

Experiences Scaling Use of Google's Sawzall

[Jeffrey D. Oldham](#)

surname at company-name.com

Google, Inc.

2011-03-13

Programming, not Theory

Not focus on theory.

- No theorems.
- No models.
- No algorithms.

Focus on users' programming of parallel systems.

- Users write code.
 - Not system developers.
- Users write tests.

Summary

Sawzall eases writing map reductions.

Structured Sawzall scales.

Parallel system API should

- separate fundamental model concepts.
 - Ex: map reduction = **map** + **reduce** + **record enumeration**
- ease writing test code.

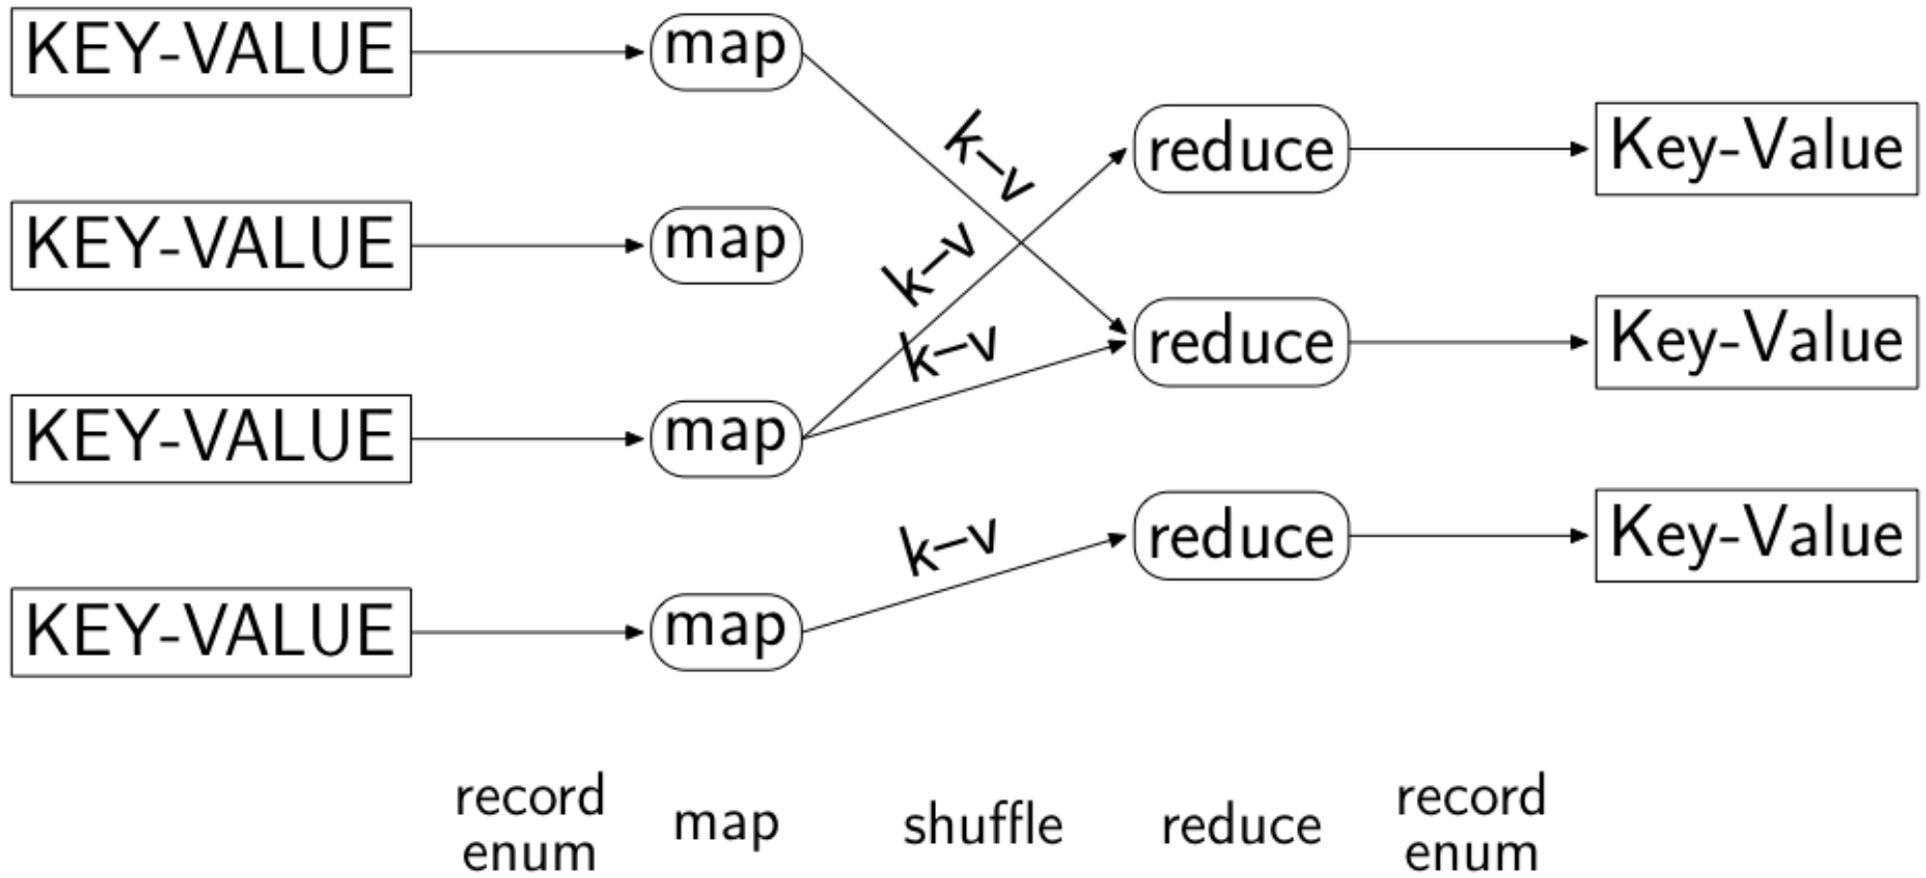
Outline

Map reductions and MapReduce

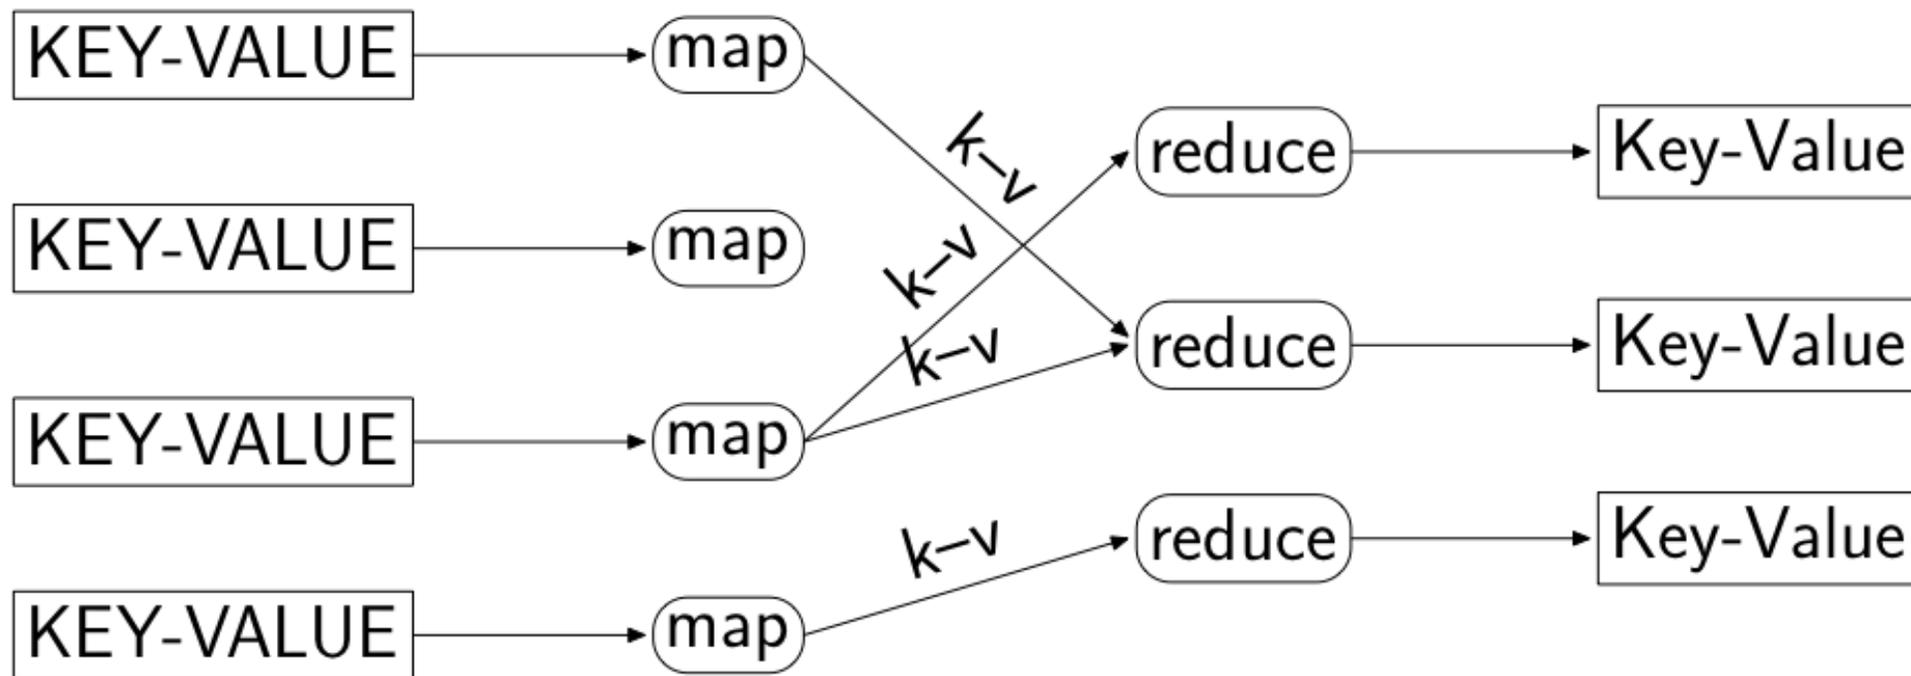
Map reductions and Saw + Sawzall

Structured Saw + Sawzall

Map Reduction



MapReduce: C++ Library



MapReduce

record
enum

shuffle

record
enum

user code

minor
config map
function

key
sharding

reduce
function

minor
config

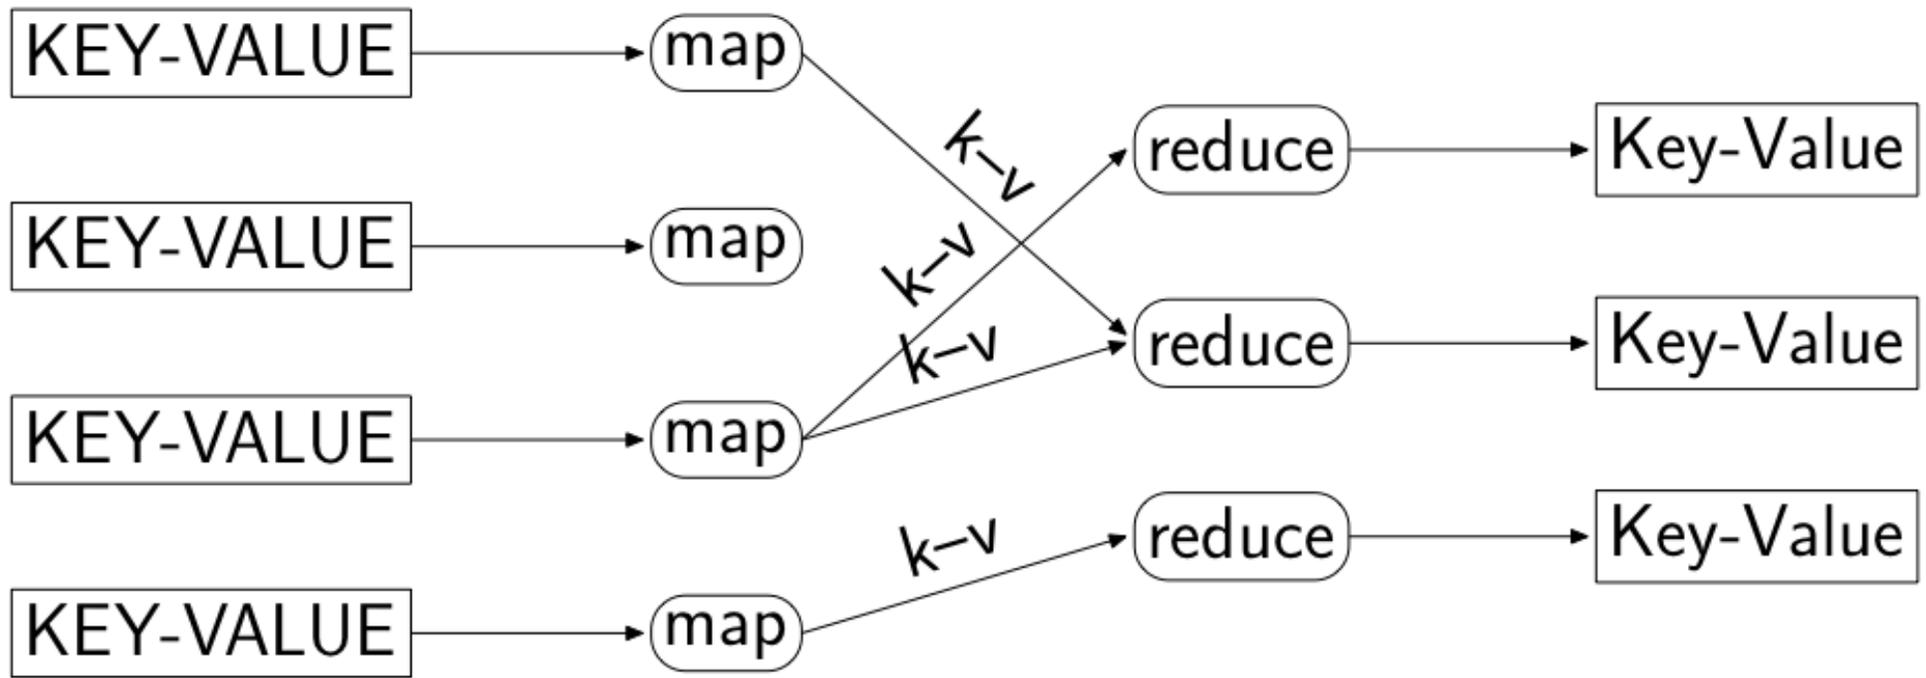
Outline

Map reductions and MapReduce

Map reductions and Saw + Sawzall

Structured Saw + Sawzall

Sawzall: Simpler Map Reductions



saw+Sawzall

record
enum

shuffle;
sharding reductions

record
enum

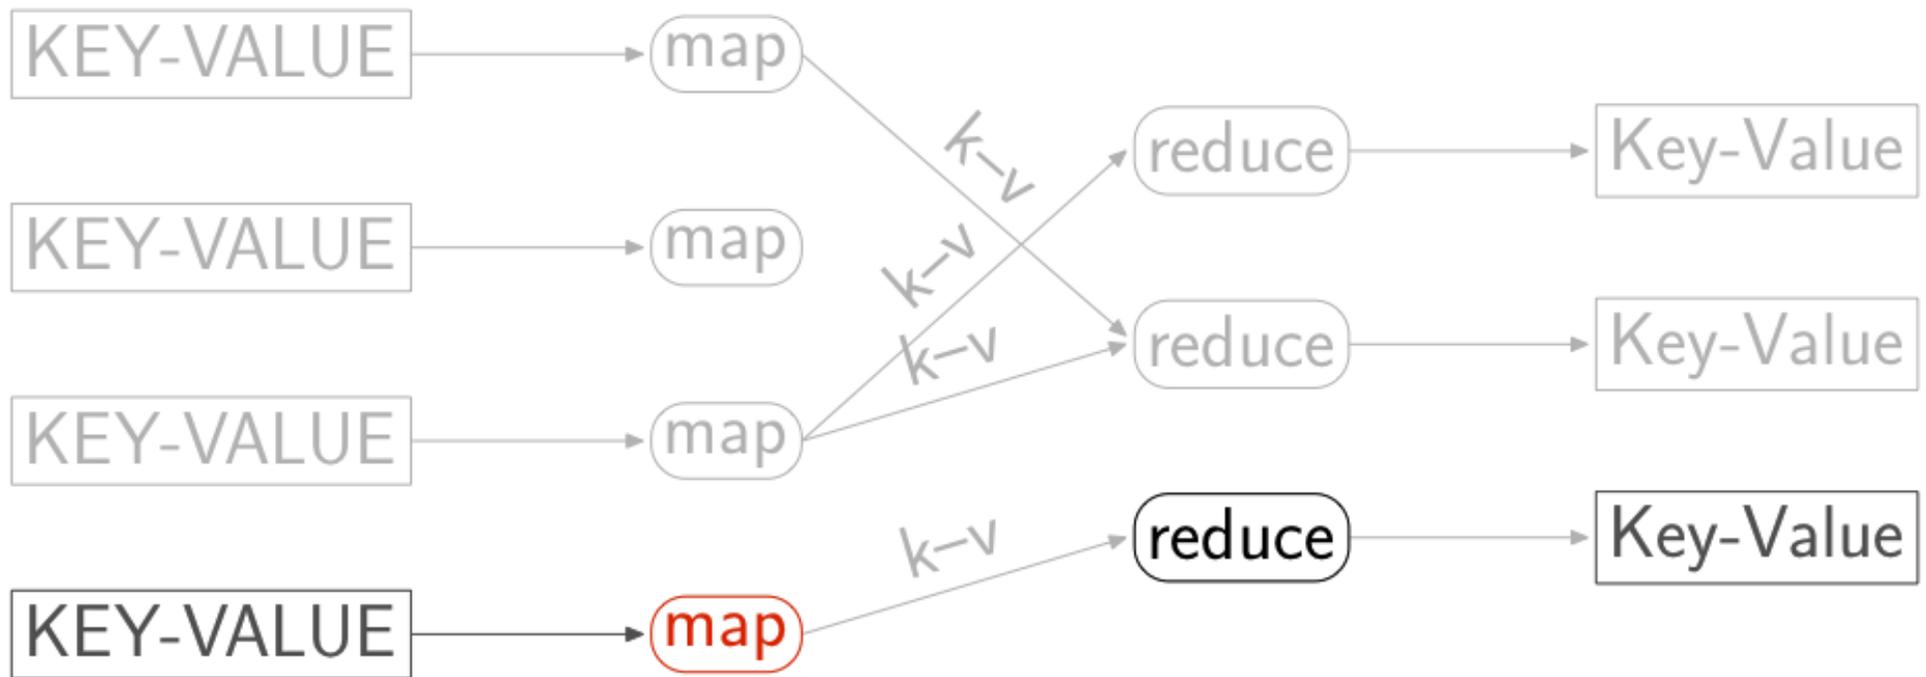
user code

trivial
config **mapping**

choose

trivial
config

Sawzall Mental Model: One Record



saw+Sawzall

record
enum

shuffle;
sharding

reductions

record
enum

user code

trivial
config mapping

choose

trivial
config

Sample Program

Compute the query number per latitude-longitude degree.

Sawzall query-location.sz1:

```
proto "querylog.proto"  
queries_per_degree: table sum[lat: int][lon: int] of int;  
log_record: QueryLogProto = input;  
loc: Location = locationinfo(log_record.ip);  
emit queries_per_degree[int(loc.lat)][int(loc.lon)] <- 1;
```

Shell code:

```
saw --program=query-location.sz1 --input=... --output=...
```

Saw + Sawzall Use

Used

- since 2003
- by 100s of Googlers
- in 1000s of programs
- to compute a lot of data
- that is directly or indirectly externally facing.

Outline

Map reductions and MapReduce

Map reductions and Saw + Sawzall

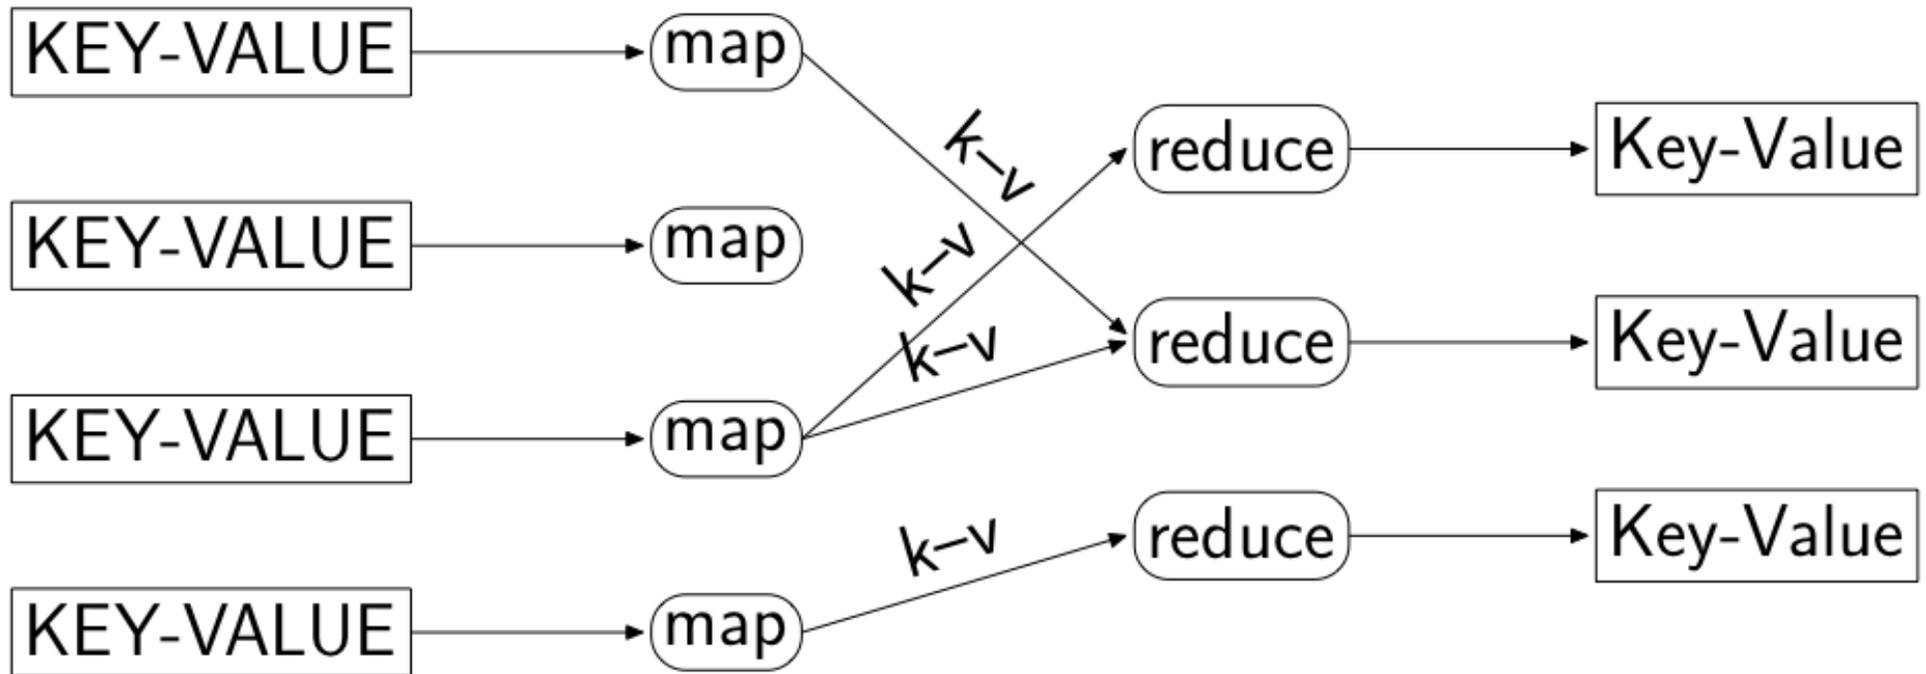
Structured Saw + Sawzall

Scaling Programs

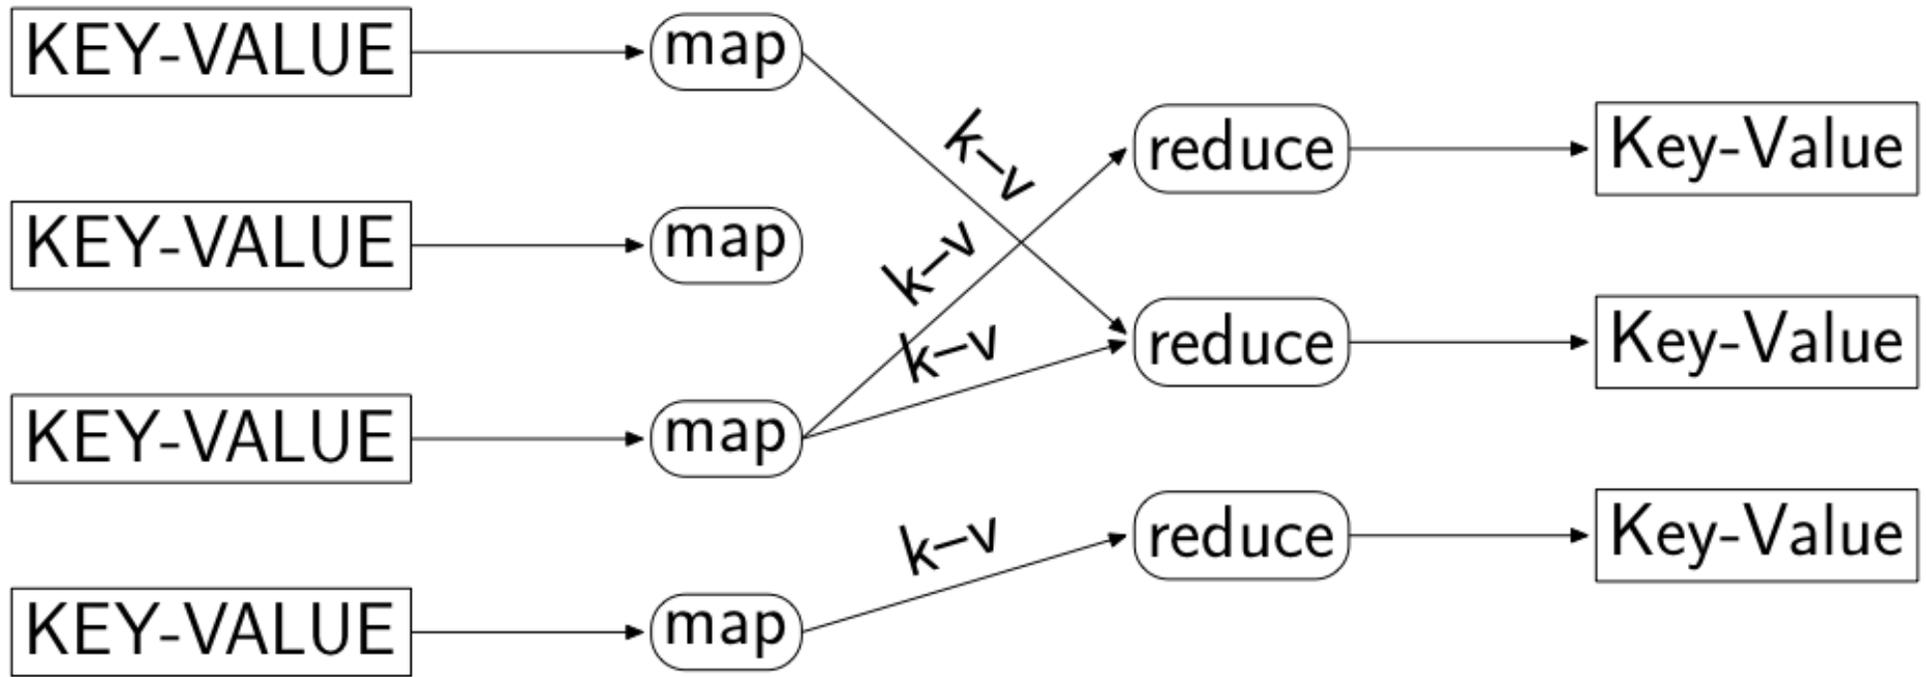
Code ecosystems support sharing tested code.

- + Sawzall function libraries have tests.
- Programs shared by copying.
- Typically untested.

Sawzall Testing Model: Map Reduction



Structured Pgms: Separate Concepts



saw+Sawzall

record
enum

shuffle;
sharding

reductions

record
enum

user code

trivial
config **map
function**

**reduce
function**

trivial
config

Sample Program

Compute the query number per latitude-longitude degree.

Sawzall query-location.sz1:

```
proto "querylog.proto"
```

```
queries_per_degree: table sum[lat: int][lon: int] of int;
```

```
log_record: QueryLogProto = input;
```

```
loc: Location = locationinfo(log_record.ip);
```

```
emit queries_per_degree[int(loc.lat)][int(loc.lon)] <- 1;
```

Shell code:

```
saw --program=query-location.sz1 --input=... --output=...
```

Structured Sample Program

Compute the query number per latitude-longitude degree.

Sawzall query-location.sz1:

```
proto "querylog.proto"
```

```
map: function(log: QueryLogProto,  
              reduce: function(int, int)) {  
  loc: Location = locationinfo(log_record.ip);  
  reduce(loc.lat, loc.lon);  
}
```

```
reduce: function(lat: int, lon: int) {  
  queries_per_degree: table sum[lat: int][lon: int] of int;  
  emit queries_per_degree[int(loc.lat)][int(loc.lon)] <- 1;  
}
```

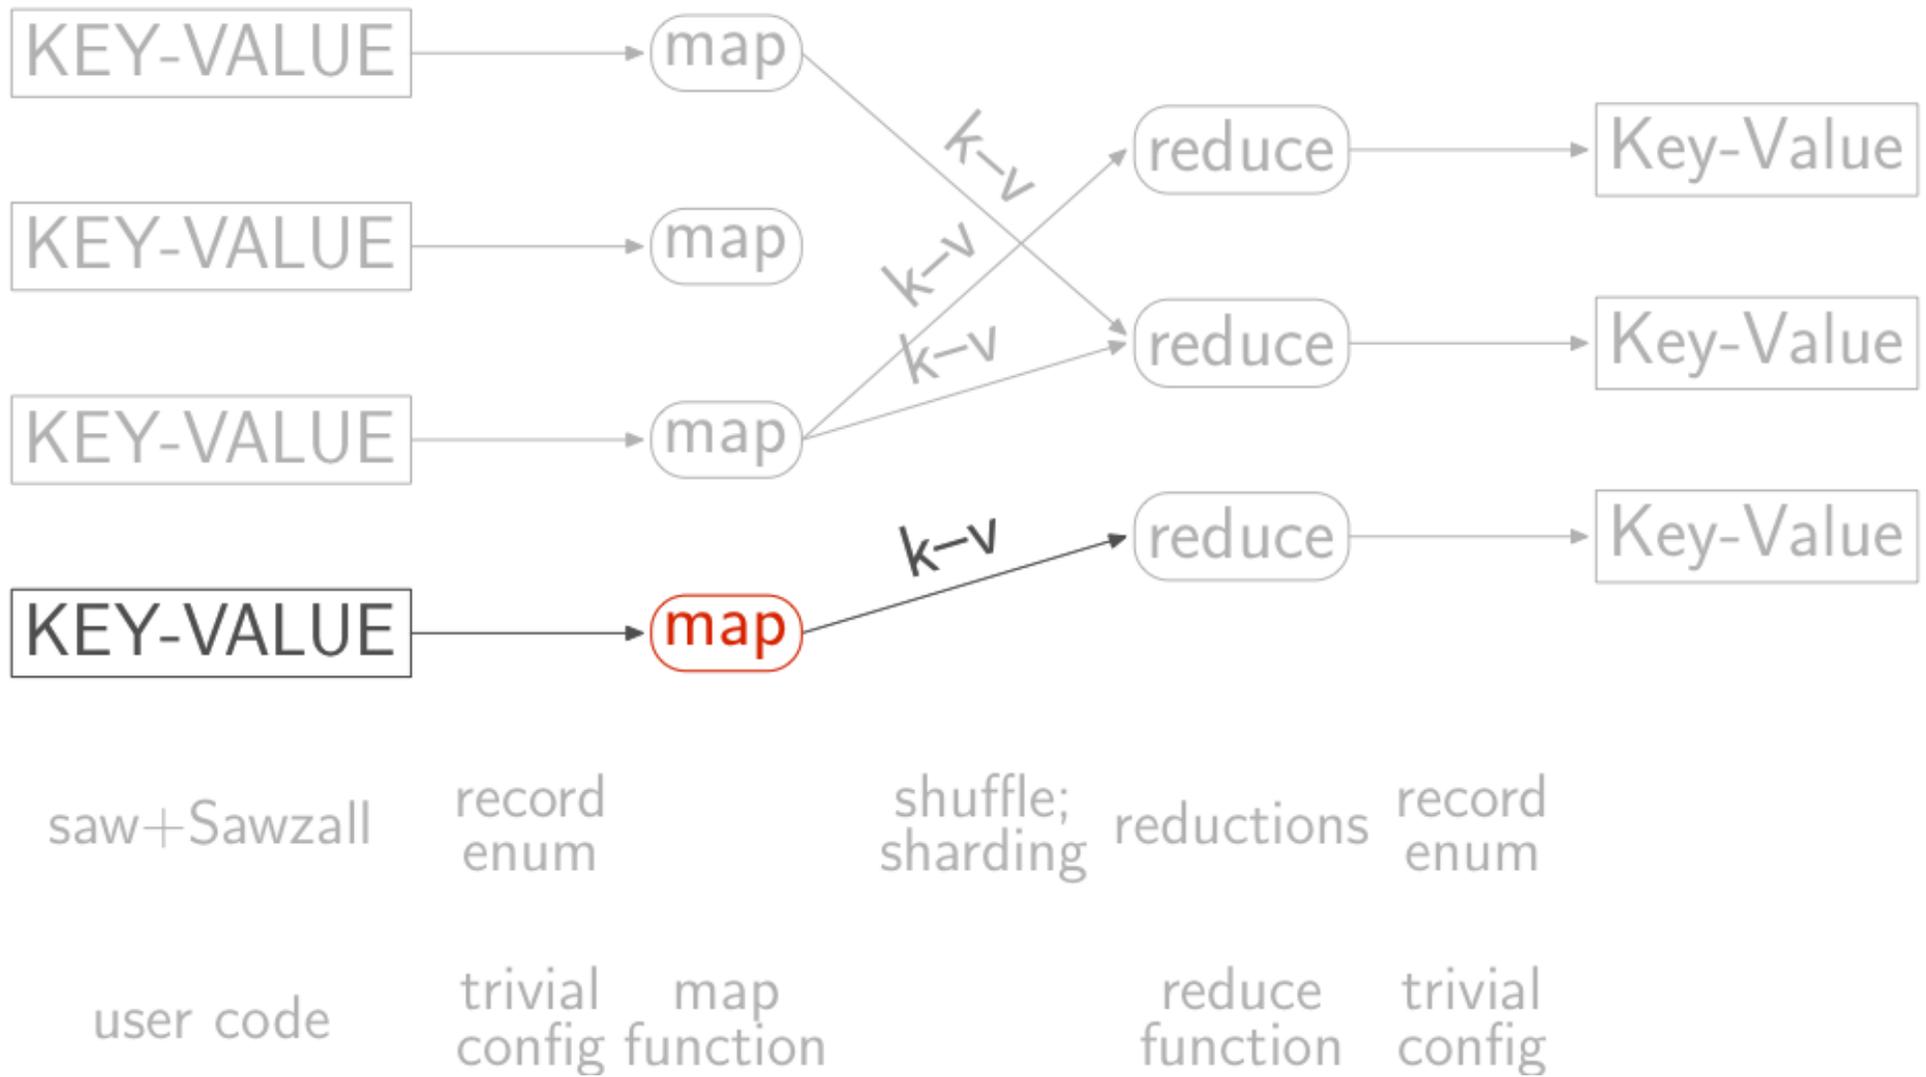
```
log_record: QueryLogProto = input;
```

```
map(log_record, reduce);
```

Shell code:

```
saw --program=query-location.sz1 --input=... --output=...
```

Structured Testing Model



Test Structured Programs

Test map functions ...

- one record at a time ...
- using mocked reduce function.
- Advantages:
 - No distributed I/O.
 - Single processor only.

Not test

- reduce functions or
- order enumeration.

Summary

Sawzall eases writing map reductions.

Structured Sawzall scales.

Parallel system API should

- separate fundamental model concepts.
 - Ex: map reduction = **map** + **reduce** + **record enumeration**
- ease writing test code.

Experiences Scaling Use of Google's Sawzall

[Jeffrey D. Oldham](#)

surname at company-name.com

Google, Inc.

2011-03-13

References

Sawzall

- [Pike et al.](#)
- [Open-source implementation](#)
- [Wikipedia article](#)

MapReduce

- Dean and Ghemawat ([2004](#), [2008](#))
- [Wikipedia article](#)