

# Accelerated As-Of Reporting, *Without The Hassle*

OSinFinance Webinar Series, December 2024



Jeremy Taylor | JUXT | @refset





# JUXT

A GRID DYNAMICS COMPANY

Delivering software engineering consultancy for financial services and beyond, since 2013.

<https://juxt.pro>

JUXT | A GRID DYNAMICS COMPANY

TEAM

FINANCIAL

DELIVERY

GET STARTED

CAREERS

BLOG

XTDB

GET IN TOUCH

## FINANCIAL SERVICES

- STRUCTURED PRODUCTS
- STRATEGIC RISK PLATFORM
- POST TRADE MANAGEMENT
- CLOUD MIGRATION BLUEPRINTS
- COMMODITY TRADING

## Structured Products

JUXT's software development experts were selected to complement the inhouse engineering team, to build the award-winning, publicly accessible Structured Notes platform for a major Tier 1 investment bank.

Entirely data-driven, the platform can rapidly onboard new Structured Products with zero code changes. The user-interface automatically adapts to the shape of the data required for each individual product.

The platform supports calendar deals, bespoke pricing, booking, and full post trade analysis, including for products booked outside of the platform. It provides comprehensive language localization, with an interface for conforming all fields and menu items to regional & desk-specific standards.

The platform has won industry awards for over 5 consecutive years, including 'Best Structured Product Technological Solution', and Bankers Investment Banking Award.

*"Best dev team I have ever worked with. You are so good at what you do and you made my job easy."*

*Product Owner*

## Technologies

Java

Elasticsearch

JSON Schema

React

ClojureScript

## Awards





# Common problems observed

- ✗ High costs to deliver timely auditability and reproducibility across bespoke and off-the-shelf software
- ✗ Non-trivial ETL/ELT work required to provide basic reporting and analytics over line-of-business applications
- ✗ Inflexible SQL databases struggle to accommodate complex domain models with evolving schemas



# Common workarounds for these problems

- ✗ Excessive logging - not easy to retrieve or analyze
- ✗ Database snapshots - costly, infrequent and delayed
- ✗ Change Data Capture - excessive duplication of schema and queries in a dedicated analytics system
- ✗ Custom in-database versioning - complex schema designs risk silent data loss and reduce agility




# Common solutions delivered, using XTDB

*(fintech and beyond!)*



We are bringing  
**decision intelligence**  
to cardiovascular  
care.



THE **LINUX** FOUNDATION PROJECTS



XTDB OMRS Repository Connector



Rijksoverheid



OpenKat - Vulnerability Analysis Tool

2019

Now

📄 [xtdb / xtdb](#) Public

An immutable SQL database for  
application development, time-travel  
reporting and data compliance.  
Developed by @juxt

🔗 [xtdb.com](#)


📜 MPL-2.0 license

⭐ 2.6k stars 🍴 170 forks



# Why does As-Of Reporting matter?

Managers,  
auditors,  
regulators



Evolving data from multiple sources

Decision makers (...AIs?)

- ✓ Auditability (i.e. tracing decisions to data)
- ✓ Proving value (e.g. best execution)
- ✓ Accurate insights into business performance



DATABASE

DEC 09, 2024

# Bitemporal TraderX - augmenting a sample trading system

Adding XTDB and Clojure into a sample microservices trading application for the FINOS 2024 Tech Sprint



**Michal Pisanko**  
Software Engineer

The screenshot displays a trading application interface. At the top, a timeline shows various events labeled v9 through v22. Below this, there are two tables: 'Trades' and 'Positions'. The 'Trades' table has columns for STATE, UNIT PRICE, and other trade-related fields. The 'Positions' table has columns for SECURITY, QUANTITY, MONEY IN/OUT, MARKET VALUE, and CALC. To the right of the tables, there are three code snippets in Clojure, each enclosed in an orange box. The first snippet defines a function 'select-points-in-time' that queries trades for a specific account. The second snippet defines a function 'account-prices-at' that queries stock prices for a specific account. The third snippet is a query for positions. Arrows point from the code snippets to the corresponding data in the tables.

```
(def select-points-in-time
  "select distinct(_valid_from) as start
  from trades
  for all valid_time
  where account_id=?
  order by start")
```

```
(if as-of
  ["select _id as id, security, trade, value, quantity,
  from positions
  for valid_time as of ?
  where account_id = ?
  order by security,
  _valid_from desc,
  _valid_to desc nulls last"
  (parse-date-time as-of)
  account-id])
```

```
(def account-prices-at
  "select sp._id as ticker, sp.price
  from stock_prices for valid_time as of
  join positions p on sp._id = p.security
  where p.account_id = ?")
```

finos / **traderX** Public  
generated from [finos-labs/project-blueprint](#)

Notifications Fork 61 Star 59


<> Code Issues 30 Pull requests 15 Discussions Actions Projects Security Insights

## Enhance TraderX Database Performance and Scalability #187

lucaborella89 started this conversation in Ideas

lucaborella89 on Jul 2 Maintainer

### TraderX Support for Bitemporal Data



**TraderX**  
Support for  
Bitemporal Data

Category

Ideas

Labels

use case

2 participants



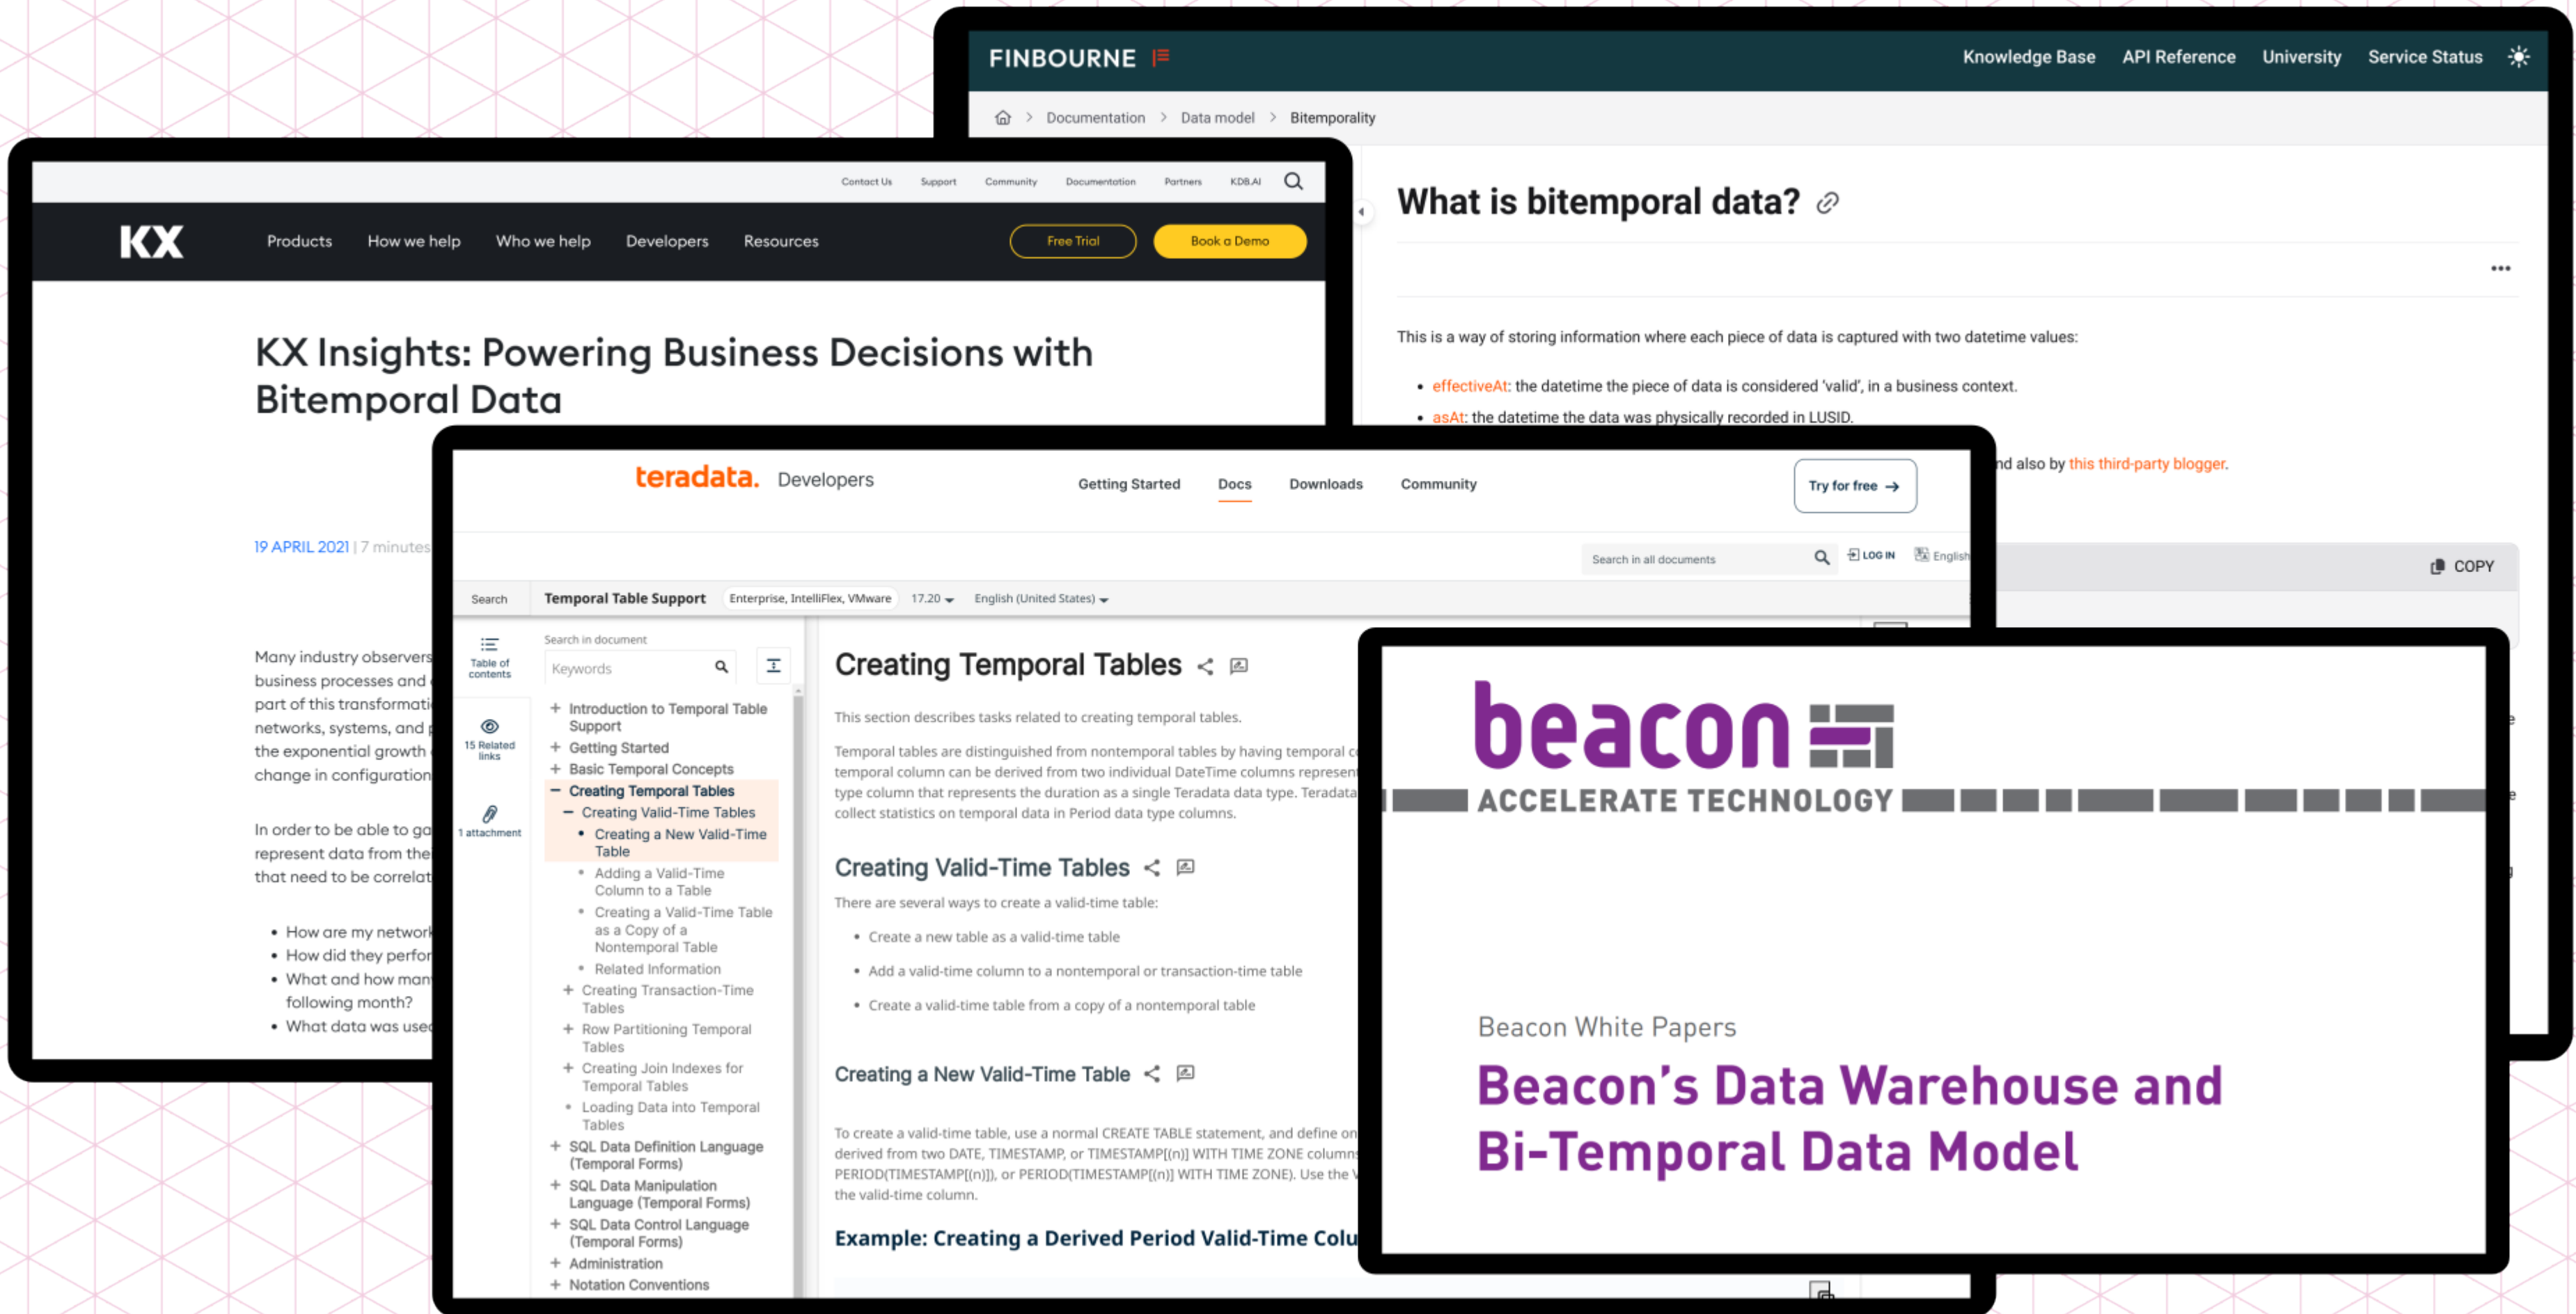
## Joining FINOS: The Fintech Open Source Foundation

If you missed the news, JUXT is now a member of [FINOS](#)! The Fintech Open Source Foundation


<https://www.juxt.pro/blog/bitemporal-traderx/>




# There are many approaches to meeting As-Of requirements





Platform ▾ Solutions ▾ Customers Resources ▾ Developers ▾ Pricing

 ▾ [CONTACT SALES](#) [START FOR FREE](#)

BLOG

CATEGORY ▾

PRODUCT AND TECHNOLOGYMAY 13, 2024

# Accelerate Your Time Series Analytics with Snowflake's ASOF JOIN, Now Generally Available

```
SELECT t.stock_symbol, t.trade_time, t.quantity, q.quote_time
FROM trades t ASOF JOIN quotes q
      MATCH_CONDITION(t.trade_time >= quote_time)
      ON t.stock_symbol=q.stock_symbol
ORDER BY t.stock_symbol;
```

STOCK_SYMBOL	TRADE_TIME	QUANTITY	QUOTE_TIME
SNOW	2023-10-01 09:00:05.000	2000	2023-10-01 09:00:05.000
SNOW	2023-10-01 09:00:05.000	1000	2023-10-01 09:00:05.000
SNOW	2023-10-01 09:00:10.000	1500	2023-10-01 09:00:10.000

# Accelerate Your Time Series Analytics with Snowflake's ASOF JOIN, Now Generally Available

At Snowflake, we're committed to helping our customers achieve their business goals. That's why we are excited to announce our new, highly performant Time Series feature that

ASOF JOIN is a type of join that pairs  
on the left side of the join, the opera

```
SELECT t.stock_symbol, t.trade_time, t.quantity, q.quote_time
FROM trades t ASOF JOIN quotes q
  MATCH_CONDITION(t.trade_time >= quote_time)
  ON t.stock_symbol=q.stock_symbol
ORDER BY t.stock_symbol;
```

STOCK_SYMBOL	TRADE_TIME	QUANTITY	QUOTE
SNOW	2023-10-01 09:00:05.000	2000	2023-
SNOW	2023-10-01 09:00:05.000	1000	2023-
SNOW	2023-10-01 09:00:10.000	1500	2023-

Support

Do you have time series data that you want to join, but the timestamps don't quite match? Or do you want to look up a value that changes over time using the times in another table? And did you end up writing convoluted (and slow) inequality joins to get your results? Then this post is for you!


One of the problems that AsOf Joins are used to solve is finding the value of a varying property at a specific point in time. This use case is so common that it is where the name came from: *Give me the value of the property as of this time.*

## ASOF Join — The "Do What I Mean" of the Database World



Dealing with time-series data means dealing with events — pieces of data that say "*this* thing happened at *this* moment". Describing "the thing" involves users, merchants, physical sensors, and so on. We denote each of these with an ID, so an event is basically a record full of IDs. To do anything meaningful with it, you must expand the ID into the info you have on that thing/person/organization. You keep this in some tables, so naturally you want a JOIN to bring it all in.

However, this info changes over time, and you have to keep the full history on it. You need the history so that you can pull in the particular info that was valid for an event *as of* the time it occurred, like this:



### ASOF JOIN illustrated



`asof left join` = `left join lateral` sugar + execution speedup

```
/* Lateral */ I
select
  users.user_id,
  users.last_review_datetime,
  latest_events.event_id,
  latest_events.event_datetime,
  latest_events.event_type
from users
  left join lateral (
    select event_id, event_datetime, event_type
    from events
    where users.user_id = events.user_id
      and users.last_review_datetime >= events.event_datetime
    order by events.event_datetime desc
    limit 1
  ) as latest_events on true
order by users.user_id, latest_events.event_id
```

	␣user_id ␣	␣last_review_datetime ␣	␣event_id ␣	␣event_datetime ␣	␣event_type ␣
1	1	2023-12-15 03:20:00	<null>	<null>	<null>
2	2	2024-01-24 03:30:00	32	2024-01-23 15:00:00	logout
3	3	2024-02-01 04:10:00	38	2024-01-31 21:00:00	login
4	4	2024-02-04 02:50:00	<null>	<null>	<null>

See also "row\_number () over"

```
/* DuckDB */
select
  transactions.date,
  transactions.account,
  transactions.amount,
  exchange_rates.rate,
  transactions.amount * exchange_rates.rate as amount_usd,
from transactions
  asof left join exchange_rates
  on transactions.currency = exchange_rates.from_currency
  and exchange_rates.to_currency = 'USD'
  and transactions.date >= exchange_rates.date
order by
  transactions.date,
  transactions.amount
```



# Bitemporal Databases




Figure 3: Historical Relation




Figure 4: Temporal Relation



**SAFE**

**& CORRECT**



# Sources of **essential** complexity:



**Sources of essential complexity:**

**Your business logic.**



# Sources of incidental complexity:

Legacy Systems Integration  
Regulatory Compliance and Audit Requirements  
Complex Security and Access Control  
Data Consistency and Synchronization Across Systems  
High Availability and Low Latency Requirements  
Changing Business Requirements  
Vendor and Third-party Dependencies  
Inconsistent Data Models  
Cross-jurisdictional Requirements  
Distributed Teams and Communication Challenges  
Version Control and Change Management

Dependency on Batch Processing Systems  
Complex Infrastructure and Cloud Migration  
Tooling Fragmentation  
Unclear or Shifting Responsibilities  
Lack of Unified Documentation  
Strict Performance Optimization Requirements  
Testing and QA Challenges  
Multi-language Support  
Complex Release Management  
Governance and Risk Management Overheads  
Cultural and Organizational Resistance

...



# Sources of incidental complexity:

Legacy Systems Integration

Regulatory Compliance and Audit Requirements

Complex Security and Access Control

Data Consistency and Synchronization Across Systems

High Availability and Low Latency Requirements

Changing Business Requirements

Vendor and Third-party Dependencies

Inconsistent Data Models

Cross-jurisdictional Requirements

Distributed Teams and Communication Challenges

Version Control and Change Management

Dependency on Batch Processing Systems

Complex Infrastructure and Cloud Migration

Tooling Fragmentation

Unclear or Shifting Responsibilities

Lack of Unified Documentation

Strict Performance Optimization Requirements

Testing and QA Challenges

Multi-language Support

Complex Release Management

Governance and Risk Management Overheads

Cultural and Organizational Resistance

...



**Bitemporal =  
Ingestion timestamp +  
Reporting timestamp**




Bitemporal =

Ingestion timestamp +

Reporting timestamp


Safe ("immutable")  
*The timeline of the system*



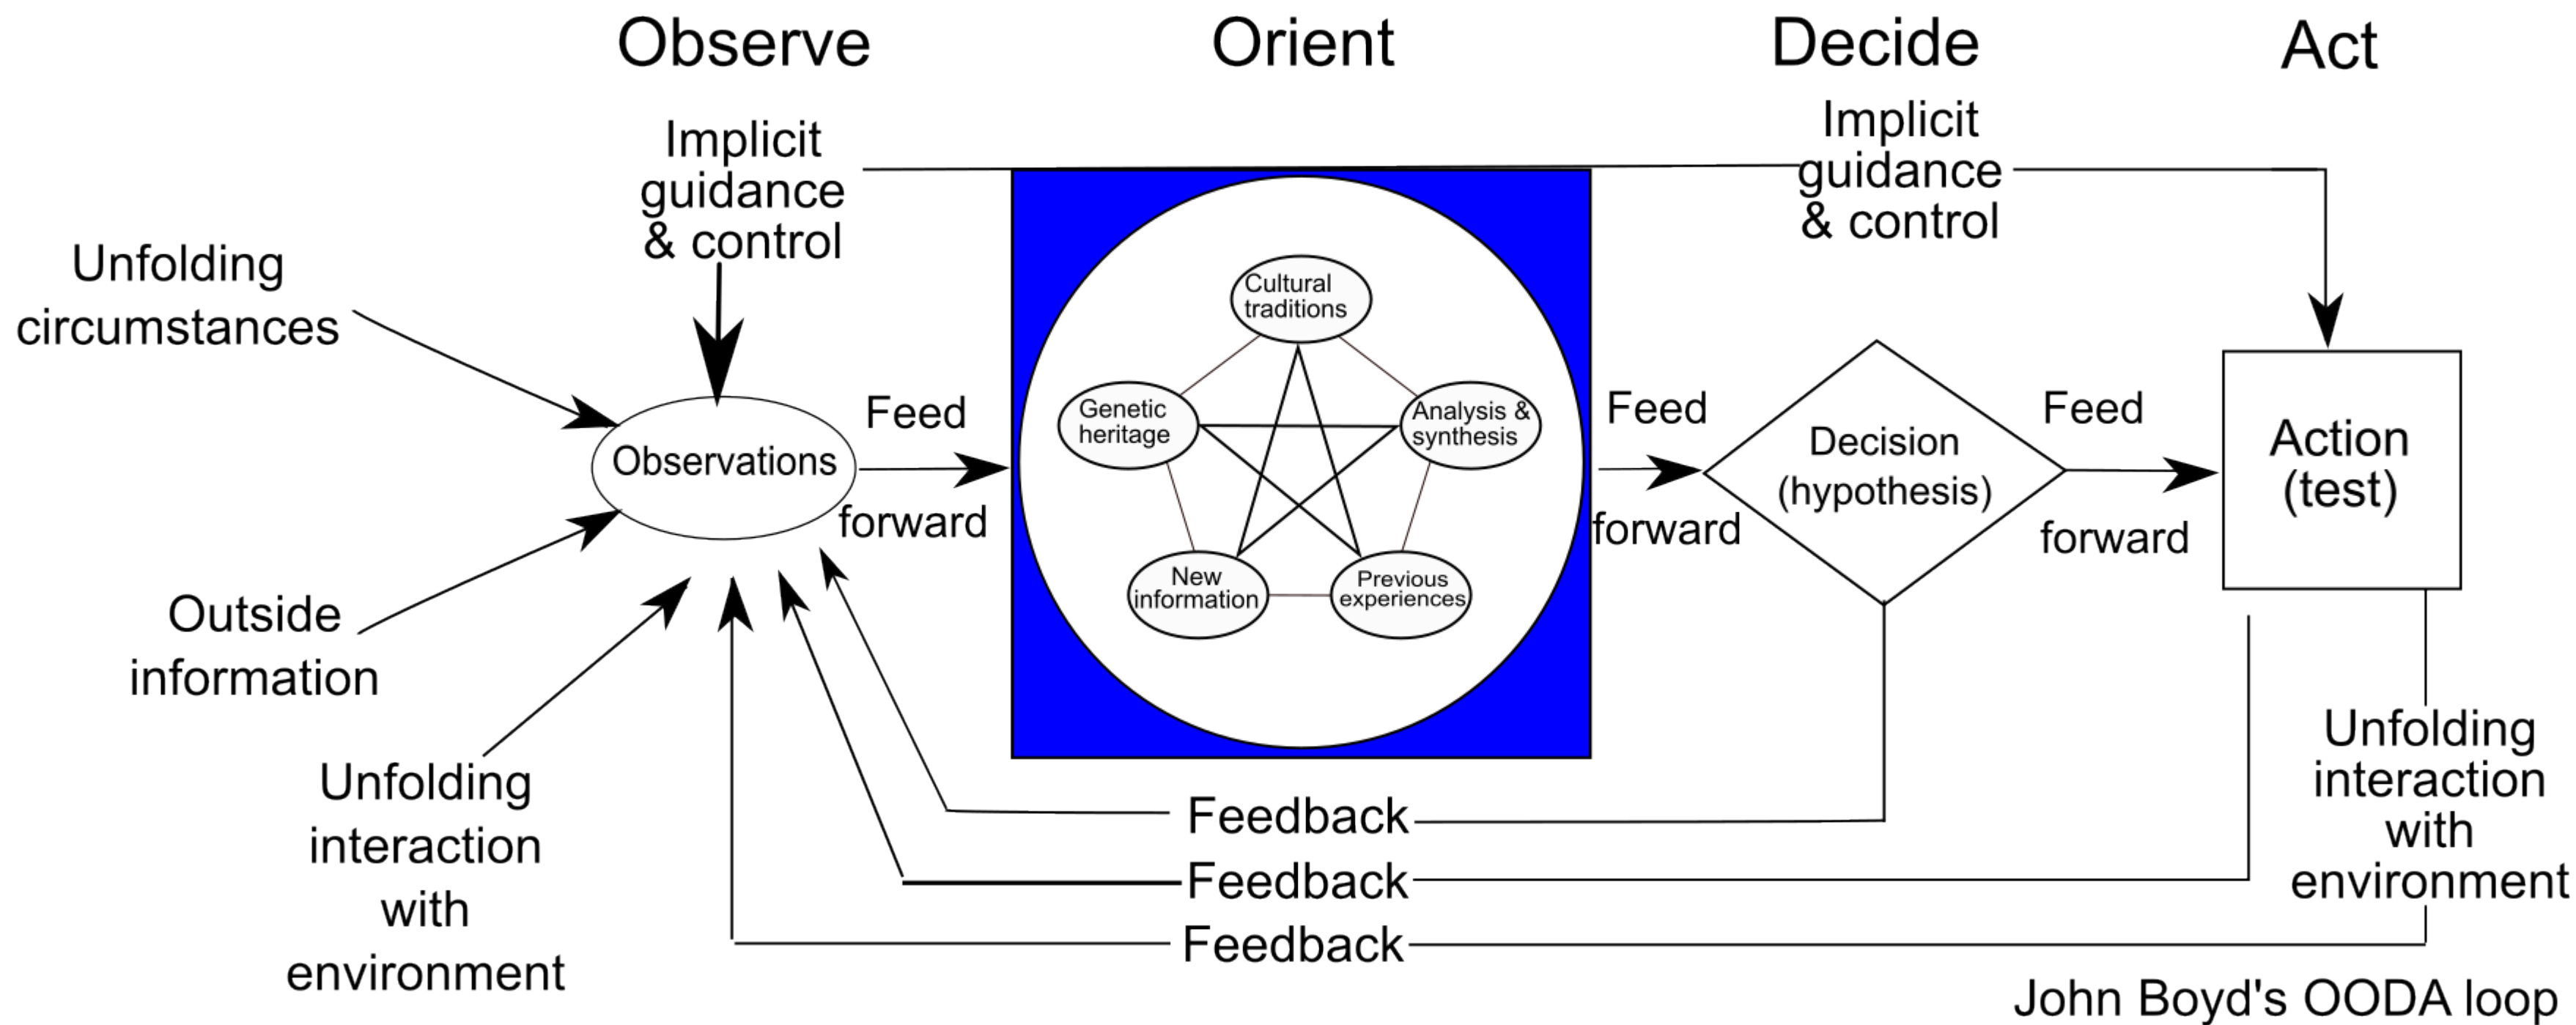
Correct  
*The timeline of the data*













# Challenge 1: Capturing Derivative Trades

TRS on QIS & QIS / Cashflow Management etc

## Business Challenge

Our clients would like to build define and execute swap wrappers within our platform.



## Technical Challenge

Redefine the wheel. Danger of building technical debt and shoehorning payoffs. Bi-temporal challenges



## Future Proofing

The unknown / unknowns. What might we have to capture. What might we have to report on





FINBOURNE

LIBERATE. SIMPLIFY. CONNECT.

Products

Solutions

FINights

Company

Careers

Client hub

Contact us

Unlock the True Power of Your Data

At FINBOURNE, we help global financial services harness new value from data by going beyond traditional data management.

Extract meaningful insights, reduce system complexity and create innovative solutions from your data with FINBOURNE's Modern Financial Data Stack.

Get started

Speak to an expert

The Modern Financial Data Stack

Solutions For Global Financial Services

API-first

Cloud-native

Extensible data model

Bitemporal

Asset Managers

Alternative Asset Managers

Asset Owners

Asset Servicers

Banking & Capital Markets

Limina

USE CASES

PRODUCT

PLATFORM

SERVICES

PRICING

INSIGHTS

ABOUT

SCHEDULE MY DEMO

IMS

PMS

OMS

Compliance

Risk

IBOR

ABOR

IDM

A consolidated IBOR platform to enable scalable operating models

The Investment Book of Records (IBOR) concept was popularised a decade ago. Now, ten years later, the [benefits of establishing an IBOR are clear](#):

Improved basis for decision-making

Operational efficiency

Improved governance and oversight

Avoidance of compliance issues from different teams looking at disjointed data sets

Golden data source to power other in-house or external solutions

IBOR – Position Extraction

IBOR generated position extract on state "Contractual"

Order Placed

Trade Executed

Trade Confirmed

Trade Settled

Estimated

Known

Contractual

Physical

Impact of Order

Impact of Trade

Impact of Confirm

Impact of Settlement

Sale Trade

First Forecast

Declaration

Ex-Date

Payment

Estimated

Known

Contractual

Physical

Impact of Forecast

Impact of Declaration

Impact of Ex Date

Impact of Payment

Dividend

Fee Estimated

Fee Fixed

Invoice Received

Adjustment

Invoice Paid

Estimated

Known

Contractual

Contractual

Physical

Impact of Estimate

Impact of Fixing

Impact of Invoice

Impact of Adjustment

Impact of Payment

Fee

Historic Reconstruction

Future Projection

Time

Today (as-of)


beacon

ACCELERATE TECHNOLOGY

Beacon White Papers

Beacon's Data Warehouse and Bi-Temporal Data Model





Streaming Data Processing Systems  
Tyler Akidau, Google

There are many  
relevant tools  
for developers  
and data teams...

goldmansachs / reladomo Public

Reladomo is an enterprise grade object-relational mapping framework for Java.

Apache-2.0 license

380 stars 93 forks Branches Tags Activity

DRW

scalegenius / pg\_bitemporal Public

Bitemporal tables in Postgres

BSD-3-Clause license

139 stars 30 forks Branches Tags Activity



Documentation Community Blog Github

Colab Demo

## ArcticDB is a DataFrame Database

Built for the modern **Python Data Science** ecosystem,  
**ArcticDB** transforms your ability to handle complex real  
world data with Incredibly fast **proven Petabyte scale**.

Get ArcticDB Find the right product >

QuestDB

Product

Enterprise

Market Data

Docs

Blog

8.1.1



DOWNLOAD

Join order book data with core prices

## Advanced data correlation with ASOF JOIN

Strong time-based analysis of financial data shouldn't be complicated. Consider two tables: one composed of snapshots of recent financial data, and another composed of core prices from your order book. Use the powerful

**ASOF JOIN** extension to query across both tables in milliseconds. It's blazing fast, efficient and clean.

QuestDB SQL

```
CREATE TABLE 'market_data_snapshot' (  
  venue SYMBOL,  
  symbol SYMBOL,  
  ts TIMESTAMP,  
  bid_1 DOUBLE,  
  bid_qty_1 DOUBLE,  
  [...]  
  bid_20 DOUBLE,
```

QuestDB SQL

```
CREATE TABLE 'core_price' (  
  send_ts TIMESTAMP,  
  ccy1 SYMBOL,  
  ccy2 SYMBOL,  
  valid BOOLEAN,  
  mid DOUBLE,  
  bid DOUBLE,  
  ask DOUBLE,
```



# Is your database pulling its weight?

## SQL:2011

🌐 5 languages ▾

Article [Talk](#)

[Read](#) [Edit](#) [View history](#) [Tools](#) ▾

From Wikipedia, the free encyclopedia

**SQL:2011** or **ISO/IEC 9075:2011** (under the general title "Information technology – Database languages – SQL") is the seventh revision of the [ISO](#) (1987) and [ANSI](#) (1986) standard for the [SQL database query language](#). It was formally adopted in December 2011.<sup>[1]</sup> The standard consists of 9 parts which are described in detail in [SQL](#). The next version is [SQL:2016](#).

### New features [\[ edit \]](#)

One of the main new features is improved support for [temporal databases](#).<sup>[2][3]</sup> Language enhancements for temporal data definition and manipulation include:

- **Time period definitions** use two standard table columns as the start and end of a named time period, with [closed set-open set](#) semantics. This provides compatibility with existing data models, application code, and tools
- Definition of **application time period tables** (elsewhere called [valid time](#) tables), using the **PERIOD FOR** annotation
- Update and deletion of application time rows with **automatic time period splitting**
- **Temporal primary keys** incorporating application time periods with optional non-overlapping constraints via the **WITHOUT OVERLAPS** clause
- **Temporal referential integrity** constraints for application time tables
- Application time tables are queried using regular query syntax or using new **temporal predicates** for time periods including **CONTAINS** , **OVERLAPS** , **EQUALS** , PRECEDES , **SUCCEEDS** , **IMMEDIATELY** PRECEDES , and **IMMEDIATELY SUCCEEDS** (which are modified versions of [Allen's interval relations](#))
- Definition of **system-versioned tables** (elsewhere called [transaction time](#) tables), using the **PERIOD FOR** SYSTEM\_TIME annotation and **WITH SYSTEM VERSIONING** modifier. System time periods are maintained automatically. Constraints for system-versioned tables are not required to be temporal and are only enforced on current rows
- Syntax for **time-sliced** and **sequenced** queries on system time tables via the **AS OF SYSTEM TIME** and **VERSIONS BETWEEN SYSTEM TIME ... AND ...** clauses
- Application time and system versioning can be used together to provide **bitemporal tables**

# Illuminated Computing

Code like song

## Survey of SQL:2011 Temporal Features

2019-09-04

### Introduction

This blog post is a survey of SQL:2011 Temporal features from MariaDB, IBM DB2, Oracle, and MS SQL Server. I'm working on adding temporal features to Postgres, so I wanted to see how other systems interpret the standard.

If you're new to temporal databases, you also might enjoy [this talk](#) I gave at PGCon 2019.

In this post I cover both application-time (aka valid-time) and system-time, but I focus more on valid-time. Valid-time tracks the history of the thing “out there”, e.g. when a house was remodeled, when an employee got a raise, etc. System-time tracks the history of when you changed the database. In general system-time is more widely available, both as native SQL:2011 features and as extensions/plugins/etc., but is less interesting. It is great for compliance/auditing, but you're unlikely to build application-level features on it. Also since it's generated automatically you don't need special DML commands for it, and it is less important to protect yourself with temporal primary and foreign keys.

At this point all the major systems I survey have *some* temporal support, although none of them support it completely. On top of that the standard itself is quite modest, although in some ways it

### Paul A. Jungwirth



[Blog](#)  
[Email](#)  
[Portfolio](#)  
[Resume](#)  
[Github](#)  
[Stack Overflow](#)



### Code

```
#!/usr/bin/  
# == Synops  
#  
# [touch: F  
#  
# == Usage
```

[Skill Spy](#)  
[Alien Words](#)  
[ElectNext](#)  
[db\\_leftovers](#) gem  
[Tech Notes](#)  
[... more](#)

### Writing



[Survey of SQL:2011 Temporal Features](#)  
[Drawing Redux Form FieldArrays with Pug](#)  
[Validating FieldArrays in Redux Form](#)  
[Testing Your ActionMailer](#)



xtldb.com

5 years

v2.0 Beta

2.6k stars

Built with:



DOCS

SOLUTIONS

SUPPORT

BLOG

DEMO



# The database for our time

XTDB is an immutable SQL database built to simplify application development and data compliance

A SQL database that automatically preserves history, with comprehensive time-travel:

```
SELECT *, _valid_from
FROM product
  FOR ALL VALID_TIME
WHERE product.price > 350
```

INSERT, UPDATE and DELETE versioning enables simple as-of querying and audit:

```
SELECT *, _system_from
FROM product
  FOR SYSTEM_TIME AS OF DATE '2023-01-01'
```

## Languages



- Clojure 70.3%
- Java 13.7%
- Kotlin 12.8%
- ANTLR 0.9%

## Contributors 22



[+ 8 contributors](#)

## Features



Temporal Records



Modern SQL



Semi-structured SQL



Cloud Native



# XTDB is built to solve Data Compliance problems

- Inconsistent reporting over historical data
  - Achieve consistency without snapshots or copying
- Inadequate auditing & compliance
  - Avoid ad hoc solutions that generate additional problems
- Challenging data integration & evolution
  - Manage changing semi-structured data with *gradual* schema

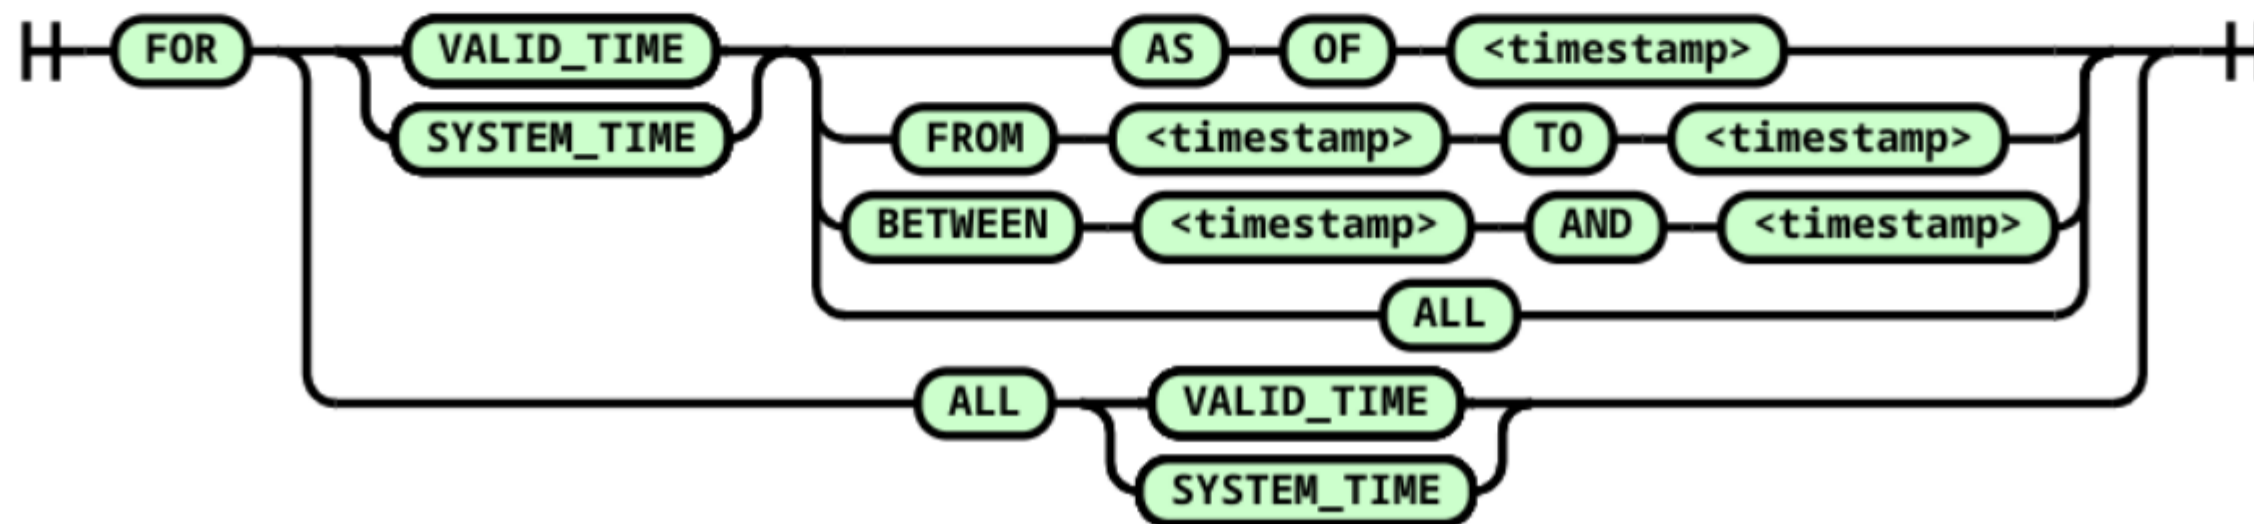
...XTDB helps when time is complex, and correctness matters.



SQL:2011 adds **bitemporal versioning** to tables:

- SYSTEM\_TIME (auditing & debugging)
- VALID\_TIME (reporting & versioning/corrections)
- Doesn't forget the past (e.g. DELETEs are 'soft')
- Various syntax to capture and query history...

**<temporal filter>**



# It's taken how long just to add two timestamps??

Scheme Evolution and the Relational Algebra  
May 1988

A database *scheme* describes the *structure* of the database; the *contents* of the database must adhere to that structure [Date 1976, Ullman 1982]. *Scheme evolution* refers to changes to the scheme of a database over time. Conventional databases allow only one scheme to be in force at a time, requiring *restructuring* (also termed *logical reorganization* [Scockut & Goldberg 1979]) when the scheme is modified. With the advent of databases storing past states [McKenzie 1986], it becomes desirable to accommodate multiple schemes, each in effect for an interval of time in the past.




Figure 3: Historical Relation




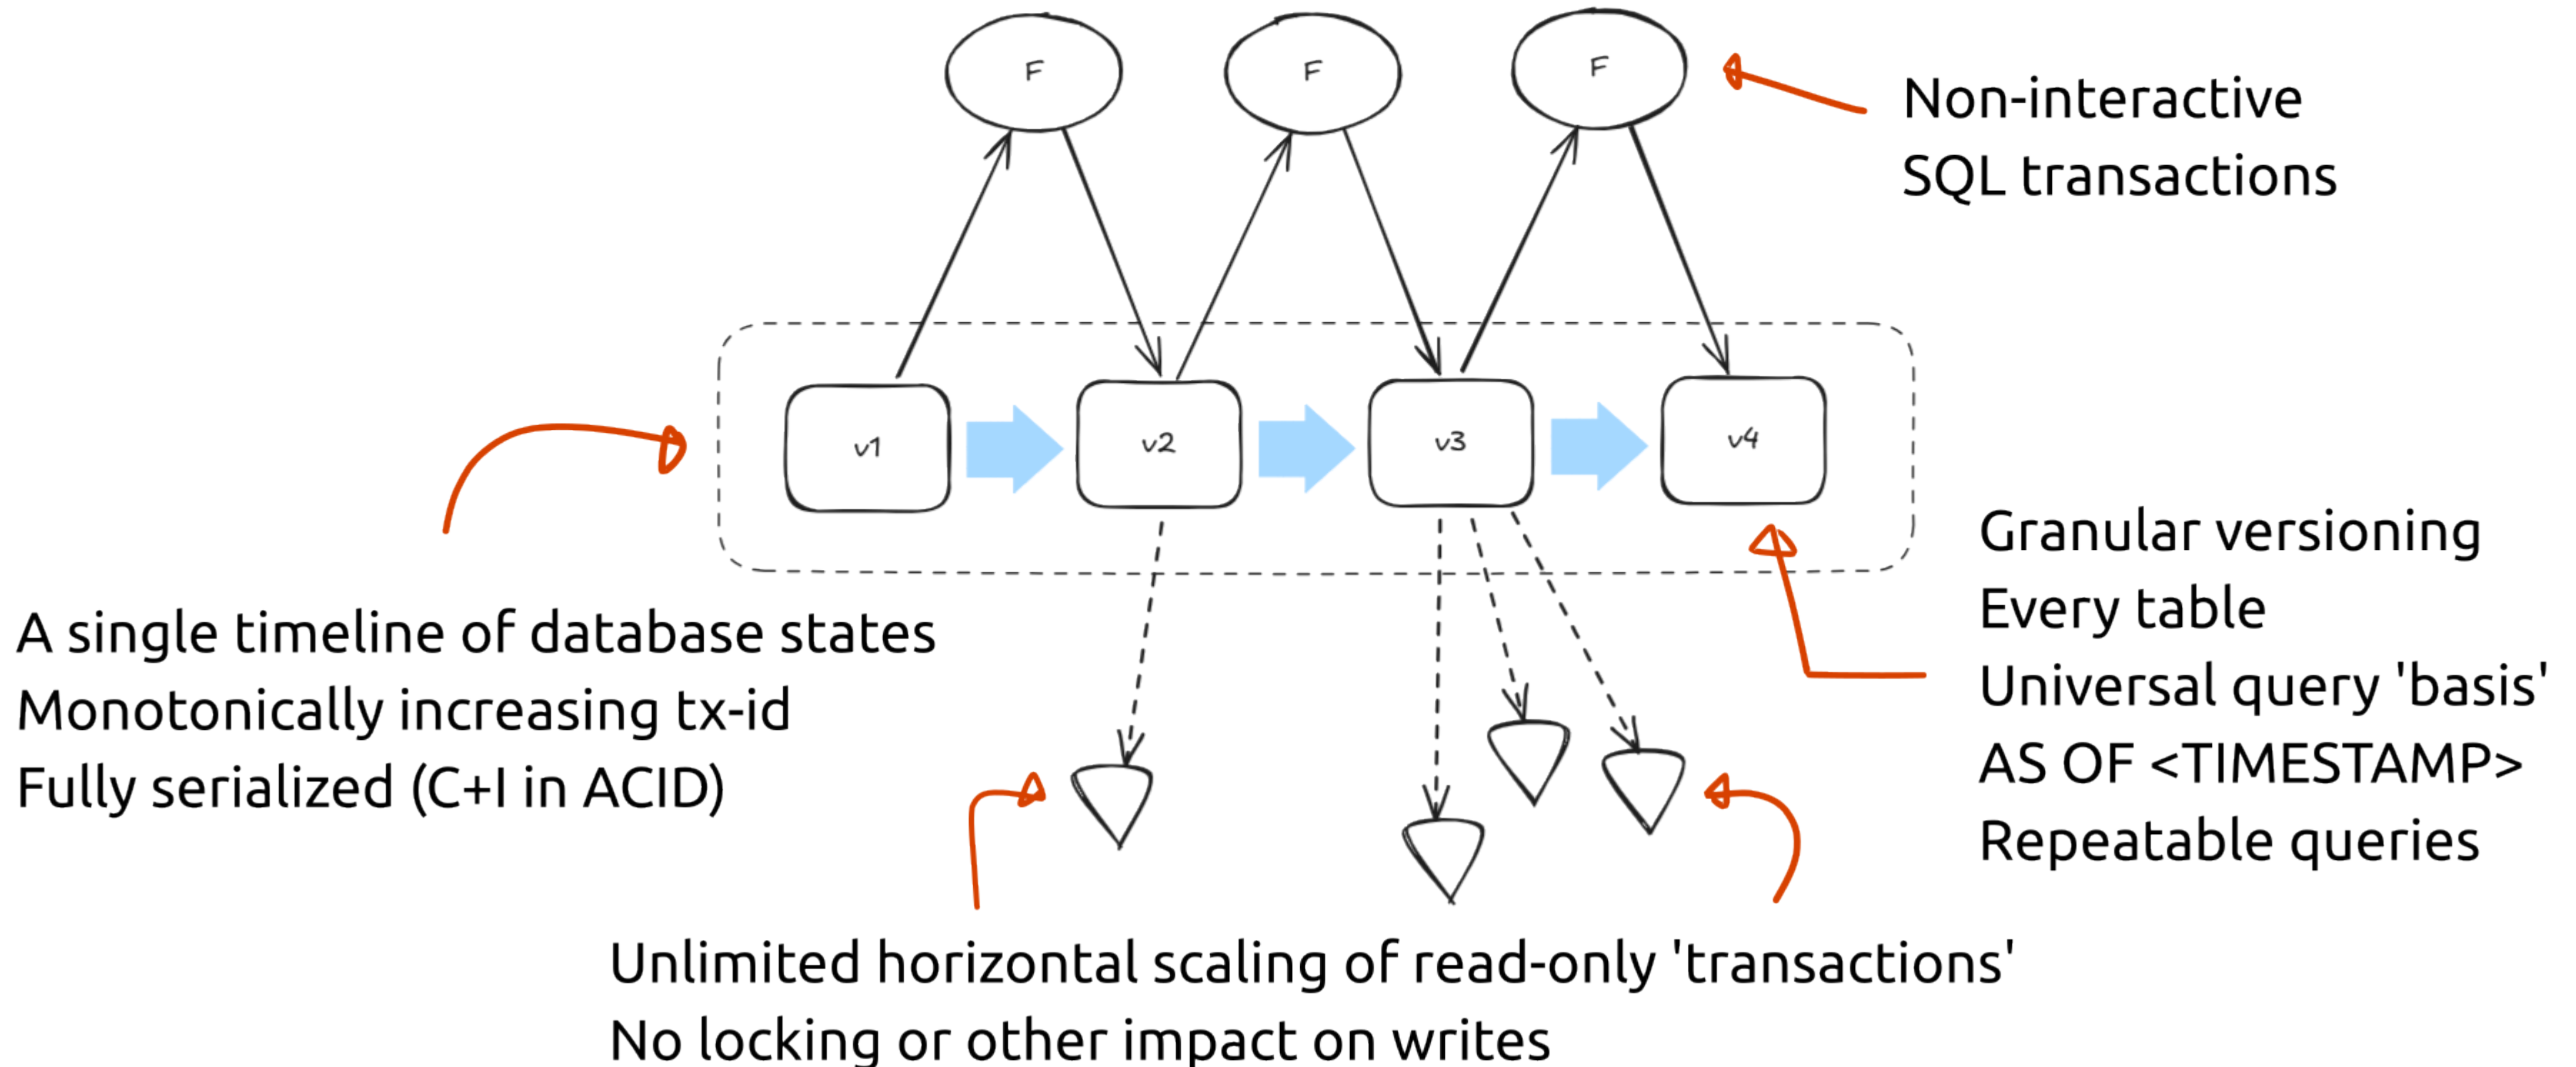
Figure 4: Temporal Relation



# A principled model of change

_id	_system_from	_system_to	_valid_from	_valid_to	value
1	"2024-09-15T20:23:25.562099Z"	null	"2024-09-15T20:23:25.562099Z"	"2029-01-01T00:00:00Z"	{"first class nested objects and":["arrays with rich types","2030-01-01"]}}
1	"2024-09-15T20:23:25.521119Z"	null	"2029-01-01T00:00:00Z"	null	"Future versioning"
1	"2024-09-15T20:23:25.473285Z"	null	"2024-01-02T00:00:00Z"	"2024-01-04T00:00:00Z"	"Historical corrections that preserve full auditability"
1	"2024-09-15T20:23:25.46007Z"	null	"2024-01-01T00:00:00Z"	"2024-01-02T00:00:00Z"	"Backfill of data"
1	"2024-09-15T20:23:25.46007Z"	"2024-09-15T20:23:25.473285Z"	"2024-01-02T00:00:00Z"	"2024-01-04T00:00:00Z"	"Backfill of data"
1	"2024-09-15T20:23:25.46007Z"	null	"2024-01-04T00:00:00Z"	"2024-09-15T20:23:25.562099Z"	"Backfill of data"
1	"2024-09-15T20:23:25.46007Z"	"2024-09-15T20:23:25.562099Z"	"2024-09-15T20:23:25.562099Z"	"2029-01-01T00:00:00Z"	"Backfill of data"
1	"2024-09-15T20:23:25.46007Z"	"2024-09-15T20:23:25.521119Z"	"2029-01-01T00:00:00Z"	null	"Backfill of data"
1	"2024-09-15T20:23:25.382963Z"	"2024-09-15T20:23:25.46007Z"	"2024-09-15T20:23:25.382963Z"	null	"Non-destructive UPDATES and DELETES, data is not forgotten"
1	"2024-09-15T20:23:25.358326Z"	"2024-09-15T20:23:25.46007Z"	"2024-09-15T20:23:25.358326Z"	"2024-09-15T20:23:25.382963Z"	"Authoritative records"
1	"2024-09-15T20:23:25.358326Z"	"2024-09-15T20:23:25.382963Z"	"2024-09-15T20:23:25.382963Z"	null	"Authoritative records"

In XTDB, SQL:2011 `SYSTEM_TIME` provides database-wide immutability ("The Epochal Time Model"):





Everything is implicitly filtered "as-of now", by default.

```
SELECT *,
       _valid_from, _valid_to,
       _system_from, _system_to
FROM docs
```

_id	_system_from	_valid_from	_valid_to	value
1	"2024-09-15T20:25:13.515361Z"	"2024-09-15T20:25:13.515361Z"	"2029-01-01T00:00:00Z"	{"first class nested objects and":["arrays with rich types","2030-01-01"]}

# Specify a filter to override the as-of now filter:


_id	_system_from	_system_to	_valid_from	_valid_to	value
1	"2024-09-15T20:23:25.562099Z"	null	"2024-09-15T20:23:25.562099Z"	"2029-01-01T00:00:00Z"	{"first class nested objects and":["arrays with rich types","2030-01-01"]}
1	"2024-09-15T20:23:25.521119Z"	null	"2029-01-01T00:00:00Z"	null	"Future versioning"
1	"2024-09-15T20:23:25.473285Z"	null	"2024-01-02T00:00:00Z"	"2024-01-04T00:00:00Z"	"Historical corrections that preserve full auditability"
1	"2024-09-15T20:23:25.46007Z"	null	"2024-01-01T00:00:00Z"	"2024-01-02T00:00:00Z"	"Backfill of data"
1	"2024-09-15T20:23:25.46007Z"				
1	"2024-09-15T20:23:25.46007Z"				
1	"2024-09-15T20:23:25.46007Z"				
1	"2024-09-15T20:23:25.46007Z"	"2024-09-15T20:23:25.521119Z"	"2029-01-01T00:00:00Z"	null	"Backfill of data"
1	"2024-09-15T20:23:25.382963Z"	"2024-09-15T20:23:25.46007Z"	"2024-09-15T20:23:25.382963Z"	null	"Non-destructive UPDATES and DELETES, data is not forgotten"
1	"2024-09-15T20:23:25.358326Z"	"2024-09-15T20:23:25.46007Z"	"2024-09-15T20:23:25.358326Z"	"2024-09-15T20:23:25.382963Z"	"Authoritative records"
1	"2024-09-15T20:23:25.358326Z"	"2024-09-15T20:23:25.382963Z"	"2024-09-15T20:23:25.382963Z"	null	"Authoritative records"

```
SELECT *,
  _valid_from, _valid_to,
  _system_from, _system_to
FROM docs FOR ALL VALID_TIME FOR ALL SYSTEM_TIME
```



The mutable VALID\_TIME dimension enables:

- backfilling history from upstream sources
- corrections
- scheduled effectivity
- time-travel for business users (~debugging)



```
INSERT INTO people (_id, name, favourite_color, _valid_from) VALUES  
(2, 'carol', 'blue', DATE '2023-01-01'),  
(2, 'carol', 'purple', DATE '2023-01-03')
```

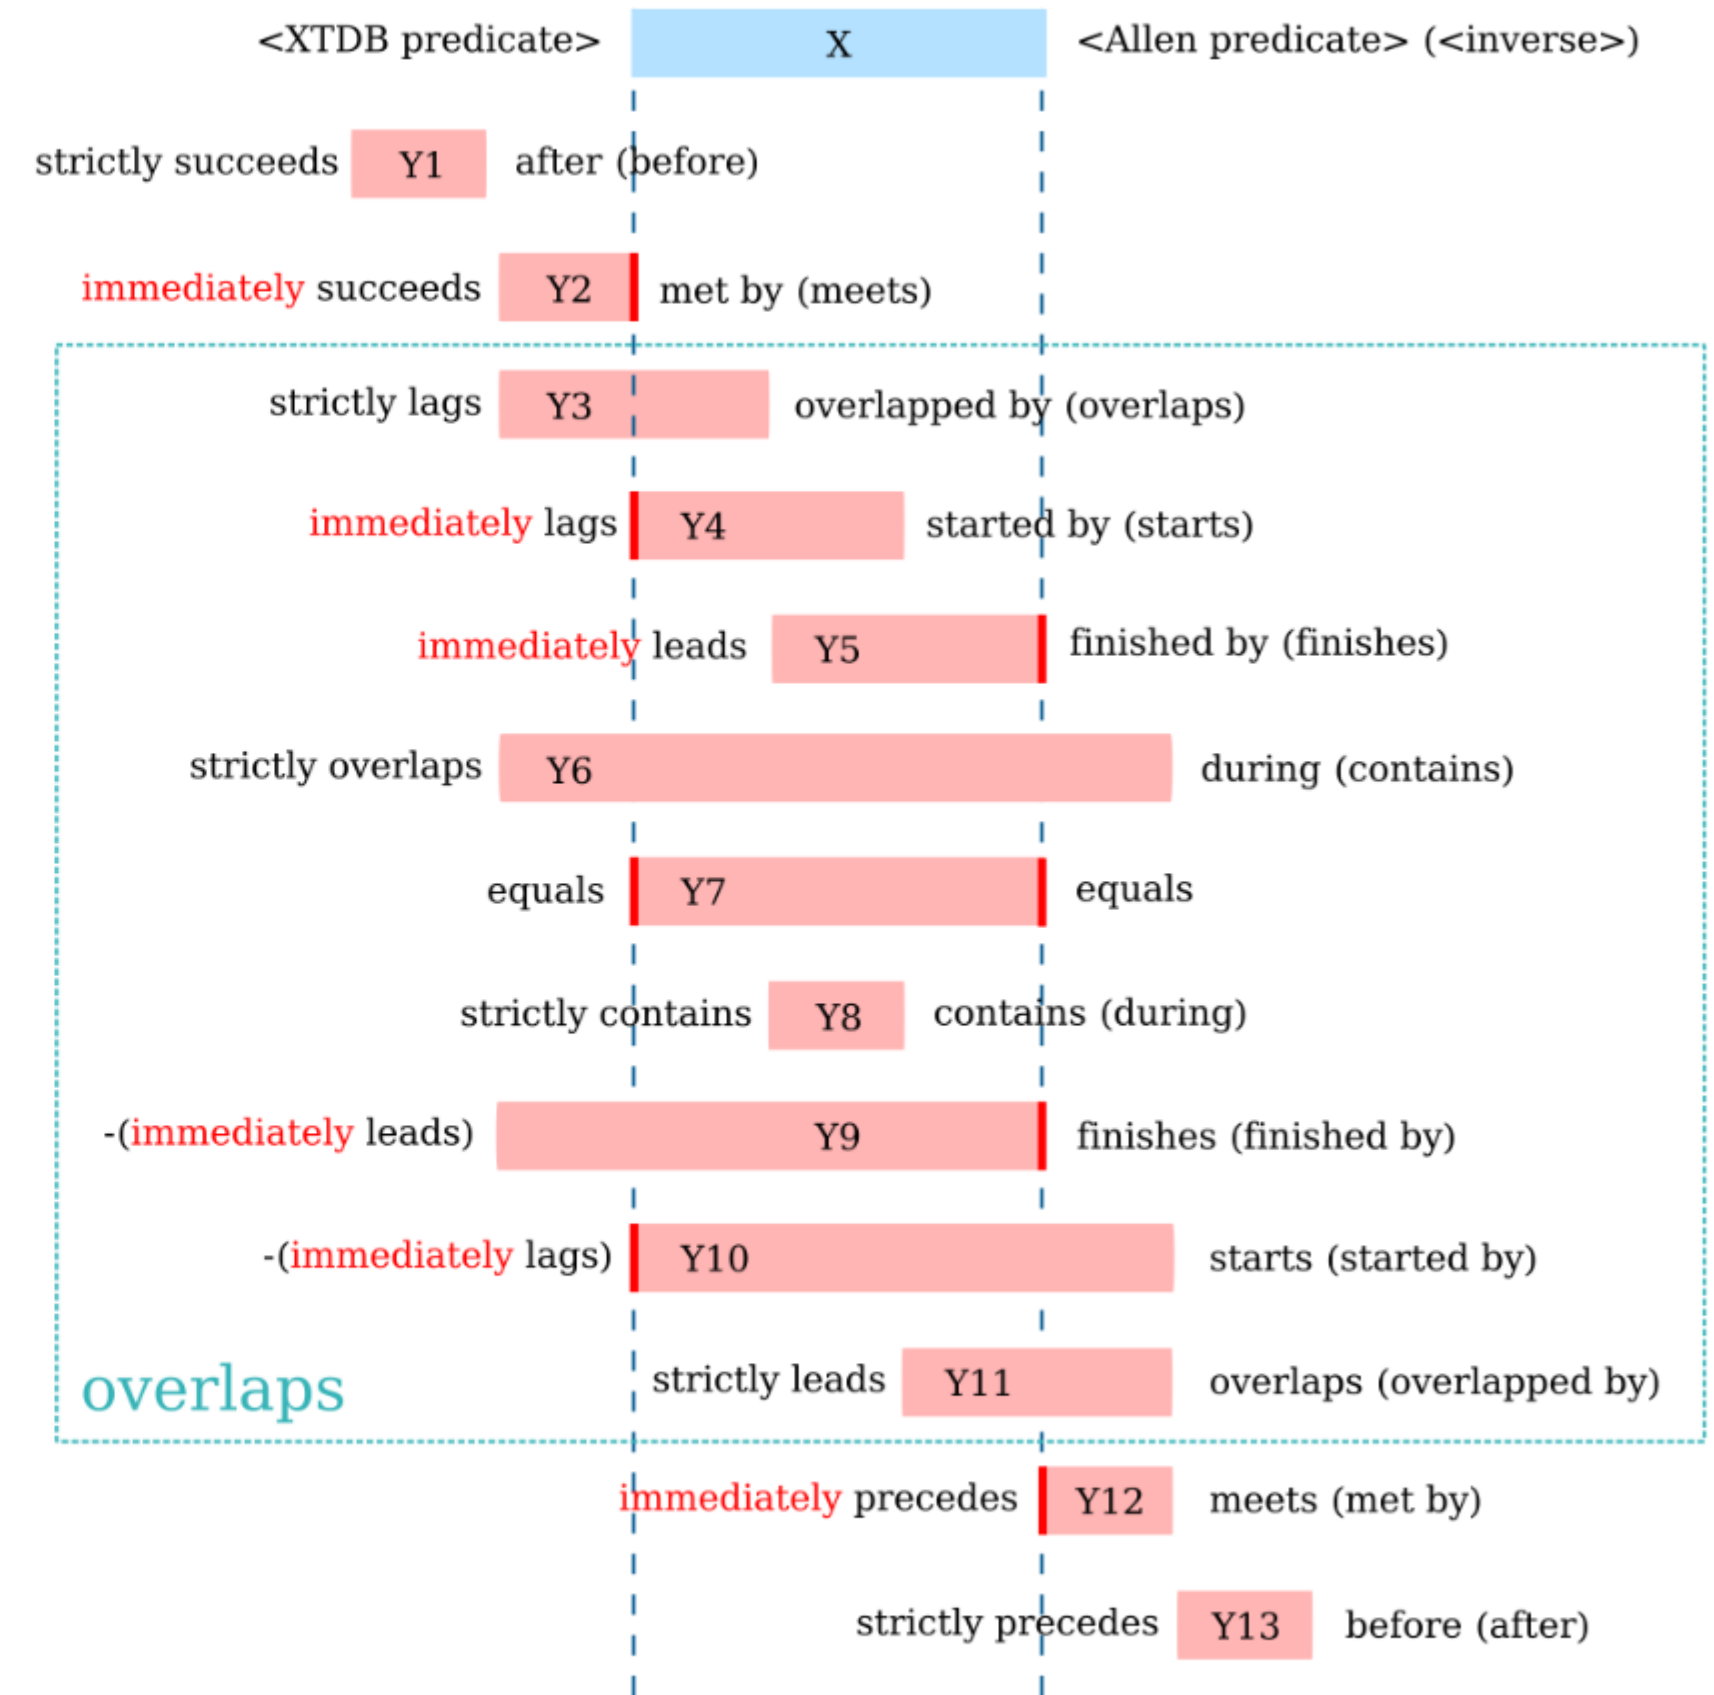
```
SELECT name, _valid_time, favourite_color FROM people FOR ALL VALID_TIME
```

_valid_time	favourite_color	name
{"from":"2023-01-03T00:00:00Z","to":null}	"purple"	"carol"
{"from":"2023-01-01T00:00:00Z","to":"2023-01-03T00:00:00Z"}	"blue"	"carol"
{"from":"2023-01-02T00:00:00Z","to":null}	"green"	"bob"

"Change Data Integration"

- The predicate “X OVERLAPS Y” in SQL:2011 is equivalent to the Boolean expression using Allen’s operators “(X overlaps Y) OR (X overlapped\_by Y) OR (X during Y) OR (X contains Y) OR (X starts Y) OR (X started\_by Y) OR (X finishes Y) OR (X finished\_by Y) OR (X equal Y)”. Note that Allen’s overlaps operator is not a true test of period overlap. Intuitively, two periods are considered overlapping if they have at least one time point in common.

"Temporal features in SQL:2011"  
(Kulkarni, Michels)





Reconciliation

# Track Complex, Evolving States & Positions

## Dividend

Forecast → Declaration → Ex-Date → Payment Schedule ....

## Fee

Estimate → Fixed → Invoiced → Adjusted → Payment

## Trade

Placed → Executed → Confirmed → Settled

Reconstruction & Reconciliation

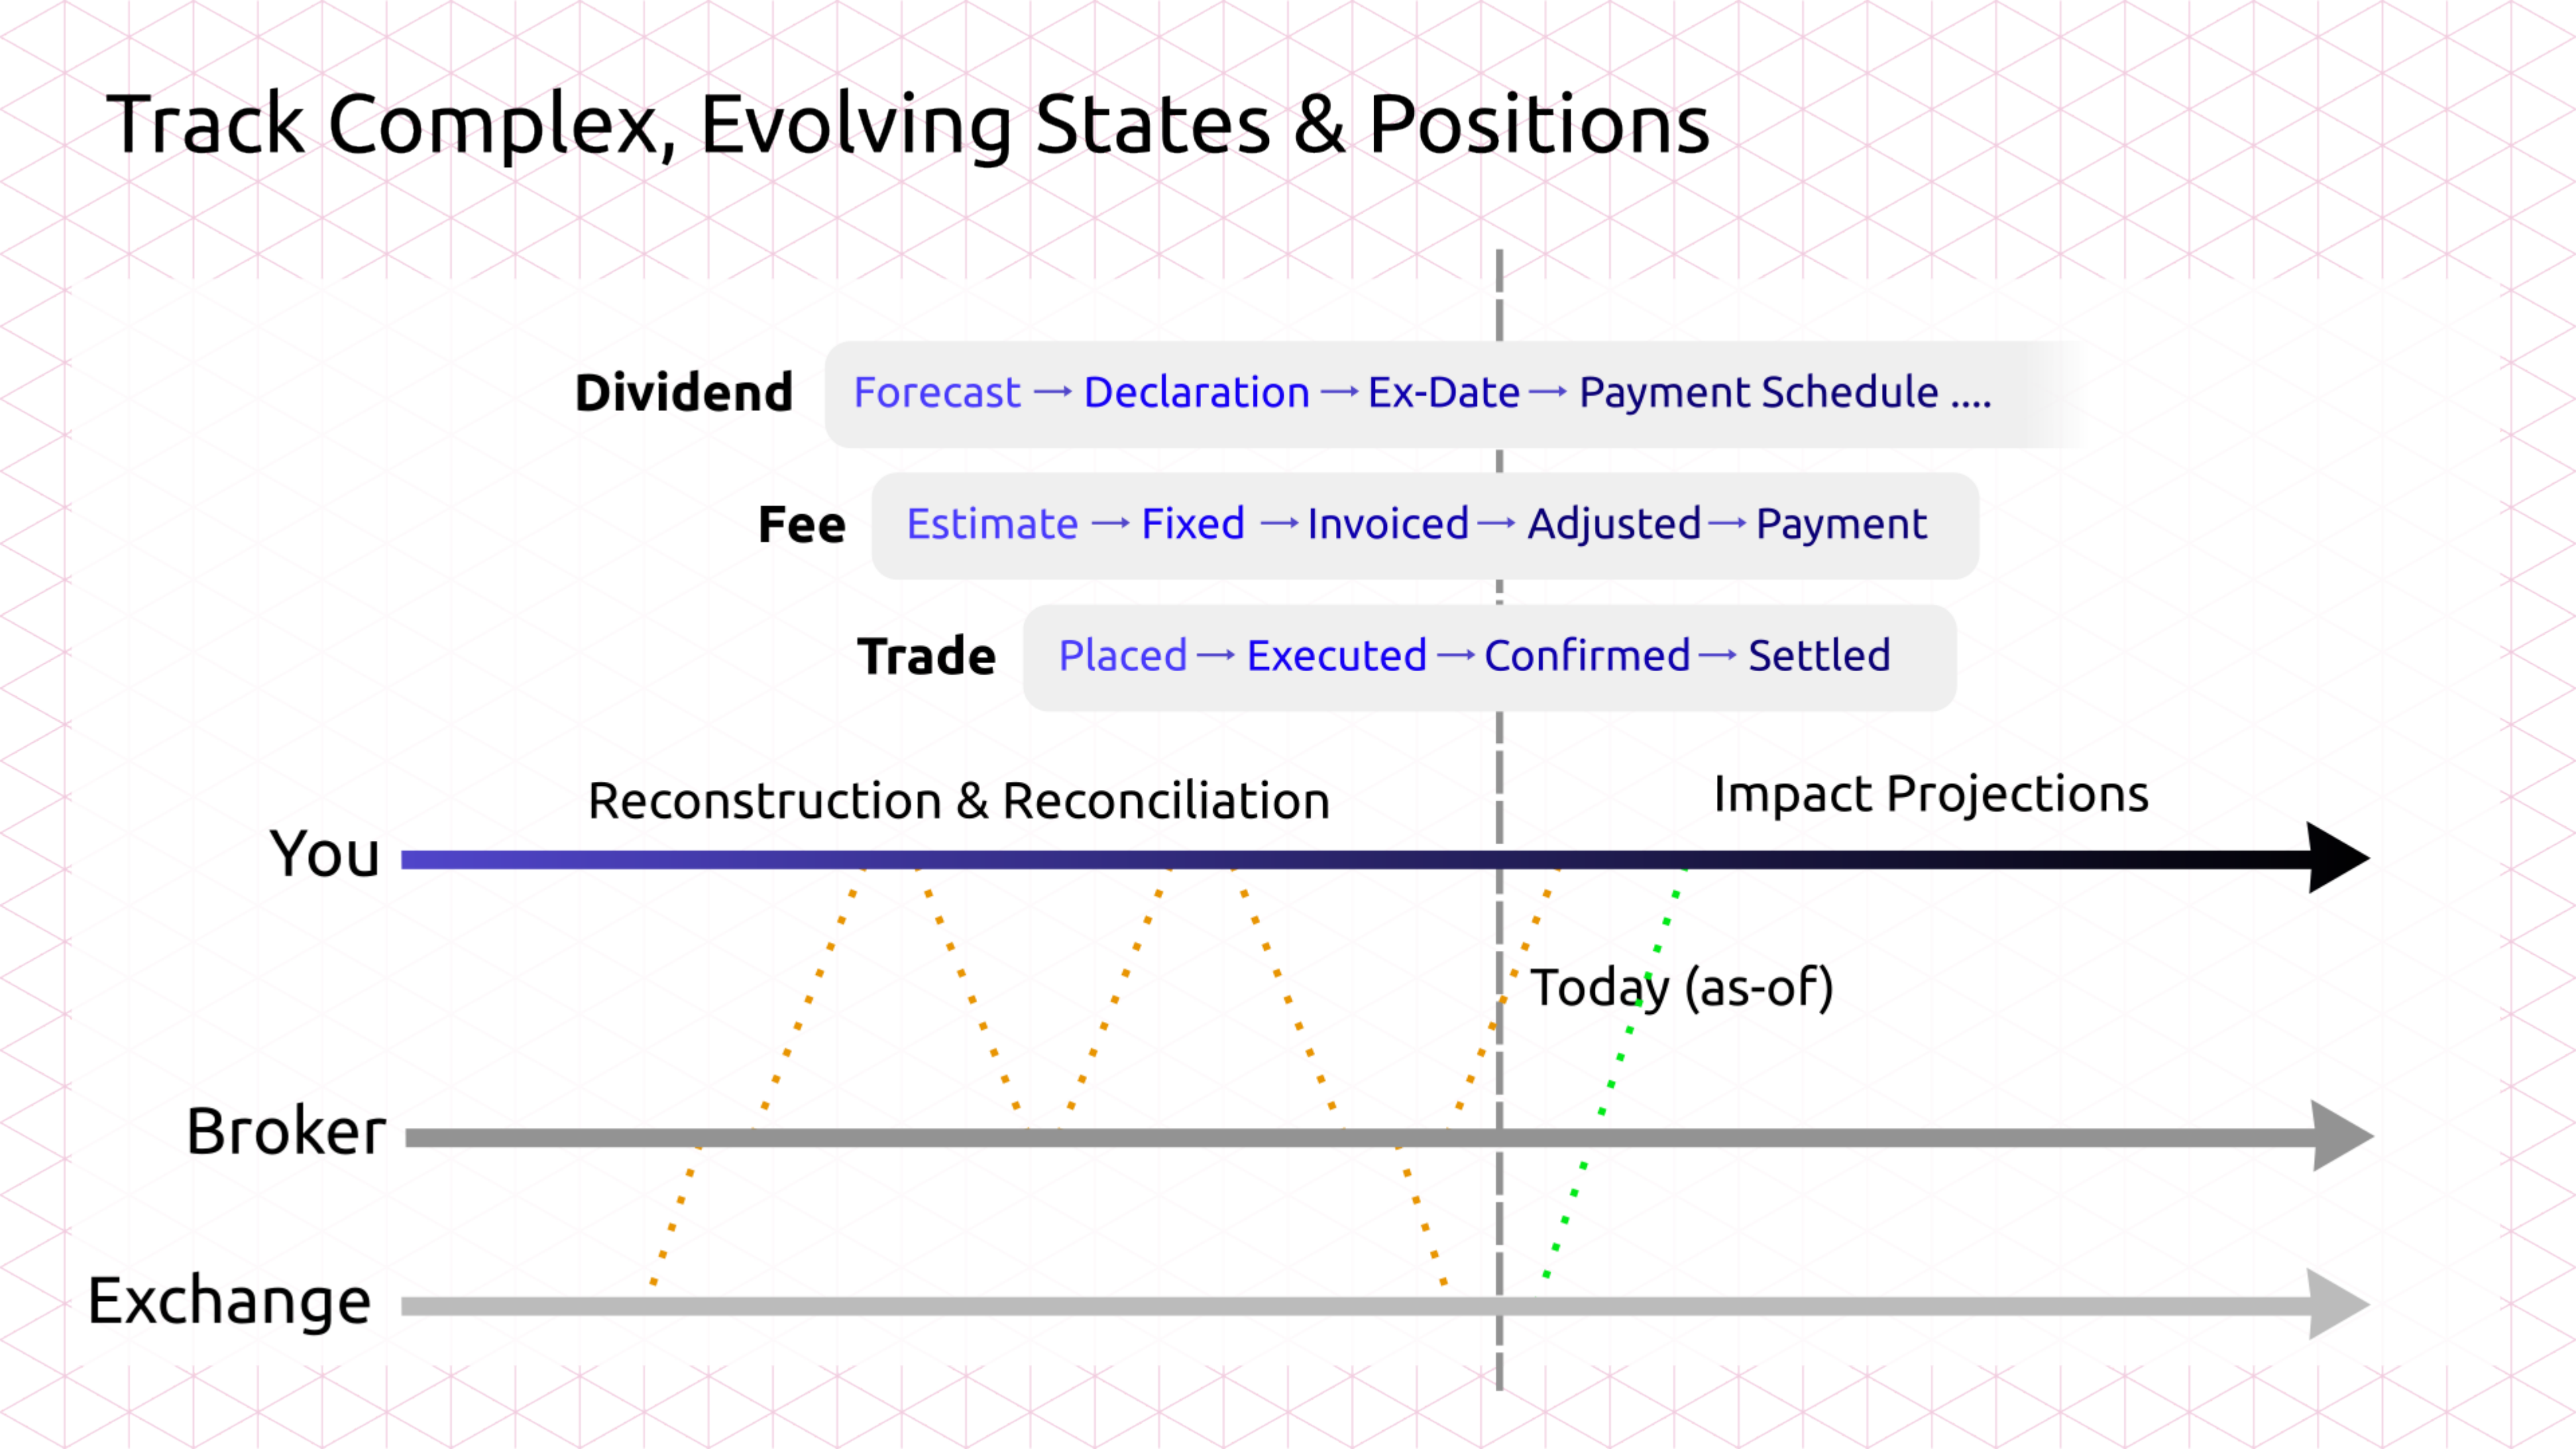
Impact Projections

You

Broker

Exchange

Today (as-of)





The background of the entire image is a repeating pattern of thin, light pink lines forming a grid of diamonds or rhombuses.

# Risk

# Time in Finance

IT systems of all kinds are plagued by many time-related complexities, from misconfigured Network Time Protocol and unexpected certificate expirations, to changing timezones and software concurrency bugs. However, systems in financial domains are additionally accutely affected by the following common requirements and challenges...

[Audit & Replay](#)

[Snapshot-free](#)

[Upstream Sources](#)

[Complex Histories](#)

[Intelligent Archival](#)

[Corrections](#)

[Schedule Change](#)

[What If](#)

## 1) Audit and Replay: what happened?

SQL databases systems typically lose information whenever UPDATE or DELETE are used. They also don't record or track the any changes to data by default.

```
INSERT INTO trades (_id, price) VALUES (1, 100);

UPDATE trades SET price = 150 WHERE _id = 1;

SELECT _id, price
FROM trades
```

Run ▶

[Open in xt-play](#)

In common SQL databases, the original price '100' is lost forever, and no information about when it was last changed is recorded. Perhaps if you're sufficiently desparate and very lucky there *might* be an old backup of the data available somewhere with a previous value.

### On this page

- Overview
- 1) Audit and Replay: what happened?

2) Snapshot-free Reporting: reproducibility without copying and result tearing

3) Interleaving Upstream Sources: whose timestamp?

4) Complex Version Histories

5) Intelligent Archival: live data and storage tiering with transparent queries

6) Corrections: curated history

7) Scheduled Effectivity: synchronize and preview future states

Automatic record expiration

Coordinated future versions

8) What If: time-travelling sources of truth



# Machine Learning

# Many people need *\*accurate\** time-travel.



## The Hopsworks Feature Store for Machine Learning

Javier de la Rúa Martínez  
Hopsworks AB, Stockholm, Sweden  
KTH Royal Institute of Technology  
Stockholm, Sweden  
jdlrm@kth.se

Alexandru A. Ormenisan  
Hopsworks AB, Stockholm, Sweden  
alex@hopsworks.ai


Kenneth Mak  
Hopsworks AB, Stockholm, Sweden  
kenneth@hopsworks.ai

Raymond Cunningham  
Hopsworks AB, Stockholm, Sweden  
ray.cunningham@hopsworks.ai

Ayushman Khazanchi  
KTH Royal Institute of Technology  
Stockholm, Sweden  
ayushman@kth.se

### ABSTRACT

Data management is the most challenging aspect of machine Learning (ML) systems. ML systems can require large amounts of historical data when training models, but in practice, data is more varied, depending on whether it is a batch or streaming system. The feature store for ML has recently emerged as a data platform for managing ML data throughout the ML lifecycle.



DuckDB



Documentation ▾

Blog

GitHub ★ 22.7k

Documentation / Guides / SQL Features

Dark Mode

1.1 (stable)

## AsOf Join

### What is an AsOf Join?

Time series data is not always perfectly aligned. Clocks may be slightly off, or there may be a delay between cause and effect. This can make connecting two sets of ordered data challenging. AsOf joins are a tool for solving this and other similar problems.

One of the problems that AsOf joins are used to solve is finding the value of a varying property at a specific point in time. This use case is so common that it is where the name came from:

*Give me the value of the property **as of this time**.*

More generally, however, AsOf joins embody some common temporal analytic semantics, which can be cumbersome and slow to implement in standard SQL.

### Portfolio Example Data Set

Let's start with a concrete example. Suppose we have a table of stock prices with timestamps:

ticker	when	price
APPL	2001-01-01 00:00:00	1
APPL	2001-01-01 00:01:00	2
APPL	2001-01-01 00:02:00	3
MSFT	2001-01-01 00:00:00	1
MSFT	2001-01-01 00:01:00	2

## Eventual Business Consistency

Executive Summary of Bi-temporality



KENT BECK  
AUG 04, 2023



92



21



6

Share

...

I'm a geek speaking to you, a technology-savvy executive, about why we are doing things in a more complicated way than seems necessary. You may have heard the word "bi-temporal". What's that about?

In a nutshell, we want what's recorded in the system to match the real world. We know this is impossible (delays, mistakes, changes) but are getting as close as we can. The promise is that if what's in the system matches the real world as closely as possible, costs go down, customer satisfaction goes up, & we are able to scale further faster.

Here's how it works.






~Low cost and ubiquitous

## PostgreSQL



The popular default



Battle hardened  
and highly extensible




Architected by  
Michael Stonebraker



Does not support bitemporal data

# Historical Cost of Computer Memory and Storage



Postgres was originally  
designed in 1986



Prof. Viktor Leis  
Technical University of Munich


What do you think about  
the future of Postgres?

"It's just a very old system that's extremely inefficient [...] there's many orders of magnitude for things to be more efficient usually [...] there's no single kind of performance bottleneck, it's just that everything over time has accumulated so much performance overhead [...]"

I don't think that it's really possible to save it."

JUXT JUXTCast S5E8 - Sane Query Languages — with Prof. Viktor Leis

Watch Later Share

 Viktor Leis  
Technische Universität München  
Verified email at in.tum.de - [Homepage](#)  
[Database Systems](#) FOLLOW

TITLE	CITED BY	YEAR
<a href="#">How good are query optimizers, really?</a> V Leis, A Gubichev, A Mirchev, P Boncz, A Kemper, T Neumann Proceedings of the VLDB Endowment 9 (3), 204-215	767	2015
<a href="#">The adaptive radix tree: ARTful indexing for main-memory databases</a> V Leis, A Kemper, T Neumann 2013 IEEE 29th International Conference on Data Engineering (ICDE), 38-49	562	2013
<a href="#">Learned cardinalities: Estimating correlated joins with deep learning</a> A Kipf, T Kipf, B Radke, V Leis, P Boncz, A Kemper CIDR	450	2019
<a href="#">Morsel-driven parallelism: a NUMA-aware query evaluation framework for the many-core age</a> V Leis, P Boncz, A Kemper, T Neumann Proceedings of the 2014 ACM SIGMOD international conference on Management of ...	373	2014

Watch on YouTube

```
$ docker run -p 5432:5432 ghcr.io/xtodb/xtodb:2.0.0b1
```

```
$ psql -h localhost
```

```
psql (16.3, server 16)
```

```
Type "help" for help.
```

```
jdt=> select xt.version();
```

```
version
```

```
-----
```

```
XTDB @ 2.0.0b1
```

```
(1 row)
```

```
jdt=>
```



# How XTDB Accelerates As-Of Reporting

*Any language*  
**Postgres Driver  
Compatibility**

*Standard SQL*  
**Transactional  
Applications**

*All History*  
**Hybrid Analytical  
Workloads**

*Any Cloud*  
**Architected for  
Object Storage**

## Common Tooling



## BI & Data Platforms



## Authentication



## ACID + Serializable Isolation

**Row-Level  
Versioning**

**Primary Key As-Of  
Temporal Indexes**

**Columnar  
Execution**

Arrow-Based  
JVM Engine



Semi-Structured  
Data Schemas

Kubernetes  
Deployment



**Join us again in January  
to hear from our users and  
discover our plans for 2025.**

## XTDB in 2025 – Mission, Users, Technology

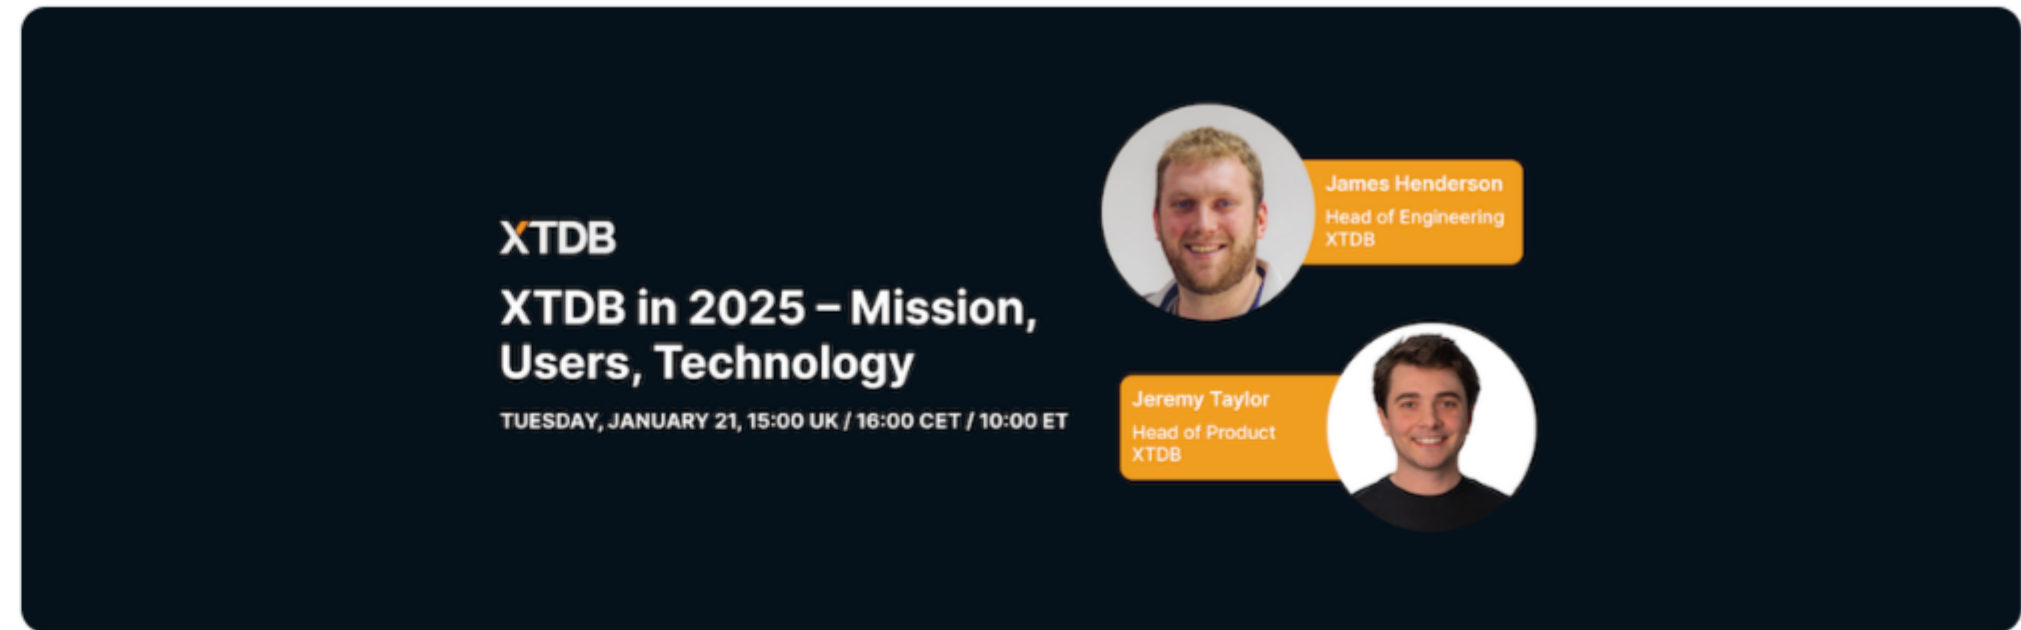
Date & Time

Jan 21, 2025 03:00 PM in [London](#)

Description

Join us for a live session with Jeremy Taylor (Head of Product) and James Henderson (Head of Engineering), where we will showcase XTDB v2, and discuss the recent progress the team has made in conjunction with our Design Partners.

Featuring: new demos, use cases and stories from our design partners, and the latest roadmap announcements!



[https://us06web.zoom.us/webinar/register/4817339038090/WN\\_QqygMHteRvmPo0O4aQDW5Q](https://us06web.zoom.us/webinar/register/4817339038090/WN_QqygMHteRvmPo0O4aQDW5Q)



Thanks for listening! Any questions?

Try out XTDB v2 live in your browser right now:  
**<https://docs.xtdb.com>**

Do you have (bi)temporal problems?  
We're looking for Design Partners - let's chat:  
**>> [hello@xtdb.com](mailto:hello@xtdb.com) <<**



Jeremy Taylor | JUXT | @refset

