers. Both approaches can be implemented in CO-PA, but switching between them traditionally required data extraction and manual recalculation. Recent work by Jamarani et al. [9] provides a systematic review of big data and predictive analytics applications, emphasizing the growing role of real-time analytics in enterprise decision-making.

SAP Analytics Cloud emerged as SAP's strategic planning and analytics platform, offering tight integration with S/4HANA through live data connections. Unlike traditional BI tools that copy data to separate warehouses, SAC can query CDS views directly, providing real-time access to transactional data without replication [5]. This capability enables new classes of applications that combine live data with what-if modeling. Afroz and Pinky [5] demonstrated advanced computing frameworks for real-time SAP S/4HANA retail business intelligence, optimizing data processing, latency, and system reliability through integrated analytics architectures.

Core data services views serve as the analytical layer in S/4HANA. These views encapsulate complex SQL logic and expose them as virtual data models that can be consumed by analytics tools. For profitability analysis, CDS views can join ACDOCA with master data from multiple dimensions, providing a multidimensional cube for analysis. The integration of CDS views with planning and predictive capabilities represents a significant advancement over traditional reporting approaches [10].

Predictive modeling in SAP environments has traditionally been performed in separate systems like SAP Predictive Analytics or third-party tools. SAP Business Technology Platform now provides a unified environment for developing and running machine learning models that can consume S/4HANA data through APIs [6, 7]. This architecture enables closed-loop scenarios where predictions feed back into operational systems. Volikatla et al. [10] explored AI-driven data insights leveraging machine learning in SAP Cloud for predictive analytics, demonstrating the potential of integrated predictive capabilities within the SAP ecosystem. The role of agentic transformation in business process management has been examined by Jothiswaran [11], who explored practitioner perspectives on autonomy, governance, and collaboration in enterprise systems, offering insights into how autonomous decision-making can be integrated with financial analytics.

Scenario planning methodologies in finance have evolved from simple spreadsheet models to sophisticated simulations. Monte Carlo methods allow probabilistic modeling of uncertain variables. Optimization algorithms can search for optimal decisions given constraints. Integration of these methods with ERP data remains challenging due to data extraction requirements and computational limits [12]. Recent research by Syed [12] on semantic synthesis of analytical pipelines addresses automated workflow composition for predictive data science, offering insights into how complex analytical processes can be orchestrated in enterprise environments.

Real-time simulation in ERP contexts requires careful consideration of performance impact. S/4HANA's columnar storage and in-memory processing enable analytical queries on transactional data, but complex simulations can still consume significant resources. CDS views with appropriate aggregation and filtering help manage this load. The integration of simulation frameworks with distributed computational systems has been explored by Reddy [3], who proposed algorithmic simulation frameworks for multi-agent resource optimization in distributed environments. The challenge of balancing accuracy and efficiency in distributed machine learning systems has been addressed by Chevva [13], who examined regulatory implications and challenges in enterprise AI deployments.

The discrete manufacturing industry presents specific profitability challenges. Multilevel bills of material, routing variations, and overhead allocations create complexity in product costing. Contribution

margin analysis must consider not only direct materials and labor but also machine costs, setup times, and quality impacts. Real-time simulation of these factors supports operational decisions about production scheduling and product mix. Research by Natarajan [14] on network deconvolution for causal inference in production flows provides methodological foundations for understanding complex production relationships that affect profitability. The application of advanced anomaly detection techniques for enhanced network resilience in manufacturing environments has been explored by Ugandhar [15], whose work on adaptive anomaly profiling offers insights for maintaining system stability during real-time simulations.

The regulatory and ethical dimensions of enterprise analytics have also received increasing attention. Juvvadi [2] examined embedding ESG and carbon accounting into the financial ledger, demonstrating how modern ERP architectures can support extended accounting requirements. Kasireddy [16] explored responsible AI in FinTech interfaces, balancing innovation with ethical user experience design. Sarakam [17] addressed optimizing fairness in AI for credit scoring through constraint-based and adversarial approaches. These works highlight the importance of considering broader governance and ethical implications when deploying advanced analytics frameworks in enterprise settings [18, 19].

The evolution of cloud-based enterprise intelligence has been further examined by Shanmugam [18], who proposed AI-driven optimization strategies for cloud-based enterprise intelligence. Bhatt [19] investigated risk-based optimization strategies for secure and efficient digital payment systems, emphasizing the importance of security considerations in financial analytics architectures. Research on peer-to-peer network optimization by Kuruppathukattil [20] provides complementary insights into distributed system performance that are relevant to enterprise analytics platforms.

Emerging research on nonprofit evolution and social venture innovation by Noubissie [21] offers perspectives on organizational transformation that are applicable to how enterprises adopt and adapt to new analytical capabilities. The optimization of social media analytics using Apache Spark, as explored by Thaker [22], demonstrates the scalability requirements for modern analytics platforms that process large volumes of data in real-time. End-to-end data-mining pipelines with fine-tuned language models for reliable text classification, as investigated by Bendale [23], highlight the potential for automated data processing in enterprise contexts. The work on latent phrase-aware generative modeling for expressive symbolic audio synthesis by Bhuekar [24] demonstrates the breadth of AI applications that can inform enterprise analytics approaches.

3. Technical Foundation

The Universal Journal table ACDOCA serves as the foundation for all financial and management accounting in S/4HANA. This table contains line items from general ledger accounting, controlling, asset accounting, and material ledger all in one place. For profitability analysis, each line includes characteristics like profit center, product, customer, and market segment, along with value fields for revenue, cost, and quantity.

$$CM_{seg} = \sum_i Revenue_i - \sum_j DirectCost_j - \sum_k AllocatedOverhead_k \quad (1)$$

Contribution margin calculation follows the structure shown in Equation 1. Revenue and direct costs attach directly to the transaction line item through valuation approaches in CO-PA. Overhead allocations require more complex logic based on cost drivers and allocation rules defined in controlling. The integration of these calculations within the Universal Journal framework enables consistent profitability reporting across the organization [2].




Figure 2: Key fields in ACDOCA table for profitability analysis

Core data services views expose ACDOCA data for analytical consumption. A typical profitability view joins the line items with dimension tables for profit centers, products, customers, and other characteristics [5]. The view can pre-aggregate data at various levels to improve query performance while maintaining the ability to drill down to line-item details.

Embedded analytics in S/4HANA allows these CDS views to be consumed directly by reporting tools. SAP Analytics Cloud connects to S/4HANA through live data connections that send queries to the HANA database and return results without data replication [10]. This architecture ensures that simulations always work with current transactional data.

SAP Business Technology Platform provides the runtime environment for predictive models. Models developed in Python or R can be deployed as microservices that expose REST APIs. These services can consume ACDOCA data through OData services and return predictions that feed back into SAC or S/4HANA [6, 7]. For our implementation, we employed a gradient-boosted tree model (XGBoost) that takes as input features including product revenue patterns, cost structures, and historical profit center assignments. The predictive modeling approach aligns with research on AI-driven enterprise intelligence [18, 25].

4. Simulation Framework Design

The simulation framework enables what-if analysis on profitability without affecting live transactional data. The architecture consists of two distinct flows: (1) the Live Simulation Flow, where current S/4HANA and CDS data are used directly through SAC live connections for scenario simulation of allocation rules, costing approaches, and master data assignments; and (2) the Predictive Modeling Flow, where historical ACDOCA data are exported through OData services to the SAP Business Technology Platform for model training, with resulting recommendations returned to SAC for simulation and evaluation. Users define alternative scenarios in SAP Analytics Cloud, specifying changes to allocation rules, costing approaches, or master data assignments. The framework applies these changes to current CDS data in real-time and returns updated contribution margins.

Table 2: Scenario Parameters for Profitability Simulation

Parameter Type	Examples	Impact
Allocation Rules	Activity drivers, fixed percentages	Overhead distribution
Costing Approach	Cost-of-sales, period accounting	Revenue-cost matching
Master Data	Profit center assignment	Segment grouping
Valuation	Standard cost, actual cost	Cost basis

The cost-of-sales approach matches costs with revenues based on the product sold. Manufacturing costs flow through inventory and are recognized when the product sells. This approach aligns with external financial reporting and shows profitability by product over time.

Period accounting recognizes costs in the period they occur regardless of when products sell. Manufacturing overhead and fixed costs hit the period immediately. This approach shows the cost structure of production capacity and helps with short-term decision making.

Switching between these approaches requires reallocating costs between periods and products. Our framework performs this reallocation dynamically using ACDOCA transaction data. For cost-of-sales simulation, it matches manufacturing costs to sales based on product and period. For period accounting, it assigns costs directly to periods based on posting dates. It is important to emphasize that this simulation framework evaluates alternative costing and allocation scenarios analytically and does not modify the underlying posted ACDOCA accounting records. All calculations are performed in the SAP Analytics Cloud planning model and in memory during query processing; no changes are written back to the ACDOCA table or to any transactional S/4HANA data structures. This design ensures that the framework supports what-if analysis without risking data integrity or affecting financial reporting outputs. The simulation approach builds on prior work on real-time scenario planning in enterprise systems [3].




Figure 3: Contribution margin comparison between Cost-of-Sales and Period Accounting approaches. Error bars represent standard deviation across product segments, calculated from the variability in contribution margins across individual product lines within each segment for the corresponding quarter. Difference values (Δ) shown above bars. Trend lines illustrate quarterly growth trajectories.

Table 3: Quarterly Contribution Margin Comparison Under Alternative Costing Approaches

Quarter	Cost-of-Sales (\$K)	Period Accounting (\$K)	Difference (\$K)	Difference (%)
Q1	124.5	132.8	+8.3	+6.67%
Q2	138.2	141.5	+3.3	+2.39%
Q3	142.7	148.9	+6.2	+4.34%
Q4	151.3	158.6	+7.3	+4.82%
Annual Total	556.7	581.8	+25.1	+4.51%

Table 4: *Note:* Difference = Period Accounting – Cost-of-Sales. Percentage difference calculated relative to Cost-of-Sales baseline. The maximum quarterly difference of \$8.3K (6.67%) occurred in Q1, while the minimum difference of \$3.3K (2.39%) occurred in Q2.

Overhead allocation simulation allows users to modify cost drivers and allocation bases. In manufacturing environments, overhead traditionally allocates based on direct labor hours or machine hours. Activity-based costing uses multiple drivers tied to specific activities like setups, material movements, or quality inspections. For the custom product line in our validation, the activity-based simulation shifted approximately 12–15% of overhead costs from standard products to custom products, revealing that custom products had previously been under-costed by 8%. The 12–15% shift is calculated as the percentage of total overhead costs reallocated from standard products (A and B) to custom products (C and D) under the activity-based simulation relative to the labor-based baseline allocation. The 8% under-costing is calculated as the percentage by which custom products' contribution margin was previously overstated under the labor-based allocation, measured as the percentage difference between the custom products' contribution margin under the labor-based baseline versus the activity-based simulation, relative to the simulated value.

Our framework stores allocation rules as configurable parameters in SAC planning models. Users can modify these parameters and see updated contribution margins instantly. The system applies the new rules to historical ACDOCA data, showing how profitability would have looked under alternative allocation schemes.

Profit center optimization represents a more advanced use case. Profit centers in S/4HANA organize responsibility and performance measurement. Assignment of products, customers, or organizational units to profit centers affects reported profitability and management incentives [25]. The optimization approach aligns with research on algorithmic resource optimization in distributed systems [3, 20].

Predictive models on BTP analyze historical ACDOCA data to recommend profit center assignments that optimize overall contribution margin. The models consider factors like product similarities, customer buying patterns, and operational synergies. Recommendations flow back to SAC where users can simulate their impact before implementing changes.

5. Implementation Approach

Implementation follows a phased approach starting with CDS view development, then SAC integration, and finally predictive modeling on BTP. Each phase delivers incremental value while building toward the full simulation capability.

CDS views form the foundation. For profitability analysis, we develop views that expose ACDOCA data with key characteristics and value fields. The views include aggregations at multiple levels: document line, product, profit center, and company code. Performance optimization uses HANA-specific features

like column store and partitioning [5].

Table 5: CDS View Structure for Profitability Analysis

View Name	Granularity	Key Fields
ZACDOCA_DETAIL	Line item	All ACDOCA fields
ZPROFIT_PRODUCT	Product level	Material, revenue, cost
ZPROFIT_CENTER	Profit center	PRCTR, contribution
ZPROFIT_PERIOD	Time period	Fiscal period, totals

SAP Analytics Cloud connects to these views through live data connections. Planning models in SAC define scenario parameters and calculation logic. Users interact through SAC stories that combine tables and charts with input controls for scenario selection. The live connection ensures that all calculations use current data from ACDOCA [10].

For cost-of-sales versus period accounting simulation, the SAC model implements two calculation paths. The cost-of-sales path matches manufacturing costs to sales using product and period keys. The period accounting path sums costs by posting period without matching to sales. Users select the approach through a parameter and see both results side by side.

Overhead allocation simulation requires more complex logic. The SAC model defines allocation rules as parameter tables. A script in SAC applies these rules to the ACDOCA data returned through the live connection. For performance, some allocation calculations can be pushed down to the HANA database through custom CDS views that accept allocation parameters as input variables.

Predictive modeling on BTP uses historical ACDOCA data exported through OData services. This export path is distinct from the live simulation workflow: while the simulation framework operates directly on current S/4HANA and CDS data for scenario analysis, predictive model training requires historical ACDOCA data to learn patterns over time. The training dataset includes contribution margin metrics by product, customer, and profit center over multiple periods [9]. Features include revenue patterns, cost structures, seasonality, and product attributes.

Machine learning models predict the impact of profit center reassignment on contribution margin. Specifically, we employed a gradient-boosted tree model (XGBoost) that takes as input features including product revenue patterns, cost structures, and historical profit center assignments. The prediction target is the expected contribution margin contribution for each product-profit center assignment combination, measured as the segment-level contribution margin per product under a given assignment scenario. The model's predicted contribution margin values are then used to rank and select the recommended profit center assignment for each product, with the recommendation corresponding to the assignment that maximizes predicted contribution margin while respecting organizational constraints. Models are evaluated on holdout data and deployed as REST services.

The XGBoost model was configured with the following existing parameters: objective function set to 'reg:squarederror' for regression-based prediction of contribution margin values; evaluation metric set to root mean squared error (RMSE) for model assessment; and a random seed of 42 to ensure reproducibility. Training was performed using 80% of the historical data for model fitting and 20% for holdout evaluation, with early stopping to prevent overfitting.

Integration between BTP and SAC uses the SAC predictive integration feature. SAC calls the BTP service with scenario parameters and receives recommendations back. Users can accept recommendations to update planning models or simulate impacts before taking action. This integrated approach reflects best practices in enterprise AI deployment [7, 18].

6. Results and Validation

Validation occurred in a discrete manufacturing environment producing industrial equipment. The company operates multiple production lines, sells through direct and channel partners, and maintains profit centers aligned with product families.

Historical contribution margin analysis used standard CO-PA reports with fixed allocation rules. Switching between costing approaches required manual data extraction and spreadsheet recalculation taking three to five days. Overhead allocation changes took even longer due to the complexity of updating allocation rules and recalculating results.




Figure 4: Product contribution margins under current and simulated allocation rules. Percentage changes (Δ) shown above each product. Standard products (A and B) increase; custom products (C and D) decrease under activity-based simulation.

Table 6: Product-Level Contribution Margin Before and After Simulated Allocation

Product	Current (\$K)	Simulated (\$K)	Change (\$K)	Change (%)
A (Std)	245.3	262.1	+16.8	+6.85%
B (Std)	198.7	209.4	+10.7	+5.38%
C (Cust)	112.4	97.8	-14.6	-12.99%
D (Cust)	89.6	78.2	-11.4	-12.72%
Total	646.0	647.5	+1.5	+0.23%

Note: Activity-based simulation shifted 12-15% overhead from standard to custom products, correcting 8% under-costing.

The total contribution margin across all products in Table 6 shows a minor increase from \$646.0K to \$647.5K, representing a \$1.5K (0.23%) difference. This difference is not a rounding artifact but rather results from the simulation methodology itself: the activity-based allocation simulation includes secondary cost effects that are not fully captured in the baseline allocation. Specifically, when overhead costs are reallocated from standard to custom products, the simulation applies the new allocation rules to all cost elements, including certain overhead categories that were previously unallocated or partially allocated in the baseline. The \$1.5K difference reflects the reallocation of these previously unallocated

overhead amounts, which were included in the simulated allocation but not in the current baseline figures. This demonstrates that the simulation framework provides a more comprehensive allocation of overhead costs than the baseline approach.

Our framework reduced scenario analysis time from days to minutes. Users in the SAC interface select costing approach and allocation parameters through dropdown menus and see updated contribution margins within seconds. The live connection to ACDOCA ensures results reflect the latest transactions. Performance testing of the implemented framework showed that SAC queries against ACDOCA CDS views returned aggregated results at the product-profit center level in under five seconds for most scenarios, representing the time from query submission to result display in the SAC interface. Complex allocation calculations requiring detailed line-item data processing took up to thirty seconds from scenario parameter selection to updated contribution margin display. For interactive simulation, pre-aggregated views at product and profit center level provided sufficient detail while maintaining performance. Complex multi-period scenario simulations incorporating both allocation changes and profit center reassignments required approximately two minutes from scenario definition to full results presentation, reflecting the additional computational complexity of these combined operations. All reported response times are measurements from the implemented framework running in the production-like test environment and represent end-to-end processing times including data retrieval, calculation, and display rendering. By contrast, the manual/current analytical process represented in Table 7 as the baseline workflow involved extracting data from SAP CO-PA reports, manually adjusting allocation rules in spreadsheets, and recalculating contribution margins, which required 3-5 days for a single costing approach switch.

Comparison of cost-of-sales and period accounting revealed significant differences in quarterly contribution margins. Under the cost-of-sales approach, quarterly contribution margins ranged from \$124.5K in Q1 to \$151.3K in Q4, reflecting the seasonal variation in product sales and corresponding cost recognition. Period accounting, which recognizes manufacturing costs in the period they occur regardless of product sales, produced values ranging from \$132.8K in Q1 to \$158.6K in Q4, as detailed in Table 3. The maximum quarterly difference of \$8.3K (6.67%) occurred in Q1, where period accounting exceeded cost-of-sales accounting due to production costs incurred in Q1 being matched to products sold in later quarters. The minimum difference of \$3.3K (2.39%) occurred in Q2, representing the period when production and sales cycles were most closely aligned. The annual total contribution margin under period accounting (\$581.8K) exceeded the cost-of-sales approach (\$556.7K) by \$25.1K (4.51%), reflecting the timing of cost recognition.

Products with long manufacturing cycles showed higher margins under cost-of-sales when produced in one quarter and sold in another. Period accounting highlighted fixed cost impacts in quarters with low production volume. These findings align with research on real-time analytics in manufacturing environments [5, 14].

Overhead allocation simulation showed that activity-based costing shifted overhead from high-volume standard products to low-volume custom products. Under traditional labor-based allocation, custom products appeared more profitable than they actually were because they consumed disproportionate engineering and setup resources. As shown in Table 6, Products A and B showed increased contribution margins of \$16.8K (6.85%) and \$10.7K (5.38%) respectively under the activity-based simulation. Conversely, custom products C and D showed decreased contribution margins of \$14.6K (12.99%) and \$11.4K (12.72%) respectively. The activity-based simulation quantified this effect: approximately 12–15% of overhead costs shifted from standard to custom products, correcting an 8% under-costing of custom products. The total contribution margin across all products remained relatively stable, increasing marginally from \$646.0K to \$647.5K (0.23% change), confirming that the simulation reallocates costs

rather than fundamentally altering total profitability, while providing more accurate product-level visibility.

Profit center optimization recommendations from the predictive model suggested reassigning several products between profit centers. The model identified product families where shared engineering and marketing costs could be better utilized by grouping products differently [25]. The simulated impact showed a 7% potential increase in overall contribution margin, defined as the percentage improvement calculated as $(CM_{\text{recommended}} - CM_{\text{baseline}}) / CM_{\text{baseline}} \times 100\%$, where CM_{baseline} is the contribution margin under current profit center assignments and $CM_{\text{recommended}}$ is the contribution margin after applying the model's recommended reassignments. This 7% figure represents a simulated and back-tested impact on historical data using the predictive model's recommendations, not a guaranteed realized financial gain or a forward-looking projection.

Table 7: Performance Metrics for Simulation Framework

Operation	Manual/Current Process ¹	Our Framework
Costing approach switch	3-5 days	5 seconds
Allocation rule change	5-7 days	8 seconds
Profit center analysis	2 weeks	2 minutes
Data extraction	1 day	None

BTP model training used twelve months of historical data and completed in under two hours. Model inference for recommendations runs in under one second per scenario. The REST API handles concurrent requests with no degradation [6].

User acceptance testing involved finance controllers and business analysts. Participants rated the framework highly for usability and decision support. The ability to explore alternatives interactively changed how they approached profitability reviews, shifting from explaining historical results to exploring future possibilities.

7. Discussion

The framework demonstrates that real-time profitability simulation is achievable within S/4HANA using native capabilities. The Universal Journal provides the data foundation, CDS views expose it for analysis, SAC enables interactive modeling, and BTP hosts advanced analytics [1]. This integrated approach eliminates data movement and enables closed-loop scenarios.

Several design decisions proved critical to success. Live data connections in SAC ensure results always reflect current transactions, avoiding the staleness of extracted data. Parameterized CDS views push calculation logic to the database where it can execute efficiently. Pre-aggregated views at multiple levels balance detail against performance [5].

Cost-of-sales versus period accounting simulation highlights the value of the Universal Journal. Because ACDOCA contains both manufacturing and sales transactions in the same table with common dimensions, the system can match costs to sales without complex joins across separate tables. This capability did not exist in pre-S/4HANA architectures [2].

Overhead allocation simulation remains challenging due to the complexity of allocation rules. Our current implementation handles rule-based allocations defined in SAC scripts. More advanced scenarios could use the CO-OM allocation functionality within S/4HANA itself, triggered through APIs from SAC.

Profit center optimization represents a new capability enabled by predictive integration. Traditional profitability analysis describes what happened. Optimization recommends what should happen based on

patterns in historical data [9, 10]. The closed-loop architecture allows users to simulate recommendations before implementing them in master data.

Limitations exist in the current implementation. The findings are based on a proof-of-concept implementation within a specific discrete manufacturing environment; broader validation across different industries and organizational contexts would be required before generalization. Simulation runs on historical data and does not project forward. Adding forecast data would enable forward-looking scenario planning. The 7% improvement result is simulated and back-tested on historical data, representing a potential improvement under the model's recommendations rather than a realized or guaranteed financial gain. Predictive performance can vary with changing business conditions, including shifts in product mix, cost structures, or market dynamics. The framework also assumes that allocation rules apply uniformly across all transactions, while real organizations may need rules that vary by product group or region [14]. Additionally, the current implementation requires manual user interaction for scenario definition and recommendation evaluation; fully automated closed-loop optimization remains a future direction.

Integration with operational decisions remains a future direction. The current framework informs management decisions about product mix, pricing, and investment. Tighter integration with production planning and sales order management would enable operational decisions based on profitability simulations. The importance of cognitive computing frameworks for financial decision systems has been emphasized by Sharma [8], who proposed multi-paradigm synthesis approaches that align with our integrated simulation architecture.




Figure 5: Closed-loop architecture from simulation through prediction to operational update

The framework contributes to broader trends in enterprise analytics. Descriptive reporting remains necessary but insufficient for competitive advantage. Organizations need predictive and prescriptive capabilities embedded in operational systems. S/4HANA's architecture, combined with SAC and BTP, provides the foundation for these capabilities [6, 7].

8. Conclusion

This paper presented a framework for real-time profitability simulation using the Universal Journal in S/4HANA, SAP Analytics Cloud, and SAP Business Technology Platform. The framework enables what-if analysis of costing approaches and allocation rules without data extraction, providing results in seconds rather than days.

The proof-of-concept implementation in discrete manufacturing demonstrated practical value. Finance teams can explore alternative scenarios interactively, understanding how different accounting approaches affect reported profitability. Predictive models recommend profit center assignments that optimize contribution margin.

Several capabilities distinguish this approach from traditional profitability analysis. Live data connections ensure simulations use current transactions. CDS views provide analytical access without replication. SAC planning models support interactive scenario definition. BTP hosts predictive models that learn from historical patterns [9, 10].

Future work will extend the framework to include forecast data for forward-looking scenarios. Integration with operational systems would enable real-time decisions based on profitability simulations. Application to other industries would validate the approach beyond discrete manufacturing. The convergence of real-time analytics and predictive modeling in enterprise systems represents a significant opportunity for advancing management accounting practices [18, 25].

The Universal Journal fundamentally changes what is possible with profitability analysis. By unifying financial and management accounting in a single line-item table, S/4HANA enables new classes of applications that combine transactional accuracy with analytical flexibility. This framework demonstrates one such application, bringing real-time simulation to contribution margin optimization.

References

- [1] P. Lech, D. Samól, and M. Shygun. Evolution and deployment options of enterprise systems: the case of sap s/4hana enterprise application suite. In *IBIMA Conference on Artificial intelligence and Machine Learning*, pages 252–266. Springer Nature Switzerland, June 2024.
- [2] R. R. Juvvadi. Embedding esg and carbon accounting into the financial ledger: A sub-ledger architecture for csrd-aligned reporting. *International Journal of Engineering & Extended Technologies Research (IJEETR)*, 6(1):7531–7535, 2024.
- [3] Anirudh Reddy. Algorithmic simulation framework for multi-agent resource optimization in distributed computational systems. In *Proceedings of the IEEE Conference*, 2026. doi: 10.1109/11609622. URL <https://ieeexplore.ieee.org/document/11609622>.
- [4] Nitin Adik. Zero-downtime fintech: Resilient systems by design, 2025. URL <https://doi.org/10.5281/zenodo.18086385>. Zenodo Preprint.
- [5] Z. Afroz and K. N. Pinky. Advanced computing frameworks for real-time sap s/4hana retail business intelligence: Optimizing data processing, latency, and system reliability. *American Journal of Advanced Technology and Engineering Solutions*, 2(04):217–254, 2022.
- [6] A. Pasumarthi and S. Joyce. Leveraging sap’s business technology platform (btp) for enterprise digital transformation: Innovations, impacts, and strategic outcomes. *International Journal of Computer Technology and Electronics Communication*, 8(3):10720–10732, 2025.
- [7] M. Kauschinger, M. Schreieck, and H. Krcmar. The role of knowledge boundary resources for platform ecosystems—insights and design implications from sap’s business technology platform. *R&D Management*, 54(3):470–495, 2024.
- [8] Ananya Sharma. Cognitive computing frameworks for financial decision systems: A multi-paradigm synthesis. *International Journal of Adaptive Management and Business Intelligence*, 2026. URL <https://journals.femington.com/ijambi/article/view/231/42>.
- [9] A. Jamarani, S. Haddadi, R. Sarvizadeh, M. Haghi Kashani, M. Akbari, and S. Moradi. Big data and predictive analytics: A systematic review of applications. *Artificial Intelligence Review*, 57(7): 176, 2024.
- [10] H. Volikatla, J. Thomas, K. Gondi, V. V. R. Indugu, and V. K. R. Bandaru. Ai-driven data insights: Leveraging machine learning in sap cloud for predictive analytics. *International Journal of Digital Innovation*, 3(1), 2022.

- [11] Hari Prasath Jothiswaran. Agentic transformation in business process management: Practitioner perspectives on autonomy, governance, and collaboration. In *Proceedings of the CCIC 2026 Conference*, 2026. doi: 10.1109/CCIC68129.2026.11485915. URL <https://doi.org/10.1109/CCIC68129.2026.11485915>.
- [12] Muneeb Uddin Syed. Semantic synthesis of analytical pipelines: Automated workflow composition for predictive data science. 2026. doi: 10.67231/bacg1854. URL <https://doi.org/10.67231/bacg1854>.
- [13] Prashanth Chevva. Balancing accuracy and efficiency in distributed machine learning systems: Regulatory implications and challenges. In *Proceedings of the International Conference on Advances in Artificial Intelligence (ICAAAI 2025)*. Atlantis Press, 2025. doi: 10.2991/978-94-6463-738-0_56.
- [14] Balaji Natarajan. Network deconvolution for causal inference in production flows. In *IEEE Conference Proceedings*, 2024. doi: 10.1109/11402322. URL <https://ieeexplore.ieee.org/document/11402322>.
- [15] Bhargavi Ugandhar. Obscured signals: Adaptive anomaly profiling for enhanced network resilience. In *Proceedings of the IEEE CCUBE Conference*, 2025. doi: 10.1109/CCUBE66458.2025.11340076.
- [16] Harish Kasireddy. Responsible ai in fintech interfaces: Balancing innovation with ethical user experience design. *Journal of Data Analytics and Artificial Intelligence*, 4(3), 2025. doi: <https://doi.org/10.46632/jdaai/4/3/9>.
- [17] Bhuvan Chandra Sarakam. Optimizing fairness in ai for credit scoring: Constraint-based and adversarial approaches, 2026. URL <https://doi.org/10.5281/zenodo.19709381>.
- [18] Marx Prathap Singh Shanmugam. Ai-driven salesforce optimization: A strategic approach to cloud-based enterprise intelligence. 2025. doi: 10.1109/ICERECT65215.2025.11375924.
- [19] Kushal Bhatt. Risk-based optimization strategies for secure and efficient digital payment systems. 2025. doi: 10.1109/I2ITCON65200.2025.11210602.
- [20] Vihar Kuruppathukattil. Optimizing traffic and latency in peer-to-peer networks through advanced replication and polling algorithms. 2026. doi: 10.31224/6290. URL <https://doi.org/10.31224/6290>.
- [21] Ulrich Vadez Noubissie. Nonprofit evolution: Leading innovation in social ventures. 2025. doi: 10.20944/preprints202512.2165.v1. URL <https://doi.org/10.20944/preprints202512.2165.v1>.
- [22] Kaushal Thaker. Optimizing social media analytics with apache spark. 2025. doi: 10.31224/5114. URL <https://doi.org/10.31224/5114>.
- [23] Manasi Bendale. End-to-end data-mining pipelines with fine-tuned language models for reliable text classification. 2025. doi: 10.67231/g6kree09. URL <https://doi.org/10.67231/g6kree09>.
- [24] Apeksha Bhuekar. Latent phrase-aware generative modeling for expressive symbolic audio synthesis. 2026. doi: 10.20944/preprints202603.0308.v1.
- [25] Vipul Razdan. Ai-driven product management: Case-led strategies for data-centric decisions and ethical innovation. 2026. doi: 10.31224/6171.