Future of **Payments**

JOURNAL OF
FINANCIAL SERVICES
PUBLISHED BY: **PROJECTIVE GROUP INSTITUTE**

01
2025

Projective
INSTITUTE

JOURNAL OF FINANCIAL SERVICES

Future of Payments

Contents

8 **A launchpad for Payments innovation: the renewed RTGS service**
Victoria Cleland, Chief Cashier and Executive Director for Payments, Bank of England

16 **Payments are in change: where is the journey heading?**
Carlos Nasher, Managing Partner, Thede Consulting (part of Projective Group)
Robert-Jan Wekking, Group Practice Lead Payments, Projective Group

32 **Smarter payment algorithms for liquidity-saving in RTGS: a practitioner's guide**
Jean-Paul Lam, Associate Professor of Economics, University of Waterloo, and Chief AI of Research, Goodlabs Studio
Thomas Lo, CEO and Co-Founder, GoodLabs Studio
Donald McGillivray, Lead Applied AI Developer · GoodLabs Studio
Chris McMahon, Director, AI & Emerging Technologies Research, GoodLabs Studio
Alex Ostrovsky, AI Advisor, GoodLabs Studio

42 **Stablecoins: a new technology, or older than the Knights of Templar?**
Paul H. Kupiec, Senior Fellow and Arthur F. Burns Chair in Financial Policy, American Enterprise Institute

46 **Impact of new technologies on the cash management competitive landscape**
Koen Vierendeels, Head of Transaction Banking, Belfius
Robert-Jan Wekking, Group Practice Lead Payments, Projective Group

54 **The evolution of customer interaction and future expectations**
Jorissa Neutelings, Chief Digital Officer, ABN AMRO Bank N.V.

58 **Issuing and investing in tokenized financial assets: what to look out for and why**
Phoebus L. Athanassiou, Senior Lead Legal Counsel, European Central Bank, and Associate Professor, Goethe Universität, Frankfurt am Main

66 **Cross-border Payments beyond 2027: how to crack speedy progress**
Magali Van Bulck, Head of Public Policy (EMEA), Wise

72 **The future operating model for Payments**
Stephen Peters, Head of Enterprise Payment Solutions, FIS
Alan Verschoyle-King, U.K. Payments Practice Lead, Projective Group

86 **Regulating retail Payments: balancing the competing interest**
Emanuel van Praag, Professor of Financial Technology (FinTech) and Law, Erasmus School of Law, and Counsel, Kennedy Van der Laan

98 **From bottleneck to breakthrough: transforming payment investigations into competitive advantage**
Galitchenko Iouri, Global Products and Markets Manager, SWIFT

102 **Towards European sovereignty in retail Payments**
Diederik Bruggink, Senior Director, Payments, Digital Finance and Innovation, European Savings and Retail Banking Group (ESBG)
Douglas Lockhart, Payments and Digital Finance Advisor, European Savings and Retail Banking Group (ESBG)

110 **The future of Payments: quo vadis Europe?**
Christophe Bonte, Senior Policy Adviser, European Banking Federation

120 **Tackling authorized push payment (APP) fraud: a challenging and evolving Payments environment**
Piers Reynolds, Partner, Freshfields LLP
Laura Feldman, Barrister, Freshfields LLP, and Lecturer in Law, University of Oxford

128 **Challenges for banks: e-wallets and the application of strong customer authentication**
Jan Jans, Partner, RegCounsel Financial Services

136 **Data management and its impact on industry**
John A. Bottega, President, EDM Association

142 **Cashless future - reality or fiction?**
Tobias Trütsch, Managing Director and Head of the Center for Financial Services Innovation (FSI-HSG), University of St. Gallen (HSG), Switzerland

148 **Preparing for October 2025: implementation and impact of the EPC's Verification of Payee scheme**
Dominique Allebroeck, Product Manager, The European Payments Council

154 **Protecting self-hosted payment systems against ransomware: a comprehensive guide**
Buck Rogers, Professor in Cyber Security, University of Gloucestershire, and Cyber Advisor, Projective Group
Scott Beange, Head of Data Management and Cyber Strategy, Projective Group

162 **Transforming cross-border Payments: navigating technological innovation, geopolitical fragmentation, and auditing challenges**
David S. Krause, Emeritus Associate Professor of Finance, Marquette University
Eric P. Krause, Assistant Professor of Accounting, Iowa State University

176 **CBDCs and the digital euro: AML compliance moves closer to the end-user**
Christian Sillaber, Senior Researcher, Department of Private Law, University Bern

Dear reader,

Welcome to this, the inaugural edition of the Projective Group Institute's new Journal of Financial Services!

We're very excited to be launching our Journal at a time of immense challenge and change for the financial services industry, and we've drawn upon the deep subject matter expertise of the many renowned contributors to this edition of the Journal to bring you what we hope you'll find to be a highly stimulating series of articles.

We'll be publishing editions of the Journal on a regular basis, and each time we will be focusing on those topics our Advisory Board members, who are leading academics and executives, are telling us are of key strategic importance to them. And it's for that reason that we've chosen to focus this edition on the world of Payments.

Payments has gone from being an often-overlooked, quite slow-changing, typically back-office, function to becoming one of the most transformed - and transformative – functions in the financial services industry. And all of that has happened within the space of only a few years. Digitization, cloud native technology, financial crime and fraud, artificial intelligence, and many other recent developments, combined with ever-greater regulatory scrutiny and ever-increasing client expectations, have created the perfect storm for legacy Payments providers and have opened this area up to newer and more agile competitors.

We focus on the drivers of that transformation, the implications of those drivers for the various players in the Payments value-chain, and the key decision points firms will therefore soon face if they intend to remain leaders in this field.

I'm conscious that in creating the articles in this Journal, many individuals have spent many hours researching and writing, and I'm very grateful for their time and expertise!

I look forward to getting your feedback on our new Journal of Financial Services, and to speaking with you again soon.

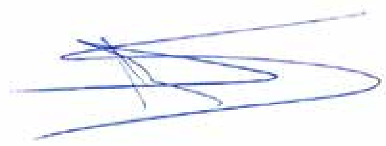

Stefan Dierckx, Founder & CEO, Projective Group

Future of Payments

Smarter payment algorithms for liquidity-saving in RTGS: a practitioner's guide

Jean-Paul Lam, Associate Professor of Economics, University of Waterloo, and Chief AI of Research, Goodlabs Studio

Thomas Lo, CEO and Co-Founder, GoodLabs Studio

Donald McGillivray, Lead Applied AI Developer, GoodLabs Studio

Chris McMahon, Director, AI & Emerging Technologies Research, GoodLabs Studio

Alex Ostrovsky, AI Advisor, GoodLabs Studio

ABSTRACT

Real-time gross settlement (RTGS) systems are the foundation of modern financial infrastructure, yet many still rely on outdated and liquidity-intensive algorithms. This paper offers a practical guide to improving liquidity efficiency through more innovative payment algorithms, with a focus on liquidity-saving mechanisms (LSMs). We survey recent innovations in queue management, graph-based methods, and AI-driven techniques, including both classical and quantum-inspired optimization. Drawing on our own research and simulations, we demonstrate how non-invasive, real-time reordering of payments can reduce peak liquidity needs without compromising system safety or speed. These insights are relevant not only for central banks and financial institutions, but also to a wide range of sectors, including logistics, operations, and resource-constrained environments.

1. Introduction

Real-time gross settlement (RTGS) systems are designed to settle high-value transactions instantly, with finality and security. However, this comes with a cost as transactions are settled one by one in real time, which often requires participants to hold large amounts of intraday liquidity. The challenge is balancing these objectives: fast, risk-free settlement on the one hand, and efficient liquidity usage on the other. Most existing RTGS designs prioritize security and immediacy, but this comes at the cost of heavy liquidity consumption.

RTGS around the world settle trillions of dollars every day. The European Central Bank's (ECB) T2 system processes over €1.8 trillion daily across more than 420,000 transactions [ECB (2024)]. In the U.K., the Bank of England's CHAPS system settles an average of £344 billion each day, with about 208,000 payments [BoE (2024a)]. In Canada, the LYNX system processed CAD 386 billion per day on average in 2024, across nearly 55,000 transactions [Payments Canada (2024)]. Despite their central role, most RTGS systems still rely on first-in, first-out (FIFO) settlement logic, which prioritizes processing speed over liquidity efficiency. Although simple, the FIFO method can waste liquidity and trigger delays that spread across the financial system. These frictions are especially costly as payment volumes continue to grow and regulatory capital pressures intensify.

The consequences of inefficient liquidity allocation go beyond balance sheets. Banks maintain billions in intraday liquidity buffers, and with prevailing interest rates, the annual opportunity costs can exceed hundreds of millions of dollars per institution. Even more concerning, delays caused by rigid settlement logic, rather than system capacity or risk, can ripple through the system. Delays by large institutions propagate through the network, leading smaller participants to face increased uncertainty. At the same time, the broader system becomes more vulnerable to liquidity shocks and gridlocks [Bech and Garratt (2003)]. These dynamics are especially dangerous under stress, when modest frictions can lead to broader instability and increase the risk of payment failures and systemic default.

The regulatory environment has only intensified these pressures. Basel III's Liquidity Coverage Ratio (LCR) and Net Stable Funding Ratio (NSFR) rules have made liquidity more expensive to hold and more closely monitored. Simultaneously, central banks are demanding both faster settlement and greater system resilience. The apparent trade-off between speed and safety has renewed interest in advanced liquidity-saving mechanisms (LSMs) that reorder payments to reduce liquidity needs while ensuring timely settlement [Byck and Heijmans (2020), Davey and Gray (2014)].

LSMs help institutions manage liquidity more effectively, and thus reduce capital needs, and operational risk. Moreover, they improve the predictability of payment flows, reduce last-minute funding pressures, and allow institutions to offer more reliable services to corporate clients. Banks that modernize in this area are more resilient in times of market stress. Those that do not may face higher costs and fall behind on regulatory and client expectations.

Central banks and payment system operators worldwide are actively implementing enhanced settlement algorithms to improve liquidity efficiency. For example, the Bank of England has argued that smarter algorithms can more efficiently match incoming and outgoing payments, helping to lower the amount of intraday liquidity needed to settle payment obligations [BoE (2024b)]. Similarly, modern payment systems, such as Singapore's FAST system, have shown that rapid settlement capabilities can coexist with sophisticated payment processing, challenging the notion that speed, and efficiency are trade-offs.

Despite growing interests, implementing advanced payment algorithms is challenging, especially in a regulated and risk-averse banking environment. Algorithms must be transparent and easy to understand, not just practical. Their goals must align with both business needs and regulatory rules. They also need to fit smoothly into systems that process trillions in payments each day, without causing any disruption. Most importantly, their decisions must be explainable, auditable, and subject to human oversight to ensure accountability.

A major obstacle in adopting more advanced payment algorithms is complexity. Liquidity efficiency depends not only on the amount of liquidity available, but also on when payments are sent and in what order they are processed. The space of possible payment ordering in a high-volume RTGS system is vast, and identifying the optimal sequence is a computationally intractable problem. As shown by McMahon et al. (2024), the payment reordering problem is an NP-hard problem, similar in complexity to the well-known Travelling Salesman problem [Lawler (1985)], meaning that as the number of payments increases, the time required to find the optimal sequence grows exponentially. In this case, no known algorithm can guarantee an exact solution within a reasonable time frame. This is why payment systems typically rely on heuristics and rule-based methods that produce good and quick results to meet real-time processing requirements. However, these rule-based mechanisms and simple heuristics, while practical, are not necessarily optimal solutions and often leave substantial efficiency gains untapped [Galbiati and Soramäki (2010), Rivadeneyra and Zhang (2022)].

That said, technological advances have created unprecedented opportunities to modernize RTGS operations. Contemporary systems can perform real-time optimization calculations, assess systemic impacts, and implement complex netting logic that would have been computationally prohibitive just a decade ago. Machine learning models can predict payment behavior with increasing accuracy, while graph-based algorithms can identify settlement sequences that minimize peak liquidity requirements. Recent research, including our work on quantum-inspired optimization, has further expanded the frontier by demonstrating how novel computing approaches can solve complex payment scheduling problems more efficiently than traditional methods.

In our recent Management Science publication [McMahon et al. (2024)], we demonstrated how quantum-inspired optimization techniques can be applied to improve liquidity efficiency in RTGS systems. We formulated the payment settlement problem as a mathematical model (Quadratic Unconstrained Binary Optimization), which represents decisions as binary variables and seeks the optimal combination to minimize liquidity usage. We evaluated our solution using both quantum annealing hardware and quantum-inspired classical solvers using Canadian payment data. Our results showed that, while current quantum hardware remains limited in scale and noise, quantum-inspired solvers outperformed traditional FIFO and heuristic approaches across a range of scenarios and simulations using historical payment data.

In scenarios where batches of 70 payments arrived simultaneously, the quantum-inspired solvers achieved an average liquidity savings of 1.36% per day compared to FIFO processing. This translated to an average daily reduction of 240 million Canadian dollars in liquidity usage, with a median savings of 128.8 million CAD. Despite the complexity, the optimization process remained fast enough to operate within the real-time constraints of RTGS systems, showing that these approaches can be practically deployed in a live environment.

Building on these methods, we have since extended our work and developed a broader set of classical algorithms for liquidity optimization, including greedy heuristics, virtual liquidity models, and network-aware prioritization. Using both synthetic and anonymized historical data, we benchmarked these methods against traditional FIFO rules. We designed our simulations to mirror real-world RTGS conditions and showed that these algorithms can significantly reduce the liquidity usage and cost of financial institutions without affecting settlement speed. The results highlight the practical benefits of our quantum-inspired approaches in improving the efficiency of high-value payment systems.

This paper provides financial professionals with a practical guide to understanding and applying advanced LSMs in RTGS systems. We focus on approaches driven by advanced computational techniques, including both classical optimization techniques and quantum-inspired methods. As RTGS platforms continue to evolve, institutions that adopt smarter liquidity management strategies will not only improve their own performance but also contribute to the resilience of the broader stability of the financial system.

2. Evolving payment algorithms: a review of modern approaches

This section outlines key developments in modern payment algorithms, focusing on the main directions the literature has taken in recent years. Rather than offering a comprehensive survey, the goal is to highlight the most influential approaches shaping current thinking. These algorithms reflect different priorities, such as speed, liquidity efficiency, and risk control. A good understanding of these methods can help financial institutions choose the tools that best suit their operations and risk profile.

2.1 Queue management and prioritization algorithms

Modern queue management has moved beyond simple FIFO rules. Many now use some form of real-time and dynamic prioritization, adjusting the order of payments as conditions change [Chen et al. (2025), BoE (2022)]. Dynamic priority scheduling algorithms adjust the payment processing order in real time based on multiple factors such as payment urgency, participants' liquidity positions, and the potential for creating settlement opportunities. Many high-value payment systems allow payments to be assigned priority categories (e.g., urgent versus normal) and even re-ordered in the central queue to improve efficiency [Bank of Canada (2022)]. One common approach is to assign a score to each payment based on how long it has been waiting, whether the sender is tight on liquidity, or if it plays a key role in the flow of funds. A high-value, time-sensitive payment might be prioritized over smaller ones, but only if it does not risk causing gridlock elsewhere in the system. One of the main challenges is to calculate these scores efficiently to maintain the stability of the system, while guaranteeing that the payment flow remains transparent and easy to audit.

An important improvement in modern payment systems is early gridlock detection [BoE (2024b), Garrat et al. (2023), Bech and Soramäki (2001)]. In the past, systems would wait until gridlock had already occurred, when payments could not move, before taking action. That often meant delays and wasted liquidity. Newer algorithms utilize graph theory to catch problems before they happen. In these systems, the payment network works as a directed graph, where banks are represented as nodes and payment instructions are represented as edges. The system scans the graph for circular dependencies, that is, loops where payments are stuck waiting on each other. When it spots one, it can reshuffle priorities or proactively activate liquidity-saving tools to prevent gridlock formation.

Some systems are beginning to utilize forecasting tools to help in determining which payments to process first. Instead of reacting only to what's in the queue at that moment, these algorithms consider expected payment flows and liquidity inflows when determining optimal processing sequences. That means a payment might be held briefly if the system expects another one to come in, which would make settlement more efficient. Machine learning can help with these predictions by analyzing historical patterns in payment timing and behavior [Vishnevsky and Gorbunova (2022)]. Using these methods in a live setting, however, is not straightforward, as it requires careful testing to avoid errors that could disrupt the system or unfairly delay some payments.

2.2 Optimization-based approaches

Researchers have employed various mathematical optimization techniques to improve liquidity efficiency in payment systems [McMahon et al. (2024)]. One approach, called Mixed-Integer Programming (MIP), treats the settlement process as a mathematical problem. The goal is to settle as many payments as possible using the least amount of liquidity, while still respecting rules around payment priorities, participant limits, and regulatory requirements [Cambe and Xing (2023)].

The problem is that optimization algorithms face a fundamental trade-off: finding the mathematically best solution takes time, but payment systems often cannot wait. MIP solvers can find optimal solutions, but solving large batches of payments may take several minutes, far too long when payments must clear immediately. Most practical systems, therefore, sacrifice mathematical perfection for speed, using faster heuristic methods that deliver solid results within the tight timing constraints that payment processing demands. Fortunately, recent advances in solver technology and hardware (like GPUs) are making real-time optimization more practical for larger problems.

Not all payment algorithms aim to minimize liquidity usage. Some are designed simply to ensure that all payments are settled on time and in line with system rules. These so-called constraint satisfaction algorithms are designed to ensure that every payment clears correctly and on time while meeting all the necessary conditions, such as timing rules, participant-specific limits, or regulatory requirements. This approach is especially useful when the focus is on reliability and compliance rather than efficiency.

In many real-world systems, though, there is more than one goal. Payment platforms often need to balance multiple priorities at once: using less liquidity, speeding up settlements, maintaining fairness across participants, and keeping decision-making transparent. That is where multi-objective optimization comes in. These methods do not try to find a single best solution. Instead, they produce a set of trade-offs called Pareto-optimal solutions, where no one objective can be improved without making another worse. This gives system operators a menu of viable options, letting them choose the right balance for the moment based on business goals or market conditions.

2.3 Graph-based approaches

Because payment systems can be represented as networks where banks are nodes and payment obligations are links, graph-theoretic tools are a natural fit for analyzing and improving settlement efficiency. One class of methods utilizes maximum flow algorithms to maximize the number of payments that can be processed through the system, given the available liquidity at each participant [Fleischman and Dini (2021)]. Well-known algorithms, such as Ford-Fulkerson and Push-Relabel, help find the most efficient way to move payments through a network given a fixed set of constraints. But in real-time payment systems, conditions are rarely constant. Payments arrive throughout the day, and liquidity positions are constantly changing. This implies that these flow solutions cannot be static and must be updated continuously. Incremental flow algorithms, which can adjust existing solutions as new information becomes available, are crucial for maintaining performance in live environments.

Another important area is cycle detection. If the system identifies a closed loop of payment obligations where Bank A owes Bank B, Bank B owes Bank C, and Bank C owes Bank A, it can often settle that loop with very little liquidity. The challenge lies in identifying these cycles quickly, without creating gridlocks or unfair outcomes for specific participants. More advanced systems use techniques such as strongly connected component analysis to focus their search on parts of the network where netting opportunities are more likely.

Graph analysis is also useful to understand how the whole payment system works. For example, network metrics, such as centrality, can show which banks play an important role in the flow of payments. Moreover, algorithms that detect communities within the network can identify groups of participants that might benefit from shared liquidity or netting arrangements. These tools are helpful not only for improving day-to-day efficiency but also for managing risk.

2.4 Machine-learning and AI approaches

Artificial intelligence (AI) is opening new possibilities for how payment systems manage liquidity and scheduling [Beckstrom (2021), Rossi and Duigenan (2021)]. AI is also used to detect anomalous transactions, which, if left undetected, can lead to operational outages and systemic risk [Desai et al. (2024)]. One key application is pattern recognition. By analyzing past settlement data, machine learning models can spot recurring behaviors such as how specific participants tend to time their payments, or how payment flows shift around major market events. These are patterns that traditional rule-based systems often miss.

More advanced deep learning models can pick up subtle relationships in high-dimensional data [Rubio et al., (2020)]. For example, recurrent neural networks (RNNs) are well-suited to capturing how payment patterns evolve, while convolutional neural networks (CNNs) can be used to detect structural features in the payment network itself. Still, these models come with trade-offs as they are complex, often hard to interpret, and must be implemented carefully to meet transparency and regulatory standards.

Another promising area is reinforcement learning (RL), which allows an algorithm to learn by trial and error [Castro et al. (2021)]. In this setup, the system receives feedback "rewards" based on outcomes such as settlement speed or liquidity use and gradually learns strategies that improve over time. But RL brings its own risks. It can be difficult to define reward functions that encourage the right behavior, and unstable policies or unexpected responses to new conditions are unacceptable in critical financial infrastructure. For this reason, RL models must be tested extensively before being used in practice.

As a forward-looking tool, predictive analytics helps payment systems anticipate the sequence of future payments. Models that can forecast incoming payment flows, liquidity shortages, or periods of network stress allow operators to adjust their processing strategies in advance before problems arise. The most advanced systems also include external information, such as interest rates, macroeconomic indicators, or market sentiment data. However, as with any predictive system, these models require ongoing monitoring to ensure they continue to function effectively as conditions change.

3. New approaches to payment reordering: classical preprocessing and quantum

While the literature offers a variety of queue management and optimization strategies for real-time gross settlement (RTGS) systems, our research has focused on bridging the gap between theoretical improvements and implementable tools that work within the technical and regulatory constraints of live systems. Our team at GoodLabs Studio has taken two different approaches to improving liquidity efficiency, both designed to work within existing system frameworks, without the need for major infrastructure changes.

The first project explored whether quantum-inspired optimization could reduce liquidity usage in Canada's high-value payment system [McMahon et al. (2024)]. We formulated the payment scheduling problem as a Quadratic Unconstrained Binary Optimization (QUBO) model, a framework well-suited for expressing combinatorial problems with large, discrete solution spaces such as payments data. Our objective was to minimize peak liquidity usage over a given batch of payments, subject to settlement constraints that mimic the operational realities of an RTGS system. To solve this problem, we used a quantum-inspired approach built on D-Wave's platform, which combines classical preprocessing and postprocessing techniques with quantum annealing executed on D-Wave's quantum hardware.

Our approach was intentionally designed to be non-intrusive and easy to integrate into LYNX's workflow. Rather than modifying the core infrastructure of the RTGS system, we decided to frame the optimization as a standalone preprocessing layer (see Figure 1). To reflect real-time operational constraints, the preprocessor was designed to handle batches of up to 70 payments at a time. For each batch, it calculated an execution order

Figure 1: Schematic of the quantum optimizer as a preprocessor to the payment system

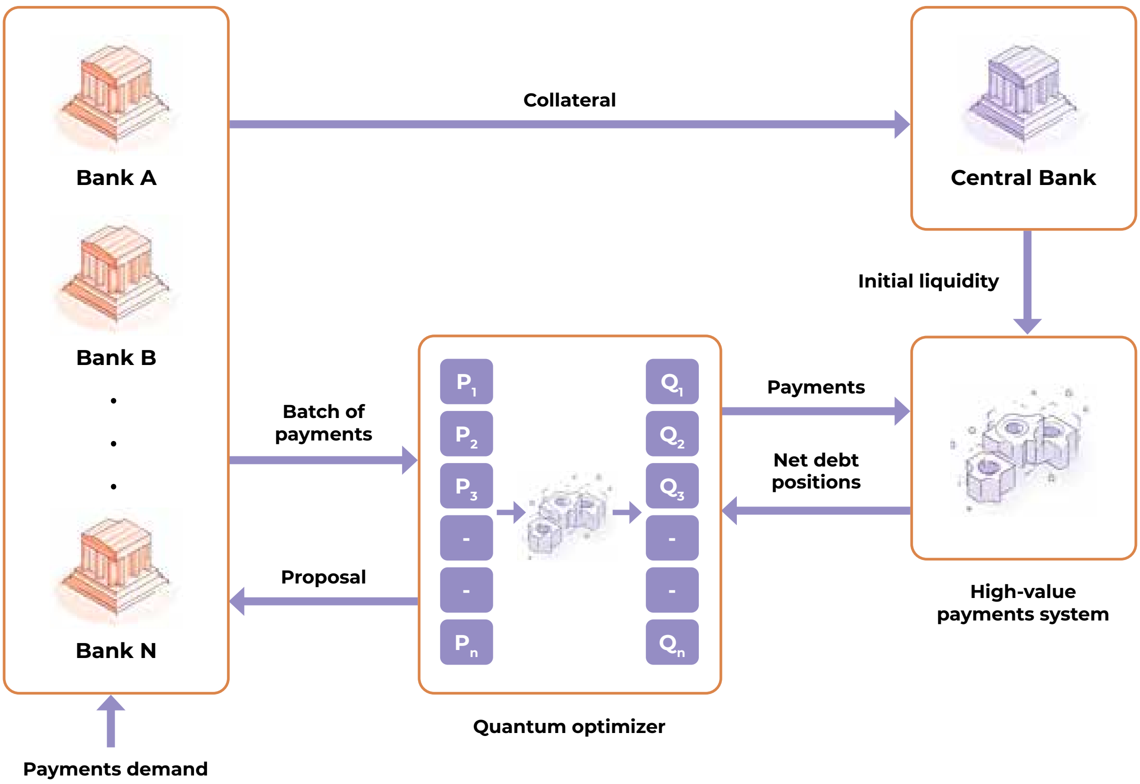

Source: Reproduced from McMahon et al (2024) that minimized liquidity usage before the payments were submitted to the system to be settled. This reordered payment list can then be voluntarily adopted by participants, without disrupting settlement protocols, affecting transaction finality, or introducing new security risks. Our approach is attractive for institutions seeking to reduce intraday liquidity needs without altering system-level infrastructure.

We compared the performance of this hybrid solver against baseline methods, including standard FIFO ordering and a more flexible FIFO-bypass strategy. The tests were conducted on historical payment data from LYNX, Canada's real-time gross settlement system, using anonymized data. Across these simulations, the quantum-enhanced scheduling algorithm consistently outperformed several benchmark algorithms, such as FIFO and others we tested. With the quantum-inspired algorithm, we achieved average liquidity savings of 1.36% per day, equivalent to 240 million Canadian dollars in liquidity on high-volume days. The median daily savings stood at 128.8 million CAD. These gains were particularly significant during periods of heavy congestion or simulated market stress, precisely when efficient liquidity management is most critical.

In our experiments, we also tested larger batch sizes of up to 140 payments. While these larger batches yielded slightly greater liquidity savings compared to smaller batches (70 payments), the marginal gains diminished as batch size increased. More importantly, the runtime grew disproportionately, making it harder to meet real-time operational requirements. For instance, solving a batch of 140 payments could take several seconds, which is acceptable in some contexts, but too slow for continuous settlement flows. As a result, we focused on 70-payment batches as a practical compromise: they offered a meaningful improvement in liquidity efficiency (up to 1.36%) while remaining fast enough to support near-real-time processing.

We also compared our quantum-inspired method against a classical optimization approach using the Solving Constraint Integer Programs (SCIP) solver for mixed-integer nonlinear programming. While SCIP can, in principle, solve the problem optimally, it proved impractical in real-world settings. For batches of 70 payments, the average runtime exceeded two hours and occasionally failed to converge after 24 hours, far beyond the limits of real-time settlement. In contrast, D-Wave's hybrid solver handled 70-payment batches in under two minutes on average, with the core optimization step taking only about five seconds.

However, we observed that as batch sizes increased beyond 140 payments, two limitations emerged. First, the total time to compile and solve the optimization problem increased sharply, approaching the limit of real-time feasibility. Second, the quality of the solutions began to degrade as the solver returned fewer feasible solutions. In some cases, the resulting liquidity usage exceeded that of the FIFO benchmark. These effects highlight the current scalability limits of quantum-inspired methods. Nevertheless, within practical batch sizes, the approach delivers meaningful efficiency gains that are currently out of reach for classical solvers.

Our more recent work takes a different approach, using conventional computing resources and classical optimization methods to achieve similar goals. In recent work, we developed a preprocessing algorithm that reorders payments using a combination of virtual liquidity levels (VLLs) and greedy optimization within sliding or tumbling time windows. The main idea is to simulate artificial liquidity constraints at the start of the day and progressively relax them, encouraging the algorithm to delay non-urgent payments just long enough to unlock better settlement sequences. Each queue of payments is then optimized by local swaps that reduce peak liquidity demand, while respecting constraints on maximum delay, both absolute and value weighted.

Even with partial information and real-time constraints, this method showed important improvements. In simulation environments that replicate typical RTGS conditions, our approach consistently reduced aggregate liquidity needs while introducing only modest delays, though in some cases the delays extended beyond what would be acceptable in an RTGS environment. Liquidity savings compared to the benchmark case were often more than 10%, typically 30% when historical payments data were used. To address the longer payment delays, we found that tuning the VLL parameter yields a clear liquidity-delay trade-off frontier, with an identifiable "sweet spot" where liquidity savings are significant, and average delay costs remain tolerable. The algorithm operates purely at the submission layer and makes no assumptions about internal settlement logic, making it deployable as an external tool. Importantly, improvements were observed across the board, not just for the largest institutions supporting broader fairness and adoption.

These two contributions share a common approach: rather than reinventing the architecture of RTGS systems, we focus on improving what participants submit and when. By making smarter use of existing flexibility in payment timing and sequencing, and by providing real-time tools to guide that decision-making, it is possible to achieve liquidity savings without disrupting the safety, fairness, or legal finality of existing settlement systems.

4. Broader applications of liquidity optimization and scheduling

While our work has focused on improving liquidity efficiency in real-time gross settlement systems, the methods and principles we have developed are also applicable across the financial sector and have numerous use-cases. Many institutions face similar challenges: how to manage constrained resources (liquidity, capital, computing) while dealing with time-sensitive and high volumes of transactions. Smarter strategies for reordering, prioritizing, or netting these transactions can yield significant gains in efficiency, risk reduction, and operational resilience.

A natural extension of our work is collateral optimization in capital markets. Large institutions often manage multiple margin calls across clearinghouses and counterparties. Firms can minimize the fragmentation of collateral and reduce overposting by sequencing obligations more strategically using the optimization techniques we have applied to real-time payments. Similarly, in securities settlement, especially in systems with multiple counterparties and delayed matching, algorithms that identify optimal settlement cycles or minimize failed trades can enhance both liquidity usage and operational throughput, particularly under settlement compression regimes or in T+1 transition contexts. In corporate treasury operations, daily cash sweeps and payment scheduling could benefit from predictive and reinforcement learning tools similar to those we applied to RTGS queues. Firms could dynamically prioritize payables and receivables to reduce daylight overdrafts or borrowing costs, especially when liquidity is tight, or markets are volatile.

Buy-side institutions, particularly asset managers and hedge funds, may find value in using our scheduling and queueing frameworks to improve order execution timing, especially for block trades or large-volume redemptions. Models that account for counterparty behavior, settlement friction, and operational latency can help reduce slippage and settlement risk. Finally, in retail and instant payment systems, where transaction volumes are high, but margins are tight, our methods could support smart routing, liquidity pooling, or delay-based netting at microsecond timescales. Even small improvements in efficiency can lead to significant savings when applied across millions of transactions.

5. Conclusion

As real-time payment volumes pressure grows, improving liquidity efficiency has become a crucial operational and strategic priority for financial and other institutions. This paper has outlined how payment reordering, whether using classical algorithms or quantum-inspired techniques, can lead to significant gains without requiring disruptive changes to the existing system. Our results demonstrate that substantial improvements in liquidity usage are possible with practical tools and minimal implementation risk. As computational methods and quantum technologies, in particular, continue to improve, we believe these innovations will become essential not only to modern payment infrastructure but also to optimization challenges across a range of industries.

References

Bardoscia, M., P. Barucca, S. Battiston, F. Caccioli, G. Cimini, D. Garlaschelli, F. Saracco, T. Squartini, and G. Caldarelli, 2021, "The physics of financial networks," Nature Reviews Physics 3, 490-507

Bech, M. L., and K. Soramäki, 2001, "Gridlock resolution in interbank payment systems," Bank of Finland Research Discussion Paper No. 9/2001

Beckstrom, J. R., 2021, "Auditing machine learning algorithms," International Journal of Government Auditing 48:1, 40-41

BoE, 2022, "Roadmap for real-time gross settlement service beyond 2024," Bank of England Consultation Paper, Bank of England, https://tinyurl.com/yu6ymac7

BoE, 2024a, "Payment and settlement statistics," Bank of England, https://tinyurl.com/mr36hzvd

BoE, 2024b, "Artificial intelligence in UK financial services," Bank of England, https://tinyurl.com/2nrwn7px

Byck, S., and R. Heijmans, 2020, "How much liquidity would a liquidity-saving mechanism save if a liquidity-saving mechanism could save liquidity? A simulation approach for Canada's large-value payment system," Payment Canada Discussion Paper

Cambe, J., and Z. Xing, 2024, "Just solve it: a simple method to improve the design and performance of liquidity-saving mechanisms," Journal of Financial Market Infrastructures 24:2, 13-37

Castro, P. S., A. Desai, H. Du, R. J. Garratt, and F. Rivadeneyra, 2021, "Estimating policy functions in payment systems using RL," Bank of Canada Staff Working Paper 2021-7

Chen, S., O. Baron, and N. Chen, 2024, "Optimal dynamic clearing for interbank payments," Management Science 71:4, 2953-2974

de Avillez, R., D. Millard, N. Bulusu, M. McNeely, L. Sabetti, K. Wai, and N. Zhang, 2022, "An overview of Lynx, Canada's High-Value Payment System," Bank of Canada Working Paper 2022-05, https://tinyurl.com/5y23t5jy

Davey, N., and D. Gray, 2014, "How has the liquidity saving mechanism reduced banks' intraday liquidity costs in CHAPS?" Bank of England Quarterly Bulletin Q2, Bank of England

Desai A., A. Kosse, and J. Sharples, 2024, "Machine learning framework for anomaly detection in HVPS," BIS Working Paper 1188, https://tinyurl.com/mr34aj48

ECB, 2024, "Target services annual report," European Central Bank, https://tinyurl.com/yf64vef4

Fleischman, T., and P. Dini, 2021, "Mathematical foundations for balancing the payment system in the trade credit market," Journal of Risk and Financial Management 14:9, 452

Gaffeo, E., L. Gobbi, and M. Molinari, 2019, "The economics of netting in financial networks," Journal of Economic Interaction and Coordination 14, 595-622

Garratt, R. J., Z. Lu, and P. Tian, 2023, "How banks create gridlock to save liquidity in Canada's large value payment system," Bank of Canada Staff Working Paper 2023-26

Lawler, E. L., J. K. Lenstra, A. H. G. Rinnooy Kan, and D. B. Shmoys, 1986, "The traveling salesman problem: a guided tour of combinatorial optimization," Journal of the Operational Research Society 37, 535–536

McMahon, C., D. McGillivray, A. Desai, F. Rivadeneyra, J-P. Lam, T. Lo, D. Marsden, and V. Skavysh, 2024, "Improving the efficiency of payments systems using quantum computing," Management Science 70:10, 7325-7341

Payments Canada, 2024, "Lynx Canada's High-Value Payment System," https://tinyurl.com/wubhd93z

Rossi, F., and J. Duigenan, 2021, "The future of finance is trustworthy AI," Payments Journal, https://tinyurl.com/5b39hc2z

Rubio, J., P. Barucca, G. Gage, J. Arroyo, R. Morales-Resendiz, 2020, "Classifying payment patterns with artificial neural networks: an autoencoder approach," Latin American Journal of Central Banking 1:1-4

Rivadeneyra, F. and N. Zhang, 2020, "Liquidity usage and payment delay estimates of the new Canadian High-Value Payments System," Bank of Canada Technical Report no. 9

Vishnevsky, V., and A. Gorbunova, 2022, "Application of machine learning methods to solving problems of queuing theory," in Dudin, A., A. Nazarov, and A. Moiseev, (eds.), Information technologies and mathematical modelling. Queueing theory and applications, ITMM 2021, Communications in Computer and Information Science, vol. 1605. Springer

WE MAKE THINGS HAPPEN.

Leading change in Financial Services.

BELGIUM
UNITED KINGDOM
THE NETHERLANDS
FRANCE
GERMANY
SWITZERLAND

www.projectivegroup.com

Projective INSTITUTE