

Raising the Level of Abstraction for Time State Analytics With the Timeline Framework


Henry Milner, Yihua Cheng, Jibin Zhan,
Hui Zhang, Vyas Sekar, Junchen Jiang, Ion Stoica

Conviva, Carnegie Mellon, UChicago, UC Berkeley

10 January 2023

Example Analytics Query Intents

Video



How much time did this session spend buffering while using CDN C1?

➔ Change CDN

Common feature:
Stateful context-sensitive metrics computed over continuous time


Cybersecurity



visiting website xyz.com in the last hour?

➔ BLOCK URL
Quarantine host

Manufacturing

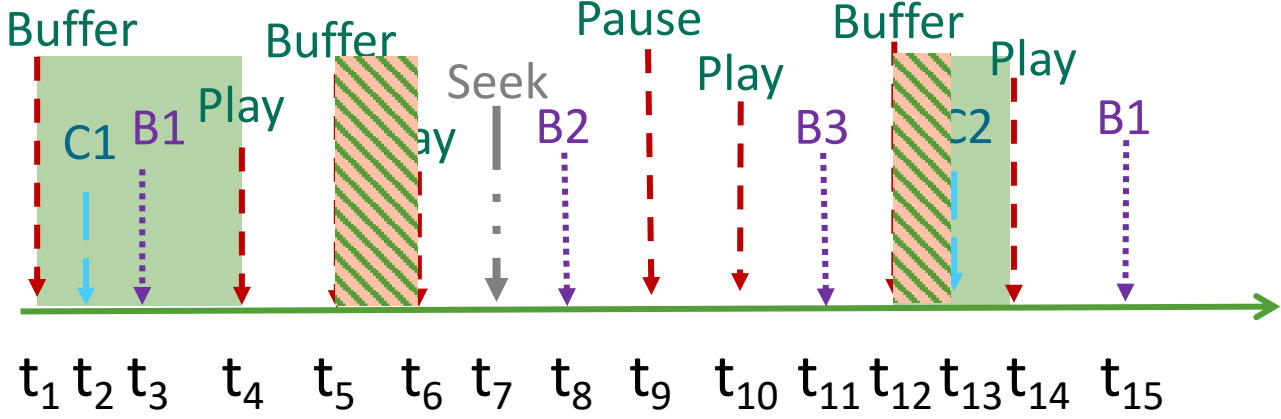


How many machines from vendor X are showing degrading health status over time?

➔ Rebalance load,
Repair

...

Streaming Video QoE: Connection-Induced Rebuffering



Stateful

Context-sensitive

Continuous time



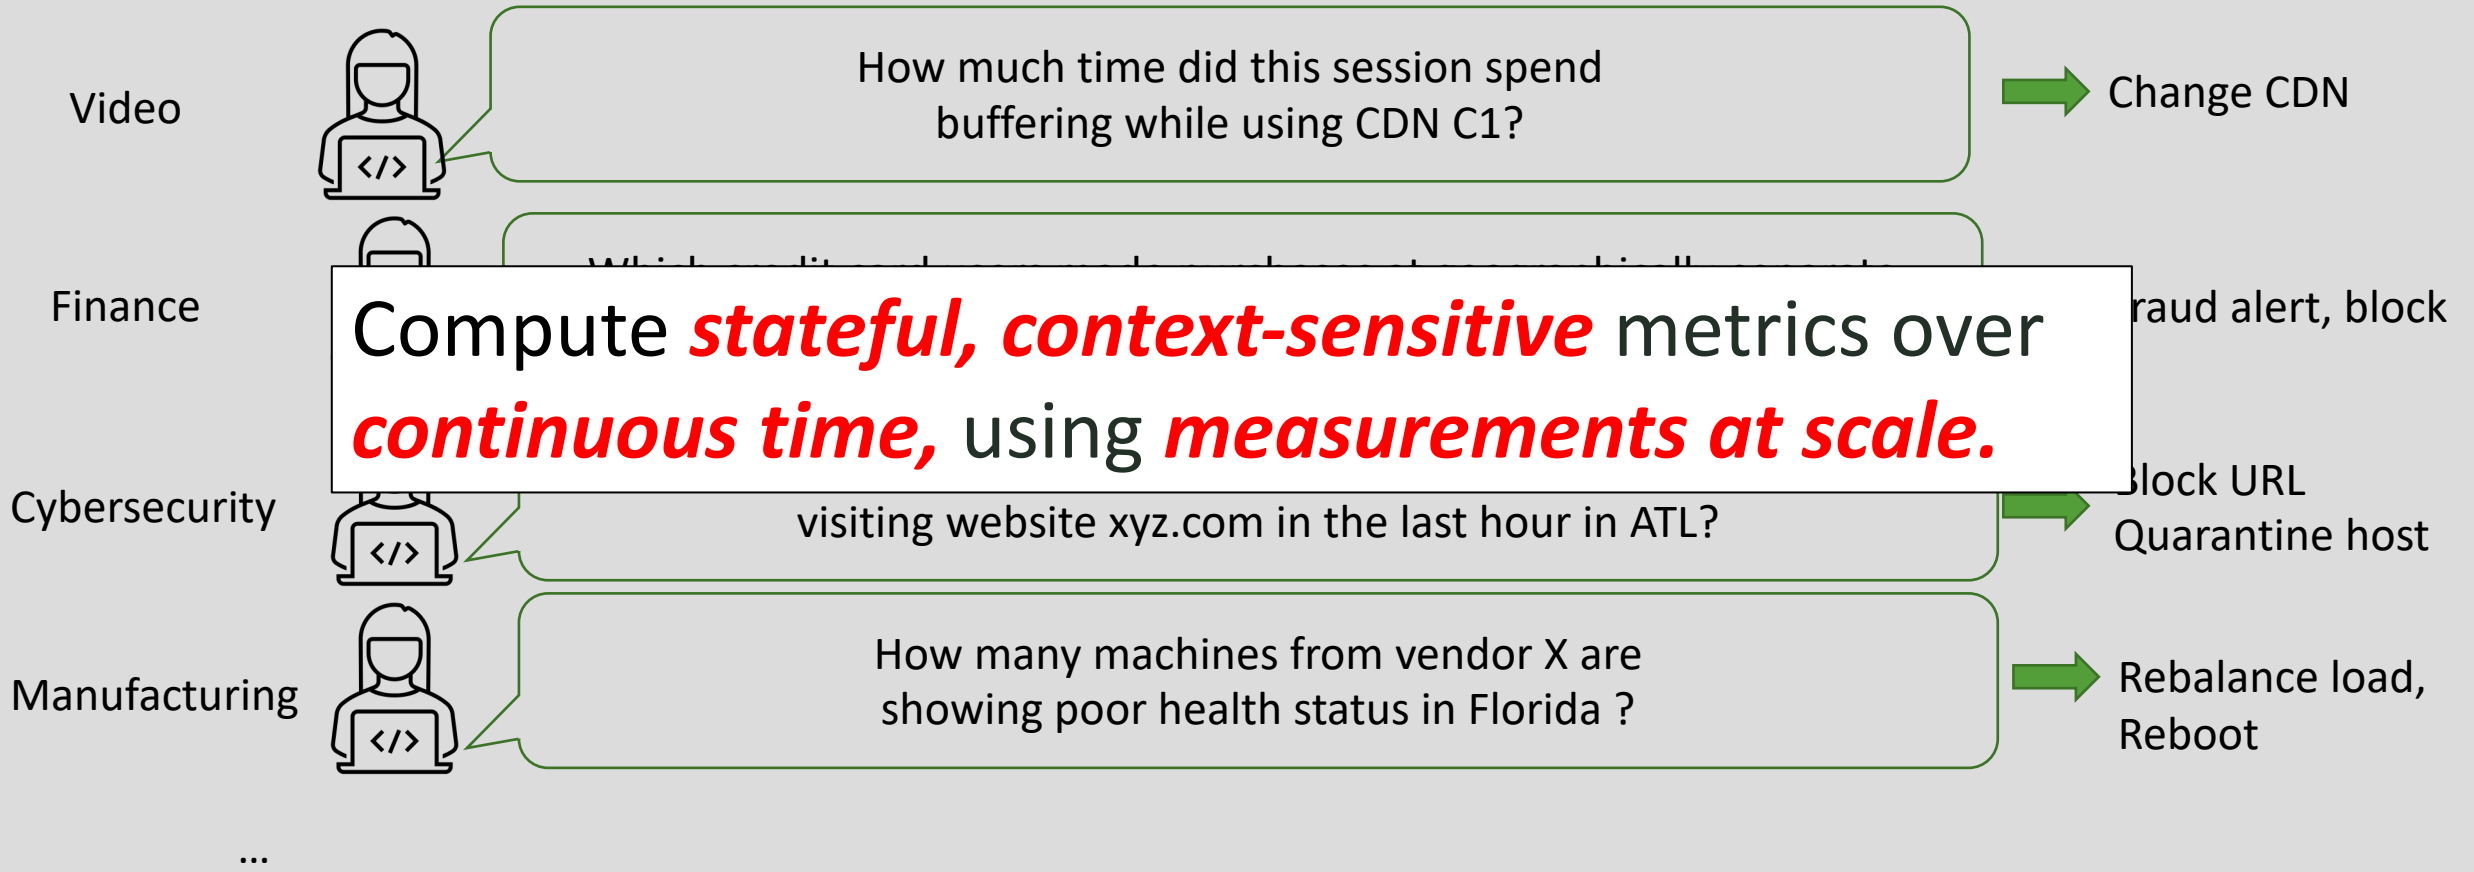
How much time did this session spend in a connection-induced rebuffering state while using CDN C1?



Count the duration where:

1. Currently buffering &
2. Play has already initialized &
3. Hasn't seeked in last 5 seconds &
4. Using CDN C1

Time-State Analytics (TSA), in a nutshell



TSA isn't served well by existing data processing systems

```
1 WITH SeekAsPlayerState(T, P) as (  
2     SELECT T, P FROM heartbeats WHERE P IS NOT NULL  
3     UNION SELECT T, "Seek_st" FROM heartbeats WHERE A IS NOT NULL  
4     UNION SELECT T + 5, "Seek_ed" FROM heartbeats WHERE A IS NOT NULL ),  
5 IgnoreBufBeforePlay(T, P) as (  
6     SELECT T, P FROM (  
7         SELECT T, P, Max(If(P == 'play', 1, 0)) OVER (PARTITION BY 1 ORDER BY T)  
8         ↪ as H  
9         FROM SeekAsPlayerState) WHERE H == True ),  
10 DuringBufferTable(T, P, DB) as (  
11     SELECT T, P, LAST(tmp1) IGNORE NULLS OVER (PARTITION BY 1 ORDER BY T)  
12     FROM (  
13         SELECT T, P,  
14         CASE P WHEN 'buffer' THEN True WHEN 'Seek_st' THEN NULL WHEN 'Seek_ed'  
15         ↪ THEN NULL ELSE FALSE END as tmp1  
16         From IgnoreBufBeforePlay ) ),  
17 DuringSeekTable(T, P, DB, DS) as (  
18     SELECT T, P, DB,  
19     (T - Max(If(P == 'Seek_st', T, 0)) OVER (PARTITION BY 1 ORDER BY T)  
20     ↪ ) < 5 as tmp2  
21     FROM DuringBufferTable ),  
22 IgnoreBufInSeek(T, P) as (  
23     SELECT T, P FROM (  
24         SELECT T, DS, IF(P == 'Seek_ed' and DB, 'buffer', P) as P  
25         FROM DuringSeekTable ) WHERE NOT (P == 'buffer' AND DS) ),  
26 WithCDNAndQuery(T, P, C) as (  
27     SELECT T, P, NULL FROM IgnoreBufInSeek  
28     UNION SELECT T, NULL, C FROM heartbeats where C IS NOT NULL  
29     UNION SELECT 2022-07-21 10:05, NULL, NULL s),  
30 Intervals(Ed, St, State, CDN) as (  
31     SELECT T, LEAD(T, 1) OVER (PARTITION BY 1 ORDER BY T), P, C  
32     FROM (  
33         SELECT T,  
34         LAST(P) IGNORE NULLS OVER (PARTITION BY 1 ORDER BY T) as P,  
35         LAST(C) IGNORE NULLS OVER (PARTITION BY 1 ORDER BY T) as C  
36         FROM WithCDNAndQuery ) )  
37 SELECT SUM(St - Ed) as result FROM Intervals  
38 WHERE Ed < 2022-07-21 10:05 AND State == 'buffer' AND CDN == 'CDN1'
```



High dev effort

High cost



Count the duration where:

1. Currently buffering &
2. Play has already initialized &
3. Hasn't seeked in last 5 seconds &
4. Using CDN C1

Our work: Timeline abstraction for Time-State Analytics

Writing time-state queries becomes intuitive visual operations

→ Reduced dev effort

Enables new opportunities for structure-aware optimizations

→ Up to 10x improvement in cost

Scope of this work: Focus on the single user-session intent modeling problem

Outside our scope: Supporting scale-out and aggregation

Outline for talk

- What is Time-State Analytics
- *Time-State Analytics not well supported by status quo*
- Introducing the Timeline abstraction
- Early Wins + Next Steps

Where's the problem?

```

1 WITH SeekAsPlayerState(T, P) as (
2     SELECT T, P FROM heartbeats WHERE P IS NOT NULL
3     UNION SELECT T, "Seek_st" FROM heartbeats WHERE A IS NOT NULL
4     UNION SELECT T + 5, "Seek_ed" FROM heartbeats WHERE A IS NOT NULL ),
5 IgnoreBufBeforePlay(T, P) as (
6     SELECT T, P FROM (
7         SELECT T, P, Max(If(P == 'play', 1, 0)) OVER (PARTITION BY 1 ORDER BY T)
8             ↪ as H
9         FROM SeekAsPlayerState) WHERE H == True ),
10 DuringBufferTable(T, P, DB) as (
11     SELECT T, P, LAST(tmp1) IGNORE NULLS OVER (PARTITION BY 1 ORDER BY T)
12     FROM (
13         SELECT T, P,
14             CASE P WHEN 'buffer' THEN True WHEN 'Seek_st' THEN NULL WHEN 'Seek_ed'
15                 ↪ THEN NULL ELSE FALSE END as tmp1
16         From IgnoreBufBeforePlay ) ),
17 DuringSeekTable(T, P, DB, DS) as (
18     SELECT T, P, DB,
19         (T - Max(If(P == 'Seek_st', T, 0)) OVER (PARTITION BY 1 ORDER BY T)
20             ↪ ) < 5 as tmp2
21     FROM DuringBufferTable ),
22 IgnoreBufInSeek(T, P) as (
23     SELECT T, P FROM (
24         SELECT T, DS, IF(P == 'Seek_ed' and DB, 'buffer', P) as P
25         FROM DuringSeekTable ) WHERE NOT (P == 'buffer' AND DS) ),
26 WithCDNAndQuery(T, P, C) as (
27     SELECT T, P, NULL FROM IgnoreBufInSeek
28     UNION SELECT T, NULL, C FROM heartbeats where C IS NOT NULL
29     UNION SELECT 2022-07-21 10:05, NULL, NULL s),
30 Intervals(Ed, St, State, CDN) as (
31     SELECT T, LEAD(T, 1) OVER (PARTITION BY 1 ORDER BY T), P, C
32     FROM (
33         SELECT T,
34             LAST(P) IGNORE NULLS OVER (PARTITION BY 1 ORDER BY T) as P,
35             LAST(C) IGNORE NULLS OVER (PARTITION BY 1 ORDER BY T) as C
36         FROM WithCDNAndQuery ) )
37 SELECT SUM(St - Ed) as result FROM Intervals
38 WHERE Ed < 2022-07-21 10:05 AND State == 'buffer' AND CDN == 'CDN1'

```



High dev effort

High cost



- Count the duration where:
1. Currently buffering &
 2. Play has already initialized &
 3. Hasn't seeked in last 5 seconds &
 4. Using CDN C1

Tabular model isn't well-suited for Time-State

Timestamp	Player State	Bitrate	CDN	Seek
t1	Buffer			
t2	Play	B1	C1	
t3	Buffer			
t4	Play	B2	C1	
t5	Buffer			
t6	Play	B3	C2	
t7	Seek			Seek
t8	Paused			
t9	Play			
t10	Play	B1		
t11	Play	B3		
t12	Buffer			
t13	Play		C2	
t14	Play			
t15	Play	B1		

Raw measurements from a video session

Player State Bitrate CDN Seek

Count the duration where:

1. Currently buffering &
2. Play has already initialized &
3. Hasn't seeked in last 5 seconds &
4. Using CDN C1



State and Context over Continuous Time is Hard

Timestamp	Player State	Bitrate	CDN	Seek
t1	Buffer			
t2	Buffer		C1	
t3	Buffer			
t4	Play			
t5	Buffer			
t6	Play			
t7	Play			
t8	Play			
t9	Paused			
t10	Play			
t11	Play			
t12	Buffer			
t13	Buffer	B3	C2	
t14	Play	B3	C2	
t15	Play	B1	C2	

Count the duration where:

- Currently buffering &
- Play has already initialized &
- Hasn't seeked in last 5 seconds &
- Using CDN C1

duration?

```
[...]
WITH SeekAsPlayerState(T, P) as (
  SELECT T, P FROM heartbeats WHERE P IS NOT NULL
  UNION SELECT T, "Seek_st" FROM heartbeats WHERE A IS NOT NULL
  UNION SELECT T + 5, "Seek_ed" FROM heartbeats WHERE A IS NOT NULL ),
[...]
DuringBufferTable(T, P, DB) as (
  SELECT T, P, LAST(tmp1) IGNORE NULLS OVER (PARTITION BY 1 ORDER BY T)
  FROM (
    SELECT T, P,
      CASE P WHEN 'buffer' THEN True WHEN 'Seek_st' THEN NULL WHEN 'Seek_ed'
        THEN NULL ELSE FALSE END as tmp1
    From IgnoreBufBeforePlay ) ),
DuringSeekTable(T, P, DB, DS) as (
  SELECT T, P, DB,
    (T - Max(If(P == 'Seek_st', T, 0)) OVER (PARTITION BY 1 ORDER BY T)
      ) < 5 as tmp2
  FROM DuringBufferTable ),
IgnoreBufInSeek(T, P) as (
  SELECT T, P FROM (
    SELECT T, DS, IF(P == 'Seek_ed' and DB, 'buffer', P) as P
    FROM DuringSeekTable ) WHERE NOT (P == 'buffer' AND DS) ),
[...]
```



t7 + 5 seconds???



Poor abstraction → Complex code

```
1 WITH SeekAsPlayerState(T, P) as (  
2     SELECT T, P FROM heartbeats WHERE P IS NOT NULL  
3     UNION SELECT T, "Seek_st" FROM heartbeats WHERE A IS NOT NULL  
4     UNION SELECT T + 5, "Seek_ed" FROM heartbeats WHERE A IS NOT NULL ),  
5 IgnoreBufBeforePlay(T, P) as (  
6     SELECT T, P FROM (  
7         SELECT T, P, Max(If(P == 'play', 1, 0)) OVER (PARTITION BY 1 ORDER BY T)  
8         ↪ as H  
9         FROM SeekAsPlayerState) WHERE H == True ),  
10 DuringBufferTable(T, P, DB) as (  
11     SELECT T, P, LAST(tmp1) IGNORE NULLS OVER (PARTITION BY 1 ORDER BY T)  
12     FROM (  
13         SELECT T, P,  
14         CASE P WHEN 'buffer' THEN True WHEN 'Seek_st' THEN NULL WHEN 'Seek_ed'  
15         ↪ THEN NULL ELSE FALSE END as tmp1  
16         From IgnoreBufBeforePlay ) ),  
17 DuringSeekTable(T, P, DB, DS) as (  
18     SELECT T, P, DB,  
19     (T - Max(If(P == 'Seek_st', T, 0)) OVER (PARTITION BY 1 ORDER BY T)  
20     ↪ ) < 5 as tmp2  
21     FROM DuringBufferTable ),  
22 IgnoreBufInSeek(T, P) as (  
23     SELECT T, P FROM (  
24         SELECT T, DS, IF(P == 'Seek_ed' and DB, 'buffer', P) as P  
25         FROM DuringSeekTable ) WHERE NOT (P == 'buffer' AND DS) ),  
26 WithCDNAndQuery(T, P, C) as (  
27     SELECT T, P, NULL FROM IgnoreBufInSeek  
28     UNION SELECT T, NULL, C FROM heartbeats where C IS NOT NULL  
29     UNION SELECT 2022-07-21 10:05, NULL, NULL s),  
30 Intervals(Ed, St, State, CDN) as (  
31     SELECT T, LEAD(T, 1) OVER (PARTITION BY 1 ORDER BY T), P, C  
32     FROM (  
33         SELECT T,  
34         LAST(P) IGNORE NULLS OVER (PARTITION BY 1 ORDER BY T) as P,  
35         LAST(C) IGNORE NULLS OVER (PARTITION BY 1 ORDER BY T) as C  
36         FROM WithCDNAndQuery ) )  
37 SELECT SUM(St - Ed) as result FROM Intervals  
38 WHERE Ed < 2022-07-21 10:05 AND State == 'buffer' AND CDN == 'CDN1'
```

Difficult to develop

Semantic bugs



- Count the duration where:
1. Currently buffering &
 2. Play has already initialized &
 3. Hasn't seeked in last 5 seconds &
 4. Using CDN C1

Poor abstraction → High cost

```
1 WITH SeekAsPlayerState(T, P) as (  
2     SELECT T, P FROM heartbeats WHERE P IS NOT NULL  
3     UNION SELECT T, "Seek_st" FROM heartbeats WHERE A IS NOT NULL  
4     UNION SELECT T + 5, "Seek_ed" FROM heartbeats WHERE A IS NOT NULL ),  
5 IgnoreBufBeforePlay(T, P) as (  
6     SELECT T, P FROM (  
7         SELECT T, P, Max(If(P == 'play', 1, 0)) OVER (PARTITION BY 1 ORDER BY T)  
8         ↪ as H  
9         FROM SeekAsPlayerState) WHERE H == True ),  
10 DuringBufferTable(T, P, DB) as (  
11     SELECT T, P, LAST(tmp1) IGNORE NULLS OVER (PARTITION BY 1 ORDER BY T)  
12     FROM (  
13         SELECT T, P,  
14         CASE P WHEN 'buffer' THEN True WHEN 'Seek_st' THEN NULL WHEN 'Seek_ed'  
15         ↪ THEN NULL ELSE FALSE END as tmp1  
16         From IgnoreBufBeforePlay ) ),  
17 DuringSeekTable(T, P, DB, DS) as (  
18     SELECT T, P, DB,  
19     (T - Max(If(P == 'Seek_st', T, 0)) OVER (PARTITION BY 1 ORDER BY T)  
20     ↪ ) < 5 as tmp2  
21     FROM DuringBufferTable ),  
22 IgnoreBufInSeek(T, P) as (  
23     SELECT T, P FROM (  
24         SELECT T, DS, IF(P == 'Seek_ed' and DB, 'buffer', P) as P  
25         FROM DuringSeekTable ) WHERE NOT (P == 'buffer' AND DS) ),  
26 WithCDNAndQuery(T, P, C) as (  
27     SELECT T, P, NULL FROM IgnoreBufInSeek  
28     UNION SELECT T, NULL, C FROM heartbeats where C IS NOT NULL  
29     UNION SELECT 2022-07-21 10:05, NULL, NULL s),  
30 Intervals(Ed, St, State, CDN) as (  
31     SELECT T, LEAD(T, 1) OVER (PARTITION BY 1 ORDER BY T), P, C  
32     FROM (  
33         SELECT T,  
34         LAST(P) IGNORE NULLS OVER (PARTITION BY 1 ORDER BY T) as P,  
35         LAST(C) IGNORE NULLS OVER (PARTITION BY 1 ORDER BY T) as C  
36         FROM WithCDNAndQuery ) )  
37 SELECT SUM(St - Ed) as result FROM Intervals  
38 WHERE Ed < 2022-07-21 10:05 AND State == 'buffer' AND CDN == 'CDN1'
```

Lacks structure:

Difficult for query engines to optimize

High cost




- Count the duration where:
1. Currently buffering &
 2. Play has already initialized &
 3. Hasn't seeked in last 5 seconds &
 4. Using CDN C1

Outline for talk

- What is Time-State Analytics
- Time-State Analytics not well supported by status quo
- *Introducing the Timeline abstraction*
- Early Wins + Next Steps

Stepping back



Timestamp	Player State	Bitrate	CDN	Seek
t1	Buffer			
t2			C1	
t3		B1		
t4	Play			
t5	Buffer			
t6	Play			
t7				Seek
t8		B2		
t9	Paused			
t10	Play			
t11		B3		
t12	Buffer			
t13			C2	
t14	Play			
t15		B1		

What's a Timeline?

“Geometric abstractions are powerful tools” – Fred Brooks


An equivalent *geometric* view:
Timeline of each measurement

Timestamp	Player State	Bitrate	CDN	Seek
t1	Buffer			
t2			C1	
t3		B1		
t4	Play			
t5	Buffer			
t6	Play			
t7				Seek
t8		B2		
t9	Paused			
t10	Play			
t11		B3		
t12	Buffer			
t13			C2	
t14	Play			
t15		B1		




Column X

Type: Step Function





Column Y

Type: Event

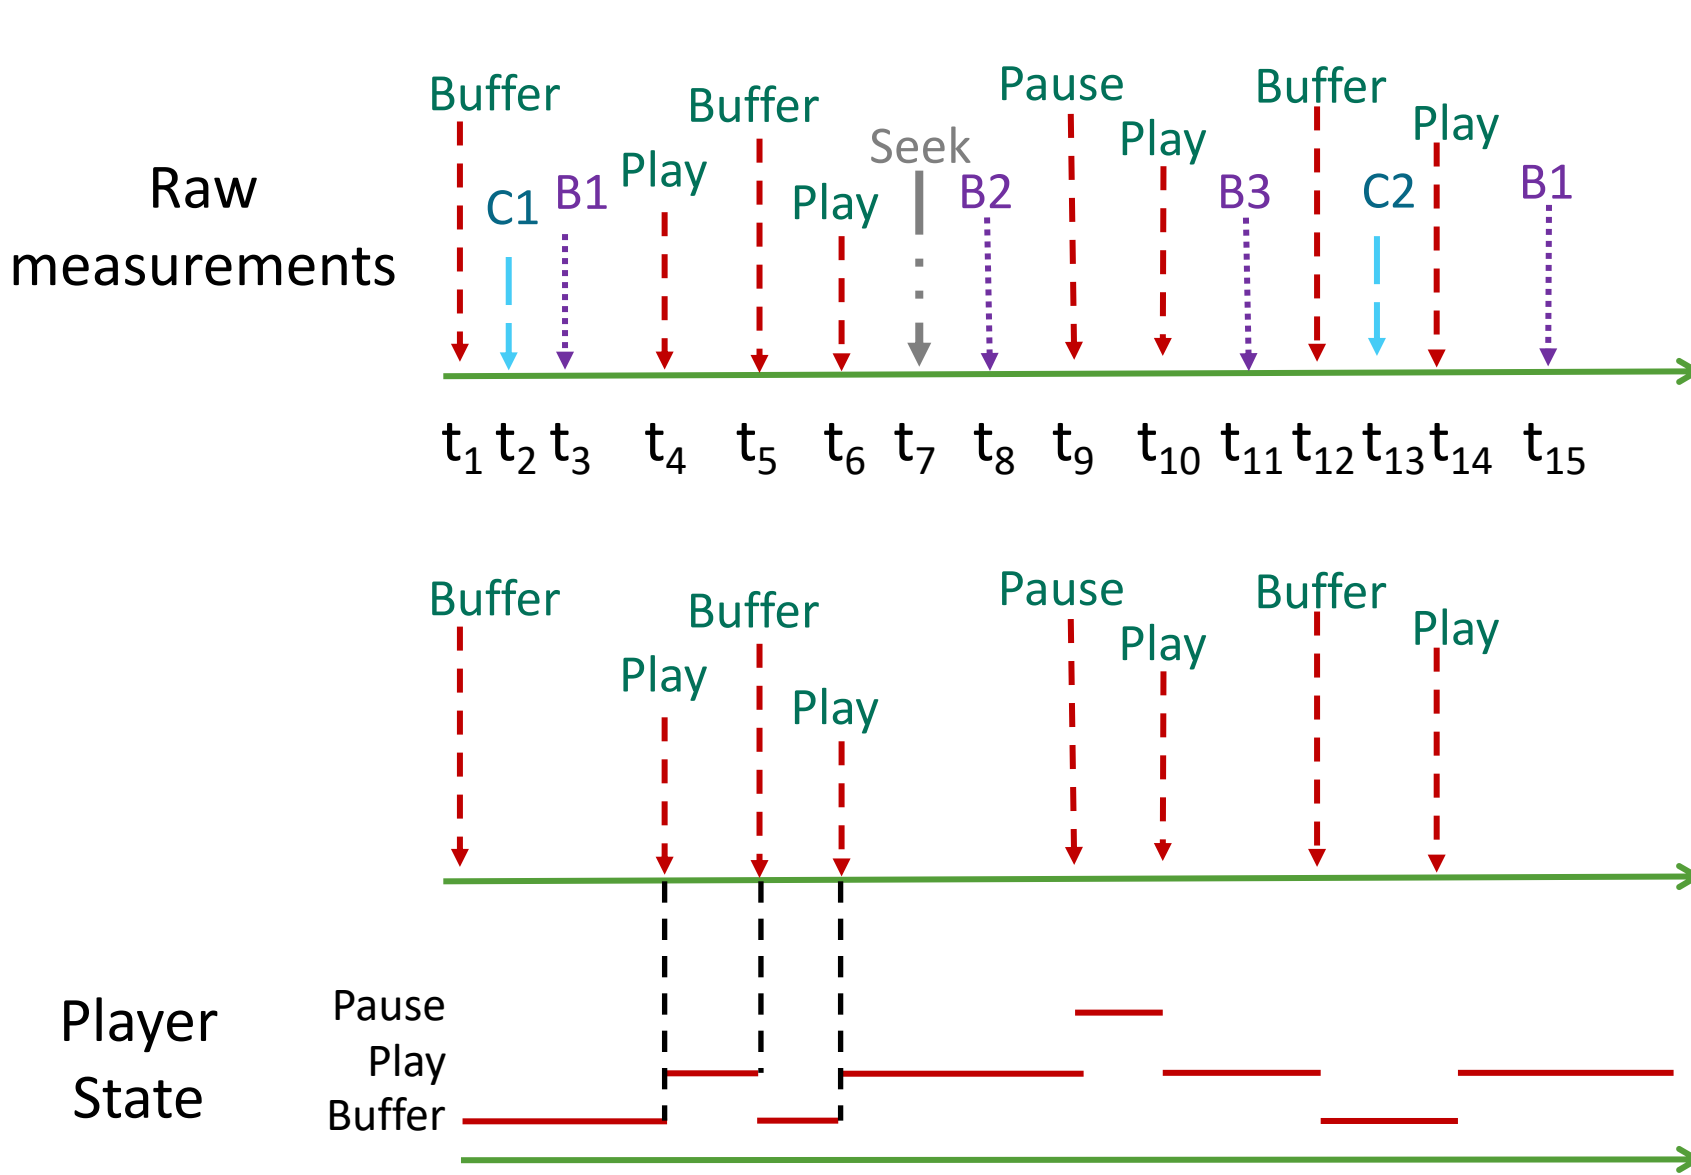


Column Z

Type: Continuous Value




Whiteboarding Timelines: the Player State over Time



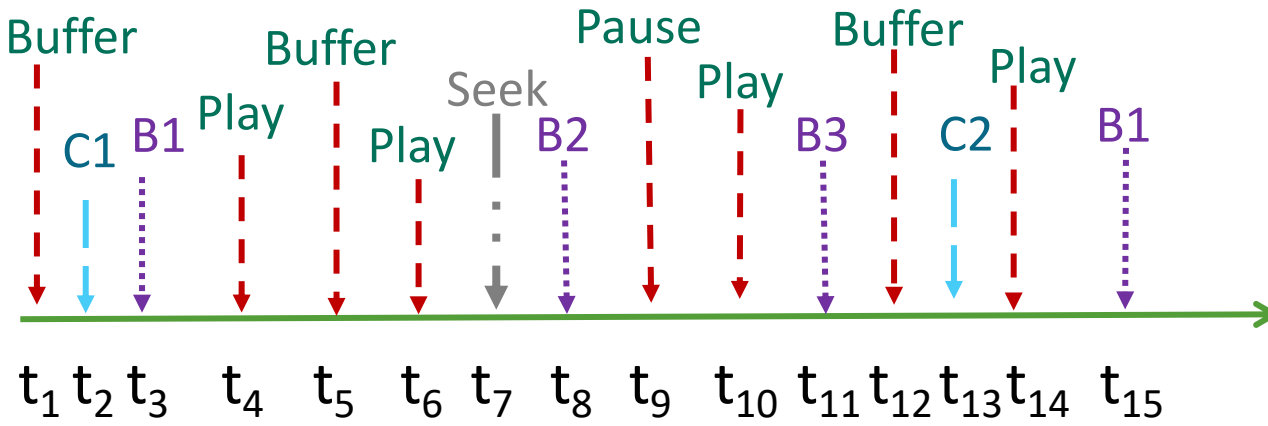
Count the duration where:

1. Currently buffering &
2. Play has already initialized &
3. Hasn't seeked in last 5 seconds &
4. Using CDN C1




Operator:
 LatestEventToState
 (PlayerState)

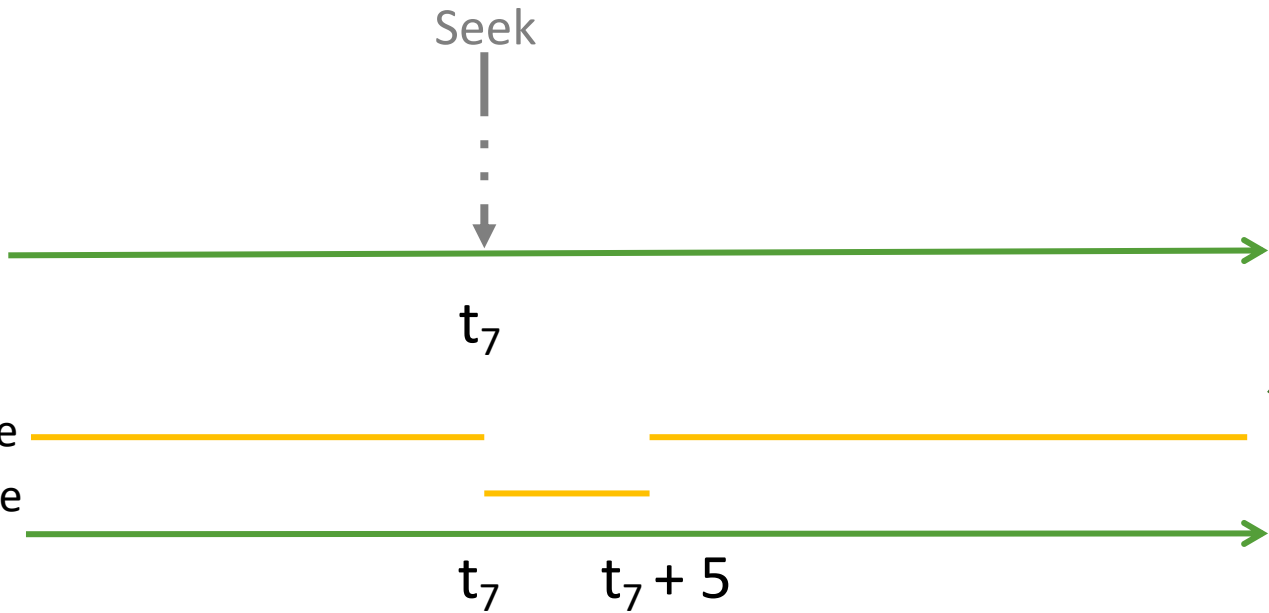
Whiteboarding Timelines: When has a Seek recently happened?



- Count the duration where:
1. Currently buffering &
 2. Play has already initialized &
 3. **Hasn't seeked in last 5 seconds** &
 4. Using CDN C1



```
Operator:
TimeSince(Seek) >
5 seconds
```




Didn't recently seek

True
False

Whiteboarding: Connection-Induced Rebuffering w/ C1


Count the duration where:

1. Currently buffering &
2. Play has already initialized &
3. Hasn't seeked in last 5 seconds &
4. Using CDN C1




Operator:
And


Connection-induced rebuffering




CDN = C1




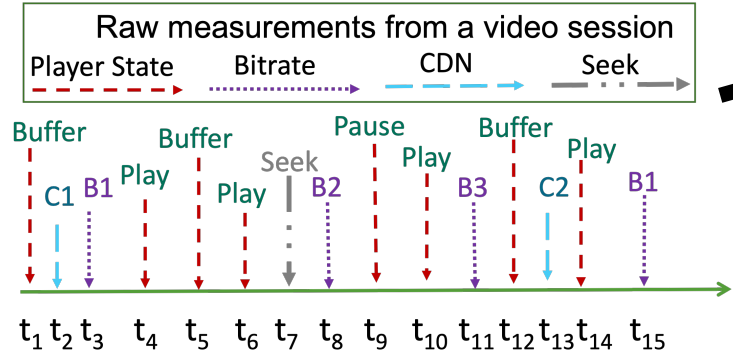
Connection-induced rebuffering while CDN = C1





Count the duration where:

1. Currently buffering &
2. Play has already initialized &
3. Hasn't seeked in last 5 seconds &
4. Using CDN C1



LatestEventToState(PlayerState) = "Buffer"

HasBeenTrue(PlayerState = "Play")




TimeSince(Seek) > 5 seconds

LatestEventToState(CDN) = "C1"


And → DurationWhere

Timeline, in a nutshell

Data abstraction with 3 types of timeline dynamics



Compositional language for defining DAG of operators



Library of operators

Timeline generalizations of classical Operators		
$==, <, >$ [constant]	Compare each update or state with a fixed value, producing True or False	
$\&, $ [timeline]	Combine 2 timelines by applying a logical operation at each point in time	
\sim	Logically invert each update or state	
Timeline-specific Operators		
TL_HASEXISTED	A StateDynamics timeline of the cumulative OR	
TL_HASEXISTEDWITHIN	As TL_HASEXISTED, but resets to False after a specified duration D without True values	
TL_LATESTEVENTTOSTATE	A StateDynamics Timeline of the latest update	
TL_DURATIONWHERE	A Numerical Timeline of the cumulative duration where the state was True	
TL_DURATIONINCURSTATE	A Numerical Timeline of the duration since the last state change	

...

Connectors with external data sources/sinks

Outline for talk

- What is Time-State Analytics
- Time-State Analytics not well supported by status quo
- Introducing the Timeline abstraction
- **Early Promise + Next Steps**

Timeline Reduced Dev Effort at Conviva

```
1 WITH SeekAsPlayerState(T, P) as (  
2     SELECT T, P FROM heartbeats WHERE P IS NOT NULL  
3     UNION SELECT T, "Seek_st" FROM heartbeats WHERE A IS NOT NULL  
4     UNION SELECT T + 5, "Seek_ed" FROM heartbeats WHERE A IS NOT NULL ),  
5 IgnoreBufBeforePlay(T, P) as (  
6     SELECT T, P FROM (  
7         SELECT T, P, Max(If(P == 'play', 1, 0)) OVER (PARTITION BY 1 ORDER BY T)  
8         ↪ as H  
9         FROM SeekAsPlayerState) WHERE H == True ),  
10 DuringBufferTable(T, P, DB) as (  
11     SELECT T, P, LAST(tmp1) IGNORE NULLS OVER (PARTITION BY 1 ORDER BY T)  
12     FROM (  
13         SELECT T, P,  
14         CASE P WHEN 'buffer' THEN True WHEN 'Seek_st' THEN NULL WHEN 'Seek_ed'  
15         ↪ THEN NULL ELSE FALSE END as tmp1  
16         From IgnoreBufBeforePlay ) ),  
17 DuringSeekTable(T, P, DB, DS) as (  
18     SELECT T, P, DB,  
19     (T - Max(If(P == 'Seek_st', T, 0)) OVER (PARTITION BY 1 ORDER BY T)  
20     ↪ ) < 5 as tmp2  
21     FROM DuringBufferTable ),  
22 IgnoreBufInSeek(T, P) as (  
23     SELECT T, P FROM (  
24         SELECT T, DS, IF(P == 'Seek_ed' and DB, 'buffer', P) as P  
25         FROM DuringSeekTable ) WHERE NOT (P == 'buffer' AND DS) ),  
26 WithCDNAndQuery(T, P, C) as (  
27     SELECT T, P, NULL FROM IgnoreBufInSeek  
28     UNION SELECT T, NULL, C FROM heartbeats where C IS NOT NULL  
29     UNION SELECT 2022-07-21 10:05, NULL, NULL s),  
30 Intervals(Ed, St, State, CDN) as (  
31     SELECT T, LEAD(T, 1) OVER (PARTITION BY 1 ORDER BY T), P, C  
32     FROM (  
33         SELECT T,  
34         LAST(P) IGNORE NULLS OVER (PARTITION BY 1 ORDER BY T) as P,  
35         LAST(C) IGNORE NULLS OVER (PARTITION BY 1 ORDER BY T) as C  
36         FROM WithCDNAndQuery ) )  
37 SELECT SUM(St - Ed) as result FROM Intervals  
38 WHERE Ed < 2022-07-21 10:05 AND State == 'buffer' AND CDN == 'CDN1'
```

**Onboarding:
Weeks → Days**


**Semantic Bugs:
Dropped by 80%**



```
1 SELECT TL_DURATIONWHERE(  
2     TL_LATESTEVENTTOSTATE(playerStateChange) = 'buffer' AND  
3     TL_HASEXISTED(playerStateChange = 'play') AND  
4     NOT TL_HASEXISTEDWITHIN(userAction = 'seek', 5s) AND  
5     TL_LATESTEVENTTOSTATE(cdnChange) = 'CDN1'  
6 ) AS result  
7 FROM heartbeats  
8 TIMELINE WITH EVENT TIME t  
9 EVALUATE AT EVENT TIME 2022-07-21 10:05:00
```

Prototype query language

Timeline Offers Reduced Cost



2-10X Faster Execution Time!

Future Outlook

- Applications to many domains
 - Cybersecurity, IoT, logistics, manufacturing, ...
- Visual interfaces to democratize TSA
- Even better performance
- Streaming implementation

Takeaways

- Growing need for *Time-State Analytics* across different domains
- Fundamentally hard problem:
Stateful, Context-Sensitive, Continuous
- State-of-art systems (e.g., streaming systems, data warehouses, RDBMS) ill suited
 - Why: Classical tabular model for data processing is ineffective for Time-State Analytics
 - Great for simple stateless filter/aggregation but not Time-State Analytics
 - High cost, low performance + High dev effort, many bugs
- Our work: Timeline → *A geometric abstraction* for Time-State Analytics
 - Early promise: Up to 10X better cost/performance AND 10X reduced effort
 - New opportunities: Generality, No-code Intents