The Cache Sketch: Revisiting Expiration-based Caching in the Age of Cloud Data Management

Felix Gessert, Michael Schaarschmidt, Wolfram Wingerath, Steffen Friedrich, Norbert Ritter
gessert@informatik.uni-hamburg.de

The Latency Problem

100 ms

Average: 9,3s

Loading…

-1% Revenue

The Latency Problem

400 ms

-9% Visitors

-1% Revenue

Google

YAHOO!

amazon.com

The Latency Problem

500 ms

Average: 9,3s

Loading…

-20% Traffic — Google

-9% Visitors — Yahoo!

-1% Revenue — amazon.com

The Latency Problem

1s

Average: 9,3s

Loading…

-7% Conversions — AberdeenGroup

-20% Traffic — Google

-9% Visitors — YAHOO!

-1% Revenue — amazon.com

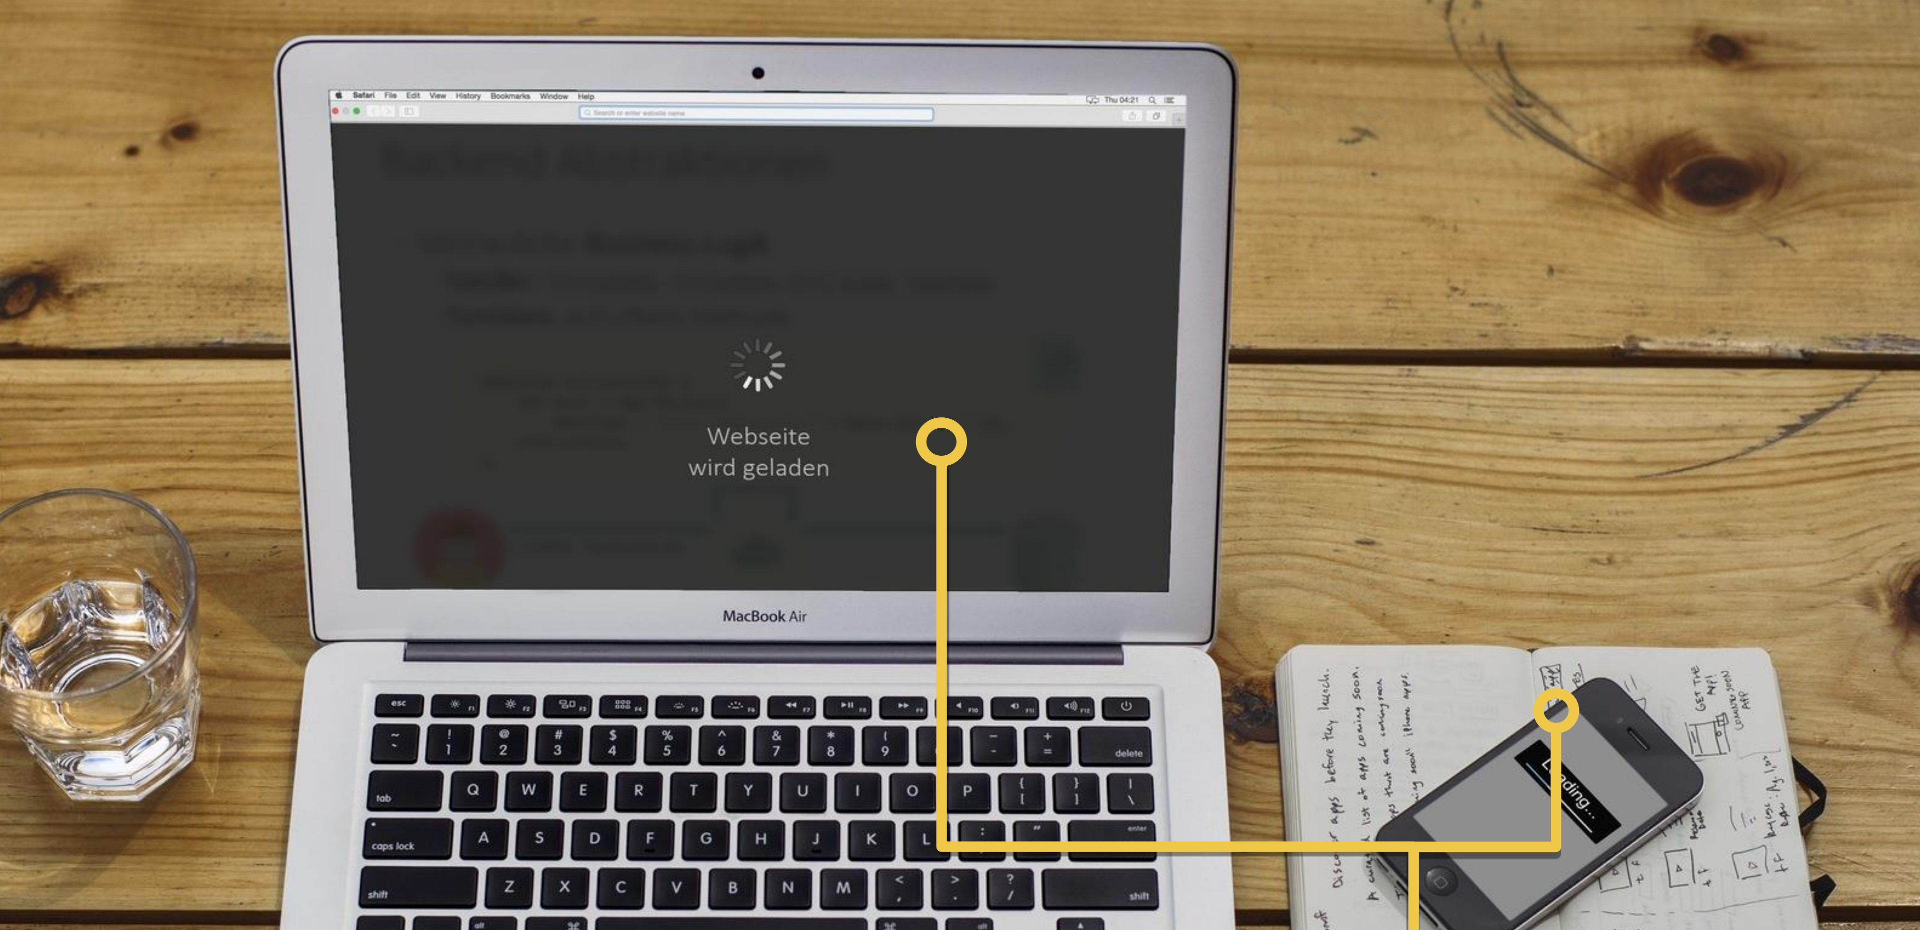

If perceived speed is such an import factor

...what causes slow page load times?

State of the art

Two bottlenecks: latency und processing

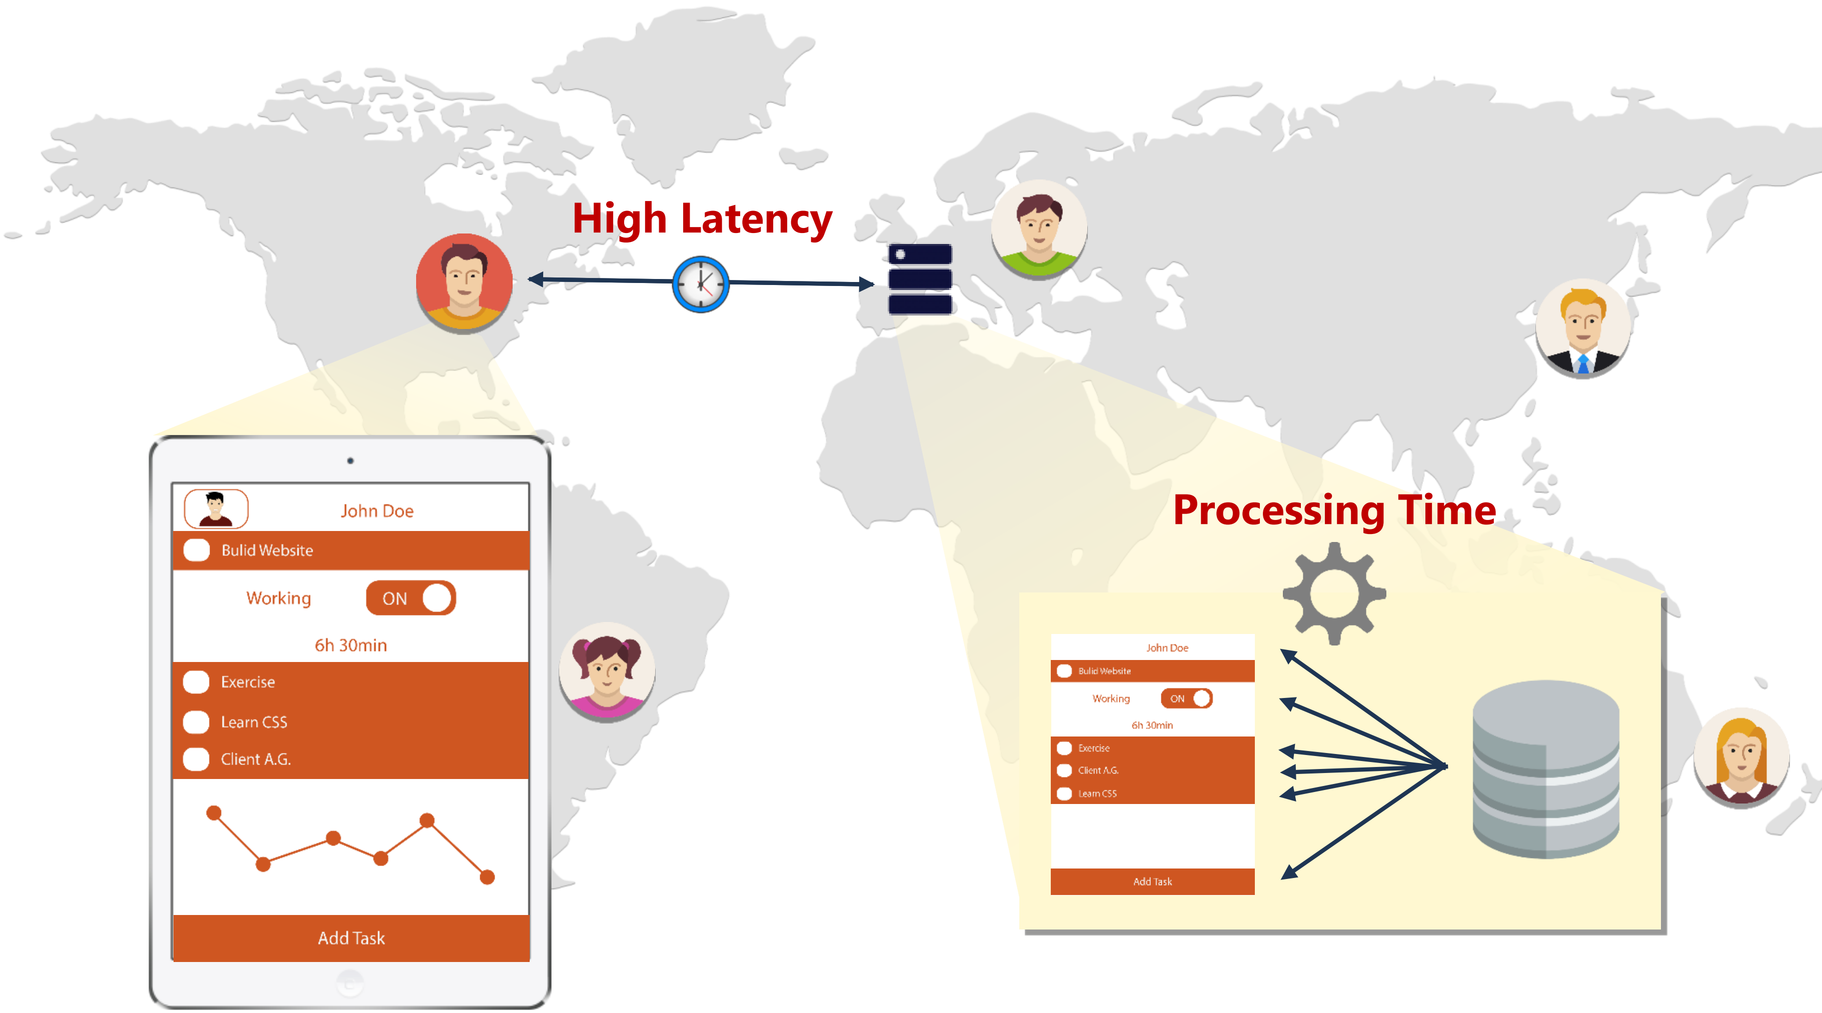

Network Latency

The underlying problem of high page load times

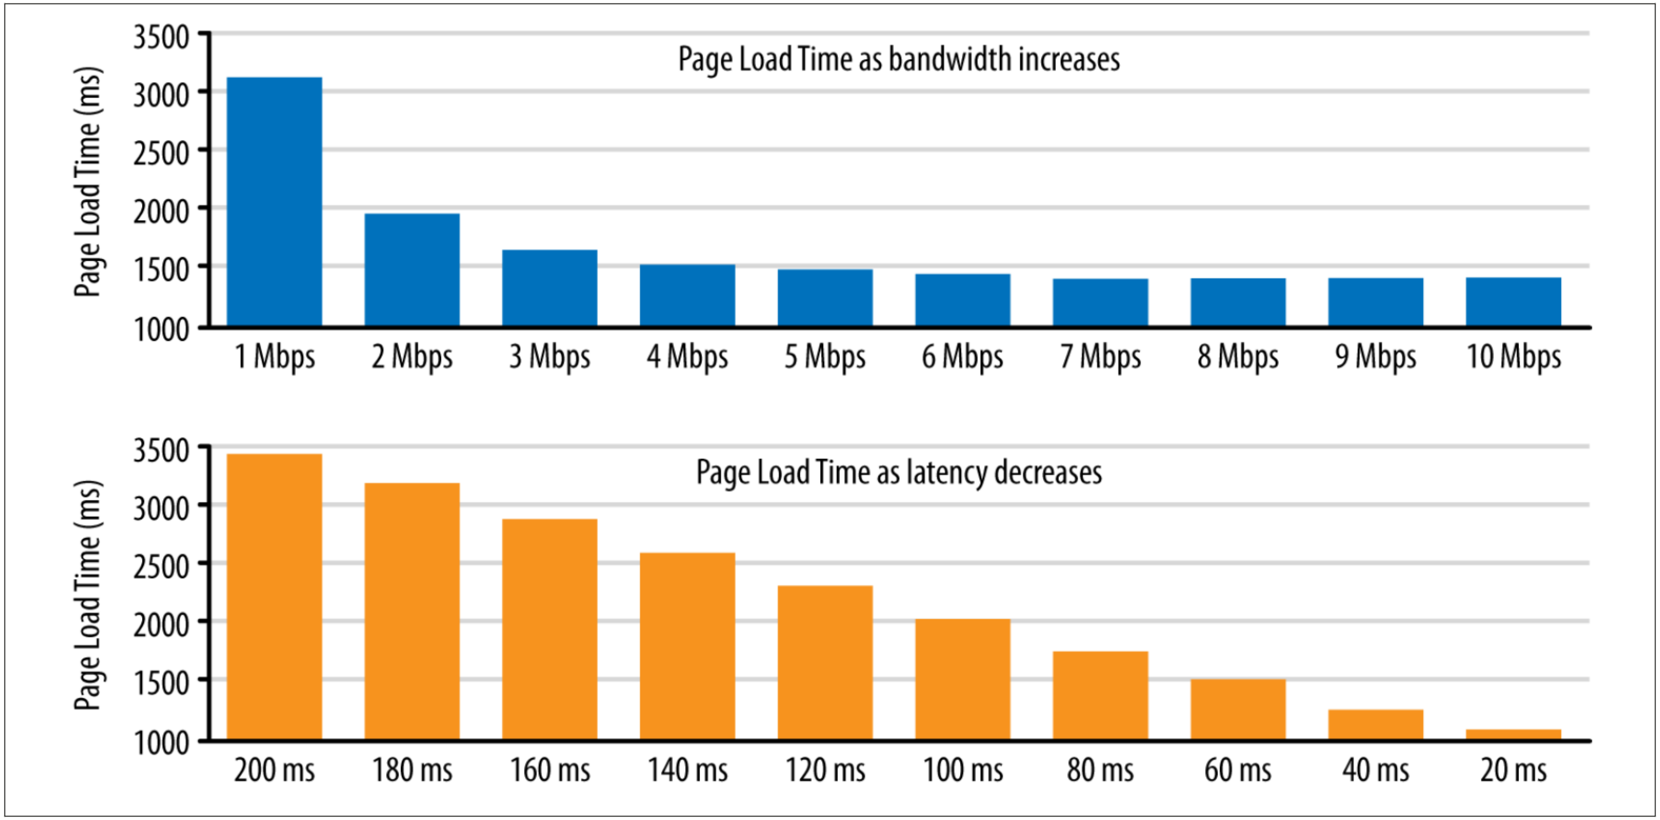

I. Grigorik, High performance browser networking. O'Reilly Media, 2013.

The low-latency vision

Data is served by ubiquitous web-caches

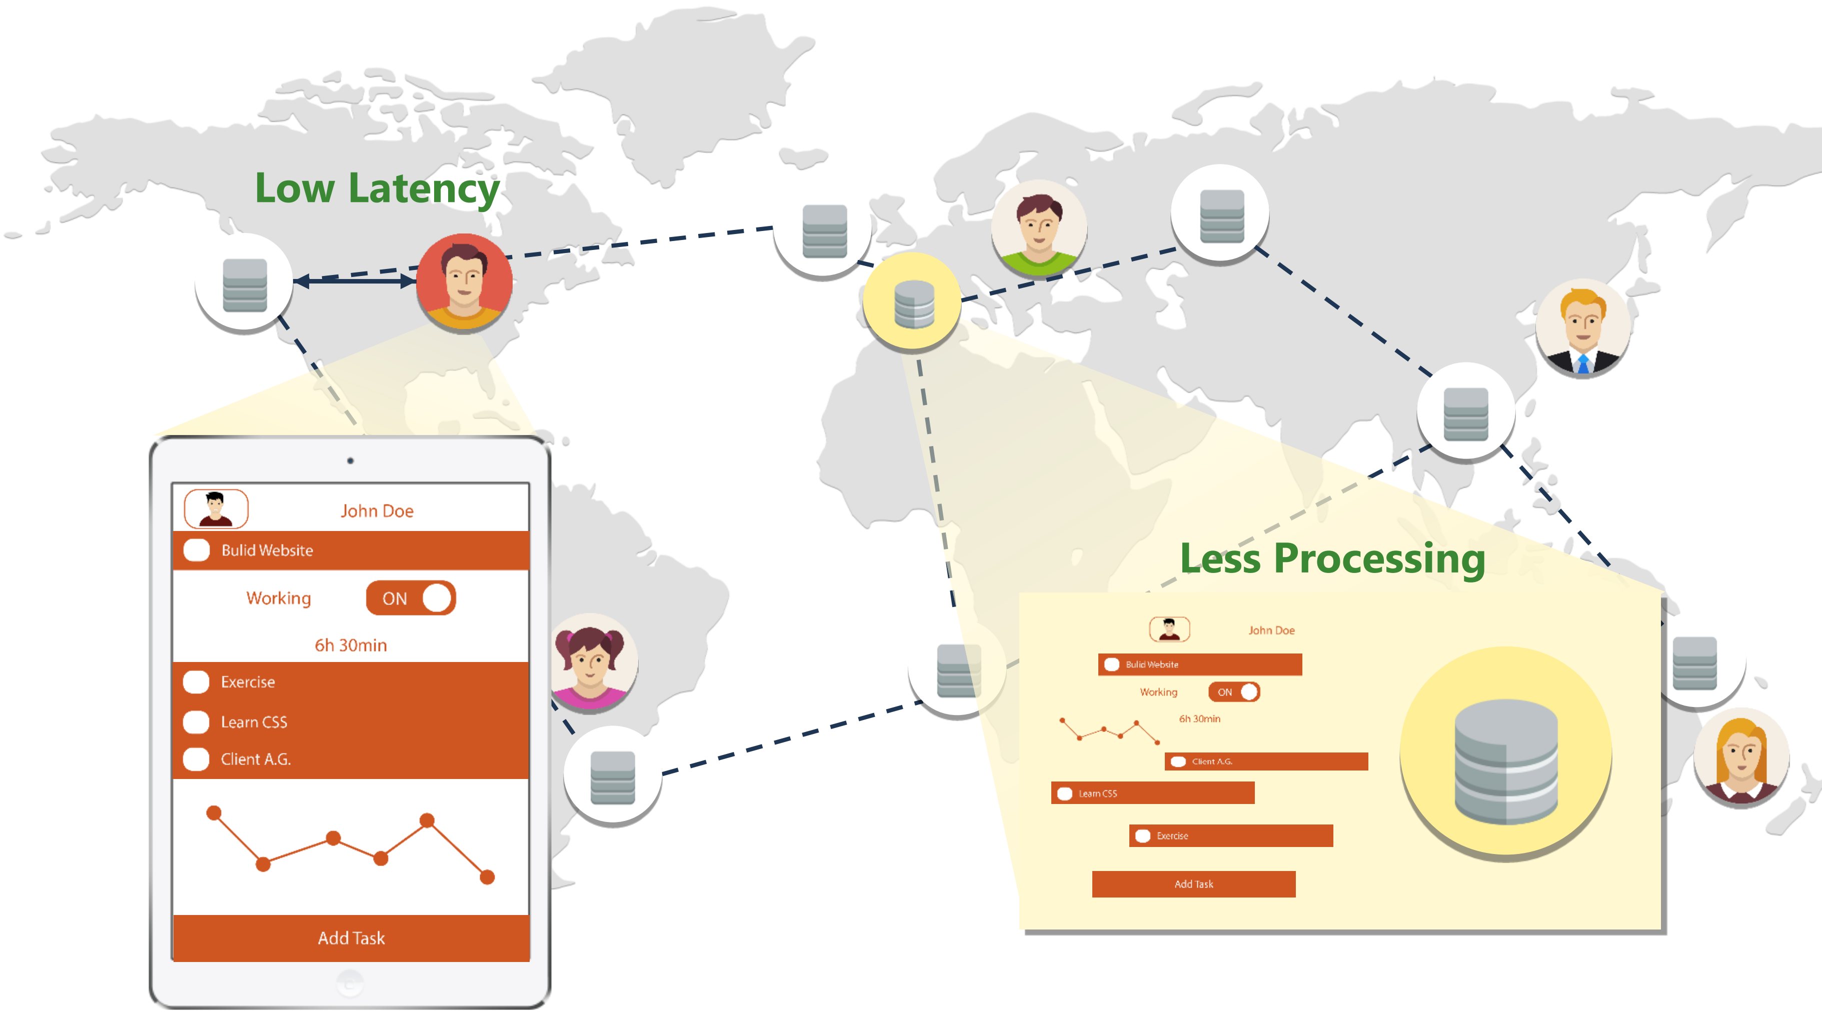

The web's caching model

Staleness as a consequence of scalability

Expiration-based

Every object has a defined Time-To-Live (TTL)

Revalidations

Allow clients and caches to check freshness at the server

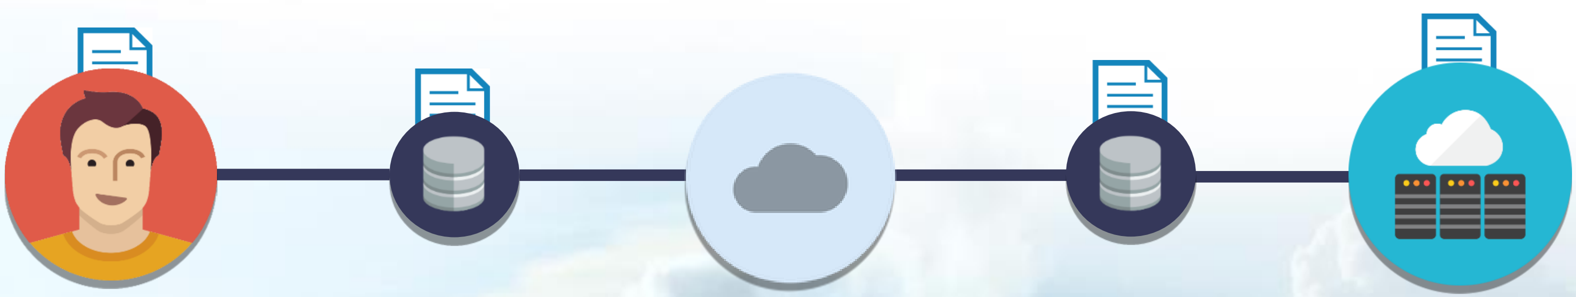

The web's caching model

Staleness as a consequence of scalability

Expiration-based
Every object has a defined Time-To-Live (TTL)

Revalidations
Allow clients and caches to check freshness at the server

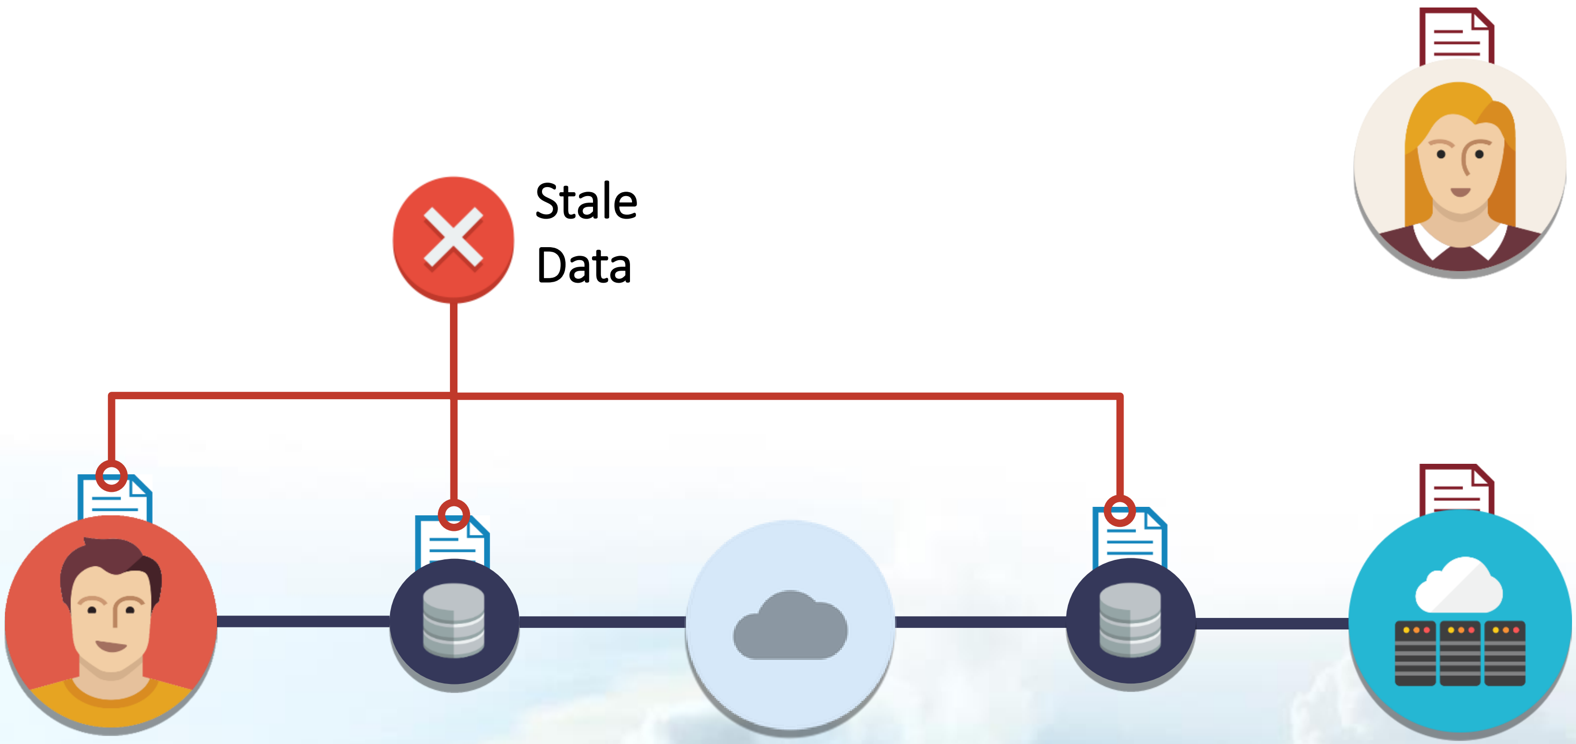

The web's caching model

Staleness as a consequence of scalability

Expiration-based

Every object has a defined Time-To-Live (TTL)

Revalidations

Allow clients and caches to check freshness at the server

Research Question:

Can database services leverage the web caching infrastructure for low latency with rich **consistency guarantees**?

Web Caching Concepts

Invalidation- and expiration-based caches

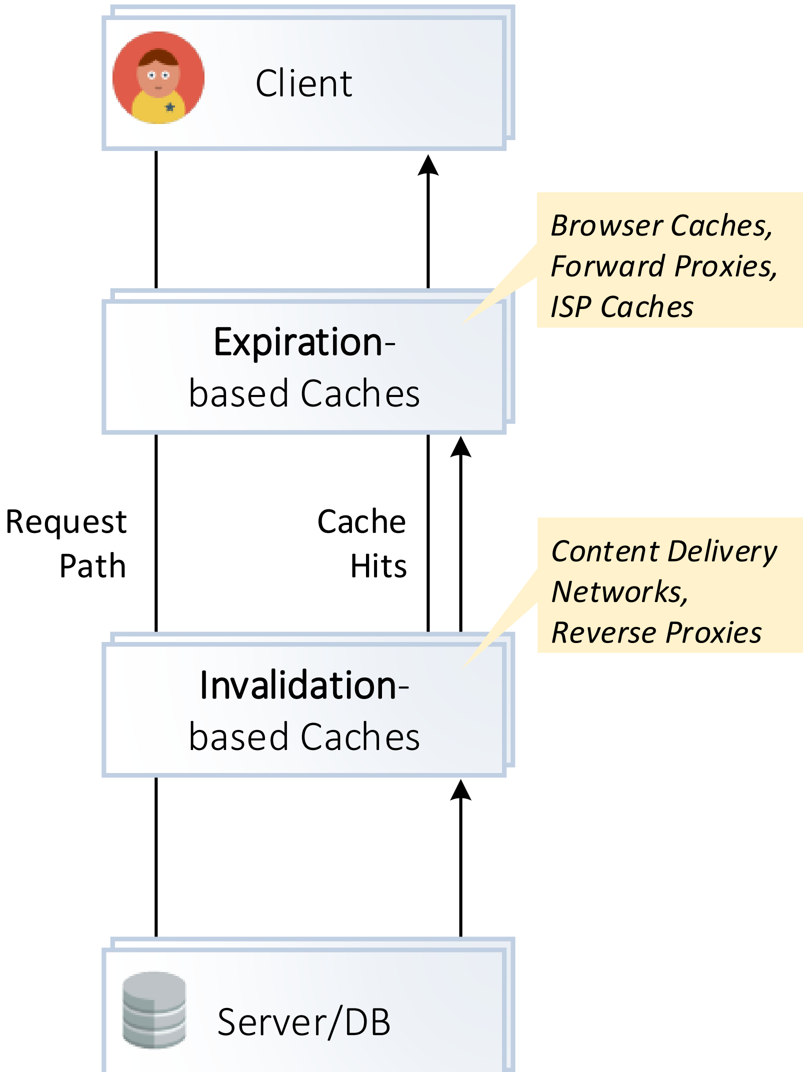

Expiration-based Caches:

- An object x is considered fresh for TTL_x seconds
- The server assigns TTLs for each object

Invalidation-based Caches:

- Expose object eviction operation to the server

The Cache Sketch approach

Letting the client handle cache coherence

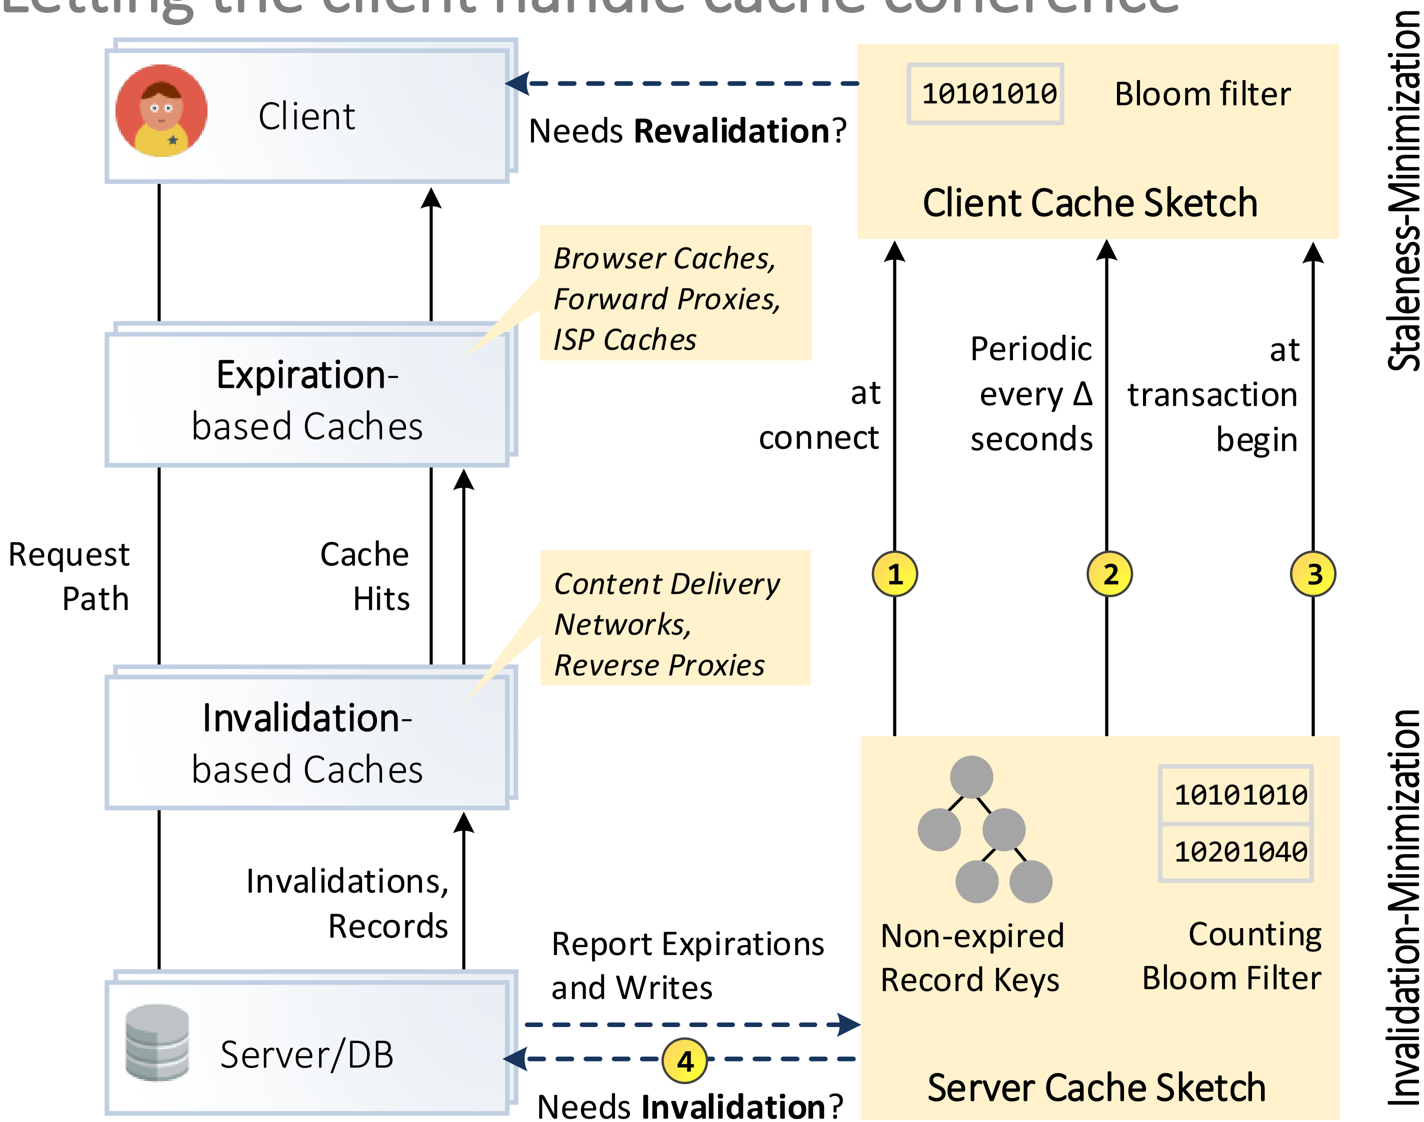

The End to End Path of Request

The Caching Hierarchy

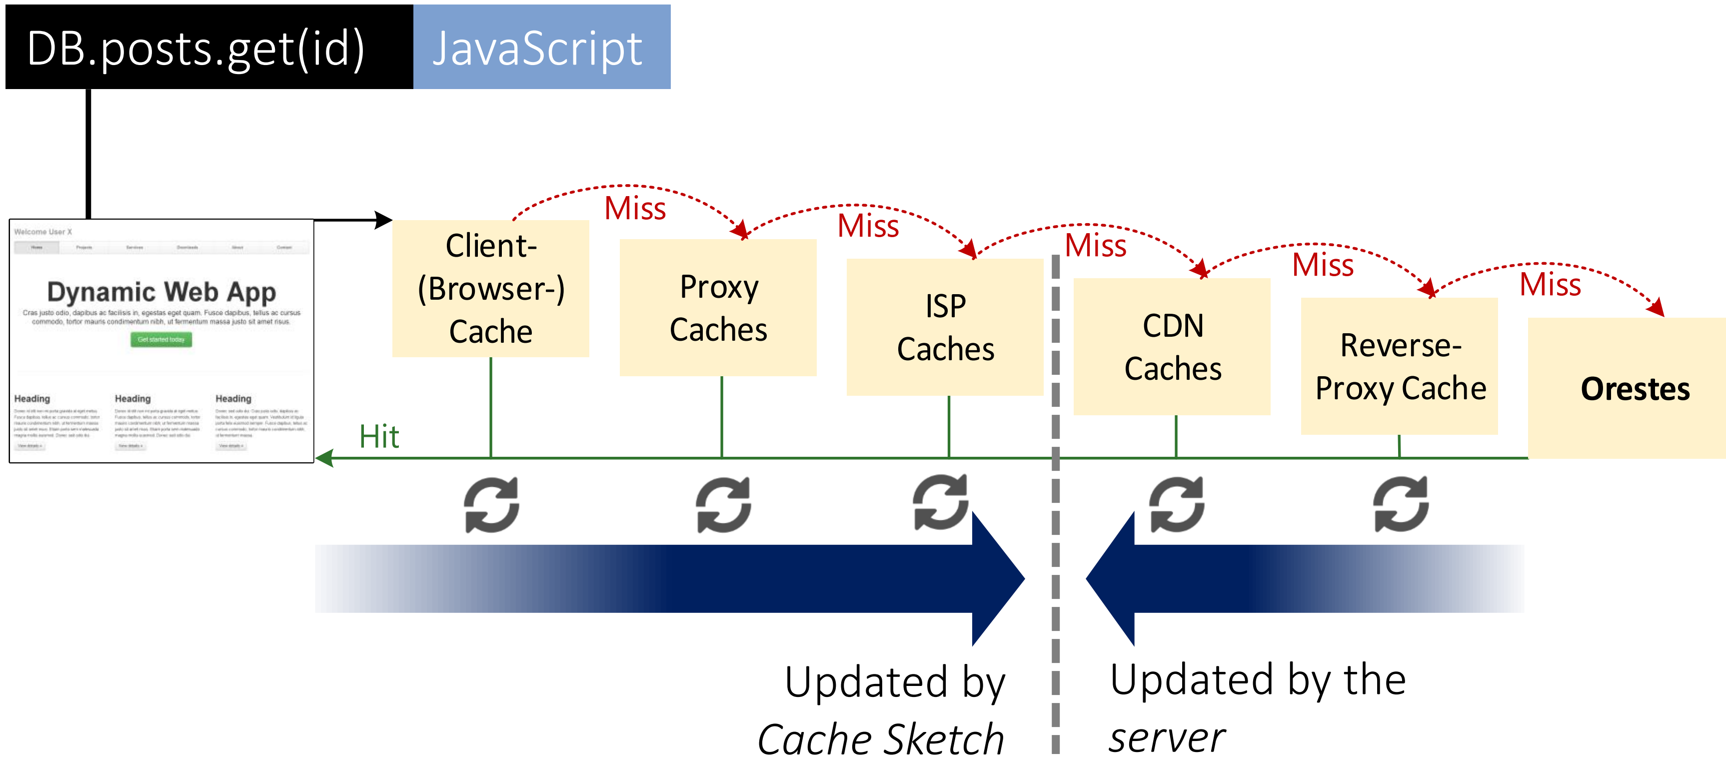

The End to End Path of Request

The Caching Hierarchy

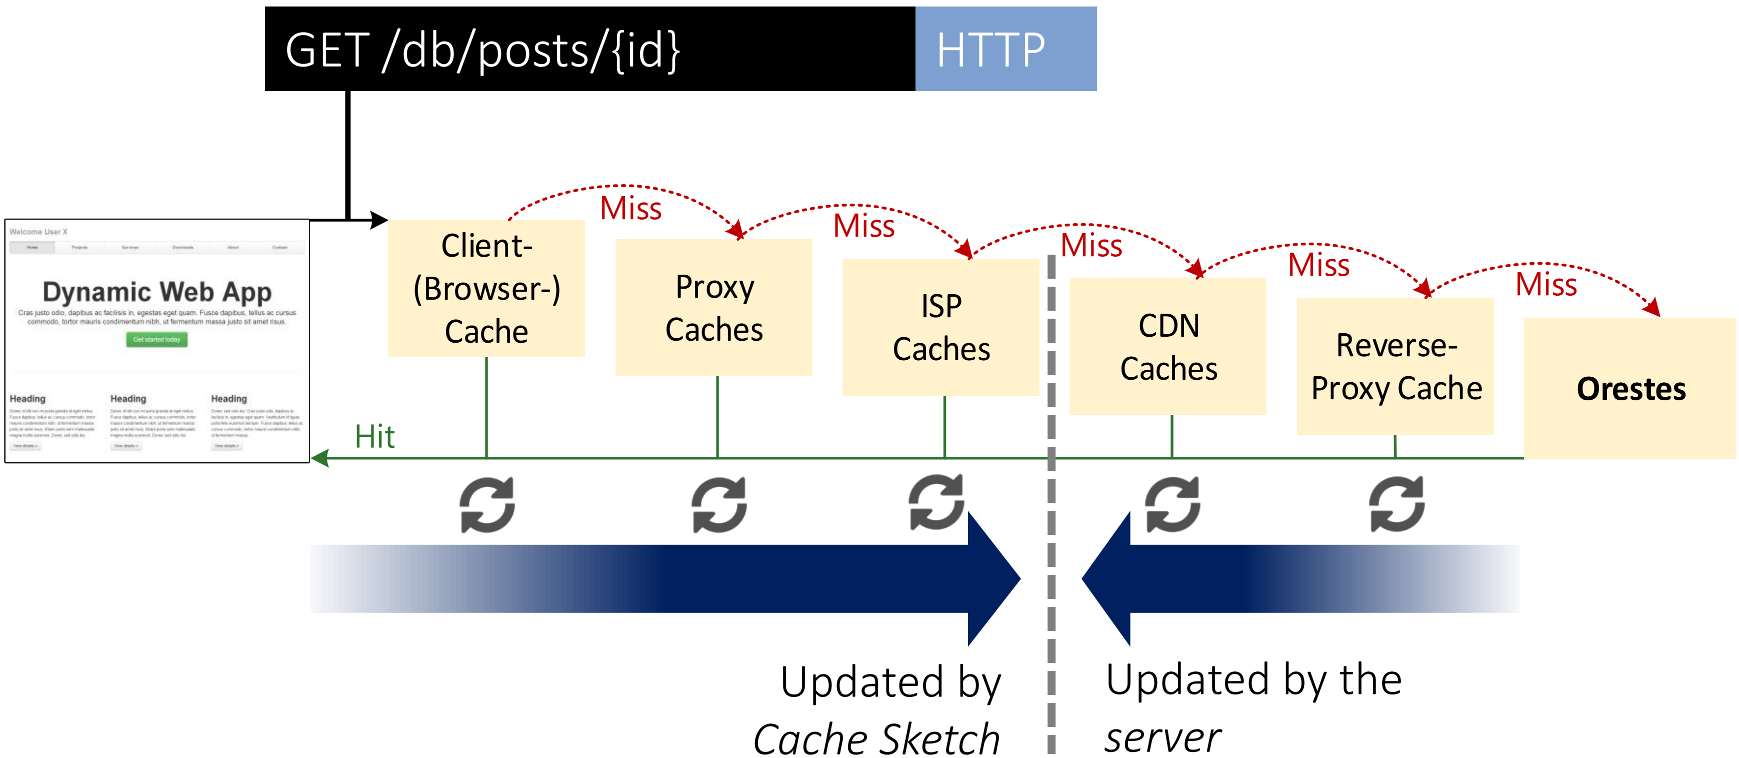

The End to End Path of Request

The Caching Hierarchy

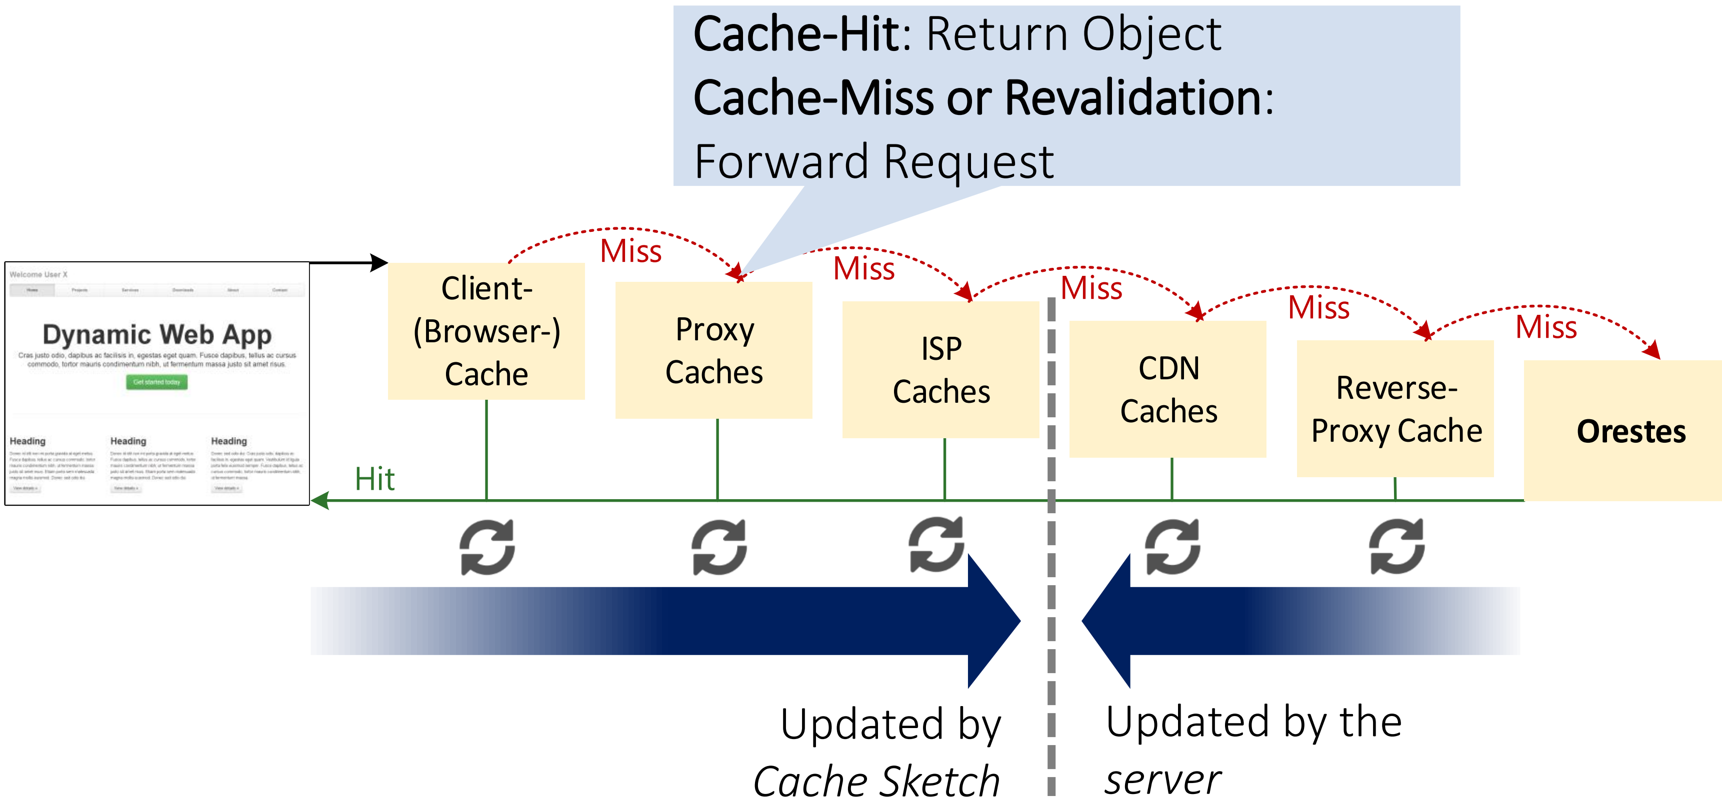

The End to End Path of Request

The Caching Hierarchy

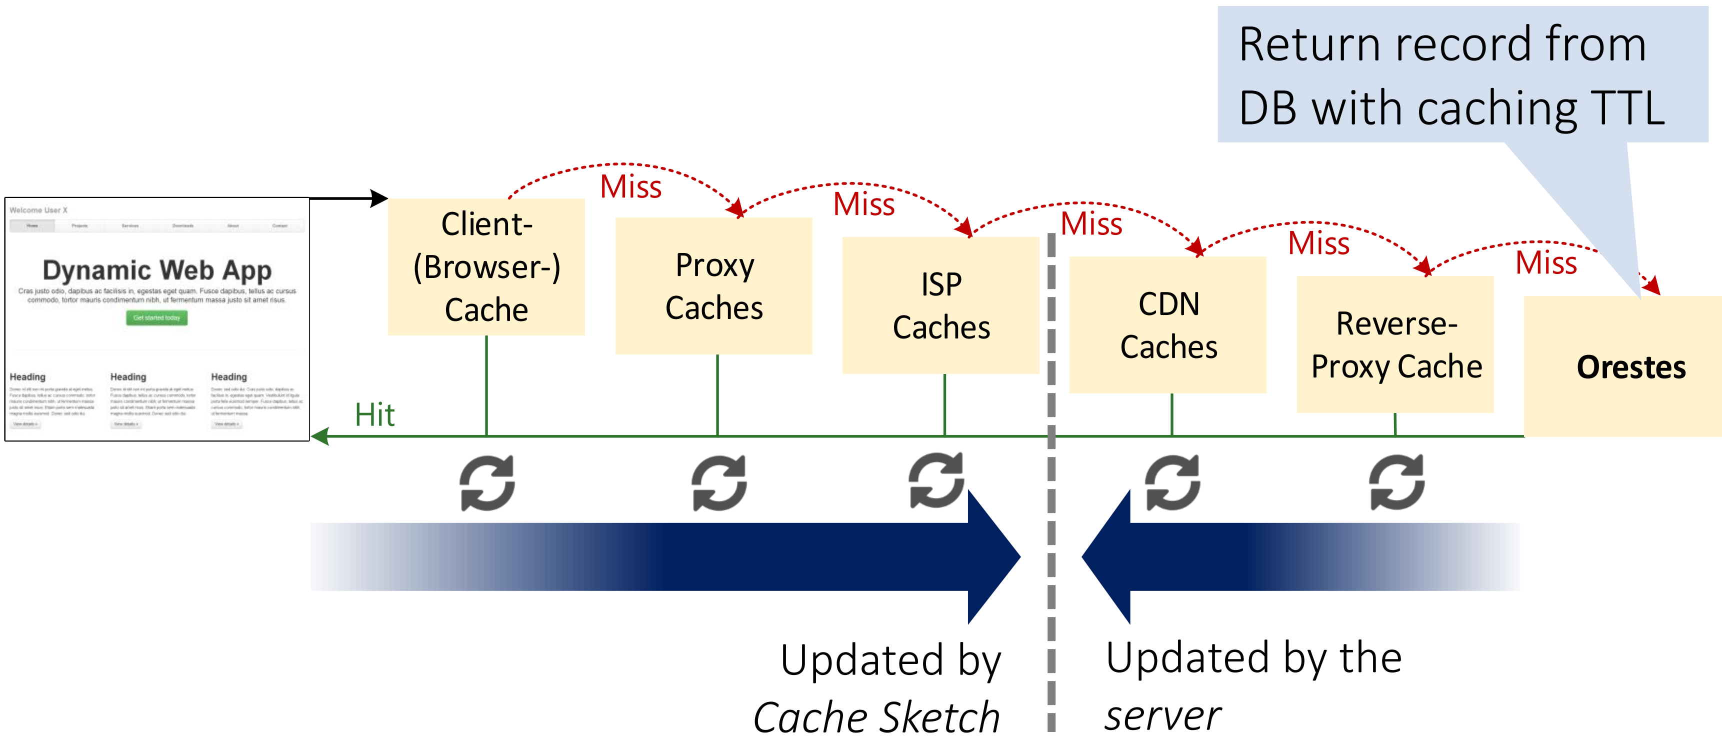

The End to End Path of Request

The Caching Hierarchy

Low Latency

Reduced Database Load

Flash-Crowd Protection

Higher Availability

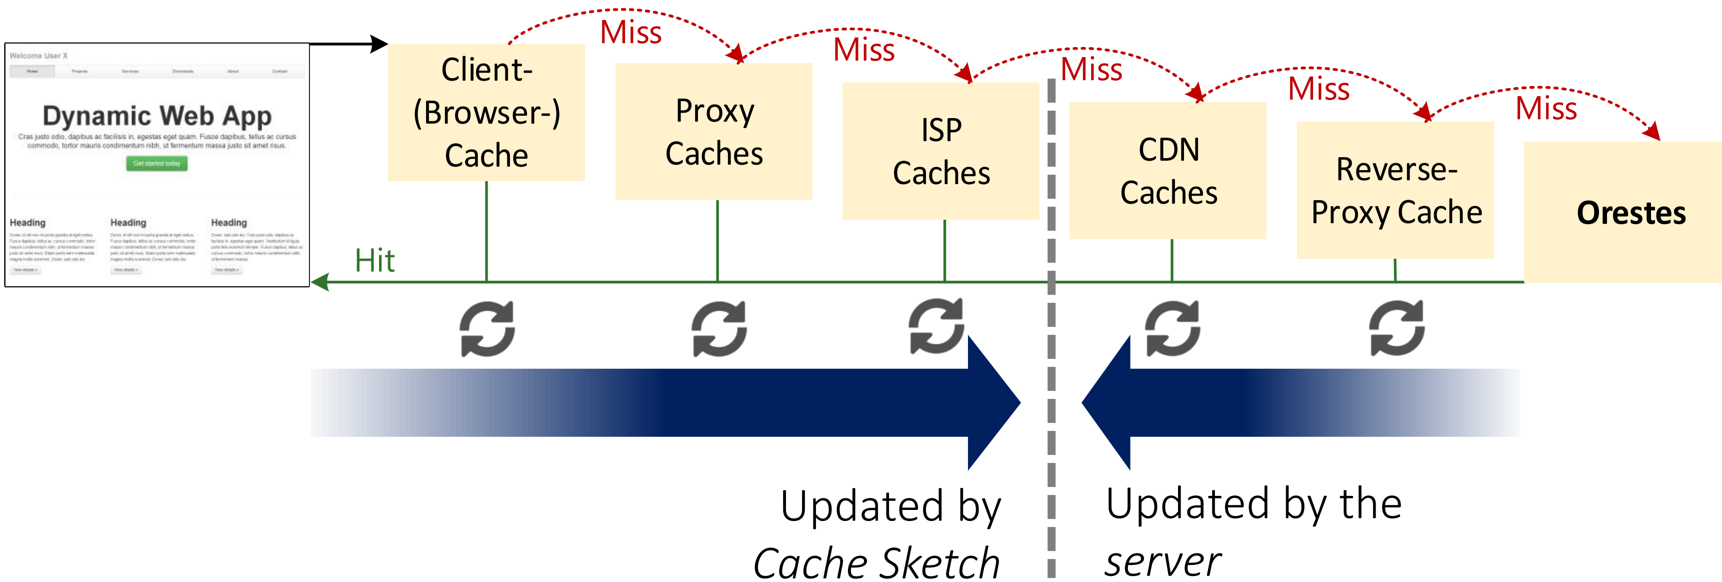

The Client Cache Sketch

- *Let* $\boldsymbol{c_t}$ *be the client Cache Sketch generated at time* **t**, *containing the key* $\boldsymbol{key_x}$ *of every record* **x** *that was written before it expired in all caches, i.e. every* **x** *for which holds:*

$$\exists\, r(x, t_r, TTL), w(x, t_w) :\ t_r + TTL > t > t_w > t_r$$

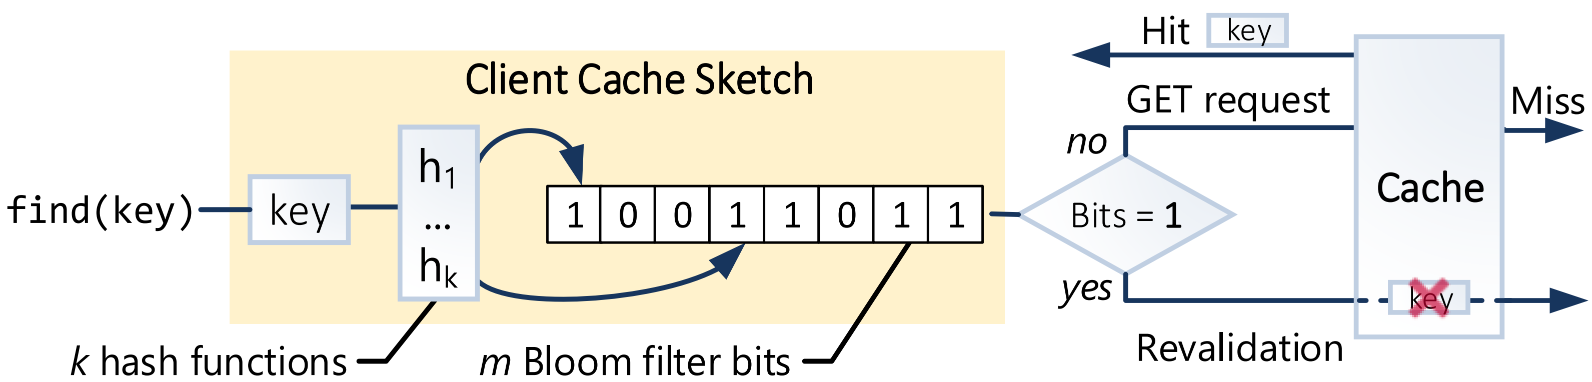

1 Slow initial page loads

- Solution: **Cached Initialization**
  - Clients load the Cache Sketch at connection
  - Every non-stale cached record can be reused without degraded consistency

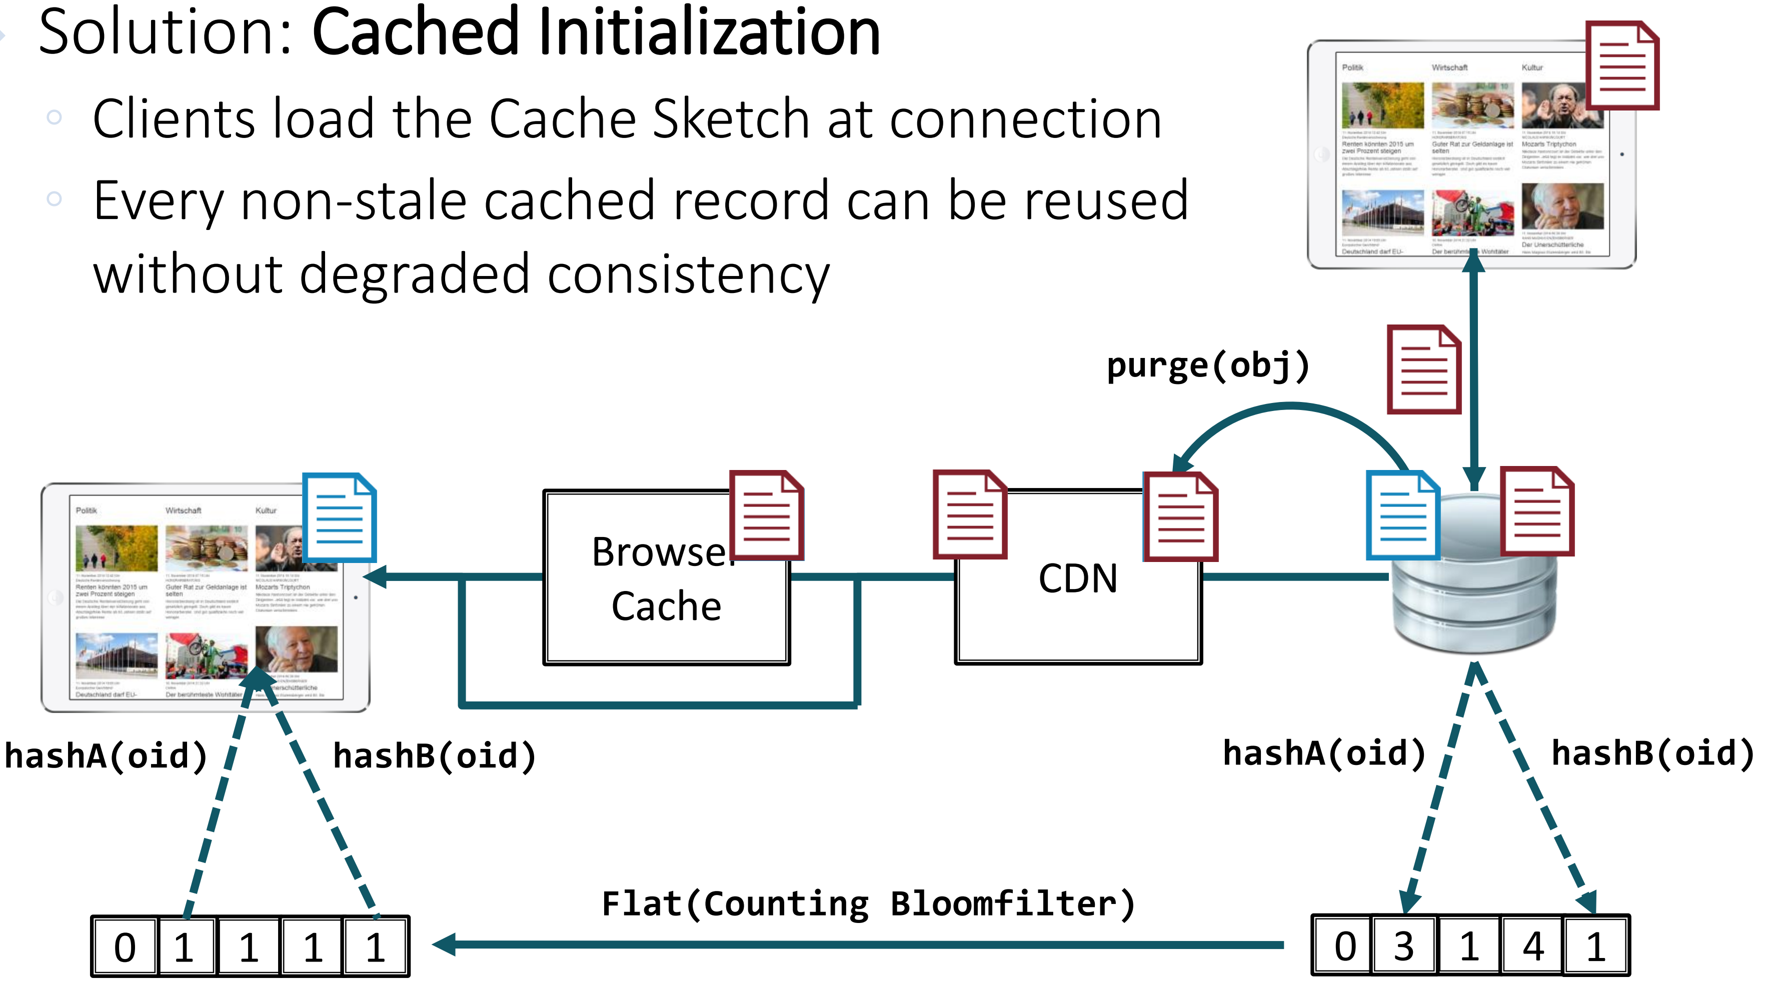

1 Slow initial page loads

- Solution: **Cached Initialization**
  - Clients load the Cache Sketch at connection
  - Every non-stale cached record can be reused without degraded consistency

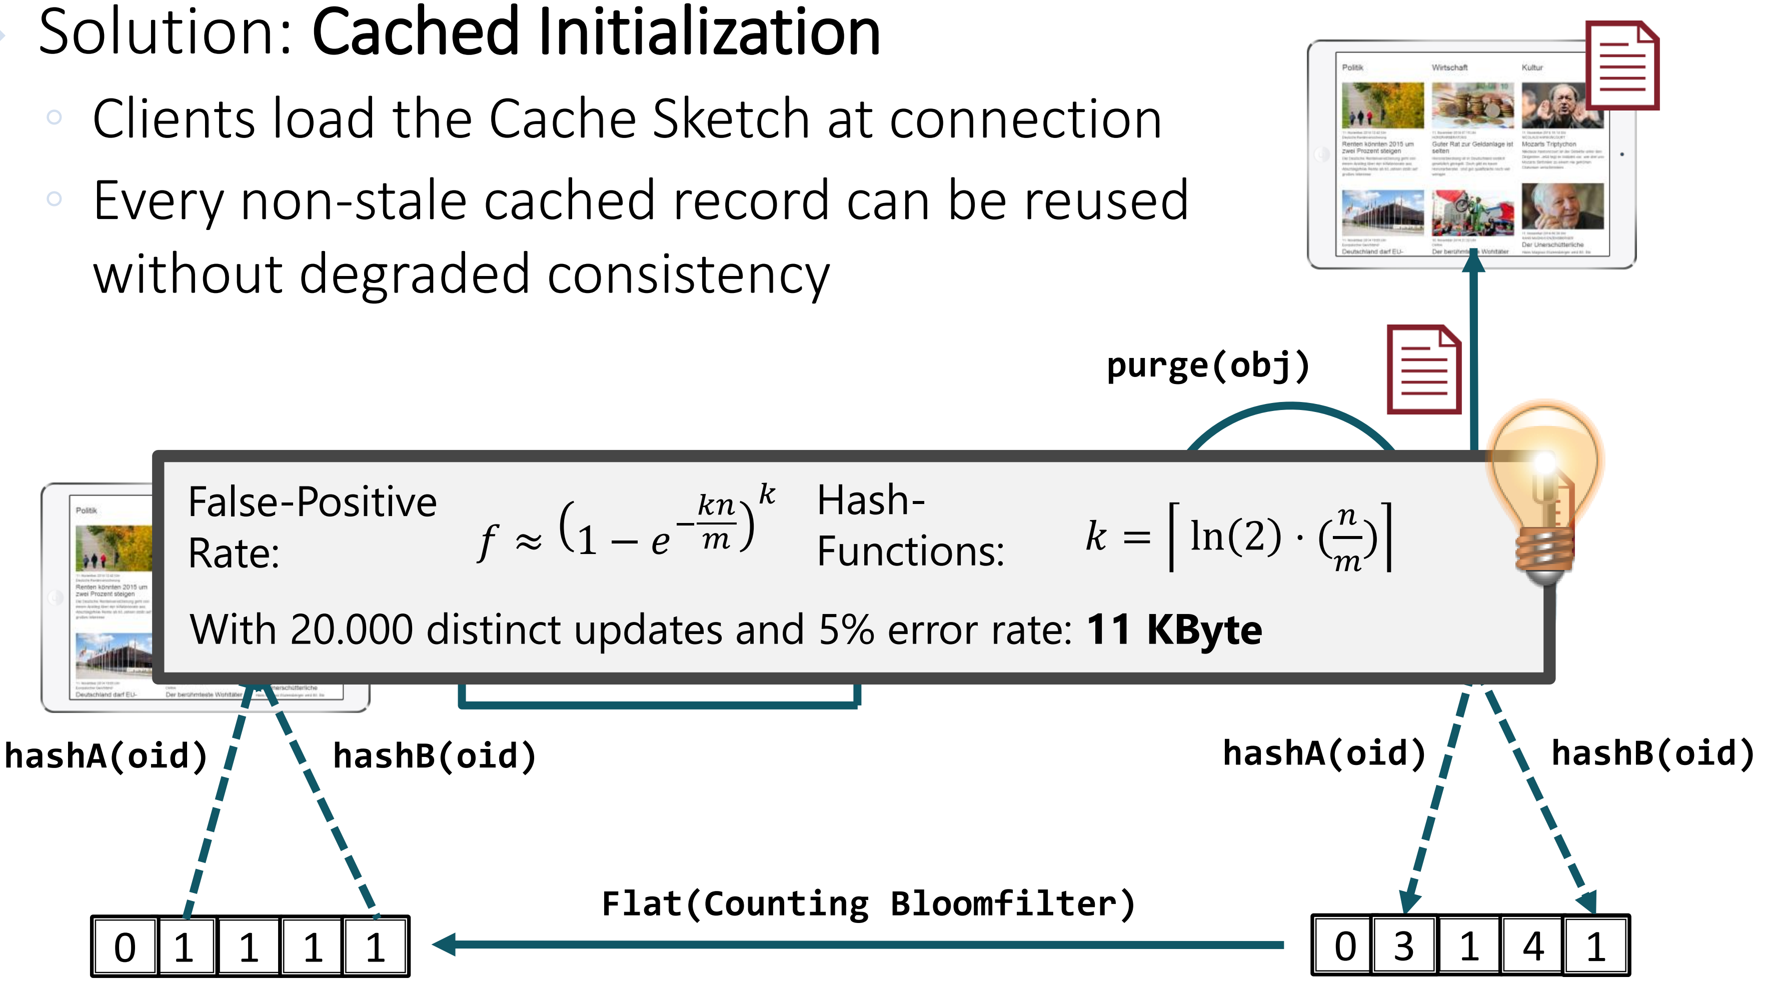

False-Positive Rate: $f \approx \left(1 - e^{-\frac{kn}{m}}\right)^k$ Hash-Functions: $k = \left\lceil \ln(2) \cdot \left(\frac{n}{m}\right) \right\rceil$

With 20.000 distinct updates and 5% error rate: **11 KByte**

2 Slow CRUD performance

- Solution: **Δ-Bounded Staleness**
  - Clients refresh the Cache Sketch so its age never exceeds Δ

  → *Consistency guarantee*: Δ-atomicity

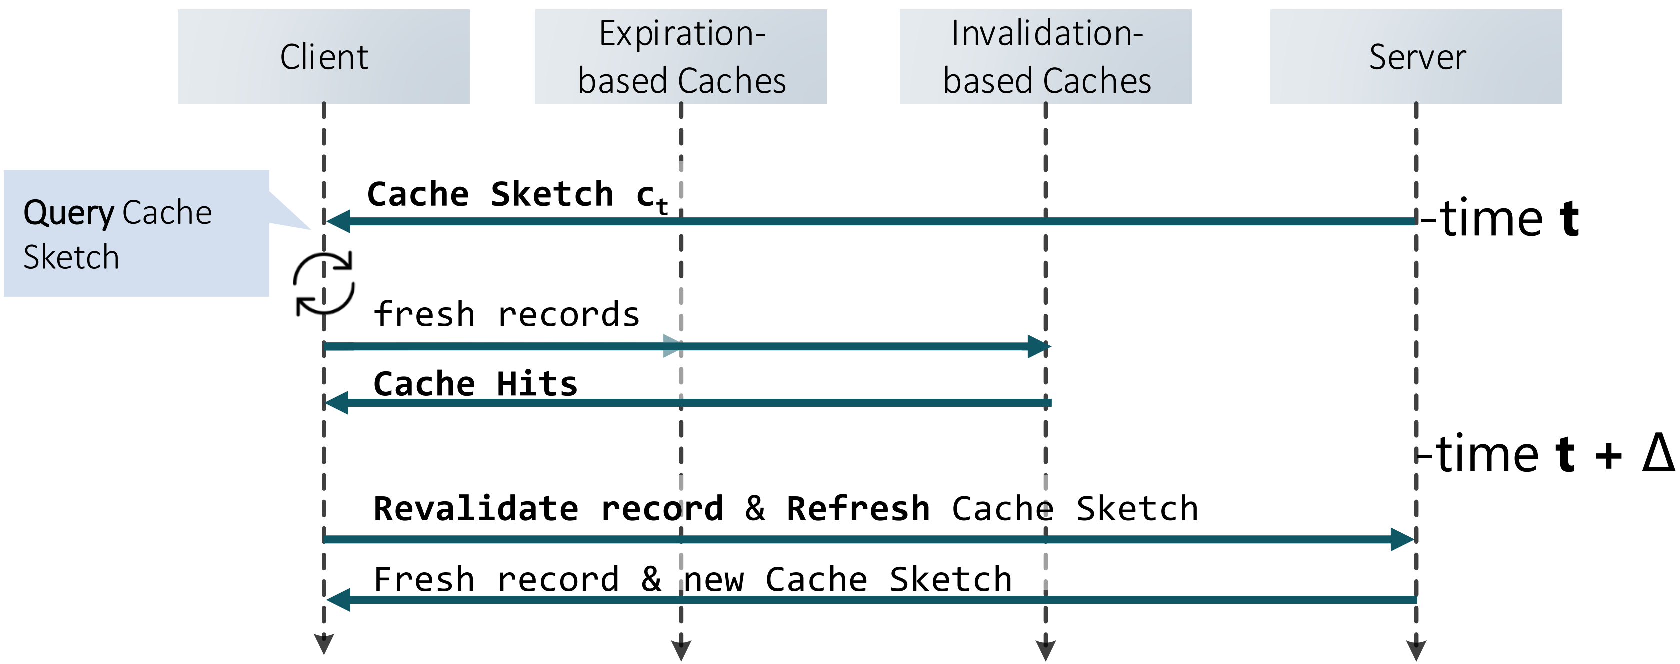

3 High Abort Rates in OCC

- Solution: **Conflict-Avoidant Optimistic Transactions**
  - Cache Sketch fetched with transaction begin
  - **Cached reads → Shorter transaction duration → less aborts**

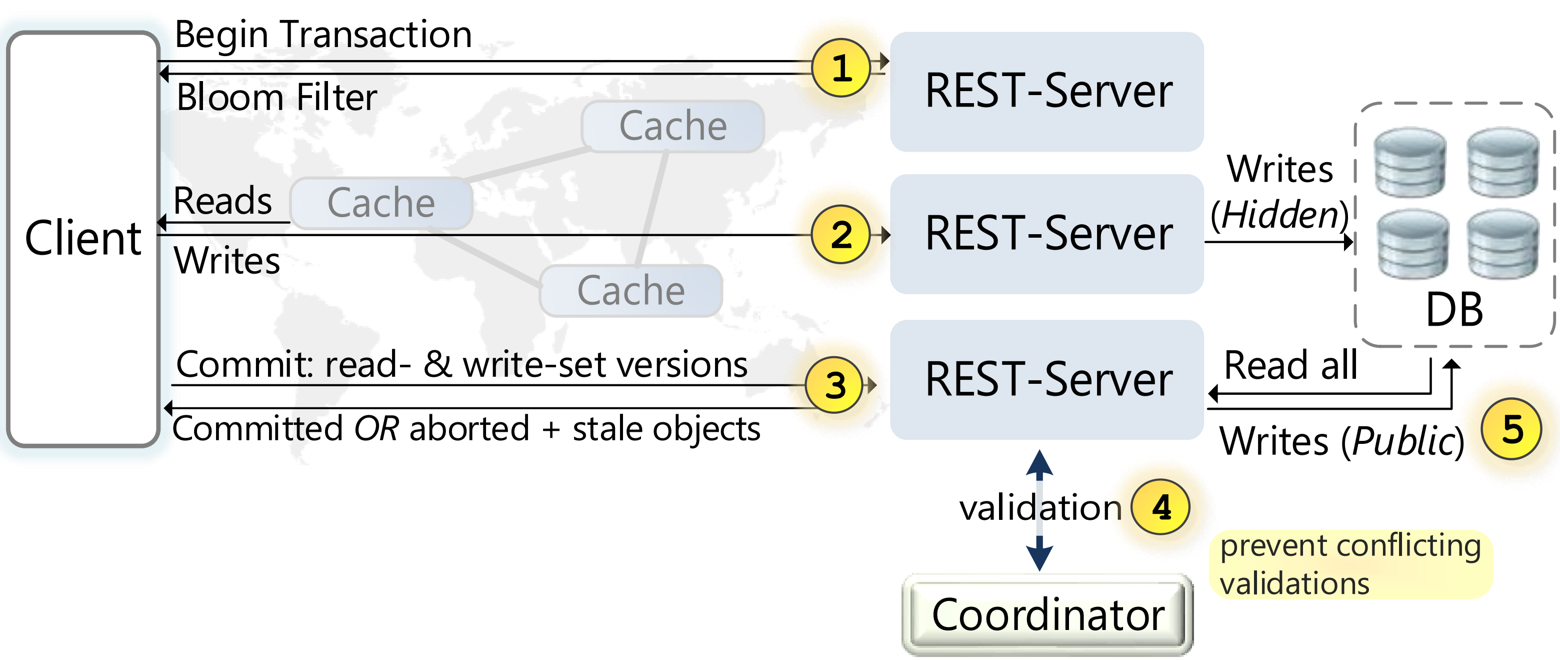

4 Costly Invalidations

- Solution: **Invalidation Minimization**
  - The server Cache Sketch tracks TTLs
  - Invalidation **only** necessary, if there are unexpired records

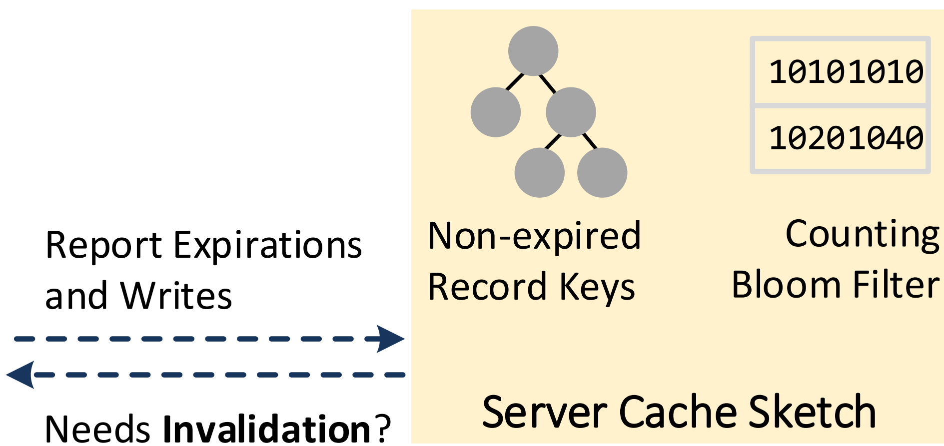

End-to-End Example

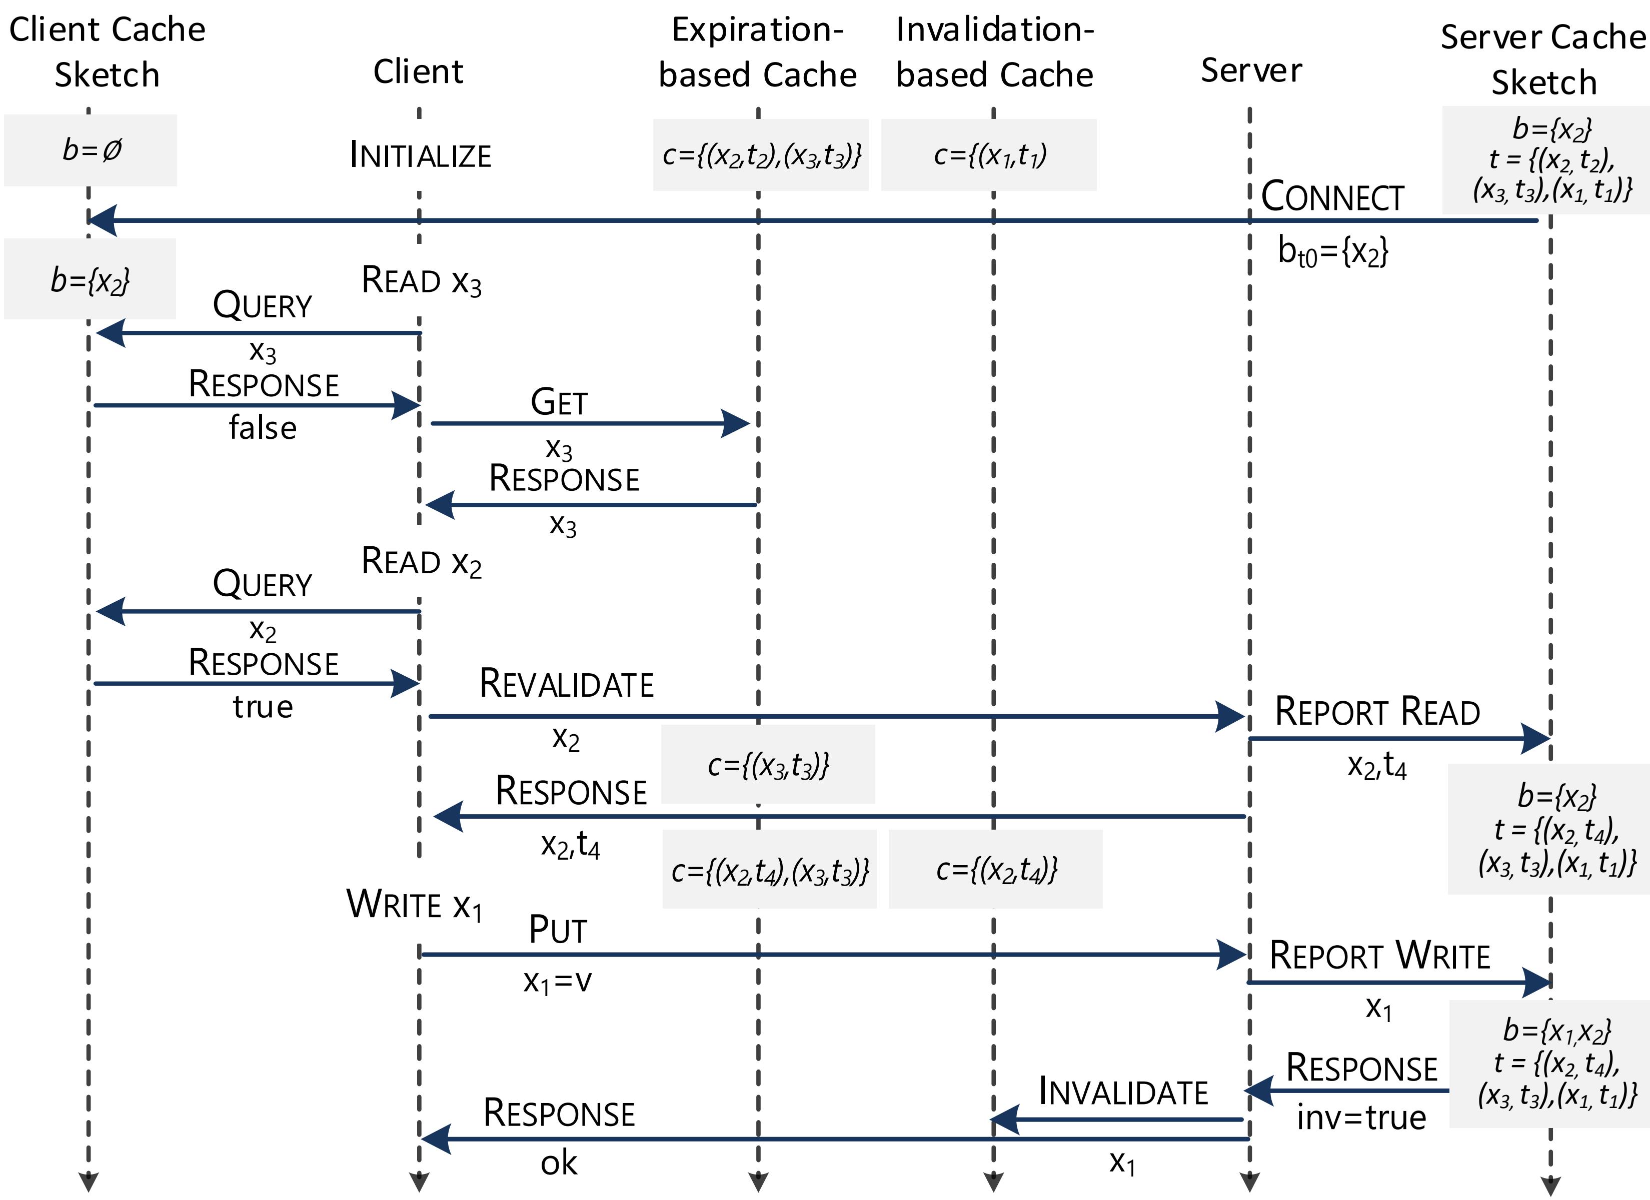

TTL Estimation

Determining the best TTL

- **Problem**: if TTL $\gg$ time to next write, then it is contained in Cache Sketch unnecessarily long
- **TTL Estimator**: finds „best“ TTL
- Trade-Off:

Shorter TTLs

- less invalidations
- less stale reads

Longer TTLs

- Higher cache-hit rates
- more invalidations

TTL Estimation

Determining the best TTL

Idea:

1. Estimate average time to next write $E[T_w]$ for each record
2. Weight $E[T_w]$ using the cache miss rate

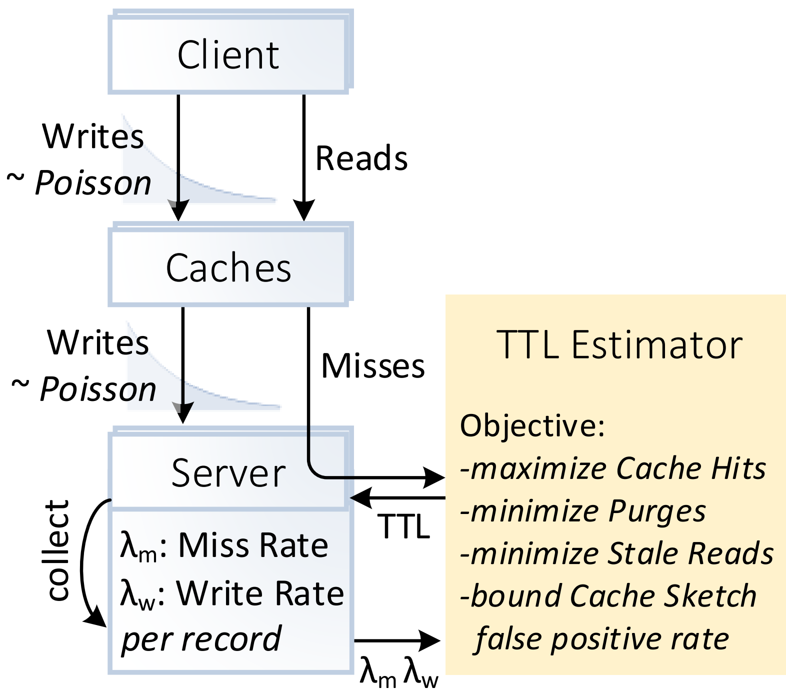

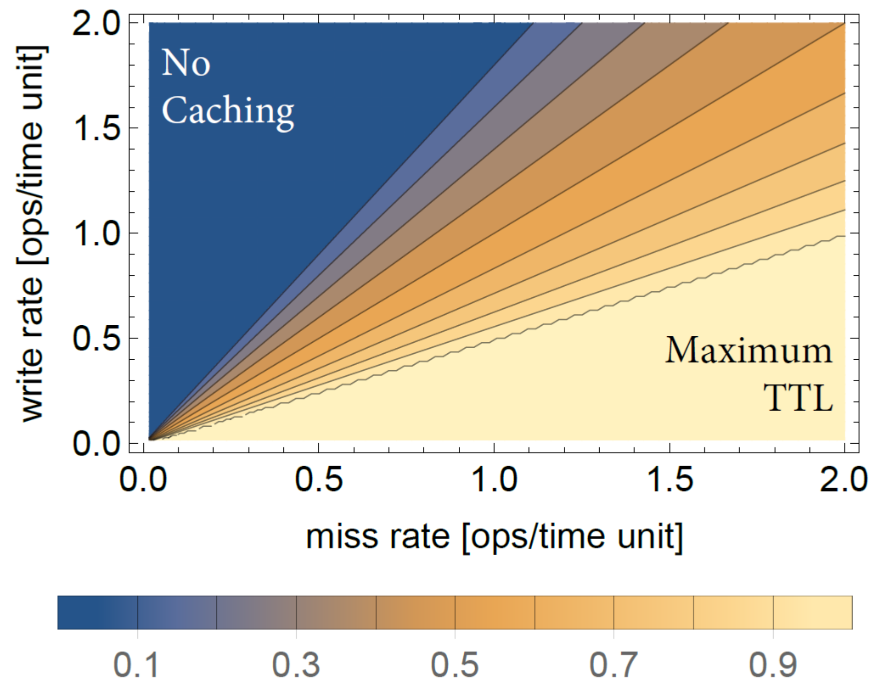

YCSB Monte Carlo Caching Simulator (YMCA)

- Goal: Analysis of arbitrary caching architectures using the standard **YCSB** benchmark
  - **Metrics to evaluate**: *Latency, Throughput, Cache Hits, Stale Reads, Invalidations*

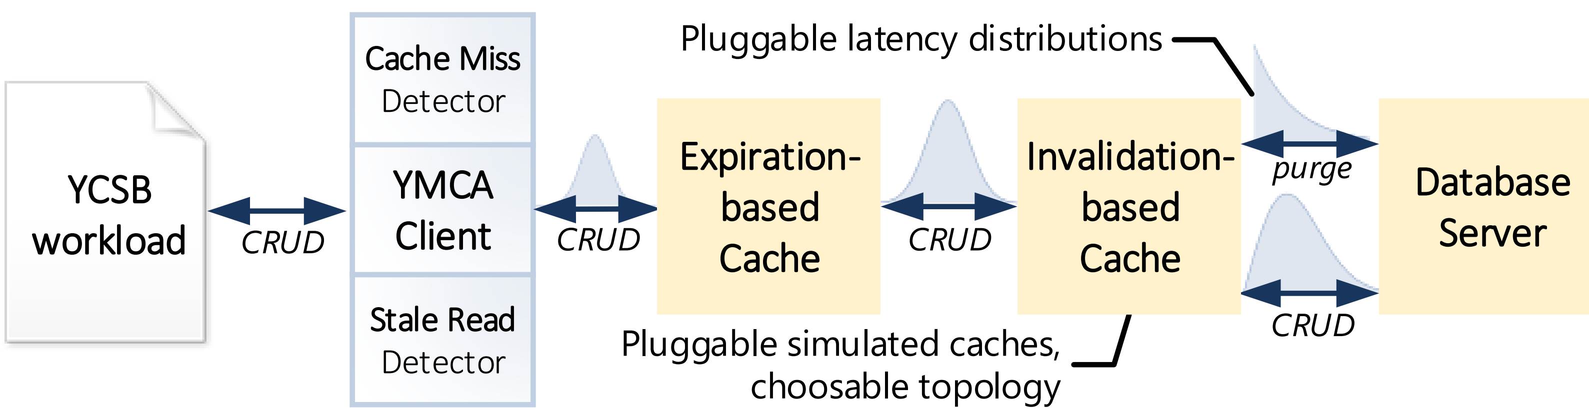

Results: Simulation & real-world

Setup:

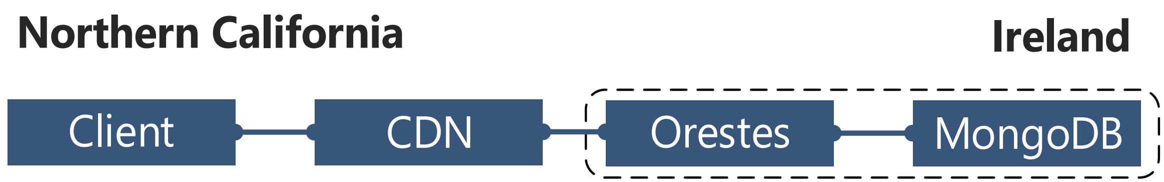

Page load times with **cached initialization** (YMCA):

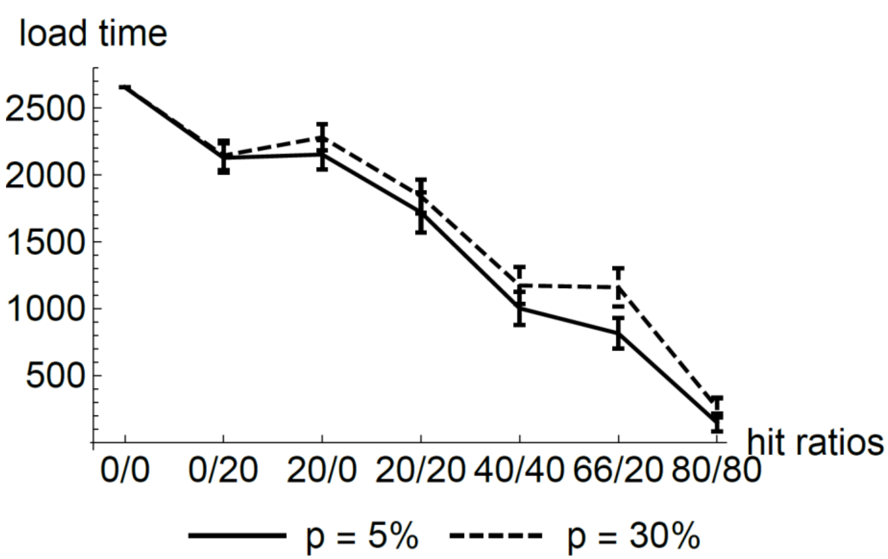

Average Latency for YCSB Workloads A and B (real):

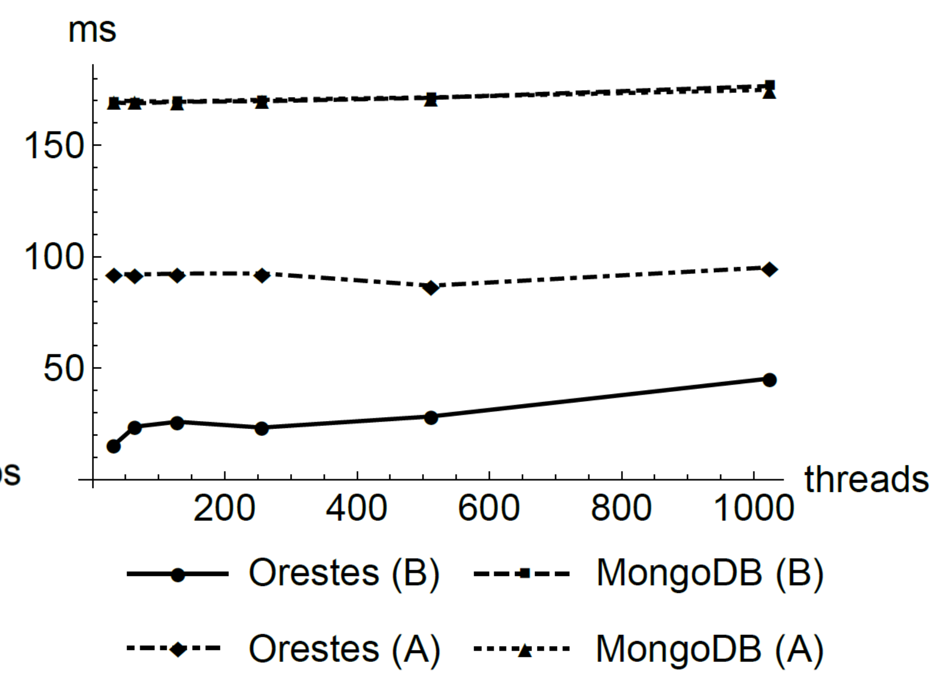

Results: Simulation & real-world

Setup:

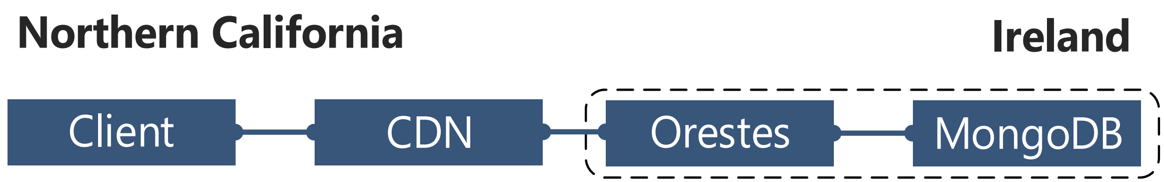

Page load times with **cached initialization** (YMCA):

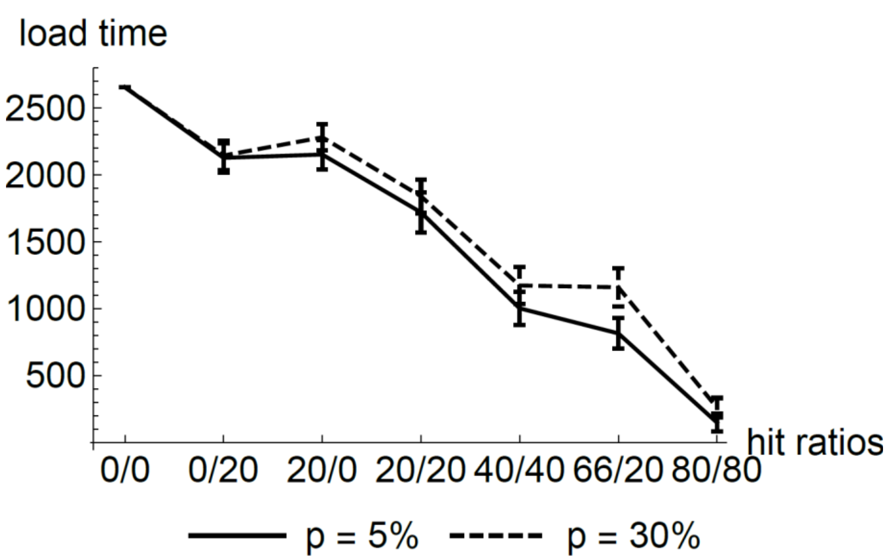

Average Throughput for YCSB Workloads A and B (real):

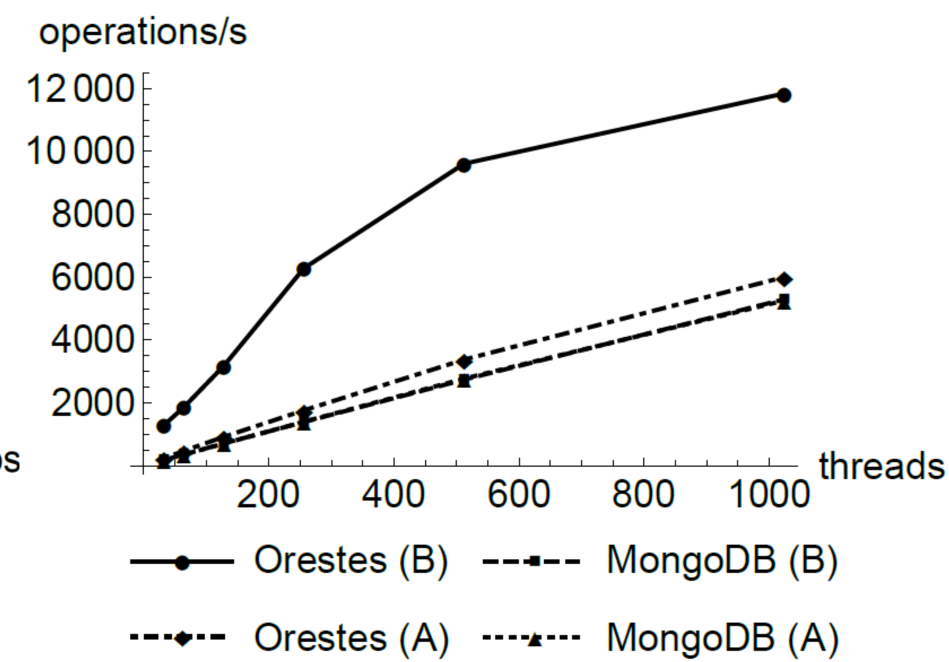

The Server Cache Sketch

Scalable Implementation

- Goal: Efficient *Generation of Cache Sketch* and *Invalidation Minimization*
- **Counting Bloom Filter** and **key** → **expiration** mapping

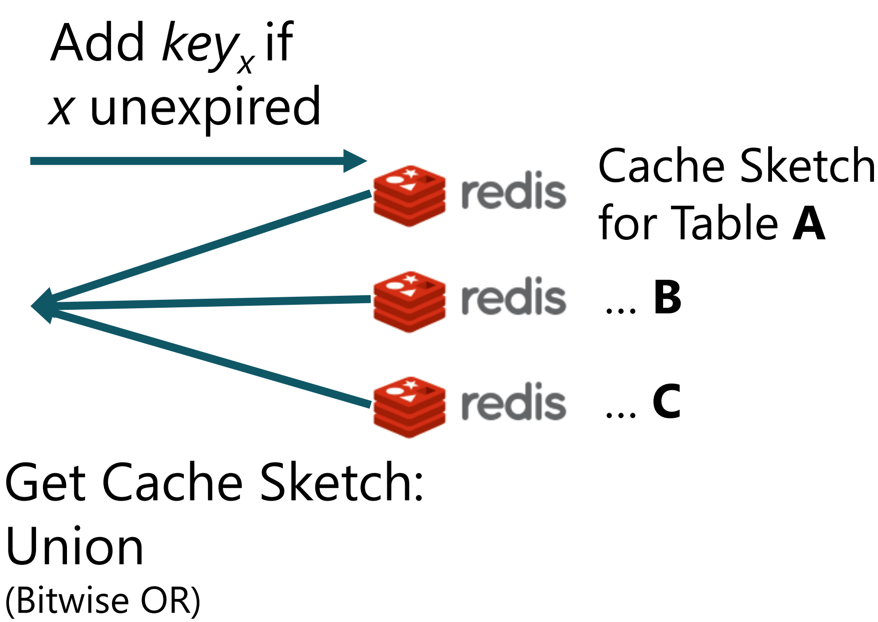

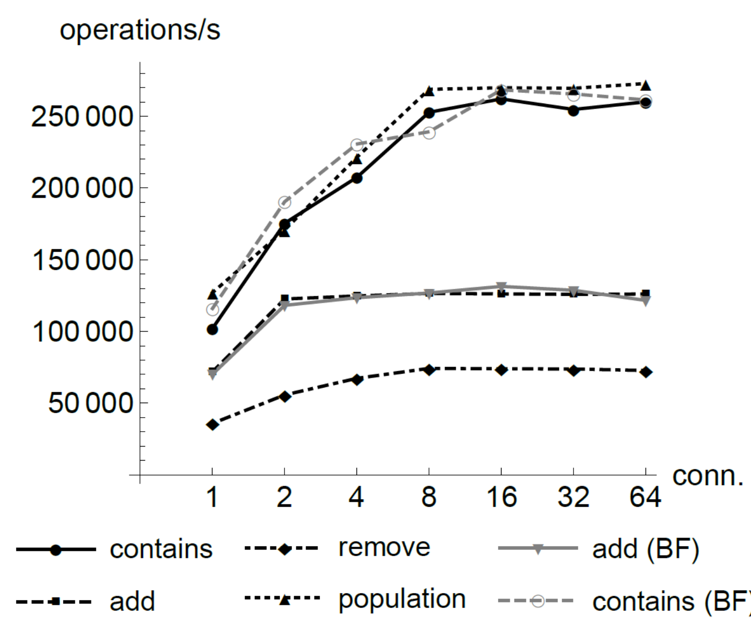

https://github.com/Baqend/Orestes-Bloomfilter

The Big Picture

Implementation in ORESTES

- Cache Sketch is part of ORESTES, a database-independent Backend-as-a-Service

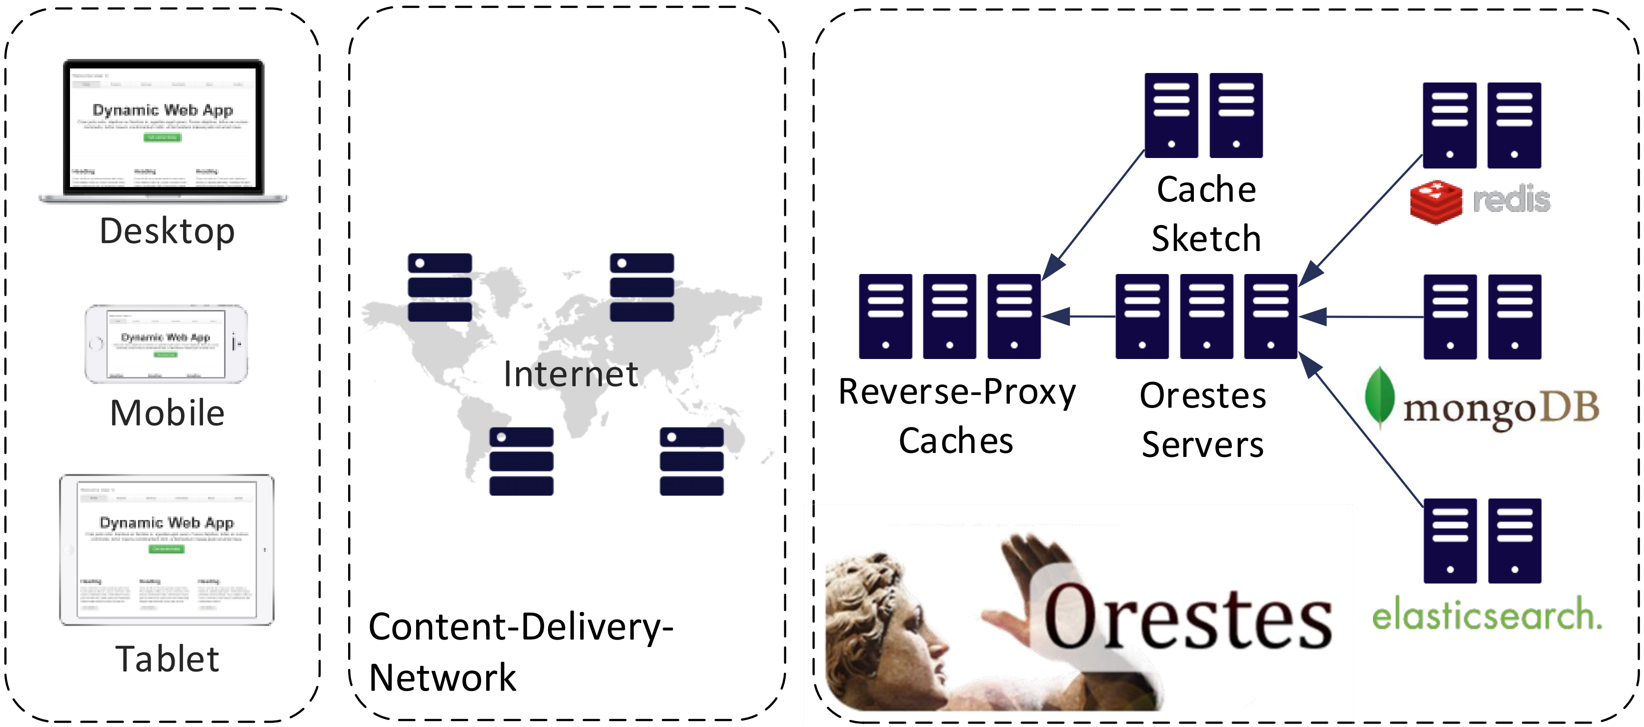

The Big Picture

Implementation in ORESTES

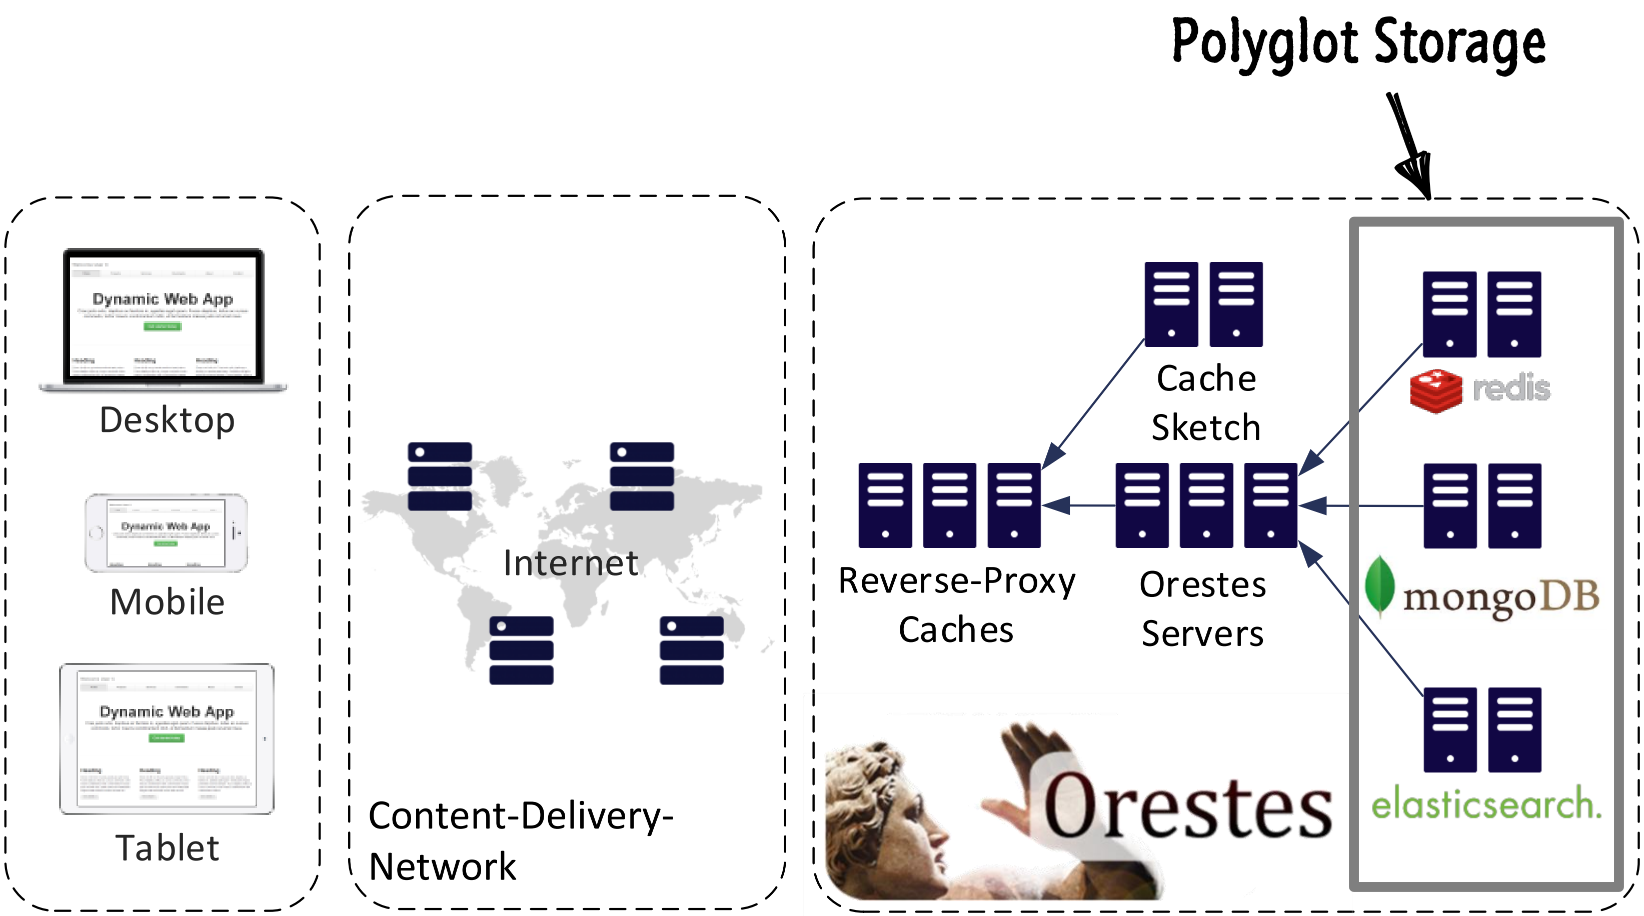

The Big Picture

Implementation in ORESTES

Database-as-a-Service Middleware:
Caching, Transactions, Schemas, Authorization, Multi-Tenancy

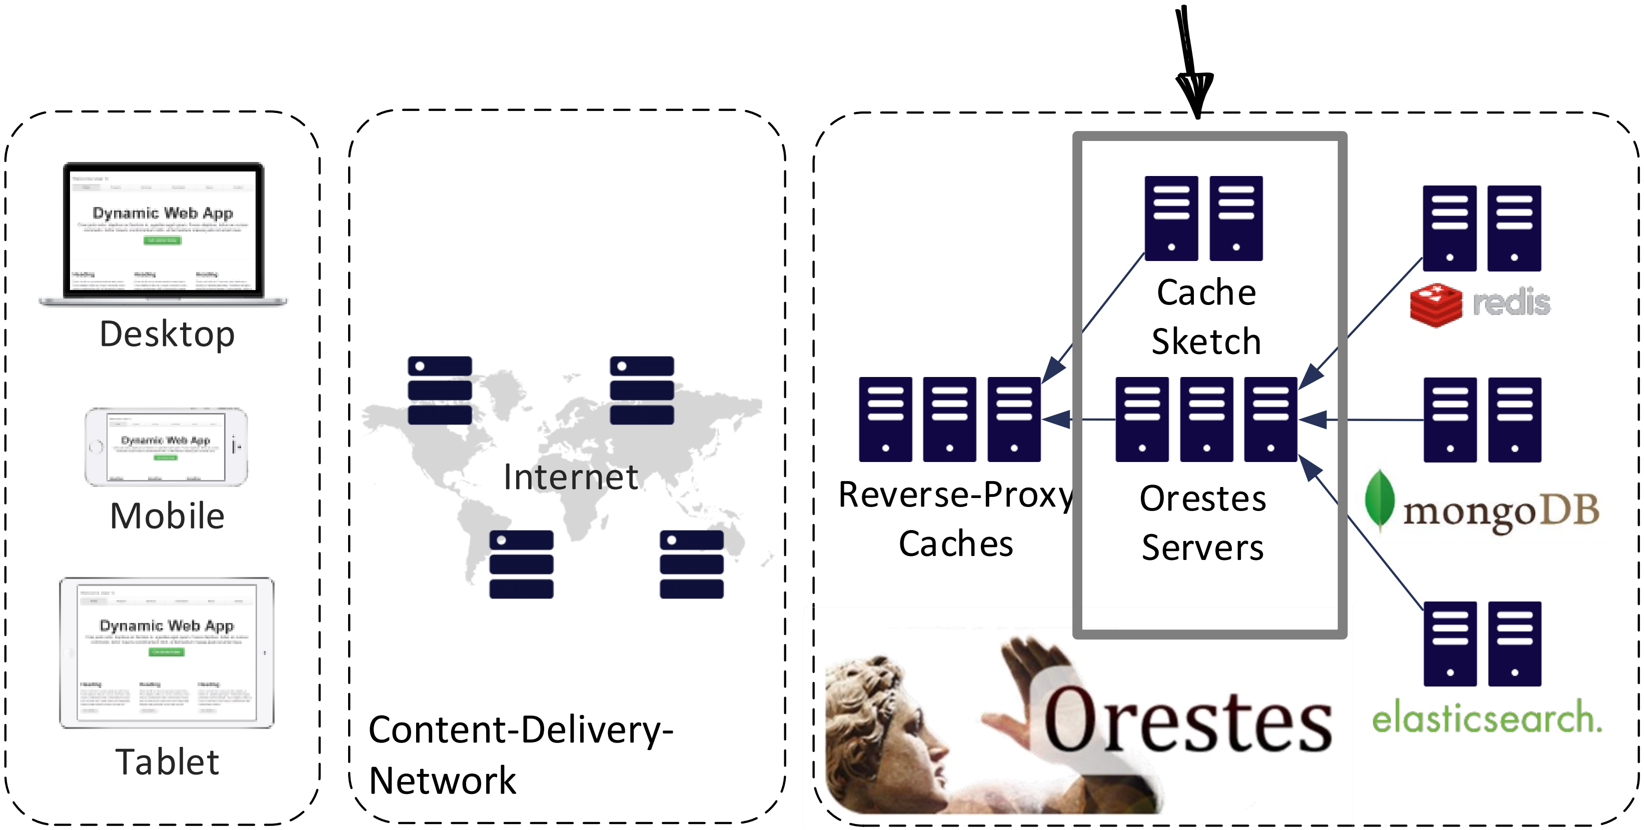

The Big Picture

Implementation in ORESTES

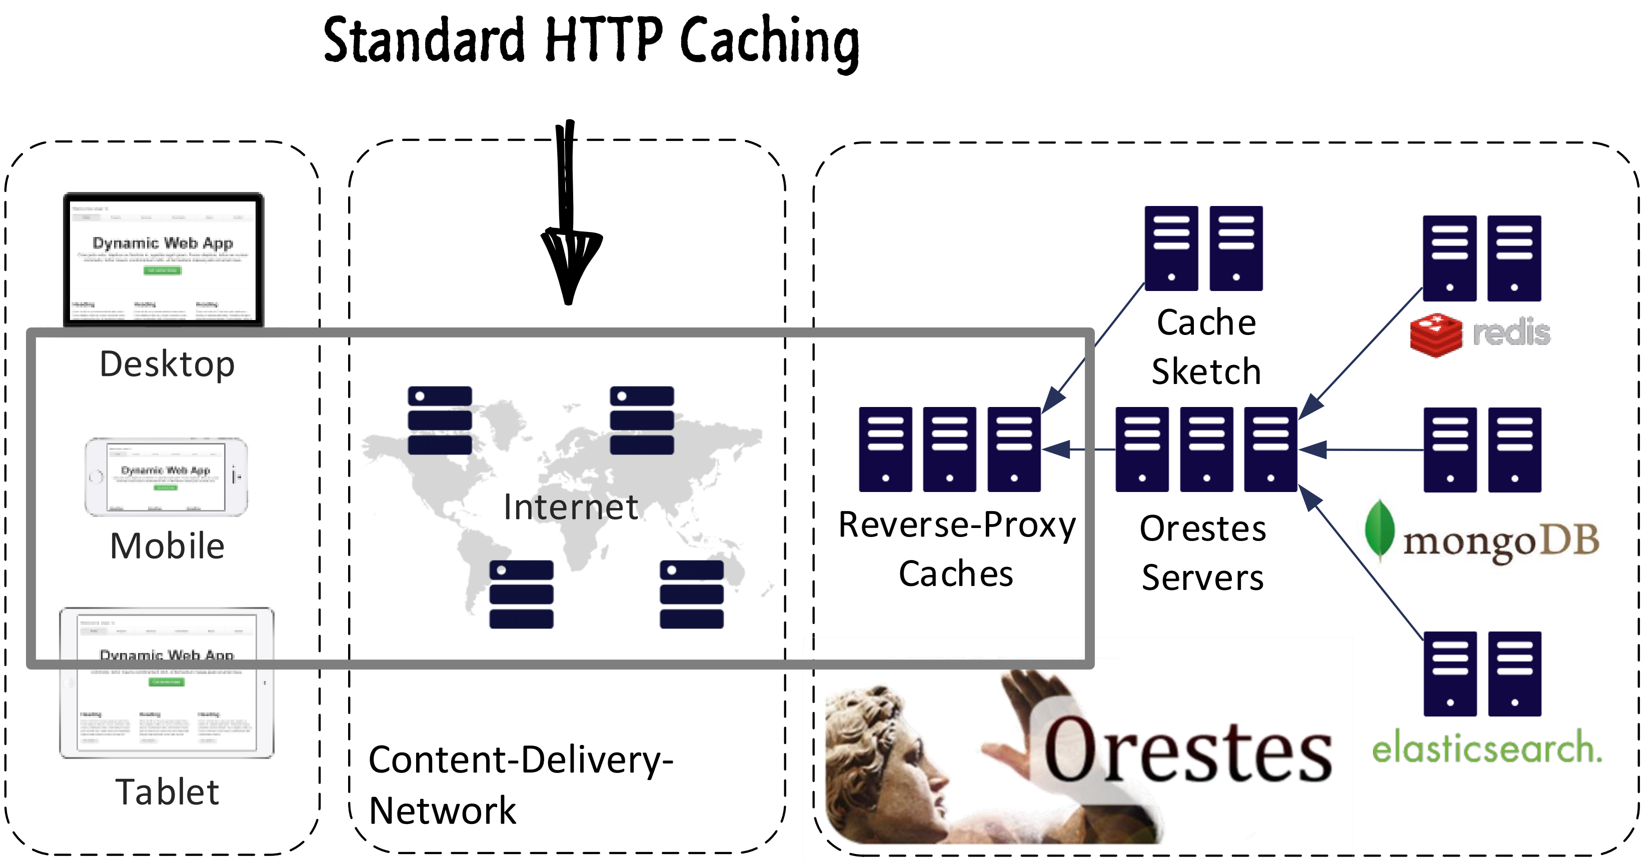

The Big Picture

Implementation in Orestes

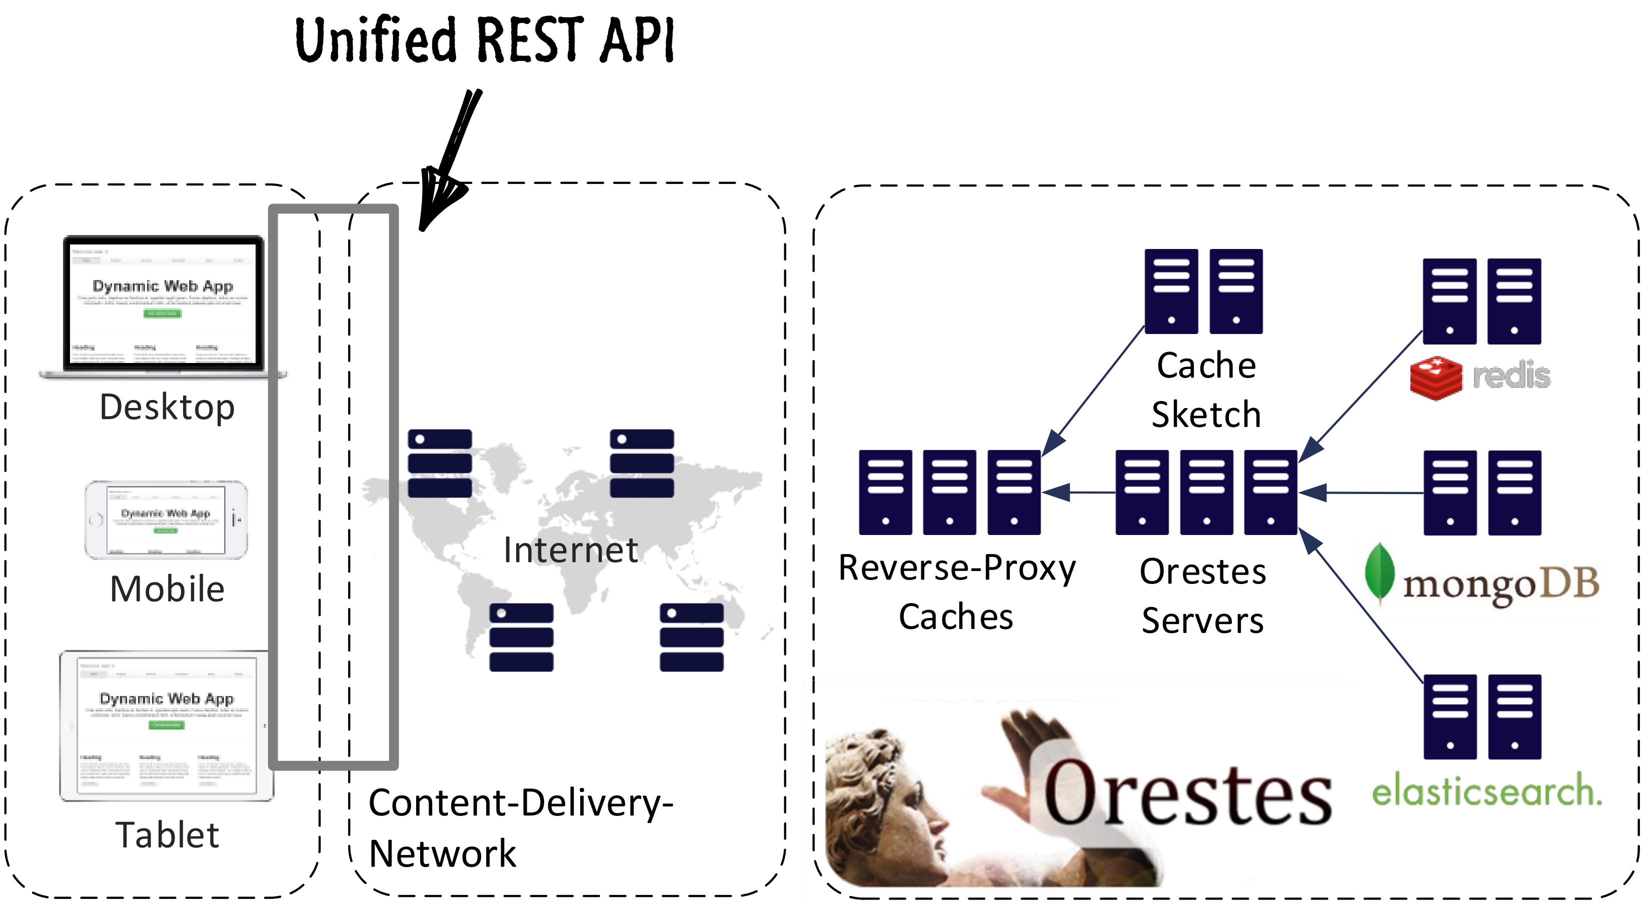

Future Work

Query-Result-Caching

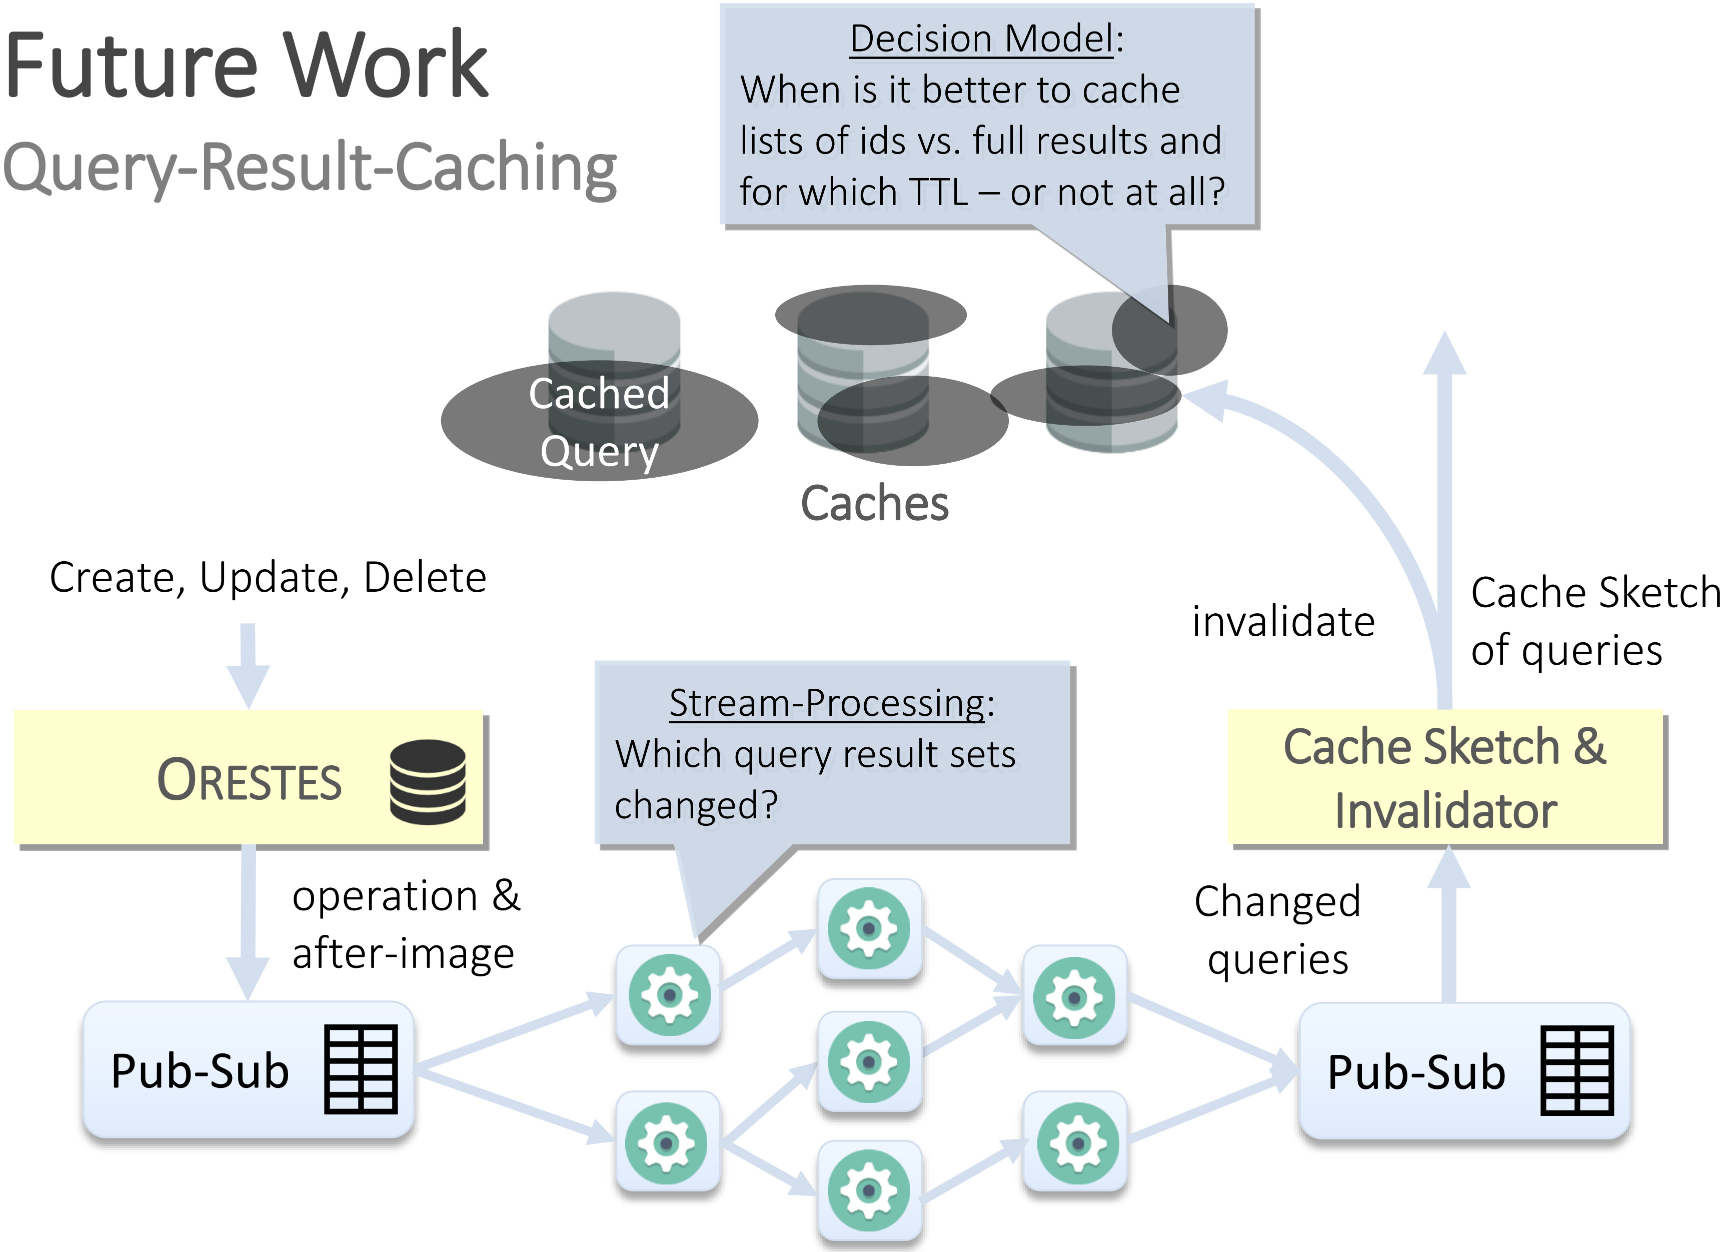


BaQend
Build faster Apps faster.

Page-Load Times

What impact does the Cache Sketch have?

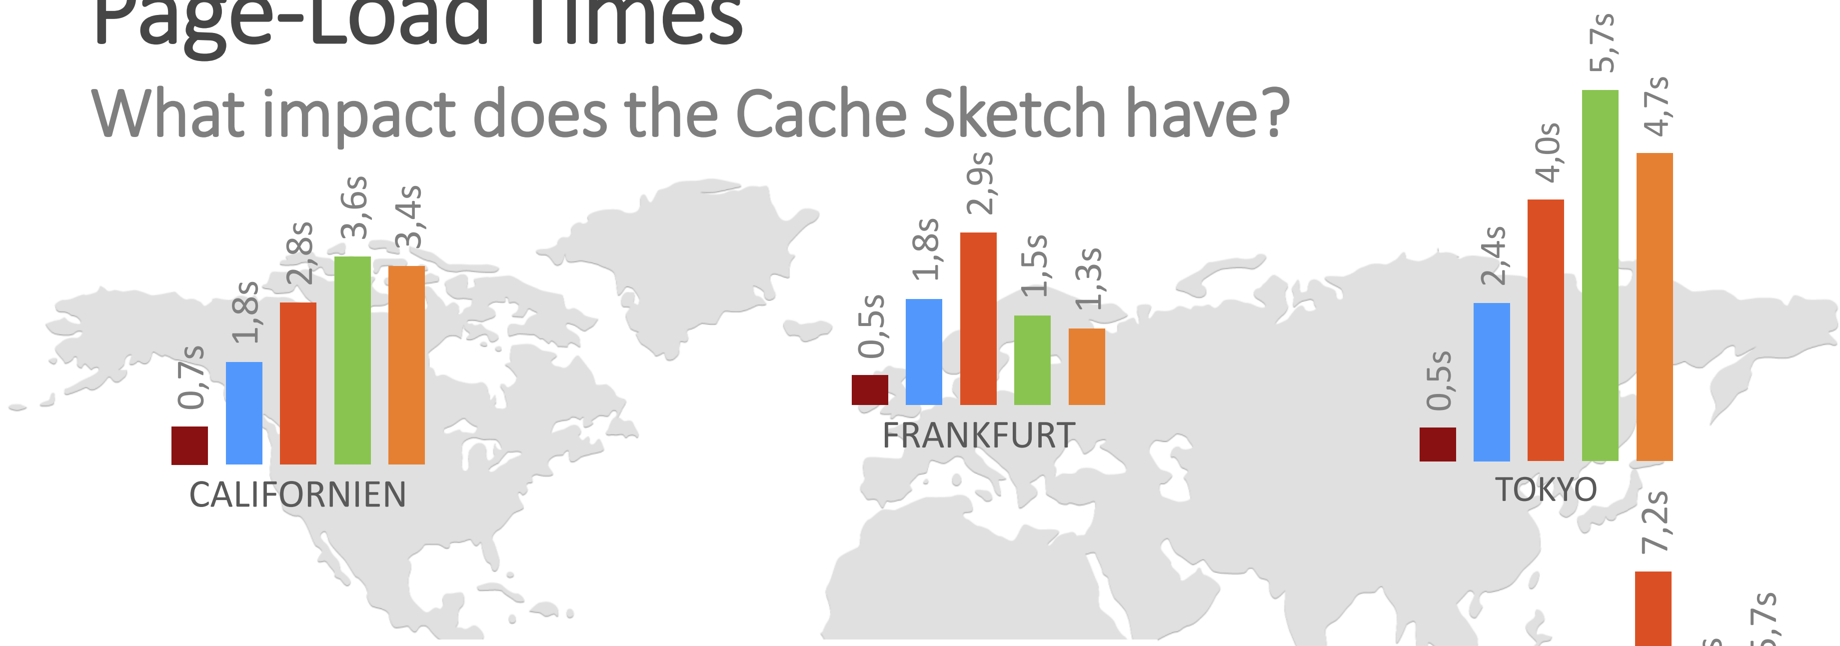

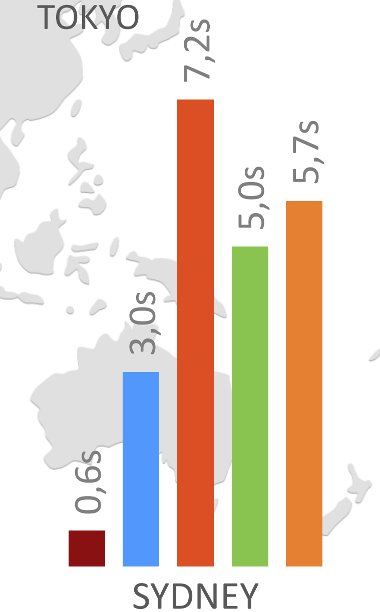

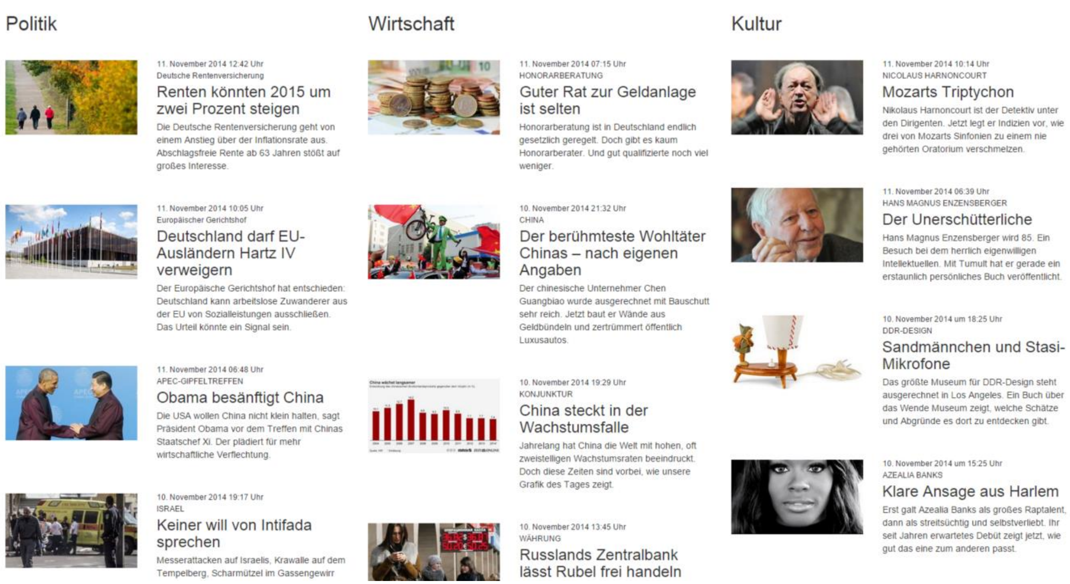

Politik

11. November 2014 12:42 Uhr
Deutsche Rentenversicherung
Renten könnten 2015 um zwei Prozent steigen
Die Deutsche Rentenversicherung geht von einem Anstieg über der Inflationsrate aus. Abschlagsfreie Rente ab 63 Jahren stößt auf großes Interesse.

11. November 2014 10:05 Uhr
Europäischer Gerichtshof
Deutschland darf EU-Ausländern Hartz IV verweigern
Der Europäische Gerichtshof hat entschieden: Deutschland kann arbeitslose Zuwanderer aus der EU von Sozialleistungen ausschließen. Das Urteil könnte ein Signal sein.

11. November 2014 06:48 Uhr
APEC-GIPFELTREFFEN
Obama besänftigt China
Die USA wollen China nicht klein halten, sagt Präsident Obama vor dem Treffen mit Chinas Staatschef Xi. Der plädiert für mehr wirtschaftliche Verflechtung.

10. November 2014 19:17 Uhr
ISRAEL
Keiner will von Intifada sprechen
Messerattacken auf Israelis, Krawalle auf dem Tempelberg, Scharmützel im Gassengewirr

Wirtschaft

11. November 2014 07:15 Uhr
HONORARBERATUNG
Guter Rat zur Geldanlage ist selten
Honorarberatung ist in Deutschland endlich gesetzlich geregelt. Doch gibt es kaum Honorarberater. Und gut qualifizierte noch viel weniger.

10. November 2014 21:32 Uhr
CHINA
Der berühmteste Wohltäter Chinas – nach eigenen Angaben
Der chinesische Unternehmer Chen Guangbiao wurde ausgerechnet mit Bauschutt sehr reich. Jetzt baut er Wände aus Geldbündeln und zertrümmert öffentlich Luxusautos.

10. November 2014 19:29 Uhr
KONJUNKTUR
China steckt in der Wachstumsfalle
Jahrelang hat China die Welt mit hohen, oft zweistelligen Wachstumsraten beeindruckt. Doch diese Zeiten sind vorbei, wie unsere Grafik des Tages zeigt.

10. November 2014 13:45 Uhr
WÄHRUNG
Russlands Zentralbank lässt Rubel frei handeln

Kultur

11. November 2014 10:14 Uhr
NIKOLAUS HARNONCOURT
Mozarts Triptychon
Nikolaus Harnoncourt ist der Detektiv unter den Dirigenten. Jetzt legt er Indizien vor, wie drei von Mozarts Sinfonien zu einem nie gehörten Oratorium verschmelzen.

11. November 2014 06:39 Uhr
HANS MAGNUS ENZENSBERGER
Der Unerschütterliche
Hans Magnus Enzensberger wird 85. Ein Besuch bei dem herrlich eigenwilligen Intellektuellen. Mit Tumult hat er gerade ein erstaunlich persönliches Buch veröffentlicht.

10. November 2014 um 18:25 Uhr
DDR-DESIGN
Sandmännchen und Stasi-Mikrofone
Das größte Museum für DDR-Design steht ausgerechnet in Los Angeles. Ein Buch über das Wende Museum zeigt, welche Schätze und Abgründe es dort zu entdecken gibt.

10. November 2014 um 15:25 Uhr
AZEALIA BANKS
Klare Ansage aus Harlem
Erst galt Azealia Banks als großes Raptalent, dann als streitsüchtig und selbstverliebt. Ihr seit Jahren erwartetes Debüt zeigt jetzt, wie gut das eine zum anderen passt.

Summary

- **Cache Sketch**: dual approach to web caching for database services
  - Consistent (Δ-atomic) *expiration-based* caching
  - *Invalidation-based* caching with minimal purges

Cached Initialization	Δ-Bounded Staleness	Conflict-Avoidant Transactions	Invalidation Minimization

- **Keys Ideas:**
  - Maintain **Bloom filter** of potentially stale objects
  - Let clients handle cache coherence through **revalidations** when an object is contained in the filter
  - Estimate the best **TTL** based on access statistics

Thank you

gessert@informatik.uni-hamburg.de
Orestes.info
Baqend.com