

# Simba: Tunable End-to-End Data Consistency for Mobile Apps

Dorian Perkins<sup>\*†</sup>, Nitin Agrawal<sup>†</sup>, Akshat Aranya<sup>†</sup>, Curtis Yu<sup>\*</sup>, Younghwan Go<sup>\*†</sup>,

Harsha V. Madhyastha<sup>‡</sup>, and Cristian Ungureanu<sup>‡</sup>

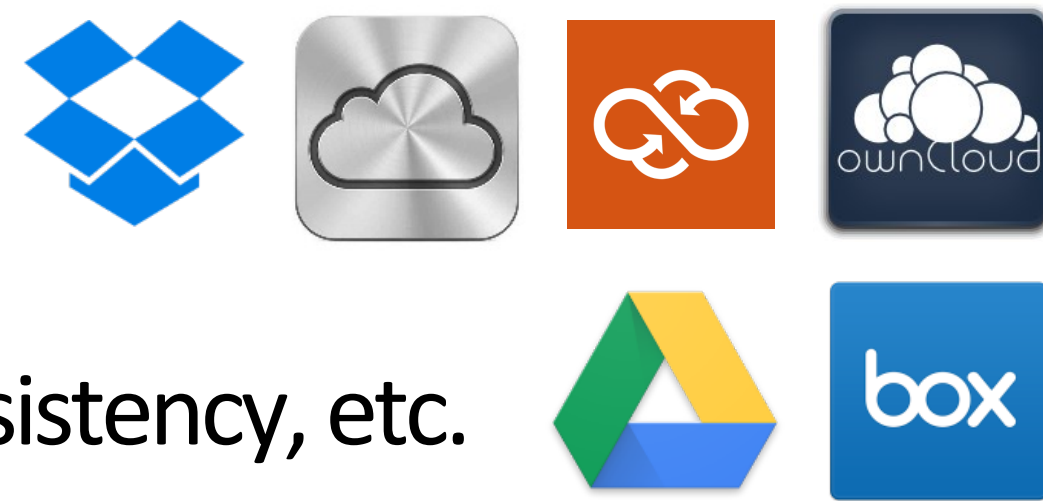
NEC Labs America<sup>†</sup>, UC Riverside<sup>\*</sup>, KAIST<sup>\*</sup>, U Michigan<sup>‡</sup>



Research Session 2:  
4/22 @ 12 PM

## Motivation

- Wide-spread use of *data-centric* mobile apps
- Data consistency is a primary requirement
  - App devs tasked with ensuring consistency of user's data
- Current solutions are *inflexible*
- Rolling own service is **difficult**
  - Failures, Conflicts, Connectivity, Consistency, etc.



## Difficult to Provide Reliability & Consistency!

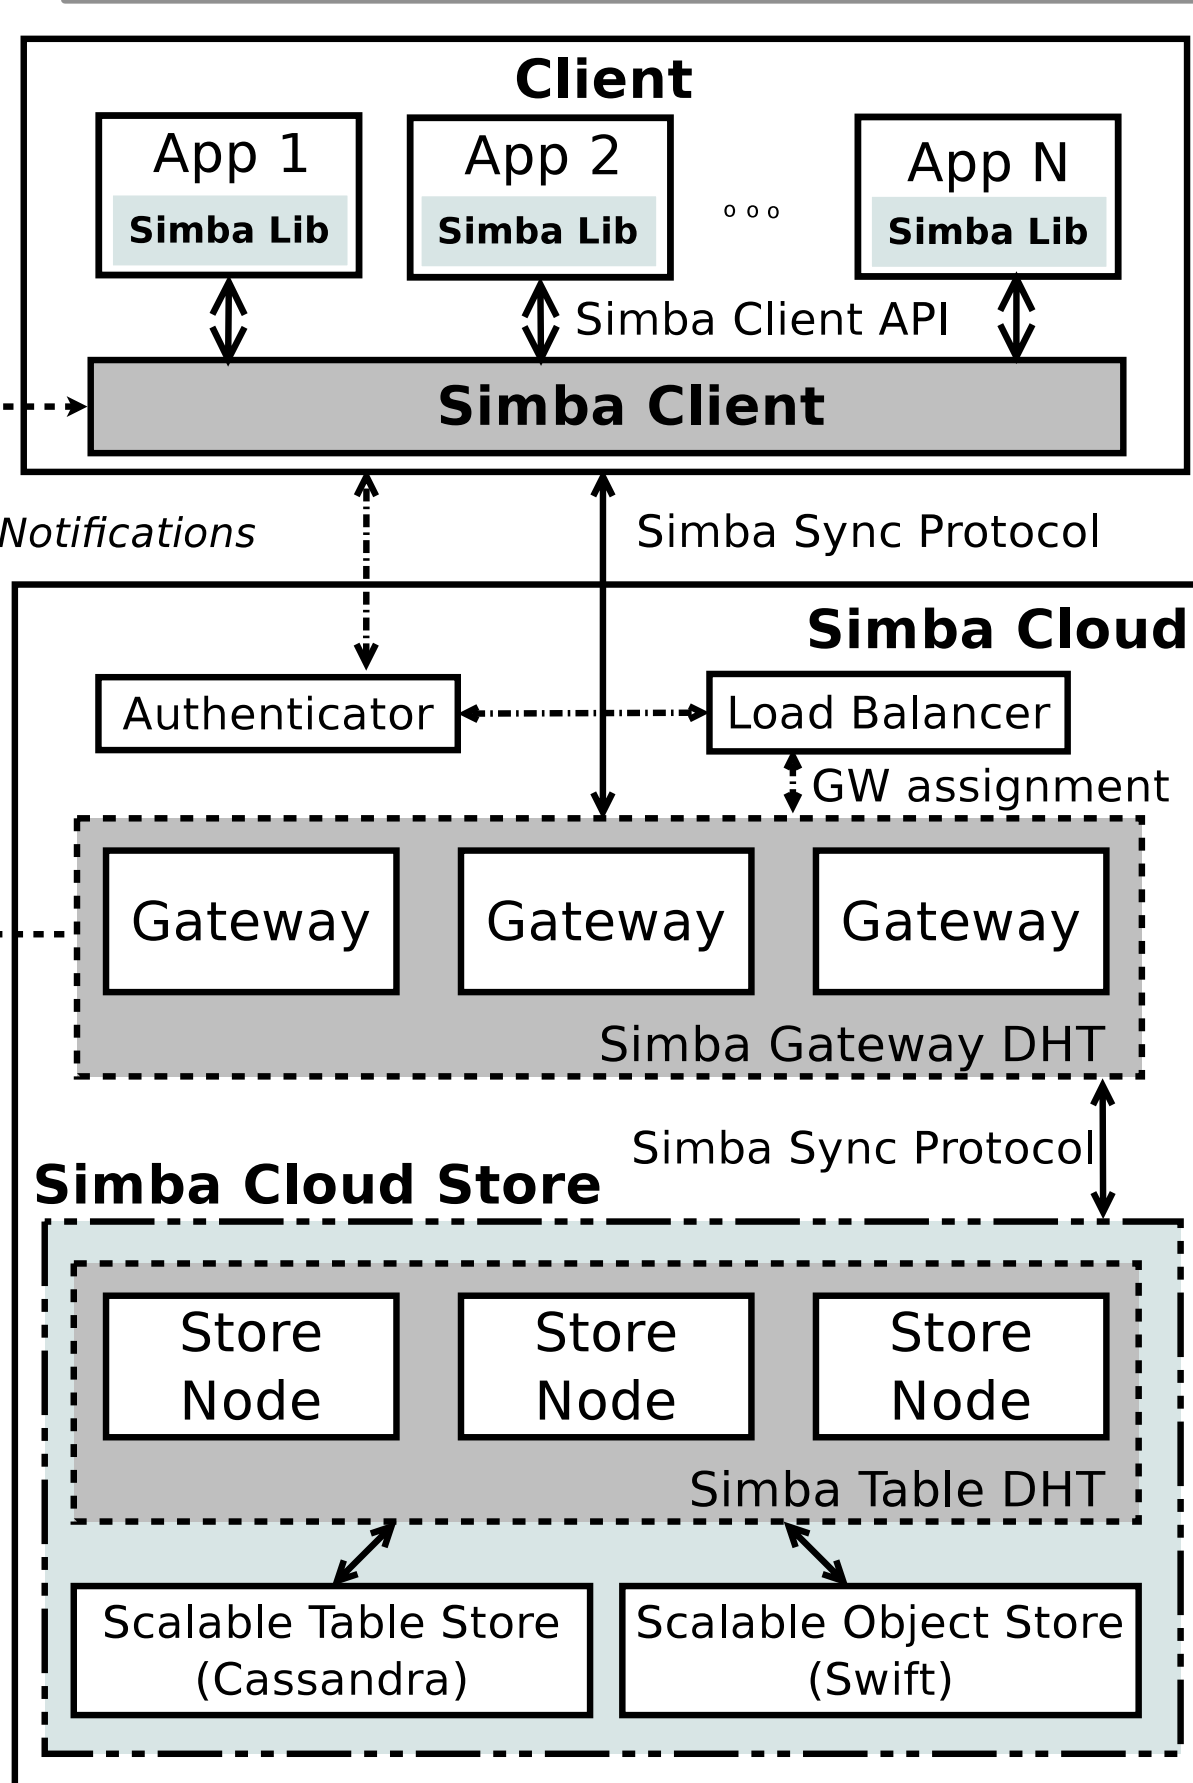
- Need to transparently/efficiently handle data sync
- Need to provide unified tabular + object data model
- Need to enable useful consistency semantics

## Simba: Cloud Infrastructure for Mobile Apps

## Consistency Study

We studied **23 popular mobile apps** and found that **half** of them exhibit **undesirable behavior!**

## Simba Design



- ### Key Features
- Simple, high-level programming abstractions
  - Transparent handling of data sync & failures
  - Atomicity across tabular & object data
  - End-to-end tunable consistency
  - Scalable architecture

Simba Client [FAST 2015]

## Study Findings

- Diverse consistency requirements
- Sync semantics often oblivious to consistency
- Limited offline support
- Inadequate error propagation
- Atomicity violations of inter-dependent data

## Simba Table

### Logical Abstraction

Name	Quality	Photo	Thumbnail
Snoopy	High	snoopy.jpg	t_snoopy.jpg
Snowy	Med	snowy.jpg	t_snowy.jpg

Tabular                      Object

### Physical Layout

#### Table Store

_rowID	_rowVersion	Name	Quality	Photo	Thumbnail
f12e09bc	780	Snoopy	High	[ab1fd, 1fc2e]	[42e11]
f12e09fg	781	Snowy	Med	[x561a, 3f02a]	[42e13]

Metadata                      Tabular                      Object

#### Object Store


ab1fd	42e11	42e13
1fc2e	x561a	3f02a

### Data Sync Abstraction

- Unified tabular + object rows
- Row-level atomicity
- Per-table consistency scheme
  - Strong, Causal, or Eventual
- Offline support
- Conflict detection/resolution

## Performance


### Efficient Syncing via Change-sets



**Change cache** enables sync of only modified row data (e.g., update of a single 64 KiB chunk in a 1 MiB object)

## Scalability

### 16 Gateways and 16 Stores, 500 ops/sec



Simba Cloud scales well with increasing tables and clients

Simba Source Code: <https://github.com/SimbaService/Simba>

Project Homepage: <http://tinyurl.com/SimbaService>