The Velocity Threshold

Refresh Interval Economics for Fragmented AMM Settlement Networks
(A network-level extension of the Jackson Liquidity Framework)

Lewis Jackson
Lewis Jackson Ventures Ltd (Research Dept)

18 December 2025

Not investment advice. This paper is research-oriented and intended for educational and analytical purposes.

Executive Summary

Automated Market Maker (AMM) liquidity pools operating under instant settlement regimes face a fundamental scaling challenge: when liquidity is fragmented across multiple pools, total capital requirements increase dramatically. However, if liquidity can be recycled rapidly between uses — characterized by a short **refresh interval** — the effective capital requirement can be reduced.

This paper derives a quantitative framework for calculating the **velocity threshold**: the refresh interval required to make a fragmented multi-pool architecture economically equivalent to a unified single-pool system. The analysis extends the Jackson Liquidity Framework (JLF) by introducing time-dependent scaling relationships for buffer requirements.

We demonstrate that the velocity threshold exhibits quadratic scaling with pool count, creating a sharp trade-off between architectural fragmentation and operational refresh speed. The framework also identifies a critical regime boundary where slippage constraints dominate time-based buffers, causing velocity improvements to lose effectiveness.

Key results:

1. **Key result 1 — velocity-adjusted fragmentation multiplier.** If per-pool reserves are dominated by time-window risk buffers (VaR / intraday depletion), then reserves scale approximately with $\sqrt{\tau}$, where τ is the refresh interval. Total liquidity across n equal pools scales as:

$$M(n, \tau; T_0) = n \cdot \sqrt{\frac{\tau}{T_0}} \tag{1}$$

   where T_0 is the baseline time horizon (eg, 1 day).

2. **Key result 2 — the velocity threshold is quadratic.** To fully neutralize fragmentation (make n pools require the same total liquidity as one pool), the refresh interval must satisfy:

$$\tau^* = \frac{T_0}{n^2} \tag{2}$$

   This is the "velocity threshold." Fragmentation grows linearly; the refresh interval must shrink quadratically to offset it.

3. **Key result 3 — velocity does not reduce the slippage floor.** JLF defines the liquidity requirement as the maximum of several constraints. The slippage constraint depends on trade size quantiles and slippage tolerance — not on time. If slippage is the binding constraint, faster refresh does not reduce required reserves. Velocity mostly compresses the buffer components.

Practical implication:

To assess whether a bridge asset architecture (such as using XRP as a settlement intermediary between fiat currencies or other digital assets) can operate without creating excessive capital requirements, three parameters must be quantified:

1. The number of distinct liquidity pools required (n)
2. Whether liquidity requirements are dominated by slippage constraints or time-based buffer constraints
3. The refresh interval (τ) that the system architecture can sustain

The velocity threshold formula ($\tau^* = T_0/n^2$) provides the critical reference point for this assessment.

1 Motivation and Problem Statement

Modern financial infrastructure is evolving toward instant, atomic settlement architectures. Tokenisation initiatives, Central Bank Digital Currency (CBDC) experiments, and blockchain-based payment systems share a common characteristic: they eliminate counterparty risk by collapsing multi-step settlement processes into single, atomic transactions. However, this architectural choice has a direct consequence — liquidity must be pre-funded and available before transactions can execute.

In traditional systems with delayed settlement, multiple transactions can be netted against each other, reducing the total amount of capital that must be locked. In instant settlement systems, each transaction requires immediate liquidity availability, and netting opportunities diminish significantly. When these settlements are facilitated through Automated Market Maker (AMM) pools, the relationship between pool depth and execution quality introduces additional nonlinear constraints.

This creates a fundamental architectural question: in a global financial system with thousands of currencies, tokens, and digital assets that need to interoperate atomically, where does the required liquidity come from, and how much is needed?

1.1 The combinatorics trap and hub architectures

Consider a global payment system with N distinct assets (currencies, tokens, CBDCs, etc.). If every asset must be able to trade directly with every other asset, the number of required liquidity pools is:

$$\text{Pairs}_{\text{full mesh}} = \frac{N(N-1)}{2} \tag{3}$$

For $N = 3{,}000$ assets, this yields 4,498,500 unique trading pairs. Each pair requires dedicated liquidity, making this "full mesh" architecture economically infeasible.

Hub architecture provides an alternative: select a single bridge asset (the "hub") and route all conversions through it. Instead of requiring Asset A $\leftrightarrow$ Asset B pools for every combination, the system only needs Asset A $\leftrightarrow$ Hub and Hub $\leftrightarrow$ Asset B pools. This reduces the scaling from $O(N^2)$ to $O(N)$:

$$\text{Pairs}_{\text{hub}} = N - 1 \approx N \tag{4}$$

For 3,000 assets, this reduces the requirement from millions of pools to approximately 3,000 pools.

However, even N pools can present a scaling challenge when liquidity must be pre-funded across all pools simultaneously. This is where the refresh interval becomes critical: if liquidity can be rapidly recycled, the effective capital requirement may be substantially lower than the sum of all pool requirements.

2 The JLF Foundations: What Scales With Time, and What Doesn't

The Jackson Liquidity Framework defines the Jackson Liquidity Requirement (JLR) as the maximum of several constraints: slippage tolerance, directional-flow Value-at-Risk (VaR), intraday peak depletion, and Basel-aligned constraints. The key point is that these constraints do not respond to "velocity" in the same way.

2.1 Two governors: size vs time

In practice, AMM liquidity has two independent governors:

- **Size constraint (slippage floor):** the pool must be deep enough to execute large trades without exceeding a slippage tolerance ϵ.
- **Time constraint (buffer requirement):** the pool must be deep enough to survive directional flow and clustering over a time horizon T before liquidity can be replenished or rebalanced.

In JLF's compact form:

VaR (imbalance over a time horizon T):

$$\text{VaR}_{0.99}(I(T)) = \mathbb{E}[I(T)] + 2.326 \cdot \sqrt{\text{Var}[I(T)]} \tag{5}$$

Slippage floor:

$$R_{\text{slip}} = \left(\frac{1-\epsilon}{\epsilon}\right) \cdot F_X^{-1}(1-\alpha) \tag{6}$$

Stability condition (JSI):

$$\frac{R_A \cdot R_B}{\sigma_I \cdot \sigma \cdot \sqrt{T}} \geq K(\epsilon, \rho) \tag{7}$$

The square-root dependence on T in the stability condition is the mathematical doorway for "refresh interval" arguments. But note: the slippage floor is not a function of T. It is a function of trade sizes and tolerance.

3 Fragmentation: What We Already Know

In JLF Case Study 3, fragmentation is simulated by splitting liquidity into multiple equal pools while holding flow parameters fixed. The results are deliberately stark: moving from one pool to ten pools increases total JLR by 900% (a 10× increase), because each pool becomes shallower and AMM slippage is nonlinear.

The JLF Simple Guide summarizes this linear scaling for equal pools as:

$$\text{JLR}_{\text{fragmented}} = n \cdot \text{JLR}_{\text{unified}} \tag{8}$$

So far, none of this includes "velocity." It treats T as fixed (eg, one business day) and asks how much liquidity must be locked in parallel across pools.

4 Refresh Interval: Turning "Velocity" Into a Variable

The refresh interval can be defined as: "how long does liquidity need to sit still before it can be used again?" In a corridor model, this maps naturally to the time horizon T used for imbalance, VaR and stability calculations.

We introduce τ as the effective refresh interval. Examples:

- In a continuous, always-on system with rapid rebalancing, τ might be seconds to minutes.
- In a batched privacy system (eg, proofs posted periodically), τ might be minutes to hours.
- In a netting-like regime, τ might approach end-of-day.

4.1 First-order scaling (buffer-dominated regime)

If the binding constraint is a time-window buffer (VaR / intraday peak / JSI stability), then reserve requirements tend to scale like $\sqrt{\tau}$ (because both VaR and the JSI condition inherit $\sqrt{T}$ terms, and the imbalance variance scales with time).

$$R(\tau) \approx R_0 \cdot \sqrt{\frac{\tau}{T_0}} \tag{9}$$

$$L_{\text{total}}(n, \tau) \approx n \cdot R_0 \cdot \sqrt{\frac{\tau}{T_0}} \tag{10}$$

$$M(n, \tau; T_0) = n \cdot \sqrt{\frac{\tau}{T_0}} \tag{11}$$

4.2 The critical caveat (slippage-dominated regime)

$$R_{\text{pool}}(\tau) = \max\left(R_{\text{slip}}, k \cdot \sqrt{\tau}\right) \tag{12}$$

$$L_{\text{total}}(n, \tau) = n \cdot \max\left(R_{\text{slip}}, k \cdot \sqrt{\tau}\right) \tag{13}$$

This creates two distinct regimes: (1) slippage-limited (velocity doesn't help), and (2) buffer-limited (velocity helps by $\sqrt{\tau}$).

5 Results: The Velocity Threshold in Numbers

To make the scaling concrete, we use a baseline day $T_0 = 86{,}400$ seconds. We compute the refresh interval τ^* required to fully neutralize fragmentation under the buffer-dominated approximation ($\tau^* = T_0/n^2$).

Interpretation: ten pools require refresh around 14.4 minutes to fully offset fragmentation; one hundred pools require ∼8.6 seconds; three thousand pools require ∼9.6 milliseconds.

Figure: τ^* vs pool count (log-log).

Figure: multiplier vs refresh interval for selected pool counts.

6 Worked Example: When Slippage Dominates, Velocity Stops Helping

JLF Case Study 1 (Poisson baseline) provides a useful decomposition: total JLR ≈ \$20.38M, driven primarily by slippage tolerance. Per-side slippage requirement is about \$10.19M, while the intraday/VaR buffers are ∼0.51M (side A) and ∼1.59M (side B).

$$\tau_c = T_0 \cdot \left(\frac{R_{\text{slip}}}{R_{\text{buffer}}} \right)^2 \tag{14}$$

Using the baseline numbers as a rough calibration: for side B, τ_c is about 40.83 day(s) (≈ 40.8 days). For any refresh interval shorter than that, the slippage floor remains the binding constraint.

Therefore, a correct "velocity thesis" for AMM architectures must also answer: what mechanism reduces effective trade size or slippage?

7 The Research Gap: We Need to Measure τ and the Binding Constraint

Public documents rarely specify target "XRP velocity." But the system design choices institutions are actively making do imply refresh intervals, and those refresh intervals determine whether we are in the buffer-limited or slippage-limited regime.

7.1 A measurement program (what we can do now)

- **Estimate a refresh interval proxy from system mechanics.** For a given venue or rail, τ can be approximated by the minimum of: (a) settlement finality time, (b) batch/proof interval, (c) liquidity rebalancing cycle, and (d) operational cut-off cycles.
- **Estimate a slippage floor proxy from trade-size tails.** If we can approximate the 99th percentile trade size (or netted trade size) and define an execution tolerance, we can estimate R_{slip} directly using the JLF slippage constraint.
- **Map regimes.** Compare τ to τ_c. If $\tau < \tau_c$, it is slippage-limited. If $\tau > \tau_c$, it is buffer-limited.
- **Network composition.** Once per-corridor regimes are known, aggregate across corridors (or across seams between venues) to estimate total bridge-asset demand.

Even outside crypto, regulators and institutions focus on intraday liquidity circulation and on the liquidity cost of immediate settlement. BCBS 248 defines monitoring tools for intraday liquidity risk management, and IMF work on financial platforms notes that immediate settlement entails liquidity costs because transactions must be pre-funded.

8 Worked Examples: Velocity, Fragmentation, and Bridge Asset Requirements

To make the theoretical framework concrete, we present a series of worked examples using the JLF baseline parameters and current market data. These examples demonstrate how refresh intervals translate into liquidity requirements and bridge asset inventory needs.

8.1 Methodological note

The Jackson Liquidity Requirement (JLR) is defined as the maximum of several constraints. For illustrative purposes, we decompose JLR into two components:

$$\mathrm{JLR}(\tau) \approx \max\left(R_{\mathrm{slip}} + R_{\mathrm{time}}\sqrt{\frac{\tau}{T_0}}, R_{\mathrm{other}}\right) \tag{15}$$

where:

- R_{slip} is the slippage floor (independent of time)
- R_{time} represents time-sensitive buffer components (VaR, intraday peaks, JSI stability)
- $T_0 = 1$ day (baseline time horizon)

This decomposition is consistent with the max-constraint definition in JLF and allows for scenario sensitivity analysis.

8.2 Example 1: Single corridor — velocity vs liquidity requirement

Using JLF Case Study 1 baseline parameters:

- JLR(1 day) = \$20.38M
- $R_{\mathrm{slip}} \approx$ \$10.19M
- $R_{\mathrm{time}} \approx$ \$10.19M

Refresh interval τ	τ/day	**JLR(τ)**	**Interpretation**
1 day	1.000	\$20.38M	Baseline
1 hour	0.042	\$12.27M	Significant reduction
5 min	0.003	\$10.79M	Most gains captured
30 sec	0.0003	\$10.38M	Near slippage floor
5 sec	0.00006	\$10.27M	Diminishing returns

Table 1: Liquidity requirement vs refresh interval for a single corridor. Velocity compresses buffer requirements but cannot reduce liquidity below the slippage floor.

Key finding: Velocity can reduce buffer-driven liquidity dramatically (from \$20.4M to \$12.3M at 1-hour refresh), but hits a hard floor once slippage dominates. This demonstrates why “speed alone eradicates fragmentation” is directionally true but not absolute.

8.3 Example 2: Fragmentation penalty with varying velocity

Using $n = 10$ equal pools with the fragmentation rule $\mathrm{JLR}_{\mathrm{total}}(n, \tau) = n \cdot \mathrm{JLR}(\tau)$:

Key finding: Even with extremely fast refresh intervals (5 seconds), fragmentation remains expensive (\$102.7M vs \$20.4M for unified). Velocity mitigates fragmentation but cannot eliminate it due to the slippage constraint floor.

Pool count n	**Refresh** τ	**Total JLR**
10	1 day	\$203.8M
10	1 hour	\$122.7M
10	5 min	\$107.9M
10	5 sec	\$102.7M

Table 2: Total liquidity requirement across 10 fragmented pools at varying refresh intervals.

8.4 Example 3: Bridge asset token requirements at current market prices

Converting liquidity requirements to XRP token counts using current market price of \$2.30 per XRP (December 2025):

Refresh τ	**JLR (USD)**	**XRP required**
1 day	\$20.38M	8.86M XRP
1 hour	\$12.27M	5.33M XRP
5 min	\$10.79M	4.69M XRP
30 sec	\$10.38M	4.51M XRP
5 sec	\$10.27M	4.47M XRP

Table 3: Single-corridor XRP inventory requirements at \$2.30/XRP across refresh intervals.

Key finding: Token inventory requirements shrink rapidly at first (from 8.86M to 5.33M XRP moving from daily to hourly refresh), then flatten as slippage constraints dominate. This is the "velocity vs inventory" tradeoff visualized in token terms.

8.5 Example 4: Seam flow estimation from incumbent payment volumes

To size potential bridge asset demand, we estimate "seam flow" — the value that must cross between different settlement venues or rails. Using incumbent (non-pilot) baselines:

- Cross-border payments: >\$190T annually (2024 estimate)
- OTC FX turnover: \$9.6T/day (April 2025 BIS survey)
- CLS PvP settlement: \$7.9T/day average (H1 2025)

We define seam flow as:

$$\text{Seam Flow} = \text{Baseline Volume} \times a \times s \tag{16}$$

where a is adoption rate into tokenized venues and s is the share requiring inter-venue settlement.

Scenario	**Adoption** a	**Seam share** s	**Daily seam flow**
Conservative	10%	10%	\$5.2B/day
Base	25%	25%	\$32.5B/day
Aggressive	50%	40%	\$104.1B/day

Table 4: Scenario bands for daily seam flow using \$190T annual cross-border baseline.

8.6 Example 5: XRP routing share and inventory requirements

If XRP captures a share b of seam flow, inventory requirements scale with refresh interval:

$$\text{Inventory}_{\text{USD}} \approx \text{Flow}_{\text{day}} \times \left(\frac{\tau}{1440 \text{ min}}\right) \times k \tag{17}$$

where k is a buffer multiplier for flow variance (using $k = 3$ for stress scenarios).

For the **Base scenario** ($32.5B/day seam flow) with 15% XRP routing share ($4.88B/day XRP-routed flow):

Refresh τ	**Inventory (USD)**	**XRP tokens at $2.30**
60 min	$610M	265M XRP
15 min	$153M	66M XRP
5 min	$50.8M	22M XRP
1 min	$10.2M	4.4M XRP

Table 5: XRP inventory requirements for $4.88B/day routing volume across refresh intervals.

Key finding: The same daily throughput ($4.88B) produces wildly different inventory requirements depending on operational velocity — from 265M XRP at hourly refresh to just 4.4M XRP at 1-minute refresh. This demonstrates the practical importance of the velocity threshold for bridge asset economics.

8.7 Combining fragmentation and seam flow: network-scale requirements

For a network with n distinct seams (inter-venue connection points), total bridge asset requirements become:

$$\text{Total Inventory} = n \times \text{Per-Seam Inventory}(\tau) \tag{18}$$

Using the Base scenario with $n = 100$ seams and 5-minute refresh:

- Per-seam inventory: 22M XRP (from Example 5)
- Total network inventory: $100 \times 22\text{M} = 2.2\text{B}$ XRP
- As percentage of supply: $\frac{2.2\text{B}}{100\text{B}} = 2.2\%$ of total XRP supply

This framework allows systematic assessment of bridge asset demand under different network topology (n), adoption (a, s), routing share (b), and velocity (τ) assumptions.

9 Bridge Asset Economics

The analysis reveals two distinct scenarios for bridge asset demand in hub architectures:

Scenario 1: High-velocity regime. If the system can sustain refresh intervals below the velocity threshold ($\tau < \tau^*$), and if time-based buffers dominate the liquidity requirement, then fragmentation penalties can be largely neutralized. In this regime, a bridge asset like XRP could facilitate large transaction volumes with relatively modest inventory requirements.

Scenario 2: Slippage-limited regime. If large trade sizes or execution quality requirements make slippage the binding constraint, then velocity improvements provide no benefit. In this regime, total liquidity requirements scale linearly with pool count regardless of refresh speed.

The critical question for any bridge asset architecture is: which regime applies? This determination requires:

1. Quantifying the operational refresh interval from system design (settlement finality, batch intervals, rebalancing cycles)
2. Estimating the slippage floor from trade size distributions and execution tolerances
3. Comparing τ to the critical refresh interval $\tau_c = T_0 \cdot (R_{\text{slip}}/R_{\text{buffer}})^2$

Bridge asset demand cannot be determined from network topology alone — operational velocity and constraint binding are equally fundamental.

10 Appendix: 2030 Global Thought Experiment

10.1 Purpose and methodological discipline

This section provides a forward-looking thought experiment for global-scale seam liquidity requirements in a mature tokenized settlement system (illustrative target: 2030). The analysis explicitly separates:

1. **Incumbent baseline flows** — measured from current global payment systems (not pilot programs)
2. **Scenario parameters** — adoption rates, seam shares, and bridge asset routing shares (explicitly labeled as assumptions)
3. **Inventory requirements** — derived from refresh interval mathematics and buffer multipliers

This is not a price prediction or demand forecast. It is a conditional sizing exercise demonstrating how the velocity threshold framework translates baseline flow data into liquidity requirements under explicit scenario assumptions.

10.2 Baseline flows: incumbent system magnitudes

We anchor the analysis to three independently measured incumbent flows:

- **Cross-border payments market:** $\approx$ \$190T annually (2024 estimates)
- **OTC FX turnover:** $\approx$ \$9.6T/day (BIS April 2025 survey)
- **CLS PvP settlement:** $\approx$ \$7.9T/day (H1 2025 average)

Using the cross-border baseline, daily flow is:

$$F_{\text{CB,day}} = \frac{190 \times 10^{12}}{365} \approx \$520.5\text{B/day} \tag{19}$$

10.3 From baseline flows to seam flows

Seam flow represents value that must cross between different settlement venues, rails, or jurisdictions. We parameterize this with two scenario variables:

- $a \in (0,1)$ — adoption rate into tokenized settlement venues
- $s \in (0,1)$ — share of tokenized volume requiring inter-venue settlement (seams)

Daily seam flow becomes:

$$F_{\text{seam,day}} = F_{\text{baseline,day}} \times a \times s \tag{20}$$

If XRP is used as a bridge asset for seam settlement, we introduce:

- $b \in (0, 1)$ — XRP's share of total seam flow

XRP-routed daily flow:

$$F_{\text{XRP,day}} = F_{\text{seam,day}} \times b \tag{21}$$

10.4 Inventory requirements from refresh interval

Let τ (in minutes) represent the **refresh interval** — the time liquidity must remain available before it can be recycled. Let $k \geq 1$ represent a buffer multiplier to account for flow clustering, peaks, and variance.

The inventory proxy becomes:

$$L_{\text{USD}}(\tau) = F_{\text{XRP,day}} \times \left(\frac{\tau}{1440}\right) \times k \tag{22}$$

This proxy captures the core insight from the Jackson Stability Invariant: liquidity stress depends on time horizon (via $\sqrt{T}$ scaling in stability boundaries), and faster reuse reduces required idle inventory.

Important limitation: This is a seam-level throughput-to-inventory proxy. Full corridor sizing remains governed by JLR = $\max\{R_{\text{slip}}, R^{\text{VaR}}, R^{\text{intraday}}, R^{\text{Basel}}\}$, meaning slippage floors and tail constraints may bind regardless of velocity.

10.5 Token count as a function of price

Let P_{XRP} represent the unit price of XRP (USD/XRP). Required inventory in token terms:

$$L_{\text{XRP}}(\tau) = \frac{L_{\text{USD}}(\tau)}{P_{\text{XRP}}} \tag{23}$$

10.6 Scenario bands: cross-border baseline

We define three illustrative scenario bands:

- **Conservative:** $(a, s, b) = (0.10, 0.10, 0.05)$
- **Base:** $(a, s, b) = (0.25, 0.25, 0.15)$
- **Aggressive:** $(a, s, b) = (0.50, 0.40, 0.30)$

Using buffer multiplier $k = 3$ and three refresh intervals: $\tau \in \{60 \text{ min}, 5 \text{ min}, 30 \text{ sec}\}$:

Scenario	τ	$F_{\text{seam,day}}$	$F_{\text{XRP,day}}$	$L_{\text{USD}}(\tau)$
Conservative	60 min	\$5.21B/day	\$0.260B/day	\$32.5M
Conservative	5 min	\$5.21B/day	\$0.260B/day	\$2.71M
Conservative	30 sec	\$5.21B/day	\$0.260B/day	\$0.271M
Base	60 min	\$32.5B/day	\$4.88B/day	\$610M
Base	5 min	\$32.5B/day	\$4.88B/day	\$50.8M
Base	30 sec	\$32.5B/day	\$4.88B/day	\$5.08M
Aggressive	60 min	\$104.1B/day	\$31.2B/day	\$3.90B
Aggressive	5 min	\$104.1B/day	\$31.2B/day	\$325M
Aggressive	30 sec	\$104.1B/day	\$31.2B/day	\$32.5M

Table 6: 2030 seam inventory requirements under cross-border baseline ($F_{\text{CB,day}} = \$520.5\text{B/day}$, $k = 3$). Inventory scales linearly with daily flow and refresh interval.

τ	L_{USD}	**at \$2.30/XRP**	**at \$10/XRP**
60 min	\$610M	265M XRP	61M XRP
5 min	\$50.8M	22M XRP	5.1M XRP
30 sec	\$5.08M	2.2M XRP	0.51M XRP

Table 7: Base scenario token requirements at different XRP prices. Higher prices reduce token count for equivalent USD liquidity.

10.7 Token requirements at illustrative price points

Converting Base scenario requirements to token counts at three price points:

10.8 Alternative baseline: FX-scale seams

If we instead anchor to CLS-scale settlement magnitudes, using $F_{\text{CLS,day}} = \$7.9\text{T/day}$ with $(a, s, b) = (0.20, 0.20, 0.10)$:

$$F_{\text{seam,day}} = \$316\text{B/day}, \quad F_{\text{XRP,day}} = \$31.6\text{B/day} \tag{24}$$

τ	$L_{\text{USD}}(\tau)$	**XRP at \$2.30**
60 min	\$3.95B	1.72B XRP
5 min	\$329M	143M XRP
30 sec	\$32.9M	14.3M XRP

Table 8: FX-scale seam inventory under CLS baseline with $(a, s, b) = (0.20, 0.20, 0.10)$ and $k = 3$.

10.9 Network-scale fragmentation effects

For a network with n distinct seams (connection points between venues), total bridge asset requirements scale as:

$$\text{Total Network Inventory} = n \times L_{\text{per-seam}}(\tau) \tag{25}$$

Illustrative network sizing (Base scenario, 5-min refresh):

- Per-seam inventory: 22M XRP (\$50.8M)
- Network with $n = 100$ seams: $100 \times 22\text{M} = 2.2\text{B}$ XRP
- As percentage of total supply: $\frac{2.2\text{B}}{100\text{B}} = 2.2\%$
- Network with $n = 500$ seams: $500 \times 22\text{M} = 11\text{B}$ XRP (11% of supply)

This demonstrates how network topology (n) and operational velocity (τ) jointly determine aggregate bridge asset requirements.

10.10 Critical constraint: JLR floors remain binding

The inventory proxy above represents a *necessary but not sufficient* condition. The Jackson Liquidity Requirement for any corridor remains:

$$\text{JLR} = \max\left\{R^{\text{slip}}, R_A^{\text{VaR}} + R_B^{\text{VaR}}, R_A^{\text{intraday}} + R_B^{\text{intraday}}, R^{\text{Basel}}\right\} \tag{26}$$

Even with rapid refresh, individual corridors cannot operate below their slippage floor or tail-risk buffers. Furthermore, fragmentation penalties scale linearly for equal pool splits:

$$\text{JLR}_{\text{fragmented}} = n \times \text{JLR}_{\text{unified}} \tag{27}$$

Therefore, the seam-level throughput calculations must be validated against corridor-level JLR constraints. Real deployment requires:

1. Full corridor-level JLR sizing for representative asset pairs
2. Venue topology optimization to minimize fragmentation
3. Verification that refresh intervals can be sustained operationally
4. Confirmation that slippage constraints do not dominate (making velocity irrelevant)

10.11 Key findings from the thought experiment

1. **Velocity creates orders-of-magnitude differences:** The same daily flow ($4.88B) requires 265M XRP at 60-minute refresh but only 2.2M XRP at 30-second refresh under Base assumptions.
2. **Network fragmentation amplifies requirements:** Moving from 100 to 500 seams multiplies total inventory by 5× under linear fragmentation.
3. **Slippage floors create hard boundaries:** Velocity improvements cannot reduce corridor requirements below slippage constraints, limiting the effectiveness of speed alone.
4. **Scenario sensitivity is extreme:** Varying adoption (a), seam share (s), and routing share (b) produces inventory estimates ranging from millions to billions of tokens — illustrating why these parameters must be measured, not assumed.

Methodological conclusion: Bridge asset demand for a mature tokenized settlement system cannot be estimated from adoption narratives alone. Rigorous sizing requires: (1) baseline flow data, (2) explicit scenario parameters, (3) operational refresh interval constraints, (4) corridor-level JLR validation, and (5) network topology mapping. This framework provides the mathematical structure for such analysis.

11 Conclusion and Next Research Deliverable

This paper establishes a quantitative framework for calculating the velocity threshold in fragmented AMM architectures. The core result — that refresh intervals must scale quadratically with pool count to neutralize fragmentation penalties — creates a sharp engineering constraint for system designers.

The framework identifies two critical parameters that govern bridge asset demand:

1. **Regime determination:** Whether the system operates in a buffer-limited regime (where velocity helps) or a slippage-limited regime (where velocity provides no benefit)
2. **Refresh interval feasibility:** Whether the system architecture can sustain the required $\tau^* = T_0/n^2$ threshold

For practical application, these theoretical results must be mapped onto real system architectures. Bridge assets like XRP, operating in tokenised settlement or CBDC interchange contexts, can be evaluated by measuring:

- Settlement finality times and batch processing intervals (determines τ)
- Trade size distributions and execution quality requirements (determines R_{slip})
- Pool count and fragmentation structure (determines n)

Next deliverable: A Refresh Interval & Slippage Regime Map for emerging tokenised settlement architectures, providing empirical estimates of τ and constraint binding across institutional payment rails.

References and Source Links (selected)

- Jackson Liquidity Framework (JLF) — Technical Documentation:
  https://www.lewisjacksonventures.com/research/jackson-liquidity-framework
- BCBS 248 — Monitoring tools for intraday liquidity management (BIS):
  https://www.bis.org/publ/bcbs248.pdf
- BIS Triennial Central Bank Survey (2025) — OTC FX turnover:
  https://www.bis.org/statistics/rpfx25.htm
- CLS Group — Settlement volume statistics:
  https://www.cls-group.com/
- IMF & G20 (Oct 2024) — Financial platforms report:
  https://www.imf.org/-/media/files/research/imf-and-g20/2024/g20-report-2024-financial-platf
  pdf
- BIS Project Mariana (2023) — AMM-based FX for wholesale CBDC:
  https://www.bis.org/publ/othp75.pdf
- CPMI (May 2022) — Extending and aligning payment system operating hours:
  https://www.bis.org/cpmi/publ/d203.pdf
- CPMI (Jul 2022) — Liquidity bridges across central banks:
  https://www.bis.org/cpmi/publ/d209.pdf
- Deutsche Bank / Project DAMA 2 litepaper (Jun 2025):
  https://cdn.mementoblockchain.com/dama2/pdfs/DAMA-2-Lite-Paper_Jun172025.pdf