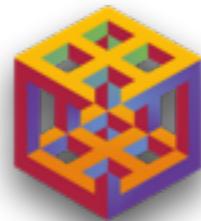


Declarative,
Secure,
Convergent
Edge Computation

Christopher Meiklejohn

Erlang Factory 2016, March 10th, 2016

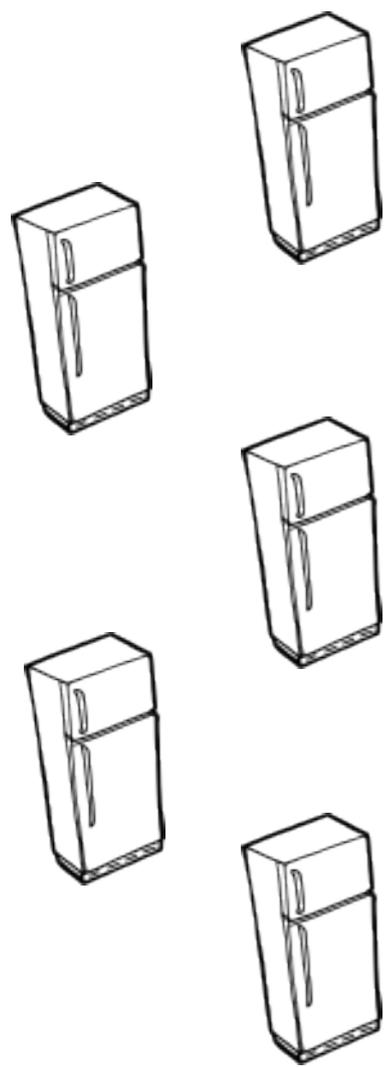


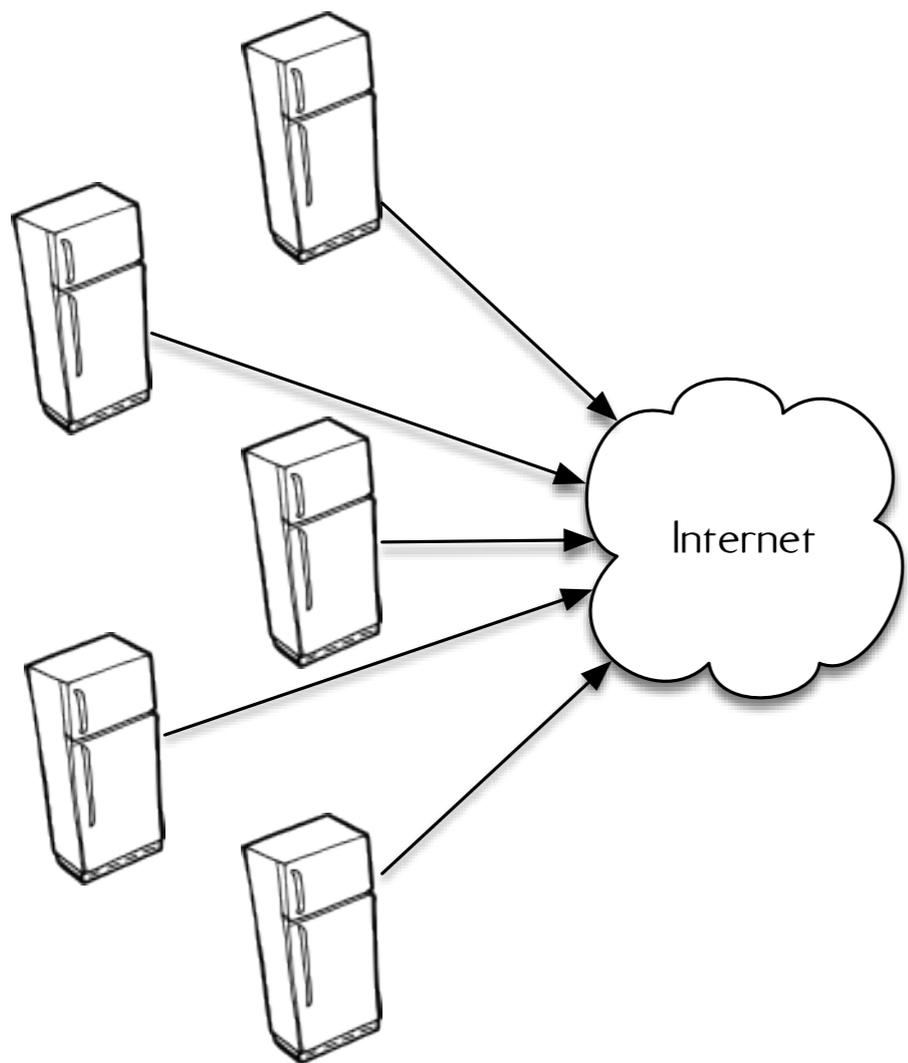
Example Application

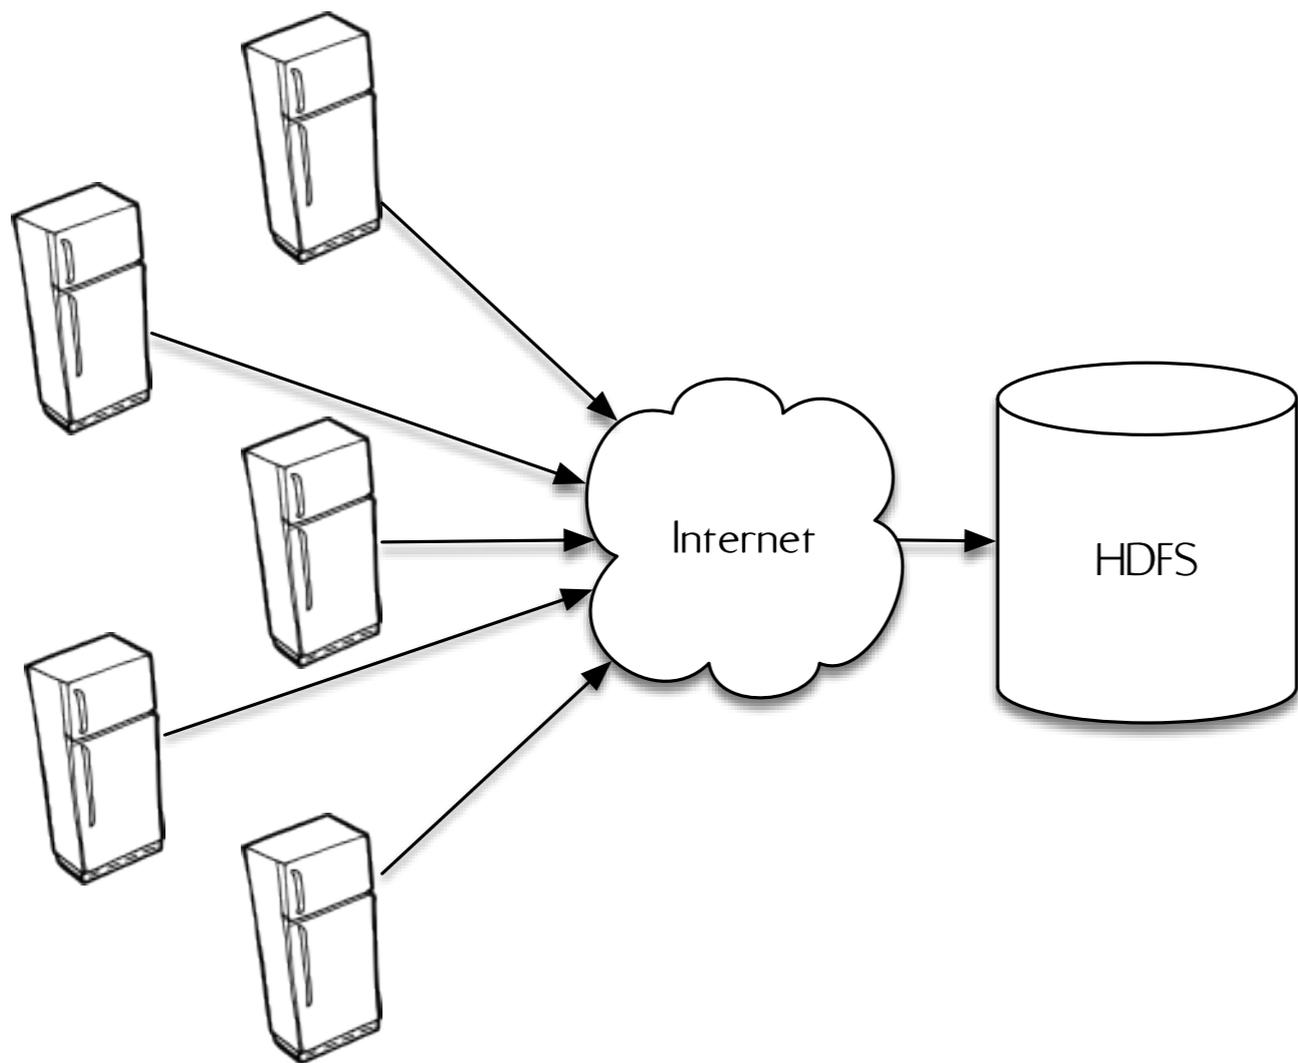
Hospital Refrigerators

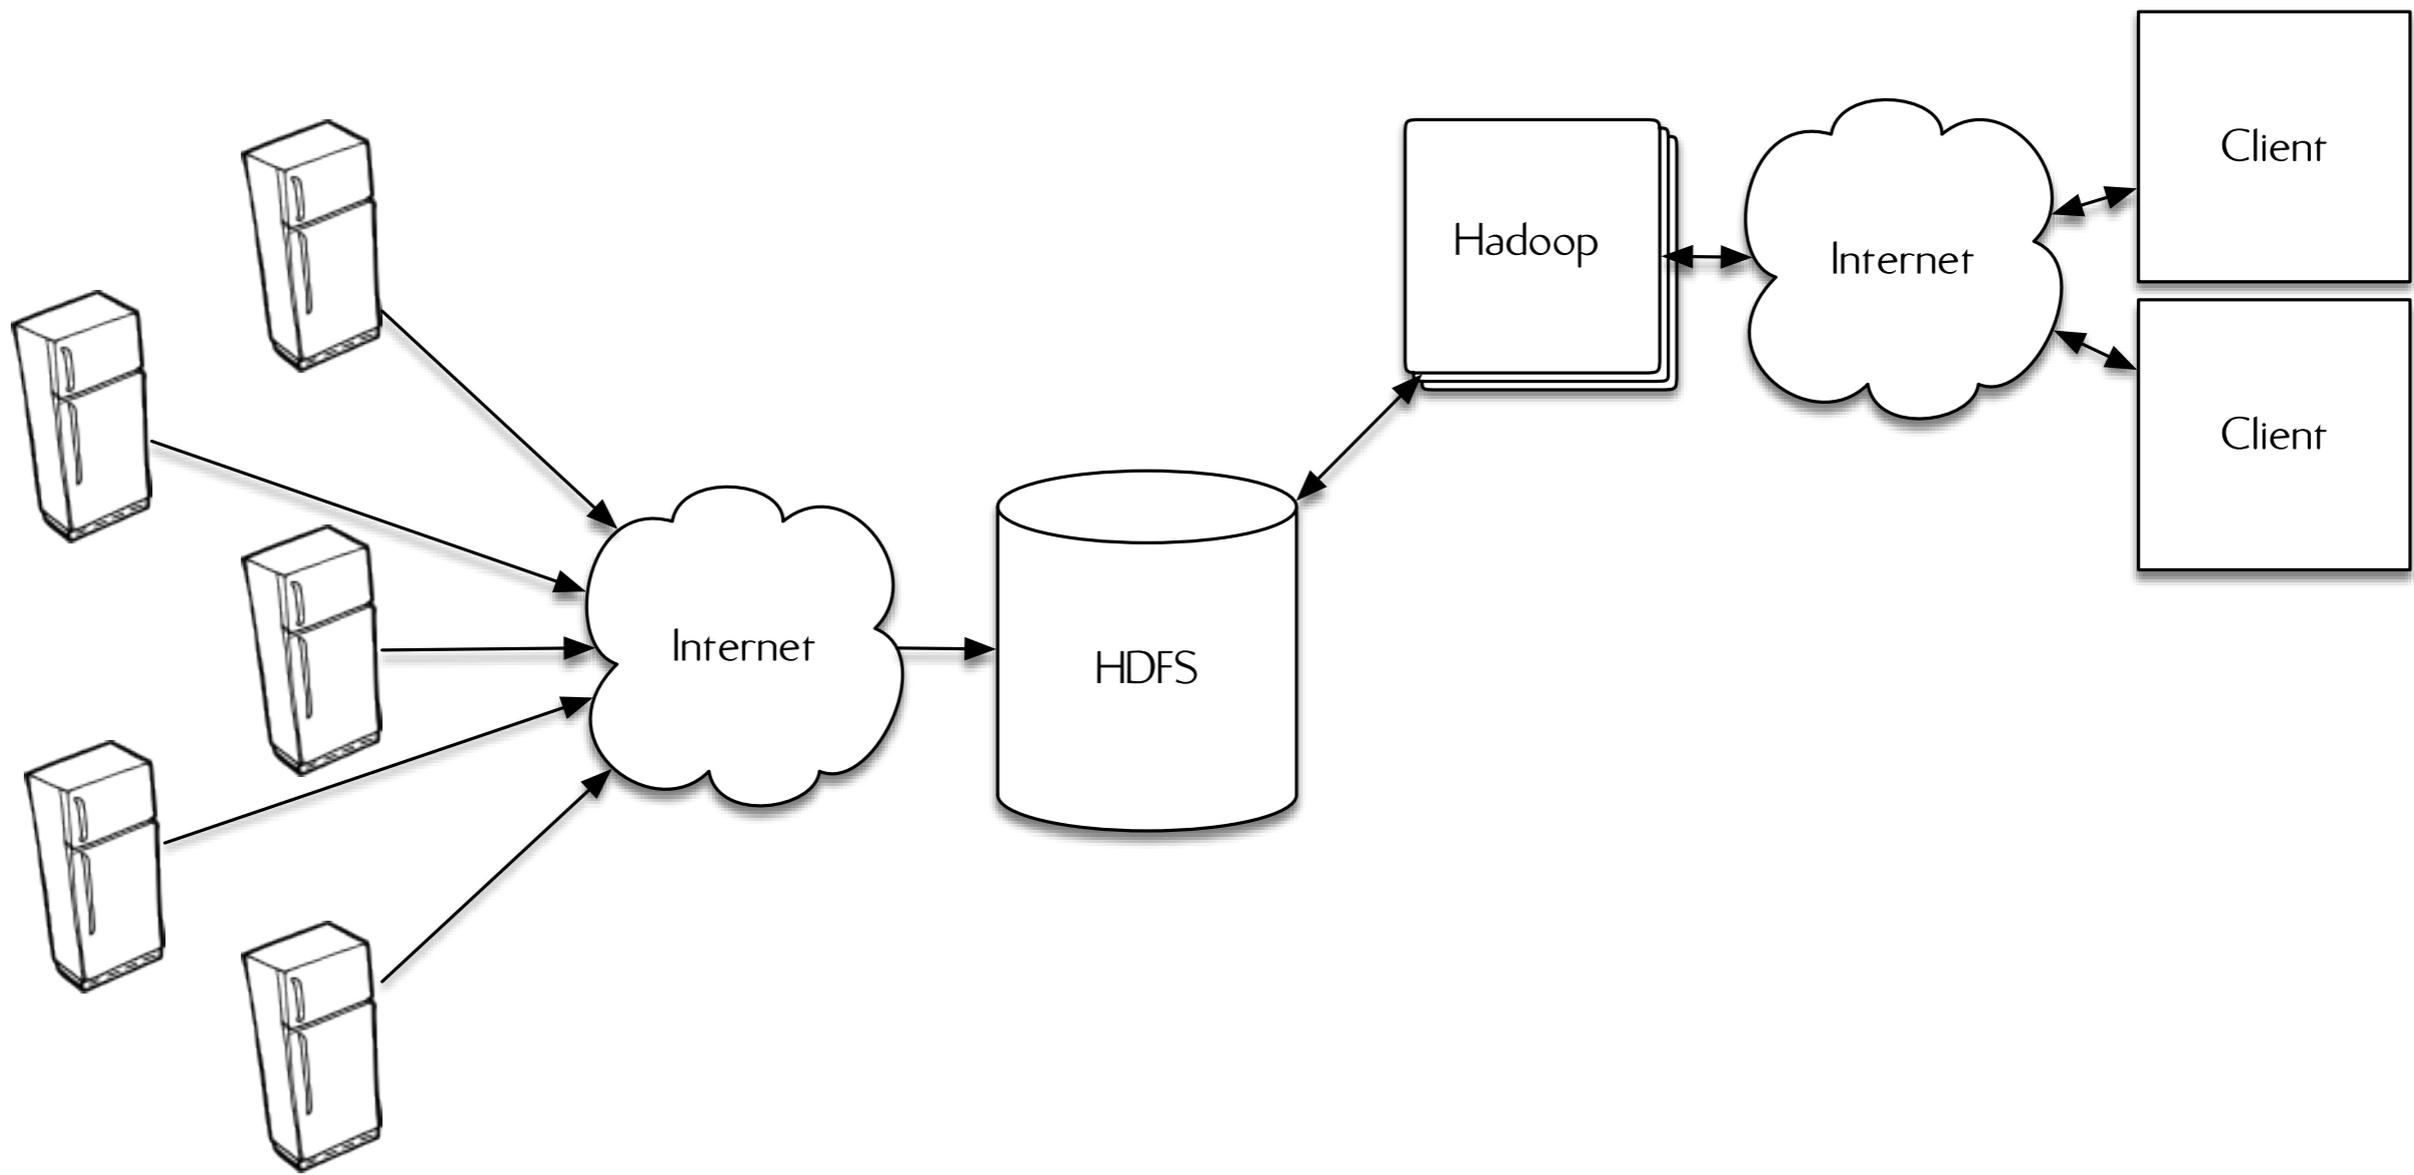
Hospital Refrigerators

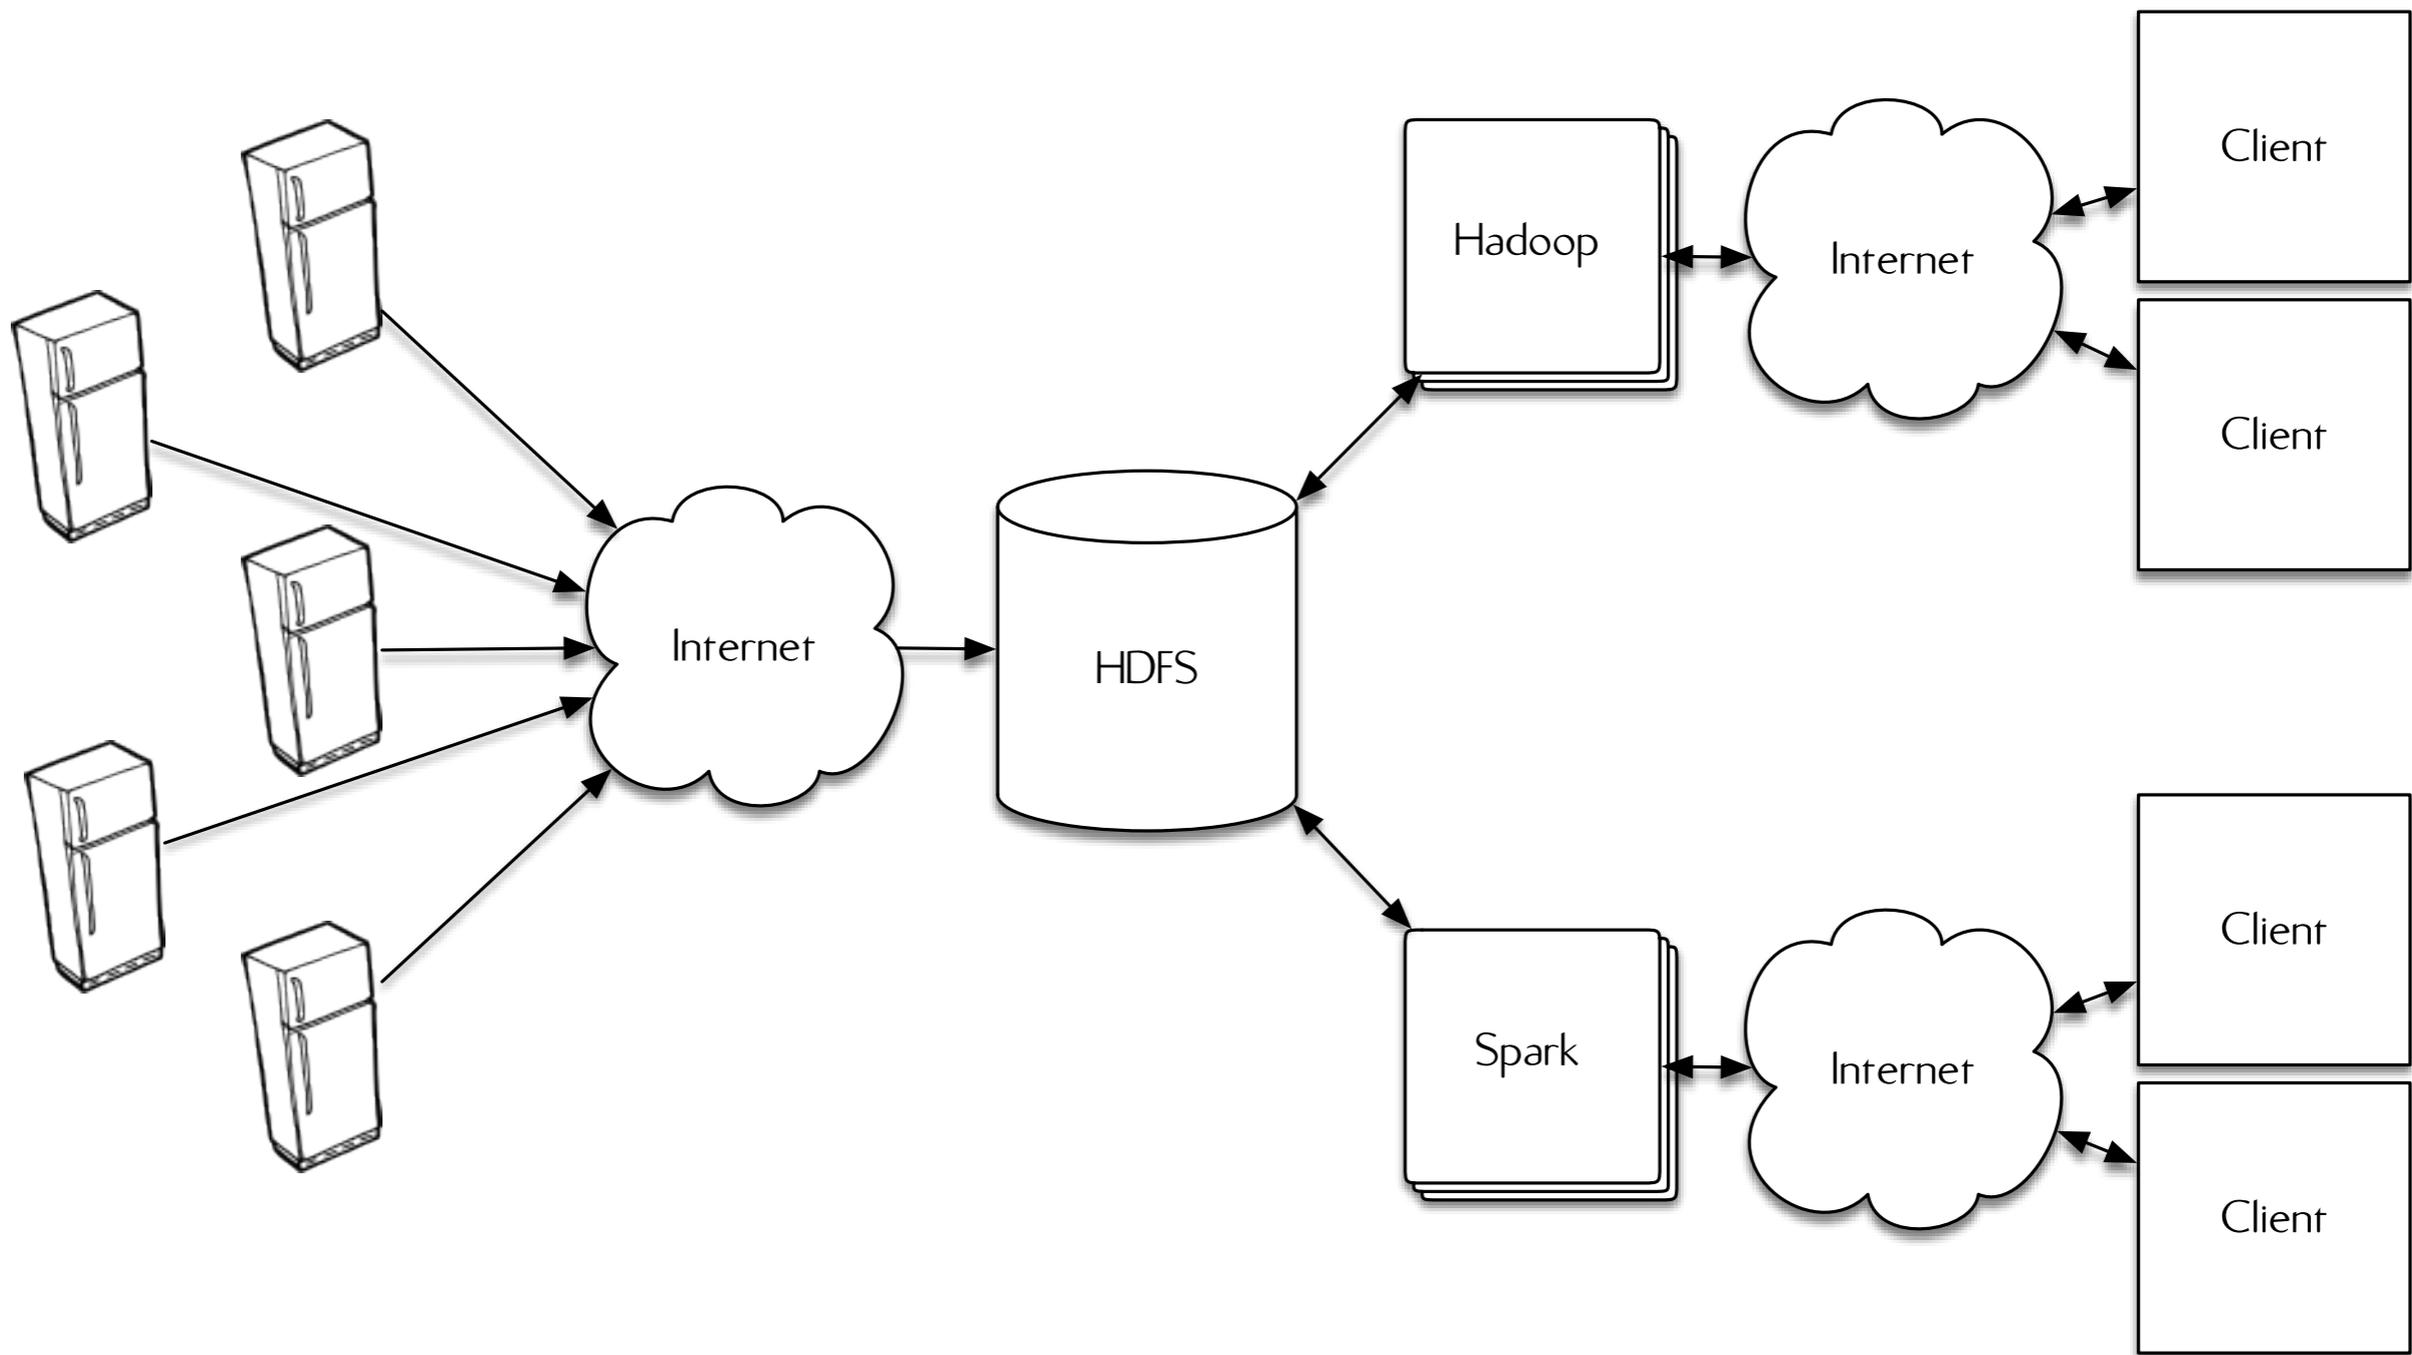
Typical Topology





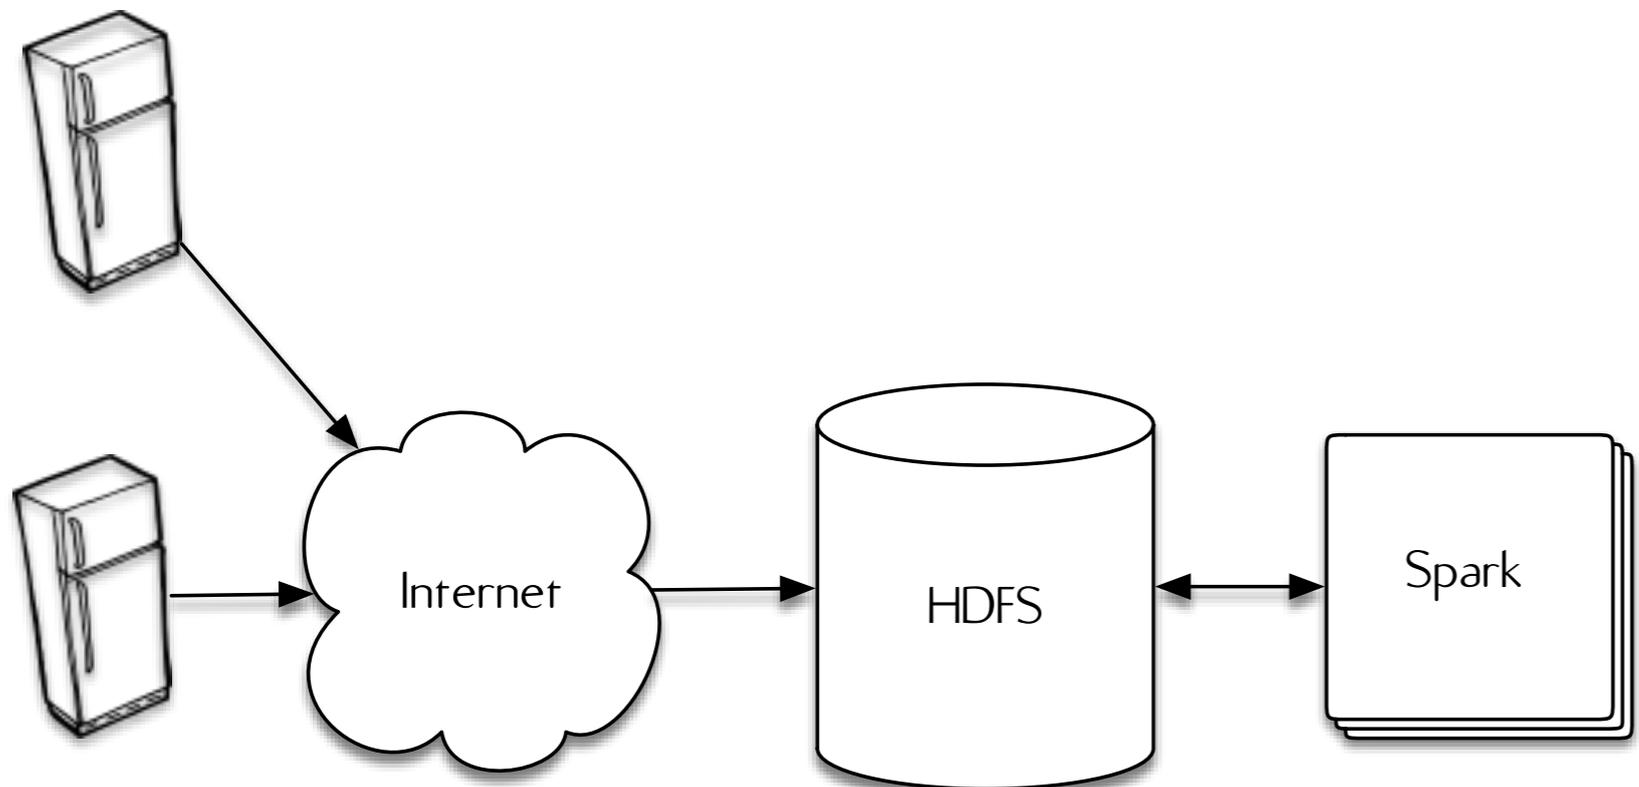


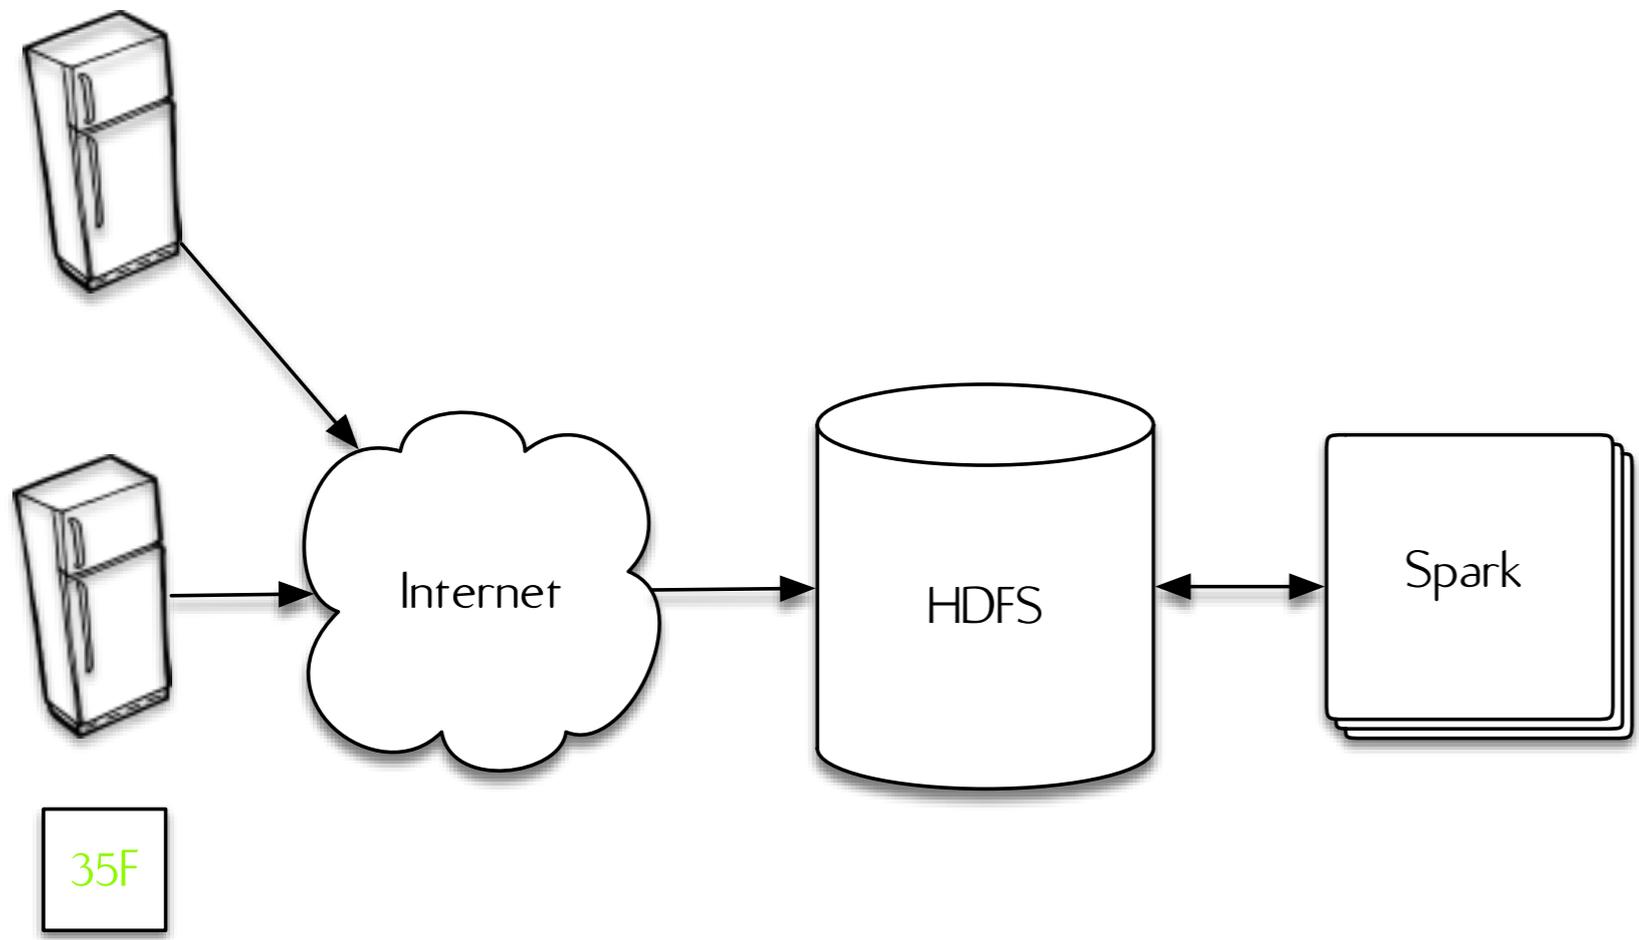




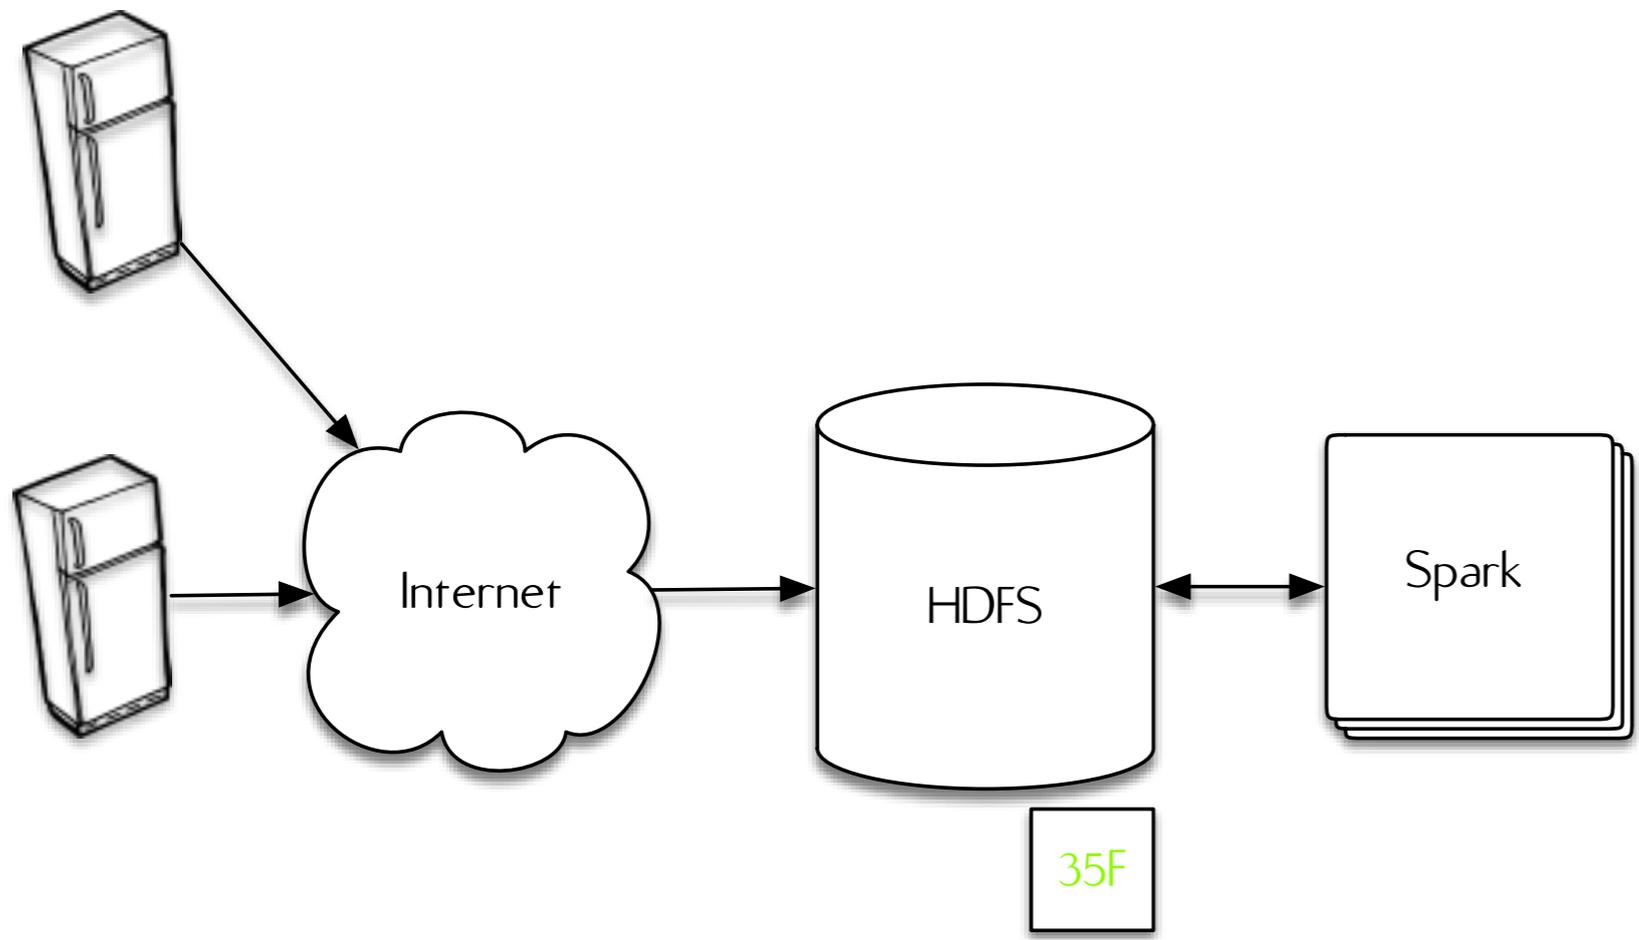
Hospital Refrigerators

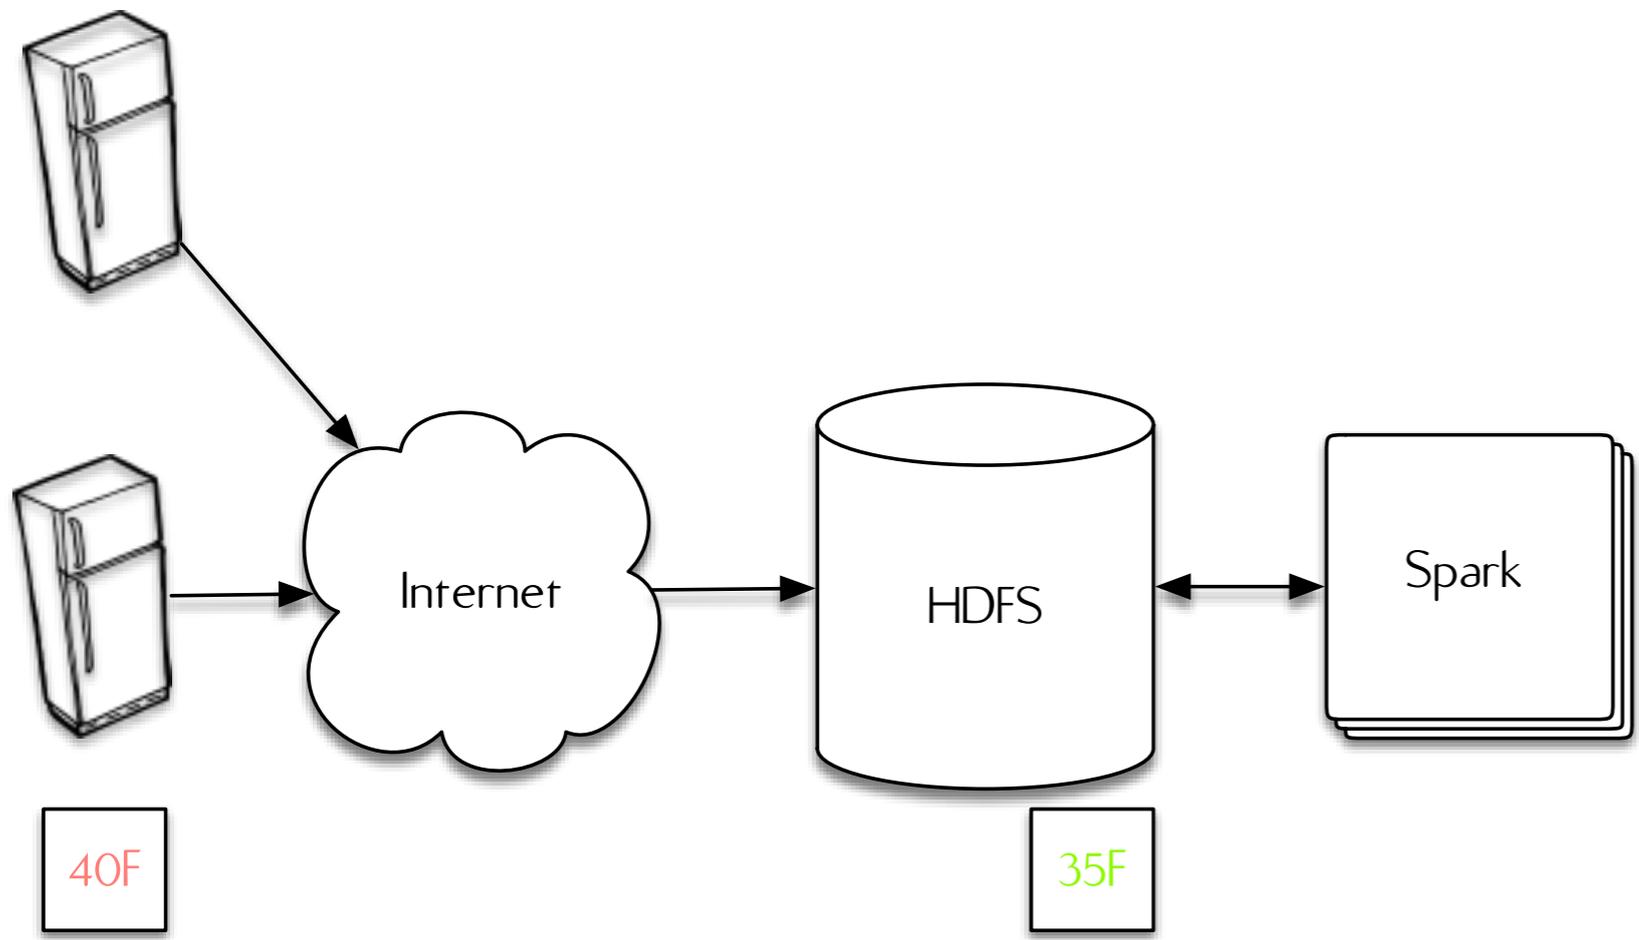
Ideal Execution

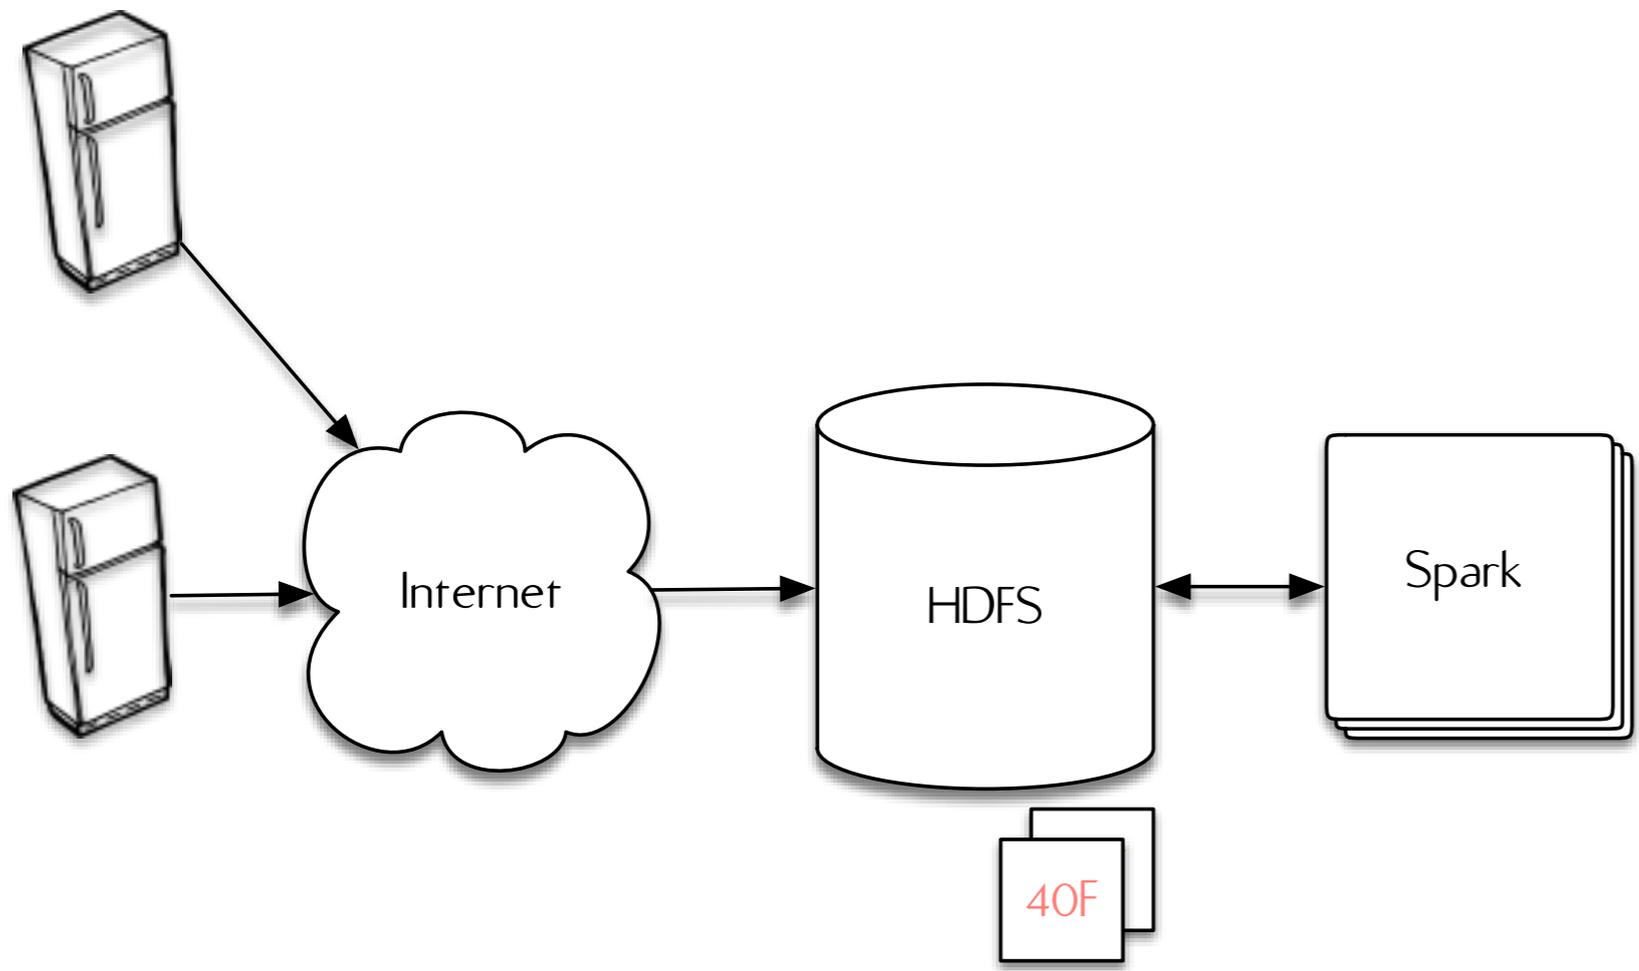


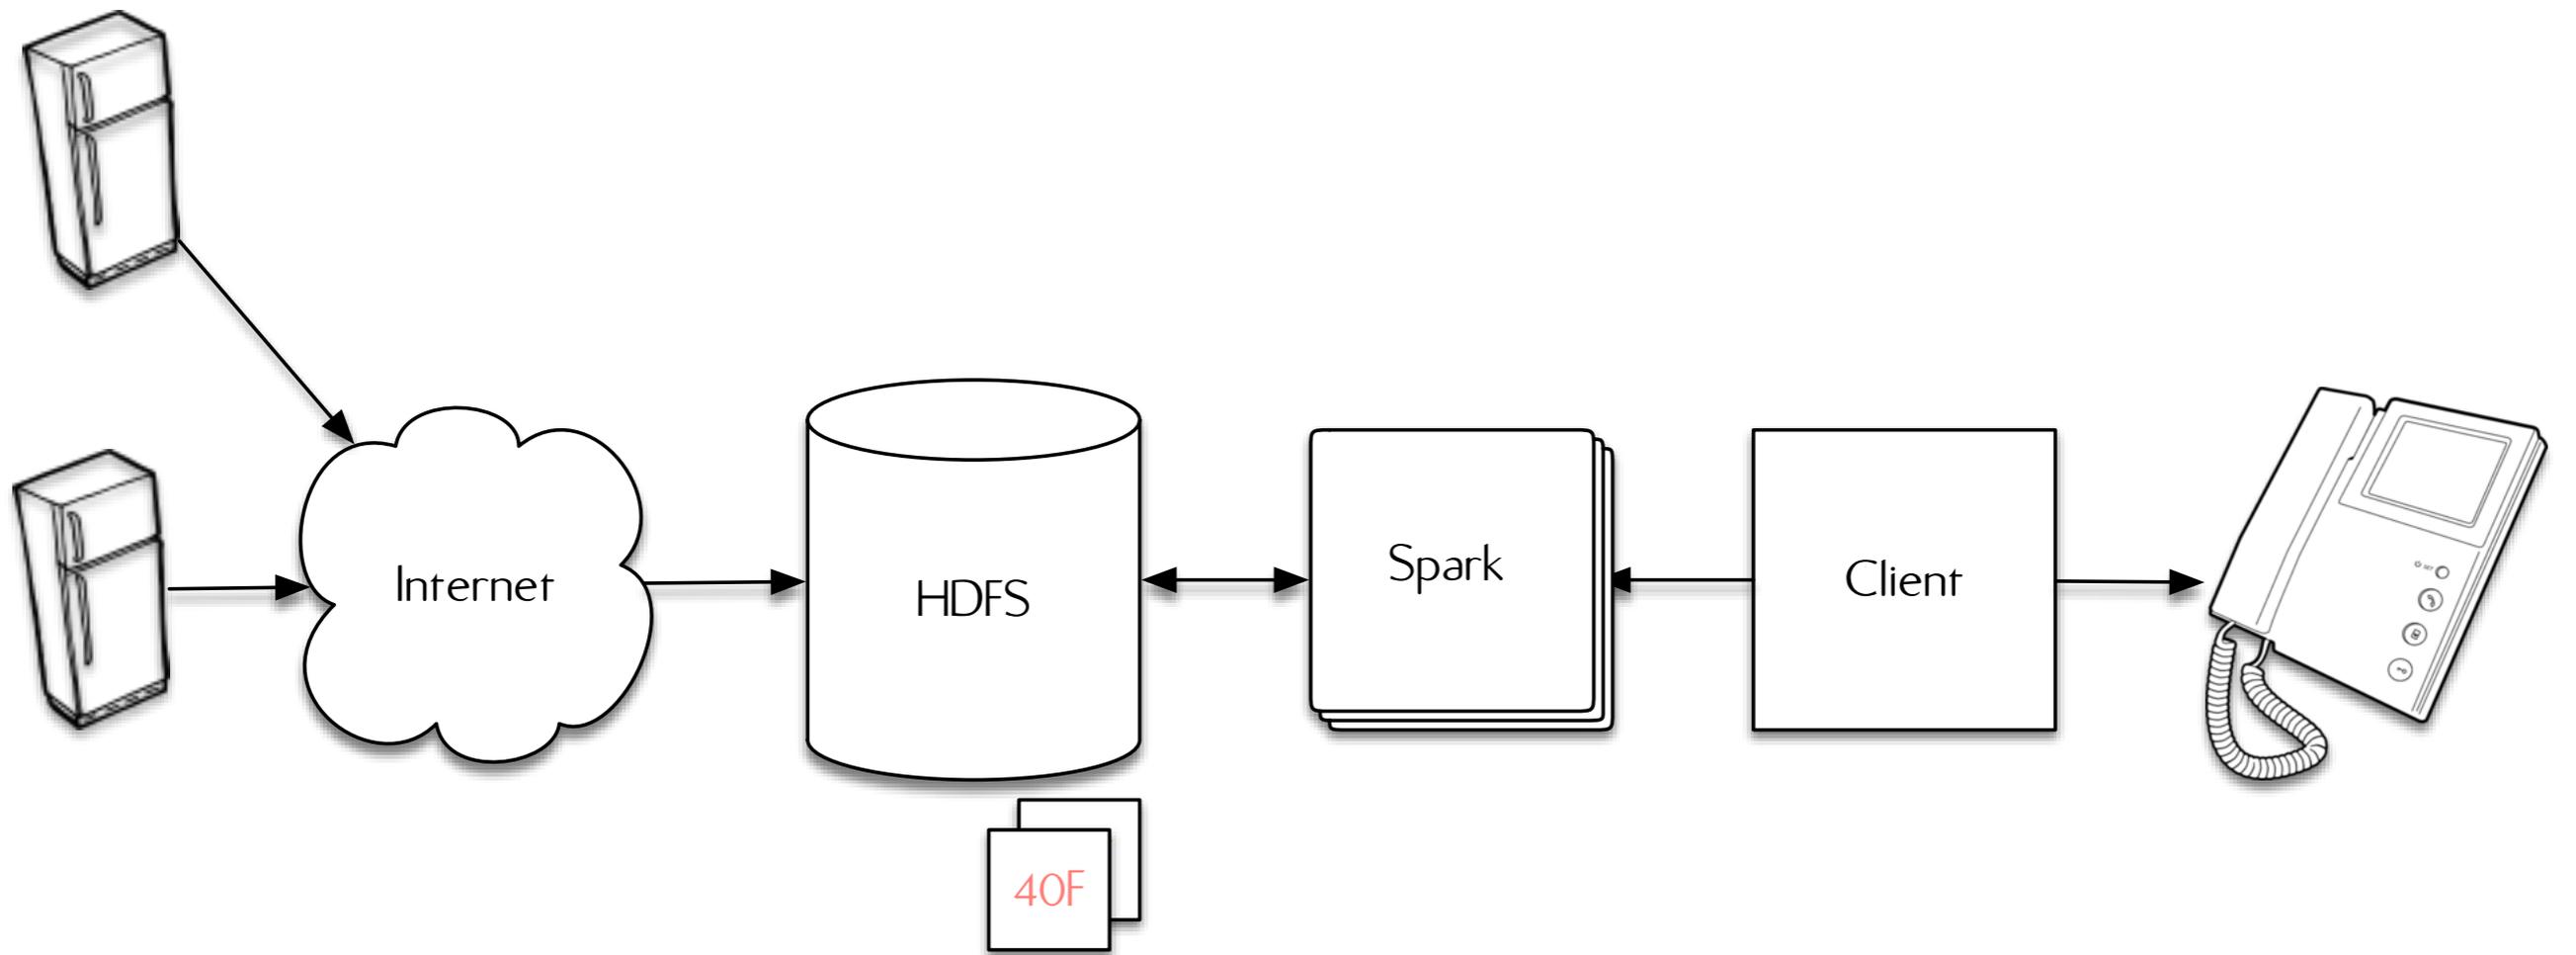


35F

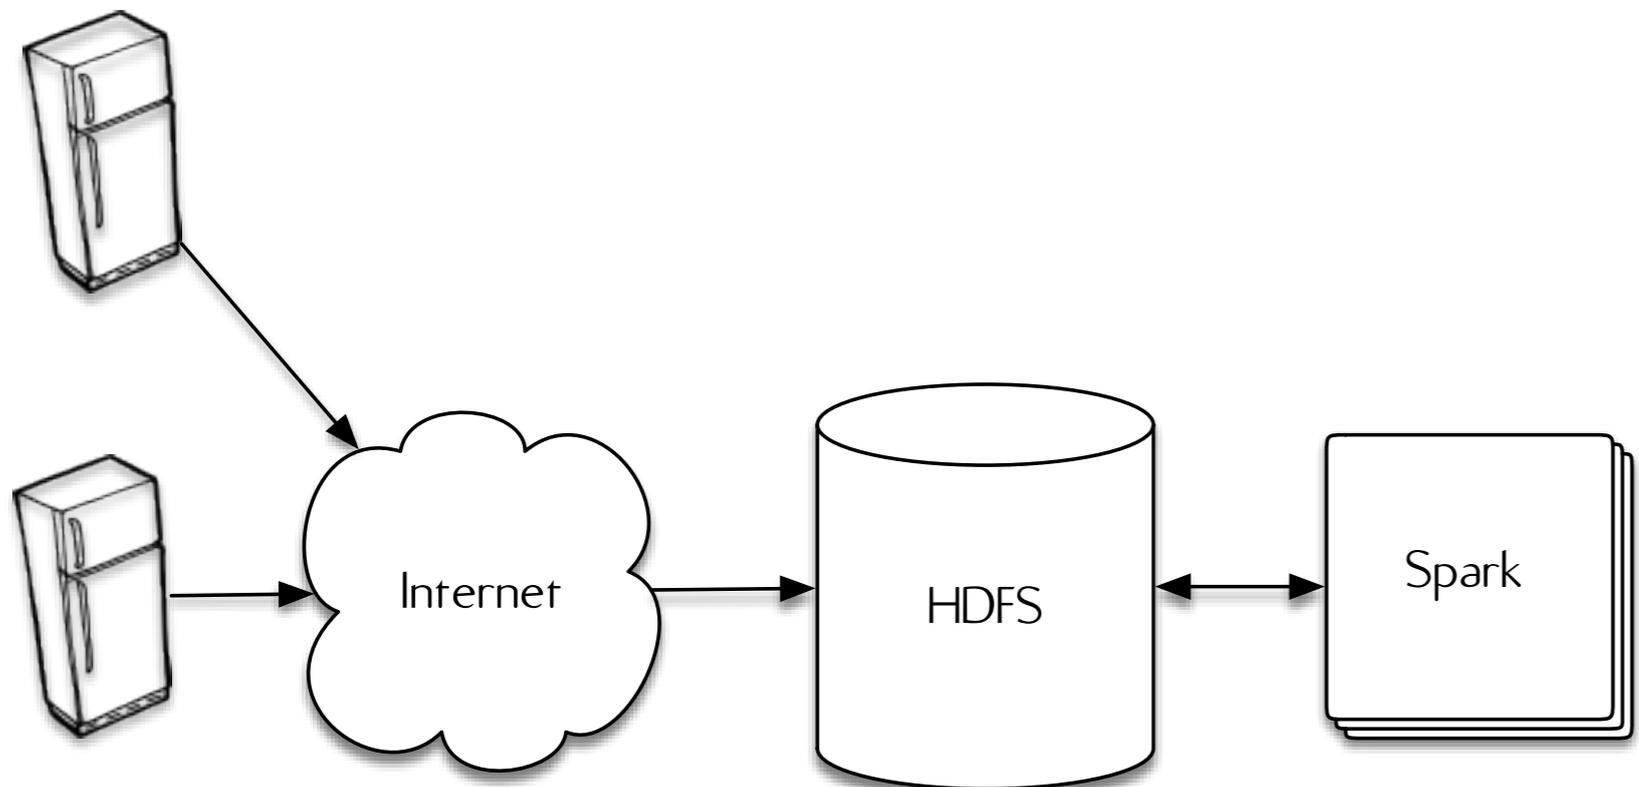


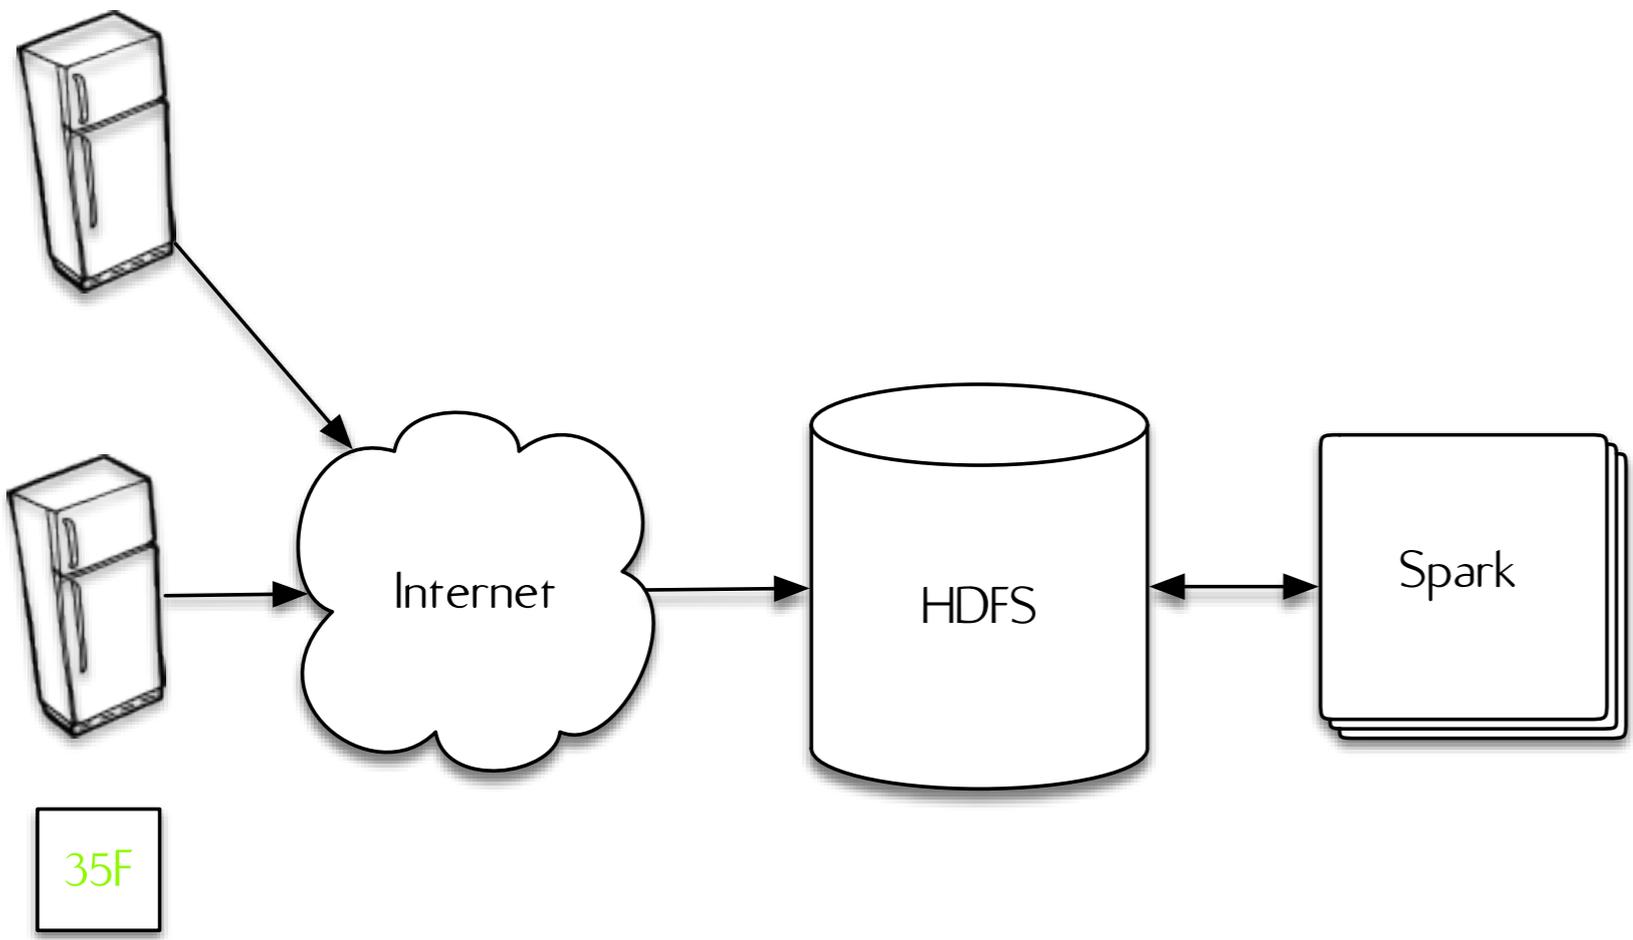




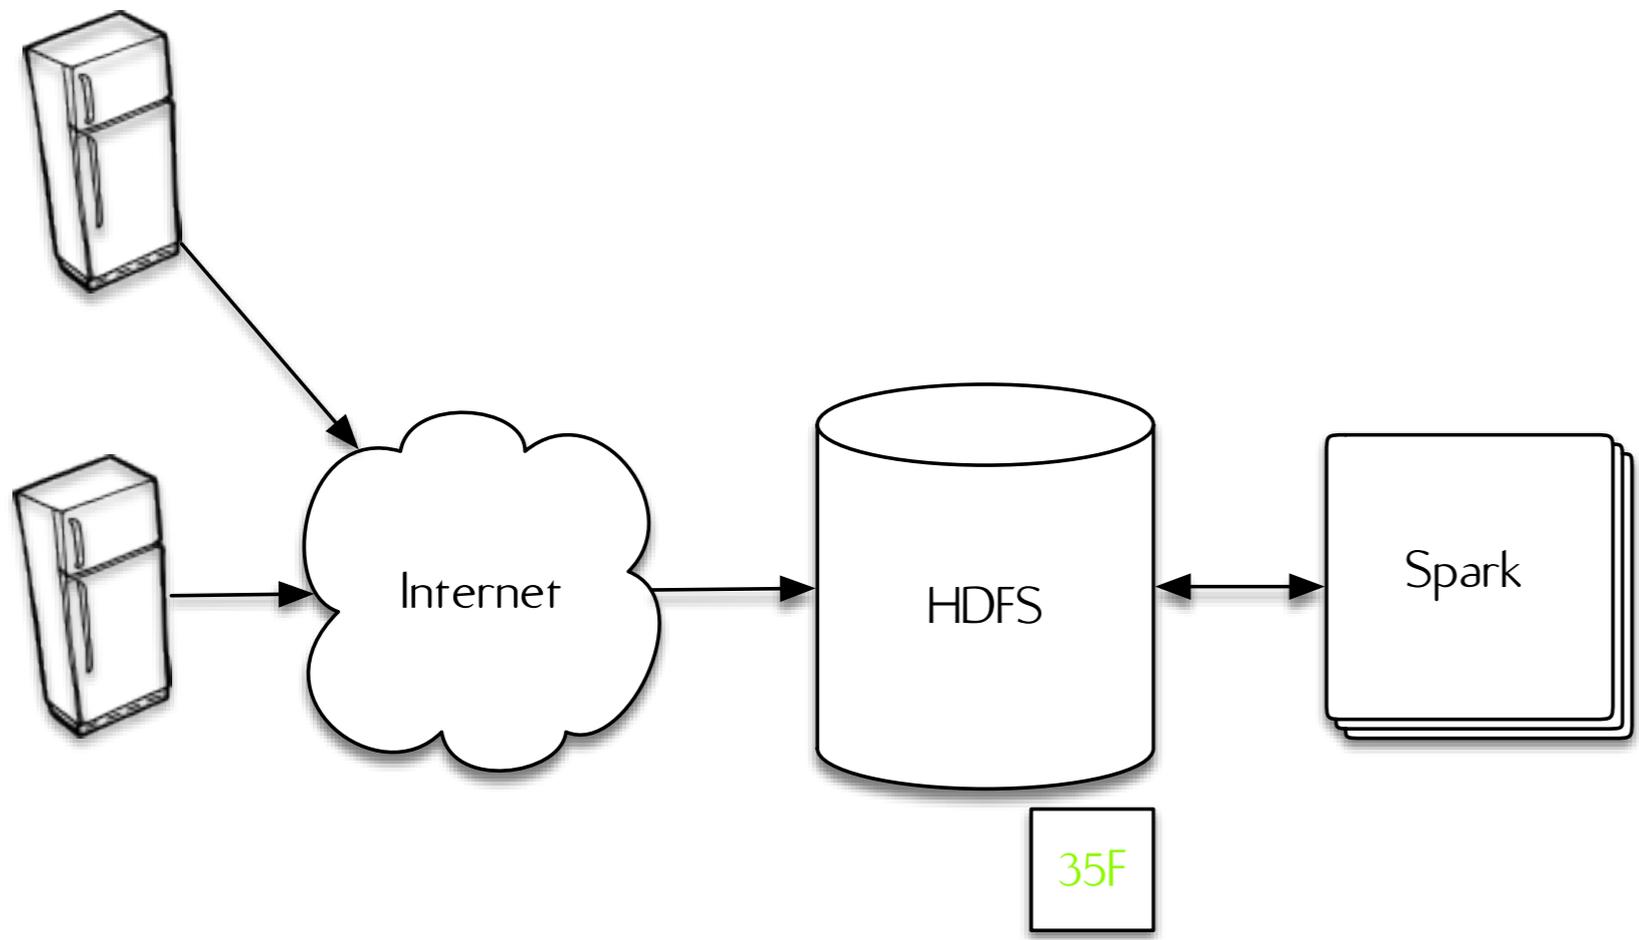


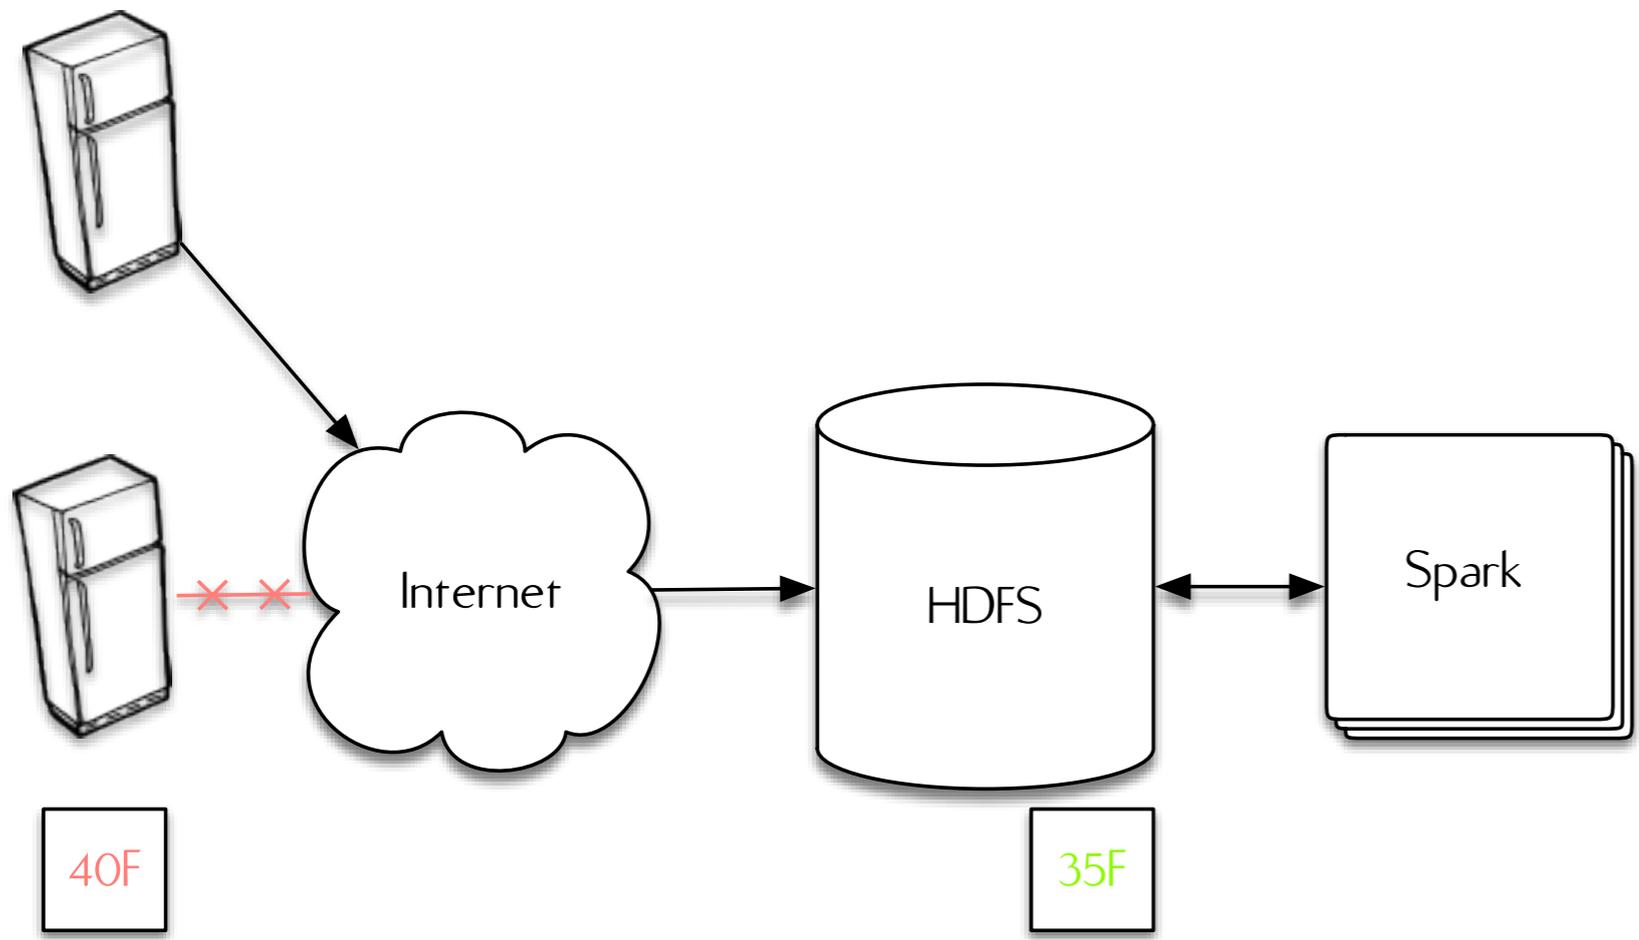
Problem Connectivity

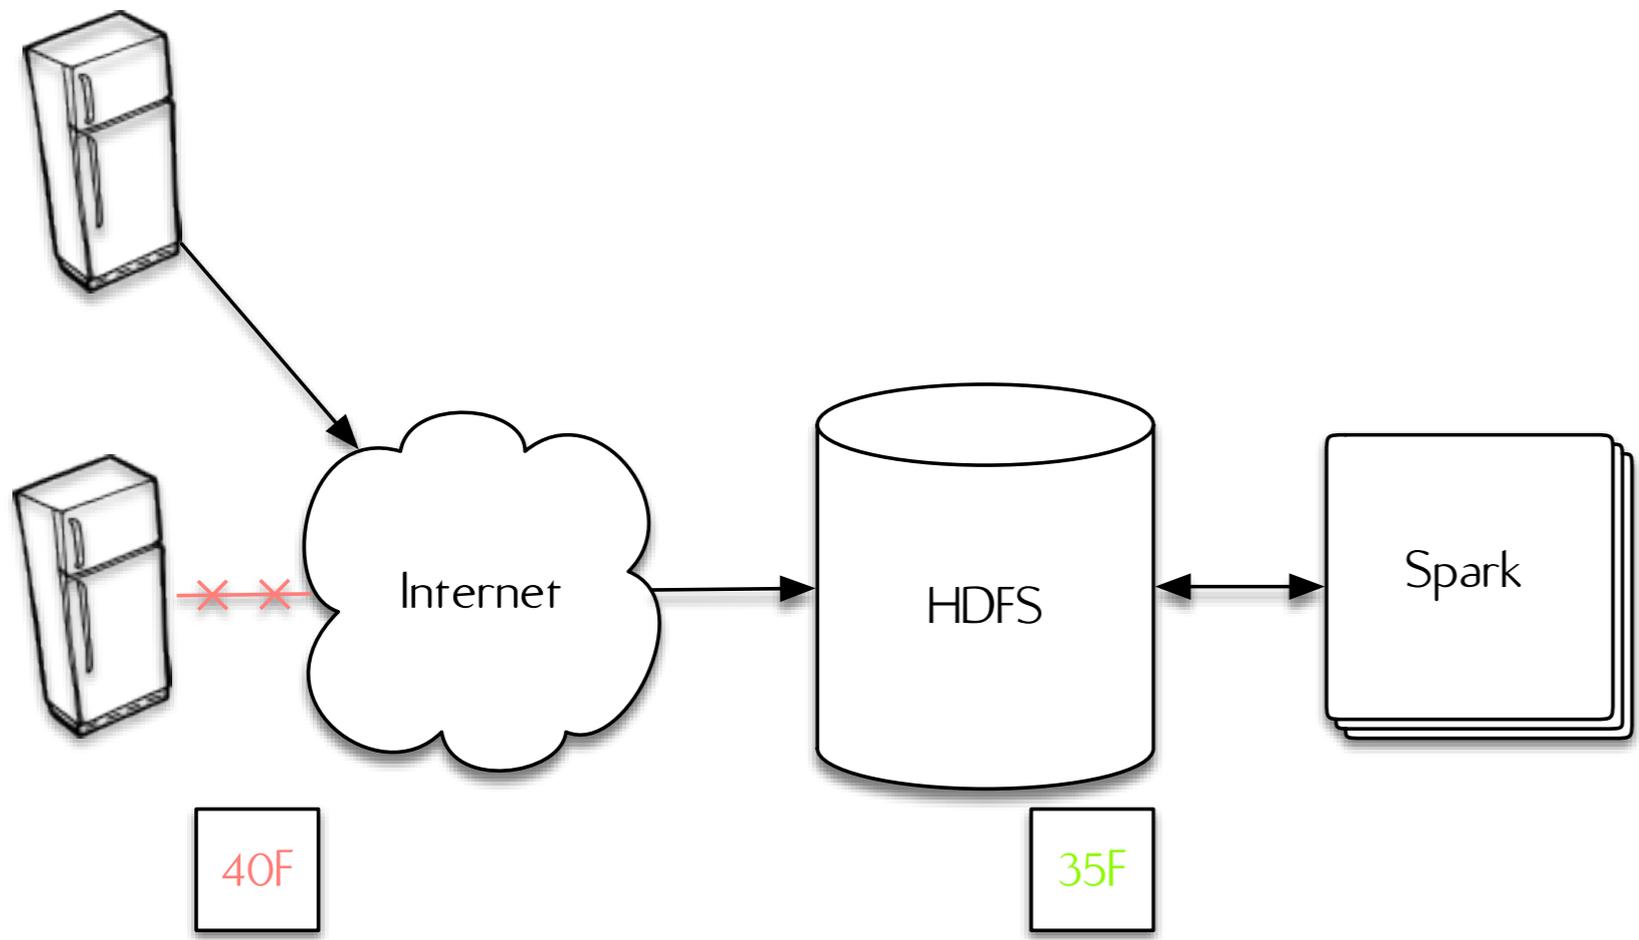


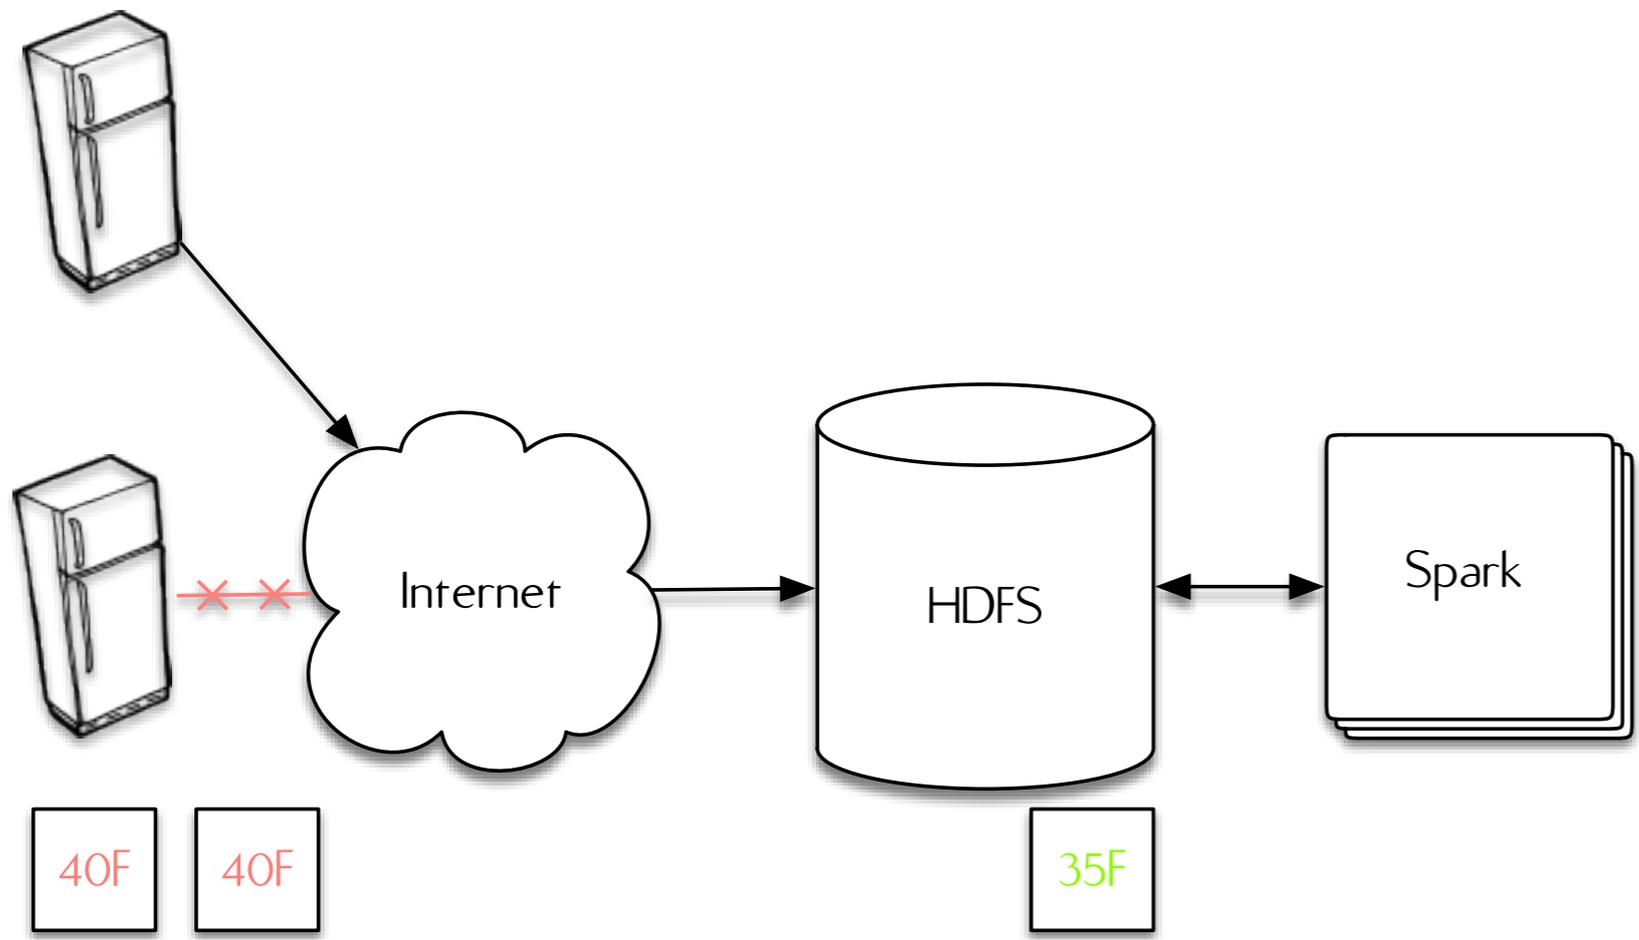


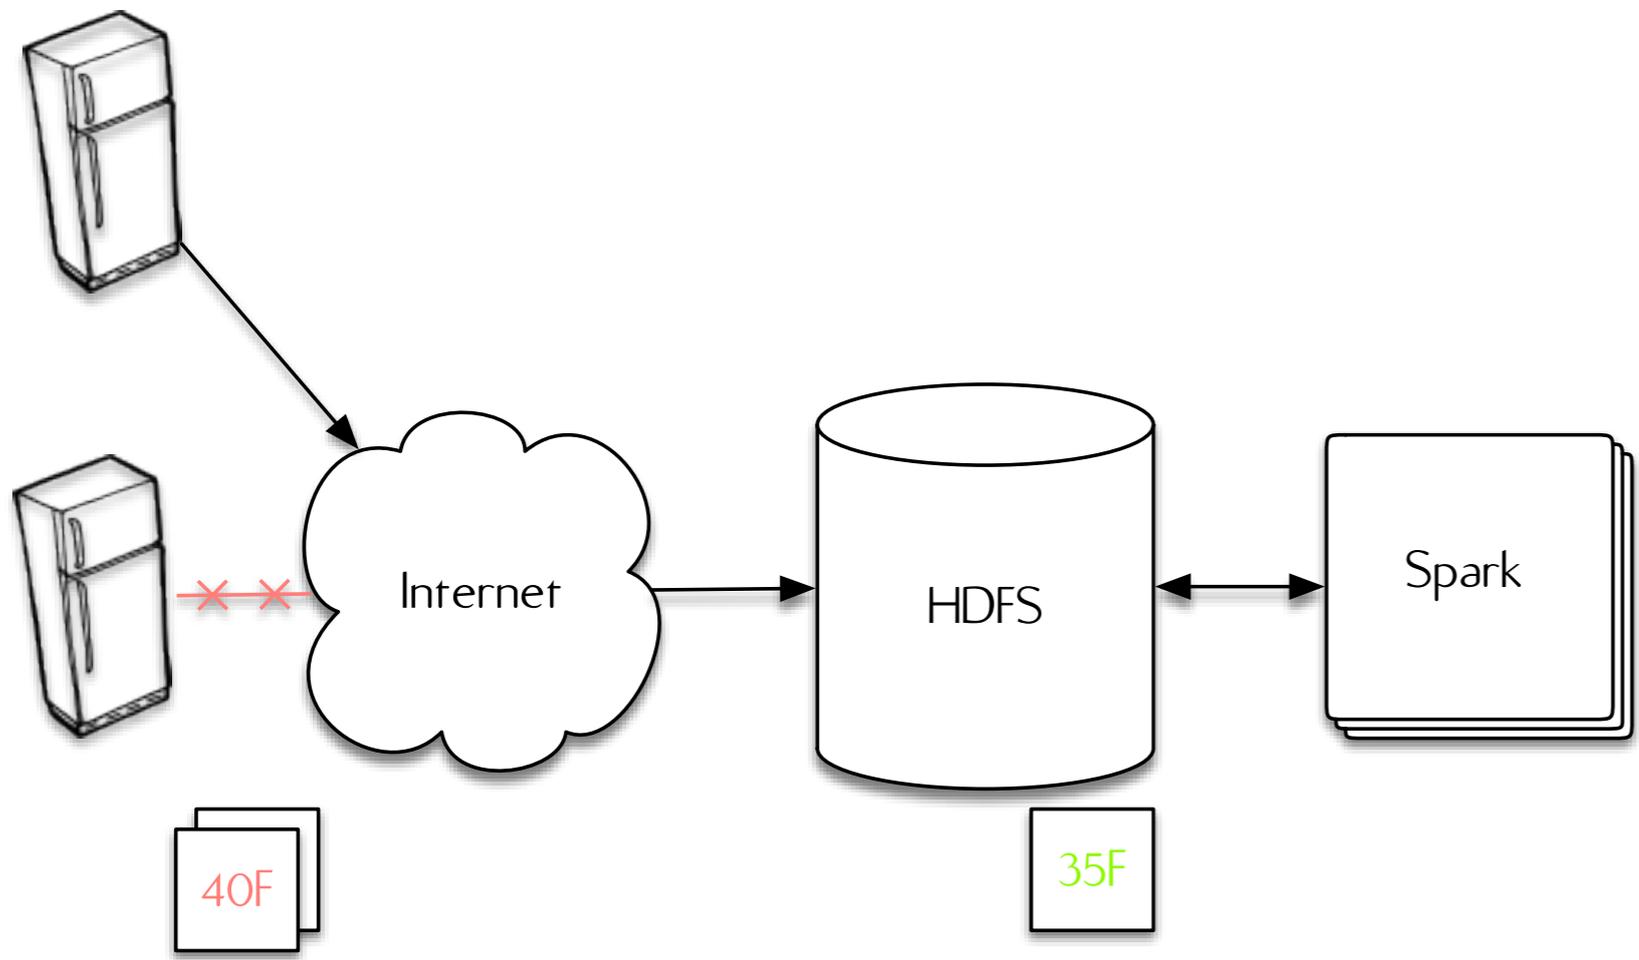
35F





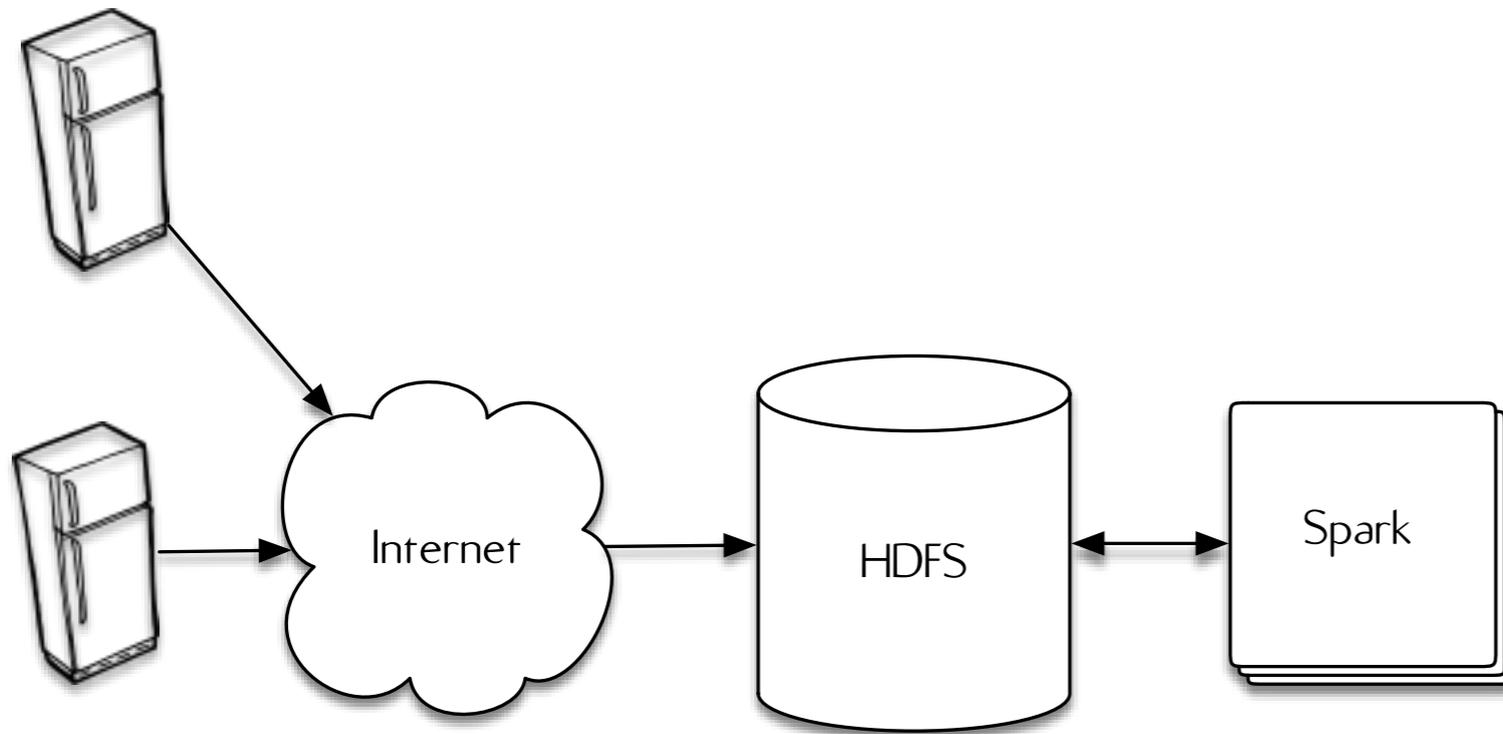


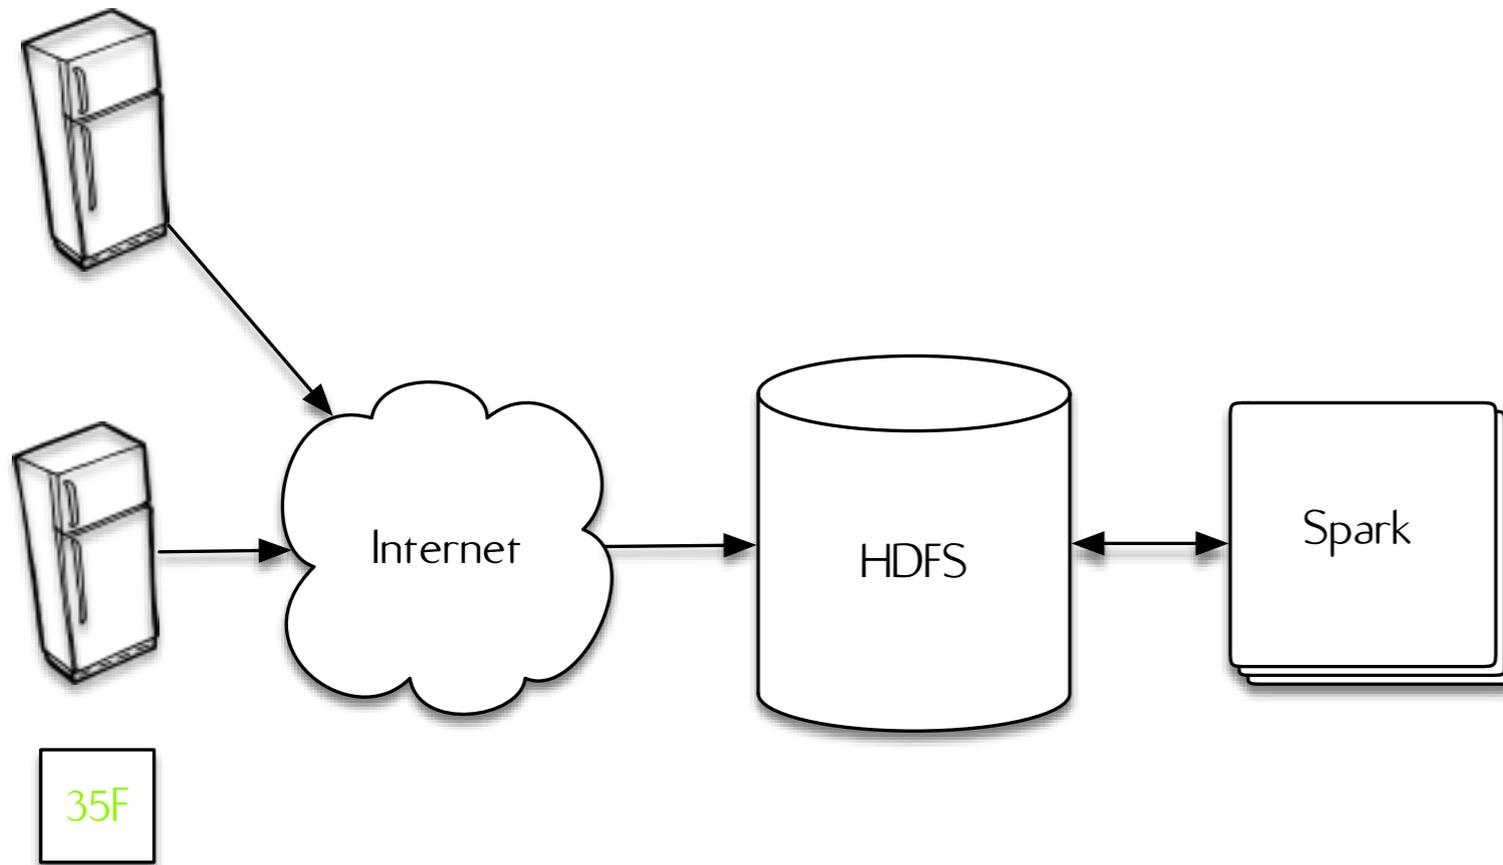


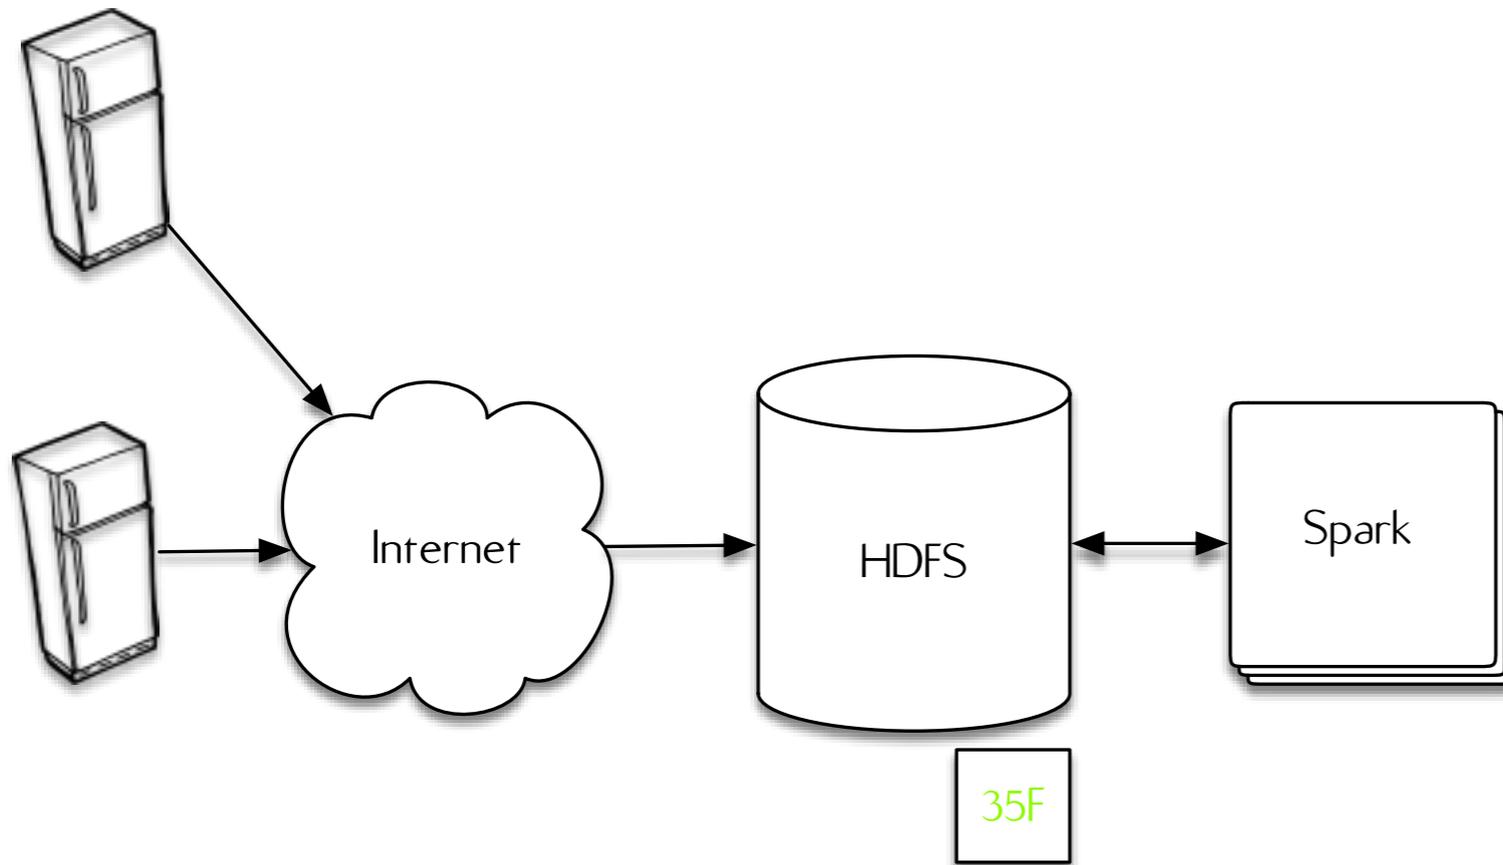


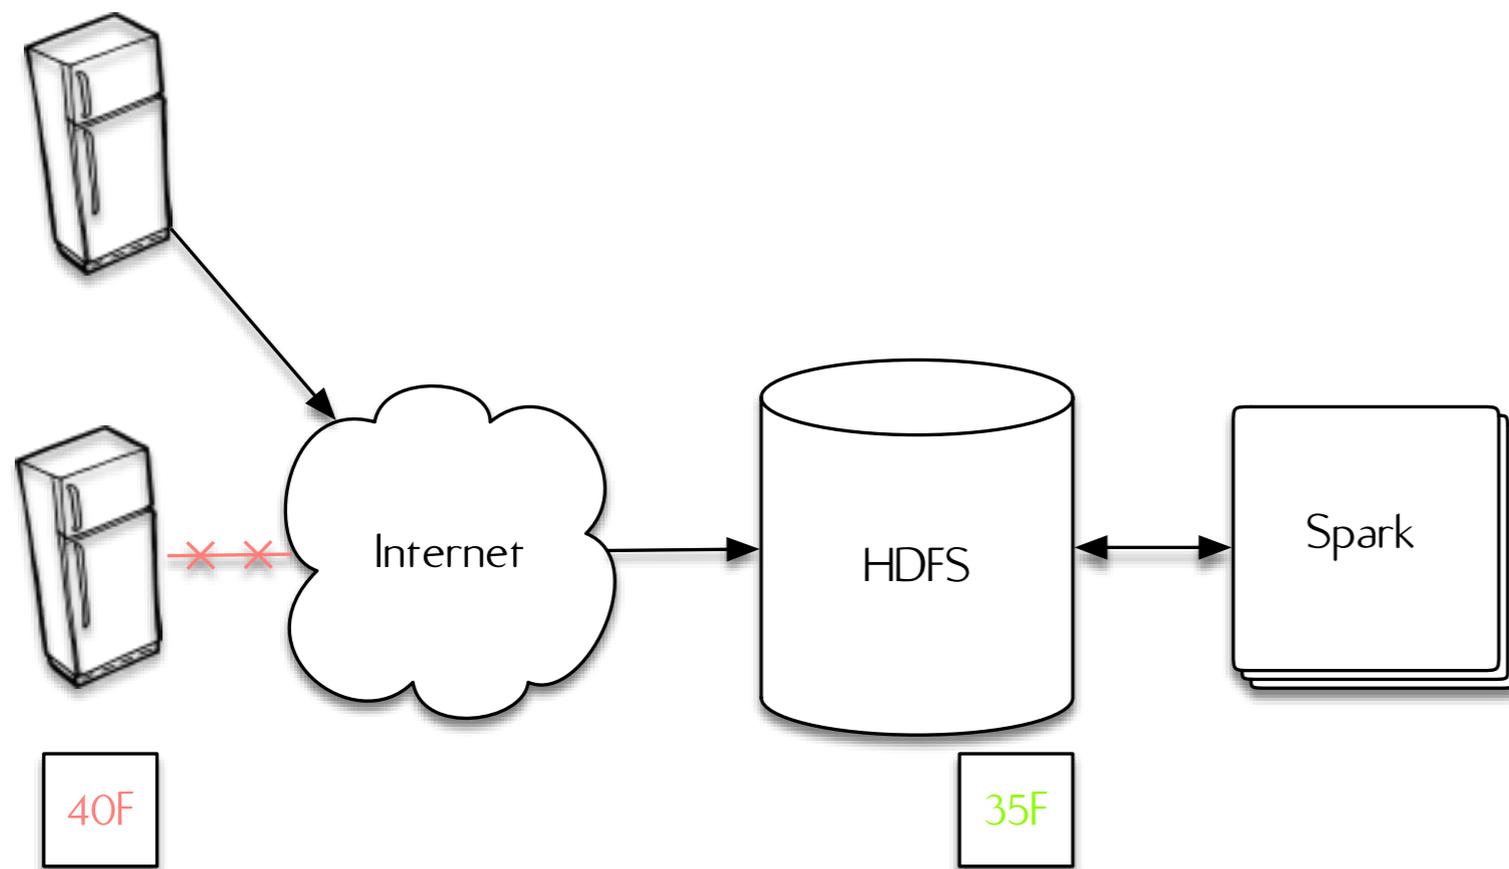
Solution

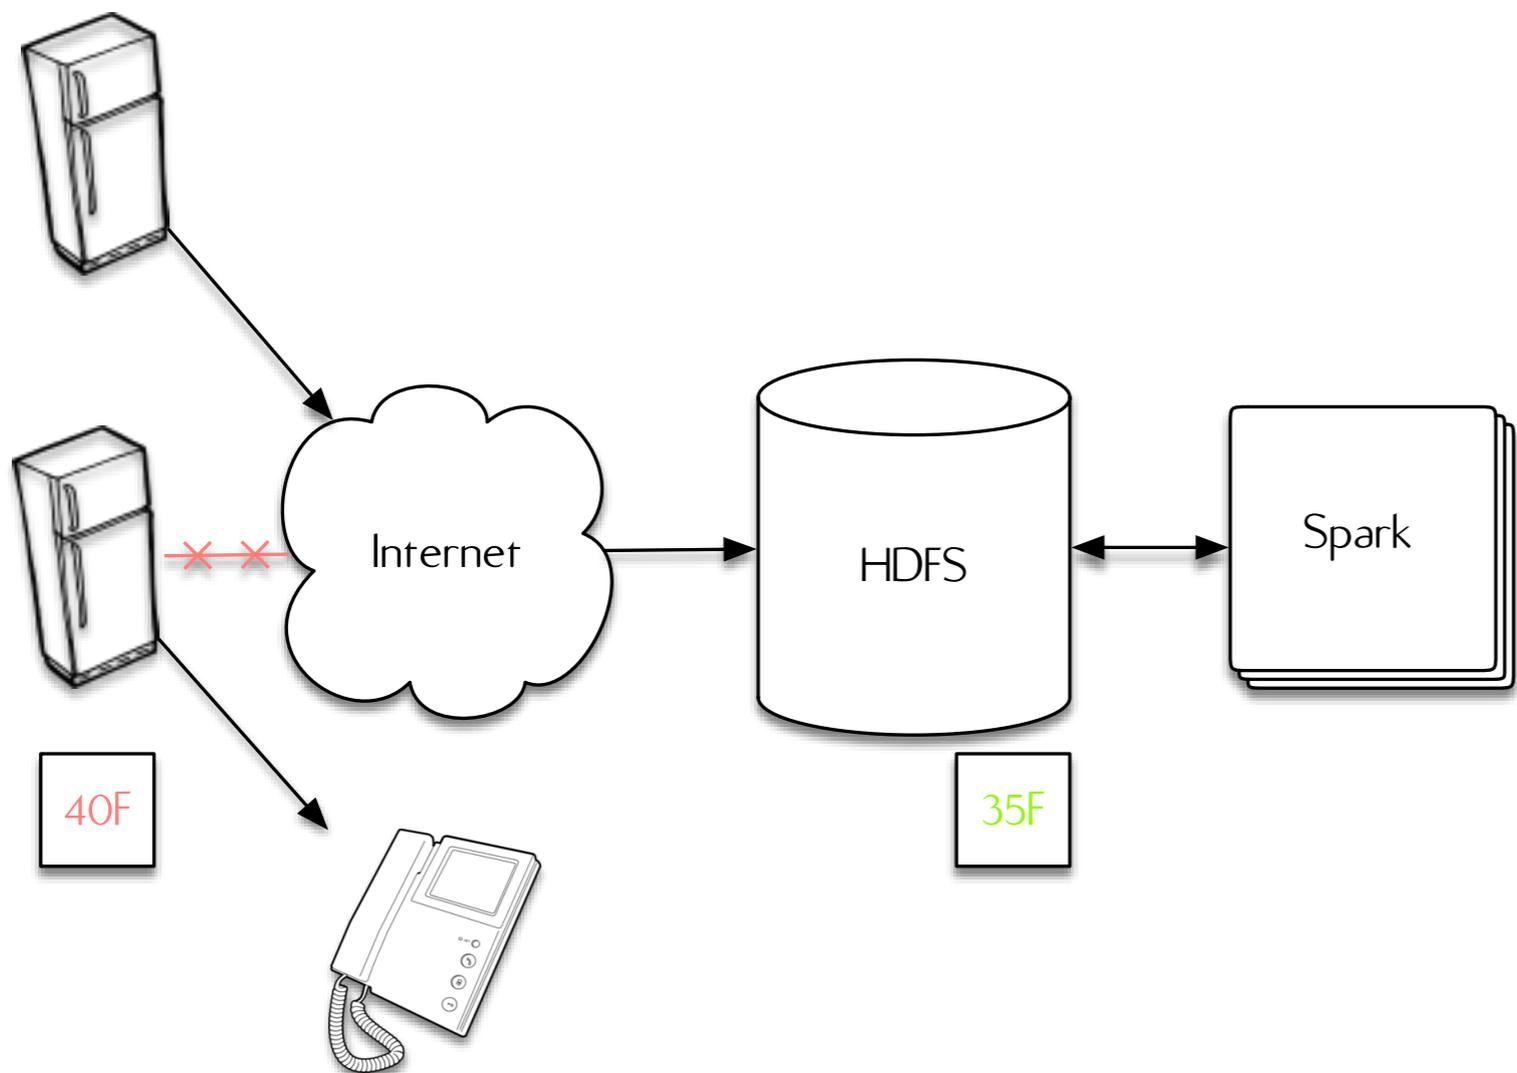
Local Decisions



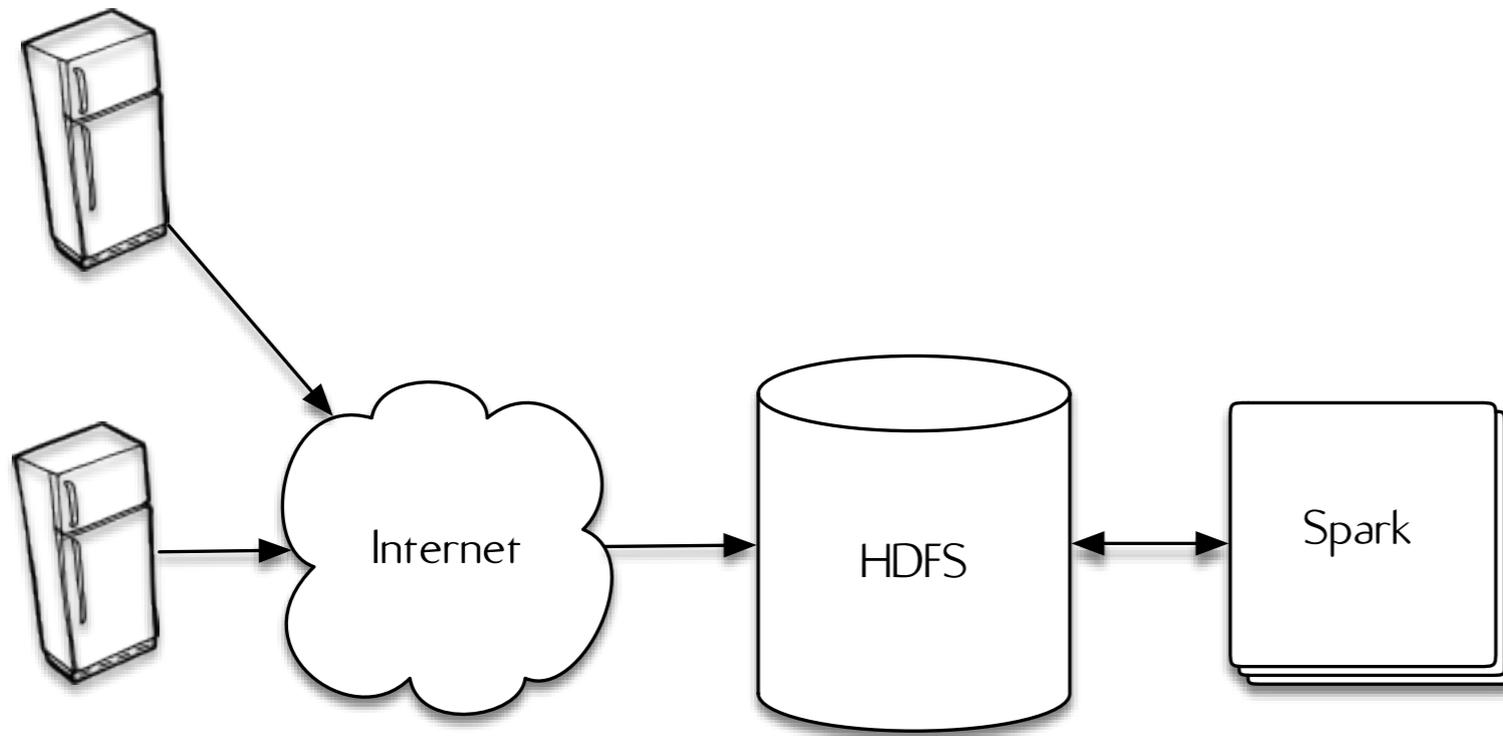


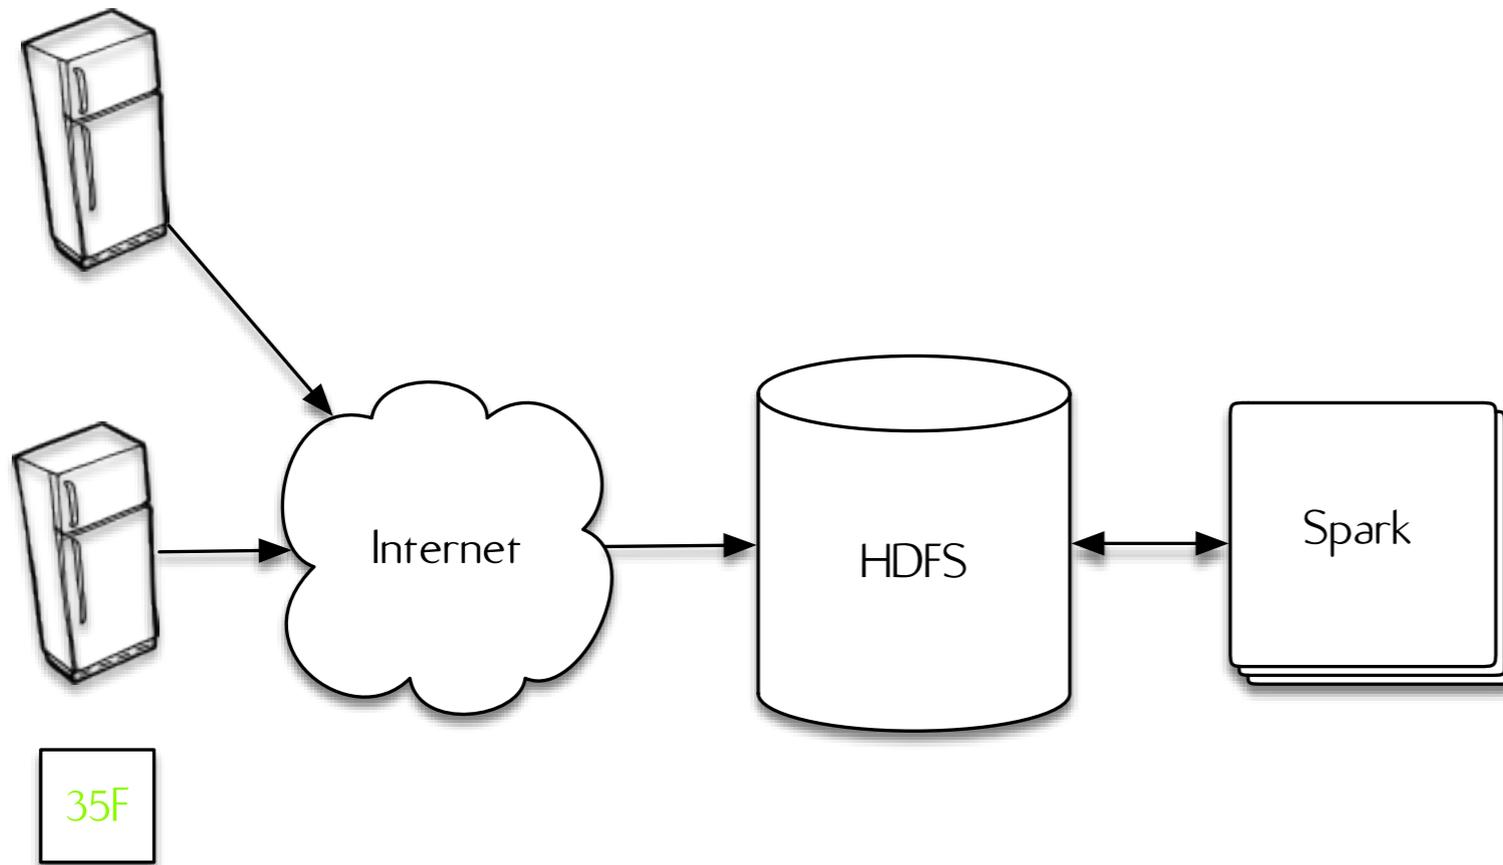


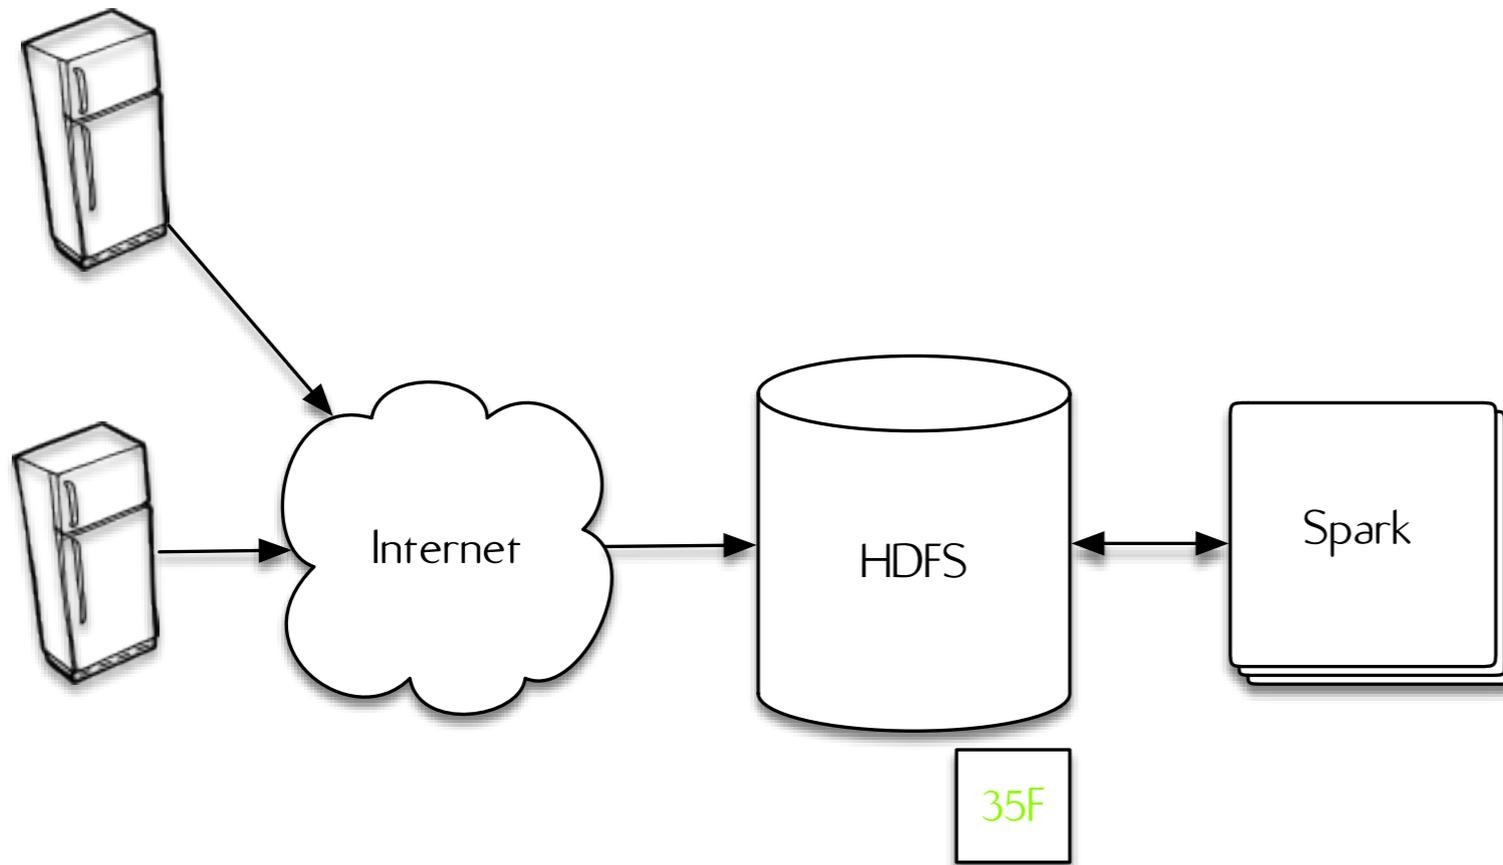


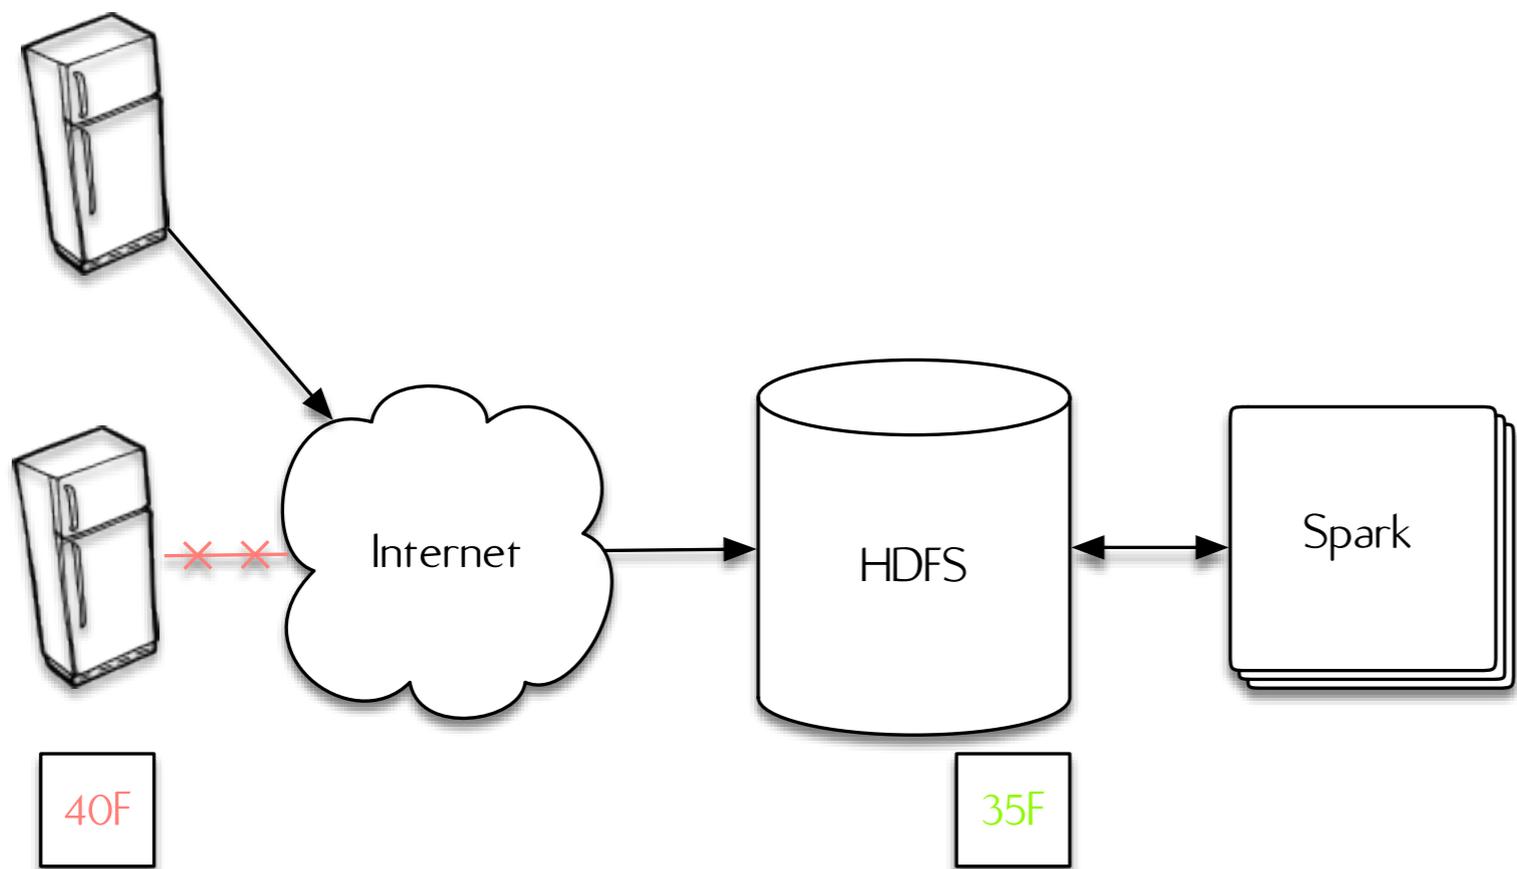


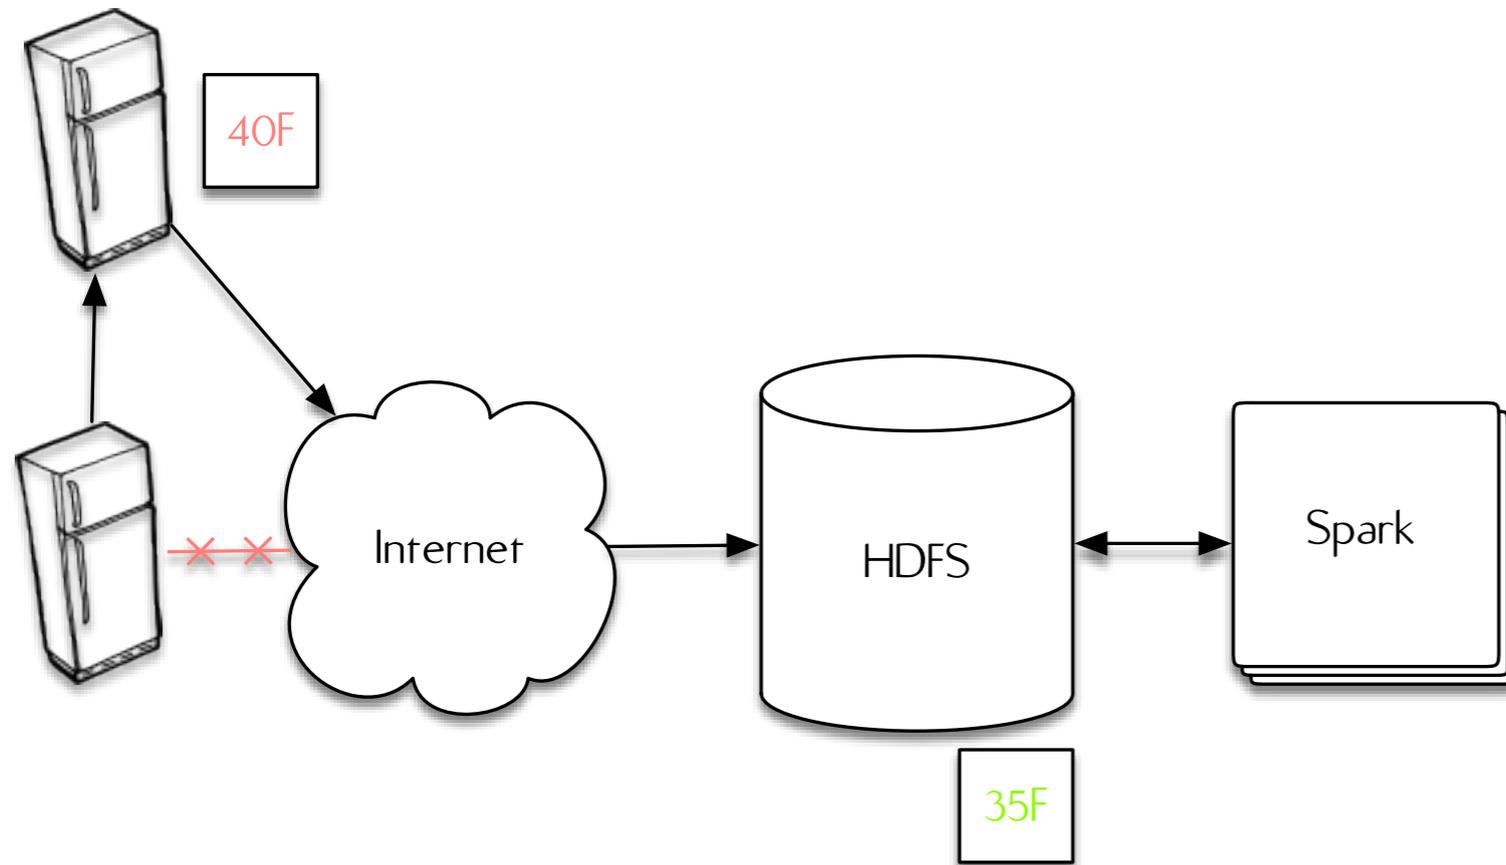
Solution
Transitive Dissemination

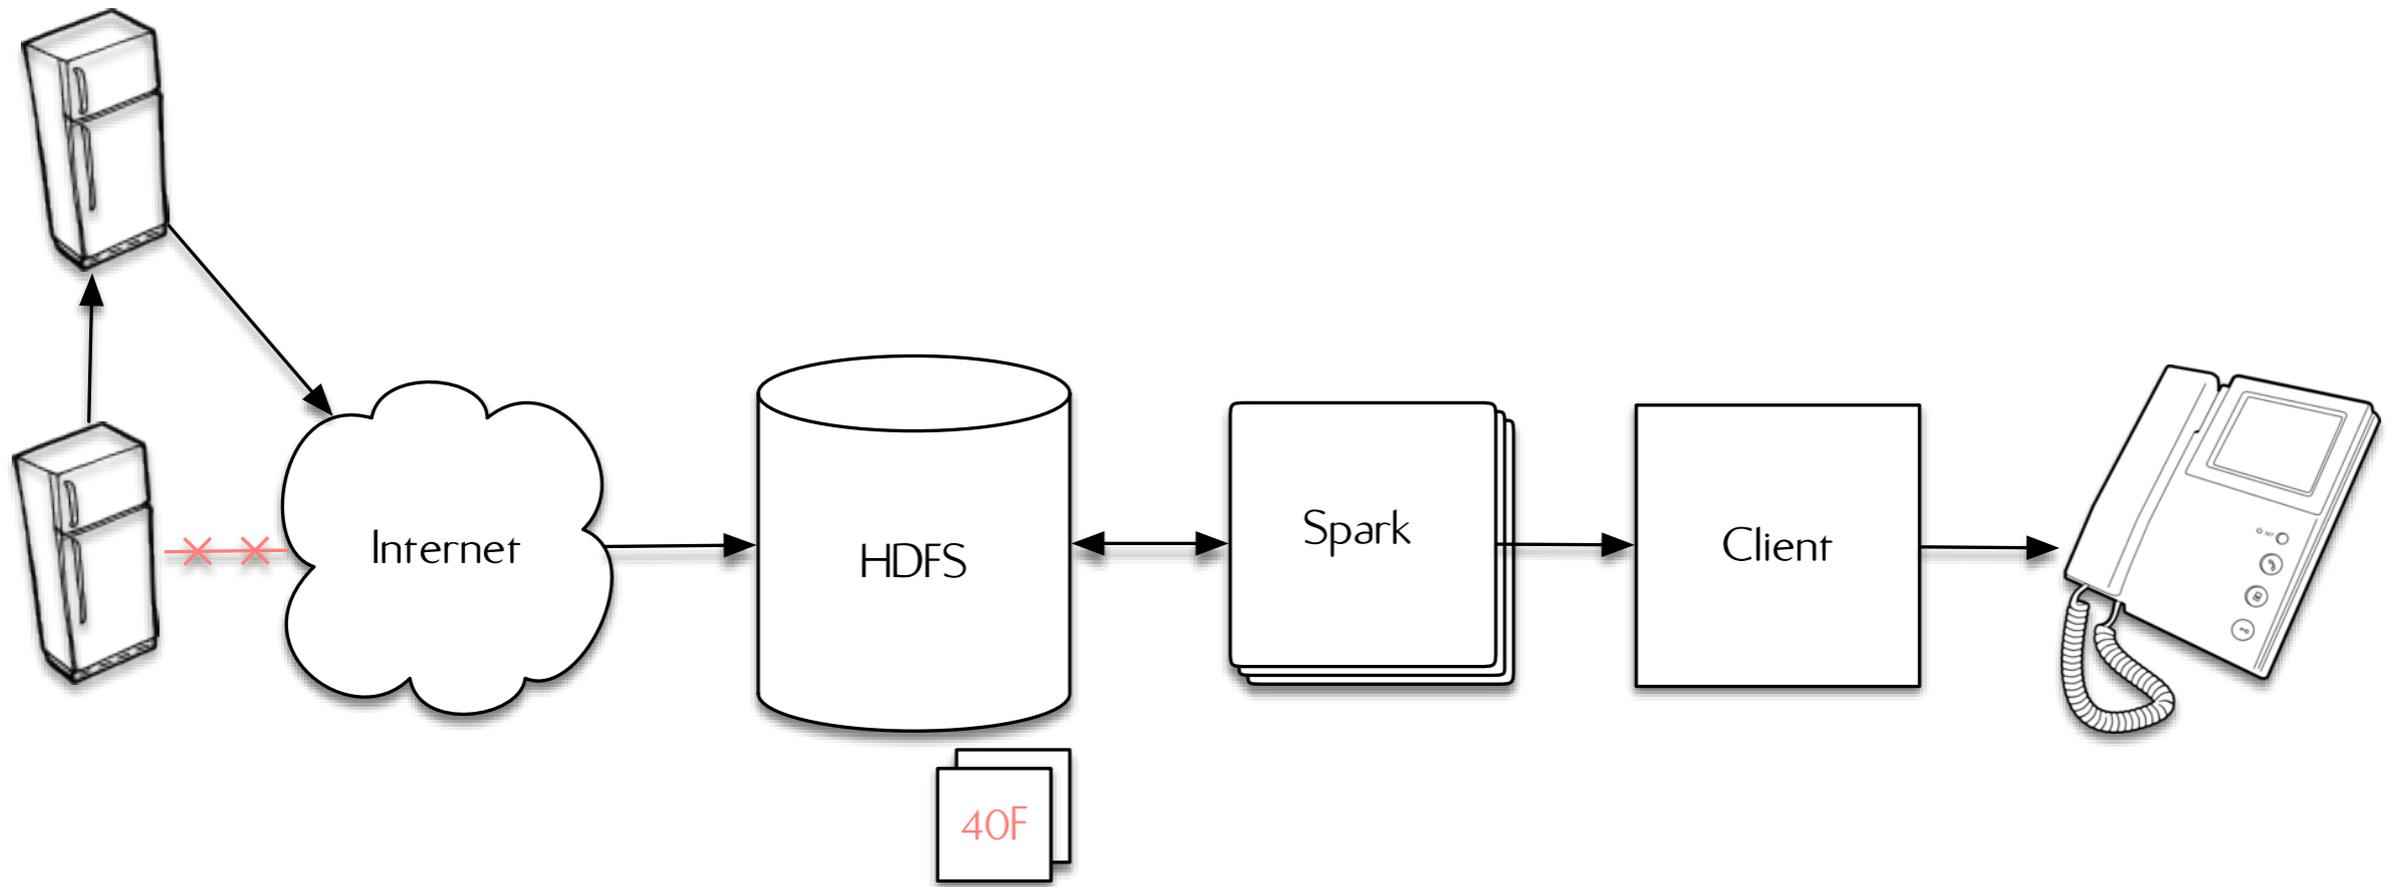




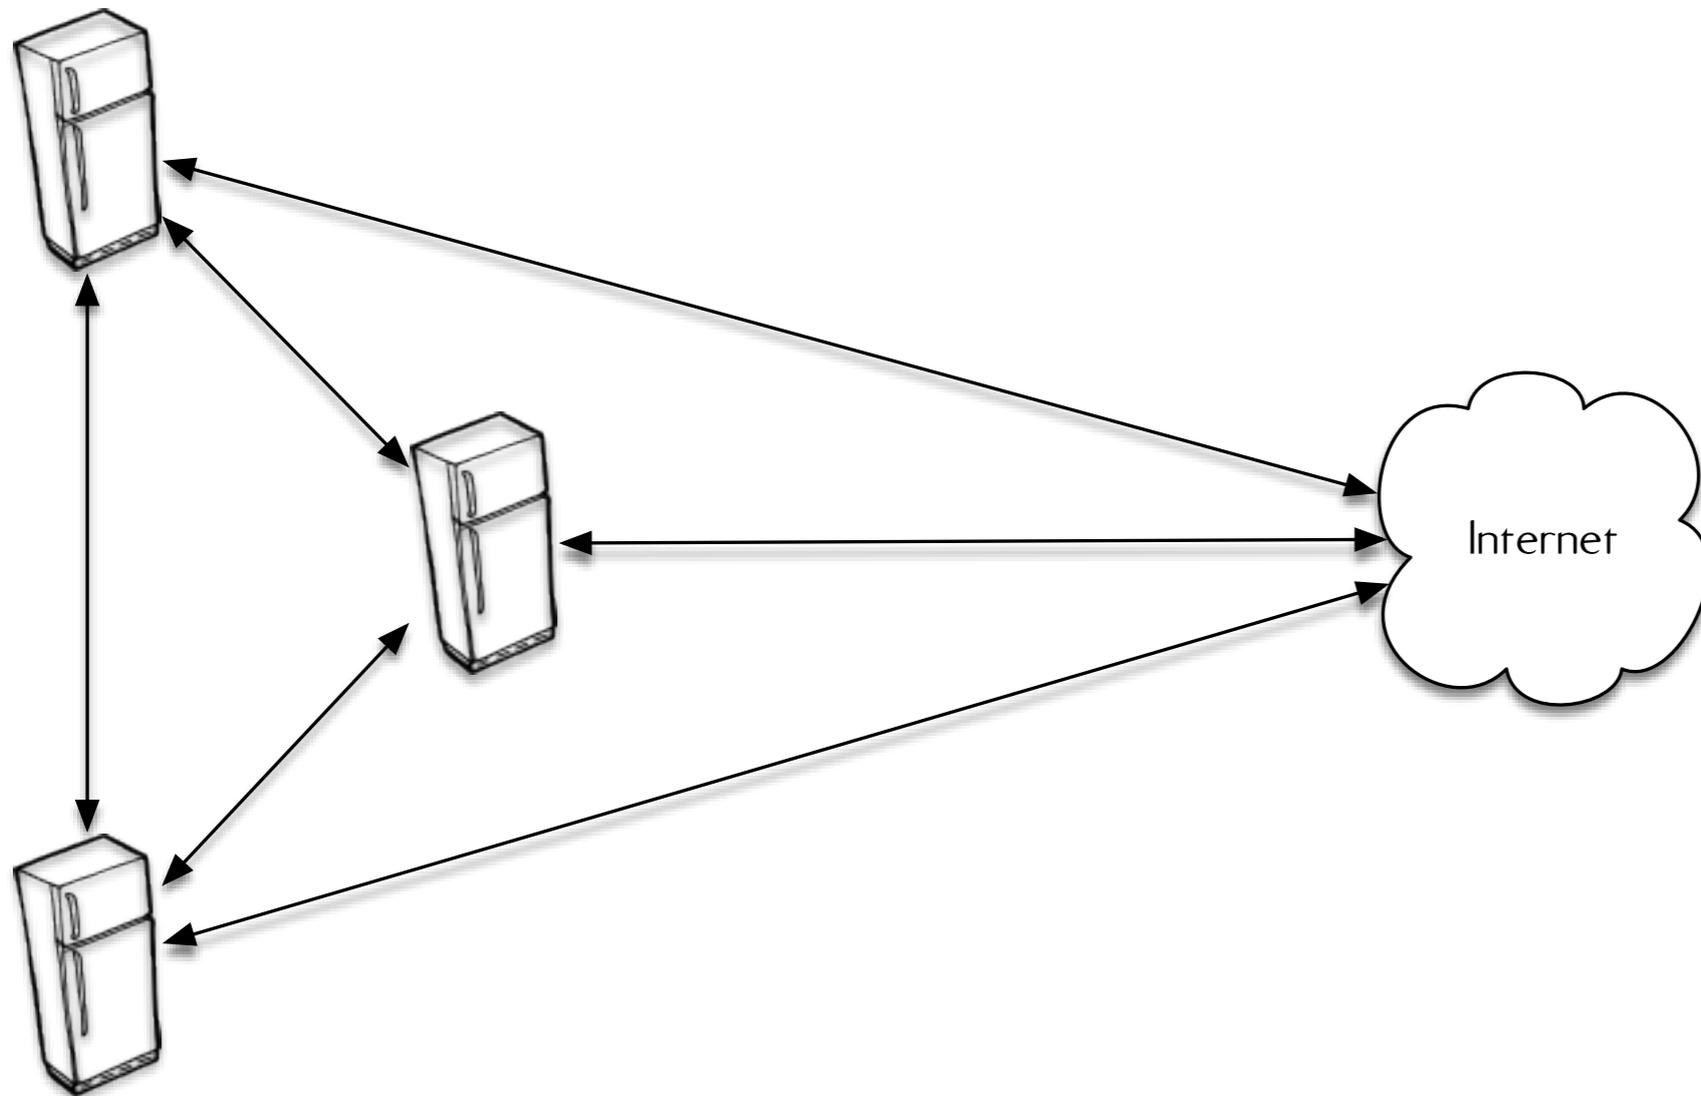


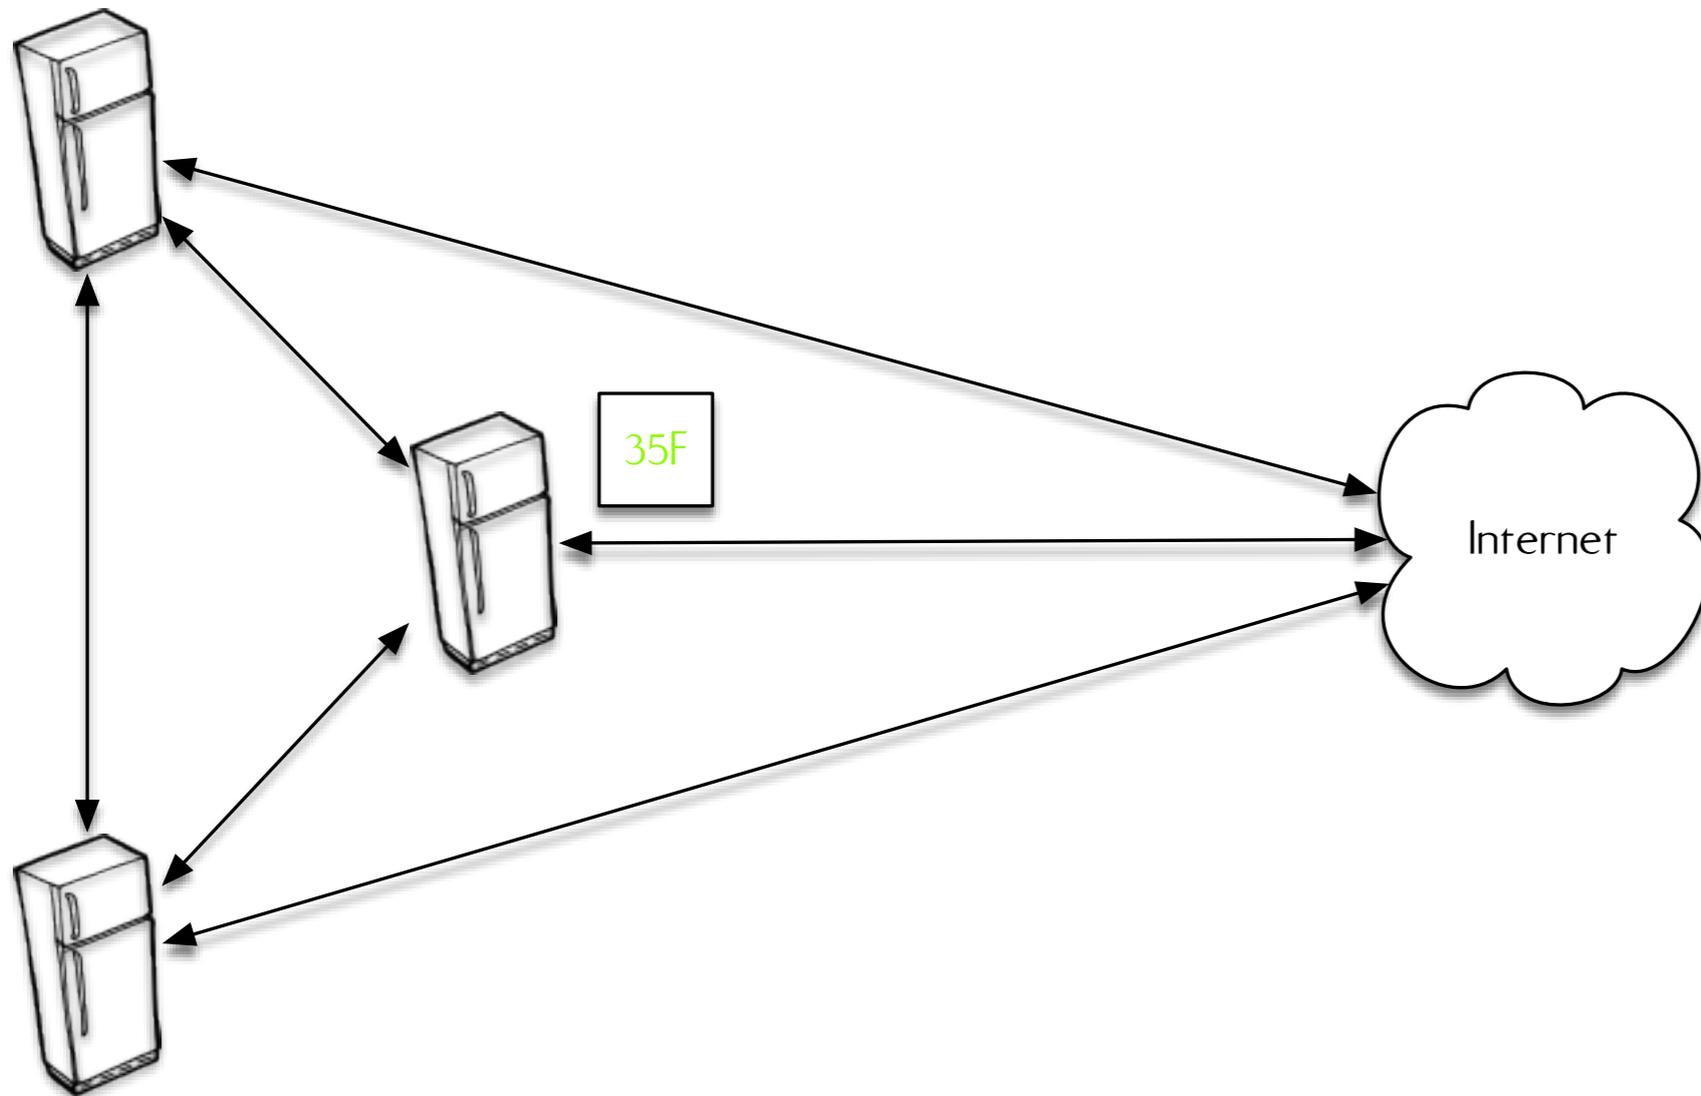


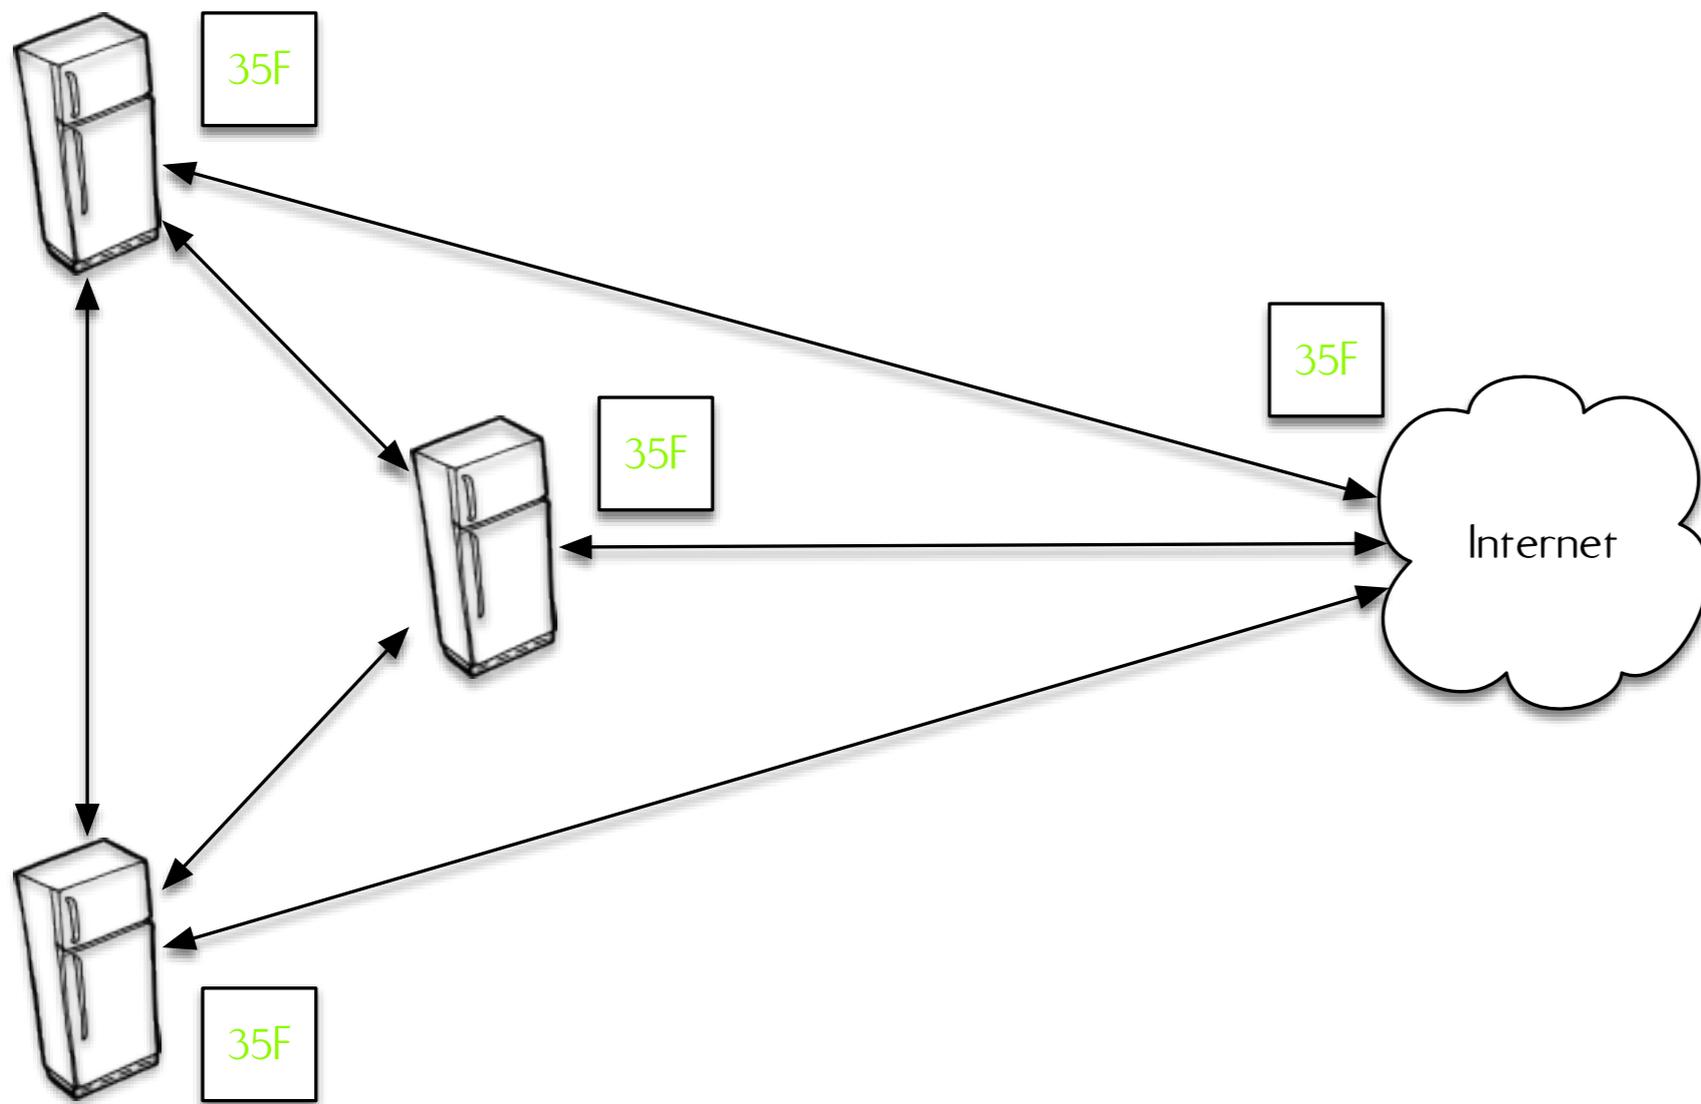


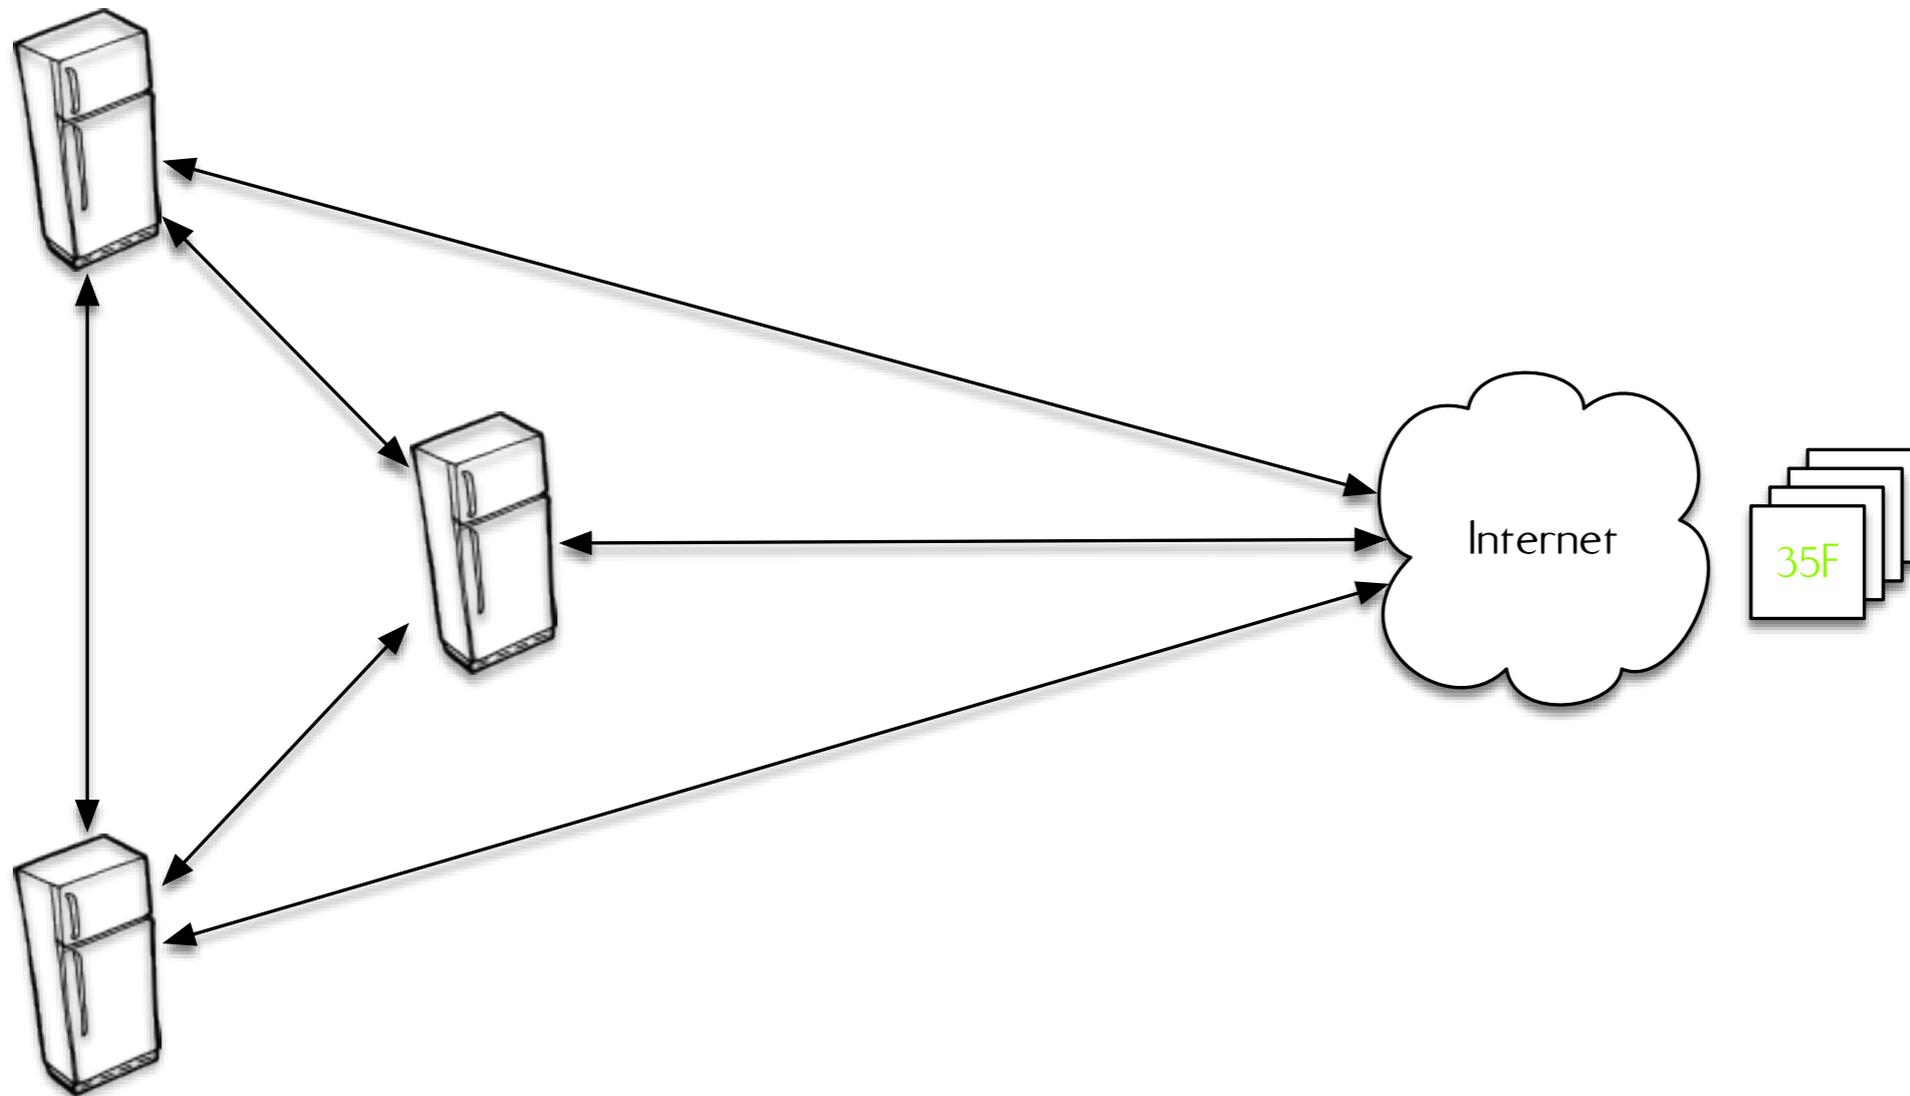


Problem State Transmission

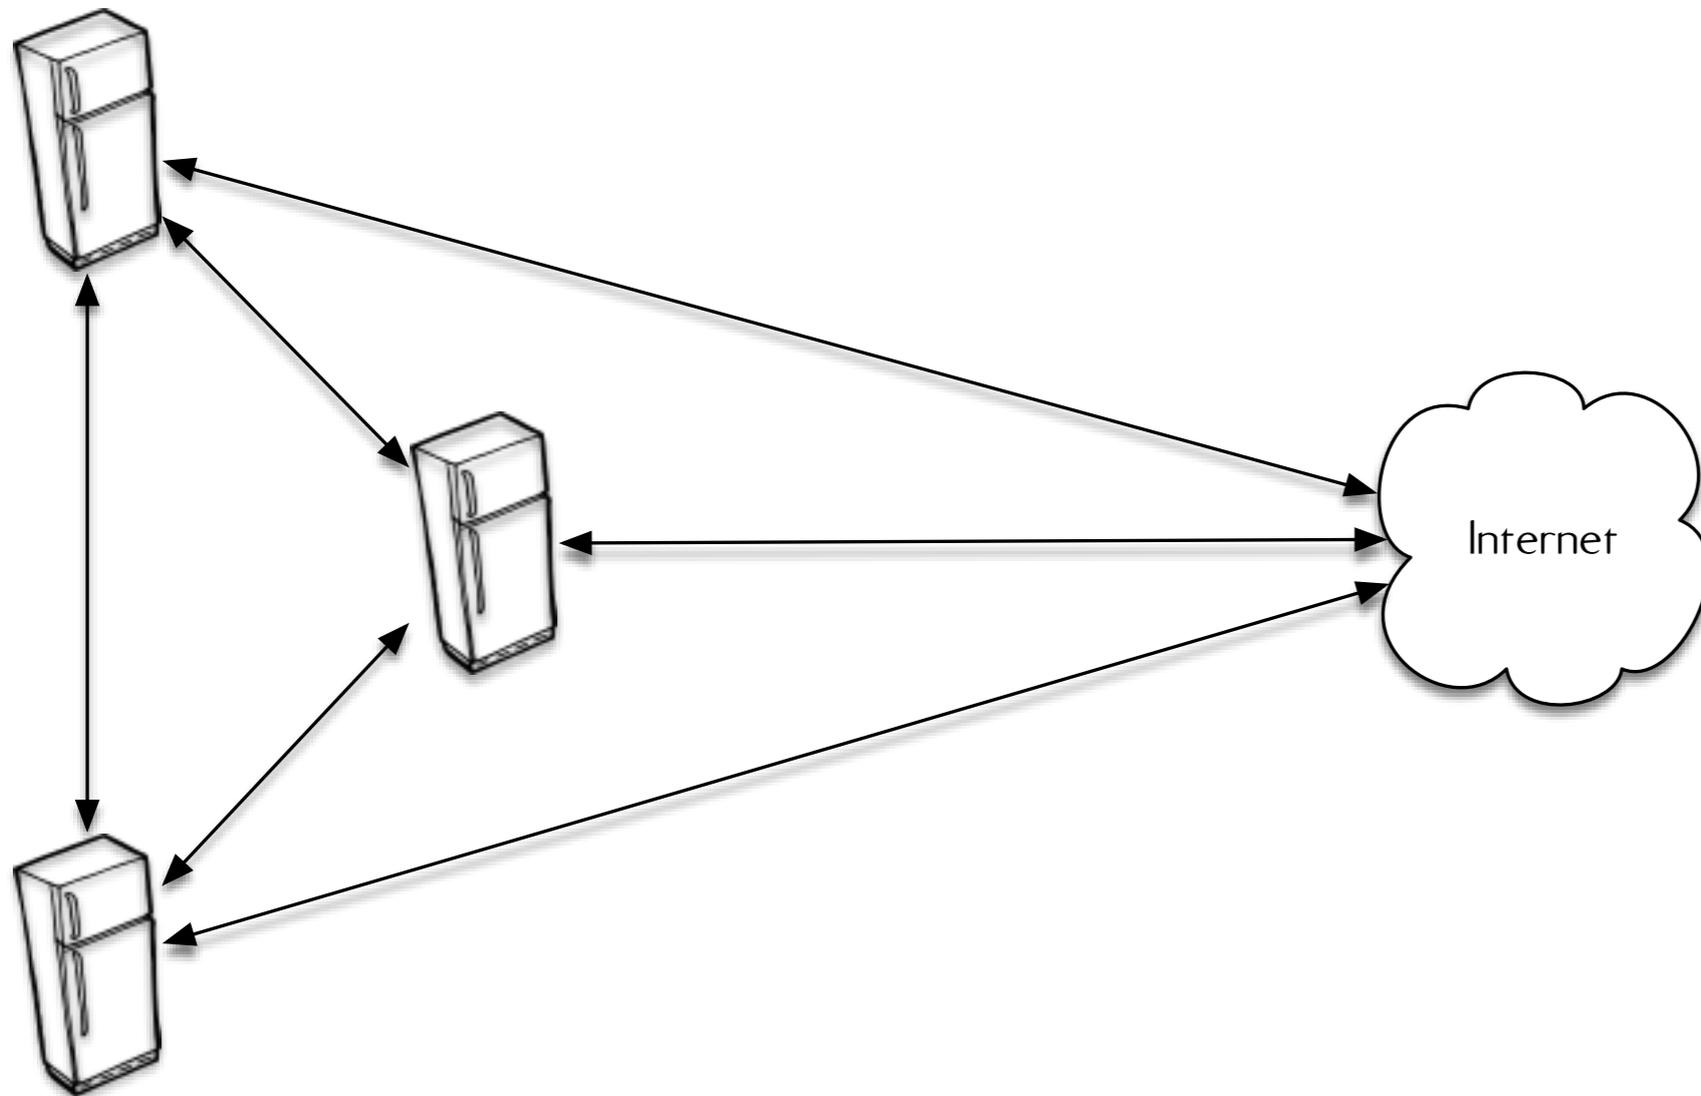


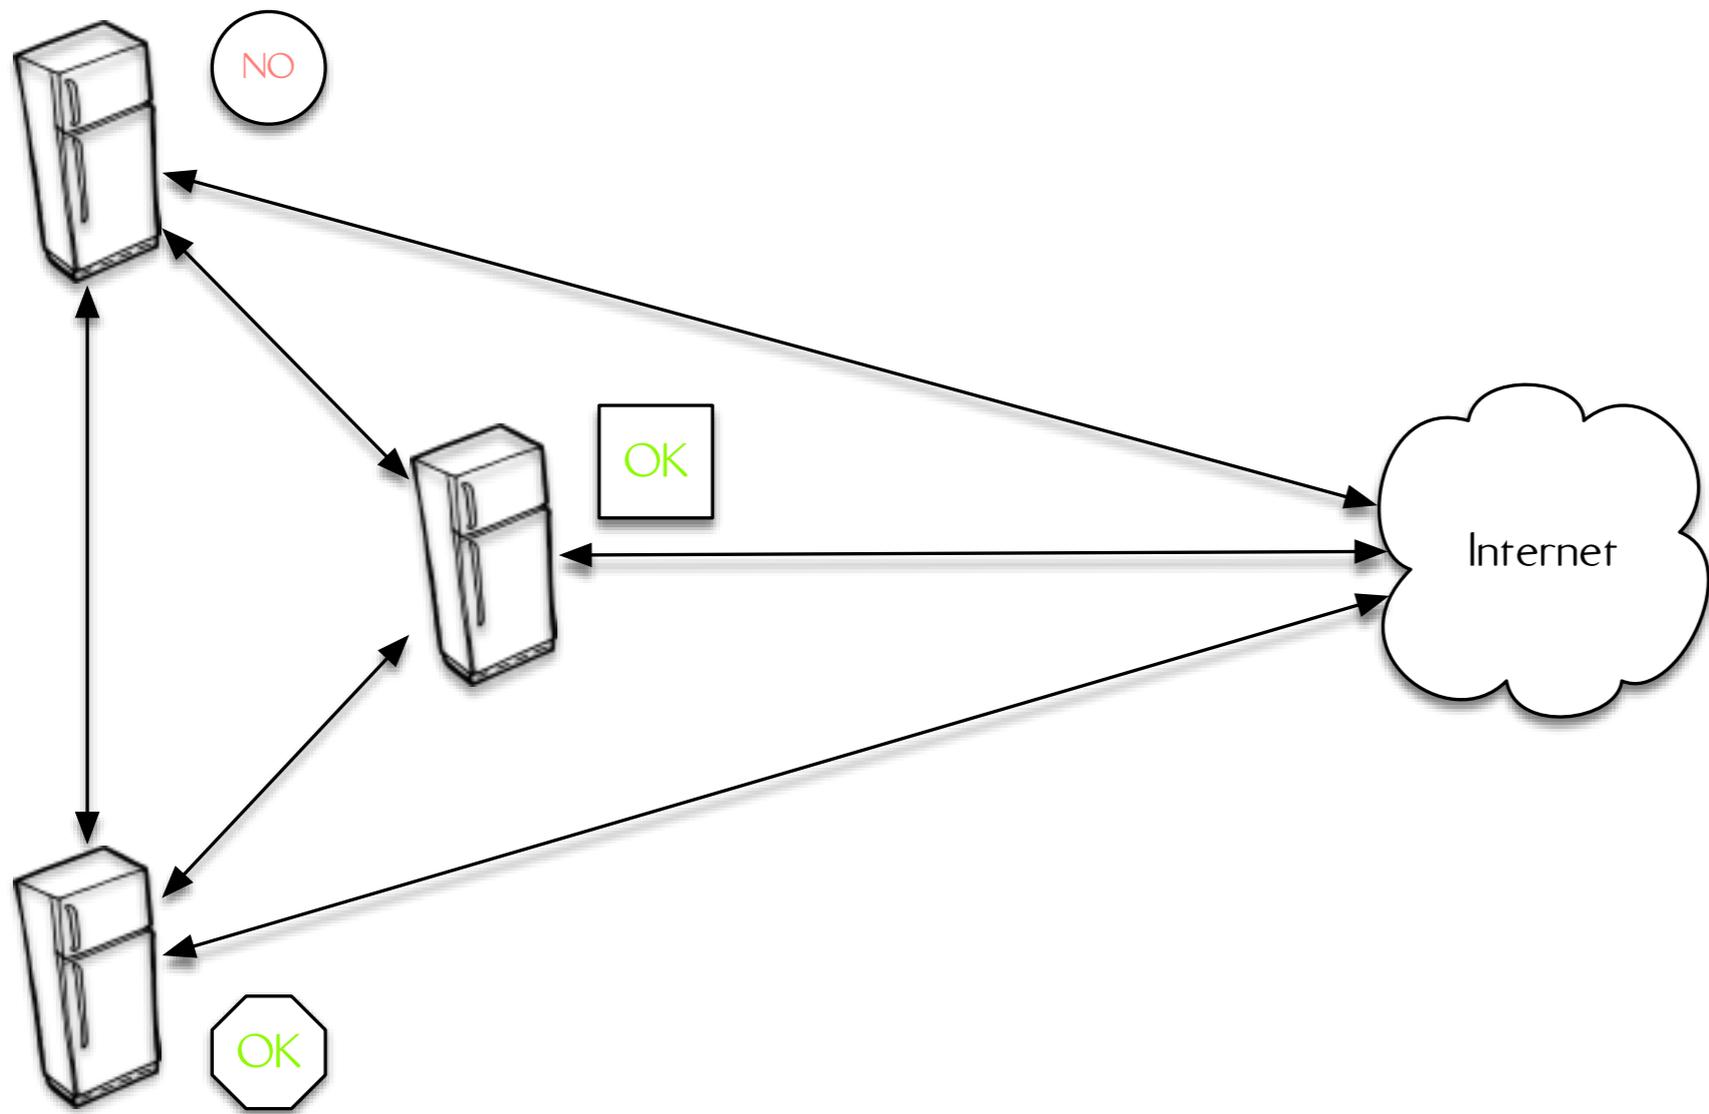


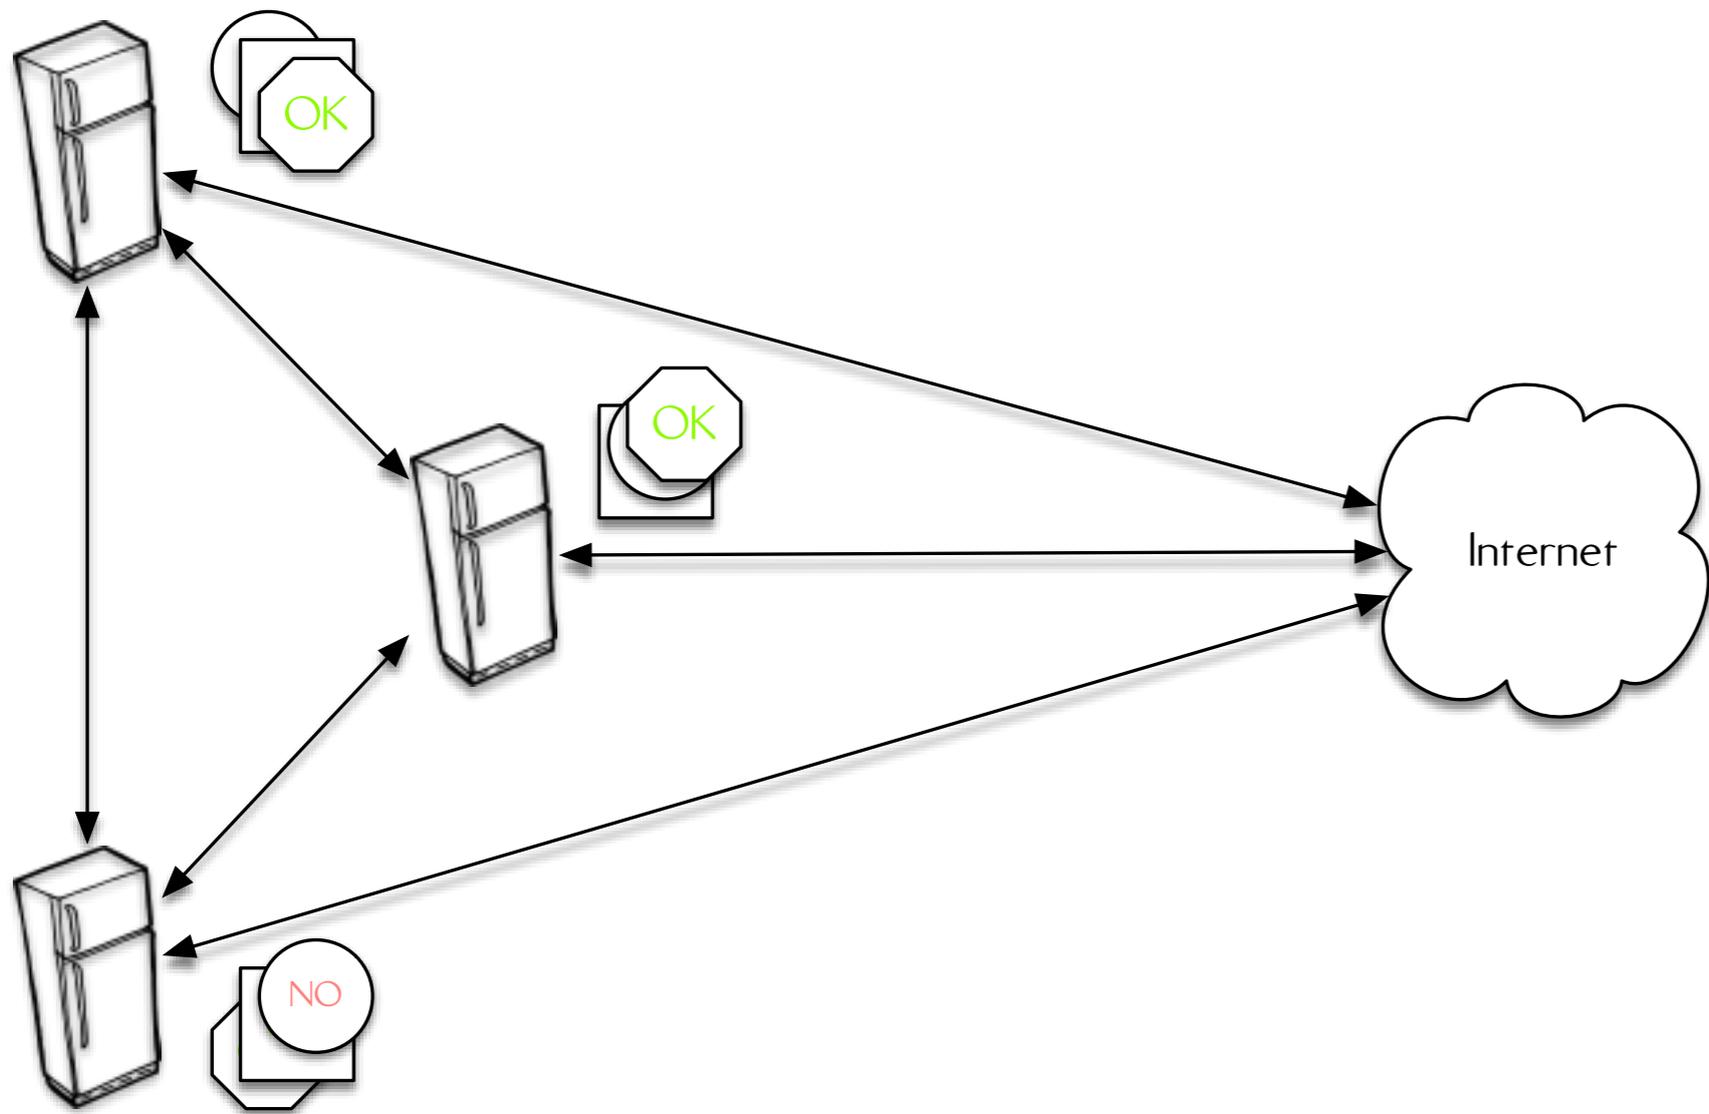


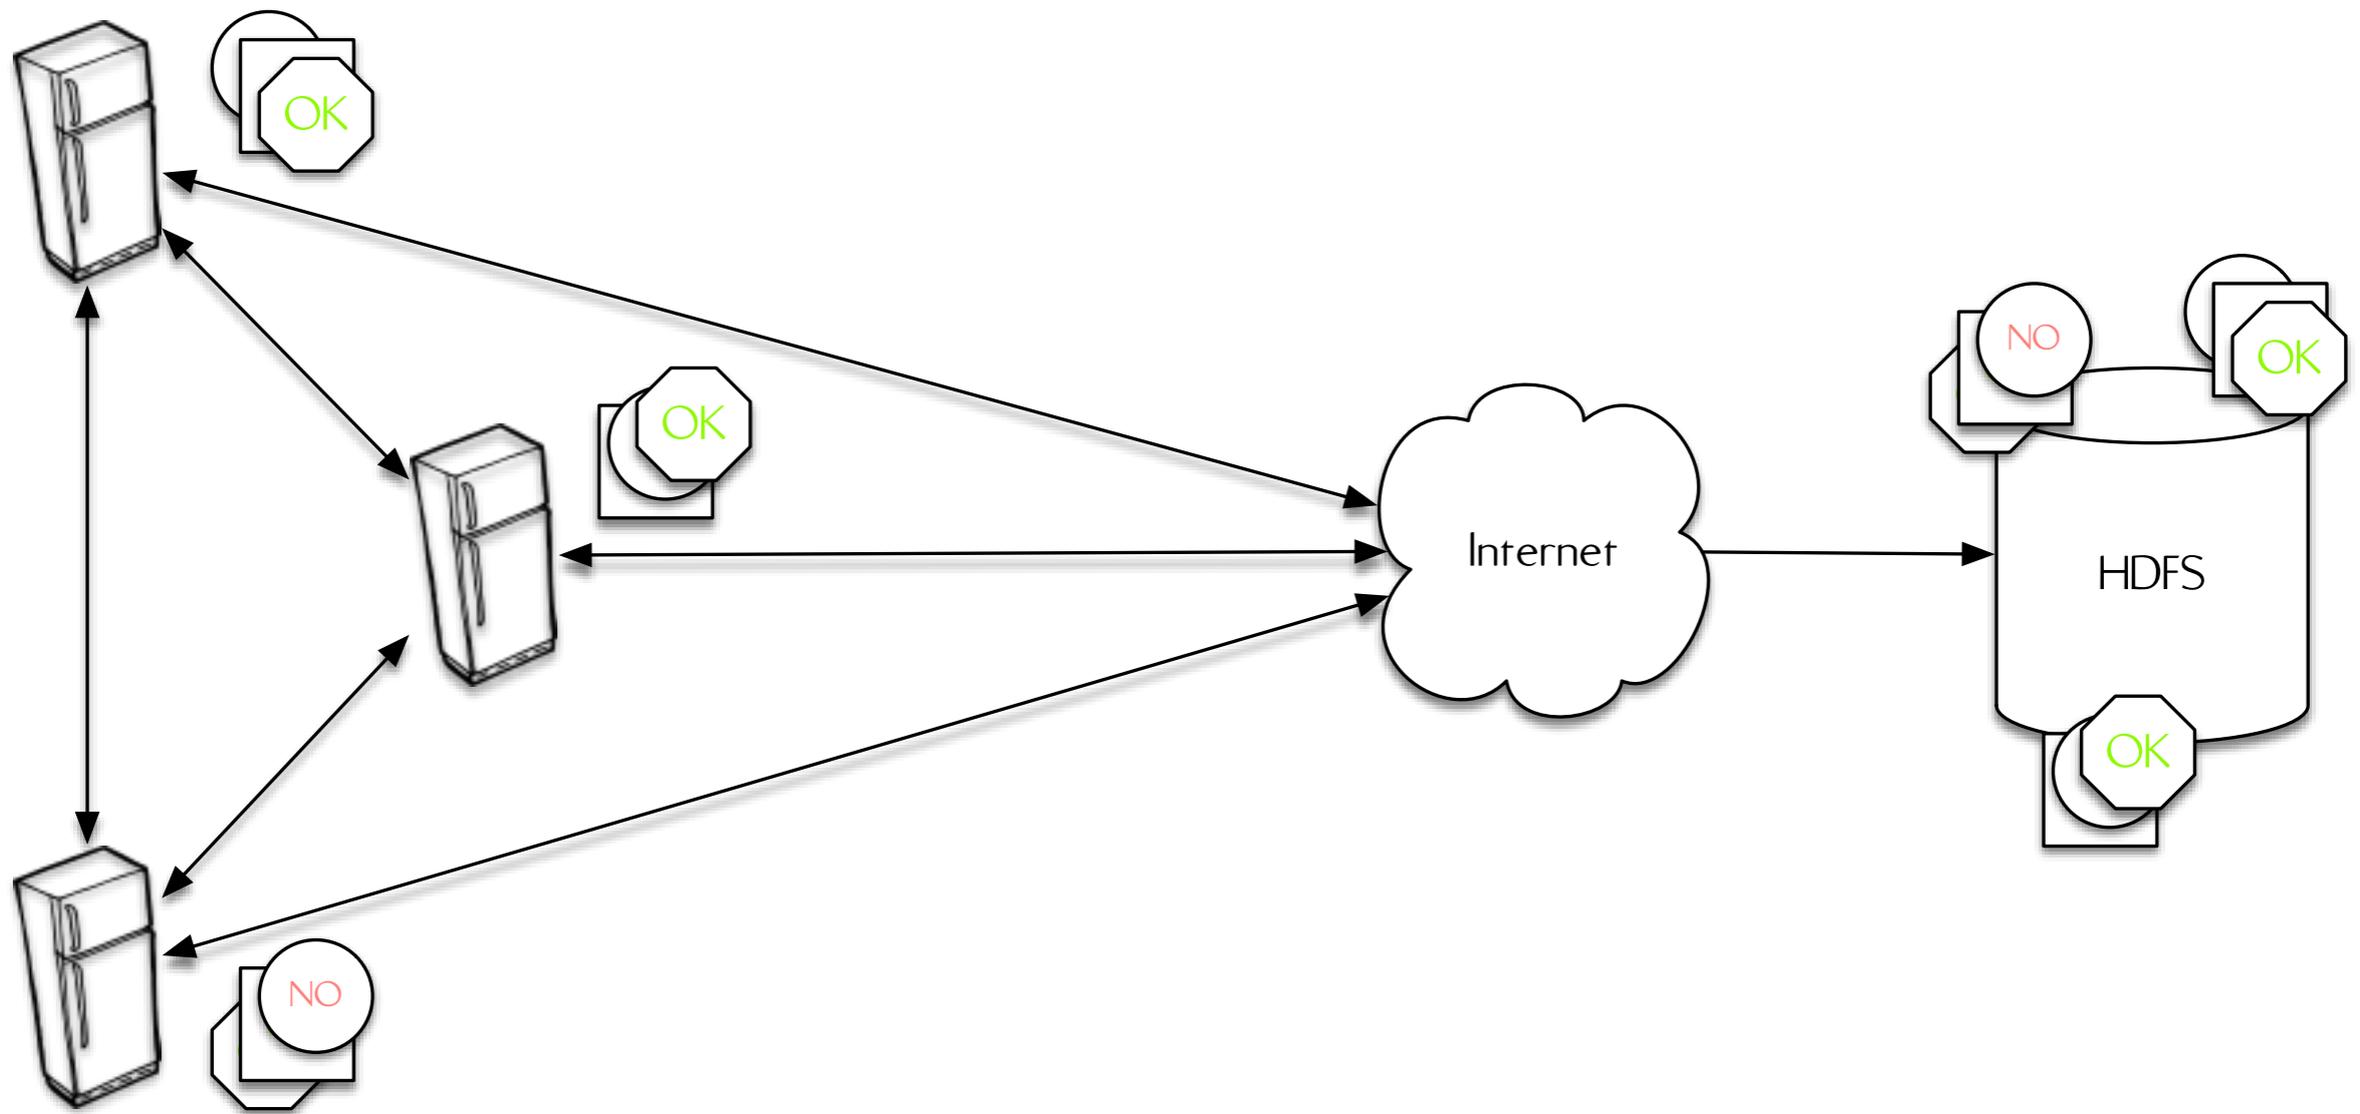


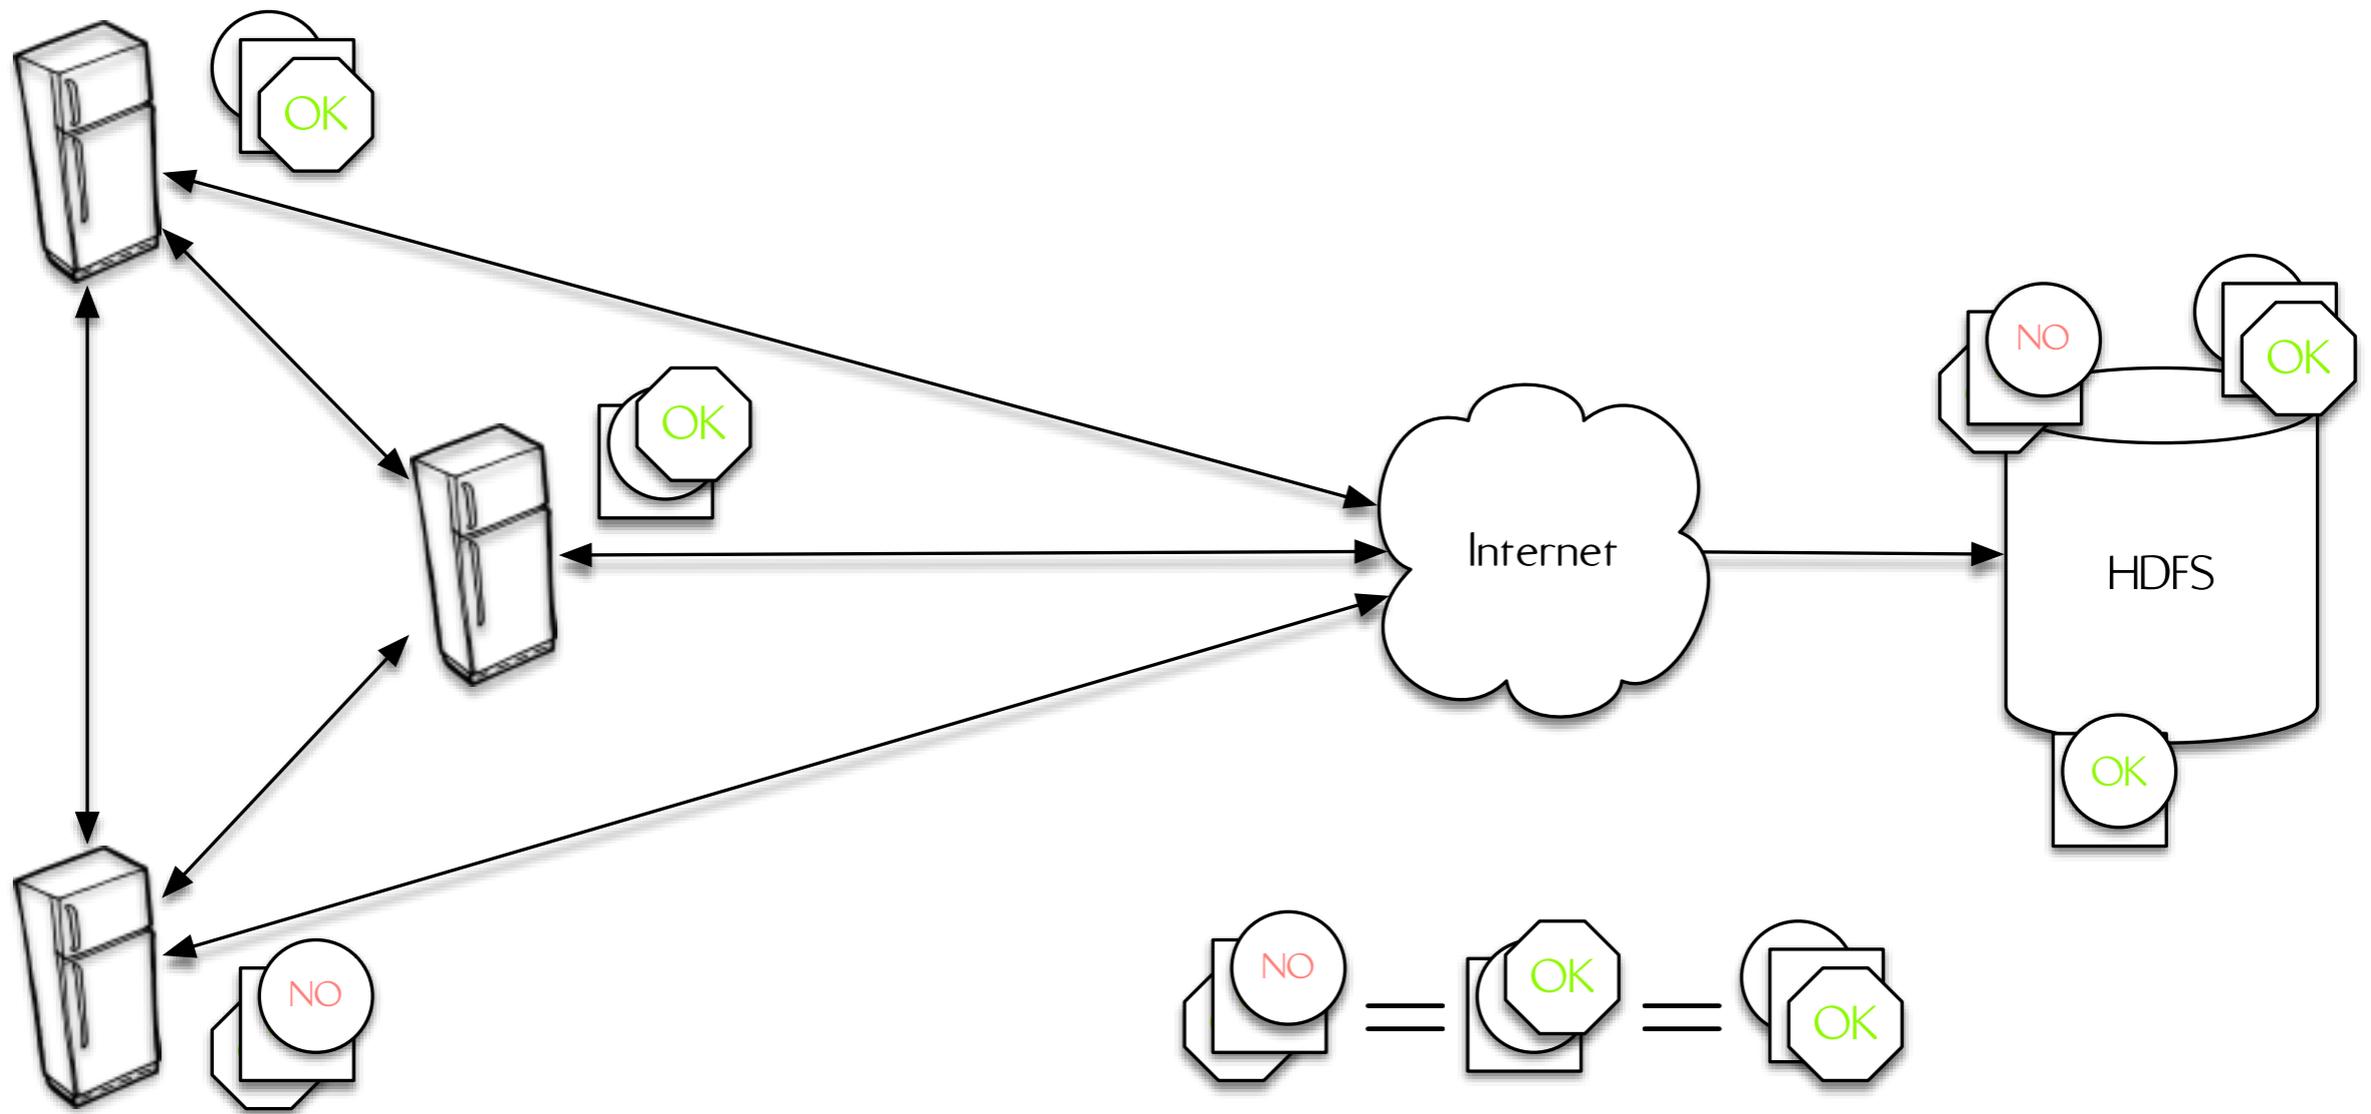
Solution
Aggregate Dissemination

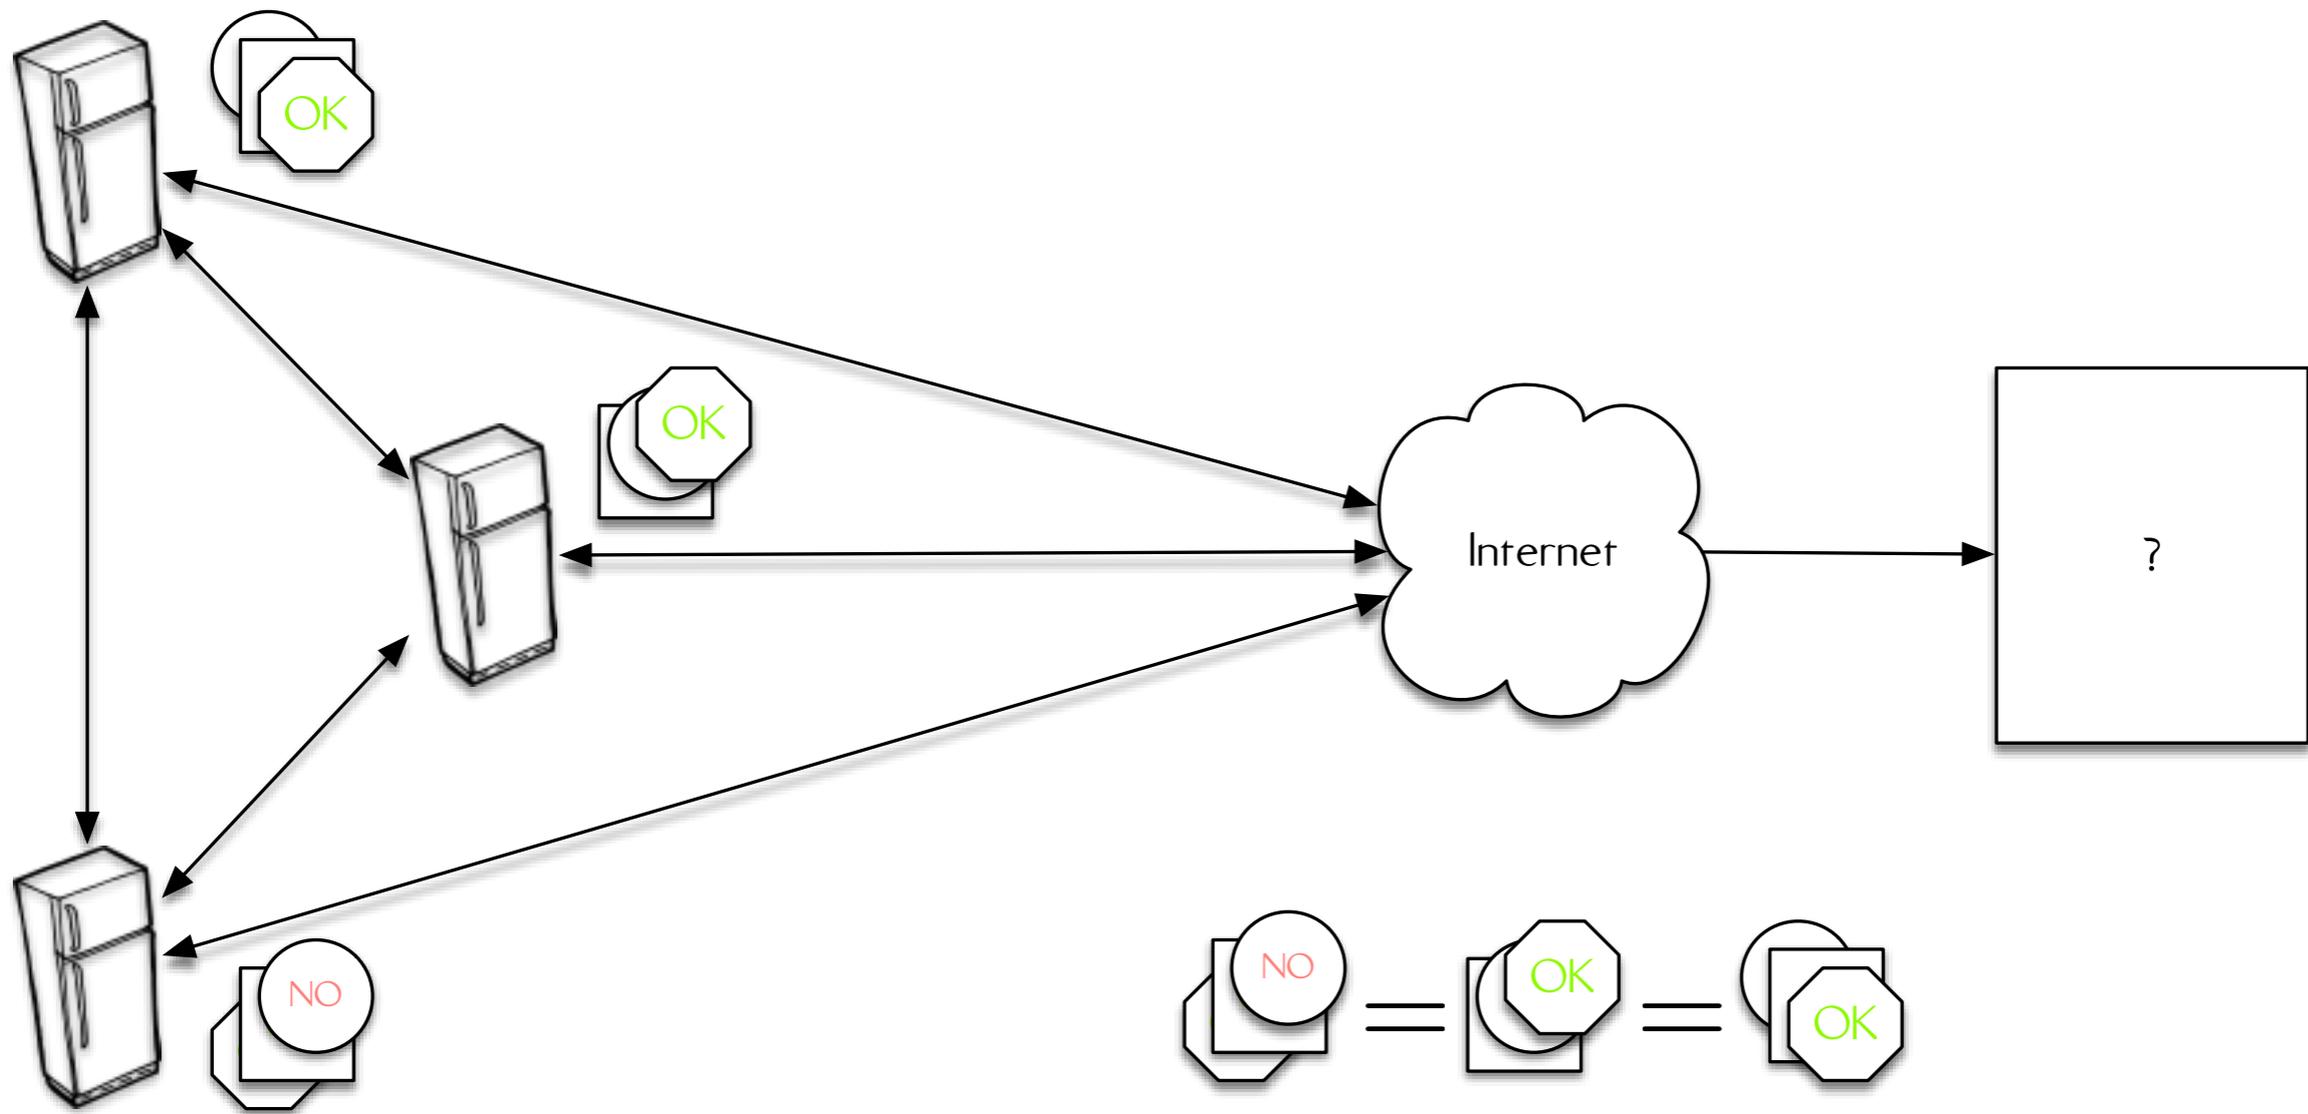


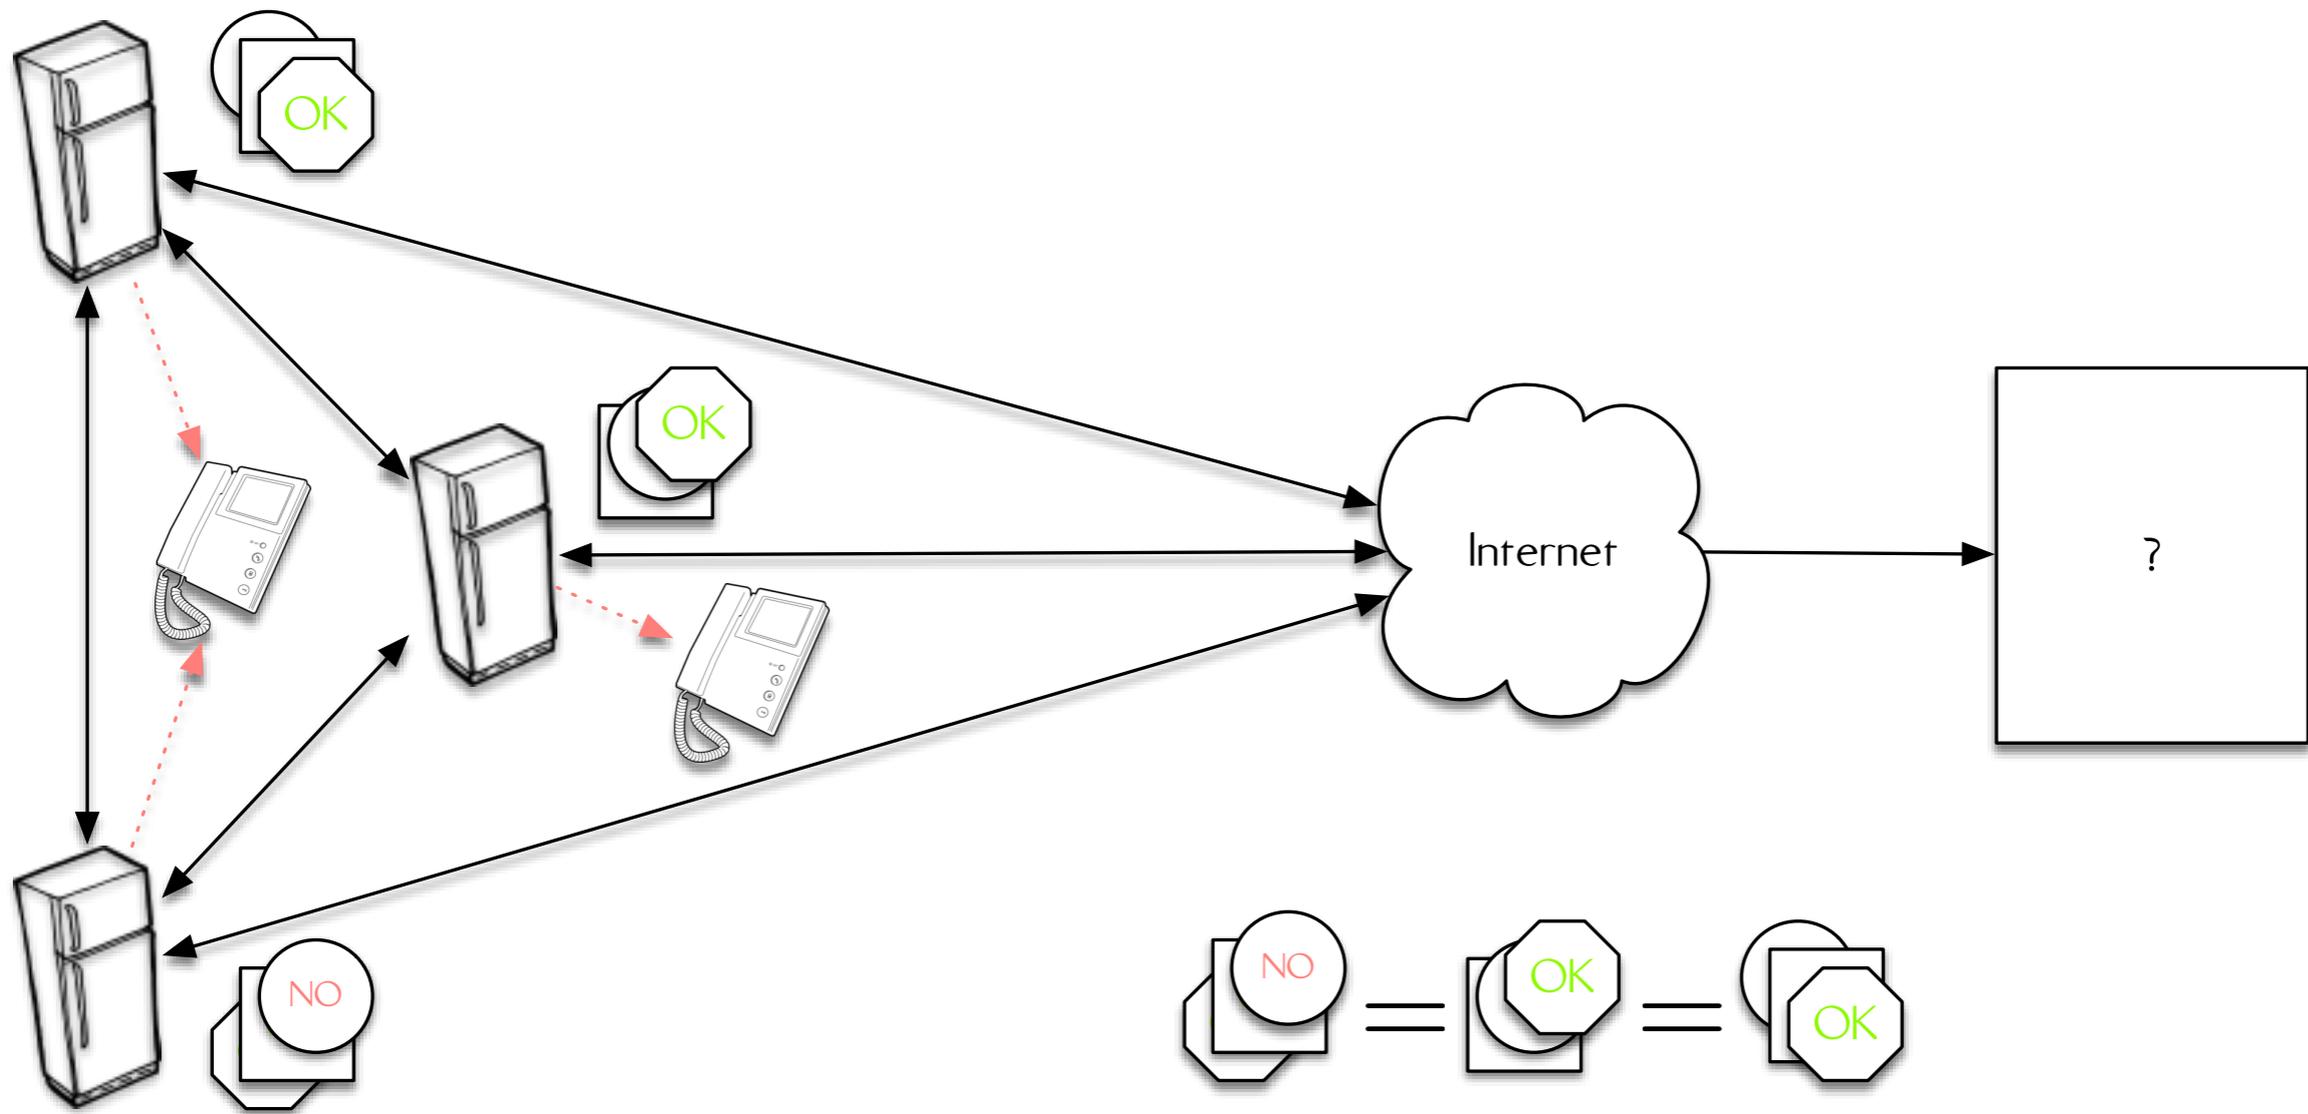












Local Computation

- **Reduce state transmission**
Perform some local computation to reduce transmitted state on the wire

Local Computation

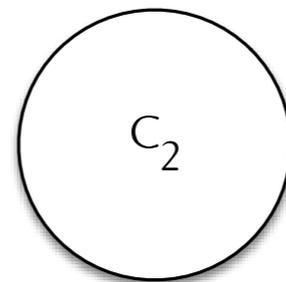
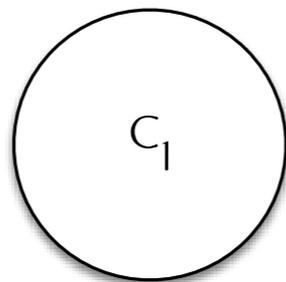
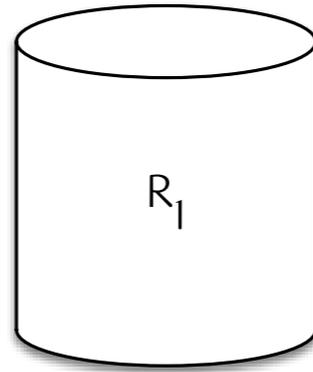
- **Reduce state transmission**
Perform some local computation to reduce transmitted state on the wire
- **Make local decisions**
Make decisions based on results of local computation

Databases

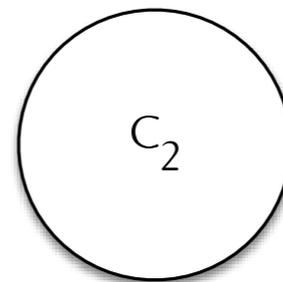
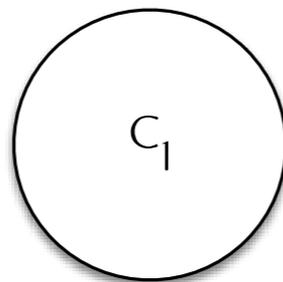
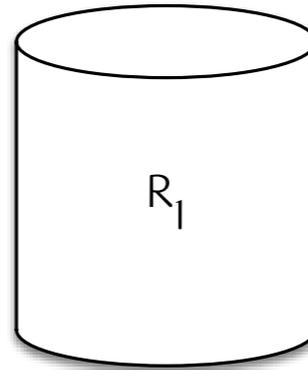
Consistency Models

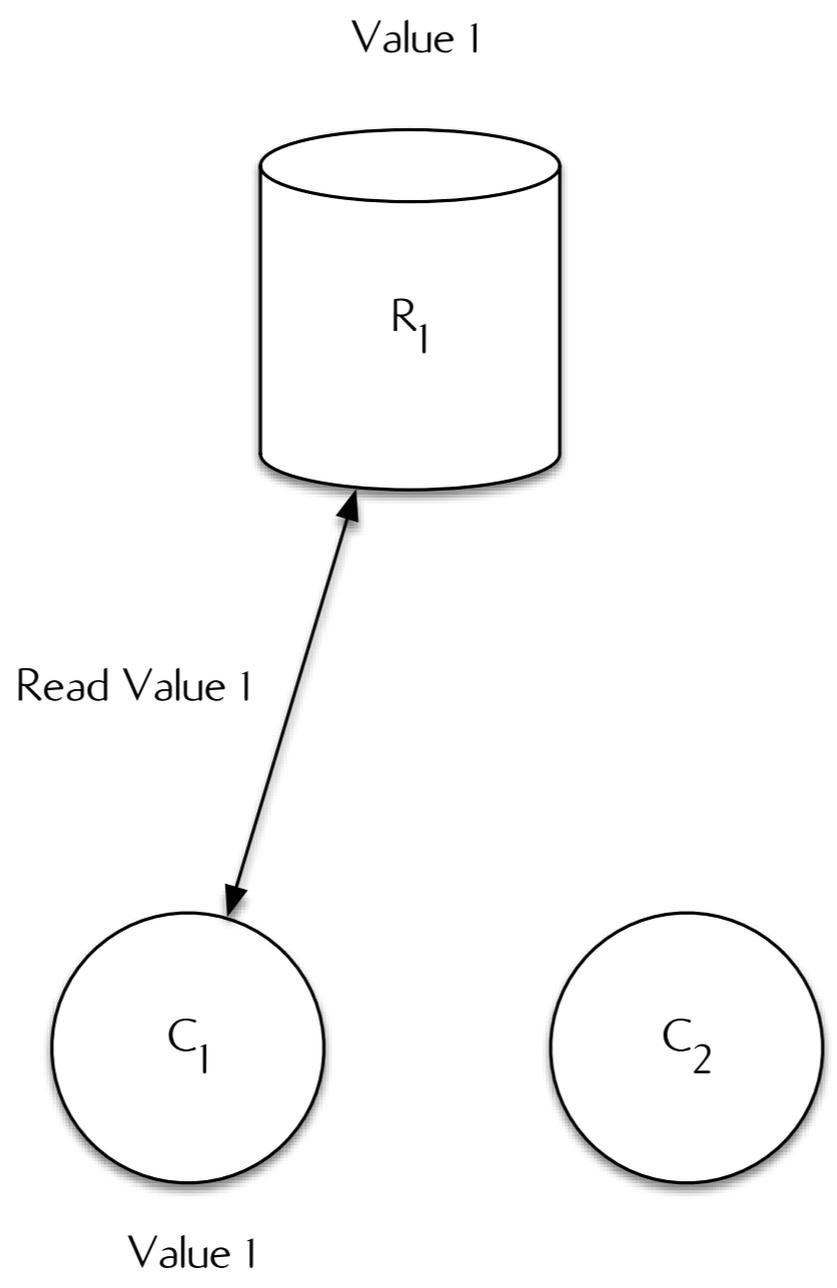
Databases

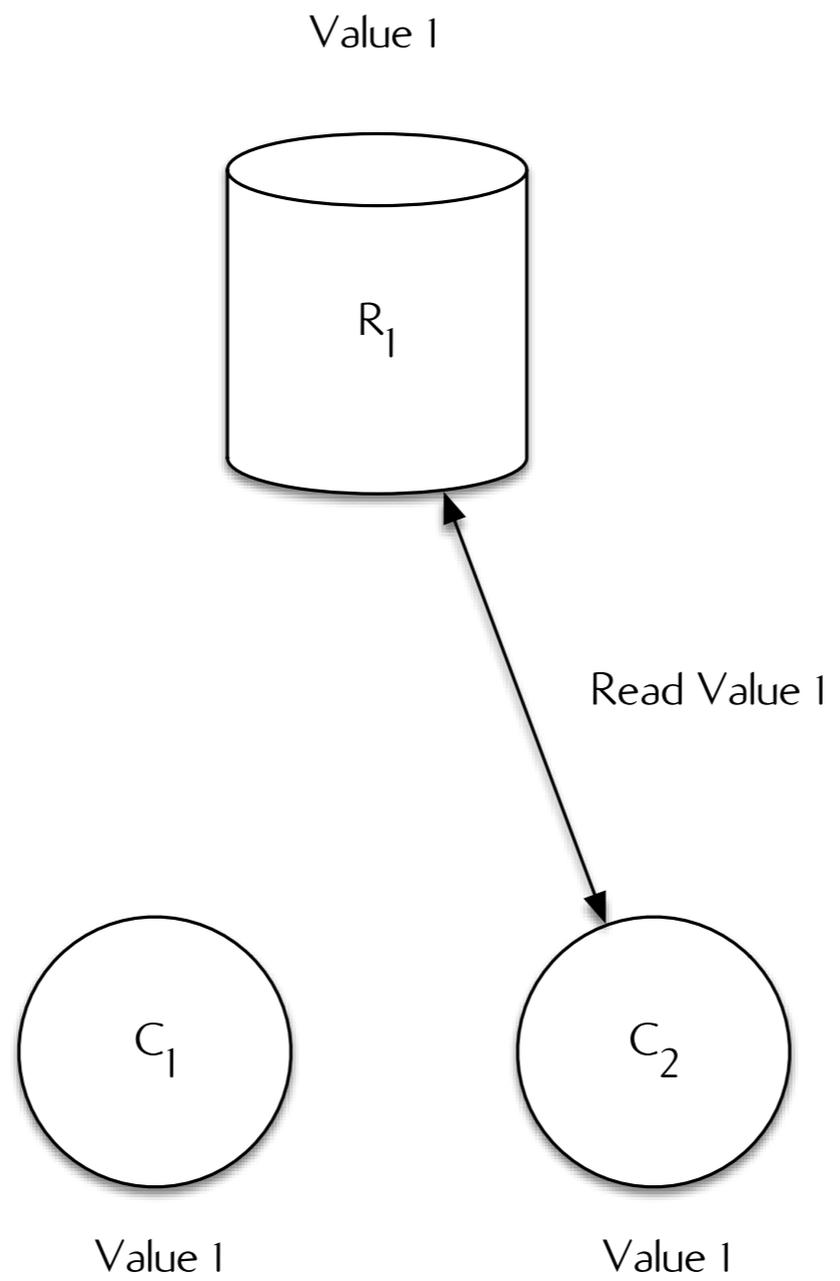
Strong Consistency

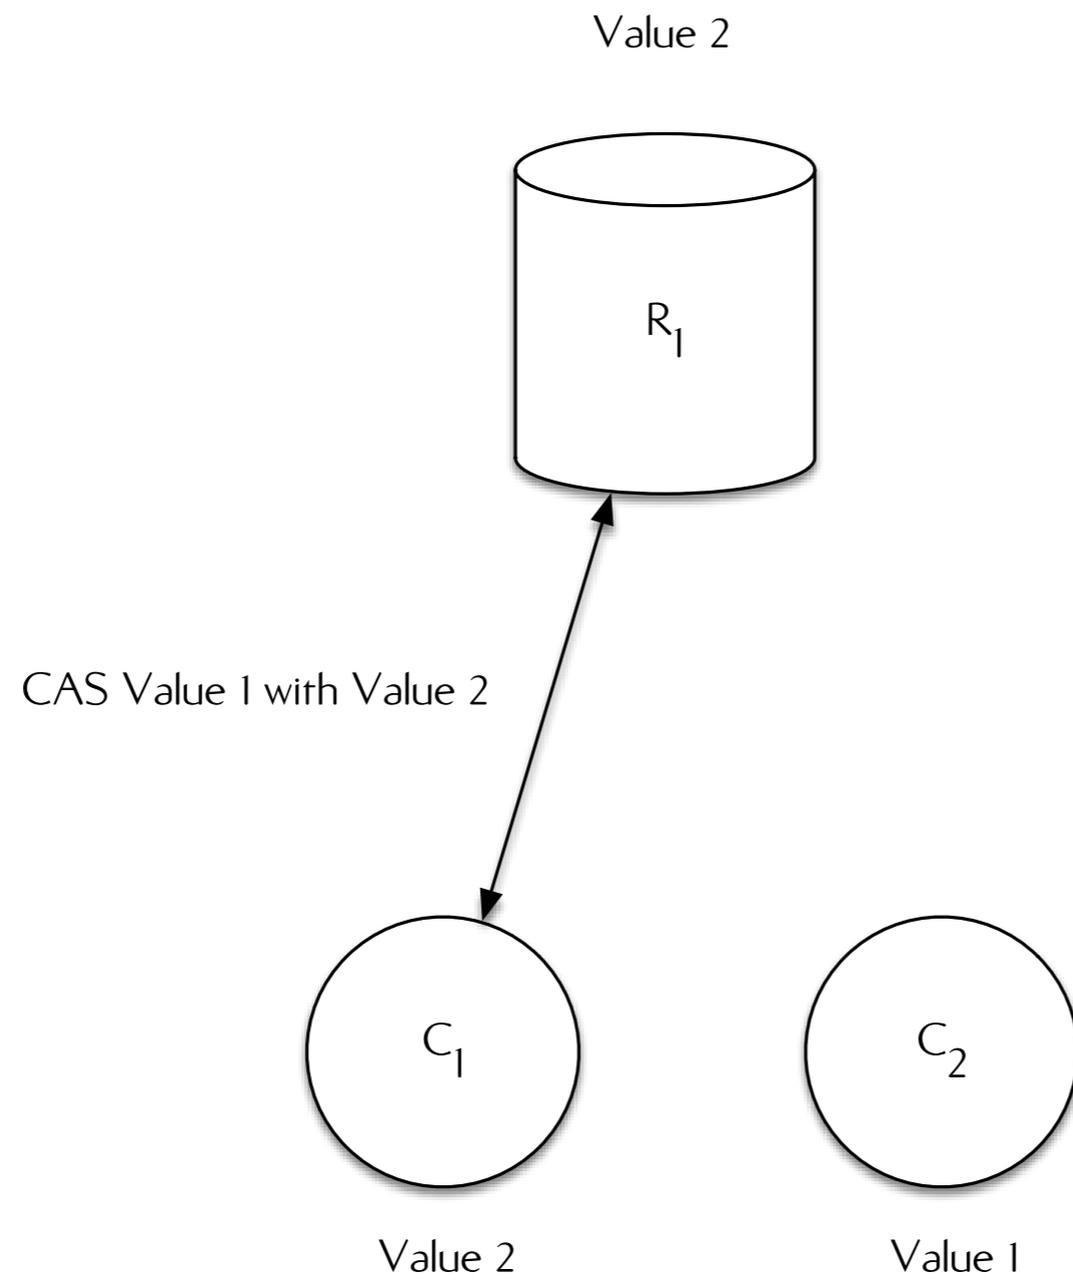


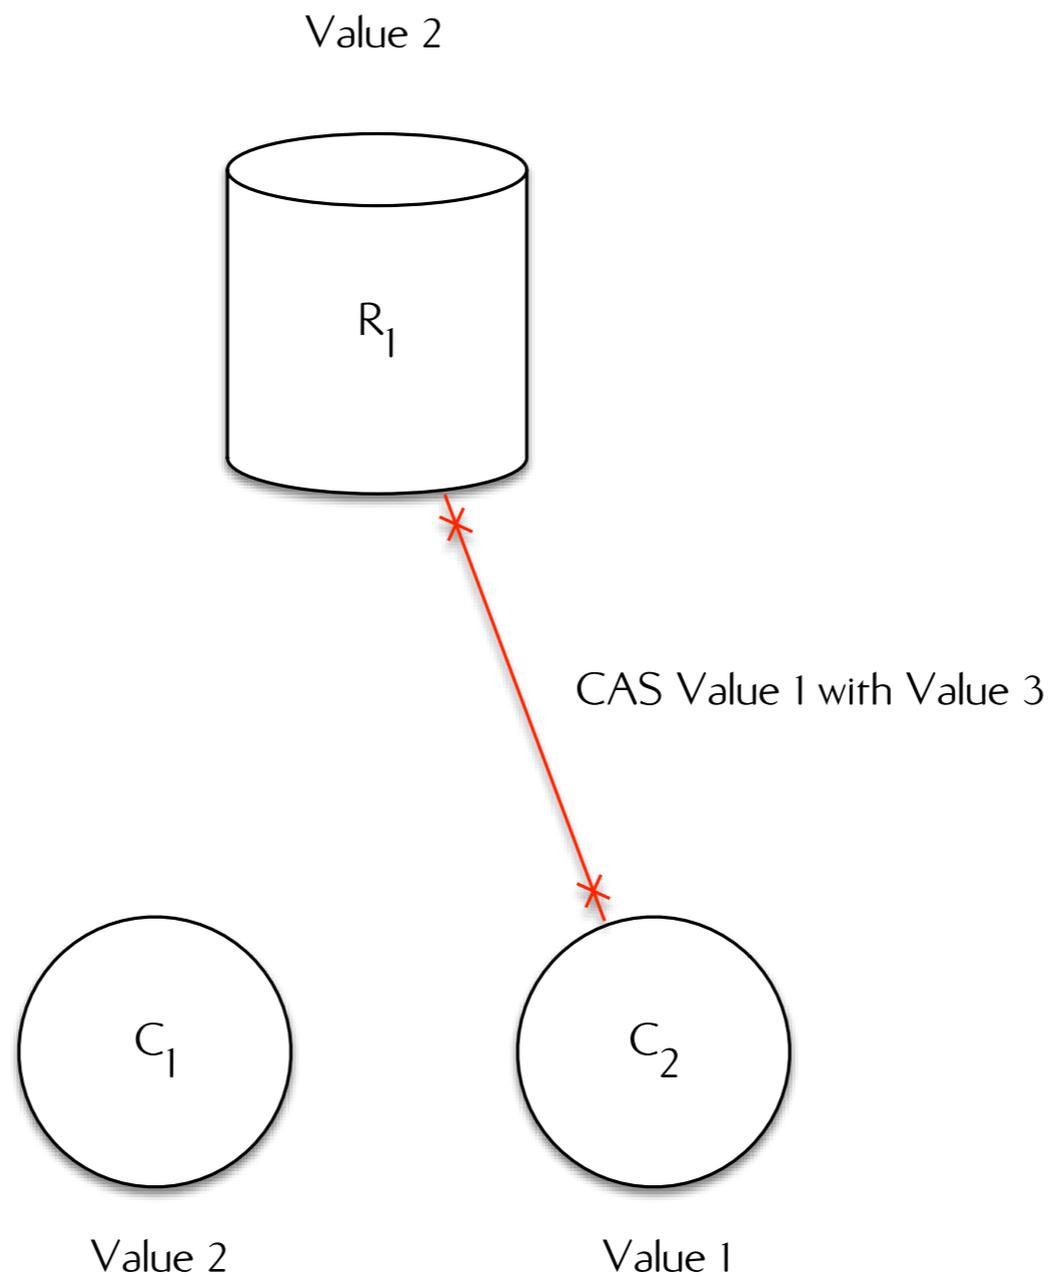
Value 1



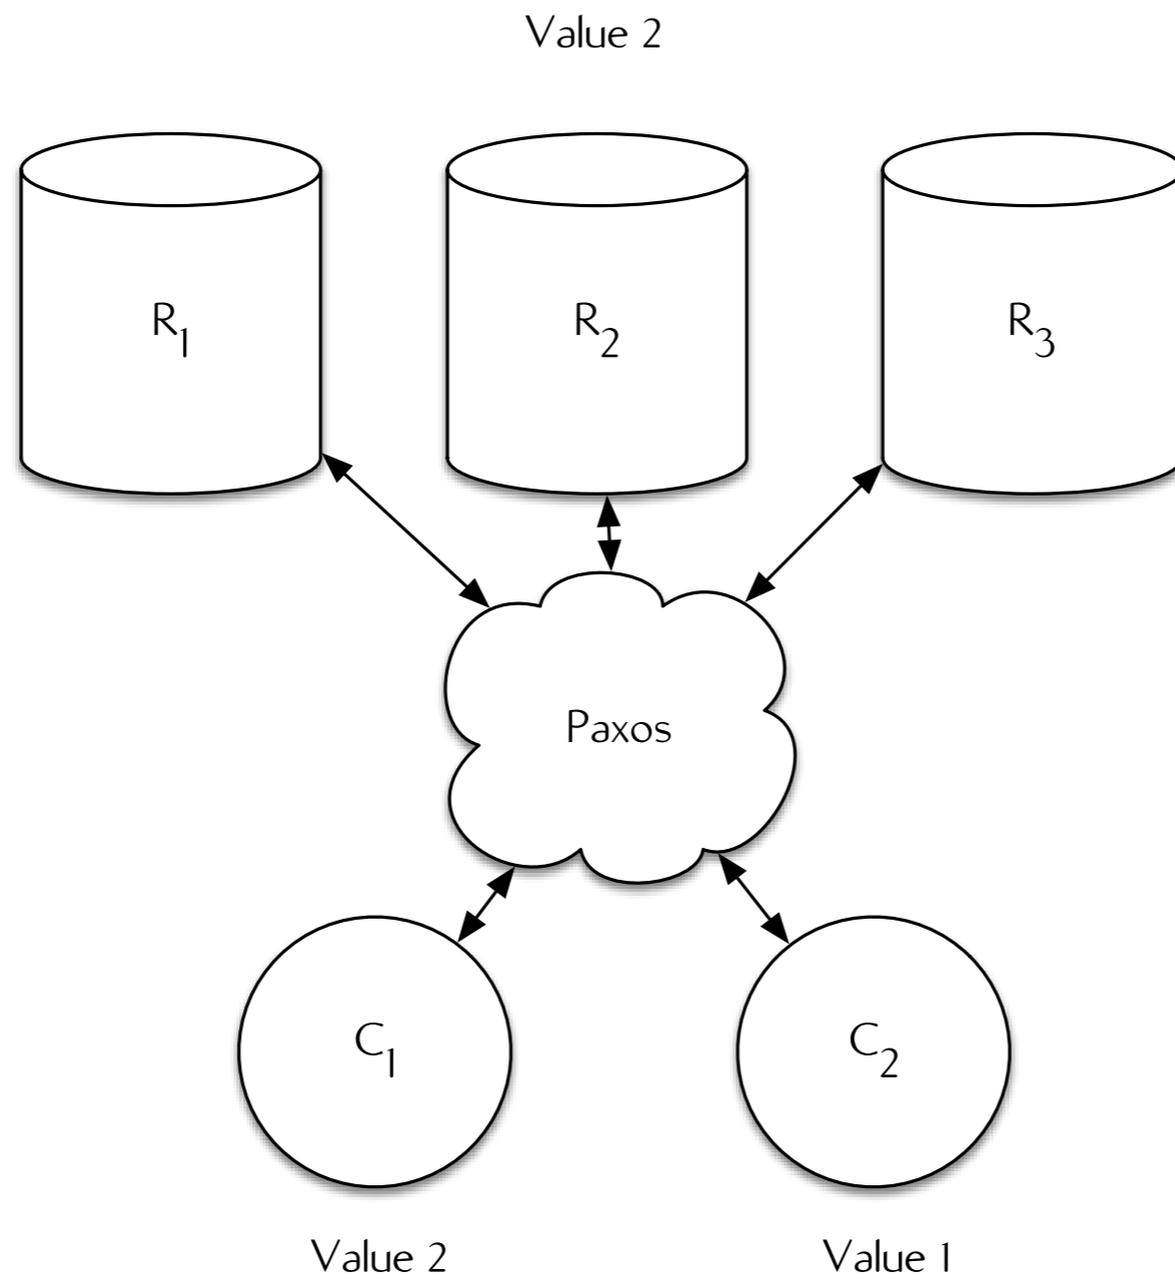






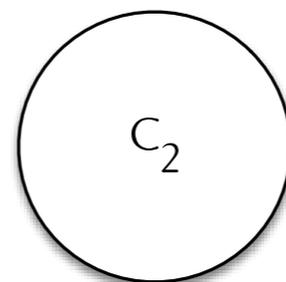
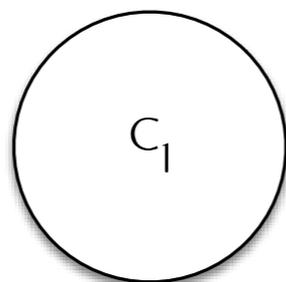
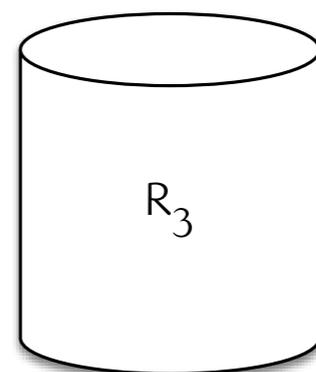
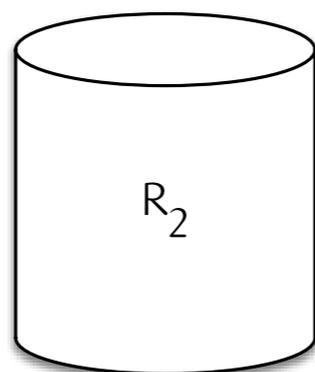
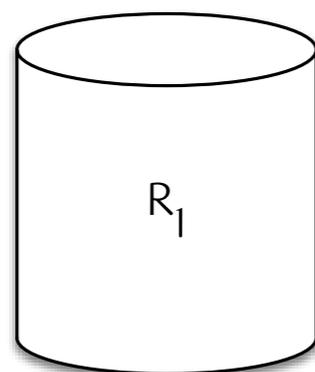


I won't diagram
the Paxos protocol

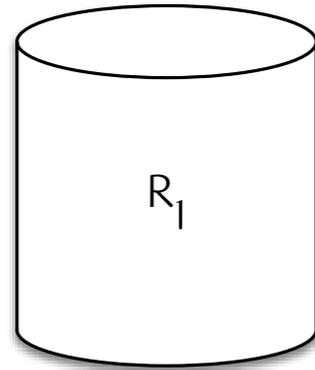


Databases

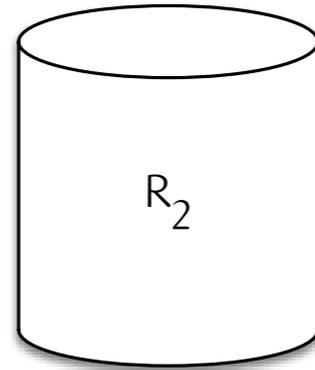
Eventual Consistency



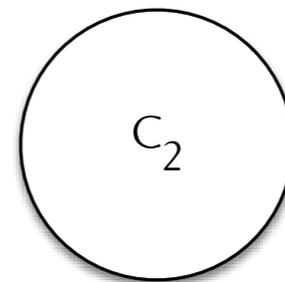
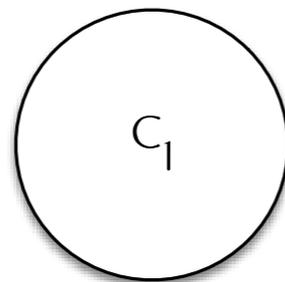
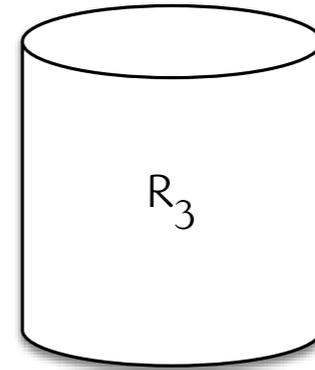
Value 1 @ [1, 0, 0]



Value 1 @ [1, 0, 0]



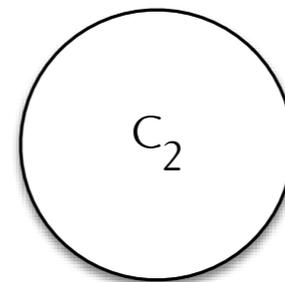
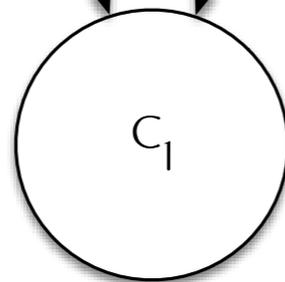
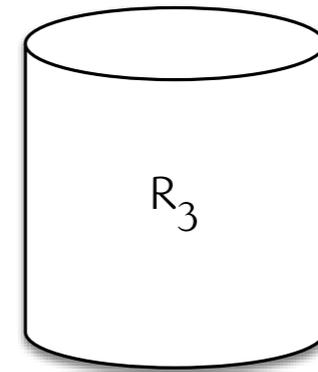
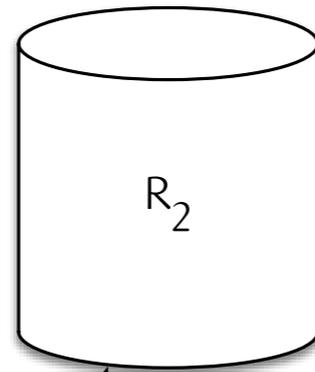
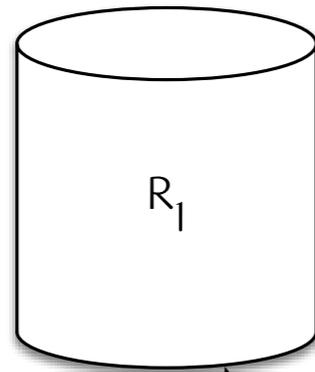
Value 1 @ [1, 0, 0]



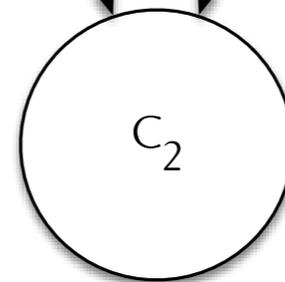
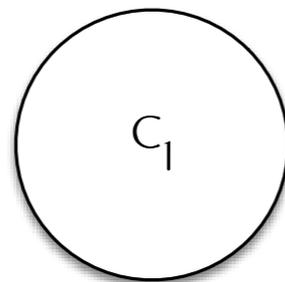
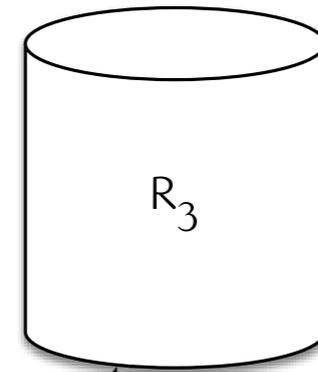
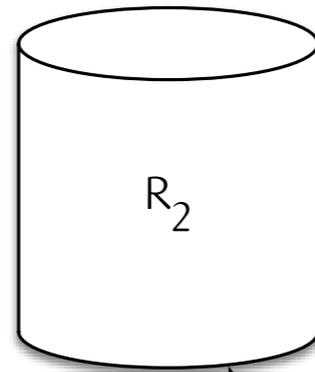
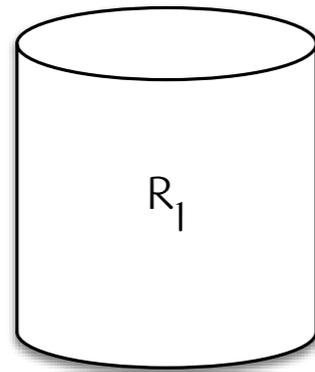
Value 1 @ [1, 0, 0]

Value 1 @ [1, 0, 0]

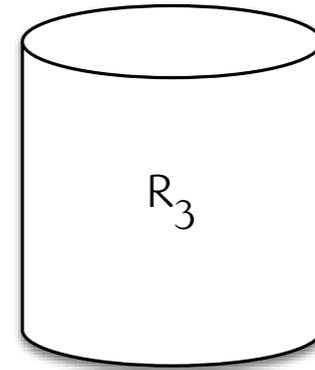
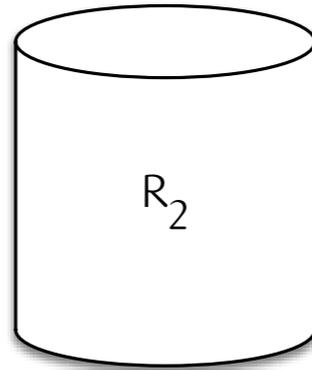
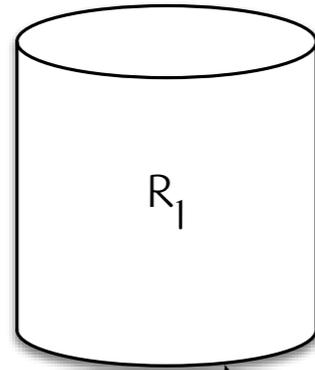
Value 1 @ [1, 0, 0]



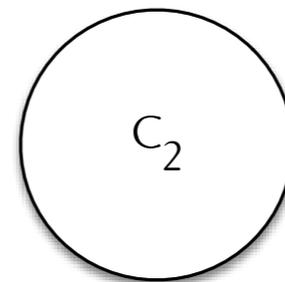
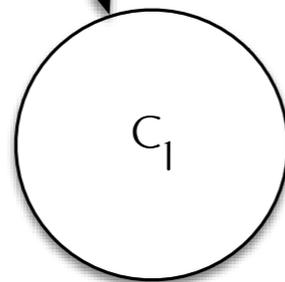
Value 1 @ [1, 0, 0]



Value 1 @ [1, 0, 0]

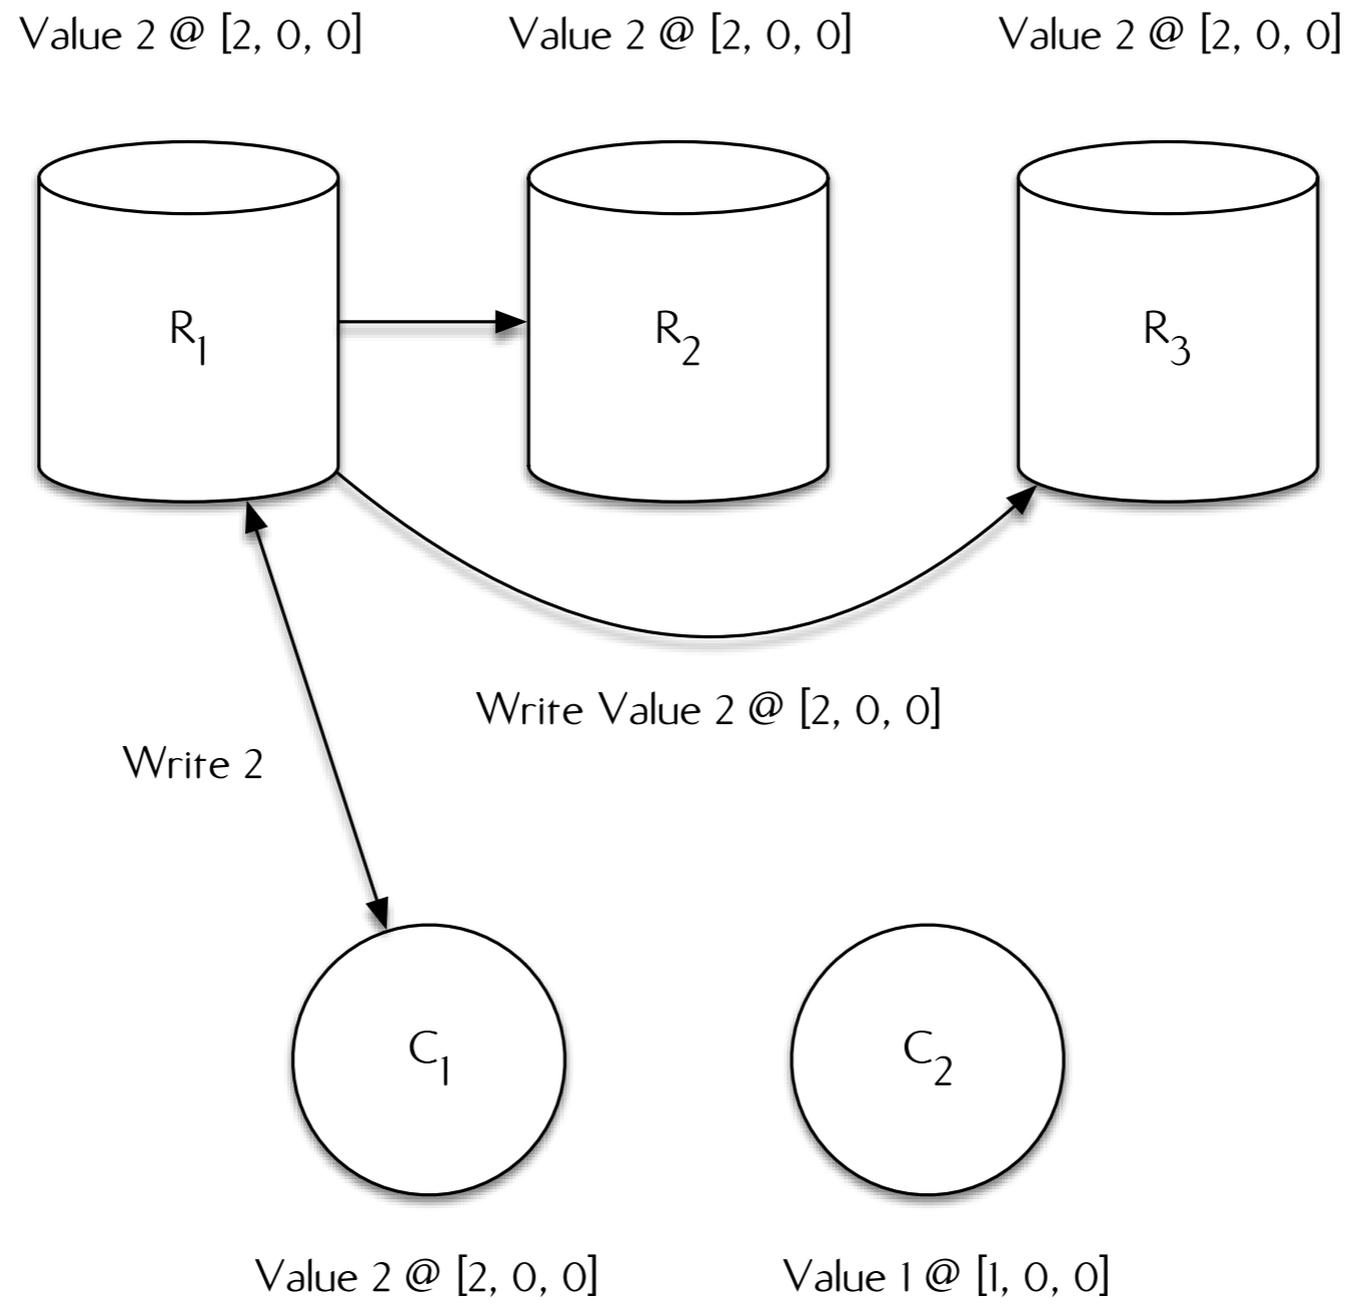


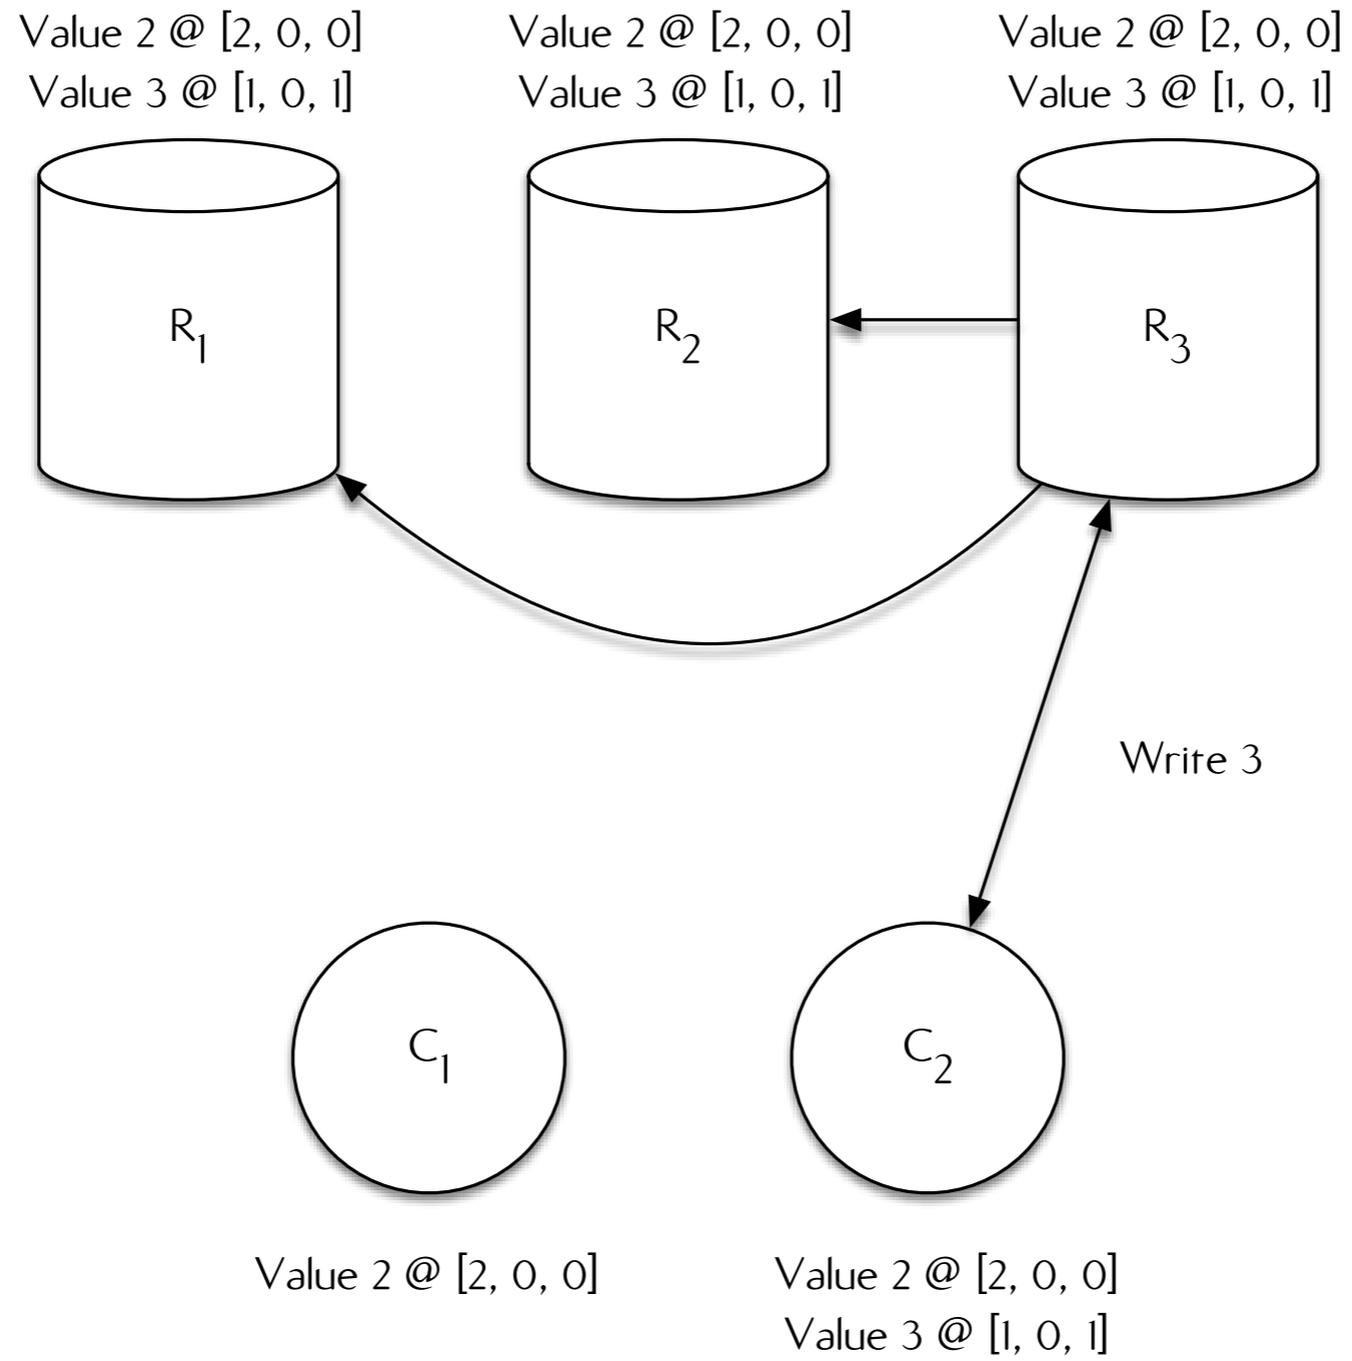
Write 2

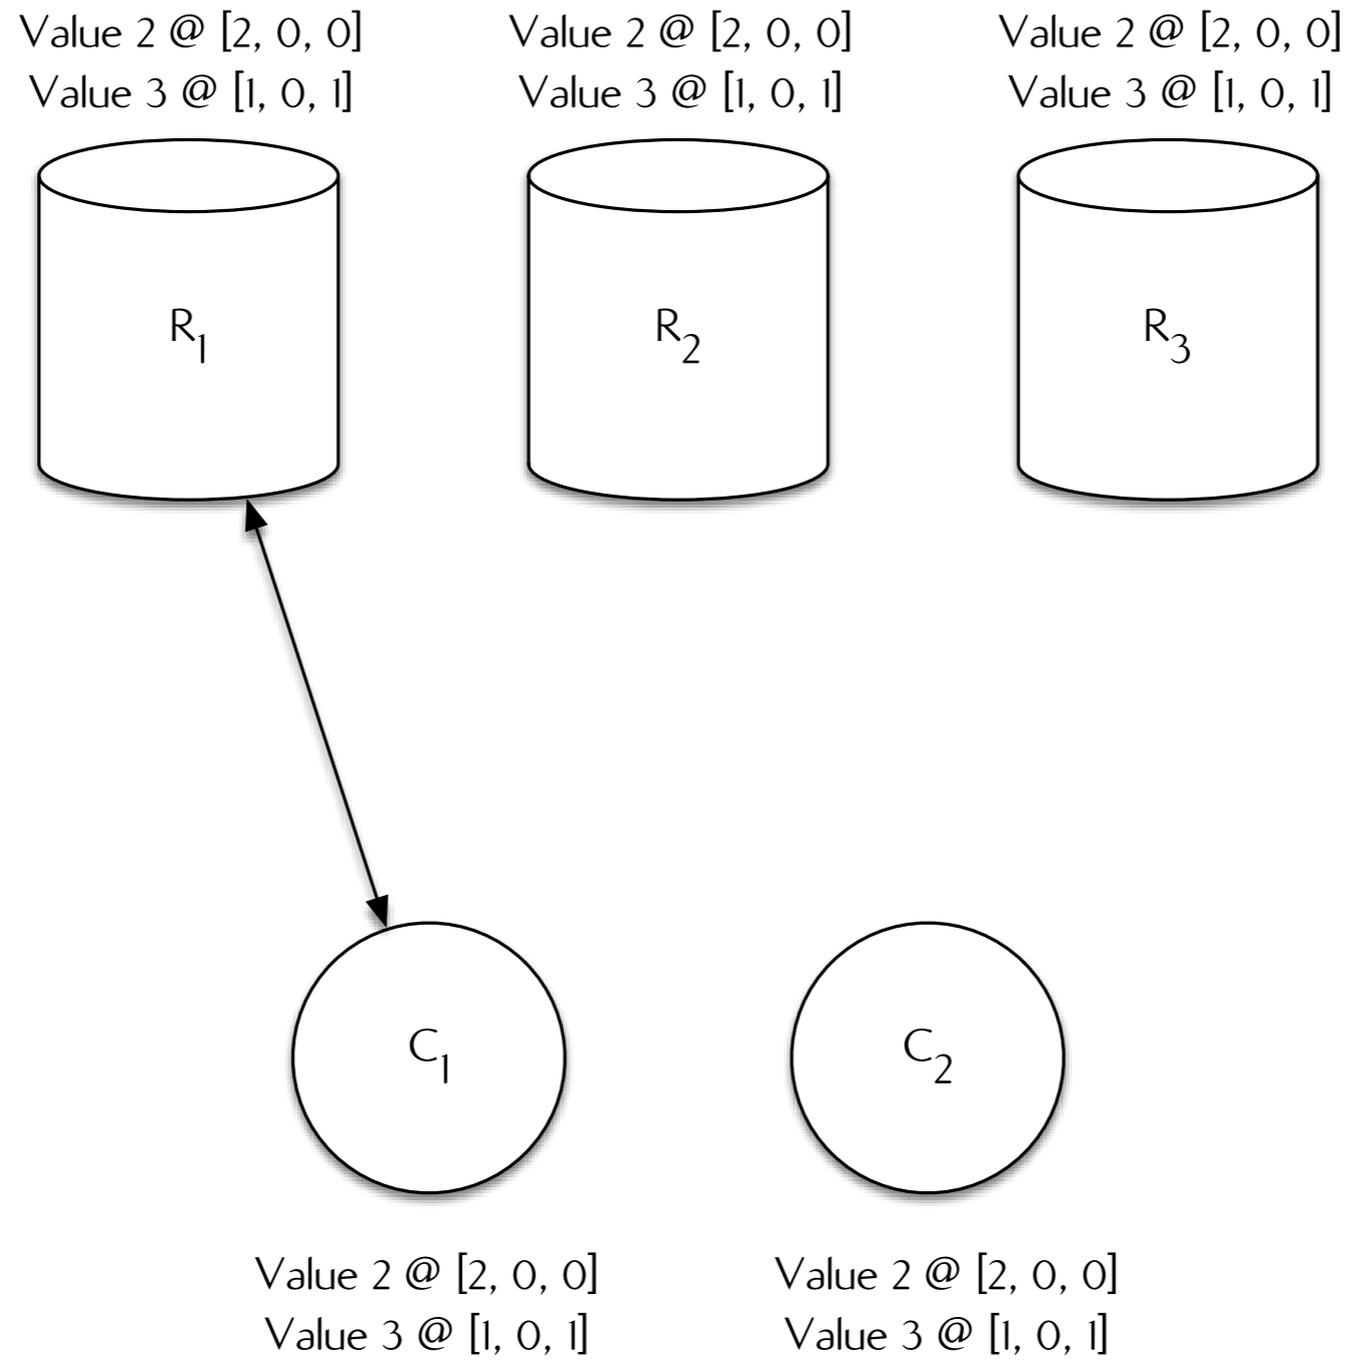


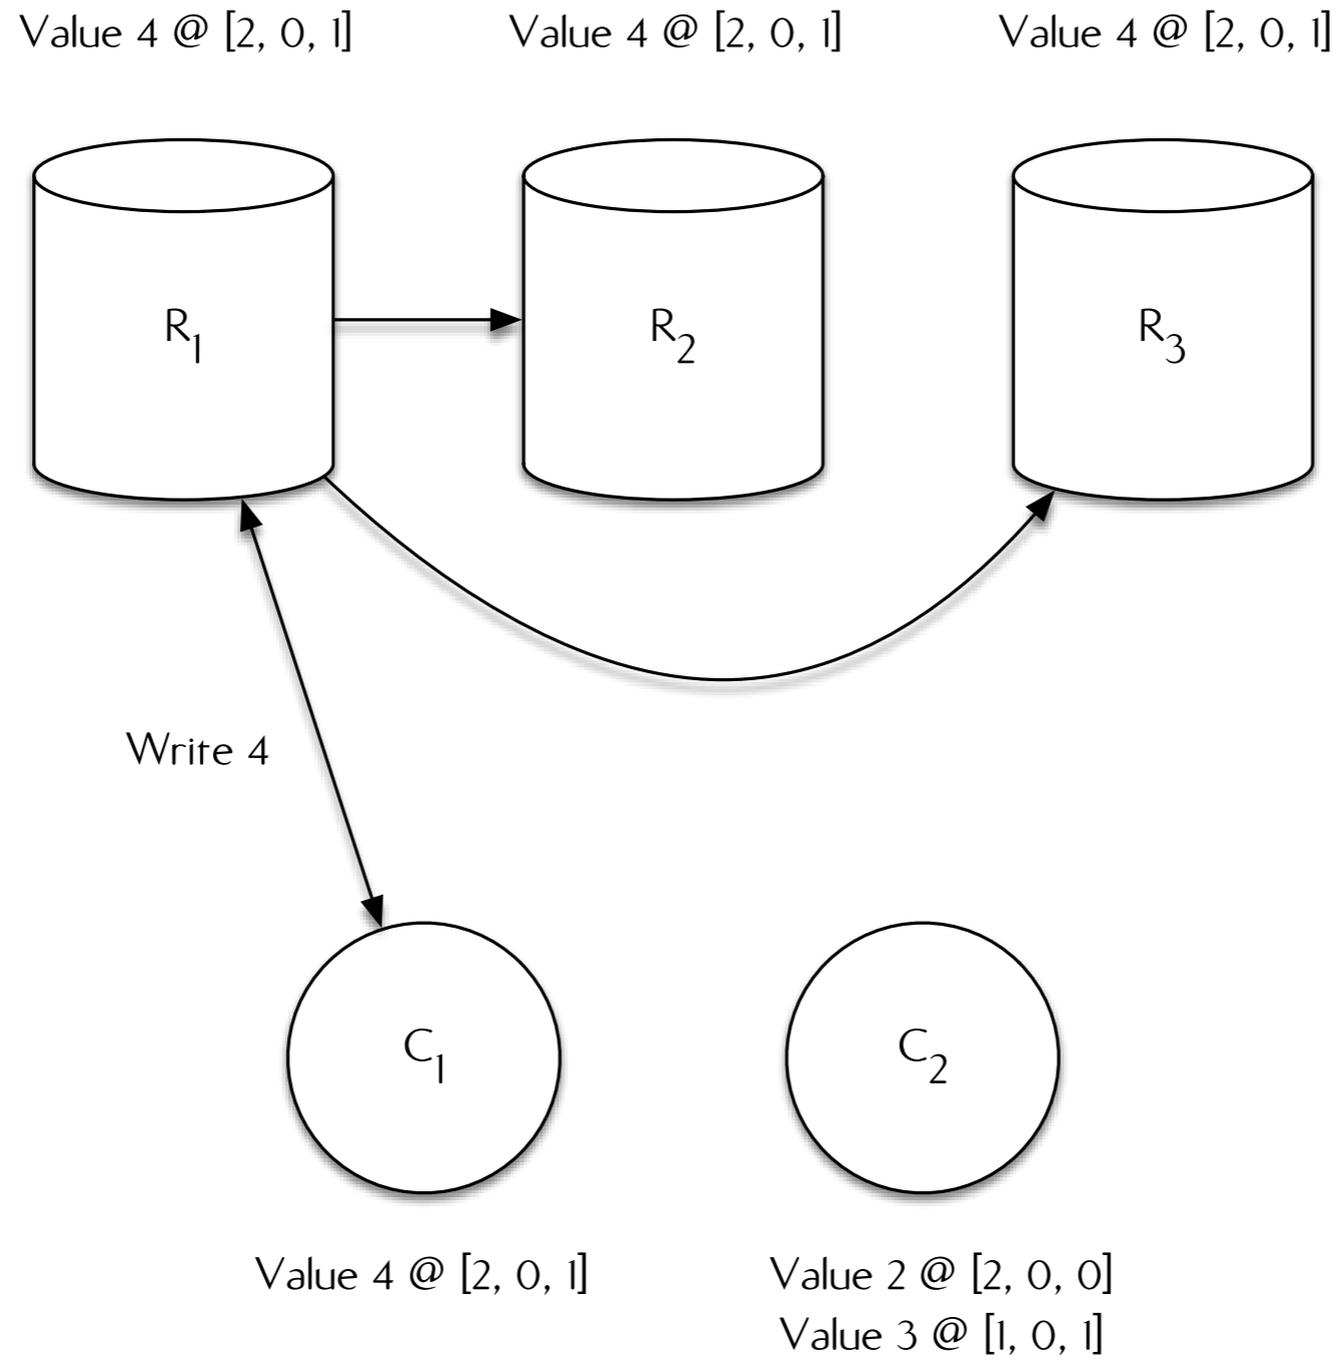
Value 1 @ [1, 0, 0]

Value 1 @ [1, 0, 0]



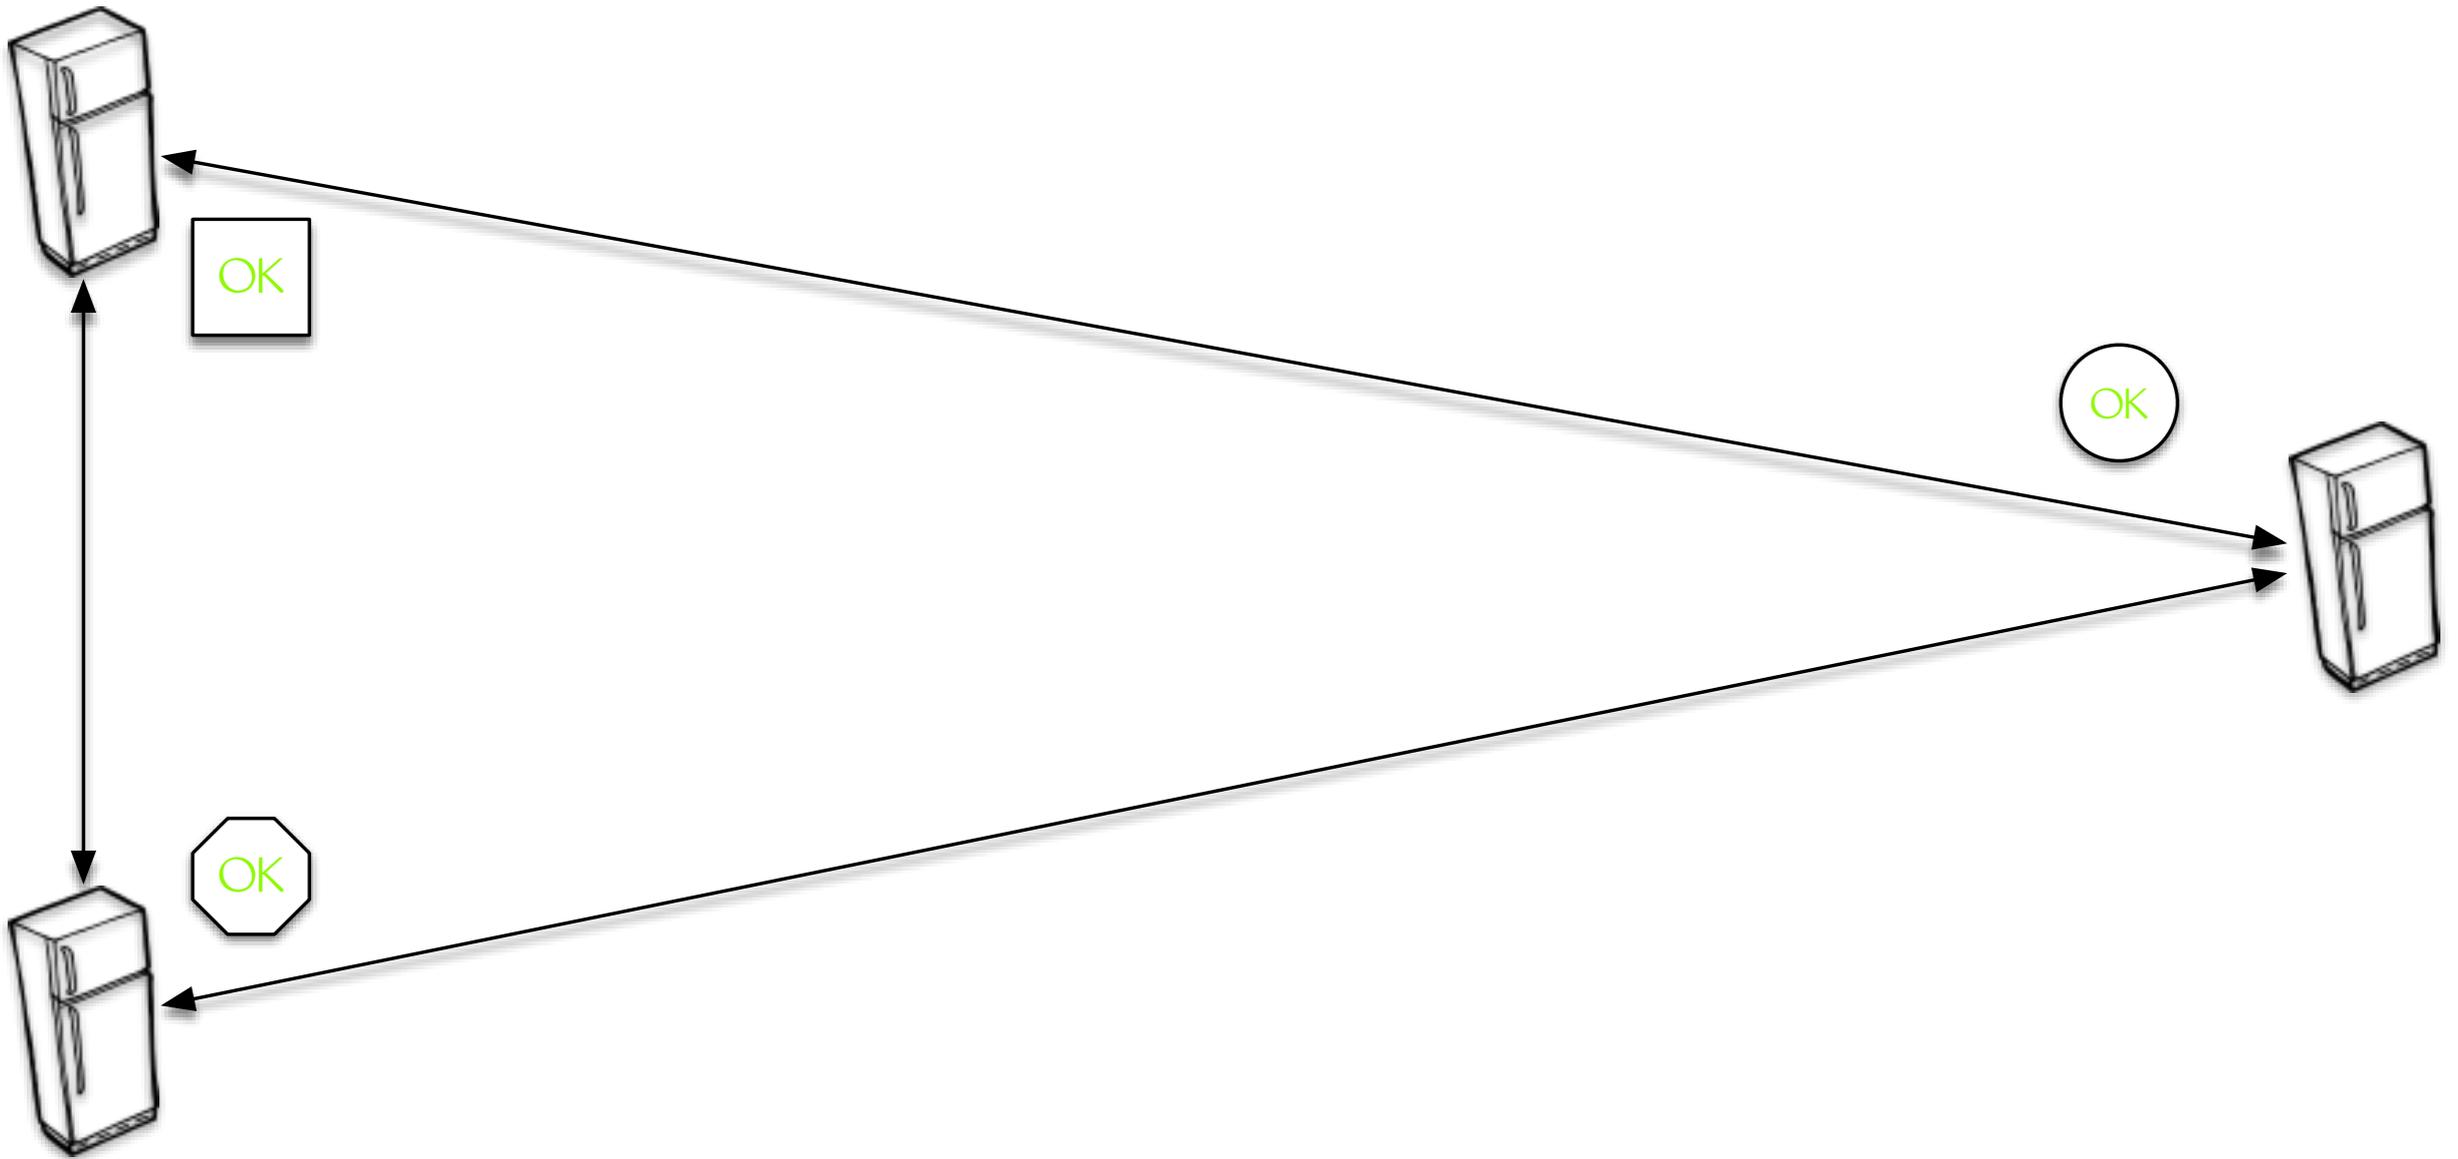


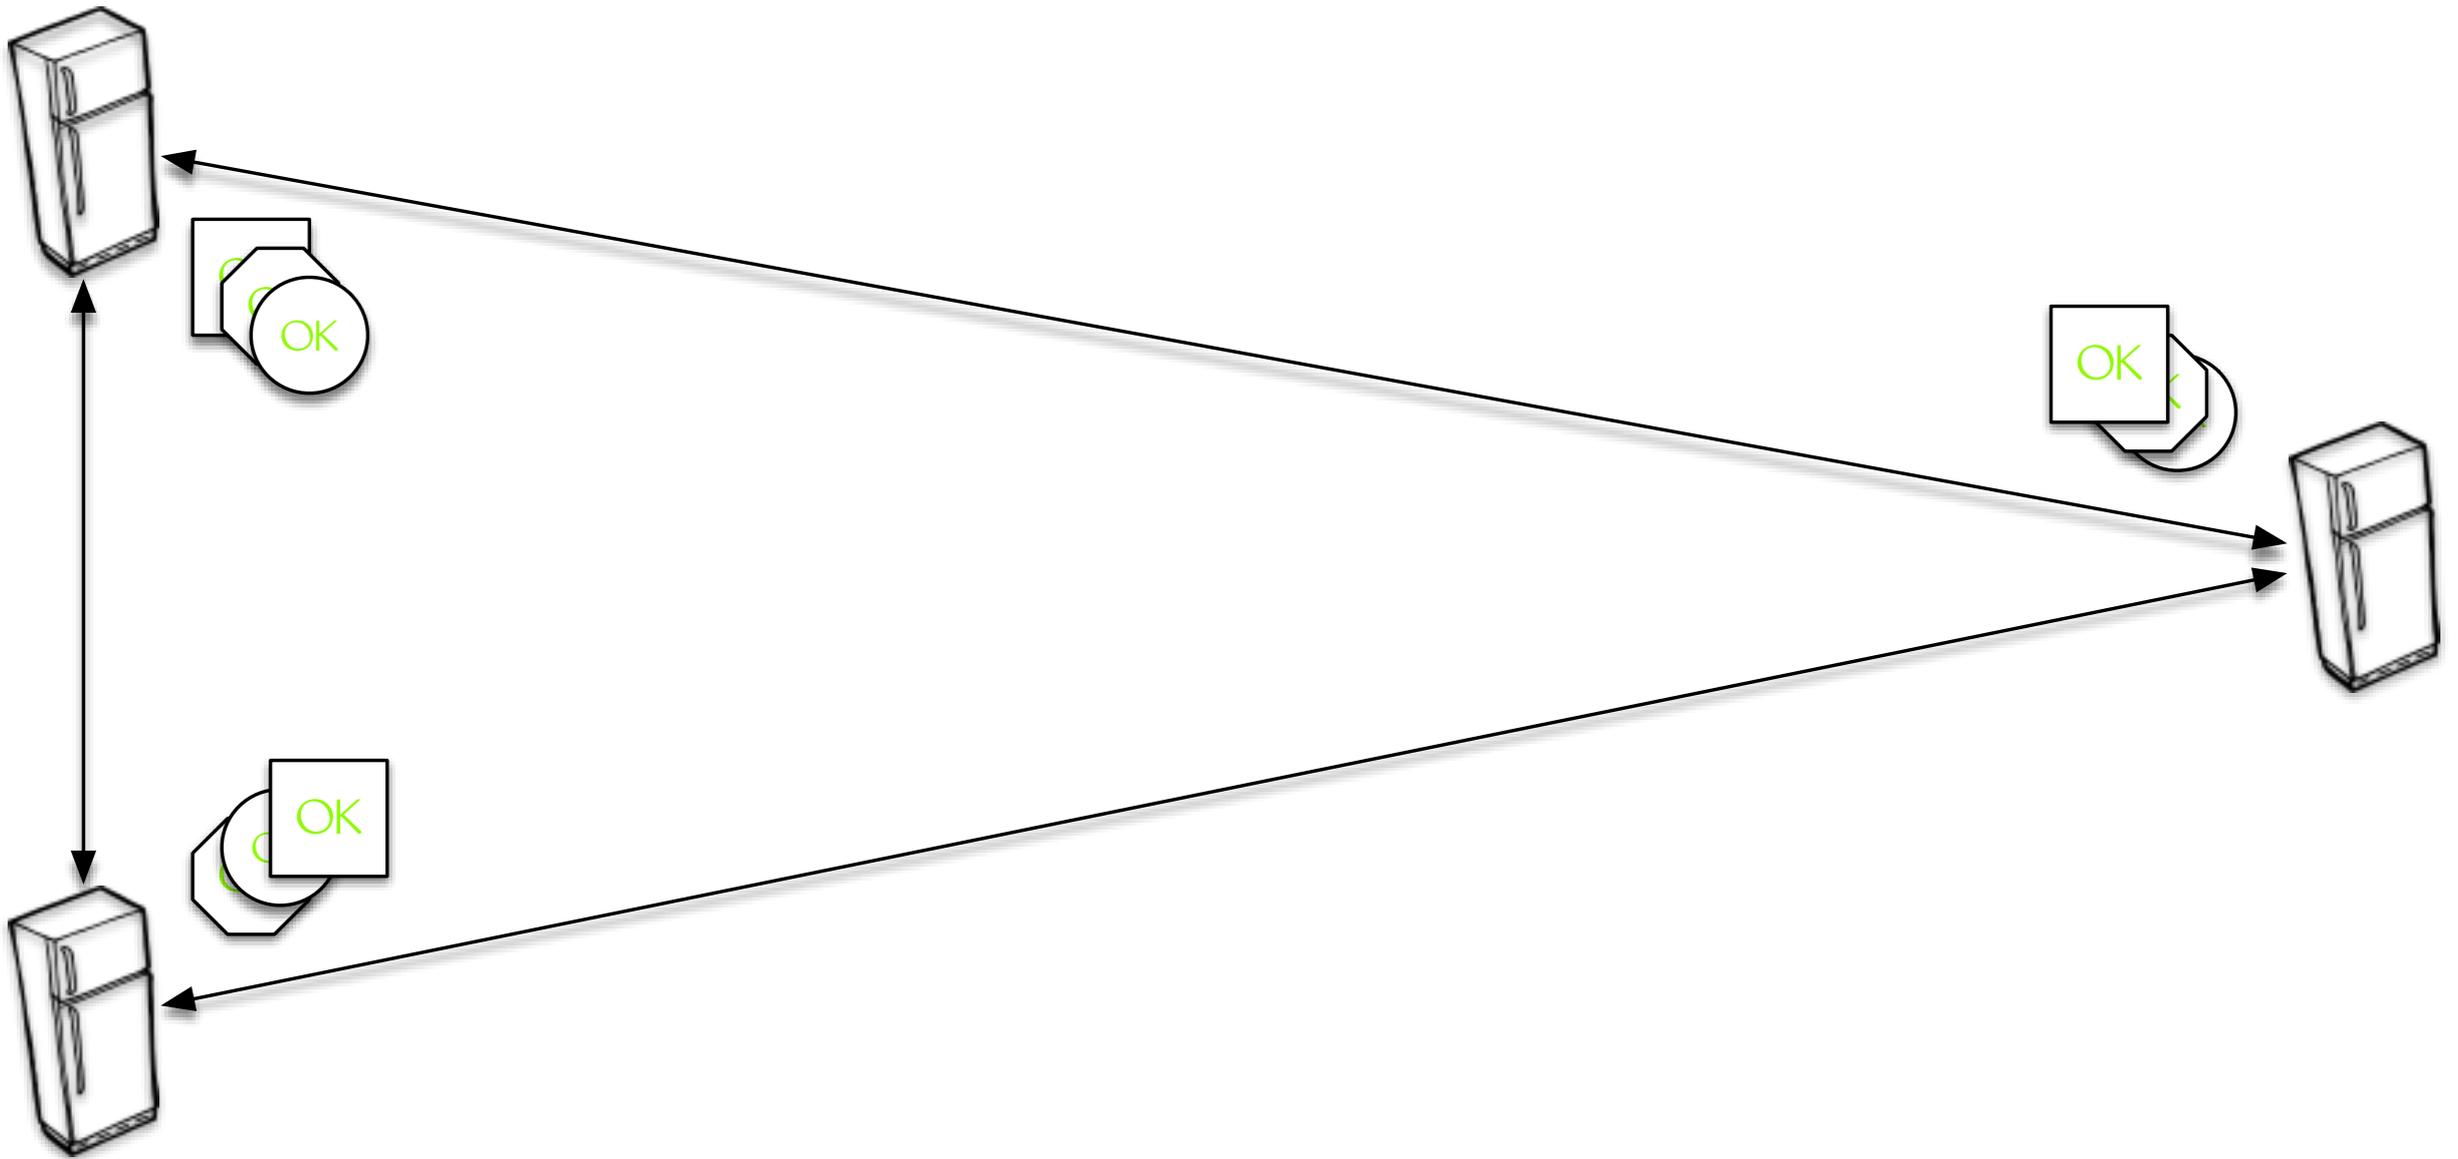


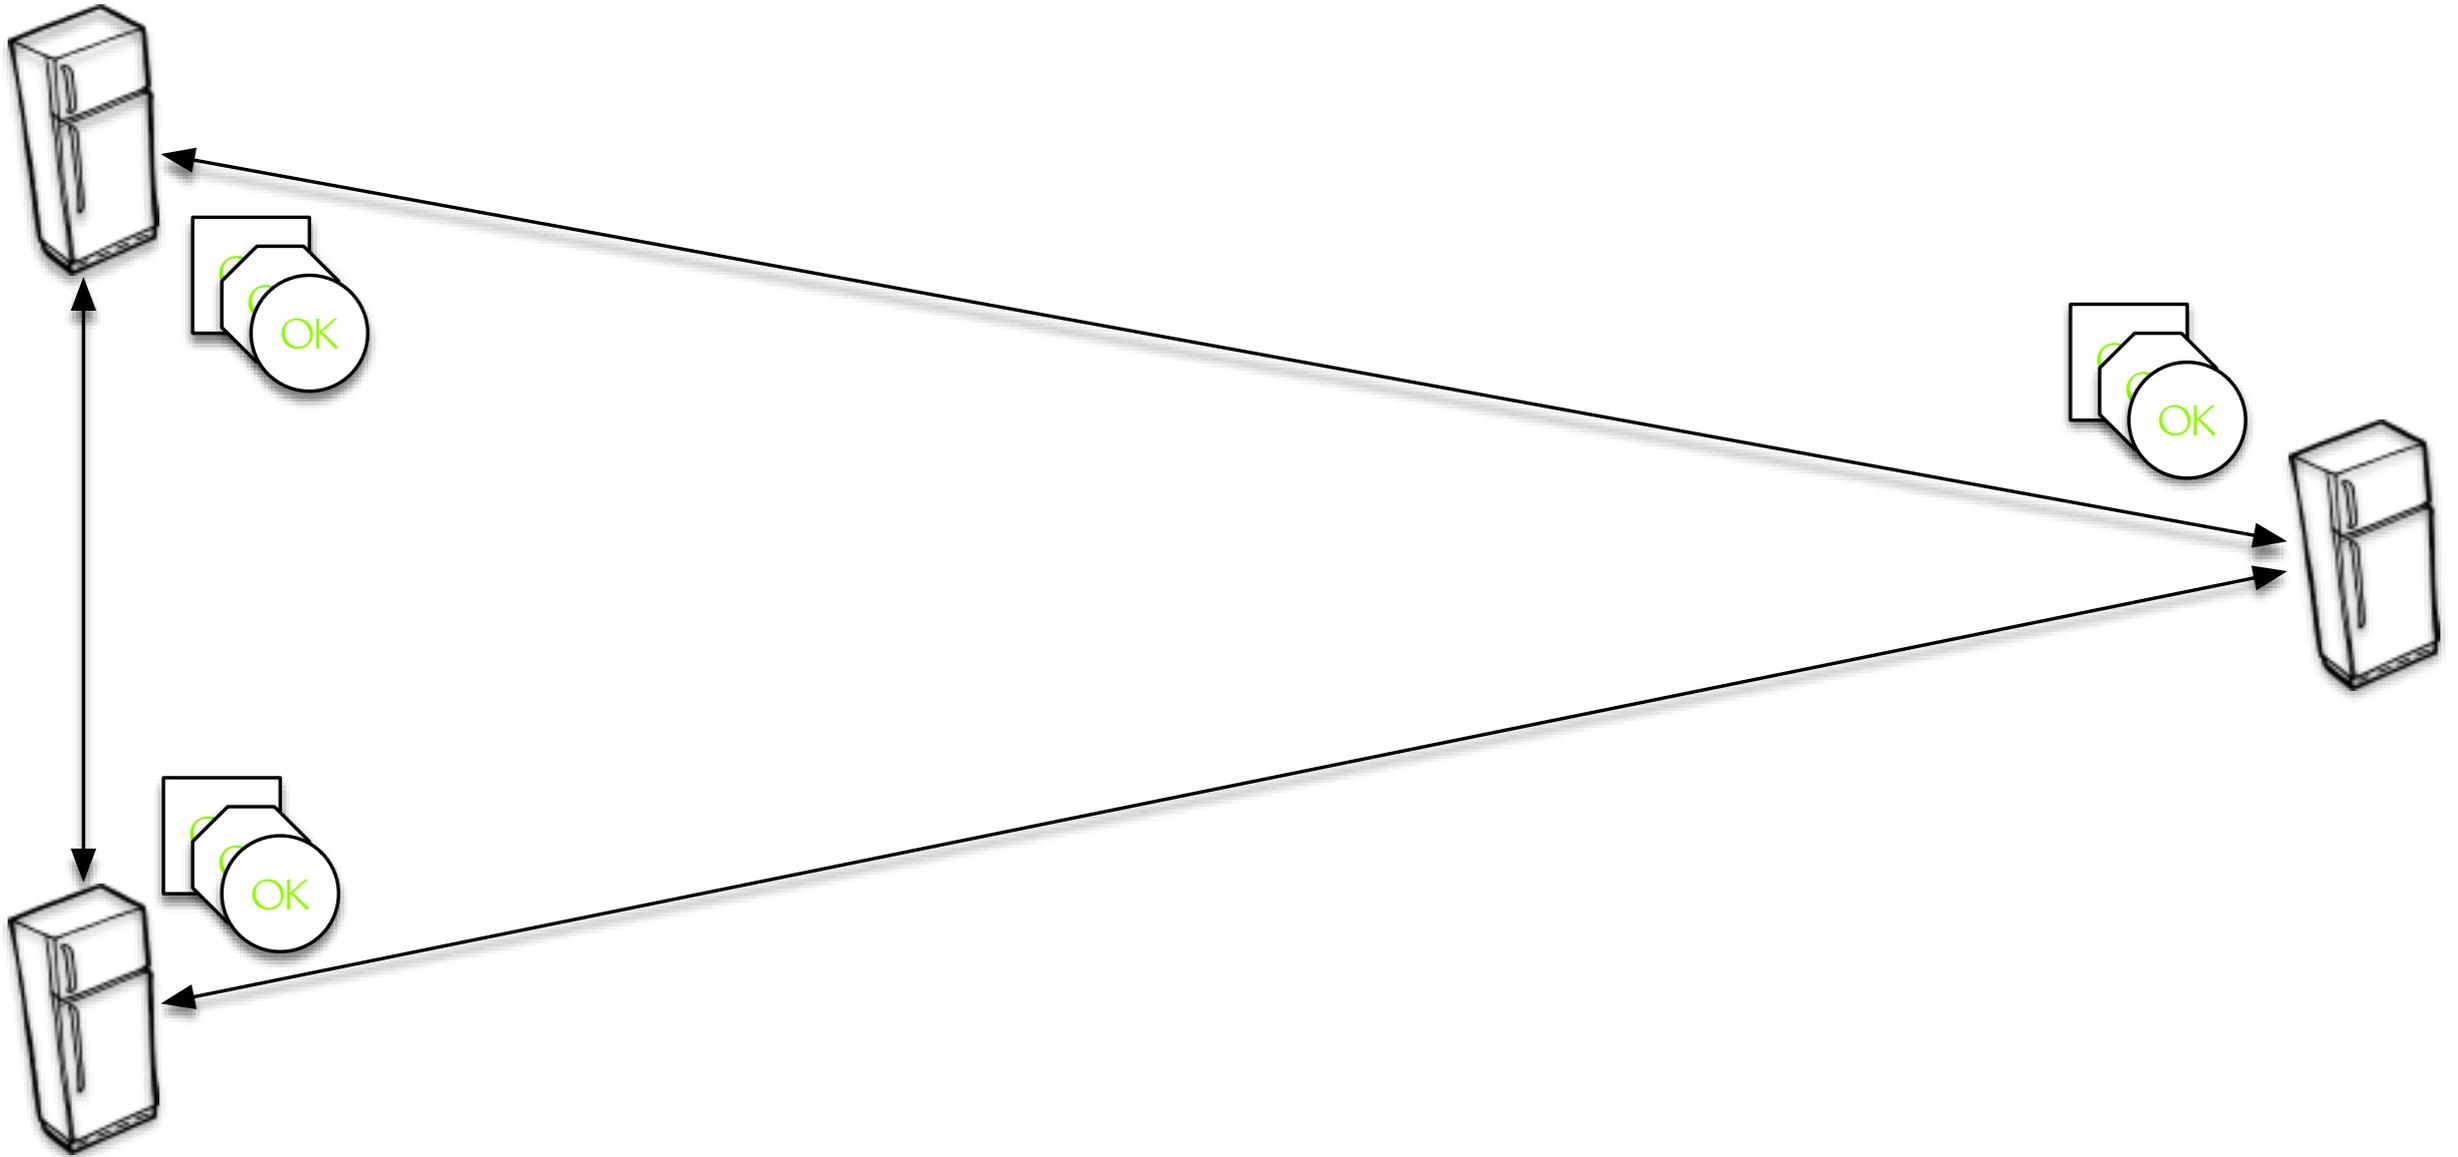


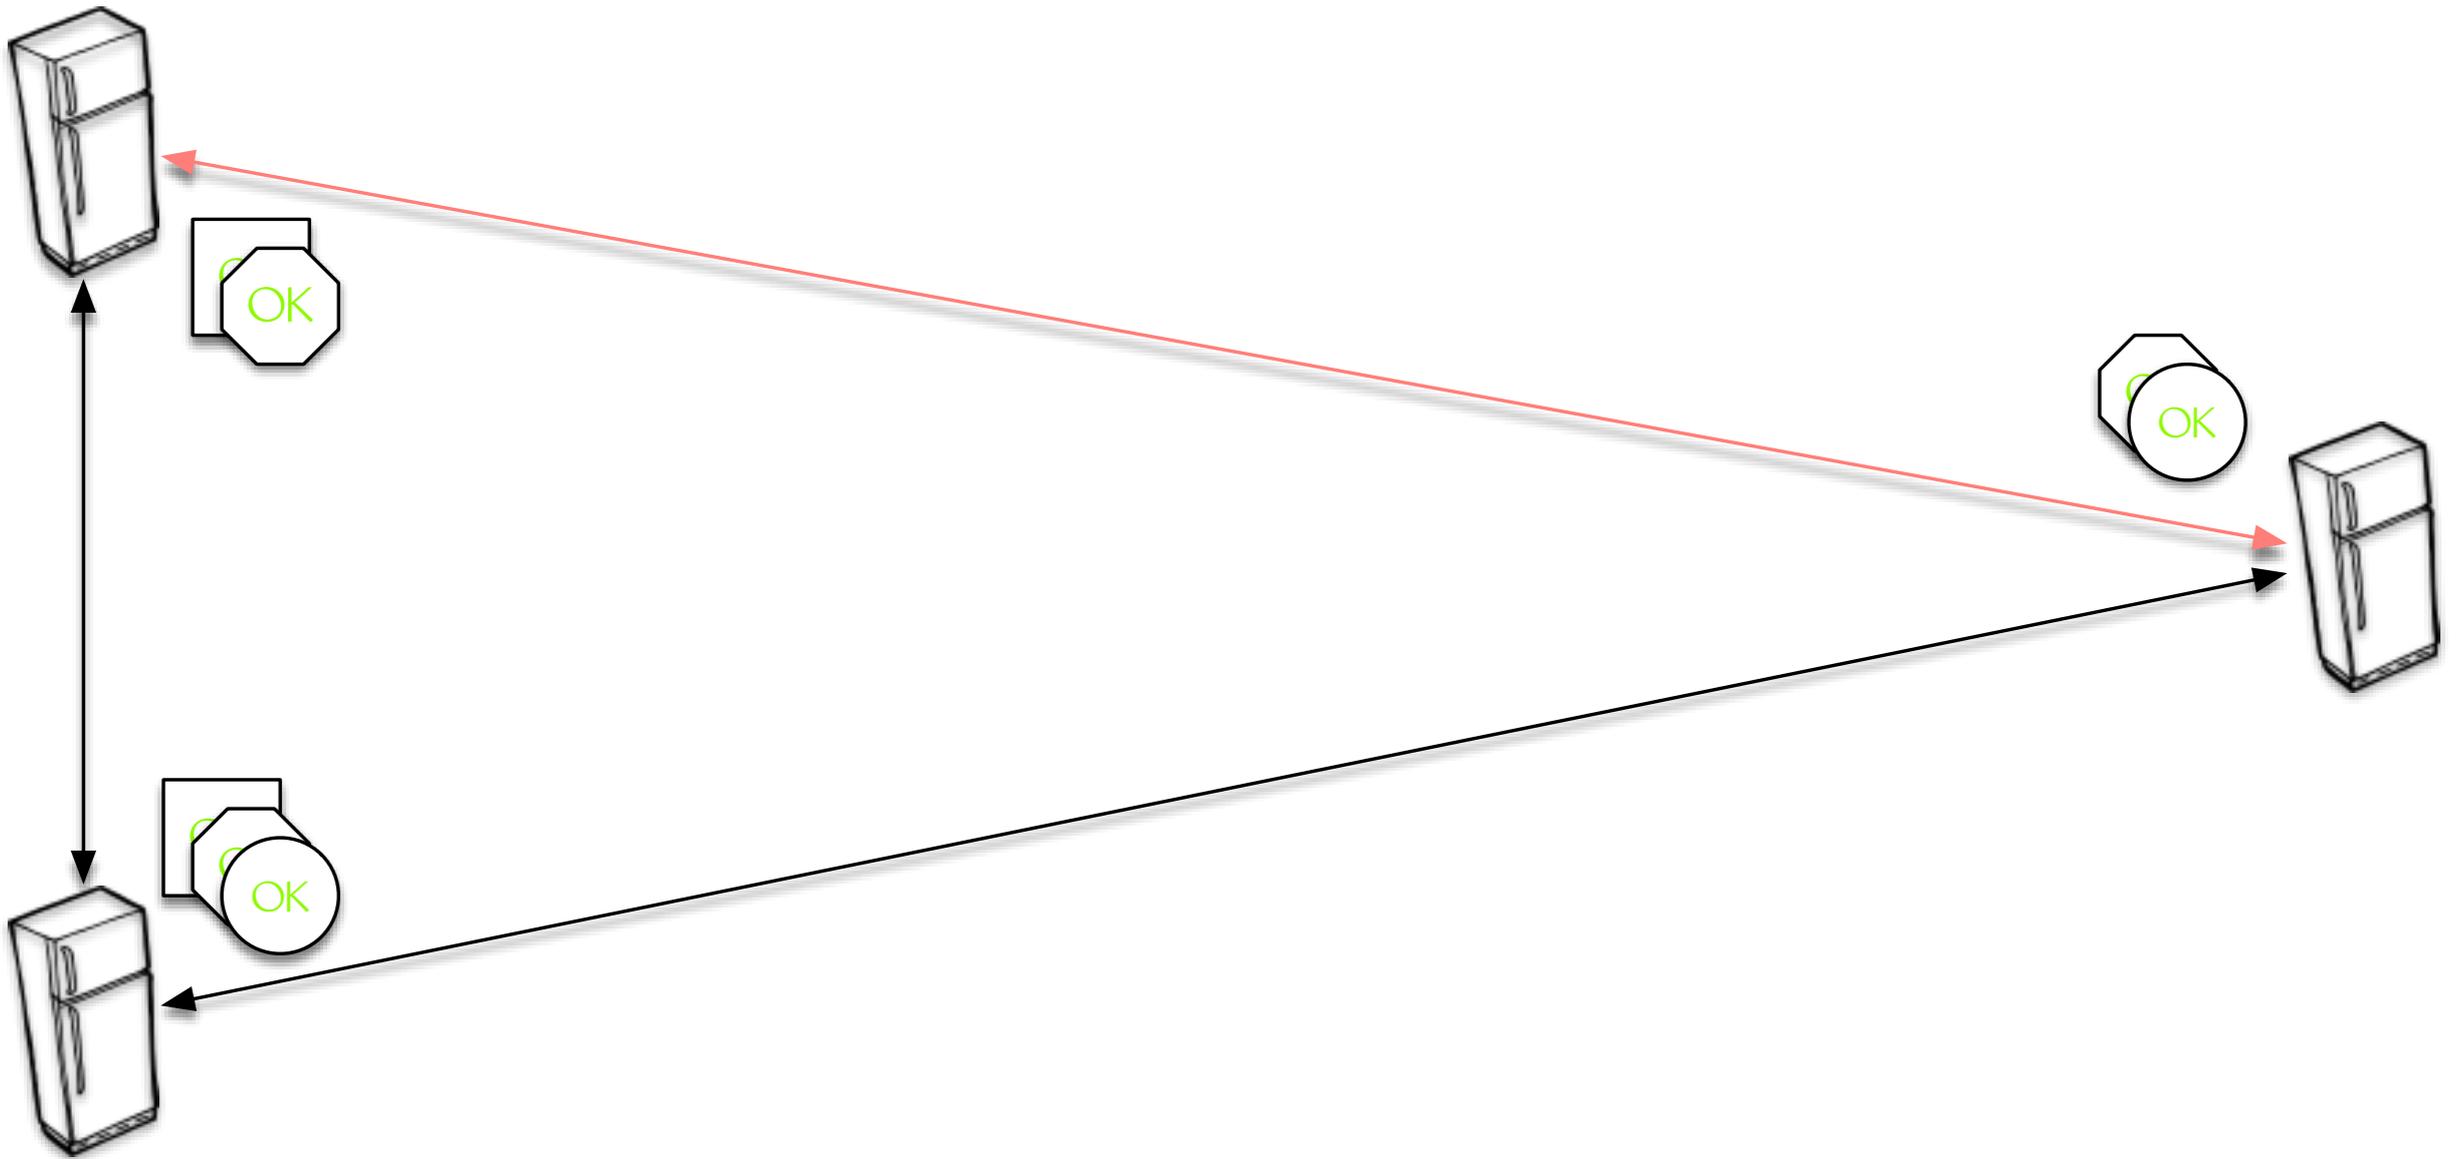
Eventual Consistency As The Model

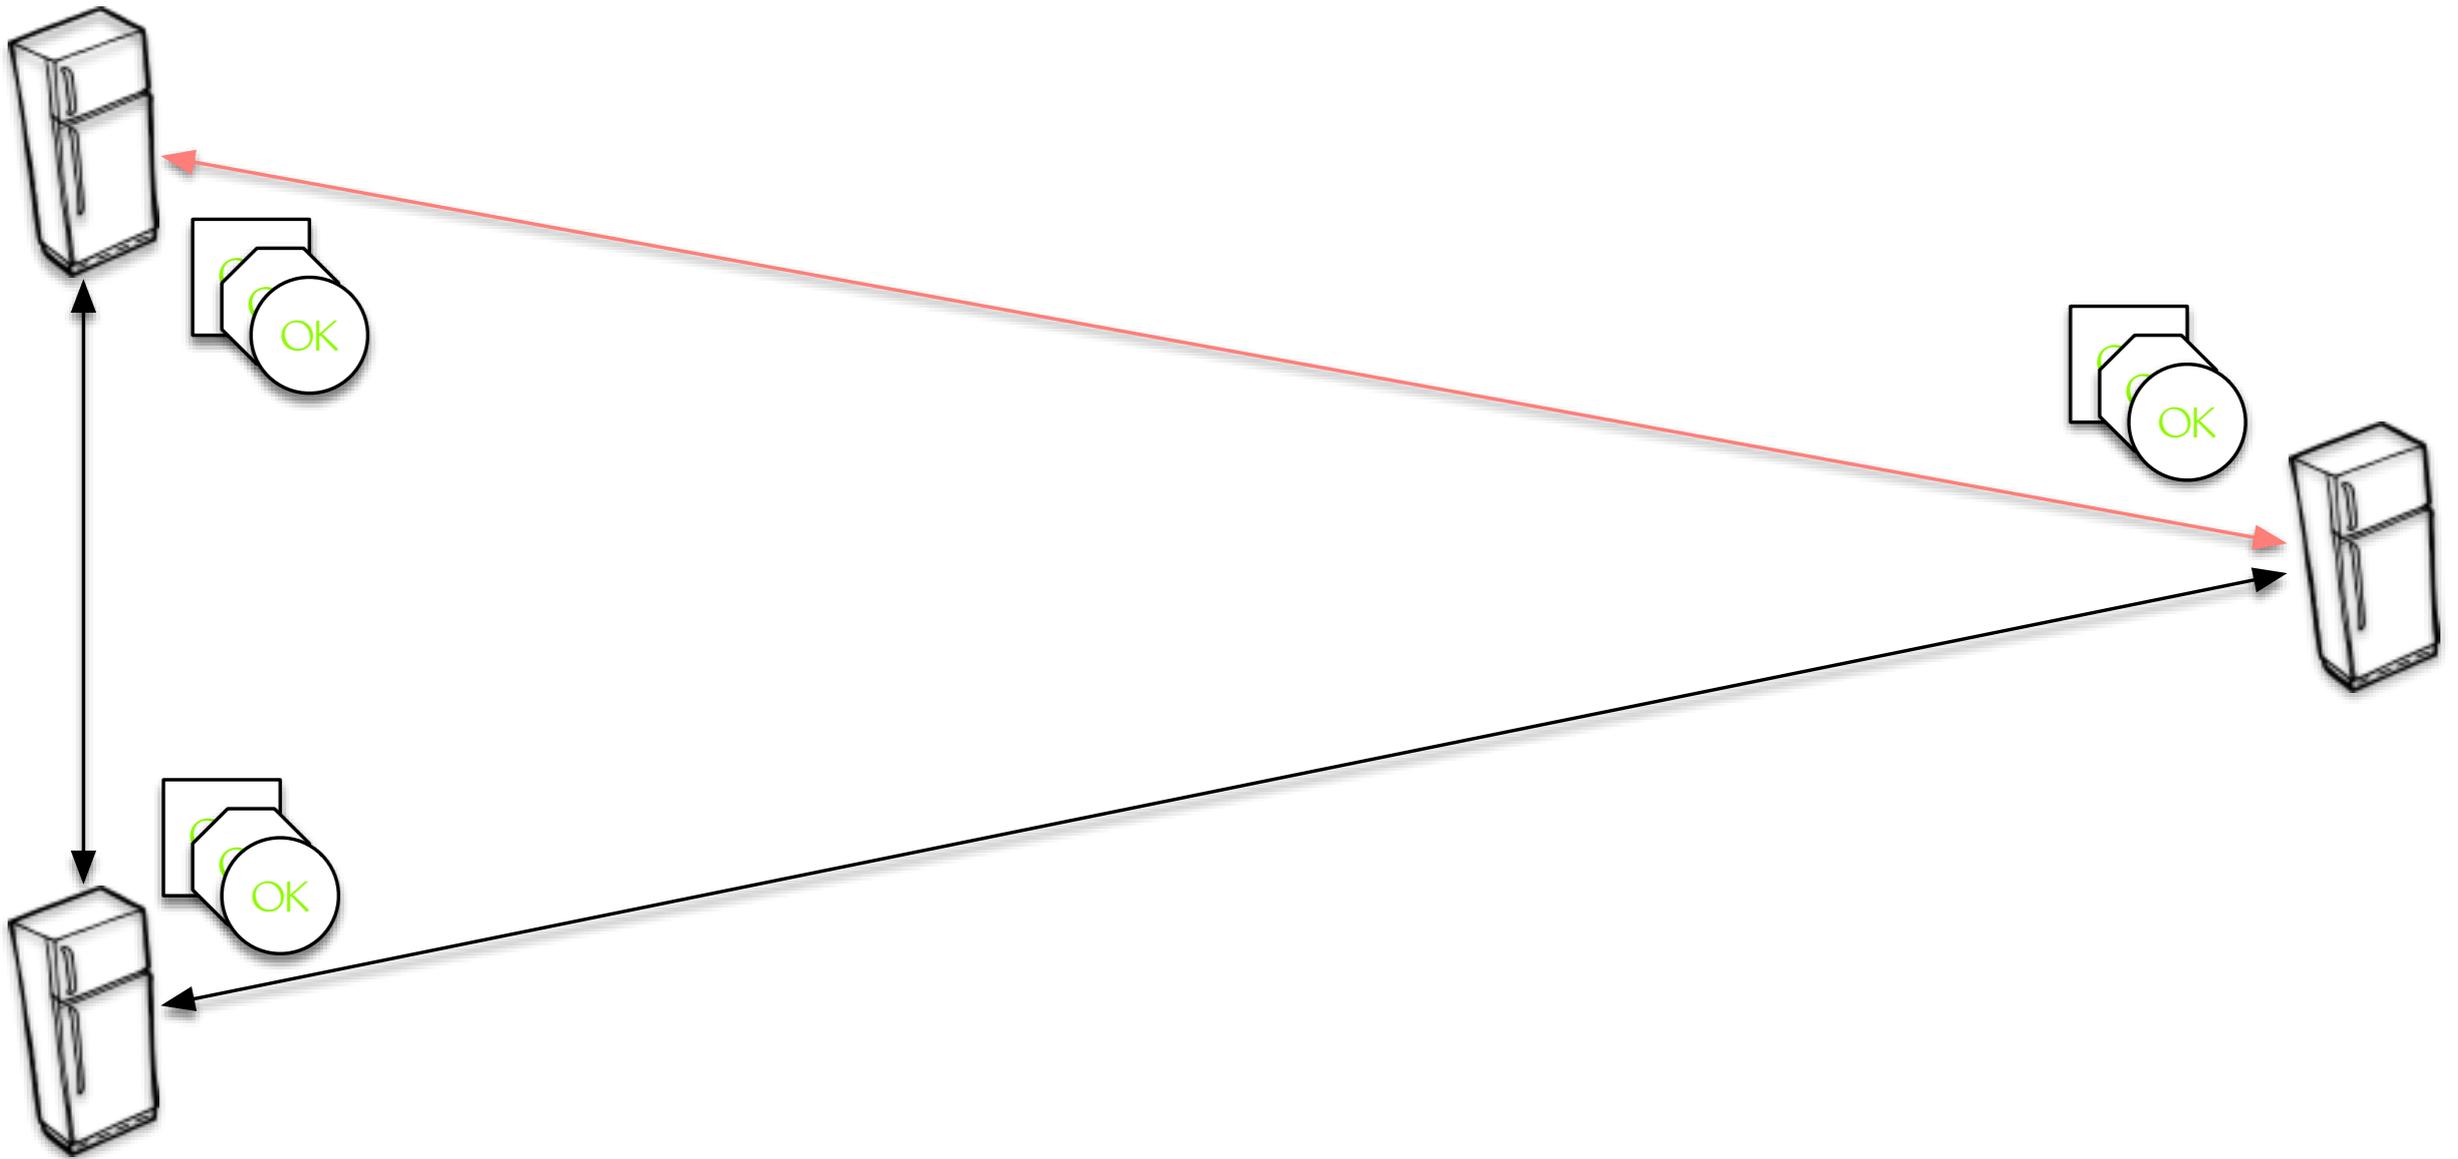
Clients
Own Their Data





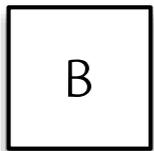
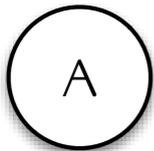
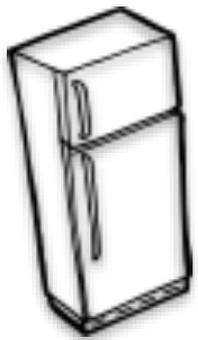
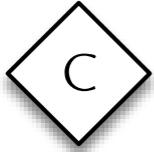
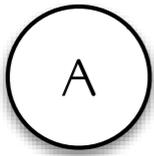
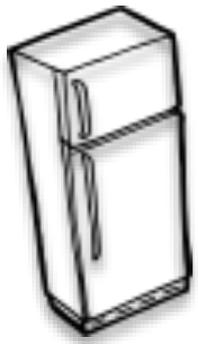


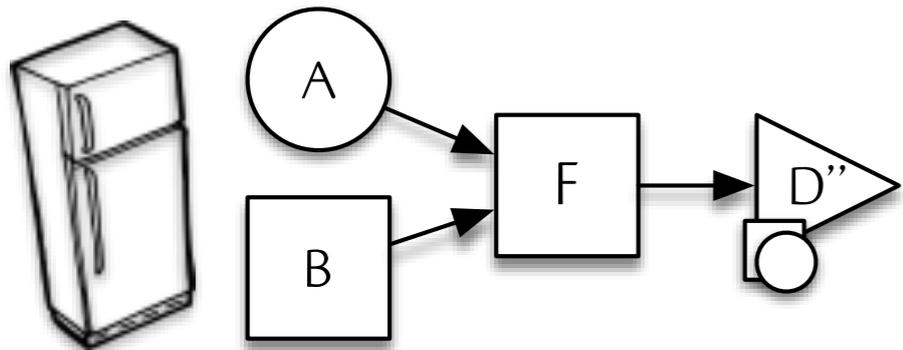
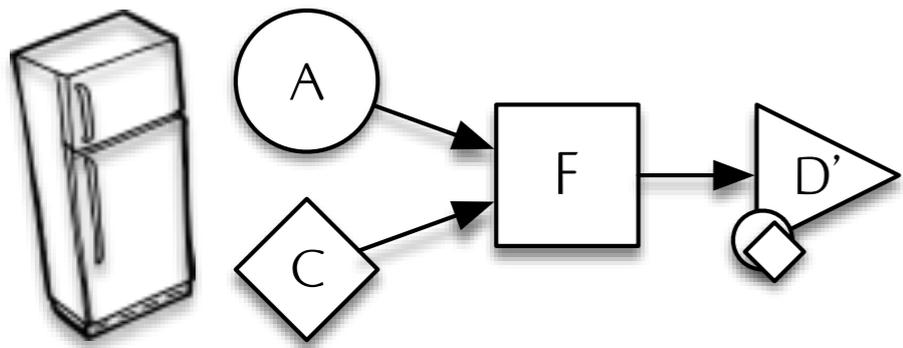


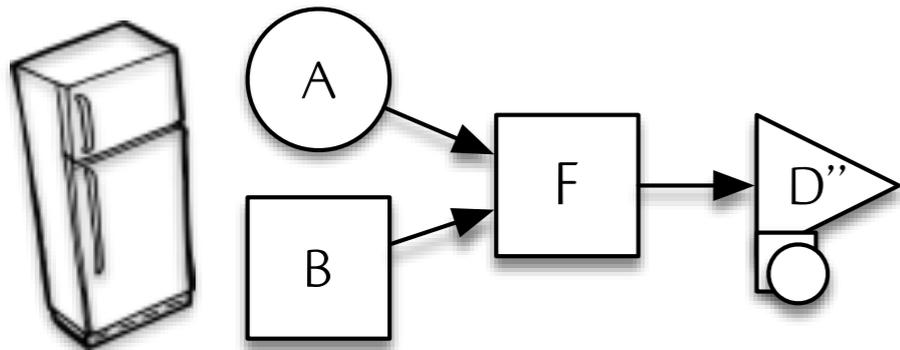
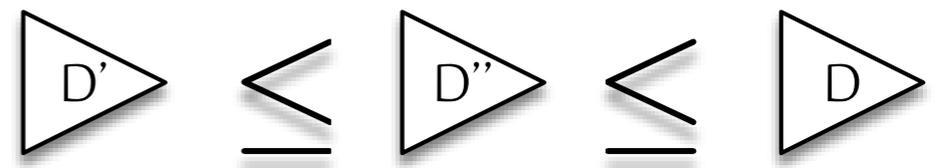
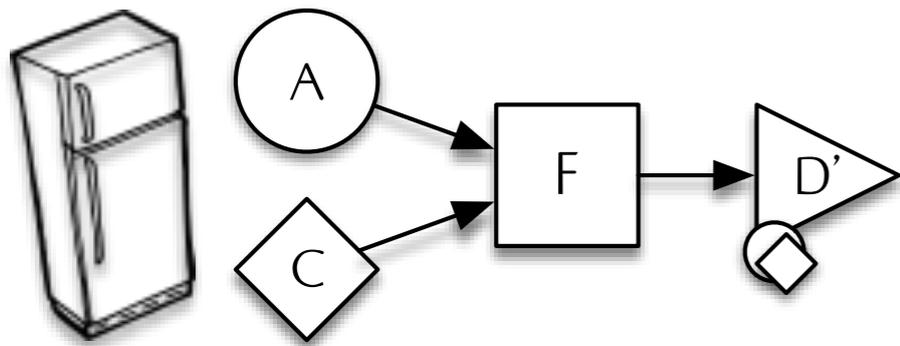


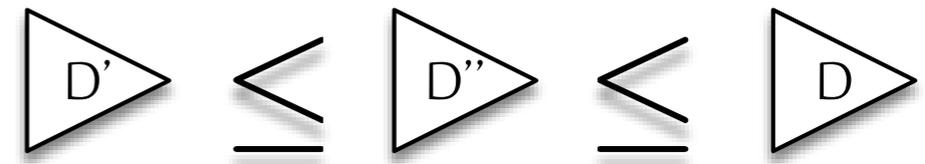
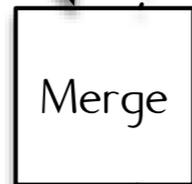
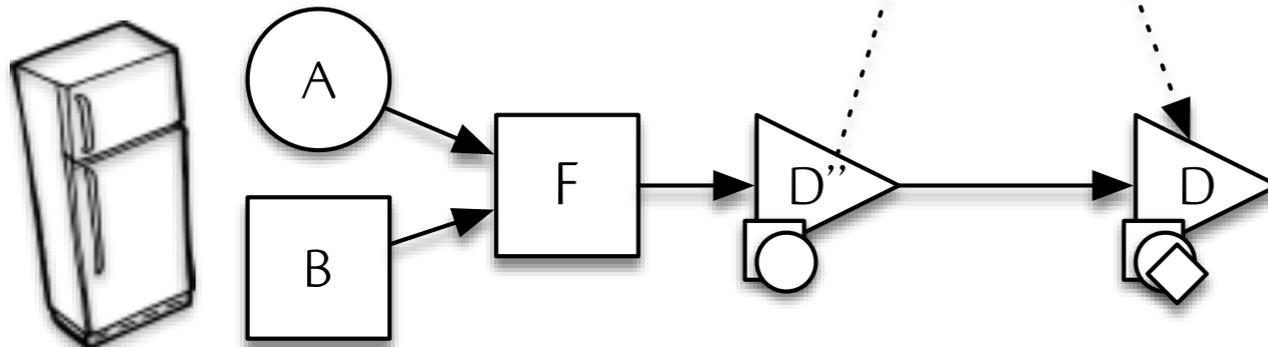
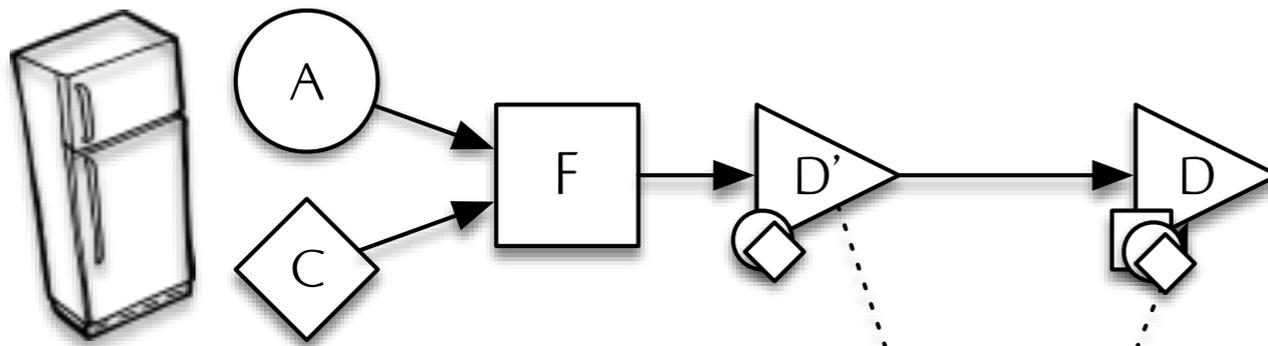
Computations

Mergability & Provenance









Example Application

Preliminary Results

Preliminary Results

- Conflict-Free Replicated Data Types
Distributed data structures designed for convergence
[Shapiro *et al.*, 2011]

Preliminary Results

- **Conflict-Free Replicated Data Types**
Distributed data structures designed for convergence
[Shapiro *et al.*, 2011]
- **Lattice Processing**
Make decisions based on results of local computation
[Meiklejohn & Van Roy, 2015]

Preliminary Results

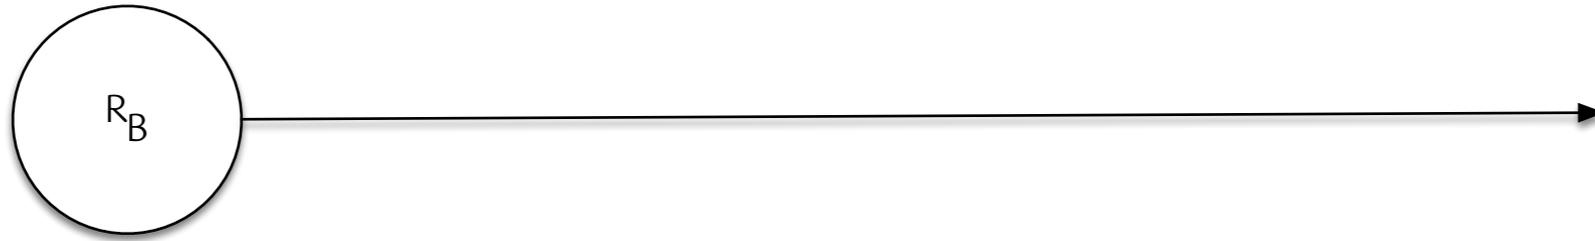
- **Conflict-Free Replicated Data Types**
Distributed data structures designed for convergence
[Shapiro *et al.*, 2011]
- **Lattice Processing**
Make decisions based on results of local computation
[Meiklejohn & Van Roy, 2015]
- **Selective Hearing**
Epidemic broadcast based runtime system
[Meiklejohn & Van Roy, 2015/2016]

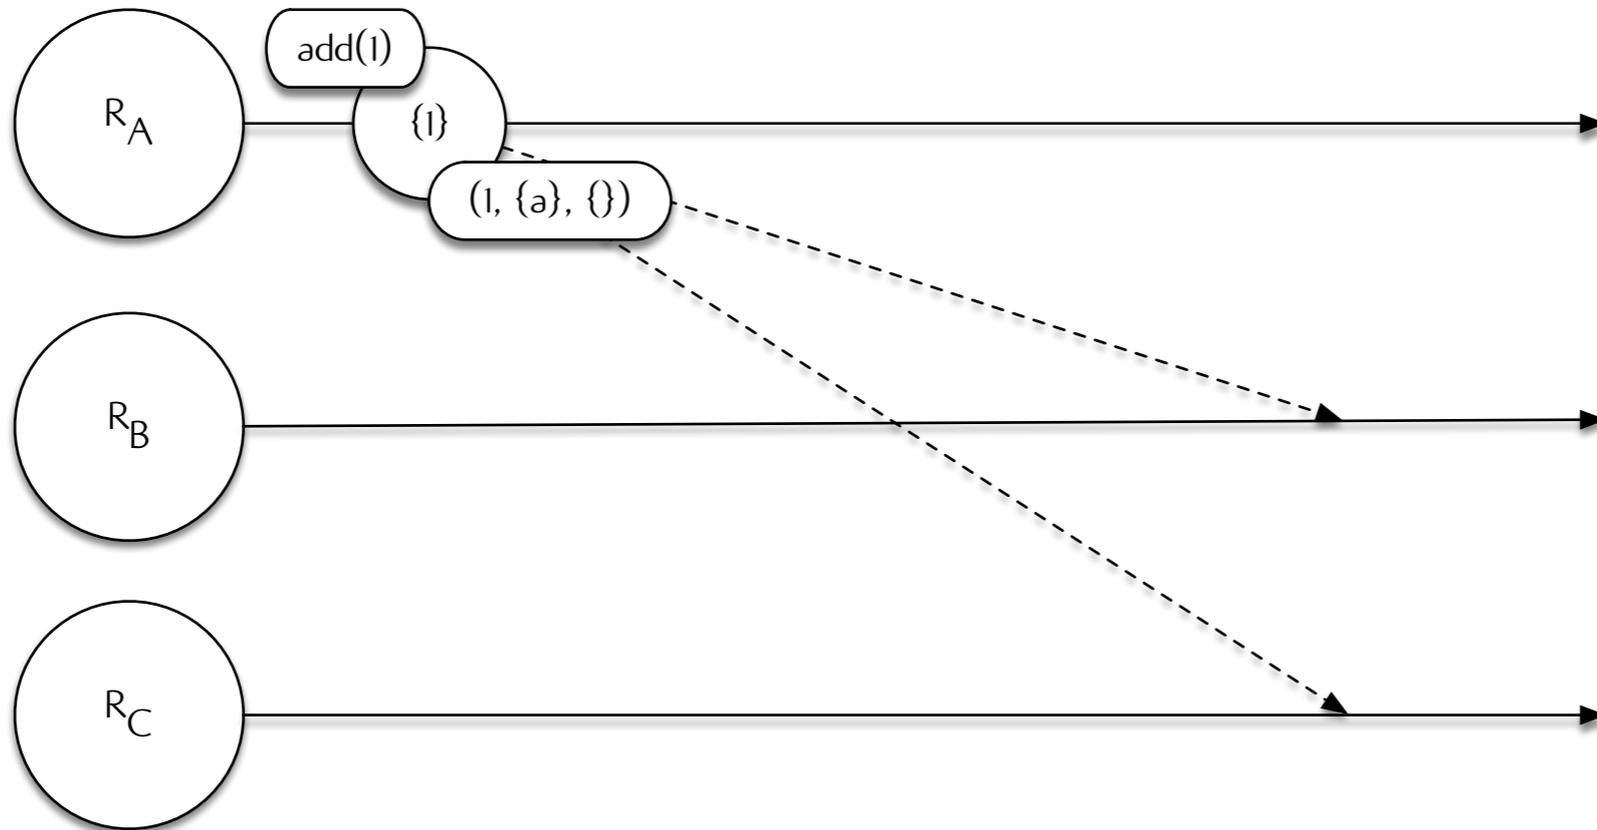
Conflict-Free Replicated Data Types

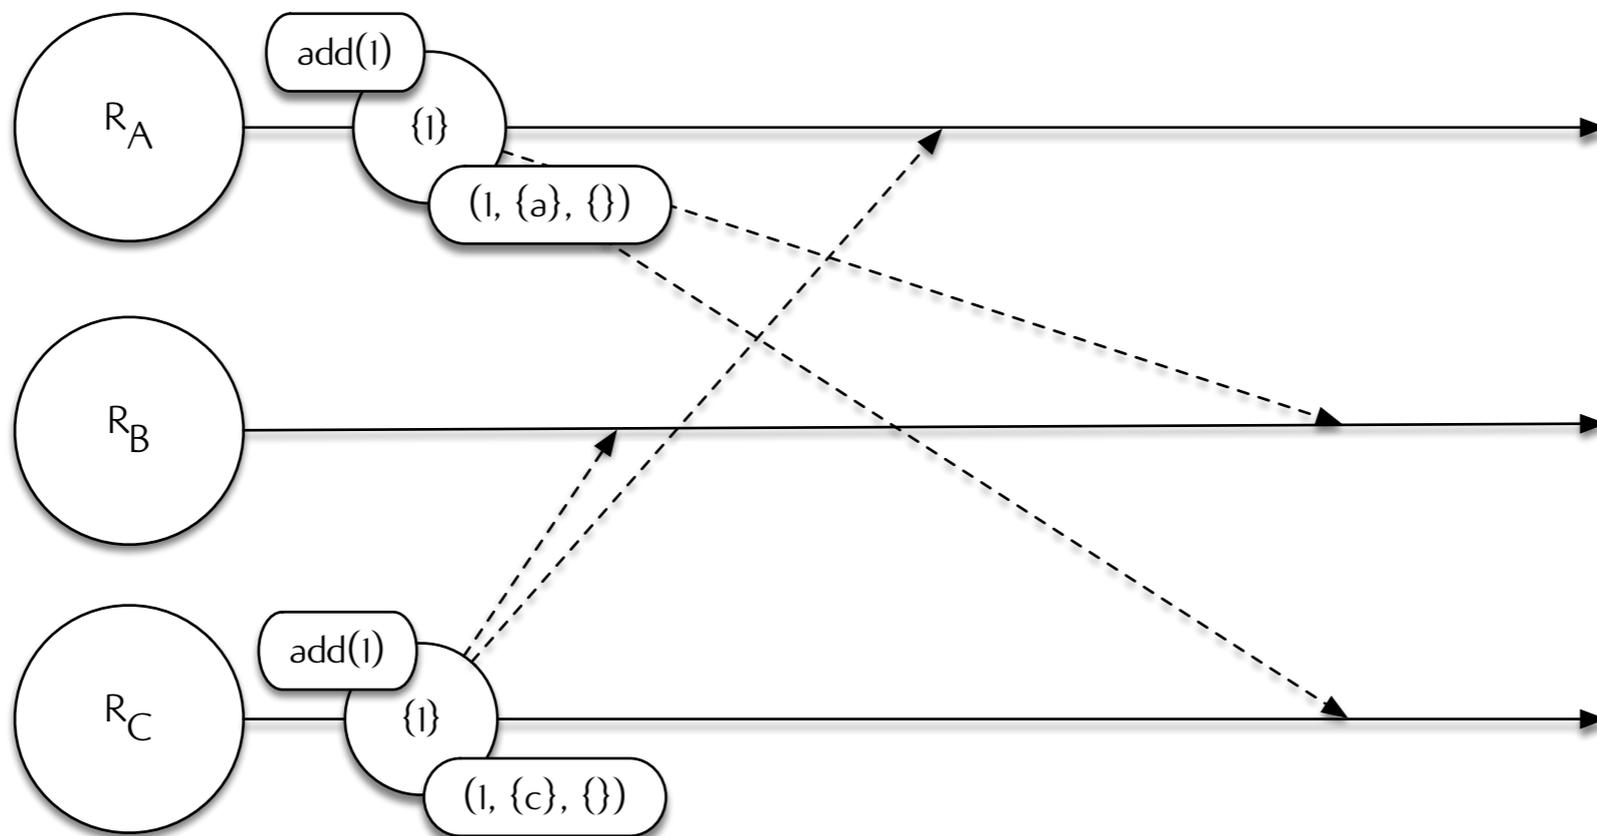
- Collection of types
Sets, counters, registers, flags, maps

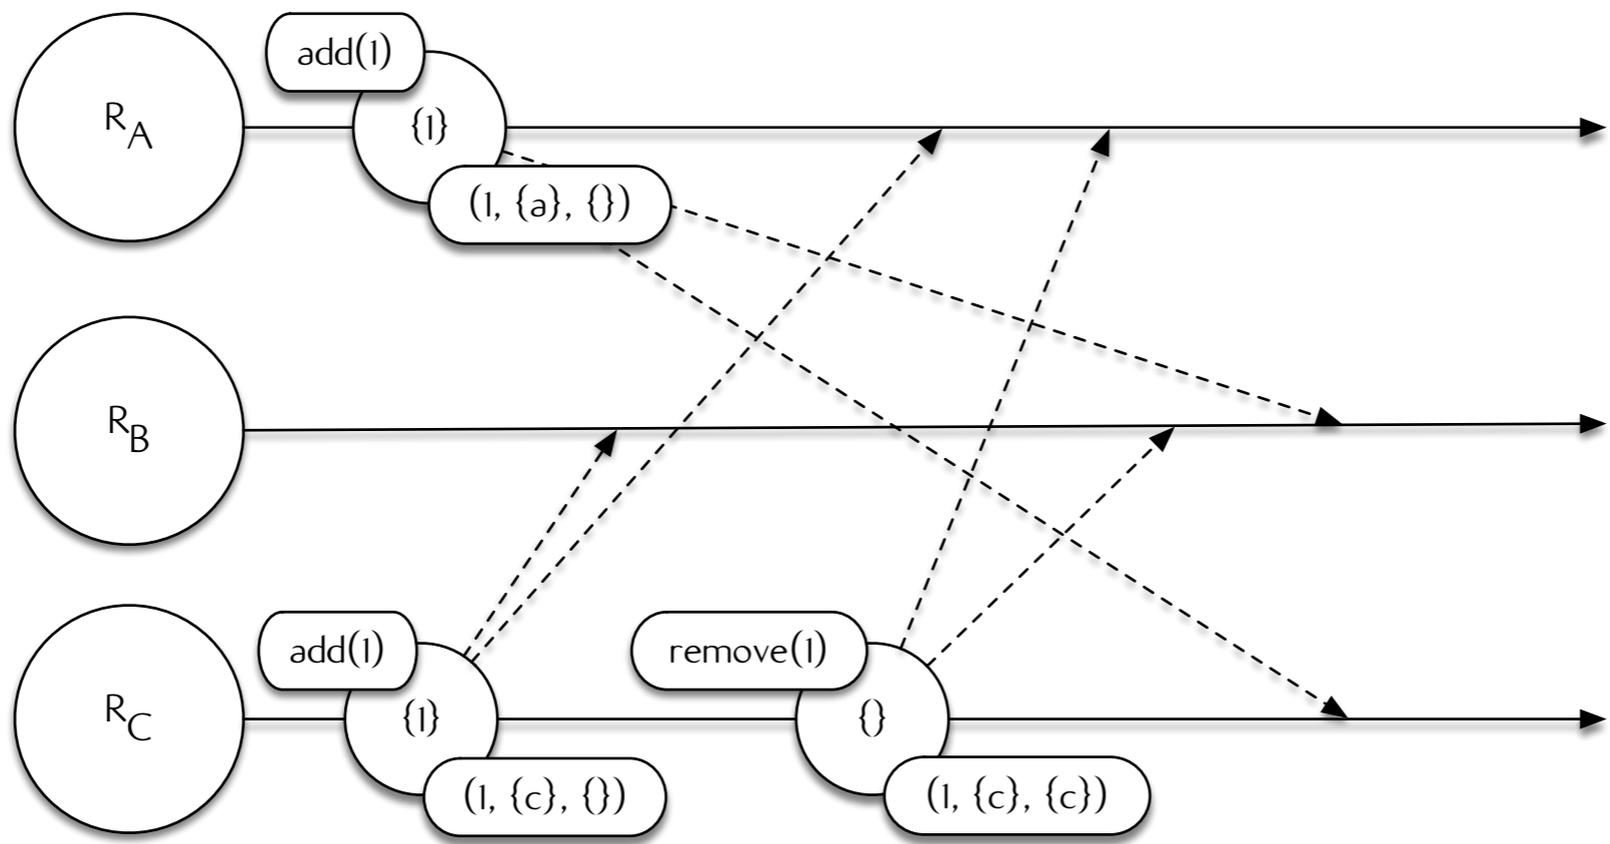
Conflict-Free Replicated Data Types

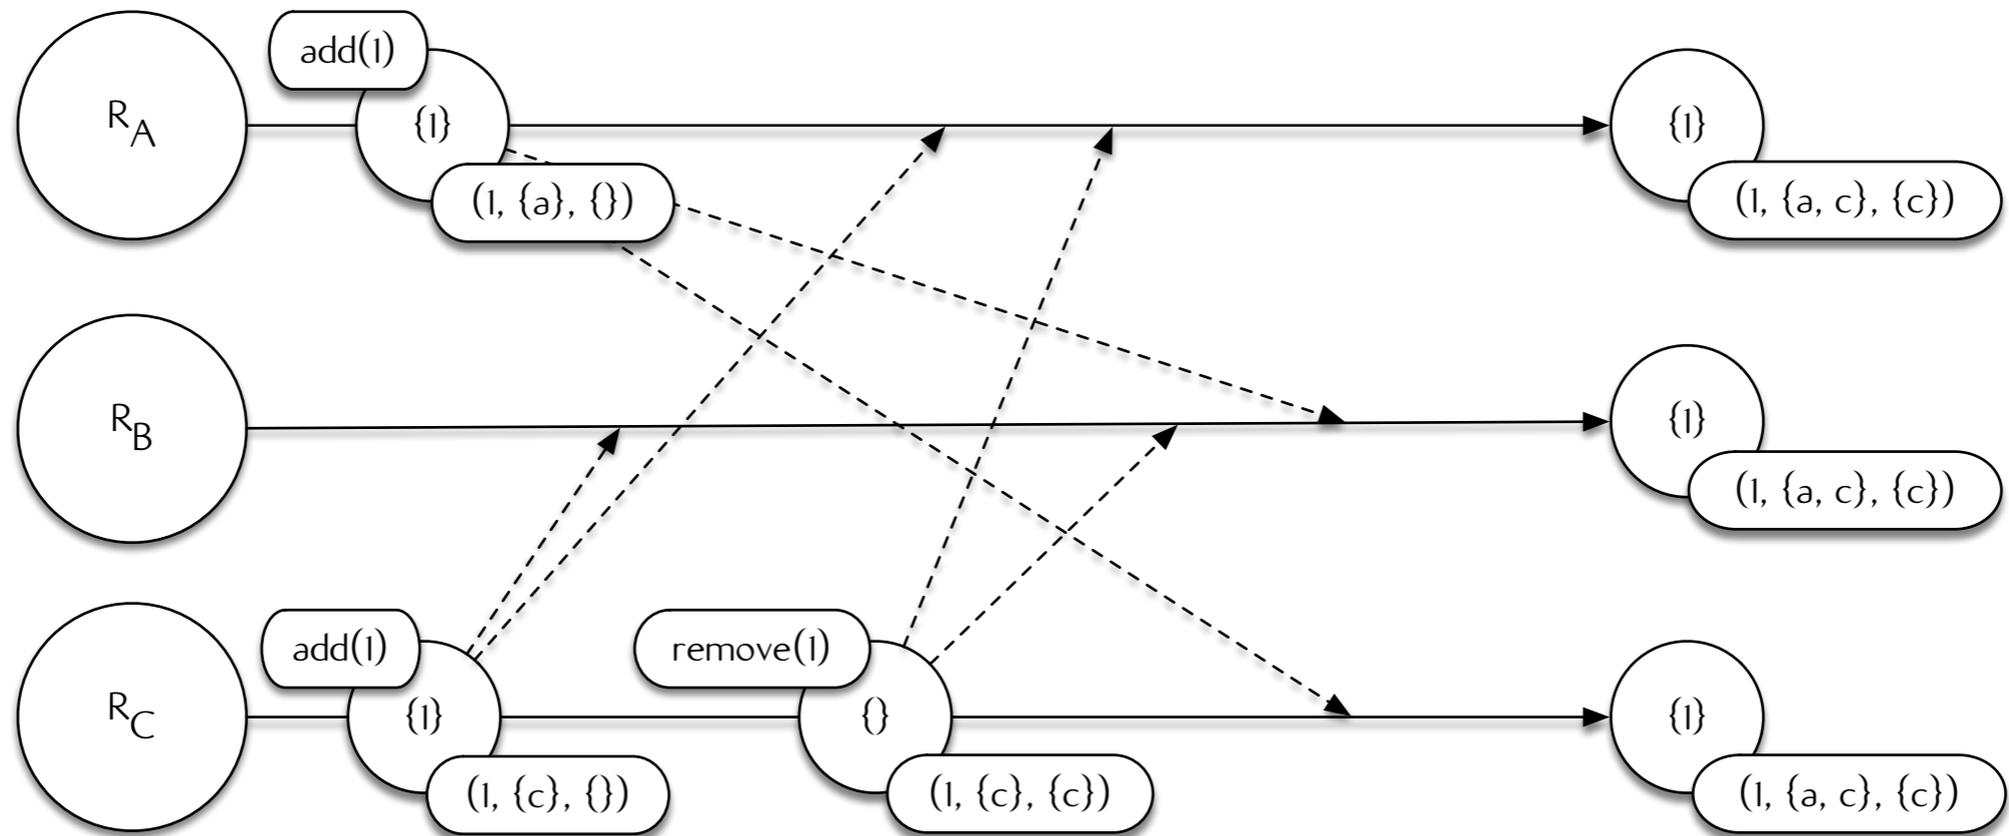
- **Collection of types**
Sets, counters, registers, flags, maps
- **Strong Eventual Consistency**
Objects that receive the same updates, regardless of order, will reach equivalent state











Lattice Processing

- Distributed, deterministic dataflow
Distributed, dataflow programming model

Lattice Processing

- **Distributed, deterministic dataflow**
Distributed, dataflow programming model
- **Convergent data structures**
Data abstraction is the CRDT

Lattice Processing

- **Distributed, deterministic dataflow**
Distributed, dataflow programming model
- **Convergent data structures**
Data abstraction is the CRDT
- **Enables composition**
Composition preserves SEC

```
%% Create initial set.  
S1 = declare(set),  
  
%% Add elements to initial set and update.  
update(S1, {add, [1,2,3]}),  
  
%% Create second set.  
S2 = declare(set),  
  
%% Apply map operation between S1 and S2.  
map(S1, fun(X) -> X * 2 end, S2).
```

```
%% Create initial set.  
S1 = declare(set),  
  
%% Add elements to initial set and update.  
update(S1, {add, [1,2,3]}),  
  
%% Create second set.  
S2 = declare(set),  
  
%% Apply map operation between S1 and S2.  
map(S1, fun(X) -> X * 2 end, S2).
```

```
%% Create initial set.  
S1 = declare(set),  
  
%% Add elements to initial set and update.  
update(S1, {add, [1,2,3]}),  
  
%% Create second set.  
S2 = declare(set),  
  
%% Apply map operation between S1 and S2.  
map(S1, fun(X) -> X * 2 end, S2).
```

```
%% Create initial set.  
S1 = declare(set),  
  
%% Add elements to initial set and update.  
update(S1, {add, [1,2,3]}),  
  
%% Create second set.  
S2 = declare(set),  
  
%% Apply map operation between S1 and S2.  
map(S1, fun(X) -> X * 2 end, S2).
```

```
%% Create initial set.  
S1 = declare(set),  
  
%% Add elements to initial set and update.  
update(S1, {add, [1,2,3]}),  
  
%% Create second set.  
S2 = declare(set),  
  
%% Apply map operation between S1 and S2.  
map(S1, fun(X) -> X * 2 end, S2).
```

Selective Hearing

- Epidemic broadcast protocol
Runtime system for application state & scope

Selective Hearing

- Epidemic broadcast protocol
Runtime system for application state & scope
- Peer-to-peer dissemination
Pairwise synchronization between peers
without a central coordinator

Selective Hearing

- **Epidemic broadcast protocol**
Runtime system for application state & scope
- **Peer-to-peer dissemination**
Pairwise synchronization between peers
without a central coordinator
- **No ordering guarantees on messages**
Programming model can tolerate message
reordering and duplication

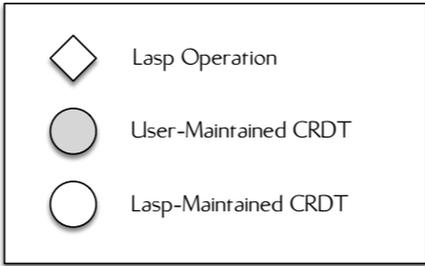
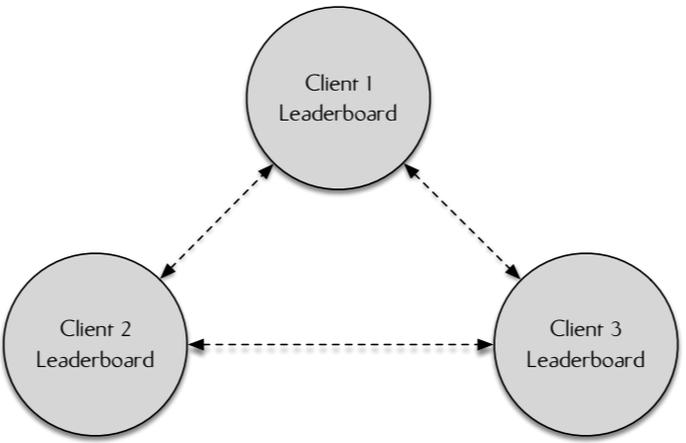
What can we build?
Leaderboard

Leaderboard

- Mobile game platform
Local leaderboard tracking top- k
highest scored games

Leaderboard

- **Mobile game platform**
Local leaderboard tracking top- k highest scored games
- **Clients will go offline**
Clients have limited connectivity and the system still needs to make progress while clients are offline



Leaderboard

- Peer-to-peer dissemination
Nodes periodically “merge” their state with a random peer

Leaderboard

- **Peer-to-peer dissemination**
Nodes periodically “merge” their state with a random peer
- **Complexity in the data type**
Each node tracks a top- k set of its own games in a bounded set

```
%% Create a leaderboard datatype.  
L = declare({top_k, [2]}).
```

```
%% Update leaderboard.  
update({set, Name, Score}, L).
```

```
%% Create a leaderboard datatype.
```

```
L = declare({top_k, [2]}).
```

```
%% Update leaderboard.
```

```
update({set, Name, Score}, L).
```

```
%% Create a leaderboard datatype.  
L = declare({top_k, [2]}).
```

```
%% Update leaderboard.  
update({set, Name, Score}, L).
```

What if we want to enhance
the behavior?

What if we want to enhance
the behavior?

**Without the creation of a
new datatype**

What can we build?

Per-User Leaderboard

Per-User Leaderboard

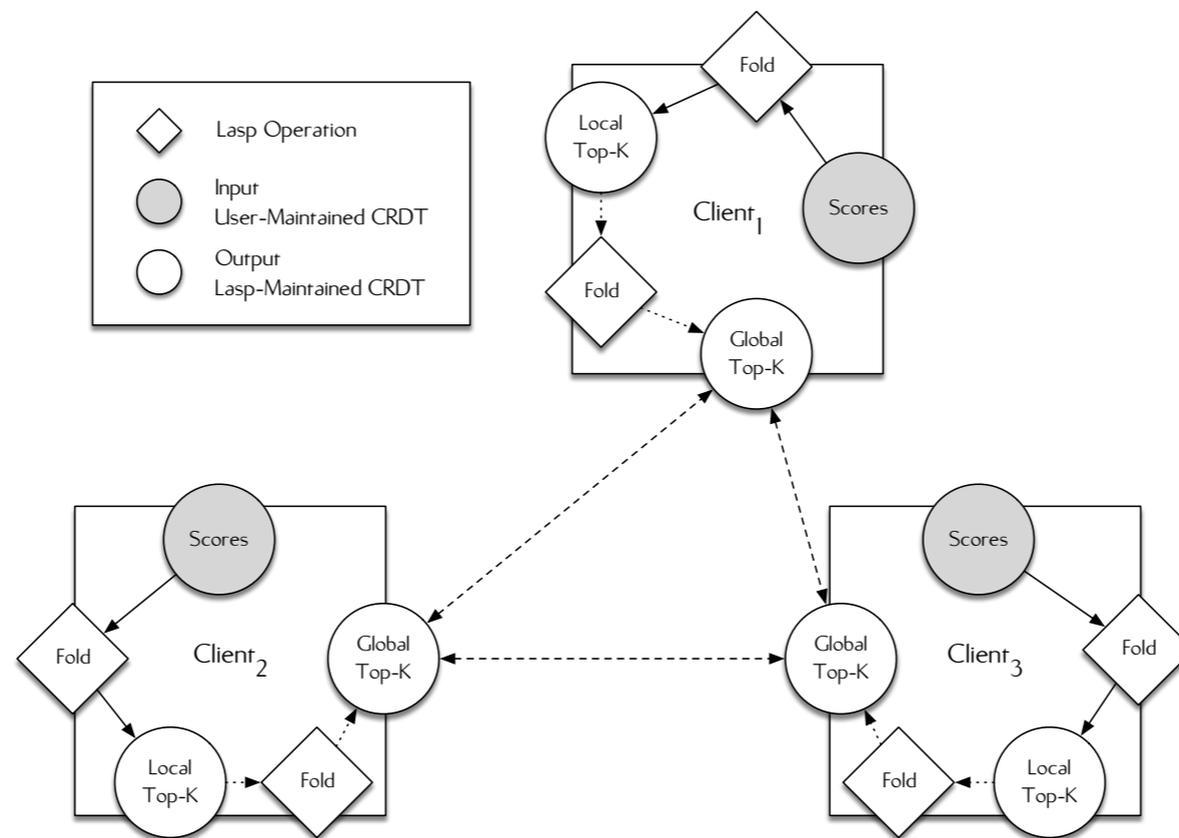
- Enhance existing design
Only the top score for each user at each device

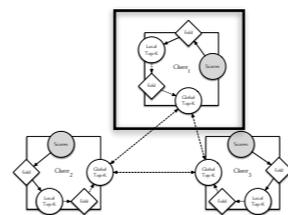
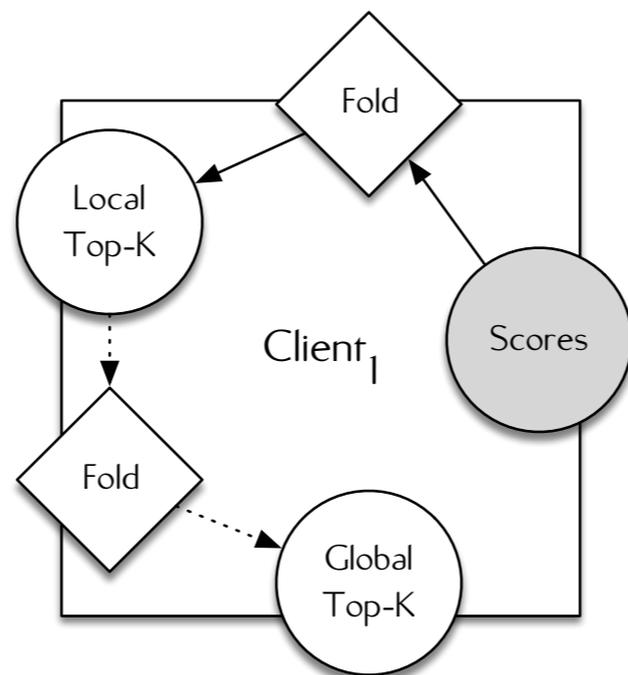
Per-User Leaderboard

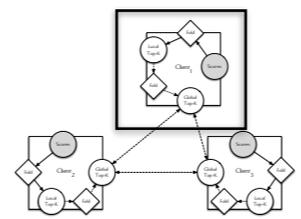
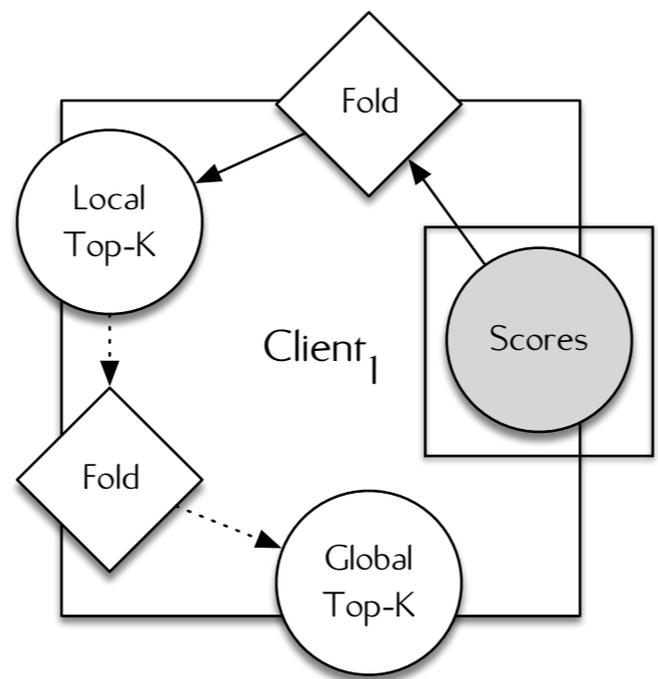
- **Enhance existing design**
Only the top score for each user at each device
- **Minimize transmitted state**
Prevent transmission of state that is not necessary to perform the computation

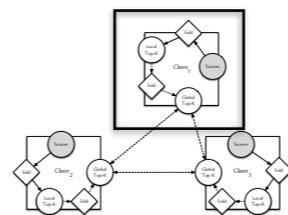
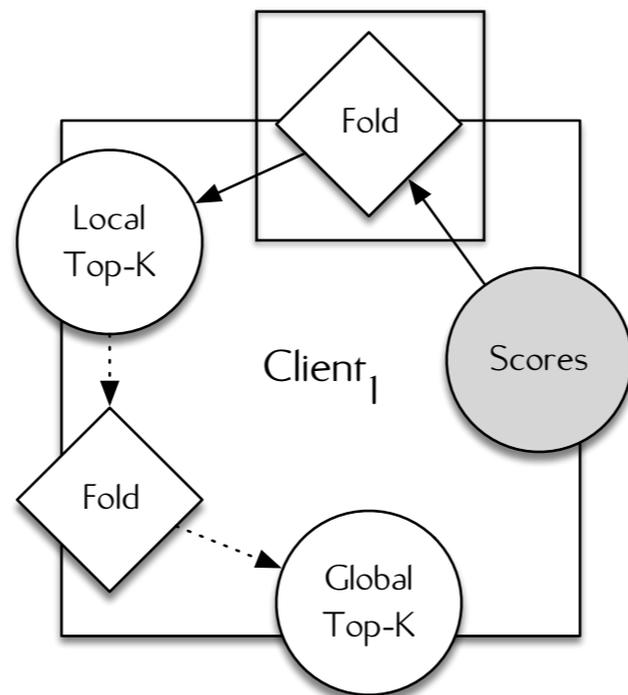
Per-User Leaderboard

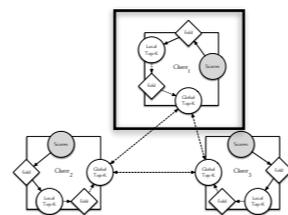
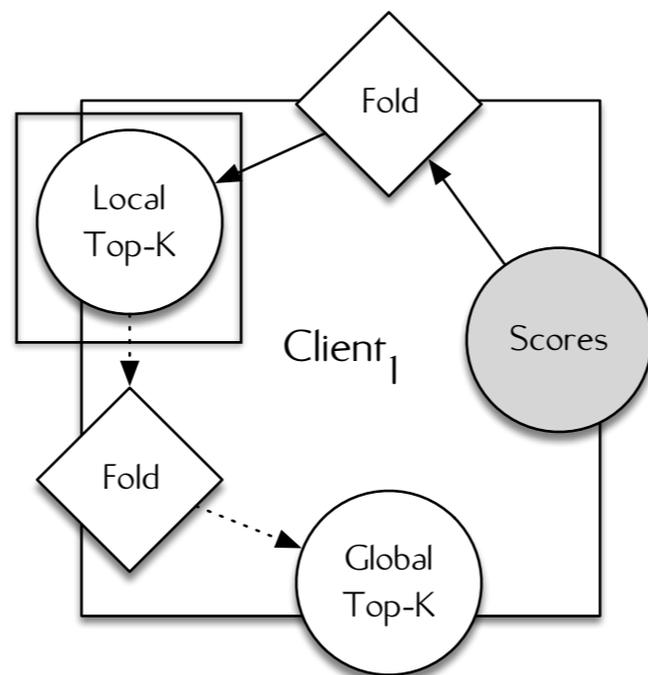
- **Enhance existing design**
Only the top score for each user at each device
- **Minimize transmitted state**
Prevent transmission of state that is not necessary to perform the computation
- **Compose data types**
Build a per-user leaderboard through the composition of existing types

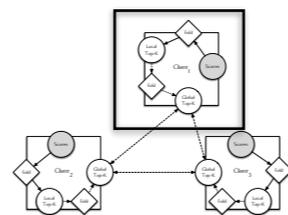
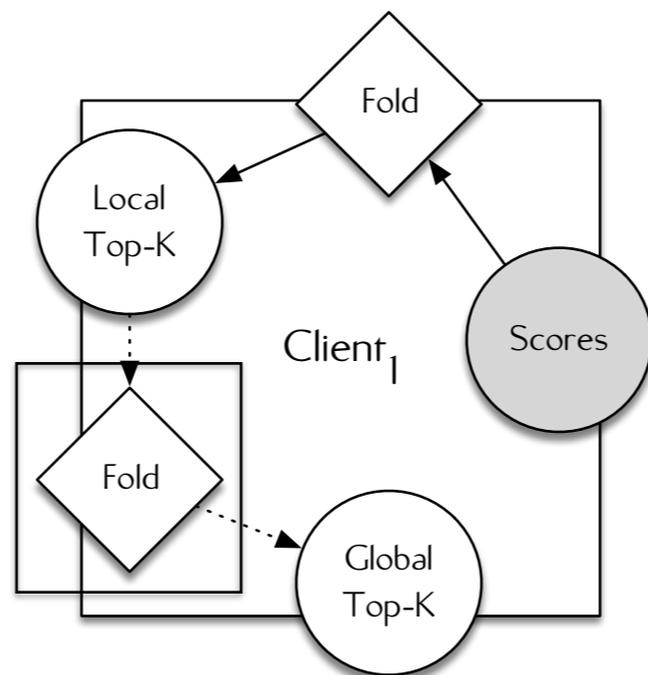


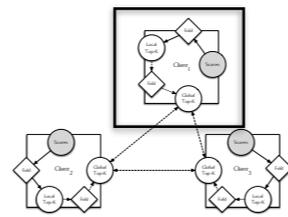
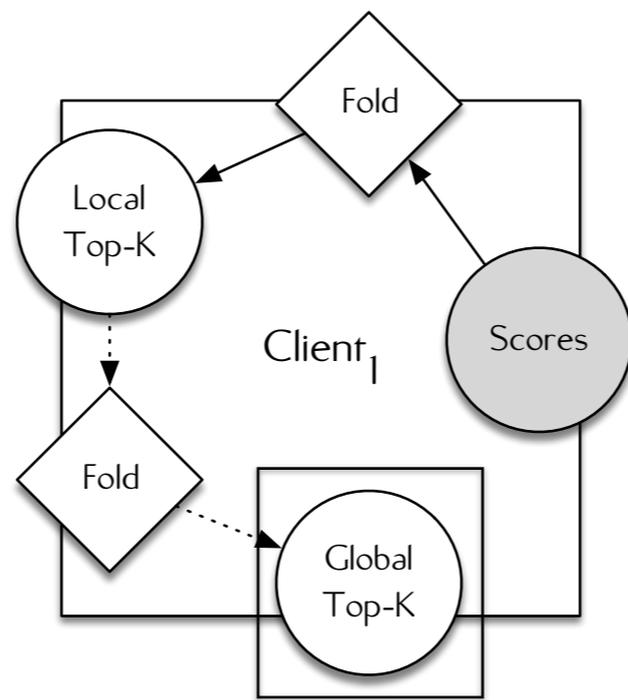


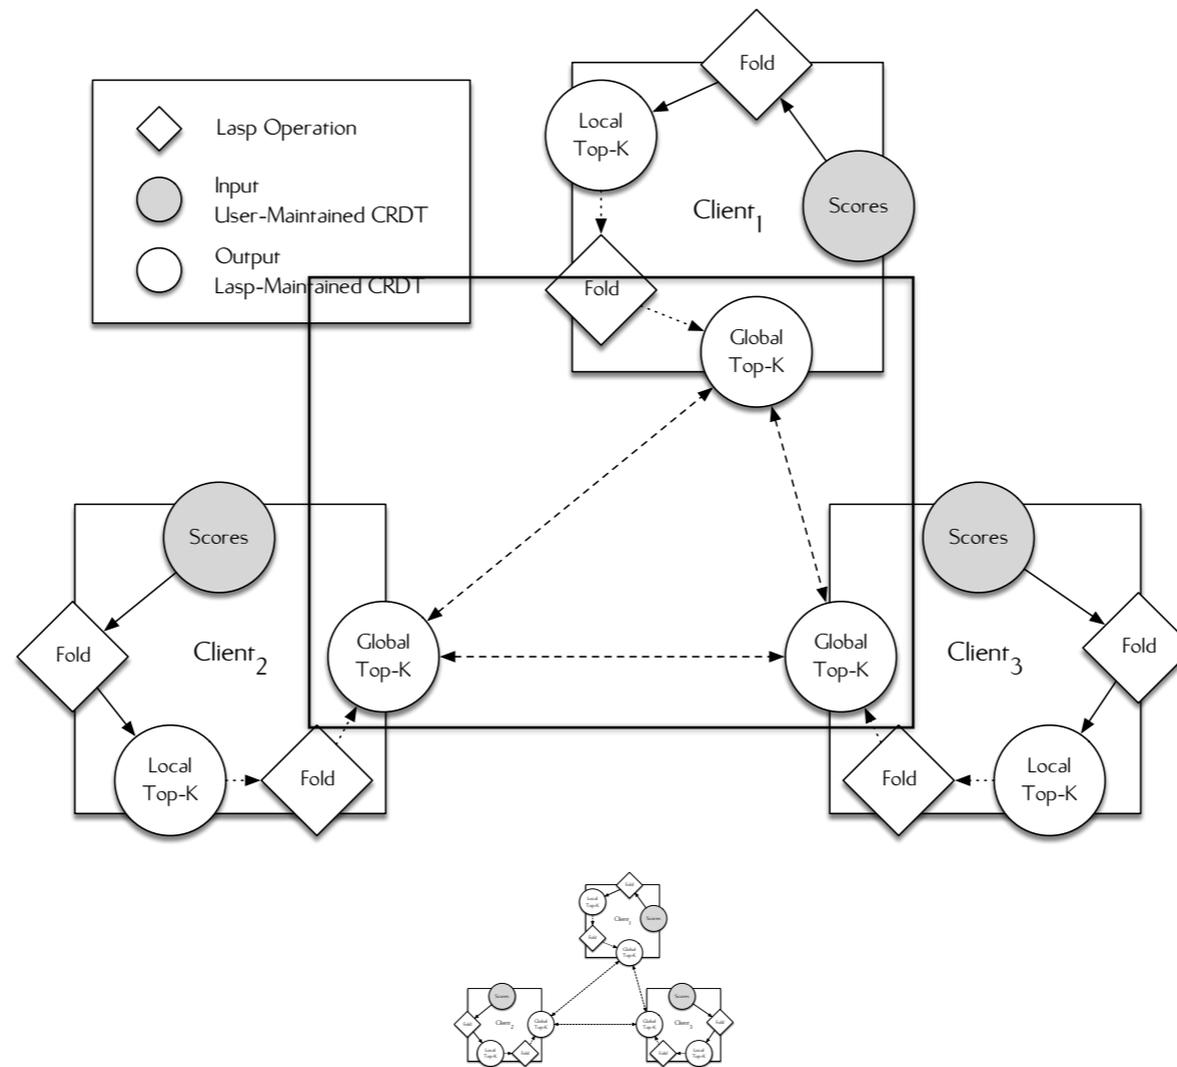












Per-User Leaderboard

- **Dynamically scoped variables**
Variable which take different values depending on where it is executing

Per-User Leaderboard

- **Dynamically scoped variables**
Variable which take different values depending on where it is executing
- **Dynamically scoped fold operation**
Perform a distributed “reduce” operation that combines the state of a dynamically scoped variables across

```
%% Create a global leaderboard.
G = declare({top_k, [10]}).

%% Create a local leaderboard.
L = declare_dynamic({top_k, [10]}).

%% Create a set of scores.
S = declare_dynamic(set).

%% Compute local top-k list.
fold(S, fun max_by_name/2, L).

%% Compute global top-k list.
fold_dynamic(L, fun max_by_name/2, G).
```

```
%% Create a global leaderboard.  
G = declare({top_k, [10]}).  
  
%% Create a local leaderboard.  
L = declare_dynamic({top_k, [10]}).  
  
%% Create a set of scores.  
S = declare_dynamic(set).  
  
%% Compute local top-k list.  
fold(S, fun max_by_name/2, L).  
  
%% Compute global top-k list.  
fold_dynamic(L, fun max_by_name/2, G).
```

```
%% Create a global leaderboard.
G = declare({top_k, [10]}).

%% Create a local leaderboard.
L = declare_dynamic({top_k, [10]}).

%% Create a set of scores.
S = declare_dynamic(set).

%% Compute local top-k list.
fold(S, fun max_by_name/2, L).

%% Compute global top-k list.
fold_dynamic(L, fun max_by_name/2, G).
```

```
%% Create a global leaderboard.
G = declare({top_k, [10]}).

%% Create a local leaderboard.
L = declare_dynamic({top_k, [10]}).

%% Create a set of scores.
S = declare_dynamic(set).

%% Compute local top-k list.
fold(S, fun max_by_name/2, L).

%% Compute global top-k list.
fold_dynamic(L, fun max_by_name/2, G).
```

```
%% Create a global leaderboard.
G = declare({top_k, [10]}).

%% Create a local leaderboard.
L = declare_dynamic({top_k, [10]}).

%% Create a set of scores.
S = declare_dynamic(set).

%% Compute local top-k list.
fold(S, fun max_by_name/2, L).

%% Compute global top-k list.
fold_dynamic(L, fun max_by_name/2, G).
```

```
%% Create a global leaderboard.
G = declare({top_k, [10]}).

%% Create a local leaderboard.
L = declare_dynamic({top_k, [10]}).

%% Create a set of scores.
S = declare_dynamic(set).

%% Compute local top-k list.
fold(S, fun max_by_name/2, L).

%% Compute global top-k list.
fold_dynamic(L, fun max_by_name/2, G).
```

Let's look at a larger
application

Let's look at a larger
application

**With visible
non-monotonicity**

What can we build?
Advertisement Counter

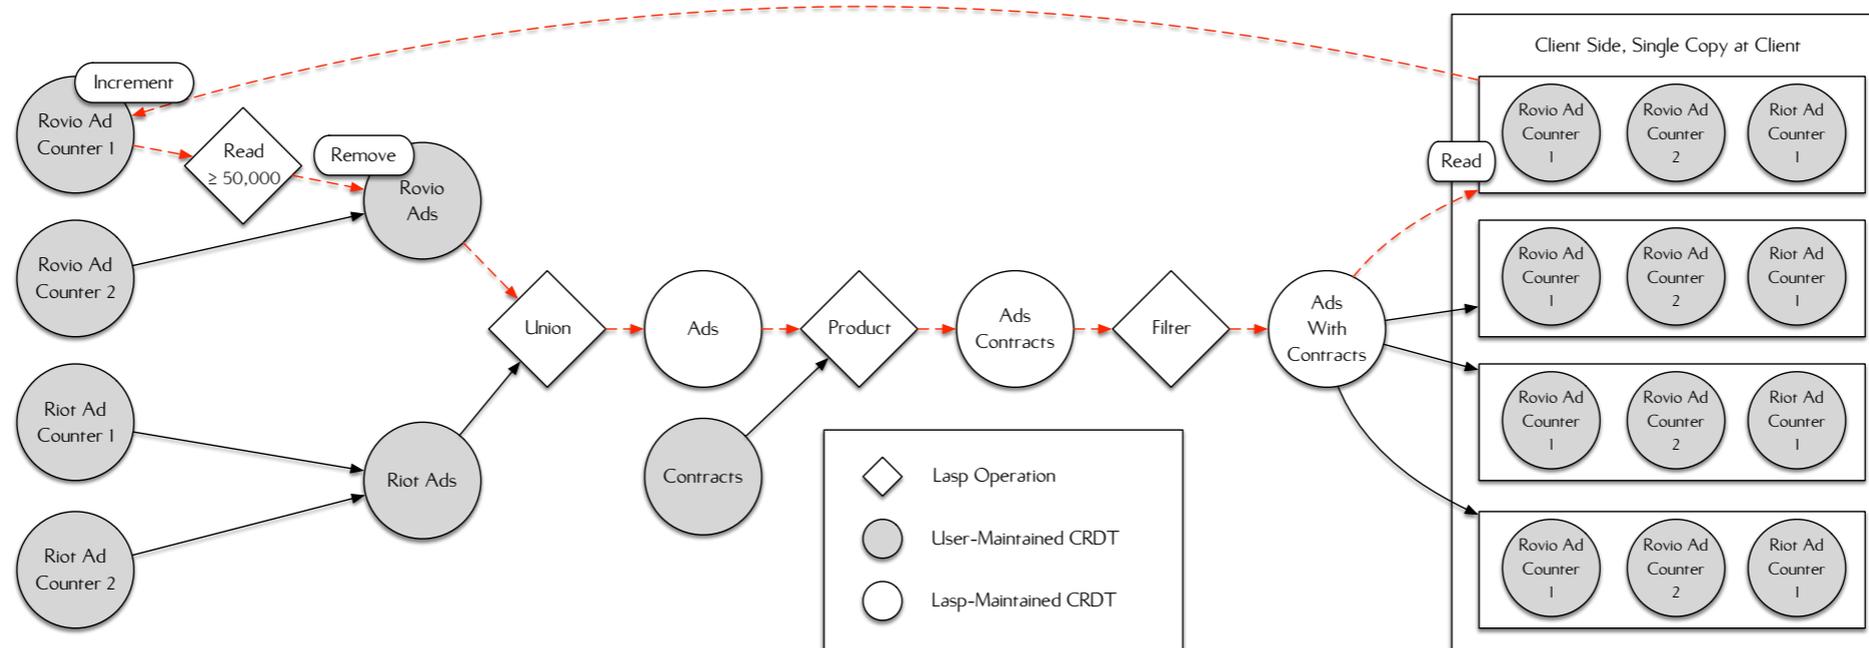
Advertisement Counter

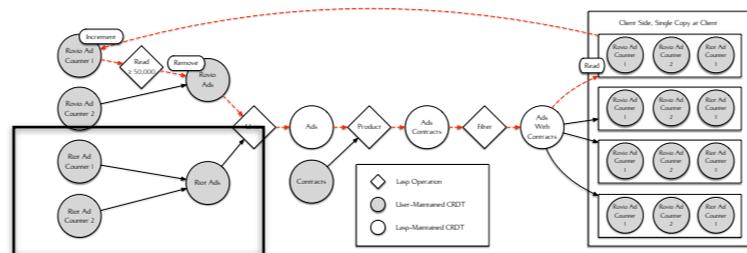
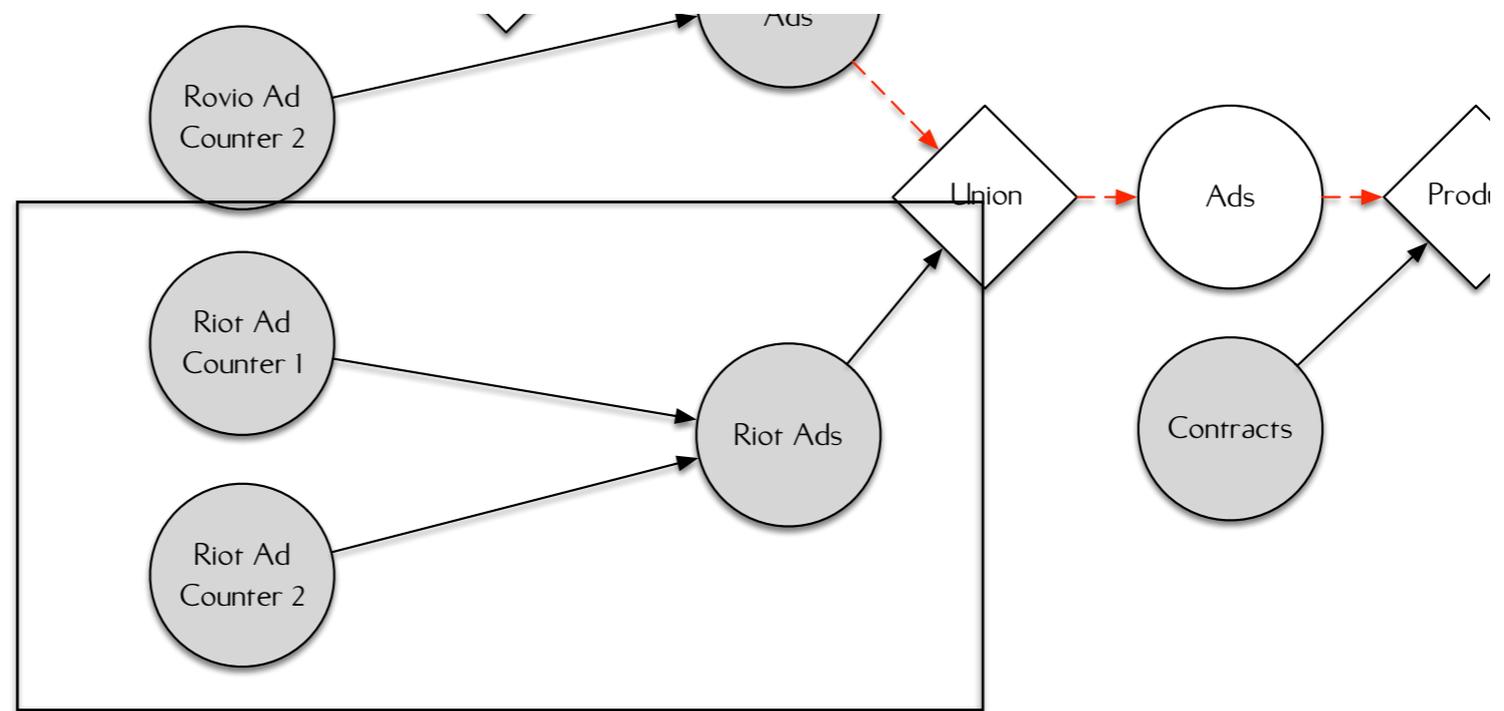
- Mobile game platform selling advertisement space

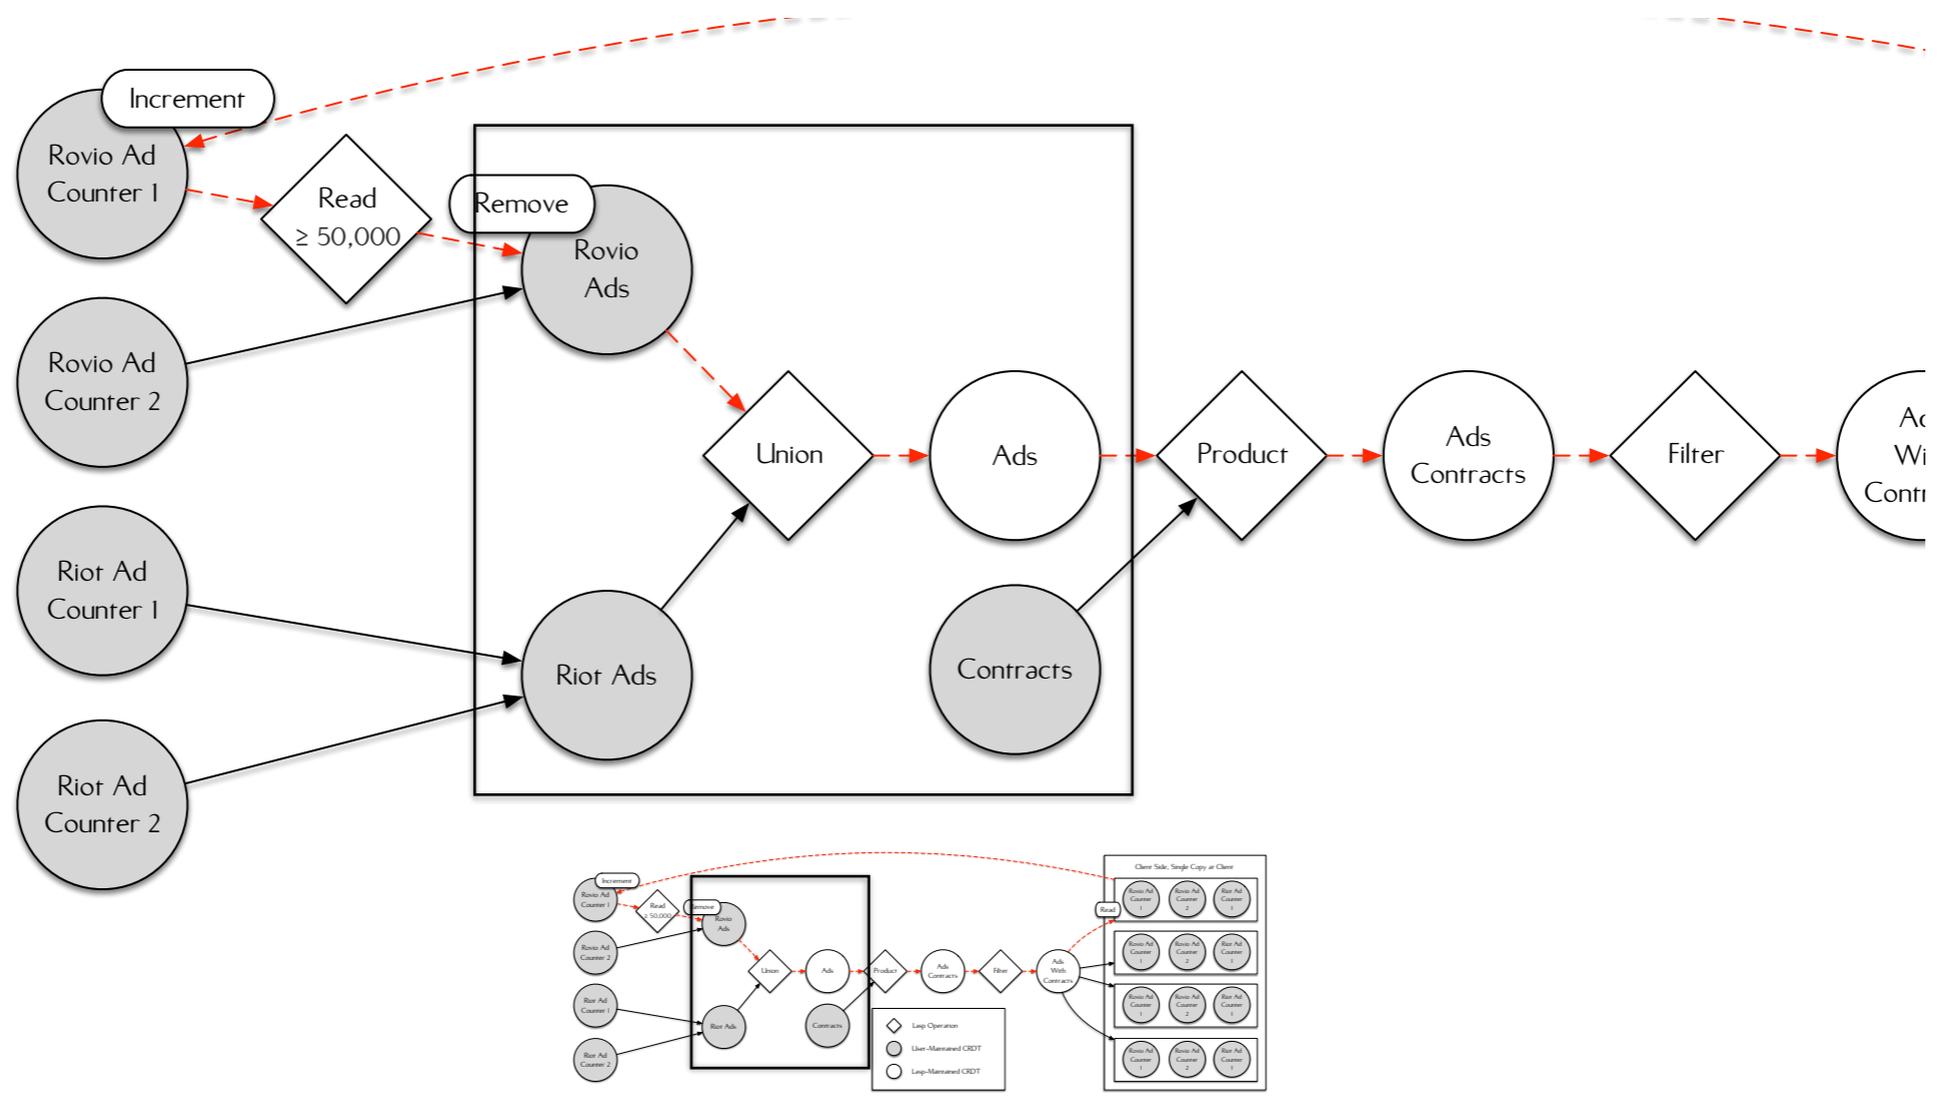
Advertisements are paid according to a minimum number of impressions

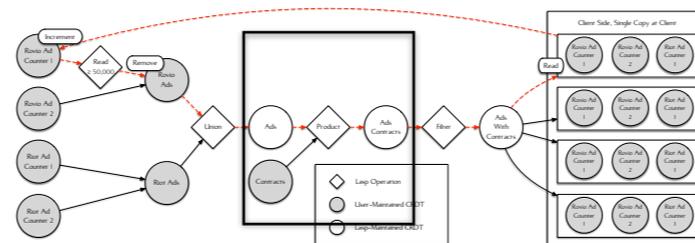
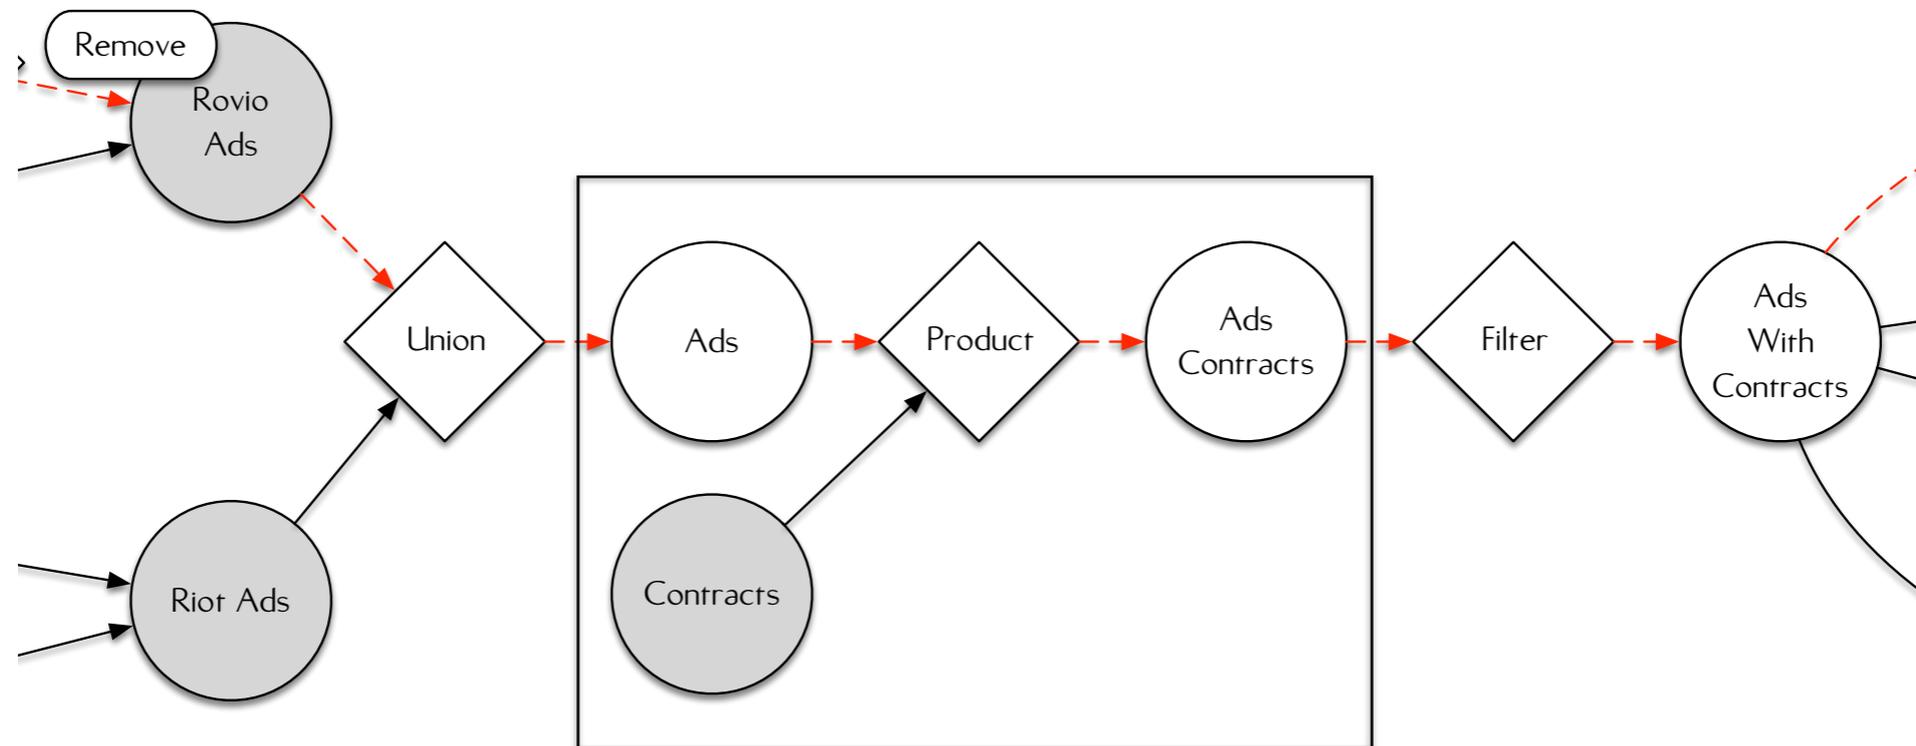
Advertisement Counter

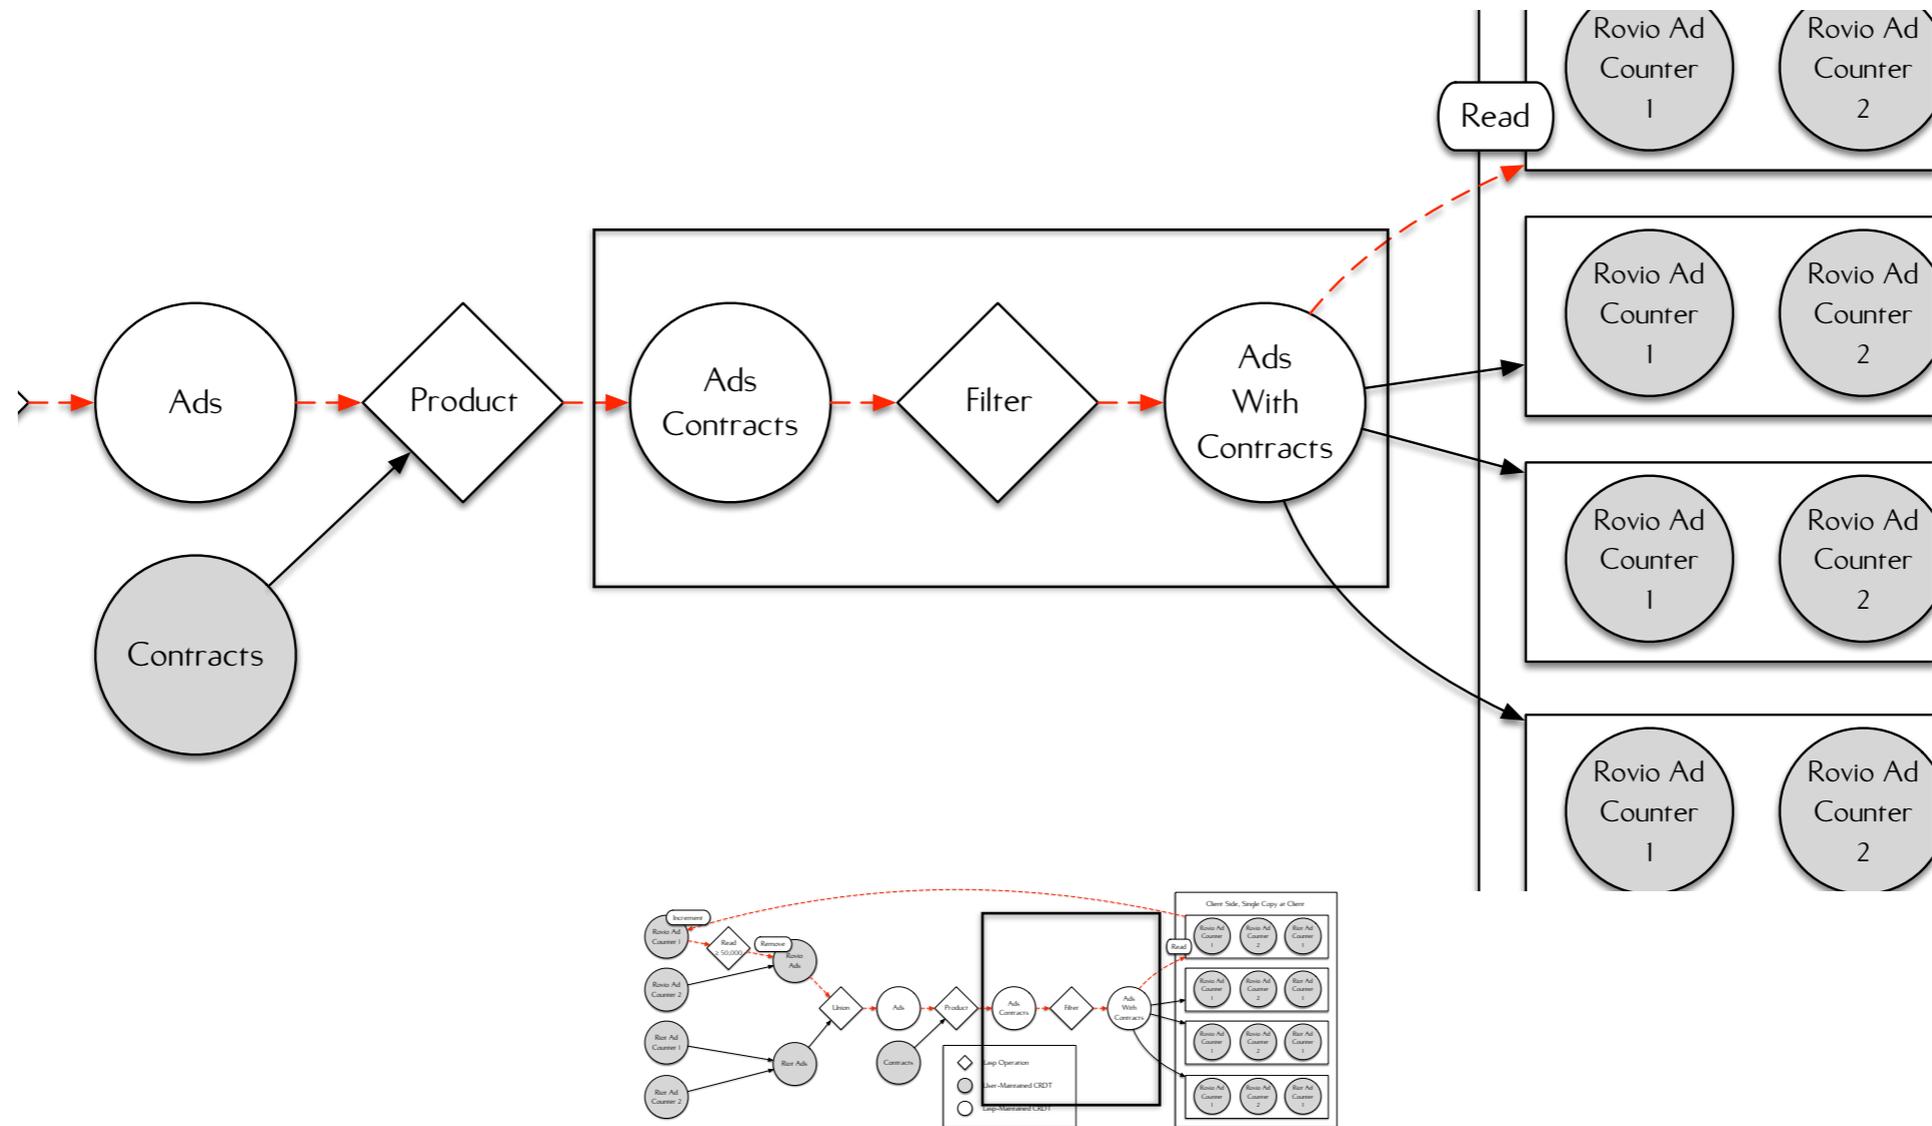
- **Mobile game platform selling advertisement space**
Advertisements are paid according to a minimum number of impressions
- **Clients will go offline**
Clients have limited connectivity and the system still needs to make progress while clients are offline

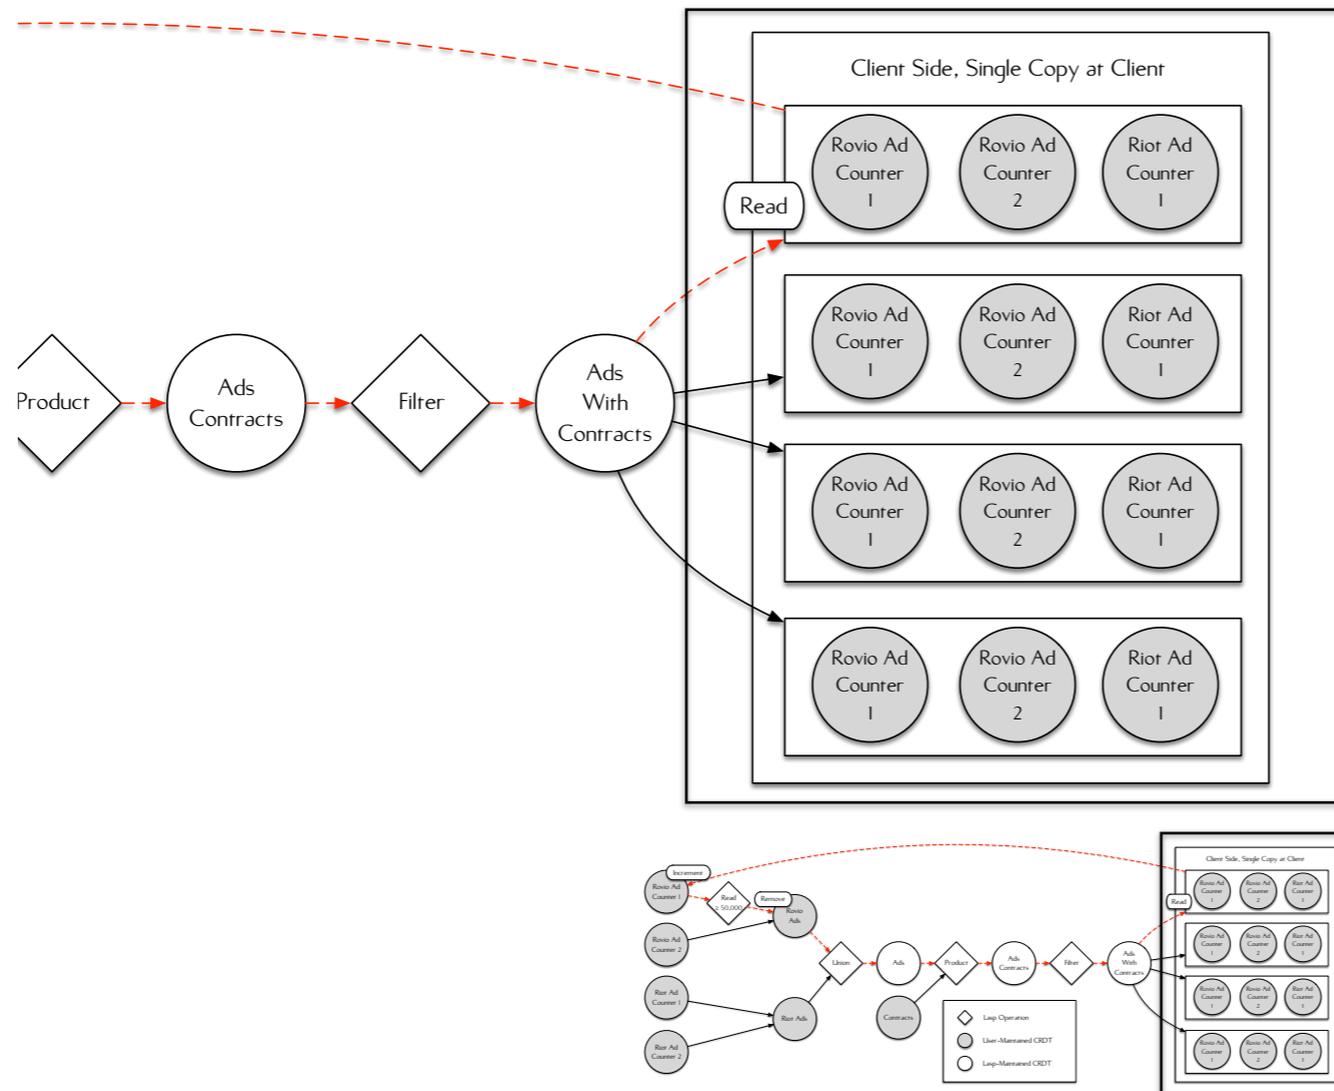


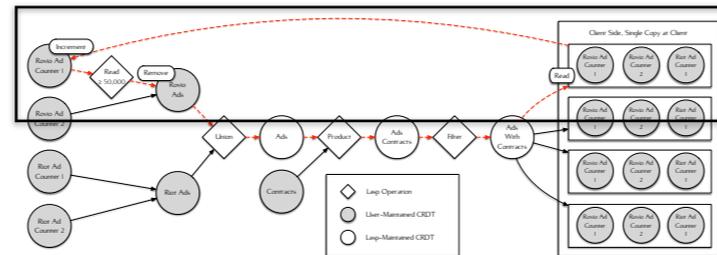
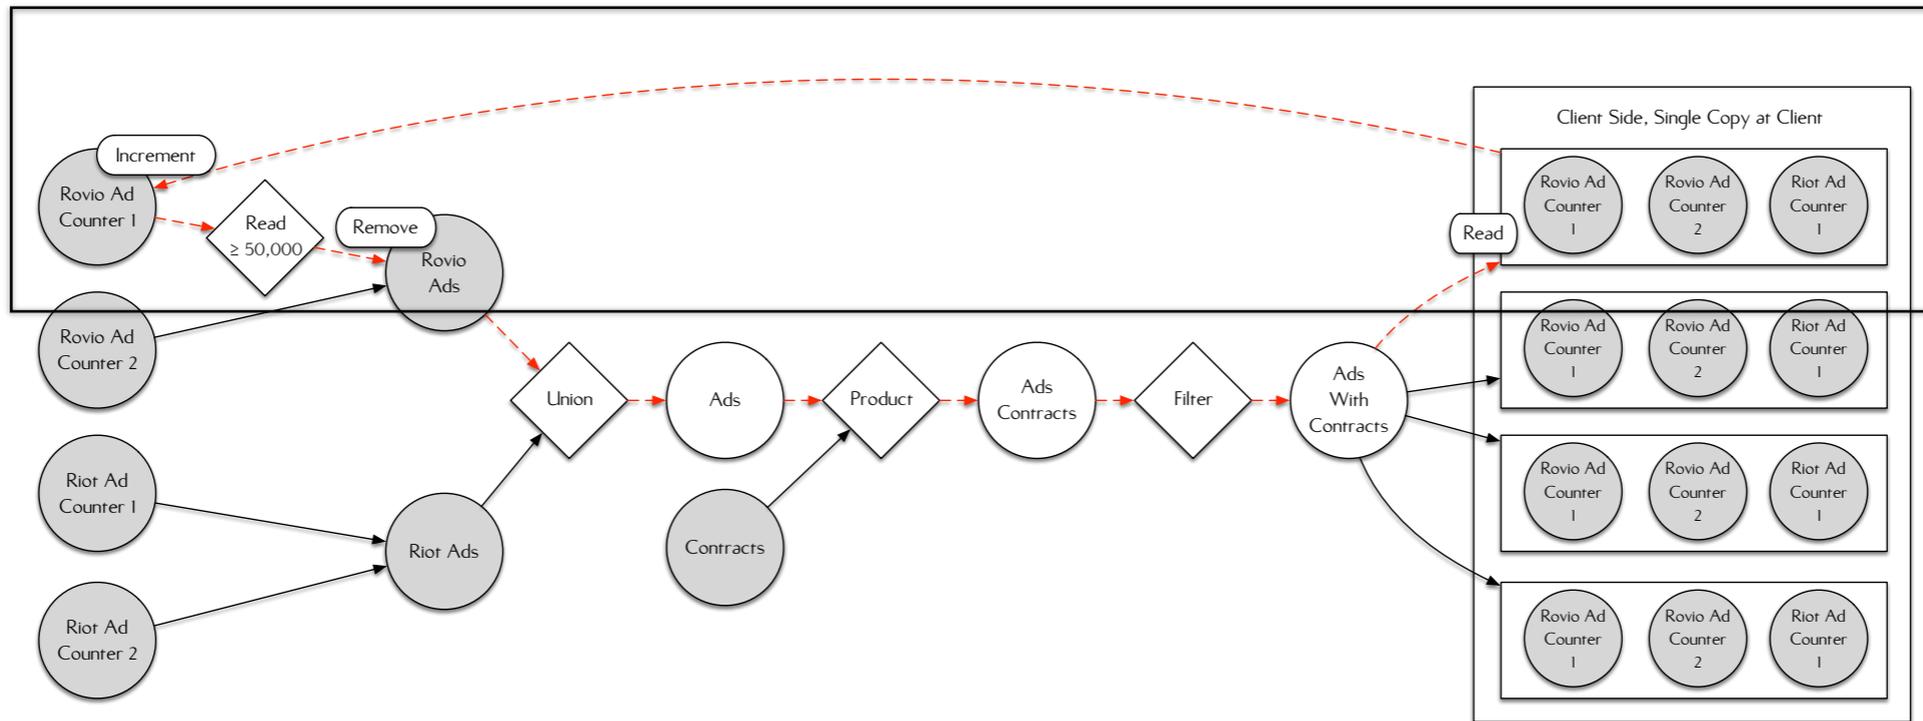


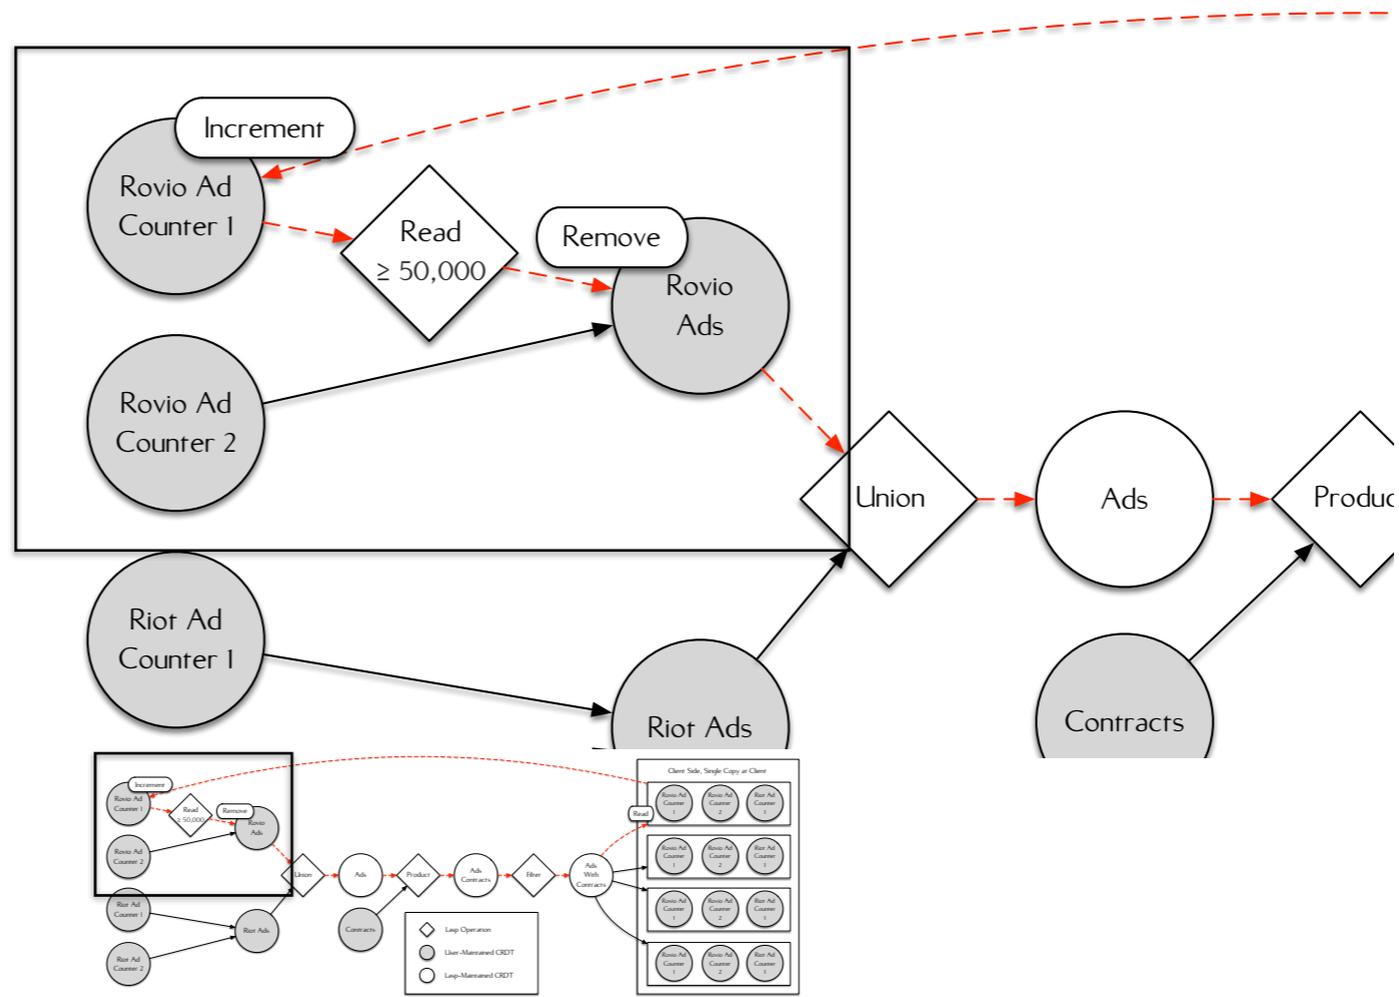


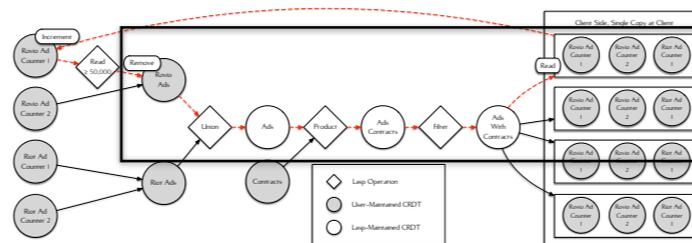
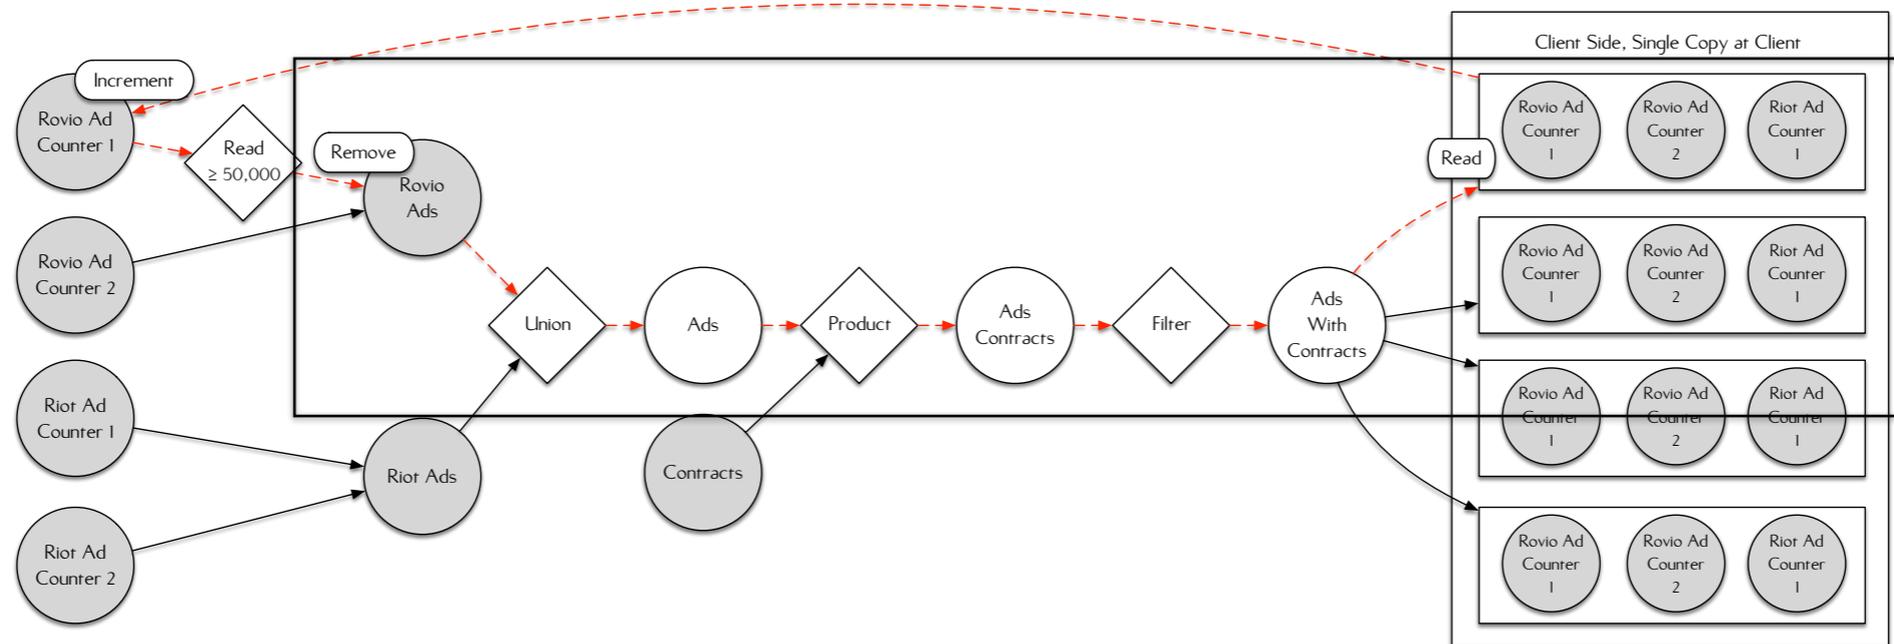












Advertisement Counter

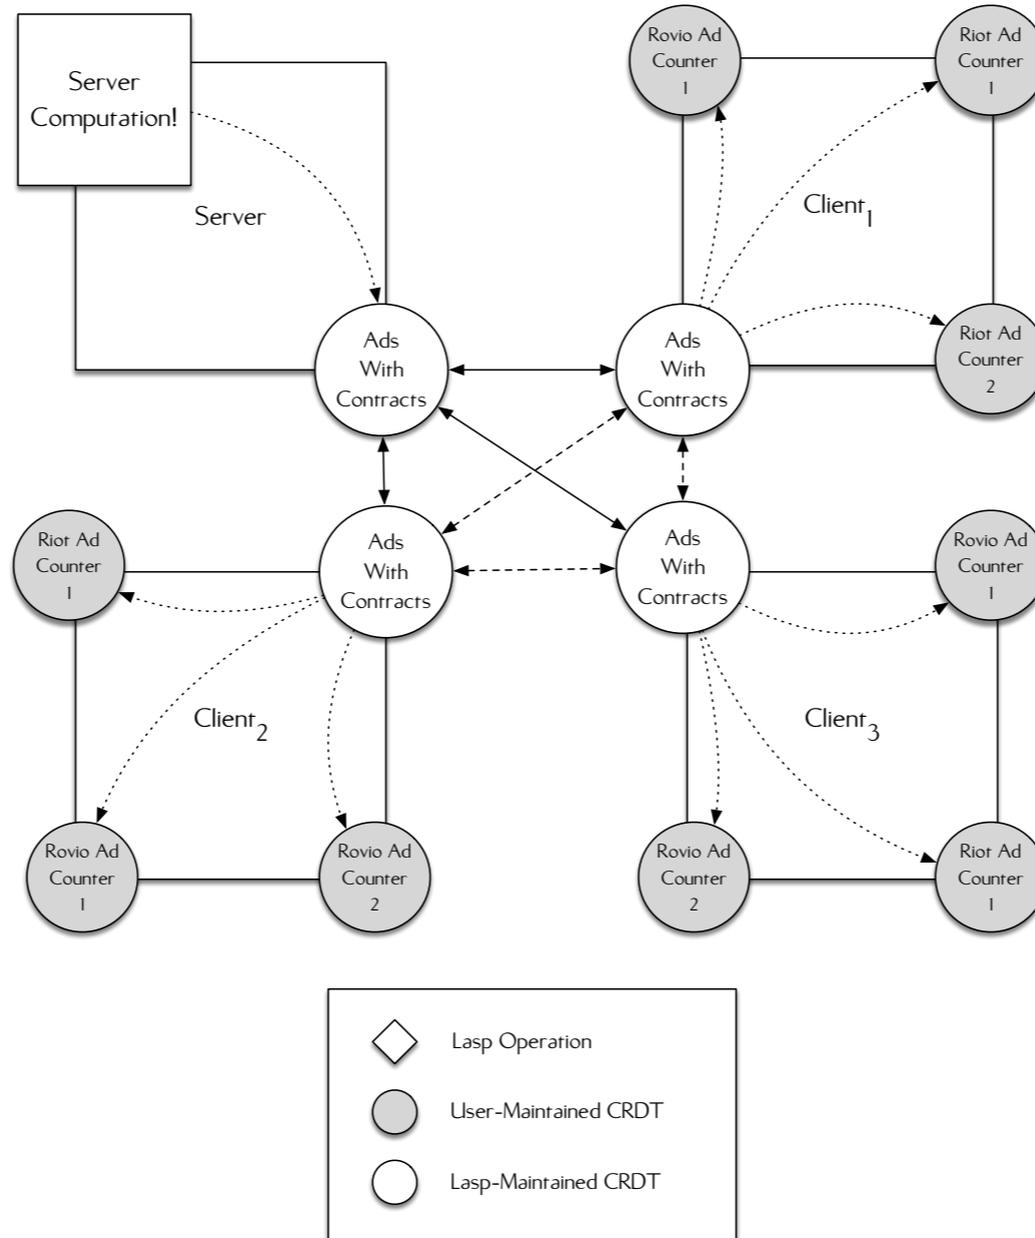
- **Completely monotonic**
Disabling advertisements and contracts are all modeled through monotonic state growth

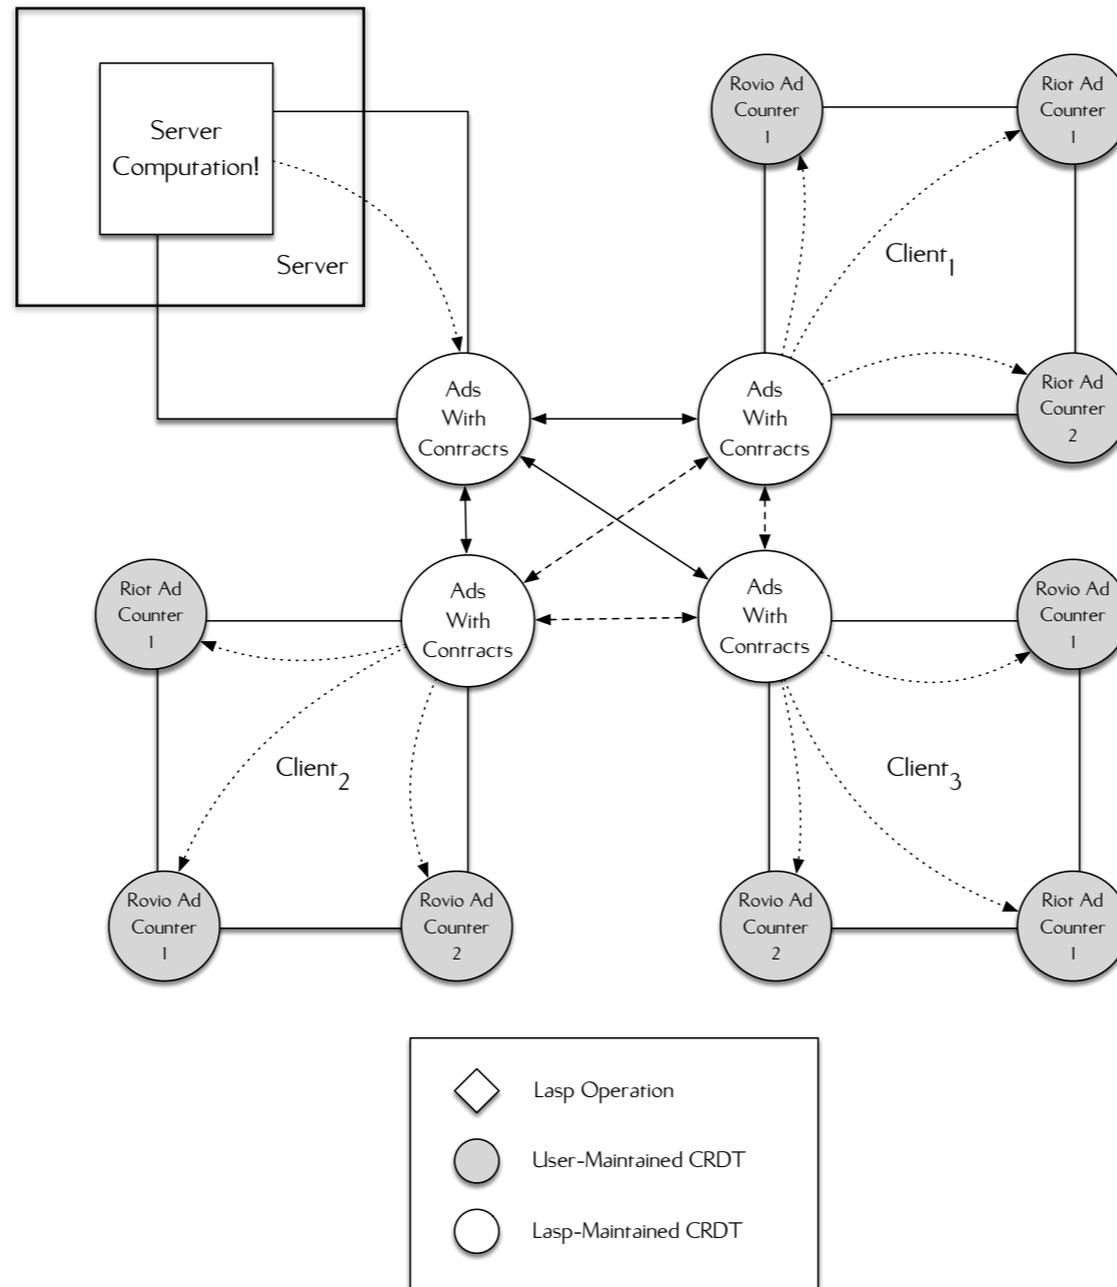
Advertisement Counter

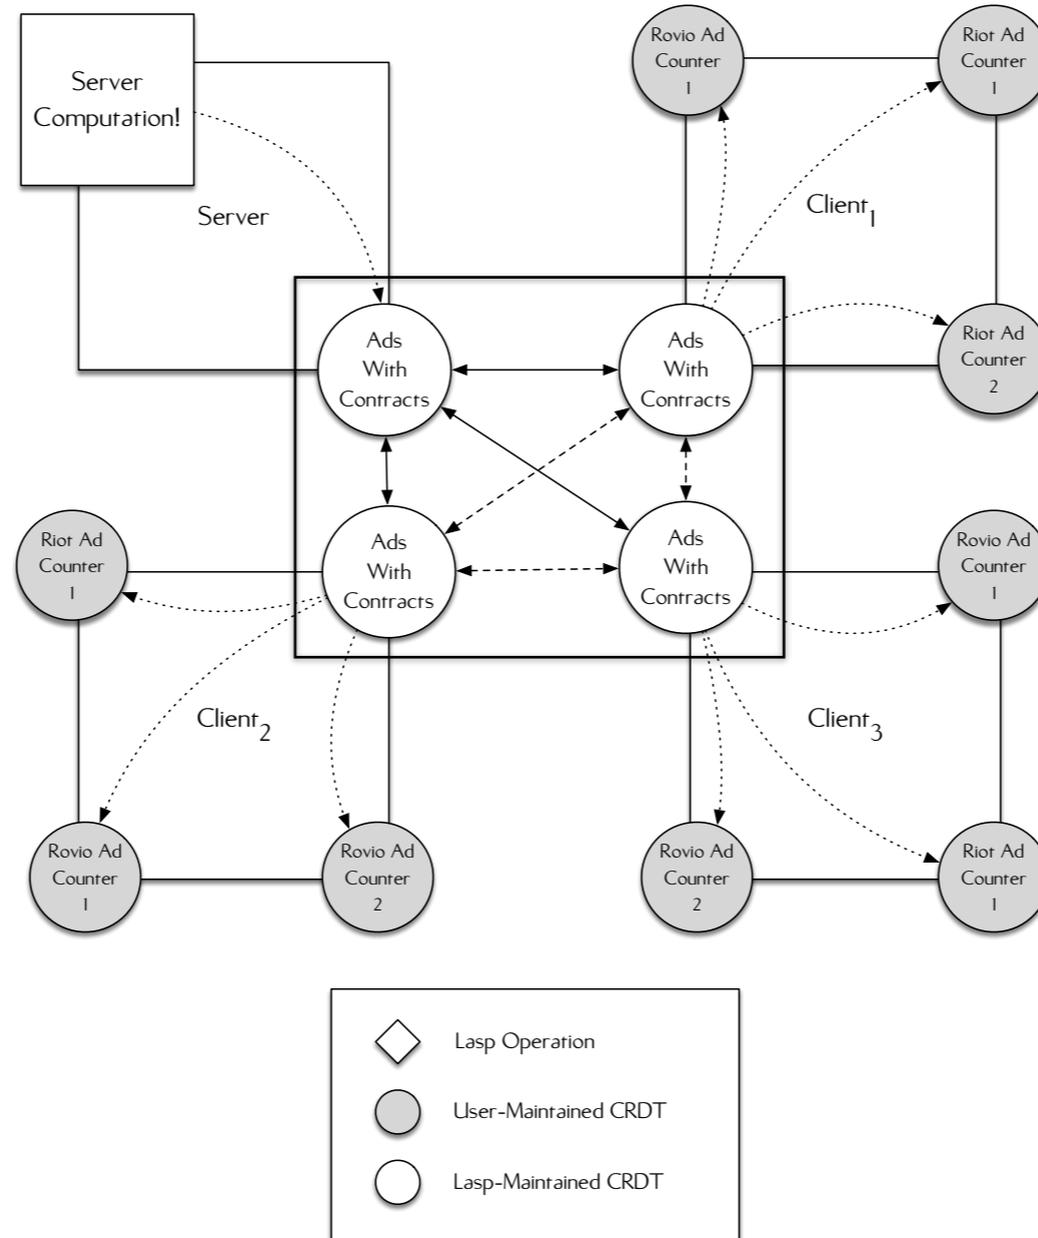
- **Completely monotonic**
Disabling advertisements and contracts are all modeled through monotonic state growth
- **Arbitrary distribution**
Use of convergent data structures allows computational graph to be arbitrarily distributed

Advertisement Counter

- **Completely monotonic**
Disabling advertisements and contracts are all modeled through monotonic state growth
- **Arbitrary distribution**
Use of convergent data structures allows computational graph to be arbitrarily distributed
- **Divergence**
Divergence is a factor of synchronization period







Advertisement Counter

- “Servers” as peers to “clients”
Servers are peers to clients that perform additional computation

Advertisement Counter

- “Servers” as peers to “clients”
Servers are peers to clients that perform additional computation
 - Any node can disable an advertisement under this model given enough information

Advertisement Counter

- “Servers” as peers to “clients”
Servers are peers to clients that perform additional computation
 - Any node can disable an advertisement under this model given enough information
- “Servers” as trusted nodes
Serve as a location for performing “exactly once” side-effects

Advertisement Counter

- “Servers” as peers to “clients”
Servers are peers to clients that perform additional computation
 - Any node can disable an advertisement under this model given enough information
- “Servers” as trusted nodes
Serve as a location for performing “exactly once” side-effects
 - Billing customers must be done at a central point by a trusted node in the system

We've build up from zero
synchronization

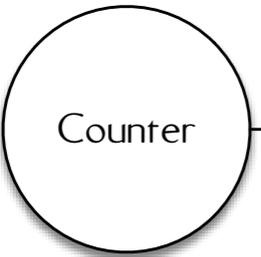
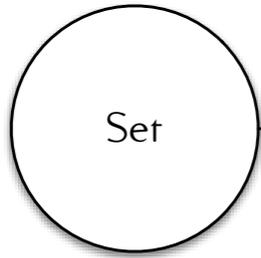
We've build up from zero
synchronization

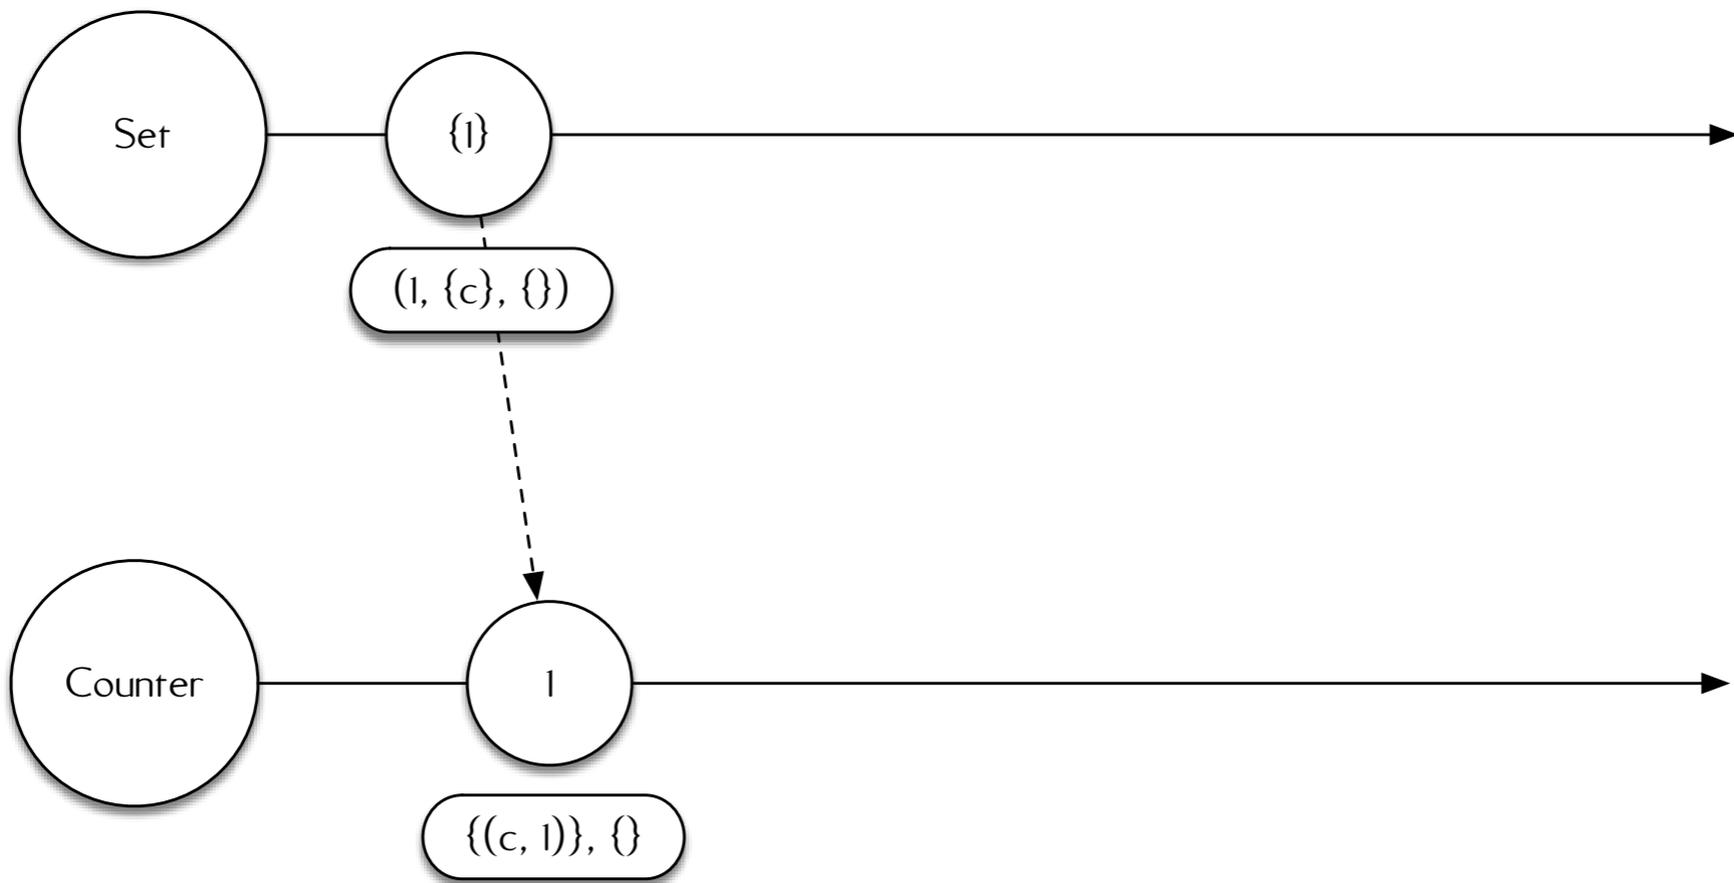
**Instead of working to
remove synchronization**

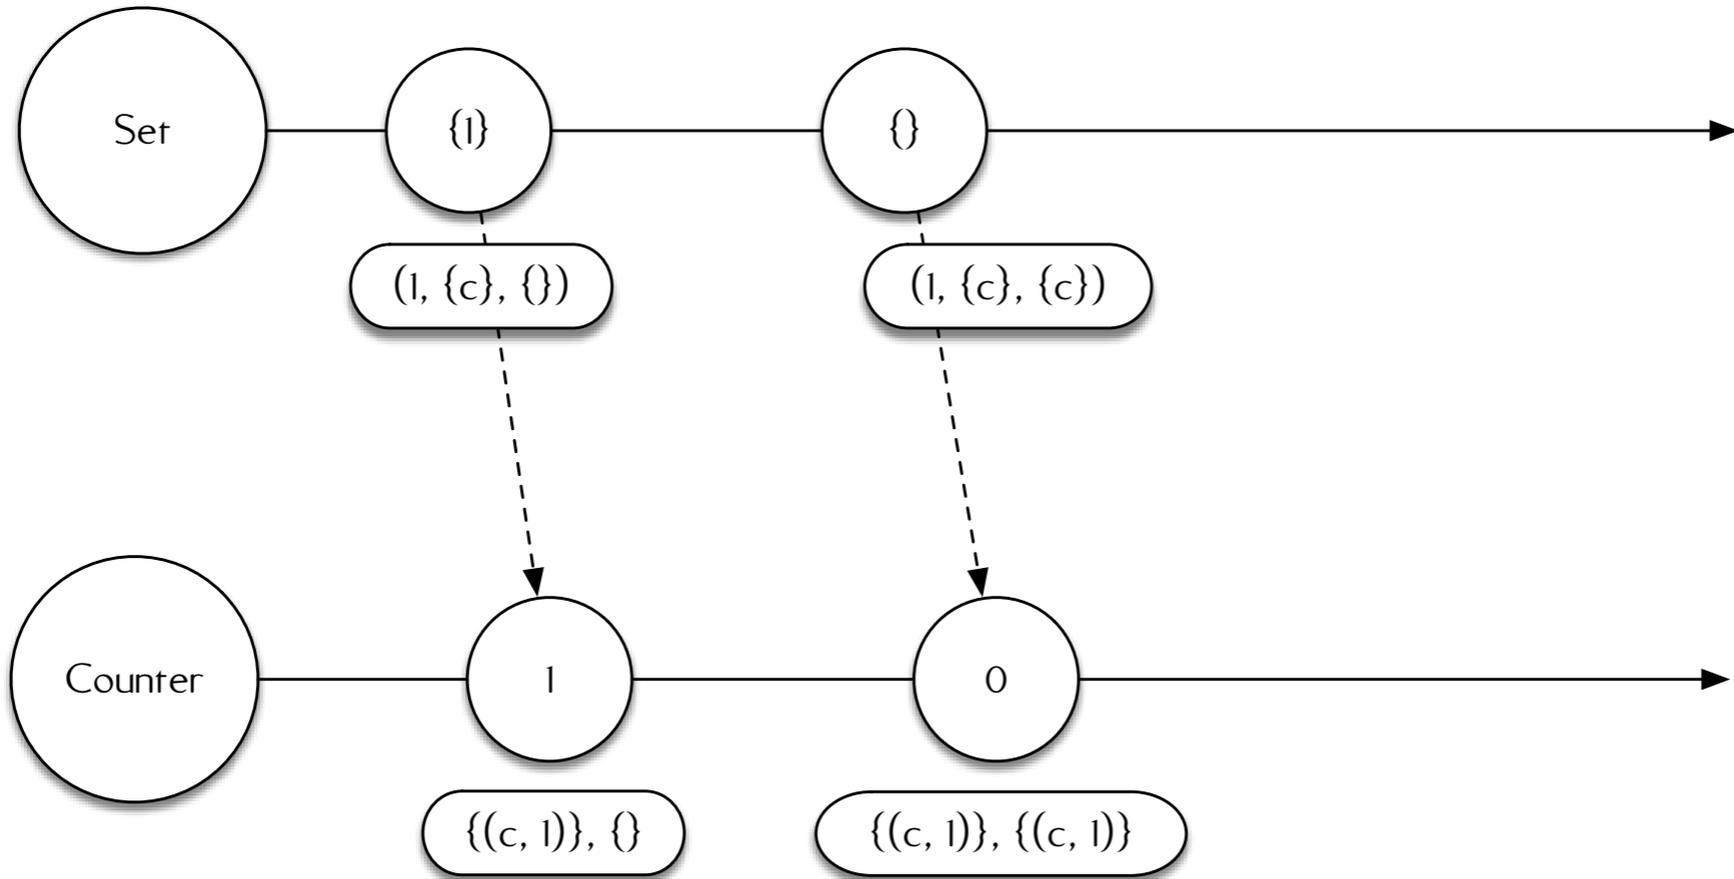
Challenges Looking Ahead

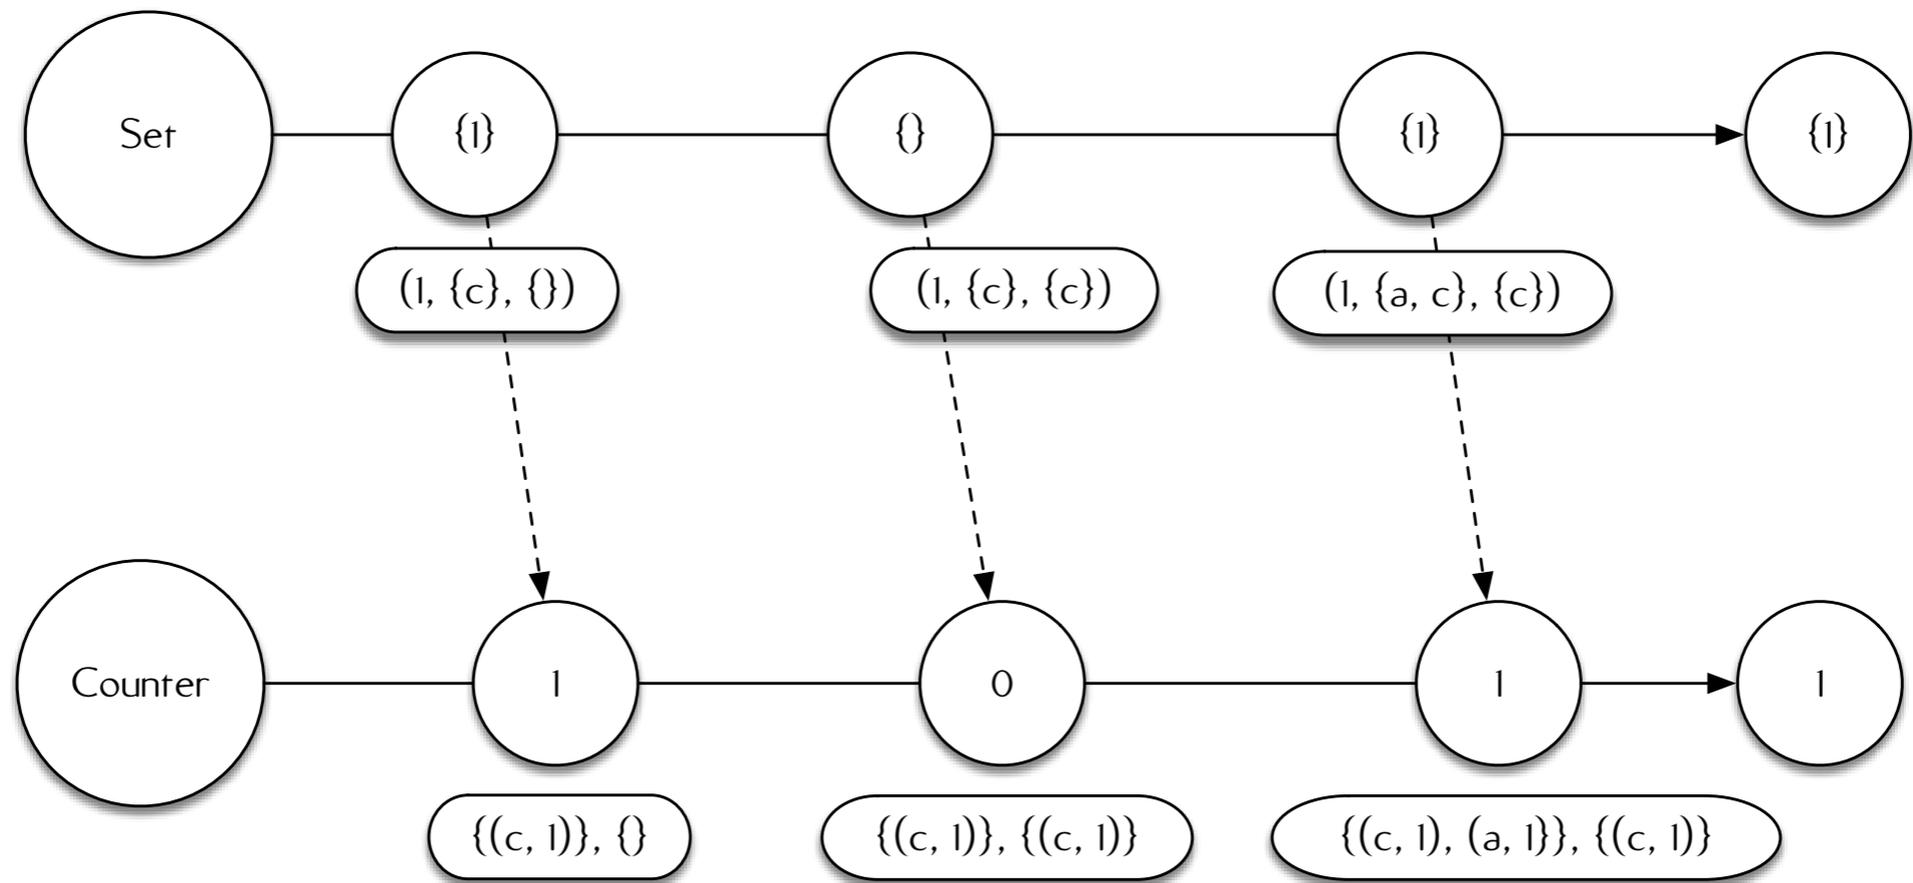
Causality

State Explosion

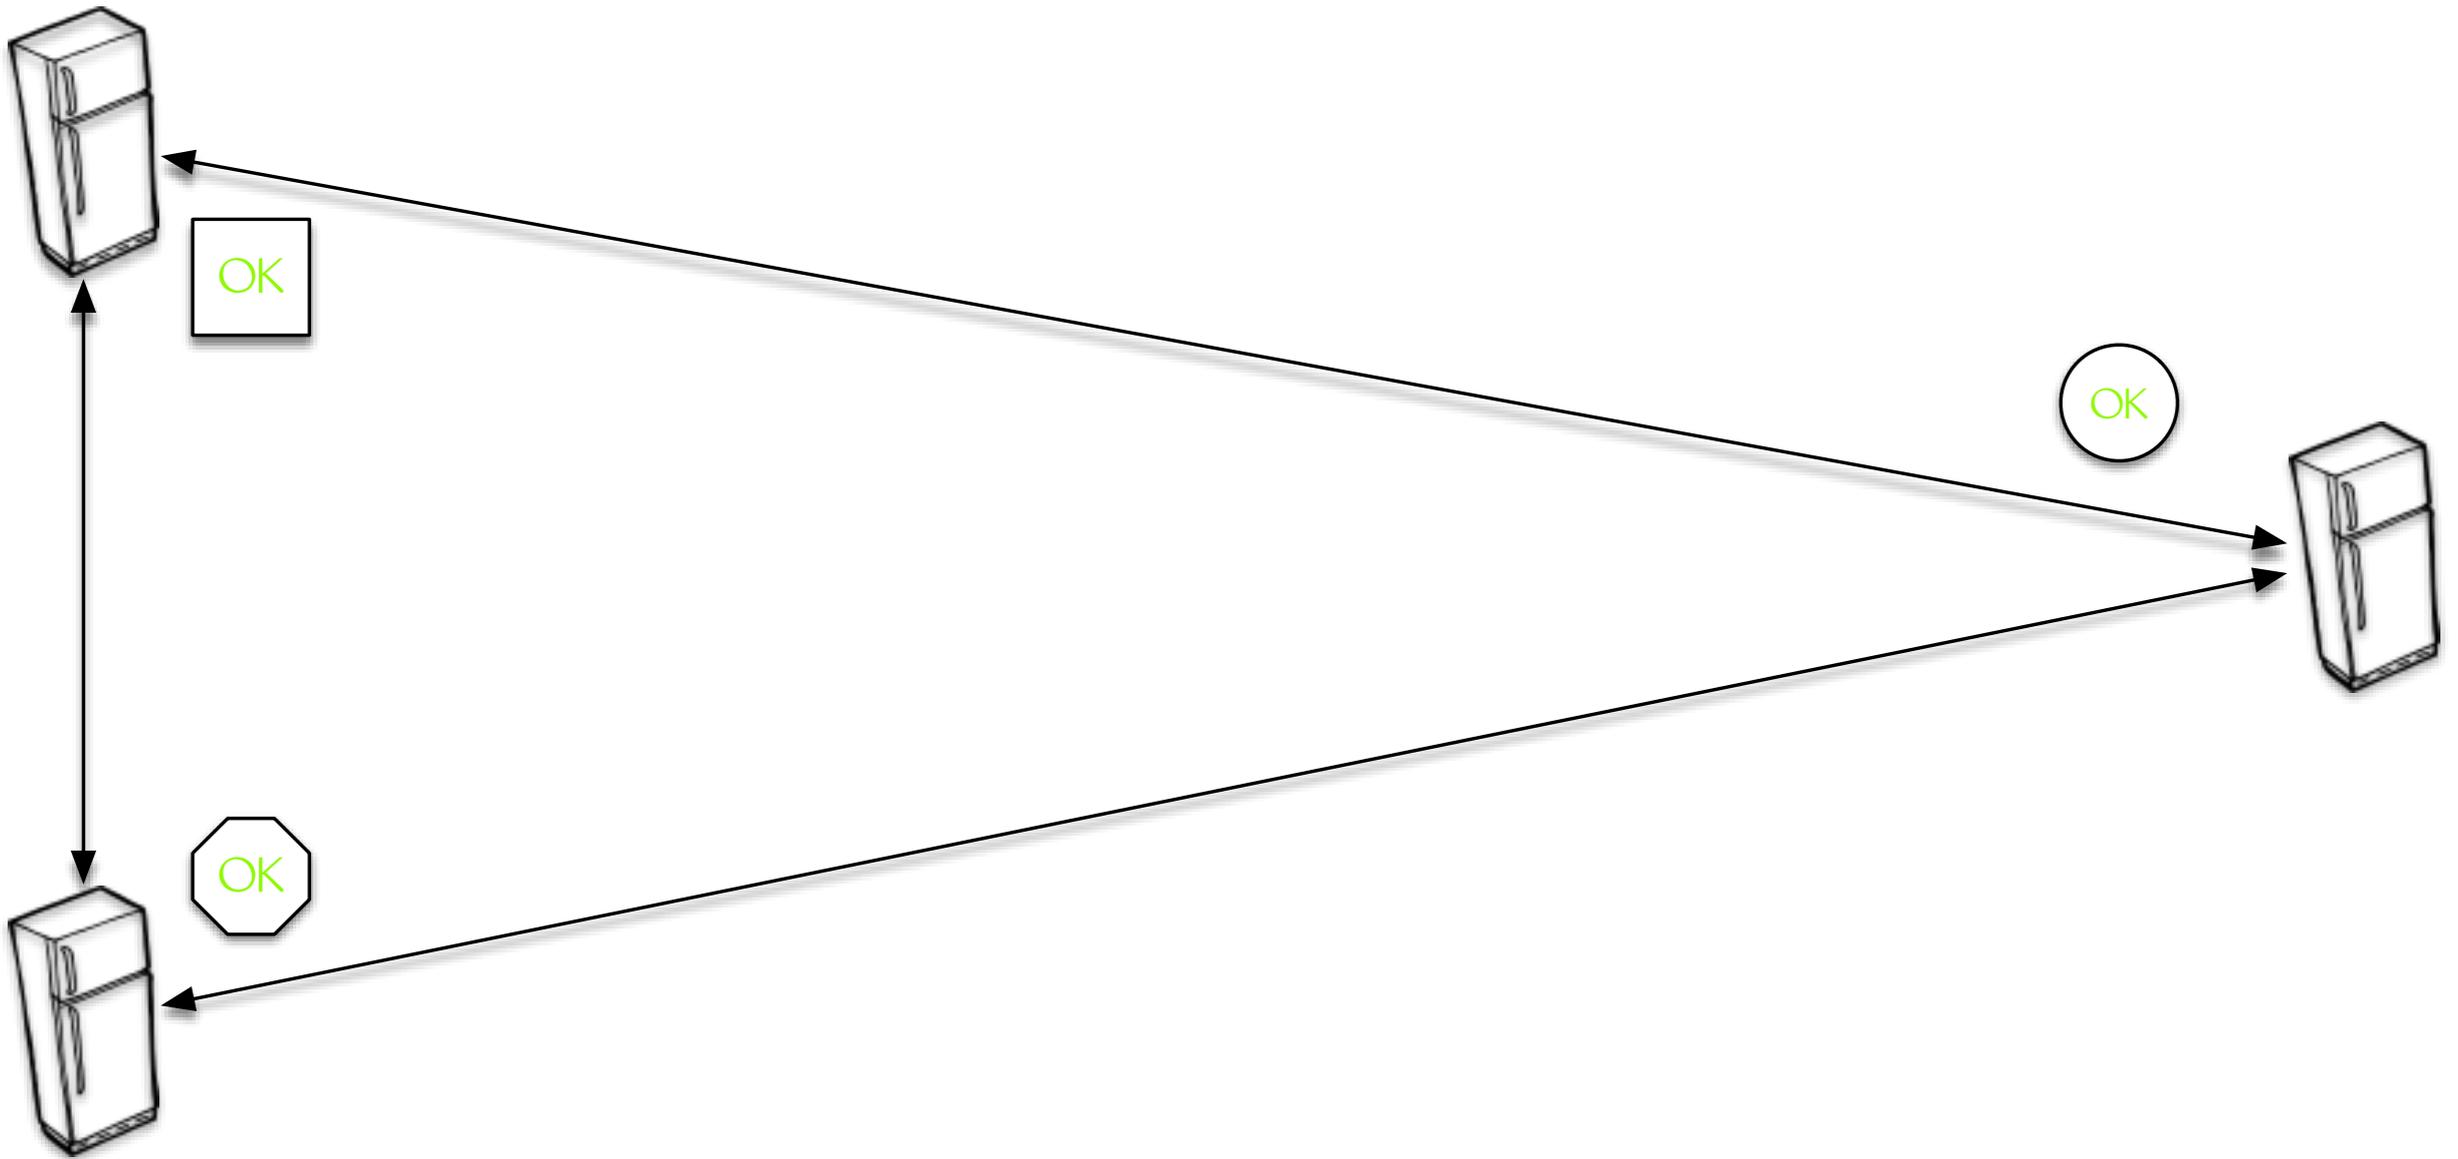








Security Computing at the Edge



Computations

Expressiveness

How restrictive is a programming model where operations must be **associative, commutative, and idempotent?**

What's new?

Since Erlang Factory 2015

- Erlang 18, rebar3, Common Test
Faster test suite, on Erlang 18 with rebar3

Since Erlang Factory 2015

- Erlang 18, rebar3, Common Test
Faster test suite, on Erlang 18 with rebar3
- No Riak Core
Plumtree distribution replaces the Riak Core distribution system

Since Erlang Factory 2015

- **Erlang 18, rebar3, Common Test**
Faster test suite, on Erlang 18 with rebar3
- **No Riak Core**
Plumtree distribution replaces the Riak Core distribution system
- **No NIFs**
Allows us to cross-compile to other platforms and not worry about NIF scheduling

Semantics Improvements

- Delta-State Based Conflict-Free Replicated Data Types
Optimized state dissemination by shipping deltas
[Almeida *et al.*, 2016]

Semantics Improvements

- **Delta-State Based Conflict-Free Replicated Data Types**
Optimized state dissemination by shipping deltas
[Almeida *et al.*, 2016]
- **Causal CRDTs**
Optimized, **garbage-free** data structure support
[Almeida *et al.*, 2016; Meiklejohn 2016, in review]

Semantics Improvements

- **Delta-State Based Conflict-Free Replicated Data Types**
Optimized state dissemination by shipping deltas
[Almeida *et al.*, 2016]
- **Causal CRDTs**
Optimized, **garbage-free** data structure support
[Almeida *et al.*, 2016; Meiklejohn 2016, in review]
- **Fold**
New semantics for a more expressive **fold** operation for
arbitrary computation over sets
[Meiklejohn 2016, in review]

Runtime Improvements

- Delta-State Based Anti-Entropy
Optimized AAE mechanism based on deltas
[Almeida *et al.*, 2016]

Runtime Improvements

- **Delta-State Based Anti-Entropy**
Optimized AAE mechanism based on deltas
[Almeida *et al.*, 2016]
- **Mesos and Docker Enabled**
Run large-scale **Lasp** clusters on **Mesos** with Marathon
[Meiklejohn and Yoo 2016, in review]

Runtime Improvements

- **Delta-State Based Anti-Entropy**
Optimized AAE mechanism based on deltas
[Almeida *et al.*, 2016]
- **Mesos and Docker Enabled**
Run large-scale **Lasp** clusters on **Mesos** with Marathon
[Meiklejohn and Yoo 2016, in review]
- **Loquat**
Epidemic broadcast, partially replicated with
Decentralized Information Flow Control
[Meiklejohn 2016, in review]

Lasp Simulator

- **Docker Containers**
Docker containers for EPMD and Lasp runtime system

Lasp Simulator

- **Docker Containers**
Docker containers for EPMD and Lasp runtime system
- **Service Discovery**
Mechanisms for clustering Erlang nodes based on either **Marathon** application definitions or **Mesos-DNS**

Lasp Simulator

- **Docker Containers**
Docker containers for EPMD and Lasp runtime system
- **Service Discovery**
Mechanisms for clustering Erlang nodes based on either **Marathon** application definitions or **Mesos-DNS**
- **Instrumentation**
Transmission instrumentation and **divergence** measurement

Lasp Simulator

- Plumtree VM-to-VM
VM-to-VM communication performed using the Plumtree epidemic broadcast protocol

Lasp Simulator

- **Plumtree VM-to-VM**
VM-to-VM communication performed using the Plumtree epidemic broadcast protocol
- **Clients-as-processes**
Multiple clients per virtual machine, acting as mobile/IoT devices that periodically simulate and can be partitioned

Lasp Simulator

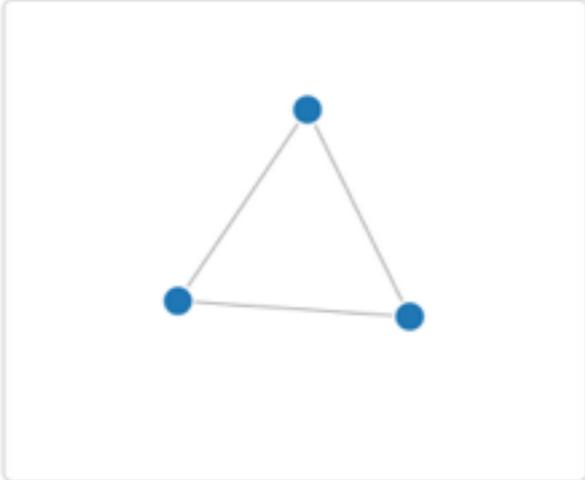
- **Plumtree VM-to-VM**
VM-to-VM communication performed using the **Plumtree** epidemic broadcast protocol
- **Clients-as-processes**
Multiple clients per virtual machine, acting as mobile/IoT devices that periodically simulate and can be partitioned
- **“Design in the small, run in the large”**
Runtime configuration change for client-to-VM ratio, allows for **single laptop design of multi-machine evaluations** of the programming model

localhost

Lasp



Topology



Controls

- [Simulate!](#)

Logs

- [divergence-false-lasp_orset-lasp_gcounter-2000000-2000-1000.csv](#)
- [divergence-false-lasp_orset-lasp_gcounter-2000000-2000-500.csv](#)
- [divergence-true-lasp_orset-lasp_gcounter-2000000-2000-500.csv](#)
- [state-client-false-lasp_orset-lasp_gcounter-2000000-2000-1000.csv](#)
- [state-client-false-lasp_orset-lasp_gcounter-2000000-2000-500.csv](#)
- [state-client-true-lasp_orset-lasp_gcounter-2000000-2000-500.csv](#)
- [state-server-false-lasp_orset-lasp_gcounter-2000000-2000-1000.csv](#)
- [state-server-false-lasp_orset-lasp_gcounter-2000000-2000-500.csv](#)
- [state-server-true-lasp_orset-lasp_gcounter-2000000-2000-500.csv](#)

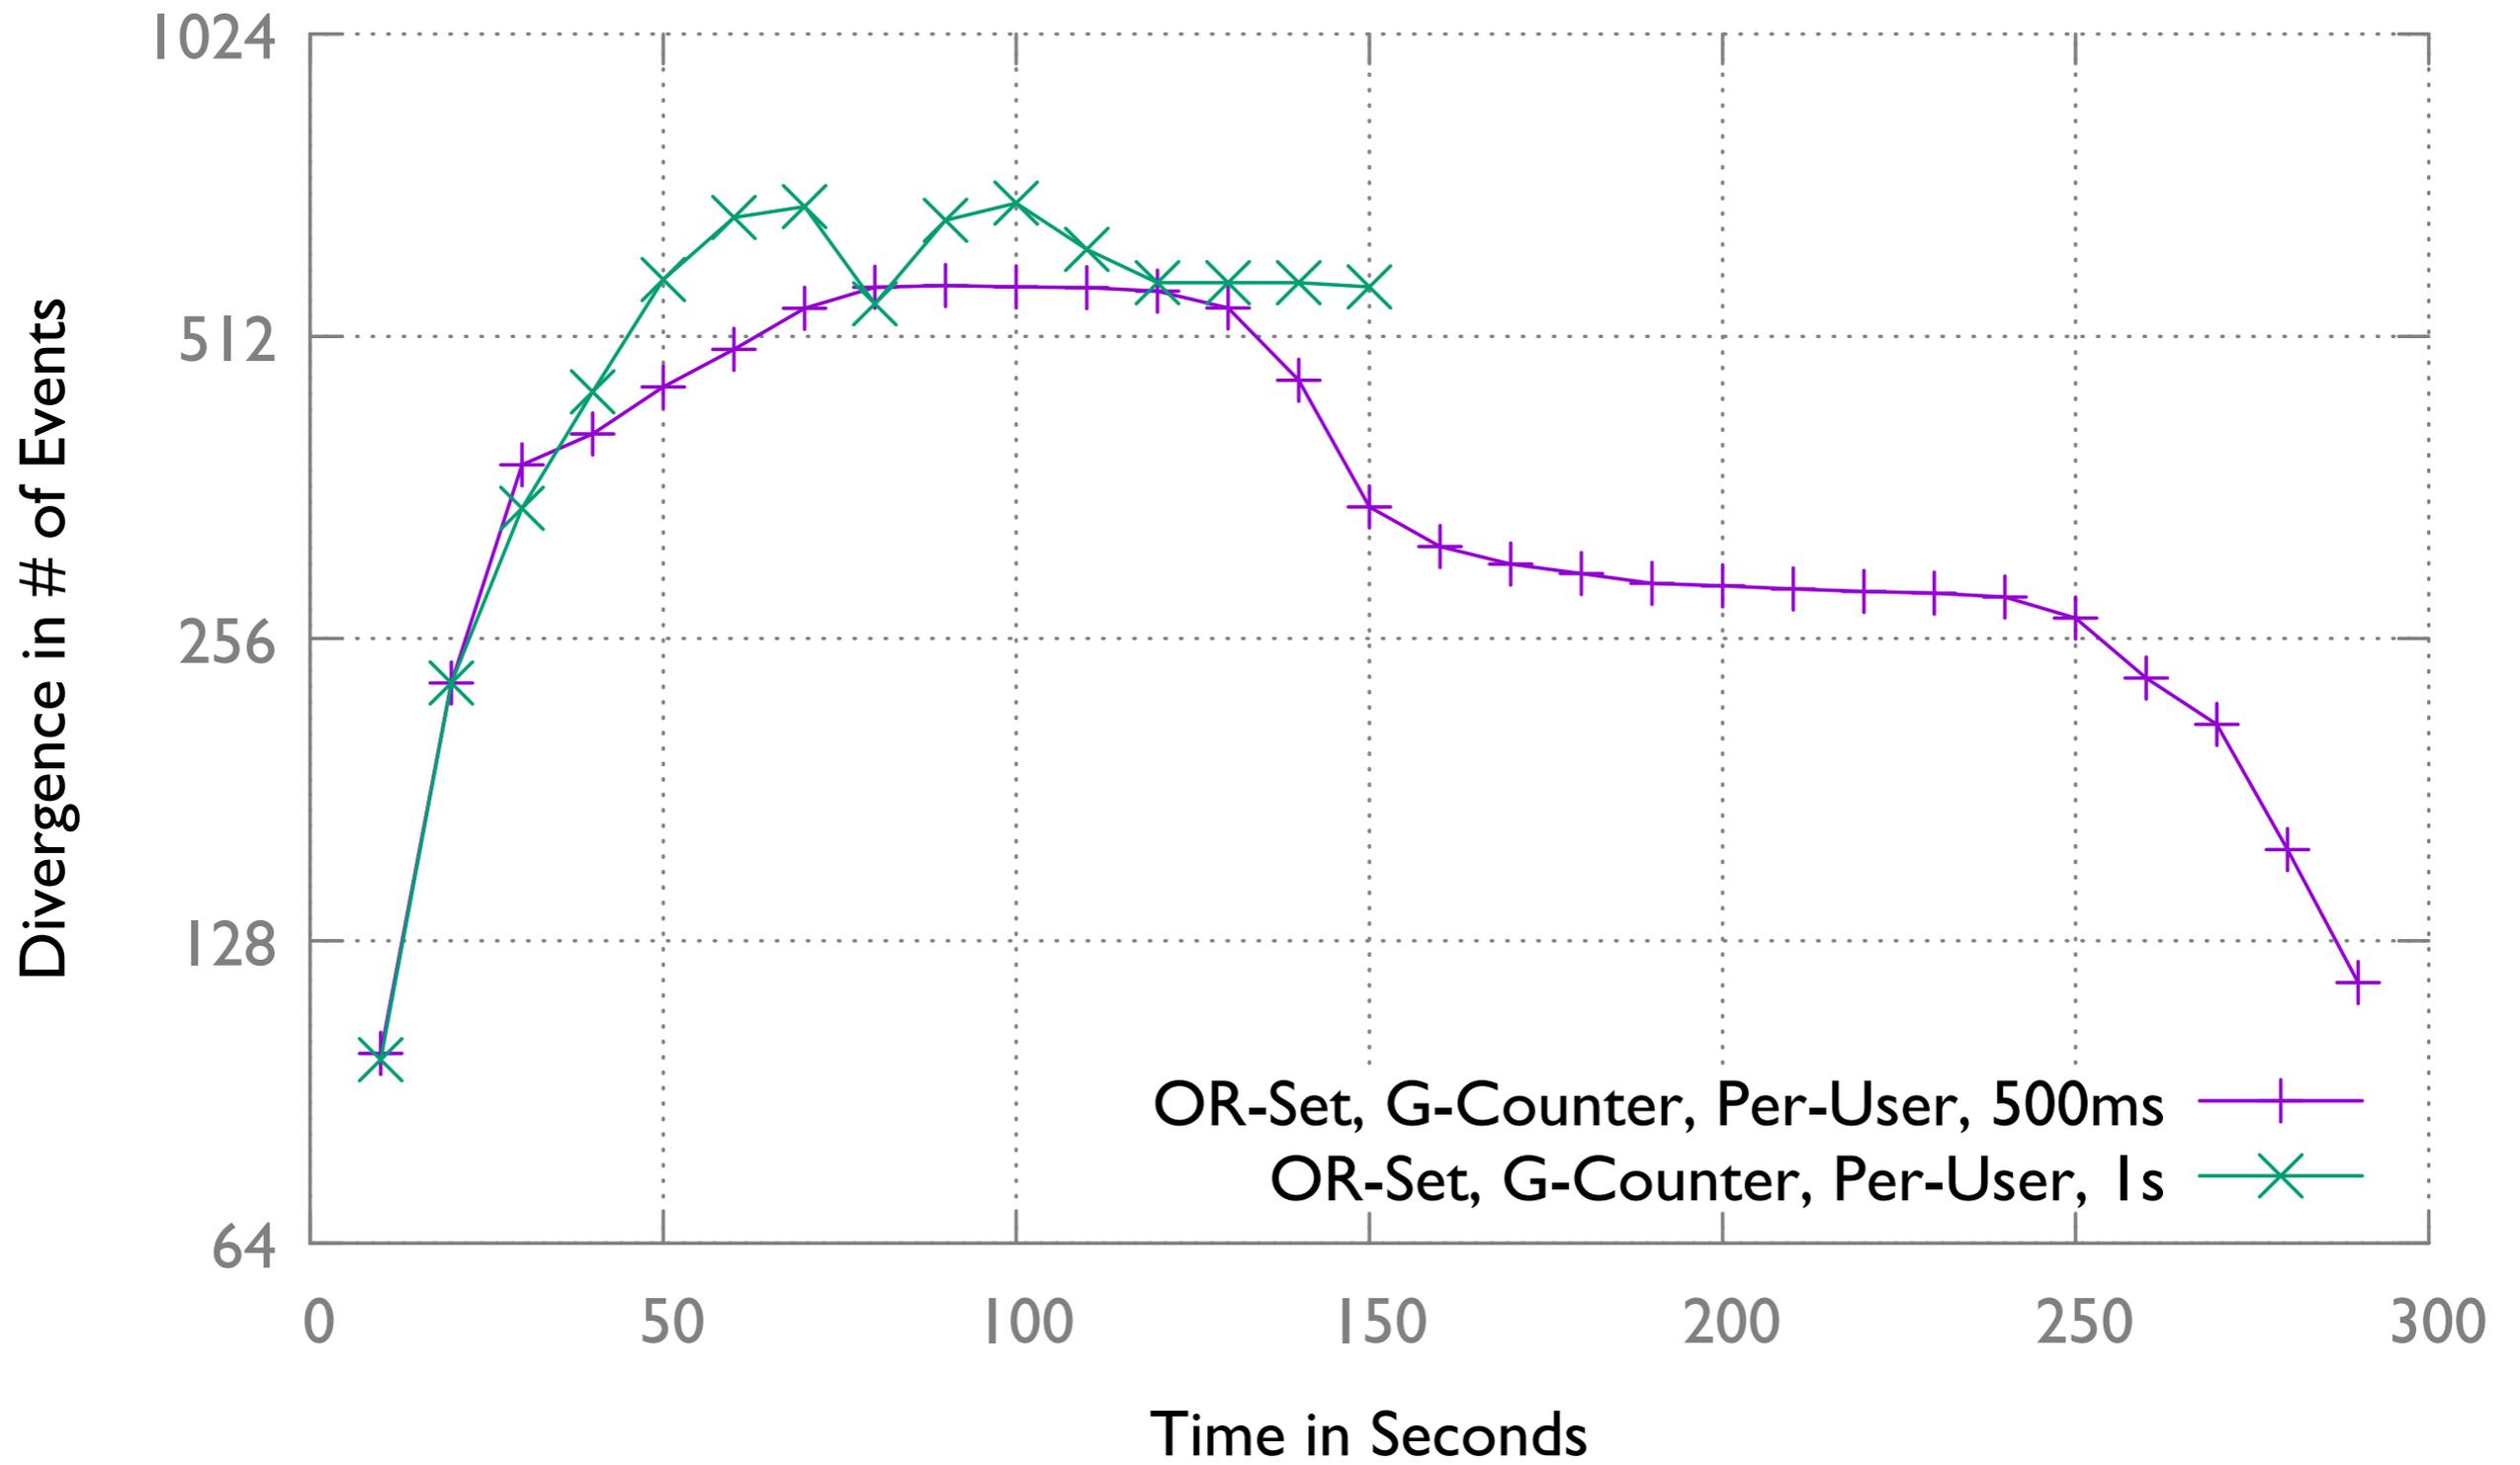
Plots

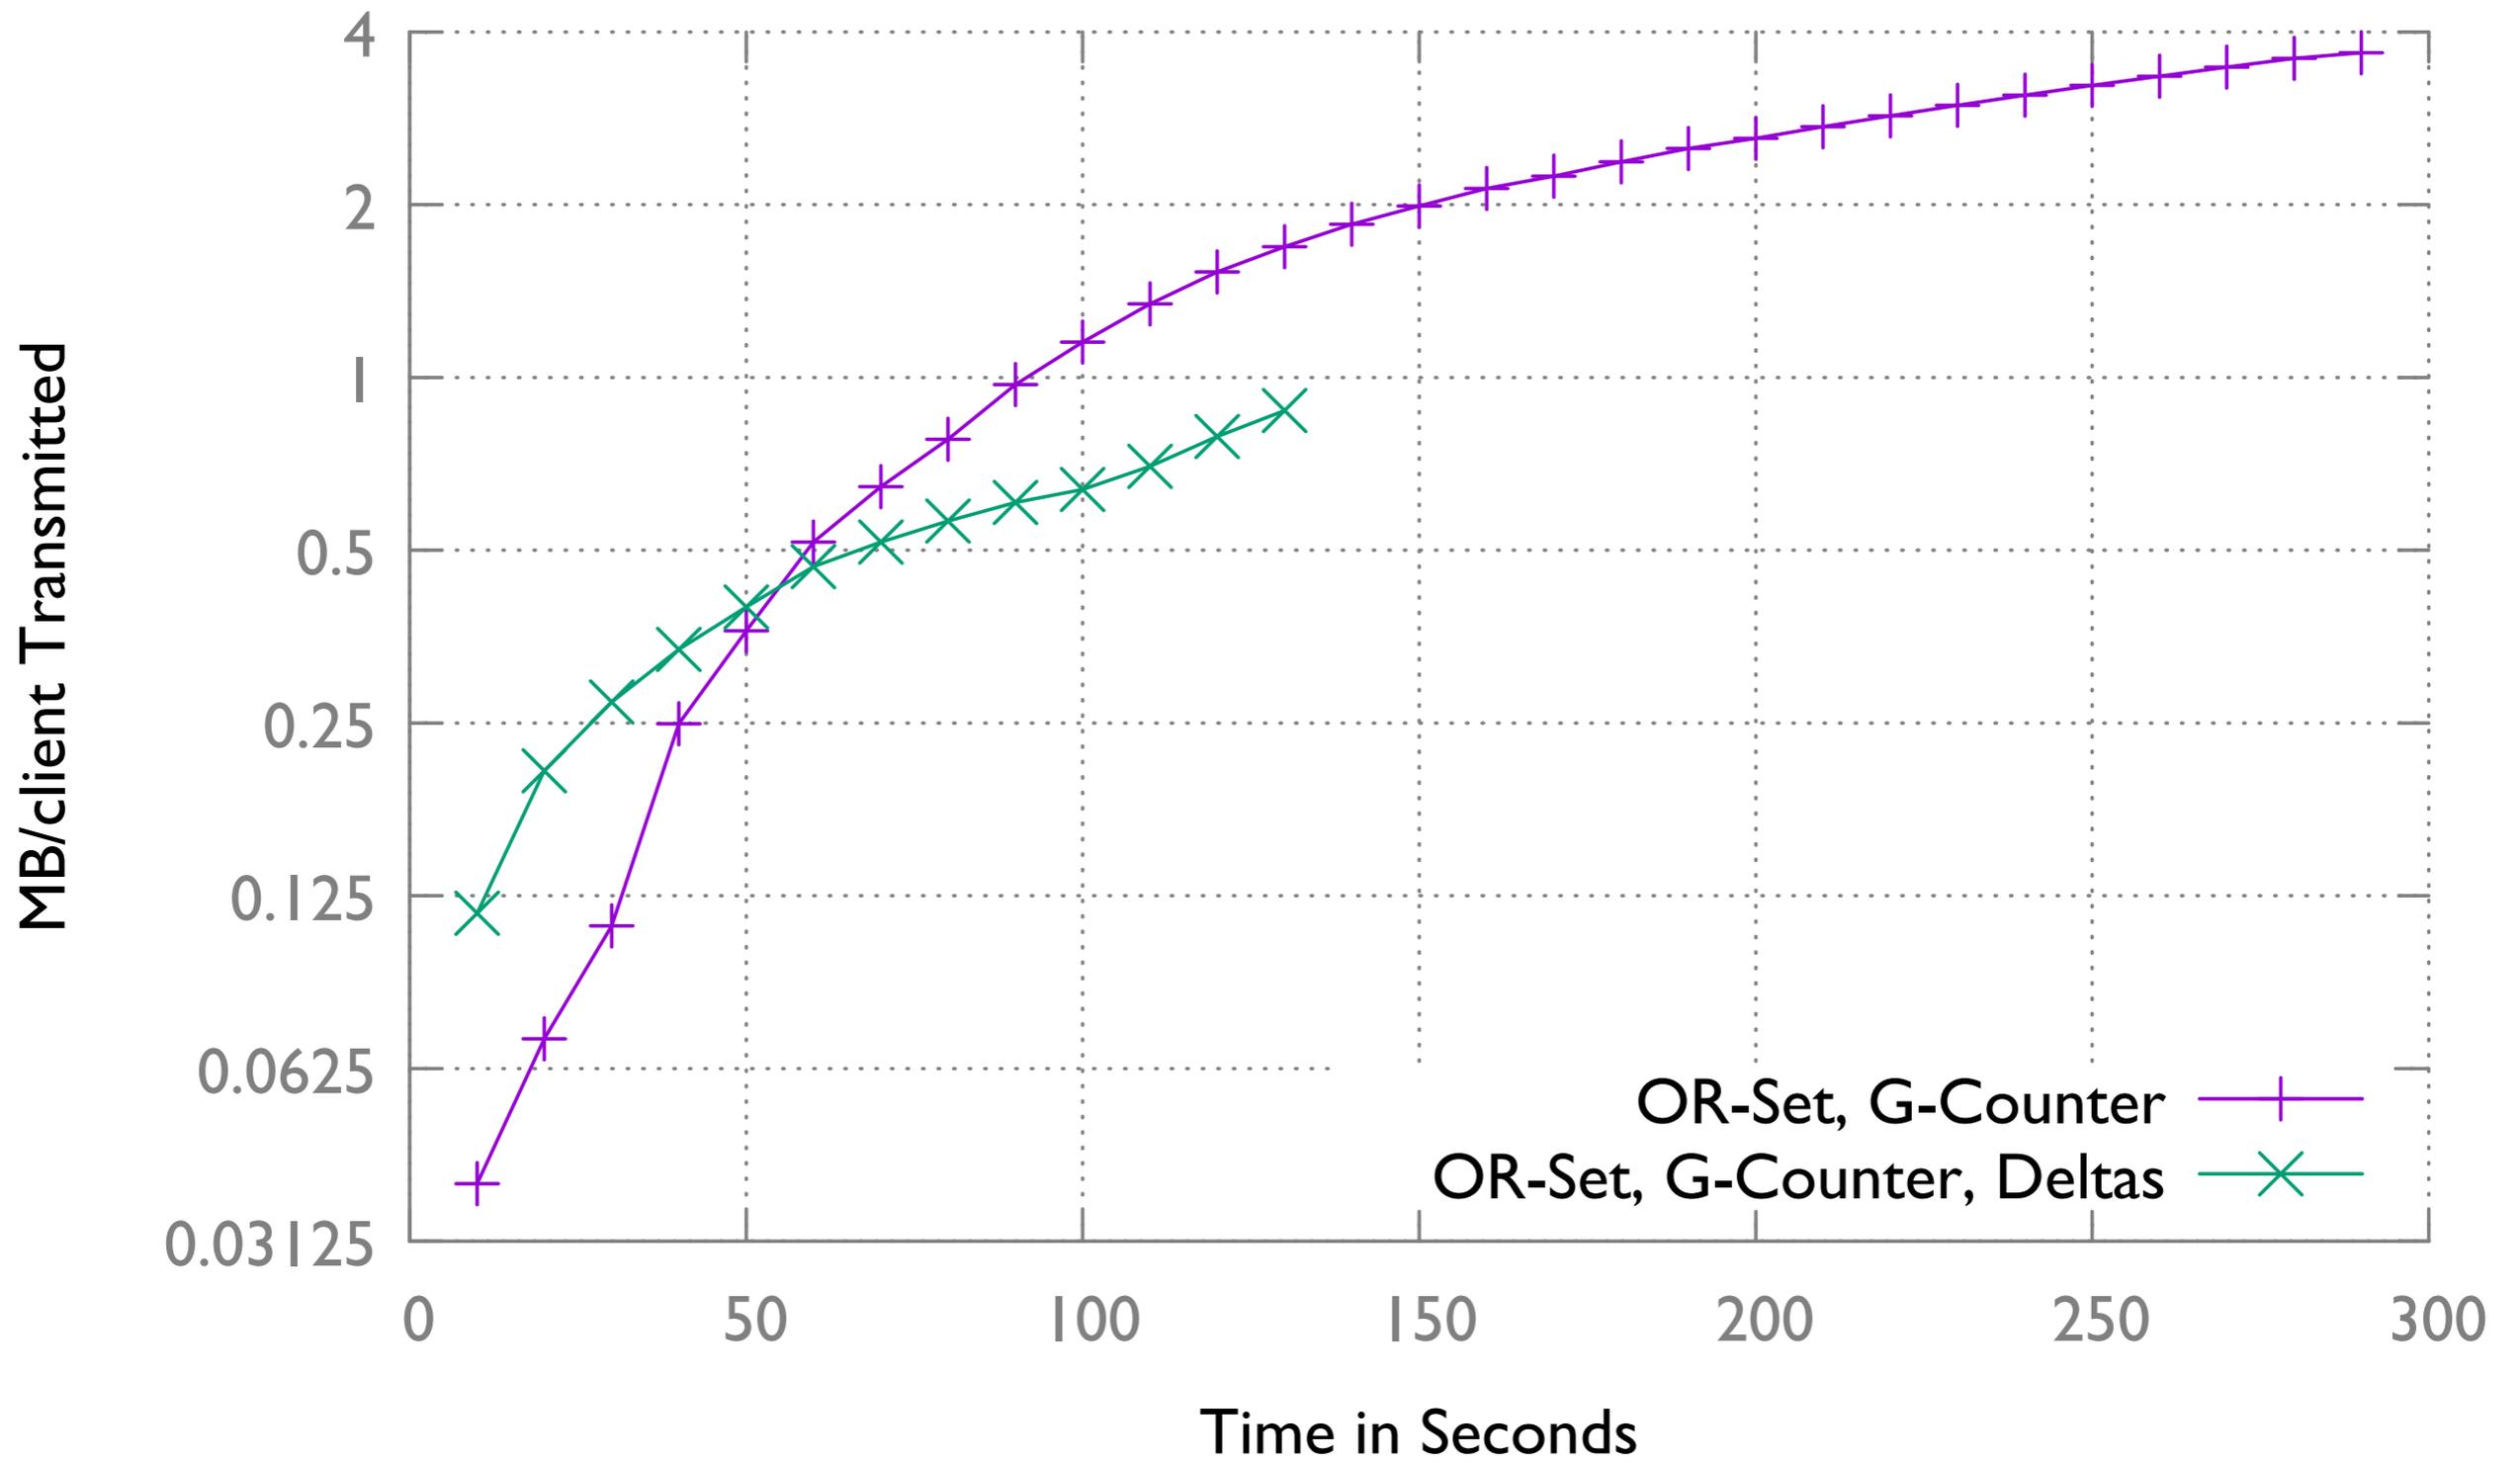
Plot 1: Divergence in # of Events vs Time in Seconds

Time (s)	OR-Set, G-Counter, Per-User, 500ms	OR-Set, G-Counter, Per-User, 1s
0	64	64
25	128	128
50	256	256
75	512	512
100	512	512
125	512	512
150	512	512
175	384	384
200	320	320
225	320	320
250	320	320
275	192	192
300	128	128

Plot 2: MB/client Transmitted vs Time in Seconds

Time (s)	OR-Set, G-Counter	OR-Set, G-Counter, Deltas
0	0.03125	0.03125
25	0.0625	0.0625
50	0.125	0.125
75	0.25	0.25
100	0.5	0.5
125	0.75	0.75
150	1.0	1.0
175	1.25	1.25
200	1.5	1.5
225	1.75	1.75
250	2.0	2.0
275	2.25	2.25
300	2.5	2.5





What's next?

Google Summer of Code (and my Ph.D.!!)

- Partial Evaluation
Optimize execution based on analysis and annotations
where we can determine local-vs-remote usage

Google Summer of Code (and my Ph.D.!!)

- **Partial Evaluation**
Optimize execution based on analysis and annotations where we can determine local-vs-remote usage
- **Optimizations**
General optimizations of the Erlang implementation of Lasp to improve performance of the runtime system

Google Summer of Code (and my Ph.D.!).)

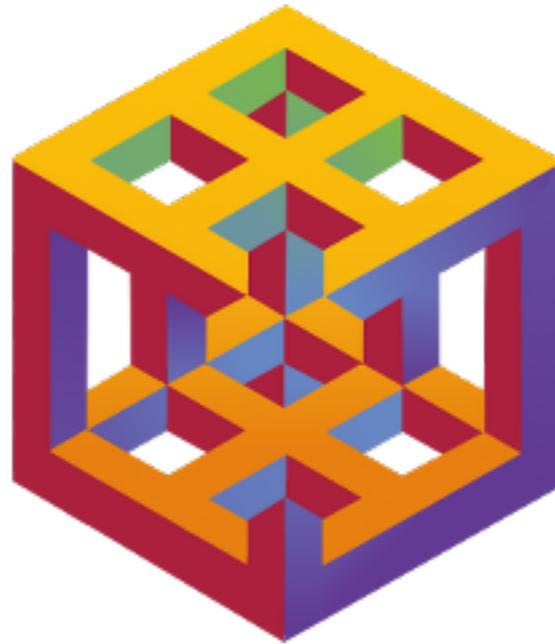
- **Partial Evaluation**
Optimize execution based on analysis and annotations where we can determine local-vs-remote usage
- **Optimizations**
General optimizations of the Erlang implementation of Lasp to improve performance of the runtime system
- **Elixir / Macros**
Provide a nicer way for working with Lasp, outside of the current syntax.

How do I learn
more?

Publications

- “Lasp: A Language for Distributed, Coordination-Free Programming”
ACM SIGPLAN PPDP 2015
- “Selective Hearing: An Approach to Distributed, Eventually Consistent Edge Computation”
IEEE W-PSDS 2015
- “The Implementation and Use of a Generic Dataflow Behaviour in Erlang”
ACM SIGPLAN Erlang Workshop '15
- “Lasp: A Language for Distributed, Eventually Consistent Computations with CRDTs”
PaPoC 2015
- “Declarative, Sliding Window Aggregations for Computations at the Edge”
IEEE EdgeCom 2016

Thanks!



Christopher Meiklejohn
@cmeik
<http://www.lasp-lang.org>