



EPL646 – Advanced Topics in Databases

Lecture 13

NoSQL Databases: CouchDB II (Semi-structured JSON DB)

**Chapter 20: Abiteboul et. Al.
+ <http://guide.couchdb.org/>**

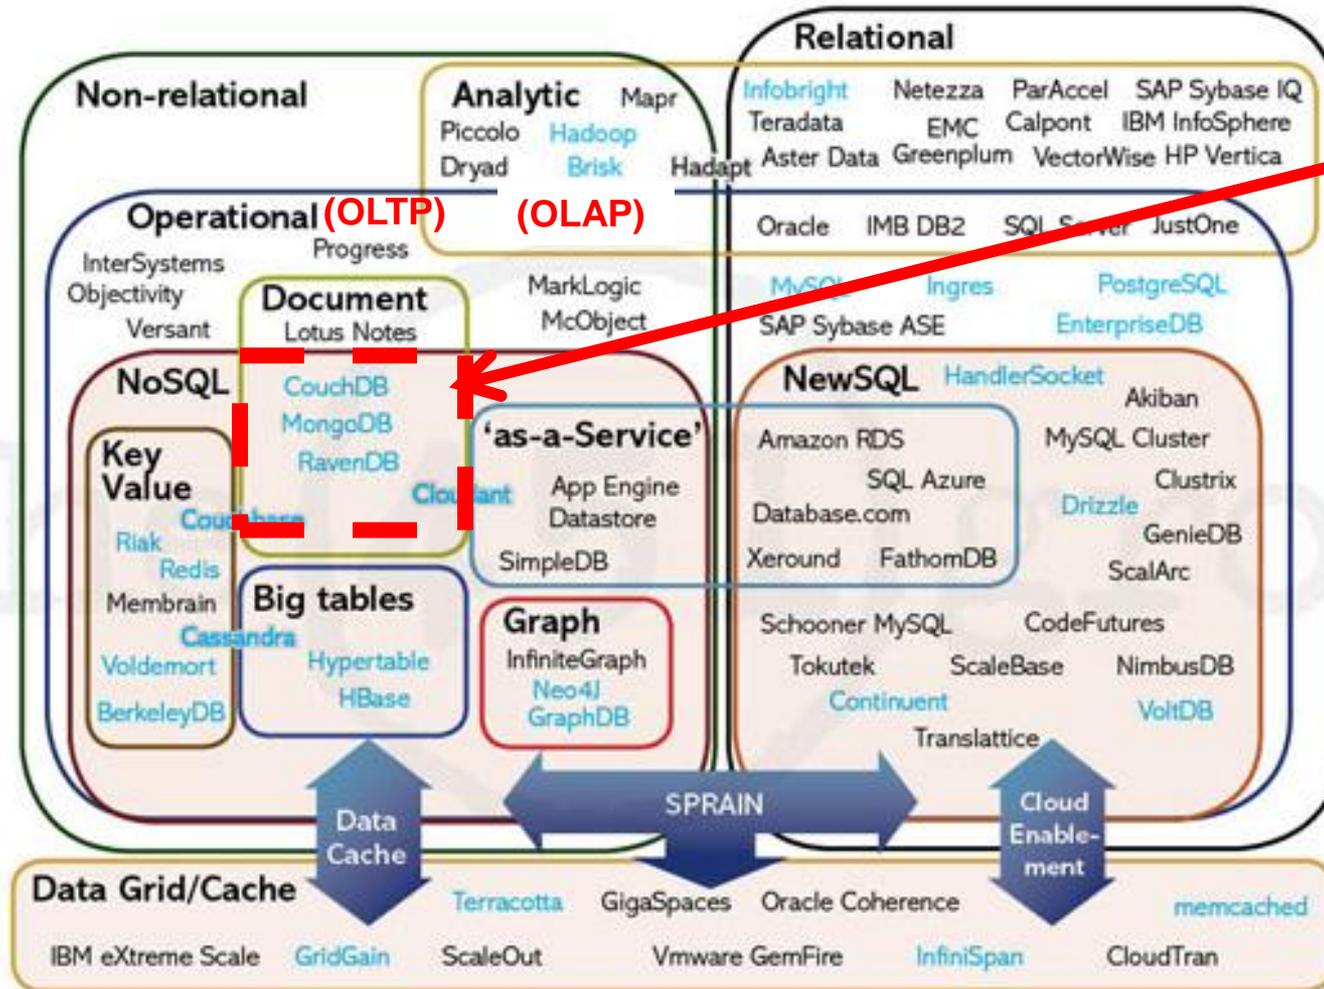
Demetris Zeinalipour

<http://www.cs.ucy.ac.cy/~dzeina/courses/epl646>

EPL646: Part B



Distributed/Web/Cloud DBs/Dstores



Lecture Focus

Venn Diagram by 451 group

<http://xeround.com/blog/2011/04/newsq-cloud-database-as-a-service>

Lecture Outline

(Introduction to Semi-structured Data)

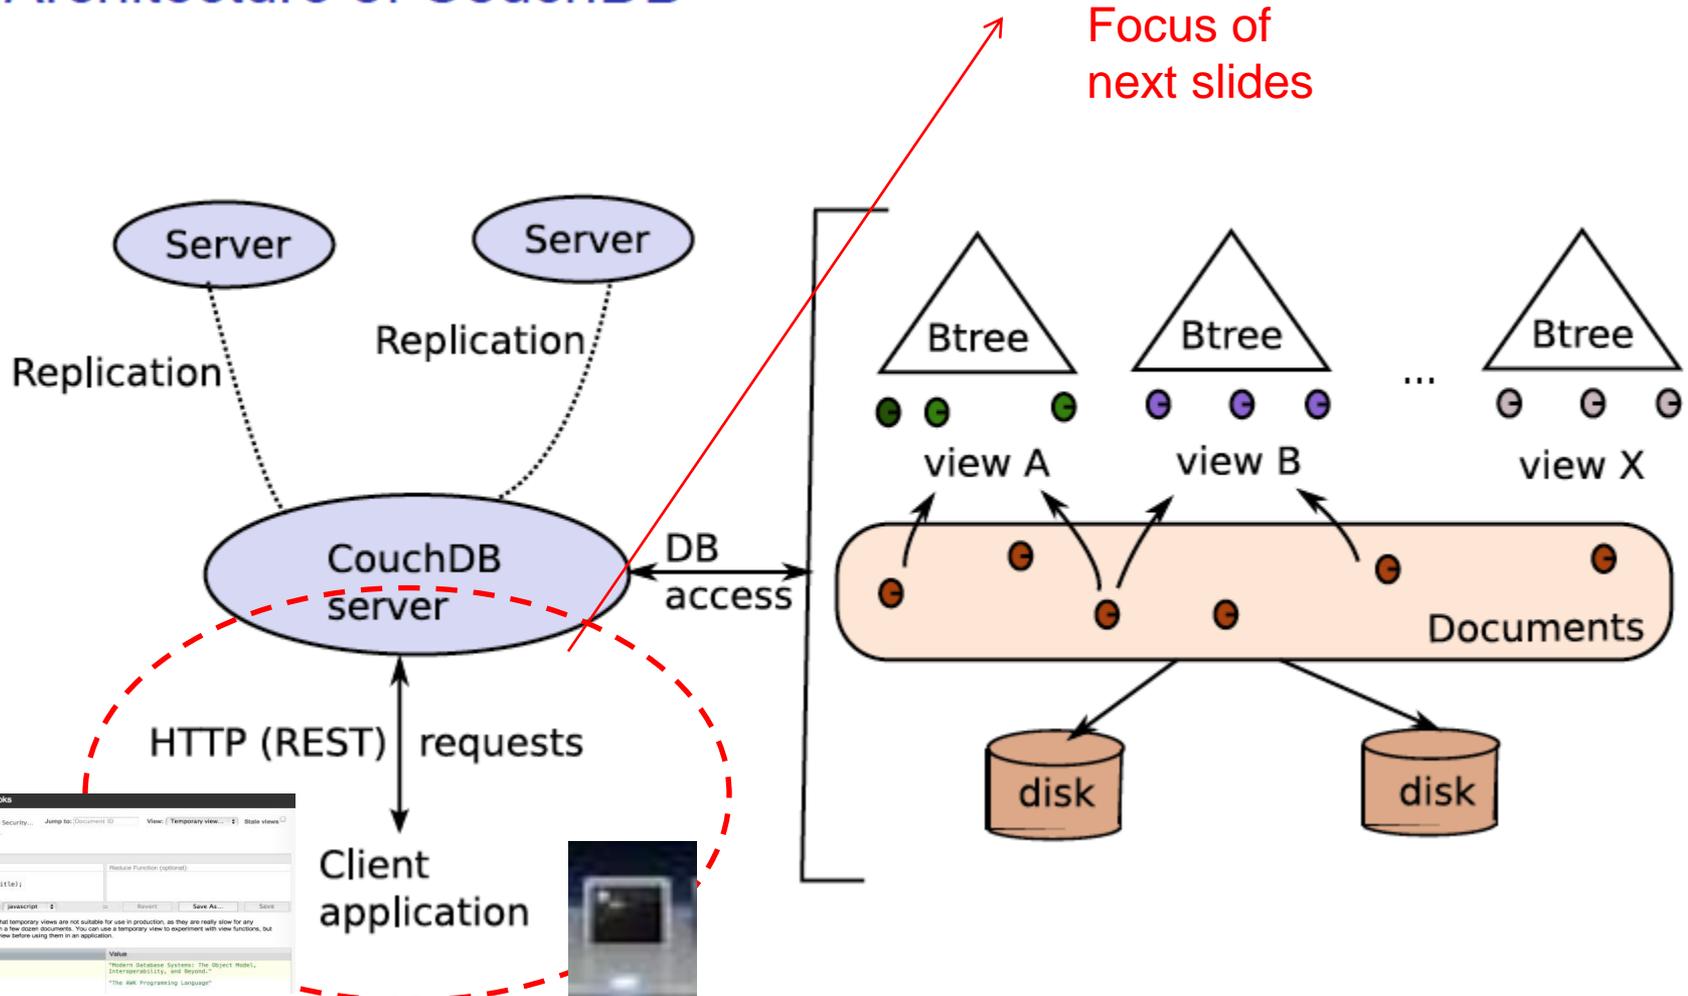


- Intro to Web2.0 & JSON Data Interchange Format
- JSON Key-Value Data Model
- CouchDB: A JSON Database (written in Erlang)
 - Using Command Line CURL/ Web-based FUTON
 - CouchDB Architecture (Btrees, Filesystem, Replication)
 - REST Principles
 - Creating DBs, Adding Docs, Updating Docs, Deleting Docs, `_ID` and `_REV` issues, Multi-Version CC (MVCC)
 - Querying Data with (Materialized) Views (Map-Reduce style in Javascript)
 - Replication and Scalability Issues

CouchDB REST Interface



Architecture of CouchDB



REST Principles



Apart©: REST principles

Roy Fielding (Univ. of Calif, Irvine PhD)
Founder of Apache HTTP Project
HTTP 1.0 (RFC1945) w/ Berners-Lee
HTTP 1.1 (RFC2616) w/ others



A Web-service dialect that enables exchanges of HTTP messages to access, create, and manage resources.
protocol as follows:

GET retrieves the resource referenced by the URI.

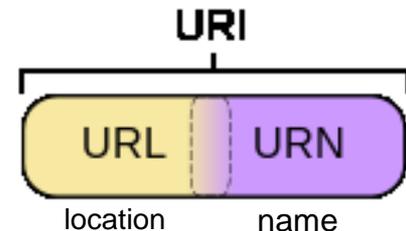
PUT creates the resource at the given URI.

POST sends a message (along with some data) to an existing resource.

DELETE deletes the resource.

GET

```
/test/demo_form.asp?name1=value1&name2=value2
```



POST

```
POST /test/demo_form.asp HTTP/1.1  
Host: w3schools.com  
name1=value1&name2=value2
```

Very convenient in a Web environment: no need to use a client library –
Documents can easily be incorporated in a Web interface.

CouchDB: CREATE DB / INSERT (with CURL)



A short interactive session

Curl: Client URL Command Line Tool

Talk to the server: send an HTTP request, get a response.

```
$ curl -X GET http://mycouch.org
{"couchdb": "Welcome", "version": "1.0.1"}
```

Create a db = put a resource (the name suffices).

```
$ curl -X PUT http://mycouch.org/myDB
{"ok": true}
```

Create a document = put a resource in a db (give the JSON document in the HTTP request).

```
$ curl -X PUT http://mycouch.org/myDB/myDoc \
-d '{"key": "value"}'
{"ok": true, "id": "myDoc", "rev": "1-25eca"}
```

Get the document after its URI:

```
$ curl -X GET http://mycouch.org/myDB/myDoc
{"_id": "myDoc", "_rev": "1-25eca", "key": "value"}
```

Retrieve all tables (e.g., show tables; in MySQL)

```
$ curl -X GET http://127.0.0.1:5984/_all_dbs
["_replicator", "_users", "books", "booksreplica", "movies", "twitter"]
```

CouchDB: `_ID` and `_REV`



Document management in CouchDB

Field	Value
<code>_id</code>	"book1.json"
<code>_rev</code>	"1-410c67caca526b476abc72e73b003605"

Each document has an **id** and a **revision number**.

Use: ``uuidgen`` to generate uniq ids during insert

UUID=128bit (32Hex digits) =
 2.4×10^{38} keys

Each update to a document creates a new version, with the same `_id` but a new revision number.

Validation functions can be assigned to a collection: any document inserted or updated must be validated by these functions (ad-hoc type-checking). *Like triggers...*

Design documents are a special type of CouchDB document that contains application code.

A **view** is a new key-document collection, specified via MAPREDUCE.

(we will see views later)

Documents can be **replicated** in other CouchDB instances.

(we will see replication later)

CouchDB: BULK LOAD (with CURL)



```
# Download the files from the web
#(or wget http://webdam.inria.fr/Jorge/files/jsonmovies.zip)
$ curl -O http://webdam.inria.fr/Jorge/files/jsonmovies.zip

# Unzip Movies
$ unzip jsonmovies.zip

# or assign unique IDs
$COUCHDB/movies/`uuidgen`

# List the files
$ ls -al | head
total 12480
-rw-r--r--@    1 dzeina  staff    218   9  ??t   2011  book1.json
-rw-r--r--@    1 dzeina  staff    222   9  ??t   2011  book10.json
-rw-r--r--@    1 dzeina  staff    197   9  ??t   2011  book100.json

# Bulk load using Bash
$ for i in `ls .`; do curl -X PUT $COUCHDB/movies/$i -d \@$i; done
{"ok":true,"id":"book1.json","rev":"1-410c67caca526b476abc72e73b003605"}
{"ok":true,"id":"book10.json","rev":"1-d0cc2ae0ab3211314a65a5c5244df221"}
{"ok":true,"id":"book100.json","rev":"1-2cfe83eea8cad920cfd66755ac78b46f"}
```

CouchDB: UPDATE / DELETE (with CURL)



Updating data

Updating in COUCHDB = adding a new version.

COUCHDB applies a **Multi-version concurrency control** protocol which requires that you send the version that must be updated:

```
$ curl -X PUT $IP/movies/tsn?rev=1-db1261 -d @newDoc.json \  
-H "Content-Type: image/jpg" \  
{ "ok": true, "id": "tsn", "rev": "2-26863" }
```

```
  "_attachments": {  
    "d1.tiff": {  
      "content_type": "image/tiff",  
      "revpos": 6,  
      "digest": "md5-2SYLdVnPnGNXPz10As8X5g==",  
      "length": 123776,  
      "stub": true  
    }  
  }  
}
```

Image separated from json

Deletion is obtained with DELETE.

```
$ curl -X DELETE $COUCHADDRESS/movies/tsn?rev=2-26863 \  
{ "ok": true, "id": "tsn", "rev": "3-48e92b" }
```

A new version has been created! (logical deletion).

CouchDB: UPDATE / DELETE (MVCC Explained)



- **Multi-Version CC (MVCC)** uses timestamps or increasing IDs to achieve transactional consistency.
 - MVCC provides each user with a snapshot of the database
- The database holds multiple versions (e.g., t0, t1 for Obj1)

Time	Object 1	Object 2	
t1	"Hello"	"Bar"	Latest Version
t0	"Foo"	"Bar"	

- Assume that a **T1** is a long-running READ xact started with state "t1"
- Also assume that **T2** is a WRITE transaction that creates state "t2"

T1 continues to have access to the t1 state! (thus, the necessary ISOLATION)

Time	Object 1	Object 2	Object 3	
t2	"Hello"	(deleted)	"Foo-Bar"	Latest Version
t1	"Hello"	"Bar"		
t0	"Foo"	"Bar"		

For **serializability / recoverability**: Multiversion Histories need to follow similar ideas to Basic TO (timestamp ordering) and Strict TO (TO w/ commits).

CouchDB: UPDATE / DELETE (with CURL)

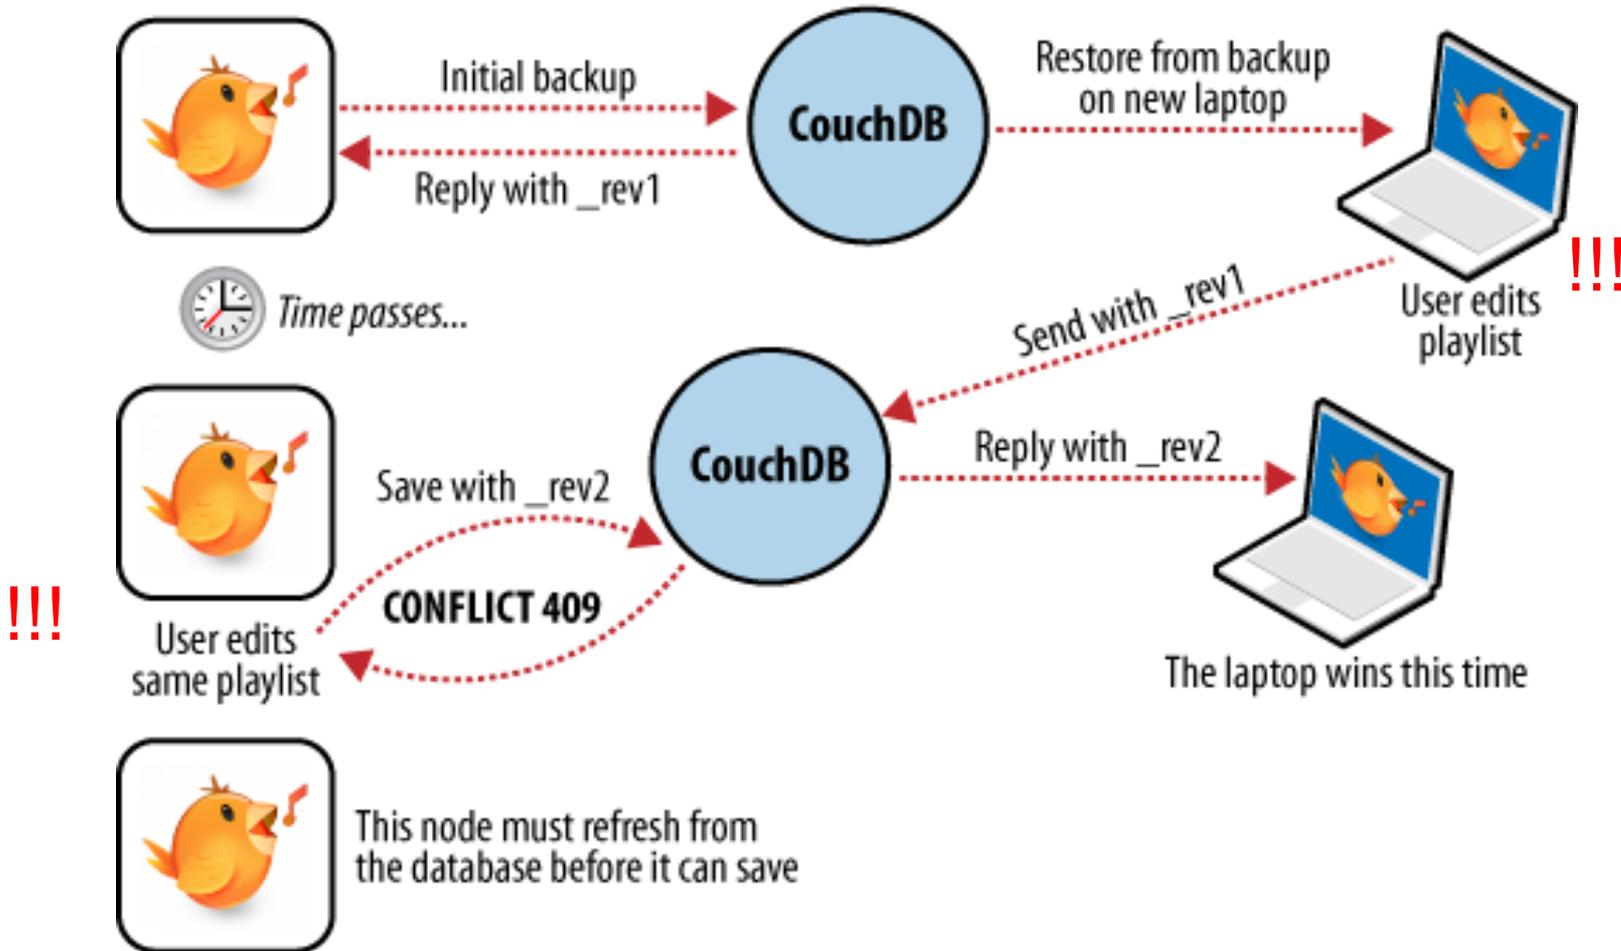


```
# UPDATE document "4C4F2E4C-E1AC-4C80-B90E-A9D0BCB638C8"
$ curl -X PUT $COUCHDB/googlebooks/4C4F2E4C-E1AC-4C80-B90E-
  A9D0BCB638C8?rev=1-1d974c2aadd42b0b8699678d61a0e4ff -d
  @potter.json
{"ok":true,"id":"4C4F2E4C-E1AC-4C80-B90E-
  A9D0BCB638C8","rev":"2-40d4b6bf3530e6af3a84904652ce9a8c"}
# Now DELETE PRELAST version "rev=1-
  1d974c2aadd42b0b8699678d61a0e4ff"
$ curl -X DELETE $COUCHDB/googlebooks/4C4F2E4C-E1AC-4C80-B90E-
  A9D0BCB638C8?rev=1-1d974c2aadd42b0b8699678d61a0e4ff
{"error":"conflict","reason":"Document update conflict."}
# DELETE LAST version "rev=2-40d4b6bf3530e6af3a84904652ce9a8c"
$ curl -X DELETE $COUCHDB/googlebooks/4C4F2E4C-E1AC-4C80-B90E-
  A9D0BCB638C8?rev=2-40d4b6bf3530e6af3a84904652ce9a8c
{"ok":true,"id":"4C4F2E4C-E1AC-4C80-B90E-
  A9D0BCB638C8","rev":"3-a77d6120602b51cbf9b8663c1ee8f9e3"} 13-11
EPL646: Advanced Topics in Databases - Demetris Zeinalipour (University of Cyprus)
```

CouchDB: UPDATE / DELETE (Example)



Syncing playlists between multiple Songbird clients



CouchDB: SELECT QUERY (expressed as Javascript)



Processors

Views in CouchDB

A **view** is the result of a **MAPREDUCE** job = a list of (*key, value*) pairs.

+ Temporary View: executed on demand (e.g., through Futon) – good for development

+ Permanent (Materialized) Views: also called design documents, accessible through URI

Views are **materialized** and indexed on the key by a B+tree.

A MAP function

*we will study the **Map-Reduce Programming Model**, i.e., founded on **BSP** (**Bulk Synchronous Parallelism**) more extensively in the next lecture.*

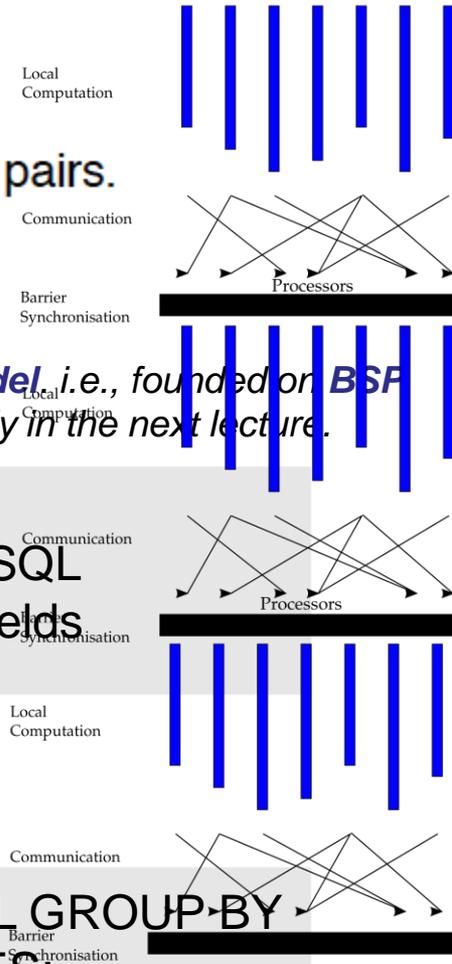
```
function (doc)
{
  emit(doc.title, doc.director)
}
```

Similar to SQL
SELECT fields

A REDUCE function

```
function (key, values) {
  return values.length;
}
```

Similar to SQL GROUP-BY
AGGREGATES;

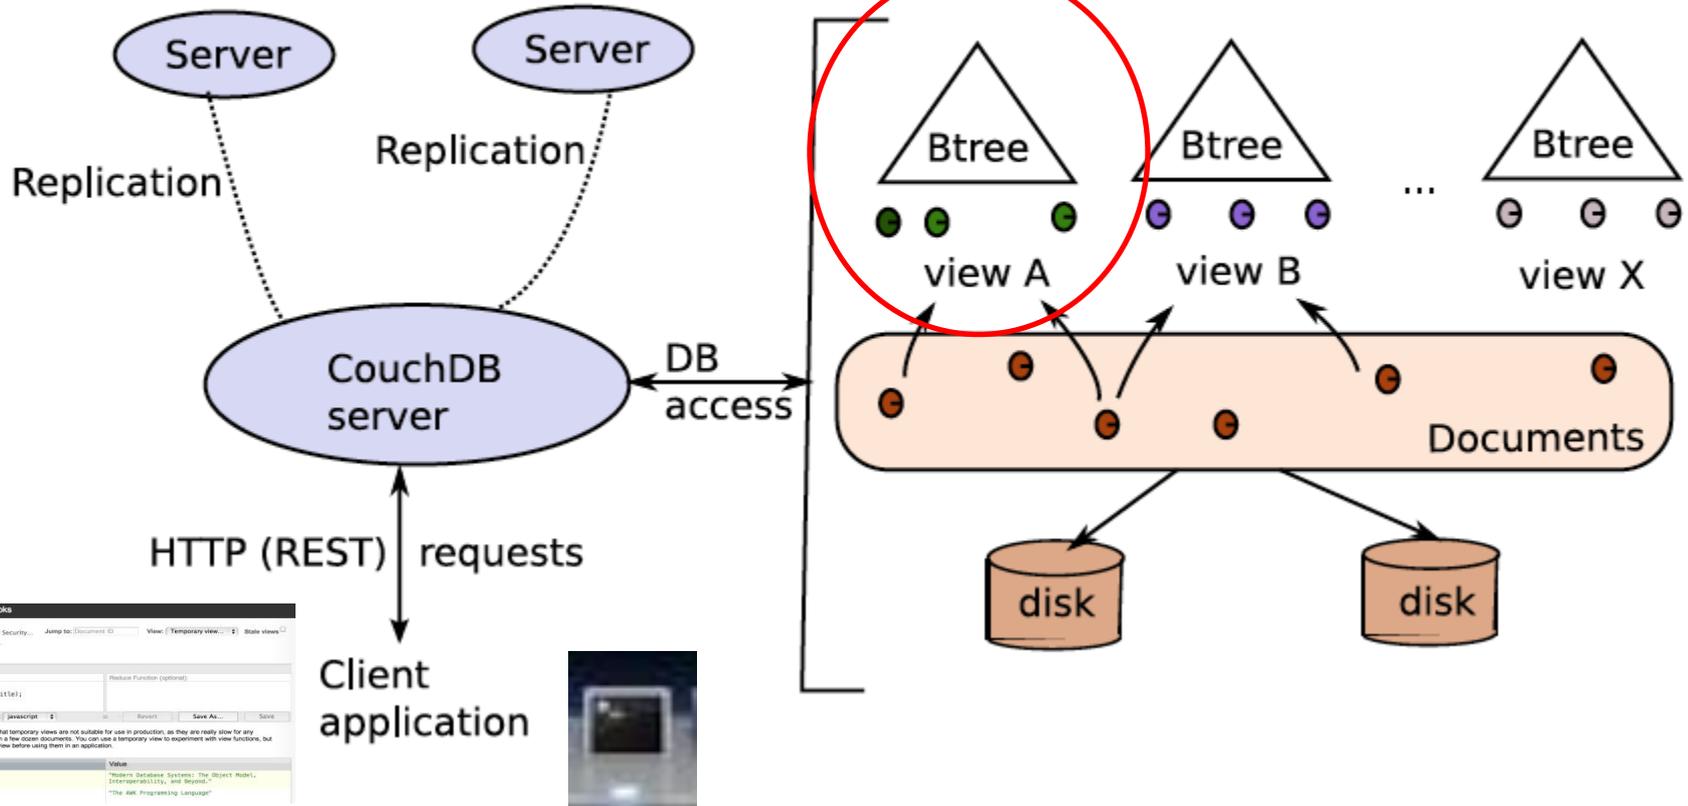


CouchDB: A JSON Database (Architecture)



Architecture of CouchDB

B+tree Key: [key,docid]
 (Materialized view => on update tree is updated as well)



CouchDB: Filesystem Layout (Datastores and Materialized Views)



Index of /Users/dzeina/Li

Name	Size	Date
[parent directory]		
.books_design/		10/29/12 11
.delete/		10/23/12 3
.movies_design/		10/23/12 4
_replicator.couch	4.1 kB	10/23/12 9
_users.couch	4.1 kB	10/23/12 9:40:05 AM
books.couch	3.7 MB	10/29/12 11:41:48 AM
booksreplica.couch	764 kB	10/29/12 11:41:48 AM
movies.couch	4.1 kB	10/23/12 3:40:27 PM
twitter.couch	4.1 kB	10/23/12 10:50:02 AM

Index of /Users/dzeina/Library/Application S

Name	Size	Date Modified
[parent directory]		
1fb02d640d642a6272bfa5334d2e3f42.view	188 kB	10/29/12 11:18:59 AM
20ab140f1492382baf3aafbab426f2d7.view	56.1 kB	10/23/12 4:12:14 PM
33abe94d1a00f46506f2c0b015540db1.view	8.1 kB	10/29/12 11:19:29 AM
3b5f338a51e55052f5513da8a6bd64a8.view	216 kB	10/23/12 4:52:51 PM
451385dfae51196393f2ceb3a2b780fa.view	4.1 kB	10/29/12 11:19:29 AM
45b8fde80881d168a44595ff5cc90ea1.view	16.1 kB	10/29/12 11:48:33 AM
5ccddad630dd20ffd2b6fb4833080772.view	4.1 kB	10/23/12 4:58:52 PM
64a3e2d25cf3876b0fa55d78c50d7f0b.view	72.1 kB	10/29/12 11:48:33 AM
6ba736a20feacba97c220a4b7f3e02e3.view	132 kB	10/29/12 11:45:02 AM
6da619412b8ffa95126c11d2b21d27cc.view	504 kB	10/23/12 4:09:13 PM
6fac50850cee7de2f185090669defd68.view	4.1 kB	10/23/12 11:21:22 PM
70419c9fa523e611895f53b01ac694e0.view	204 kB	10/23/12 4:59:53 PM

Lecture Outline

(Introduction to Semi-structured Data)



- **SQL: SELECT * FROM Books;**

```
function(doc) {  
  emit(null, doc);  
}
```

Key ▲	Value
null ID: book1.json	{_id: "book1.json", _rev: "1-410c67caca526b476abc72e73b003605", type: "Book", title: "Modern Database Systems: The Object Model, Interoperability, and Beyond.", year: "1995", publisher: "ACM Press and Addison-Wesley", authors: [], source: "DBLP"}
null ID: book10.json	{_id: "book10.json", _rev: "1-d0cc2ae0ab3211314a65a5c5244df221", type: "Book", title: "The AWK Programming Language", year: "1988", publisher: "Addison-Wesley", authors: ["Alfred V. Aho", "Brian W. Kernighan", "Peter J.

- **SQL: SELECT pub FROM Books;**

```
function(doc) {  
  emit(doc._id, doc.publisher);  
}
```

Key ▲	Value
"book1.json" ID: book1.json	"ACM Press and Addison-Wesley"
"book10.json" ID: book10.json	"Addison-Wesley"

_ID always part of answer in Futon (but not _REV)

- **SQL: SELECT pub, typ FROM Books**

```
function(doc) {  
  emit(doc._id,  
    {"pub": doc.publisher, "typ": doc.type});  
}
```

Key ▲	Value
"book1.json" ID: book1.json	{pub: "ACM Press and Addison-Wesley", typ: "Book"}
"book10.json" ID: book10.json	{pub: "Addison-Wesley", typ: "Book"}

CouchDB: SELECT QUERY (with FUTON)



SQL Equivalent: **SELECT B.title FROM Books B;**

The screenshot shows the Apache CouchDB Futon interface for a database named 'books'. The 'View' dropdown is set to 'Temporary view...'. The 'Map Function' field contains the following JavaScript code:

```
function(doc) {  
  emit(doc._id, doc.title);  
}
```

The 'Results' table below shows the output of the query:

Key	Value
"book1.json" ID: book1.json	"Modern Database Systems: The Object Model, Interoperability, and Beyond."
"book10.json" ID: book10.json	"The AWK Programming Language"

Red annotations on the image include:

- A red circle around the 'View: Temporary view...' dropdown.
- Red text: 'Not saved yet => No Design document => no Index yet'.
- Red text: 'MAP (Query)' next to the Map Function code.
- Red text: 'Results' above the results table.
- Red circles around the Map Function code and the Results table.

CouchDB: SELECT-WHERE QUERY (with FUTON)



SQL Equivalent: **SELECT B.title FROM Books B WHERE B.publisher="AW";**

+ New Document Security... Jump to: View: **Temporary view...** Stale views

⊗ Compact & Cleanup...

⊗ Delete Database...

▼ View Code

Map Function:	Reduce Function (optional):
<pre>function(doc) { if (doc.publisher == "Addison-Wesley") { emit(doc._id, doc.title); } }</pre>	

Run Language: javascript = Revert Save As... Save

MAP (Query)

Warning: Please note that temporary views are not suitable for use in production, as they are really slow for any database with more than a few dozen documents. You can use a temporary view to experiment with view functions, but switch to a permanent view before using them in an application.

Results

Key ▲	Value
"book10.json" ID: book10.json	"The AWK Programming Language"
"book100.json"	"The Design and Analysis of Spatial Data Structures"

CouchDB: Array Object Iteration (with FUTON)



```
{ Document
  "_id": "book10.json",
  "_rev": "1-d0cc2ae0ab3211314a65a5c5244df221",
  "type": "Book",
  "title": "The AWK Programming Language",
  "year": "1988",
  "publisher": "Addison-Wesley",
  "authors": [
    "Alfred V. Aho",
    "Brian W. Kernighan",
    "Peter J. Weinberger"
  ],
  "source": "DBLP"
}
```

Map Function

```
function(doc) {
  for (i in doc.authors) {
    author = doc.authors[i];
    emit(doc._id, author);
  }
}
```

Check "View Cookbook for SQL Jockeys" for more!

<http://guide.couchdb.org/editions/1/en/cookbook.html>

Results

Key ▲	Grouping: exact	Value
"book10.json" ID: book10.json		"Alfred V. Aho"
"book10.json" ID: book10.json		"Brian W. Kernighan"
"book10.json" ID: book10.json		"Peter J. Weinberger"

CouchDB: UNIQUE Keys (with FUTON)



- If we want to make sure that a certain column is unique in the database, e.g., `UNIQUE KEY(column)`, we just make that column to become the `_id`:

`_id = uniquecolumn`

```
curl -X PUT $COUCHDB/movies/uniquecolumn -d \@file.json
```

- ***Pitfalls:***

- Uniqueness can be guaranteed only per node
- In single node: No problem!
- In master-master or master-slave replication (seen later), CouchDB will allow two identical IDs to be written to two different nodes.
 - On replication, CouchDB will detect a conflict and flag the document accordingly.

CouchDB: Array Object Iteration (with FUTON + CURL)

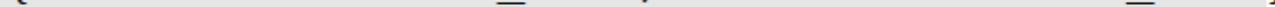


Accessing views

Here is a view (without reduce function).

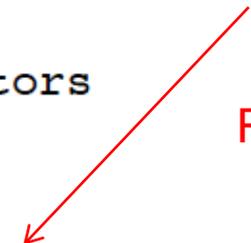
Array object iterator

```
function(doc)
{
  for (i in doc.actors) {
    actor = doc.actors[i];
    emit({"fn": actor.first_name, "ln": actor.last_name},
    }
  }
}
```

Save it in the **design document** named `examples`, and name the view `actors`. The view can be queried with:

```
$ curl $IP/movies/_design/examples/_view/actors
{"total_rows":16,"offset":0,
"rows":[
  {"id":"bed7",
   "key":{"fn":"Andrew","ln":"Garfield"},"value":"The Socia
  {"id":"91631b",
   "key":{"fn":"Clint","ln":"Eastwood"},"value":"Unforgiven
```



CouchDB: Design Documents (Apps are Documents!)



A **Design document** is a CouchDB document with an id that begins with `_design/`. These documents store the materialized view code + other code (e.g., for transforming data to HTML, etc.)

```
1 {
2   "_id": "_design/sofa", ← Determines the app URL
3   "_rev": "3157636749",
4
5
6   "language": "javascript", (for the web)
7
8
9   "validate_doc_update": "function (newDoc, oldDoc, userCtx) { ... }",
10                          Application is stored as JSON data
11
12   "views": { ← Views field stores incremental
13     "comments": { map reduce functions
14       "map": "function(doc) { ... };",
15       "reduce": "function(keys, values, rereduce) { ... };",
16     }
17   },
18
19   "shows": { ← Shows functions transform
20     "post": "function(doc, req) { ... }"
21   },
22
23   "attachments": { ← Attachments show
24     "jquery.couchapp.js": {
25       "stub": true,
26       "content_type": "text/javascript",
27       "length": 7539
28     }
29   },
30
31
32   "signatures": { ← CouchApp traces attachments here
33     "jquery.couchapp.js": "80078849ad6ca281f6993bd012c708f5",
34   },
35
36
37
38   "lib": {
39     "templates": { ← CouchApp can include
40       "post": "<!DOCTYPE html> ... </html>"
41     }
42   }
43 }
```

Like XSLT

CouchDB: Design Documents (Apps are Documents!)



CouchDB Design Documents can lead to "Standalone Web Apps"

July 2009

today < >

Sun	Mon	Tue	Wed	Thu	Fri	Sat
28 3a Amsterdam.rb	29	30 12a Kingsofcode	1 12a Kingsofcode 5a SysMap	2 6a Add new calendar item 9p Armageddon Day	3	4 6a foo 3p asf
5	6 11p hello	7	8 11a ruleza	9	10	11

Help | Shopping Cart

 Search



shawl, "x70"



Blue Kippah with Khaki border £2.99
 White Kippah with Black and Blue border £2.99
 White Kippah £2.99



Messianic Scripture Songs from Israel £8.55
 The Jewish Album - Jonathan Settel £8.55
 Silver-Plated Jerusalem Tallit Clip £9.99

hadai zuzah

- More Information
 - Shipping & Returns
 - Privacy Policy
 - Conditions of Use
 - Links & Resources
 - How to Order
 - RSS Feeds
- Currency
 - Pound Sterling (£) ↓
- Best Sellers
 - Acrylic Prayer Shawl, Blue & Silver (23"x70")
 - Blue Kippah with Khaki border
 - White Kippah with Black and Blue border

18
25
1a Corp Gonzo
9p Super 64 Tournament
1
8

CouchDB: Querying Materialized Views (with FUTON)



Querying views (Permanent | Materialized)

A view is a B+tree index. So:

```
function(doc)
{
  emit(doc.genre, doc.title) ;
}
```

is equivalent to

```
create index on movies (genre);
```

Recall the B+trees support key and range queries:

```
$ curl $IP/movies/_design/examples/_view/genre?key=\"Drama\"
{"total_rows":5,"offset":2,"rows":[
{"id":"9163", "key":"Drama","value":"Marie Antoinette"},
{"id":"bed7", "key":"Drama","value":"The Social network"}
]}
```

For range queries, send the two parameters `startkey` and `endkey`. (next slide)

CouchDB: Range Queries (with CURL)



CURL Command

```
curl -X GET  
http://127.0.0.1:5984/books/_design/authors/_view/autho  
rs?startkey=\"book980.json\"&endkey=\"books998.json\"
```

Results

View count

First result occurrence

```
$ {"total_rows":1893,"offset":1871,"rows":[  
{"id":"book980.json","key":"book980.json","value":"A. J. Kfoury"},  
{"id":"book980.json","key":"book980.json","value":"Michael A. Arbib"},  
{"id":"book980.json","key":"book980.json","value":"Robert N. Moll"},  
{"id":"book981.json","key":"book981.json","value":"Peter D. Mosses"},  
{"id":"book982.json","key":"book982.json","value":"Anne Mulkers"},  
{"id":"book983.json","key":"book983.json","value":"ller"},  
{"id":"book984.json","key":"book984.json","value":"ller"},
```

CouchDB Querying with Aggregation (Map & Reduce)



90% of queries will be MAP queries.

REDUCE allows us to carry out aggregation on a given field.

SQL:
SELECT key, sum(key)
FROM Books
GROUP BY key

Overview > books

+ New Document ? Security... Jump to: View: Stale views

⊗ Compact & Cleanup...
⊗ Delete Database...

View Code _design/examples

Map Function:	Reduce Function (optional):
<pre>function(doc) { emit(doc.publisher, 1); }</pre>	<pre>function(key, values, rereduce) { return sum(values); }</pre>

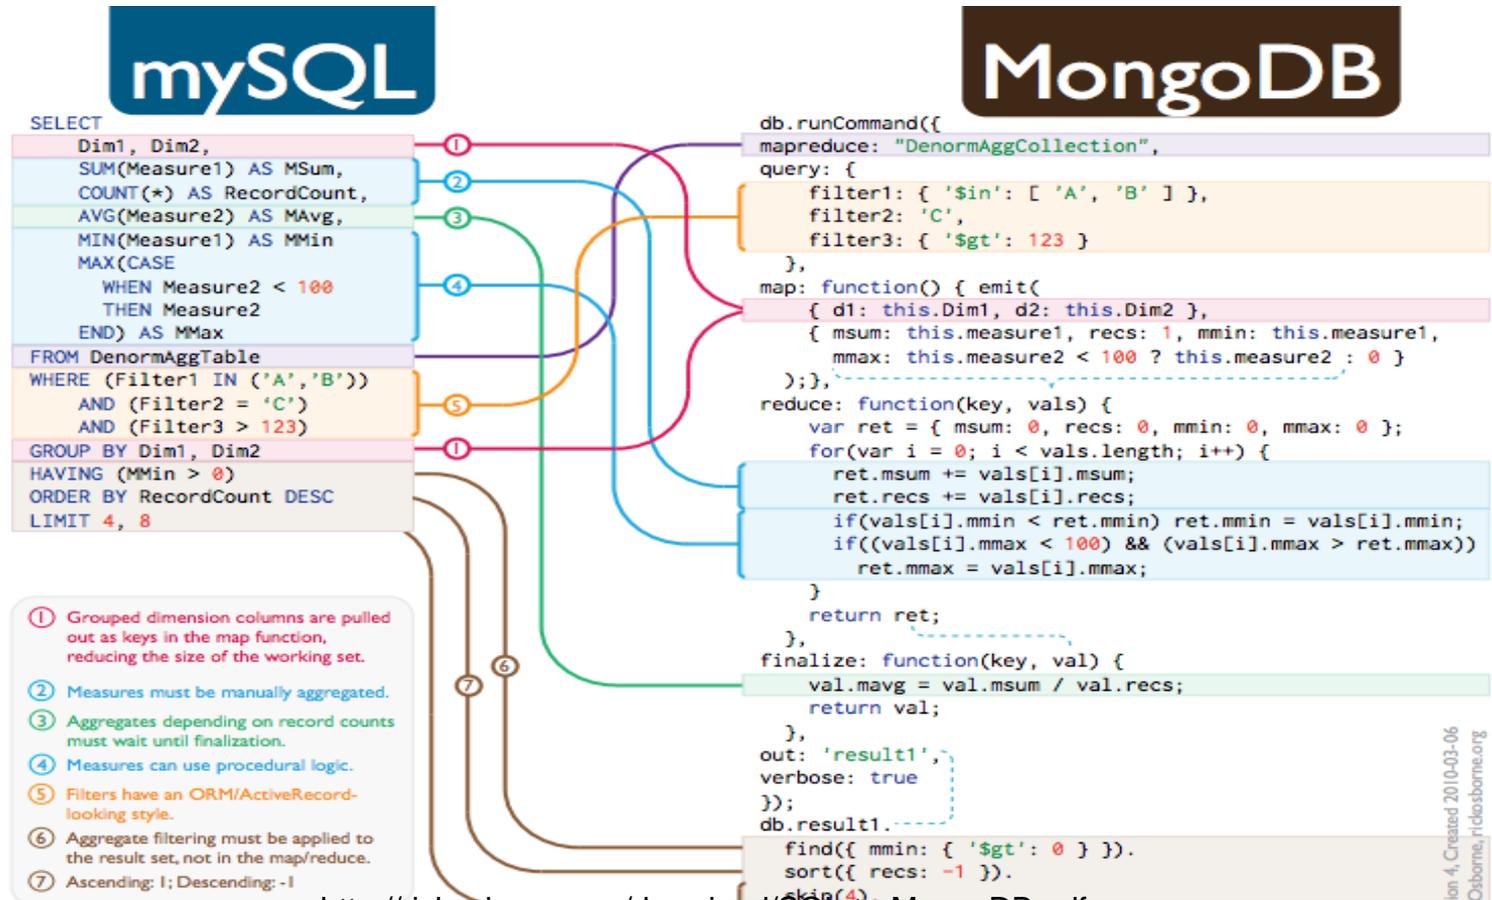
Run Language: = Revert Save As... Save

Key ▲	Grouping: <input type="text" value="exact"/>	Value	<input checked="" type="checkbox"/> Reduce
" ACM Press"		1	
" Claypool Publishers"		83	
" Francis"		1	
" Row"		1	
" Sons"		1	

Complex Map & Reduce



For a real app we could envision much more complex queries.



<http://rickosborne.org/download/SQL-to-MongoDB.pdf>

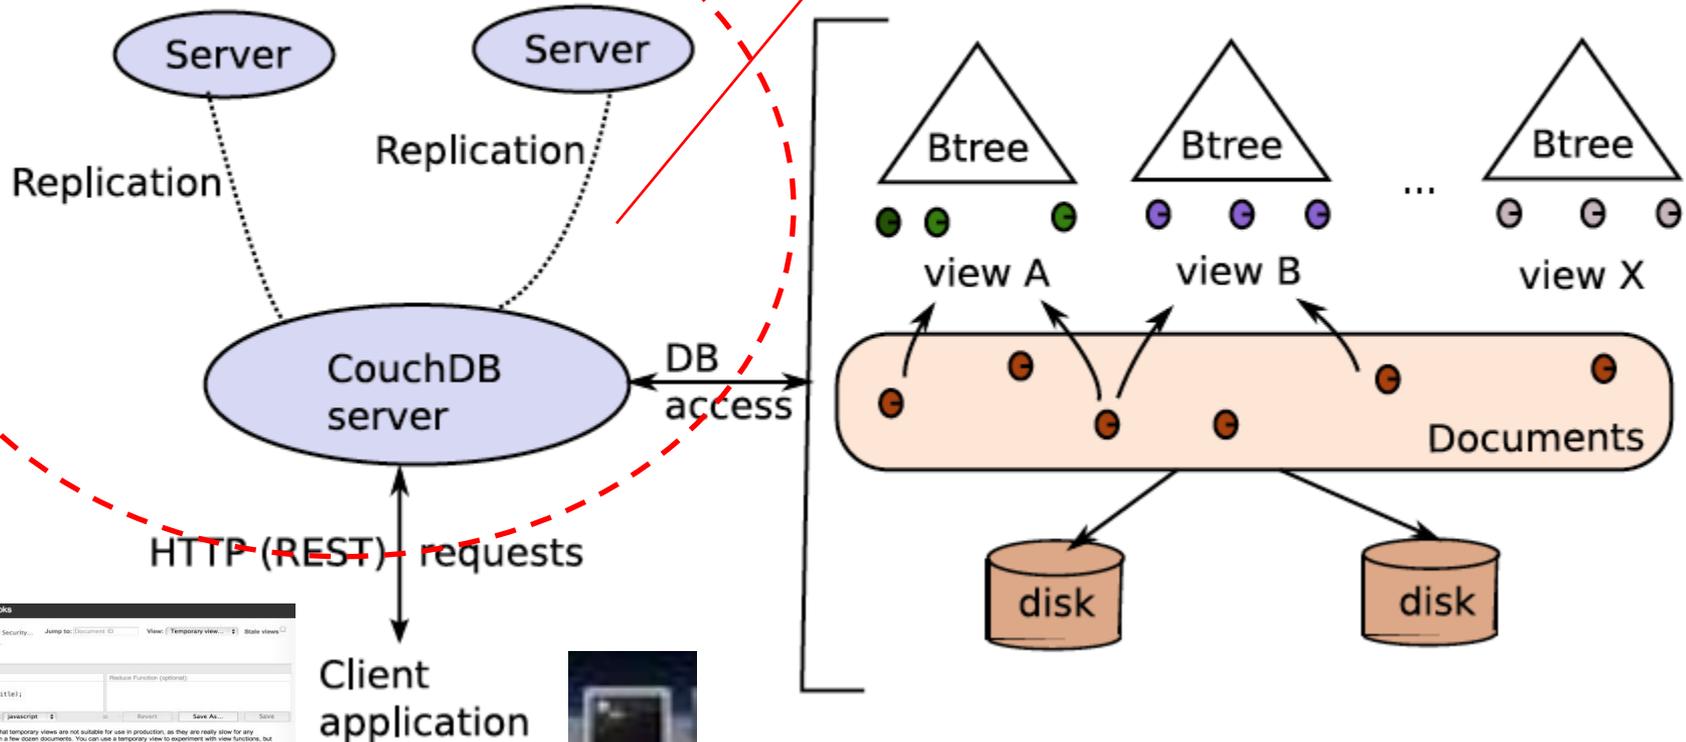
ion 4, Created 2010-03-06
 Osborne, rickosborne.org

CouchDB Replication



Architecture of CouchDB

Focus of next slides



CouchDB One-way/Symmetric Replication (with CURI)



The replication primitive

COUCHDB supports natively one-way **replication** from one instance to another.

```
curl -X POST $COUCHADDRESS/_replicate \  
  -d '{"source": "movies", "target": "backup", \  
      "continuous": true}' \  
  -H "Content-Type: application/json"
```

That's all: any change in `movies` is automatically reported in `backup`.

Two way replication can be achieved by executing the inverse statement as well.

CouchDB One-way/Symmetric Replication (with CURL)



Overview > Replicator

Replicate changes from:

Local Database: ↕

Remote database:

to:

Local database:

Remote database:

Continuous

Event

No replication

CouchD
rel

Tools

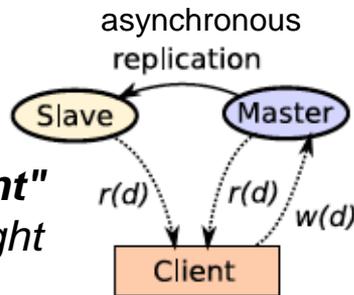
- Overview
- Configuration
- Replicator**
- Status
- Verify Installation

CouchDB Distribution (Concepts)



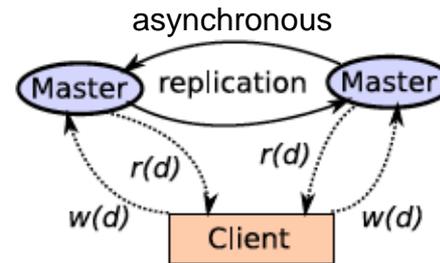
Distribution strategies

Combine a proxy that distributes requests, with the replication feature of COUCHDB.



"Eventually Consistent"
(reading from slave might be outdated)

a - Master-slave arch.



"Conflict Management is Necessary"

b - Master-master arch.

```
function(doc) {  
  if (doc._conflicts) {  
    emit(doc._conflicts, null);  
  }  
}
```

CouchDB Security



- CouchDB allows any request to be made by anyone (i.e., by default admin).
 - By default, CouchDB will listen only on your loopback network interface (127.0.0.1 or localhost) and thus only you will be able to make requests to CouchDB, nobody else.
- If you don't like that, you can create specific admin users with a username and password as their credentials.
- You could also add **validation functions** to your design documents that will make sure that the right people make the right changes.
- **Securing CouchDB is outside the scope of this lecture... The same applies to web programming with CouchDB.**