What would it take to operationalize UTXO-based settlement for central bank digital currency?

Dr. Silvio Petriconi, Architecture and Security Chapter, Digital Euro Division, Deutsche Bundesbank

Disclaimer: All views expressed here are my own and do not necessarily represent the views of Deutsche Bundesbank.

Sofia, 28. March 2025

Central Bank Digital Currency: Status of the Digital Euro Project

Oct 2021- Oct 2023

Investigation Phase

- Concept definition
- Technical exploration
- Design proposal

Nov 2023 – Oct 2025

Preparation Phase

- Finalizing the Scheme Rulebook
- Selecting Service Providers
- Learning by experimentation
- Deep dive into technological aspects, including offline functions

From Nov 2025

Next Phase

- Potentially developing a Digital Euro CBDC
- Preparation for roll-out

A decision to issue a digital euro will only be considered by the ECB once the European Union's legislative process has been completed.

Today‘s talk is NOT about the Digital Euro

- Decision whether to emit Digital Euro **has not been taken yet.** Technological design of the actual solution will be **unrelated** to today‘s talk.
- Today‘s talk summarizes the findings of extensive **independent exploration work** regarding a specific technology**:**
- *„What would be the key issues if one were to emit digital currency as UTXO tokens that are secured by cryptographic primitives?“*

Today‘s talk offers questions, not answers

- I‘d like to share some new issues that arise in the context of implementing a CBDC with cryptographic UTXO tokens.
- My hope is that this may inspire thinking and research.
- If you think that any of these issues have better solutions, or if you are generally enjoying the topic, we‘d love to hear from you!

UTXO: Unspent Transaction Outputs as cryptographic means of payment

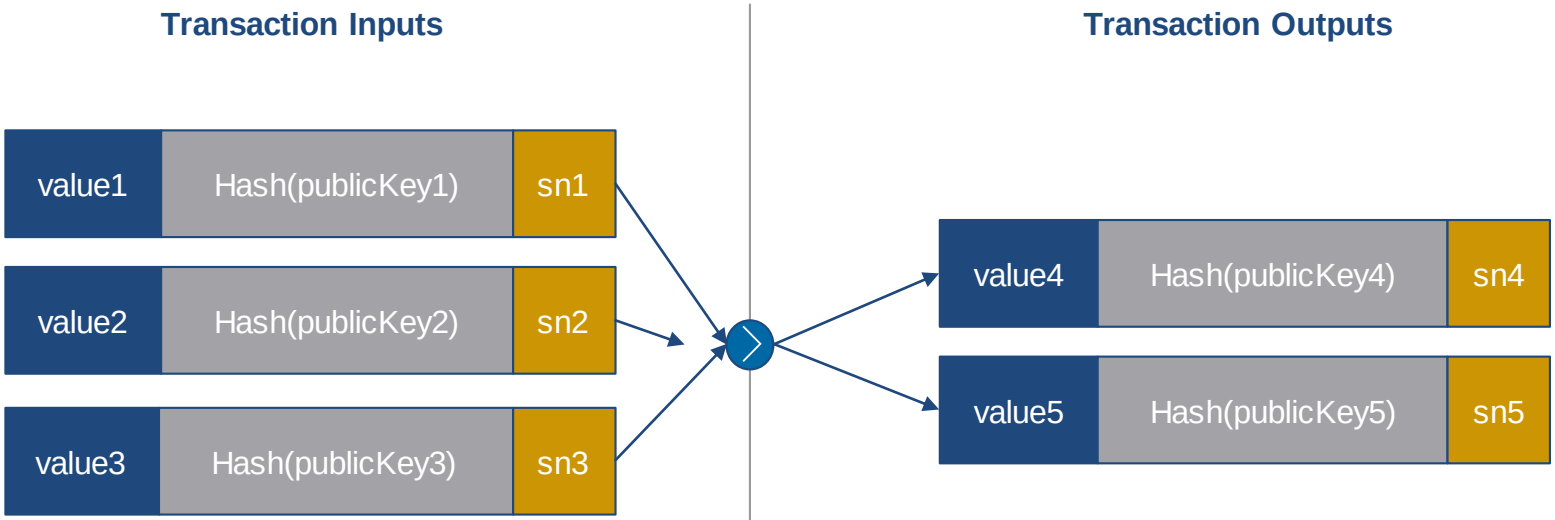

UTXO structure: tuple of

(value, commitment to unlock condition, serial no)

UTXO: Unspent Transaction Outputs as cryptographic means of payment

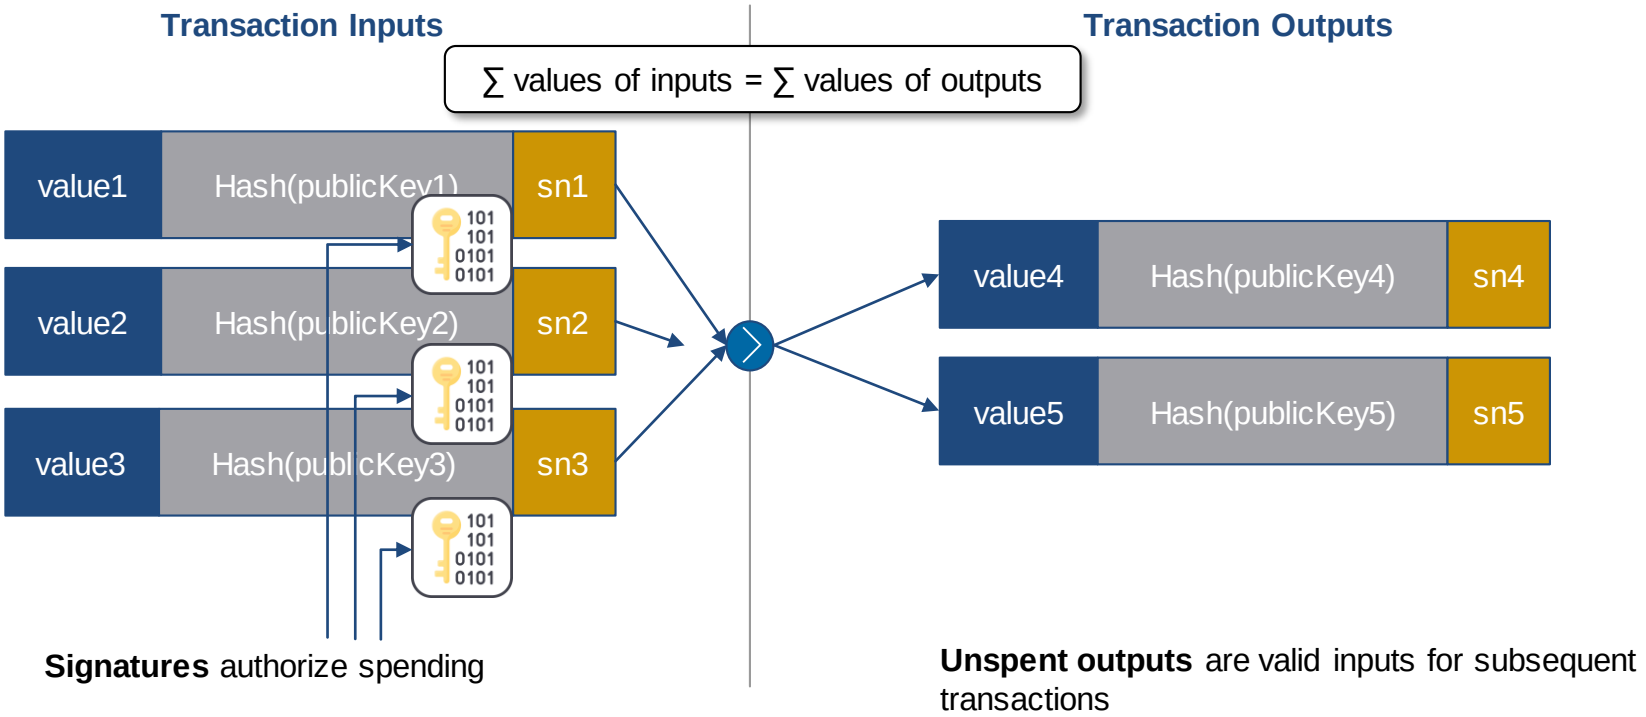

Why a UTXO data model might make sense for a digital currency (at first glance)

Some **advantages** of UTXO over the account model:

- Immutability of UTXOs helps for **better concurrency and scalability** of settlement. Little, if any, contention!
- **Higher privacy (when combined with other PETs)** thanks to one-to-many relationship between users and addresses
- Flexible **spending authorization, custody** and **interoperability** models**.**

Specifically, openCBDC by MIT DCI & Boston FED:

1. **Demonstrated >1.7 million tx/s** in UTXO-based open source settlement core.
2. Showed that **central ledger only needs to record cryptographic commitments** of unspent tokens
3. **Auditing of money supply remains nevertheless possible** when using homomorphic encryption & zk proofs.
4. **Low latency** even in geo-replicated deployment

So, what‘s the catch?!

What would it take to operationalize UTXO-based settlement for central bank digital currency?

Key aspects in which CBDCs fundamentally differ from permissionless blockchains

1. **No blockchain at all, or at least: No ledger that is publicly accessible**.
   Good reasons! Trust model, throughput, latency

2. **Must support holding limits** to prevent potentially catastrophic financial disintermediation („digital bank run“).

3. Wallets may auto-fund themselves in real time from commercial bank money sources („reverse waterfall“). Excess holdings must convert „near-instantly“ to commercial bank money („waterfall“). => **Complex, high-frequency funding and de-funding scenarios**.

4. **Acceptable latencies, and time to finality:** milliseconds, not minutes

5. **Regulatory compliance** e.g., anti money laundring and anti-fraud; end-of-day accounting of intermediary liquidity

6. **Mature technologies only:** „fail fast, fail often“ is not an option

The easy part: Prototyping a centralized UTXO settlement engine

UTXO Settlement Logic:

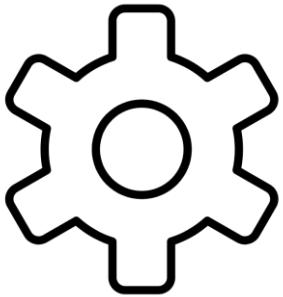

- Verify signatures
- Validate that no money is created or destroyed in a transaction:

  ∑ values of inputs = ∑ values of outputs

- Check in DB that inputs are currently unspent
- Mark inputs as spent and creates new unspent outputs

Implementation sketch:

- Distributed KV store with geo-replication support underneath the solution

  Excellent open source choice option (for experimentation): TiKV

- Stateless processors, scalable via Kubernetes; metrics in Prometheus
- Runs > 30.000 tx/s out of the box

Up to here: Easy!

A hard problem: There‘s no blockchain – how do you notify wallets?

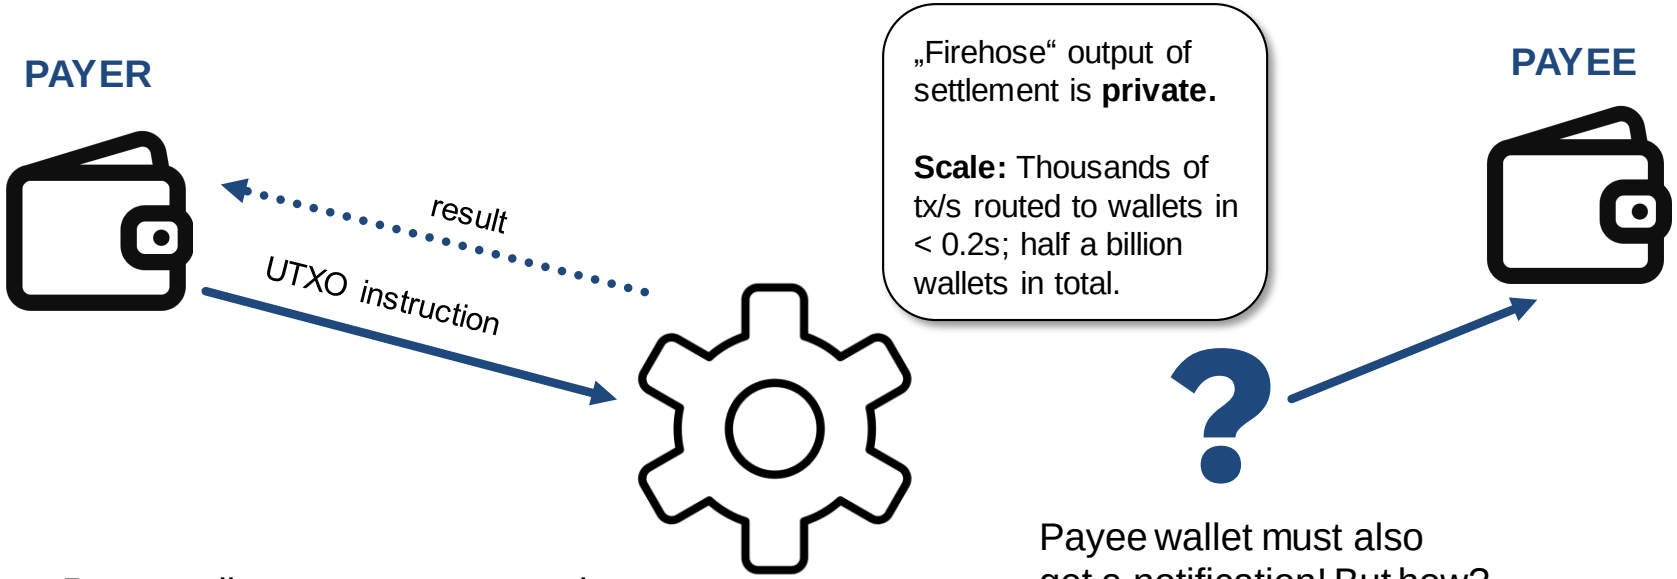

Payer wallet can connect e.g. via REST API to settlement and await its result

Payee wallet must also get a notification! But how?

For privacy, wallets can‘t directly disclose their set of addresses!

Incomplete list of how we‘ve considered to solve this

X

Today: Payee‘s bank BIC code is part of every valid payment instruction.
Commercial banks receive payment outcomes on behalf of their customers.

Please let‘s do better than this!

Incomplete list of how we‘ve considered to solve this

X **Today: Payee‘s bank BIC code is part of every valid payment instruction.**
Commercial banks receive payment outcomes on behalf of their customers.

1 **A privacy-preserving routing network**
Bloom filters, onion routing, etc – can this scale to > 10 billion addresses?

2 **Centralized DID-based service directory**
Intermediaries (or other service providers) take role of dispatchers of information towards retail wallets. Correct intermediary is identified in centralized DID document service directory.

More ideas are higly welcome!

Holding Limits

„Each wallet shall not hold more than X units of digital currency at any given point in time.“

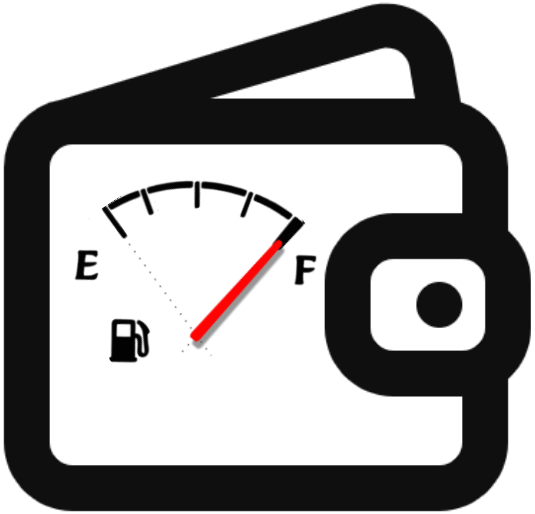

1. **Trusted third party (intermediary) controls wallet**
   Of course this works, but innovation incentives become misaligned!

2. **Cryptographic protocol to prove inventory below holding limit?**
   If wallets are trustless – how can you be sure they tell the truth and count all their coins, not just some of them?

An approach under investigation:

- Merkle tree to commit to set of assets under management
- Homomorphic cryptography (Pedersen commitments) + zk

- Issues: latencies, and lack of confidence in newer zk techniques

Much more work is needed! Confidence in the solution essential to making it viable.

High Performance Funding Wallets for Intermediaries

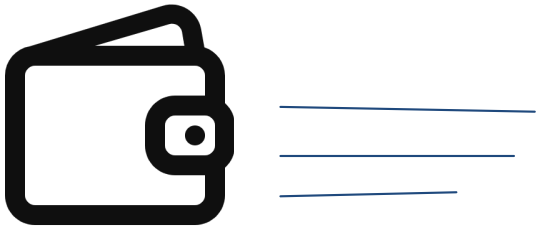

1. **Token selection is crucial to scale well and avoid dust**
   Need algorithms that provably sustain a stable distribution of token denominations under wide range of operating conditions.

2. **Token selection must be efficient and must always succeed**
   With thousands tx/s throughput required from a single wallet, token selection can‘t optimize every individual step, as in BTC wallets

3. **End-of-day balance consolidation must be supported**
   For monetary policy, consistent intermediary EoD balances are required: „simultaneous“ snapshot of all intermediary wallet amounts.

Not cryptography problems, but important for UTXO to work

The biggest open issue: Money Laundring and Fraud Detection

(this problem is not specific to UTXO)

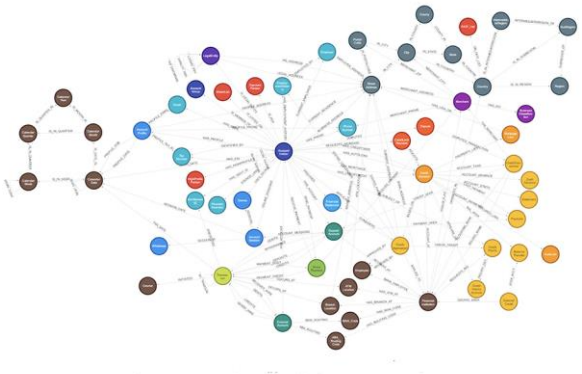

Instant Settlement attracts fraud:	Experiences (e.g. Brazilian real-time payment system) have shown that fraud can quickly become pervasive
„Heisenberg" principle of fraud & ML:	**Patterns change continuously** to evade detection. For this reason, not much public data sources, little published academic research.
Global view of all activity might be needed to identify fraud & ML:	**Graph centrality measures** or similar global properties are superior predictors but are typically constructed from a global view of all activity.
Can one do it reliably w. federated learning and MPC? **How to reveal circles in payment graphs in privacy-respecting ways?**	**<u>Your ideas here could greatly help to attain higher privacy without sacrificing security.</u>**

Thank you!

Please always feel free to reach out to me:

silvio.petriconi@bundesbank.de