



# **Scripting the cloud with Skywriting**

---

**Derek G. Murray   Steven Hand**  
**University of Cambridge**

**A universal model?**

---

**MapReduce**

**A universal model?**

---

~~MapReduce~~

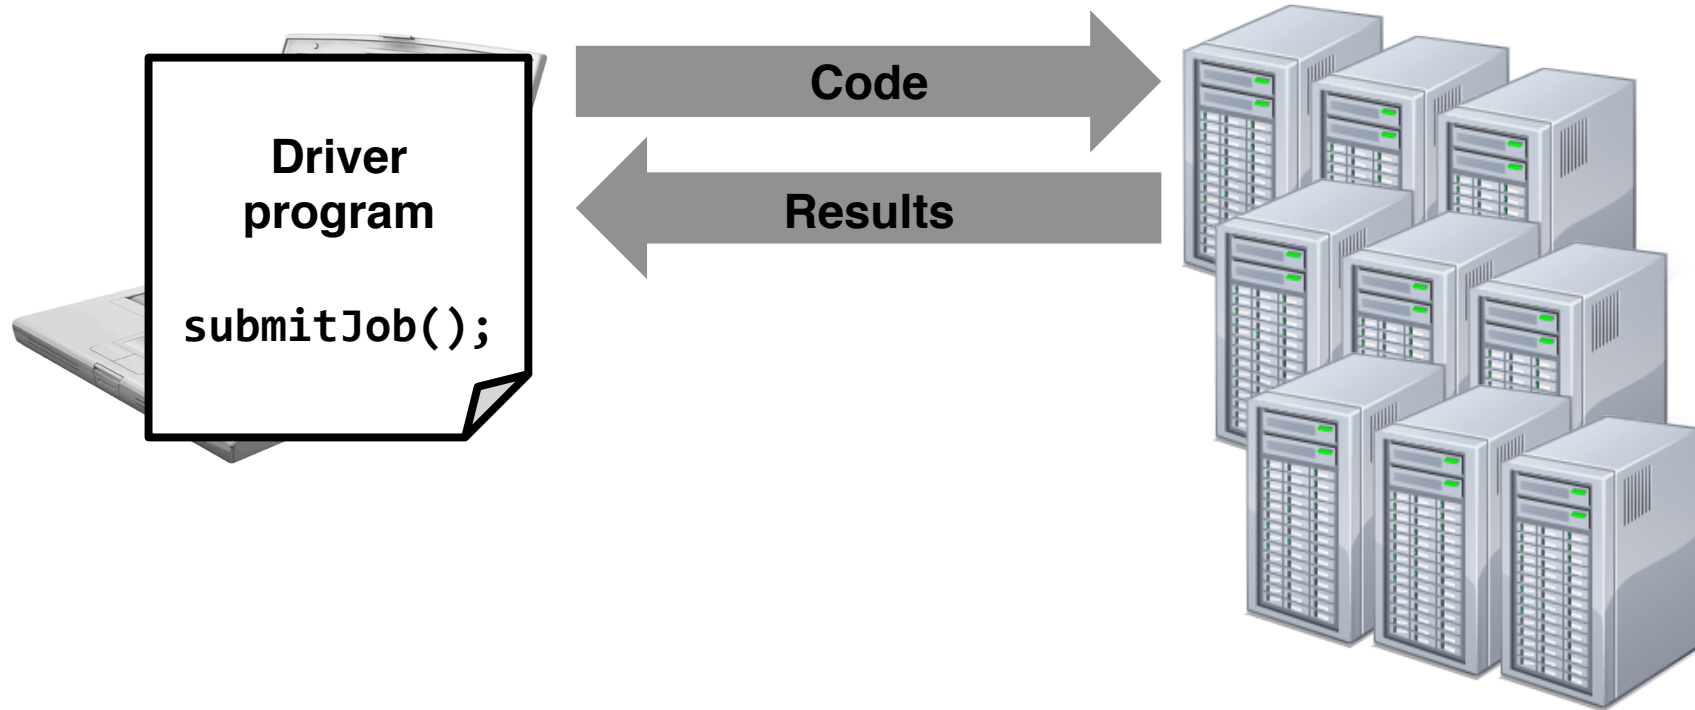
# A universal model!

---



# Move computation to the data

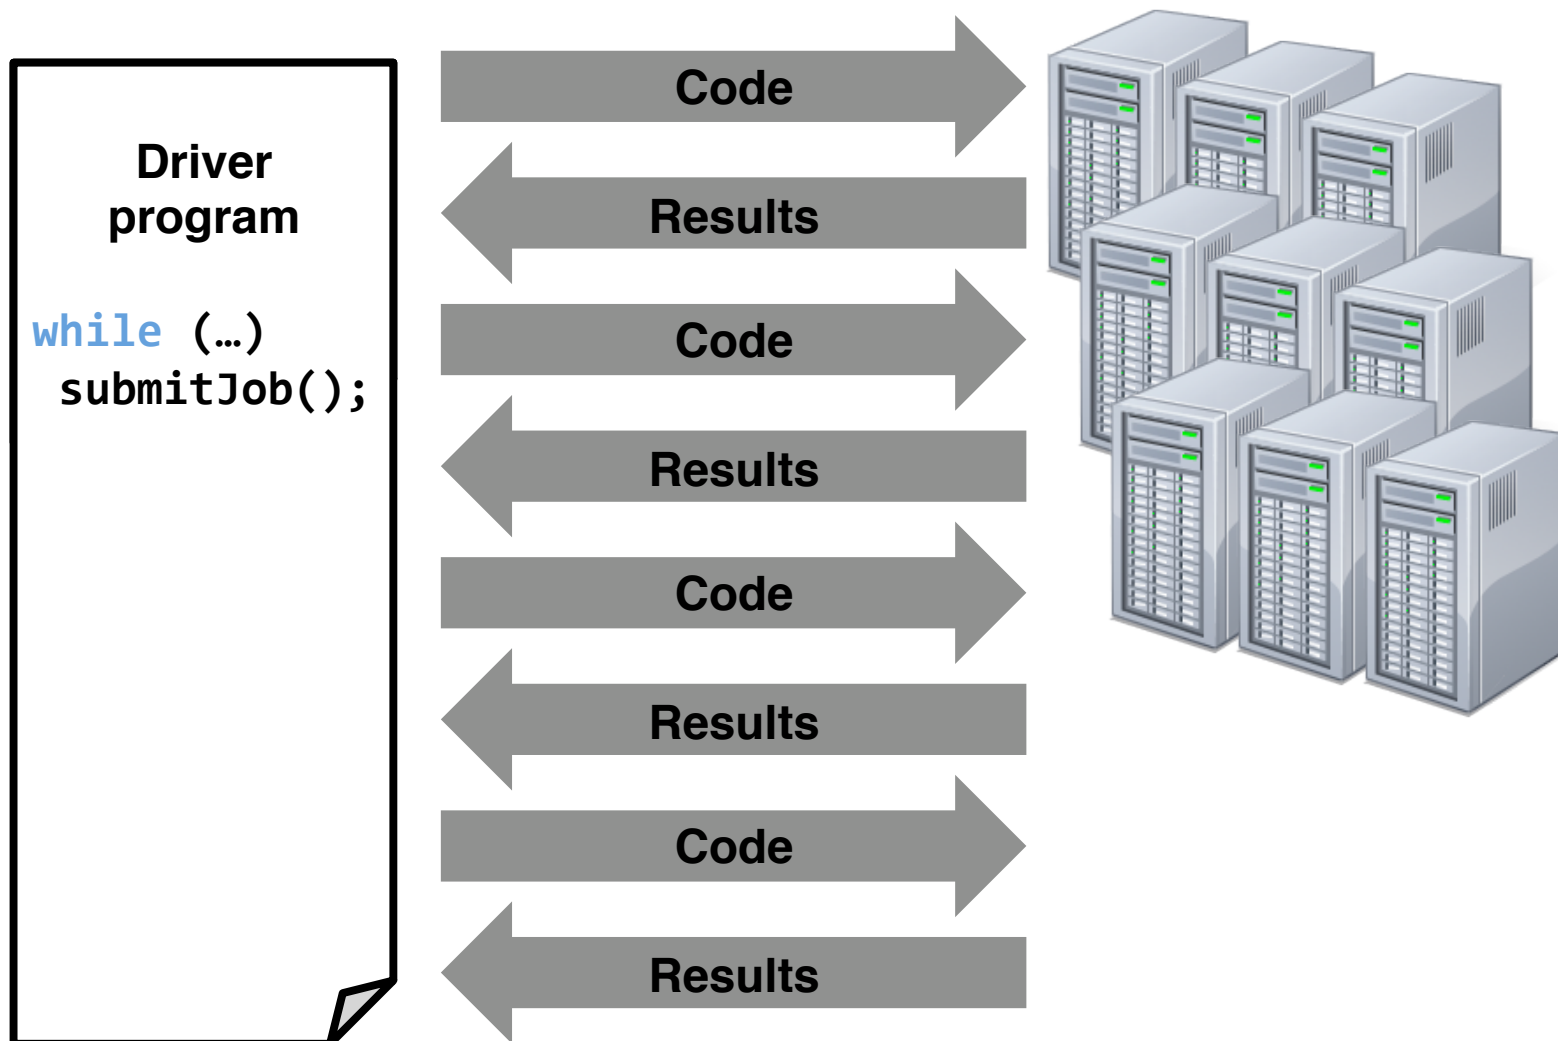
---



```
while (!converged)  
    do work in parallel;
```

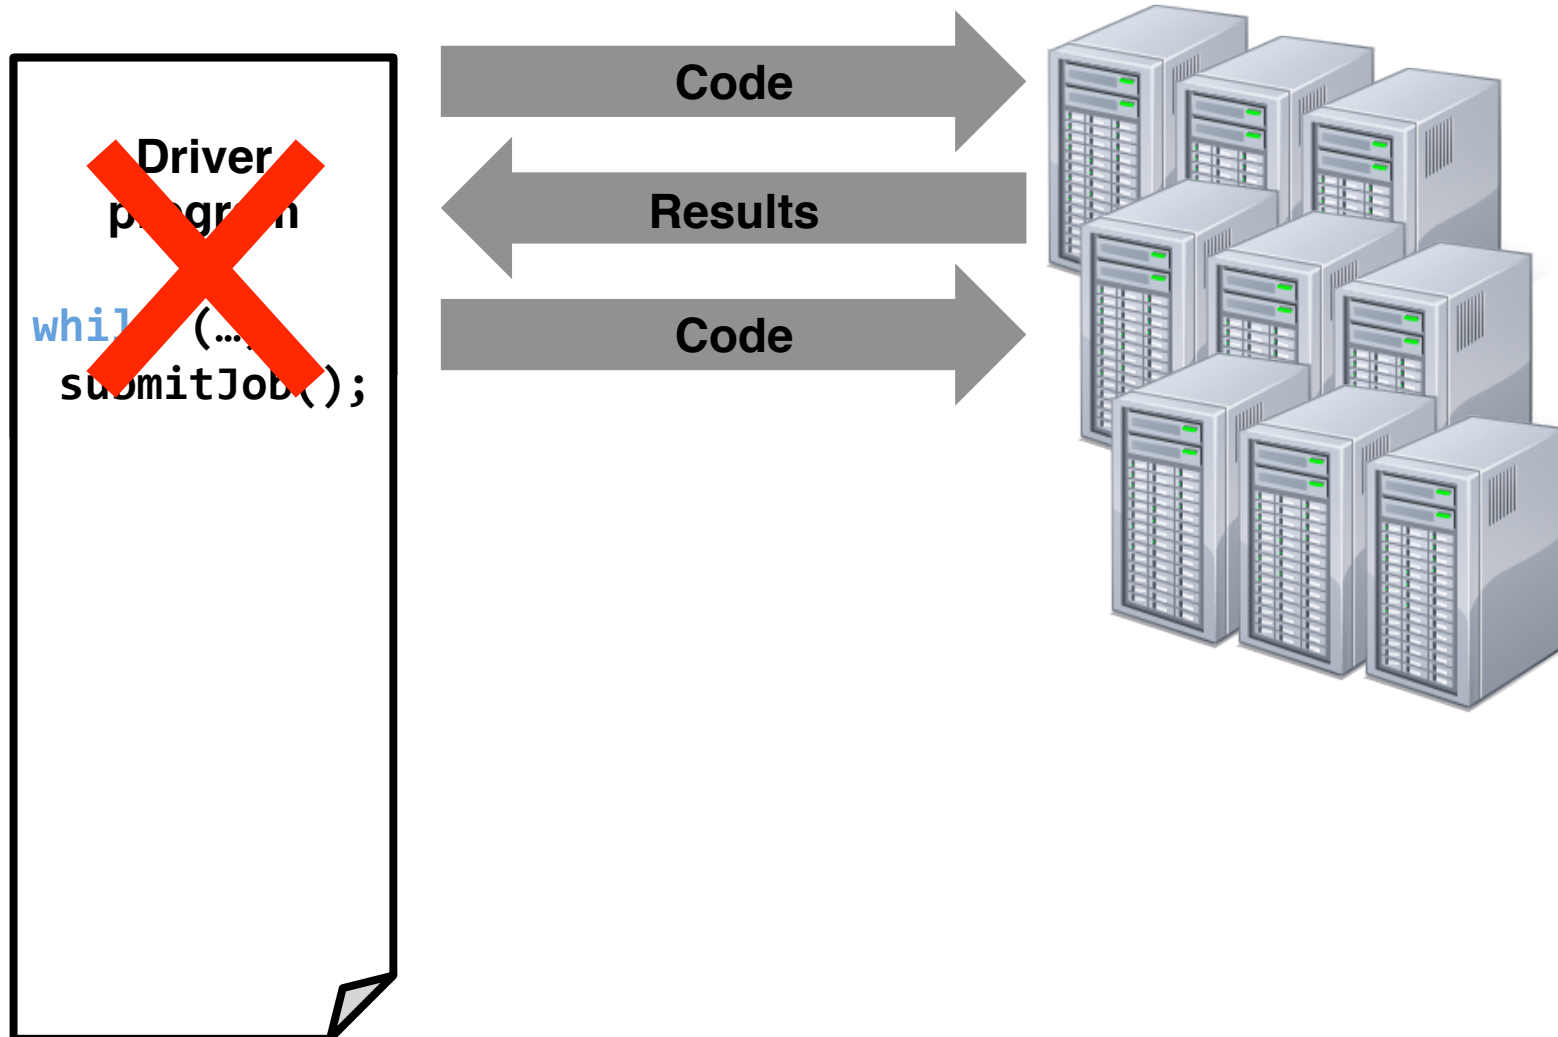
# Iterative algorithm

---



# Iterative algorithm

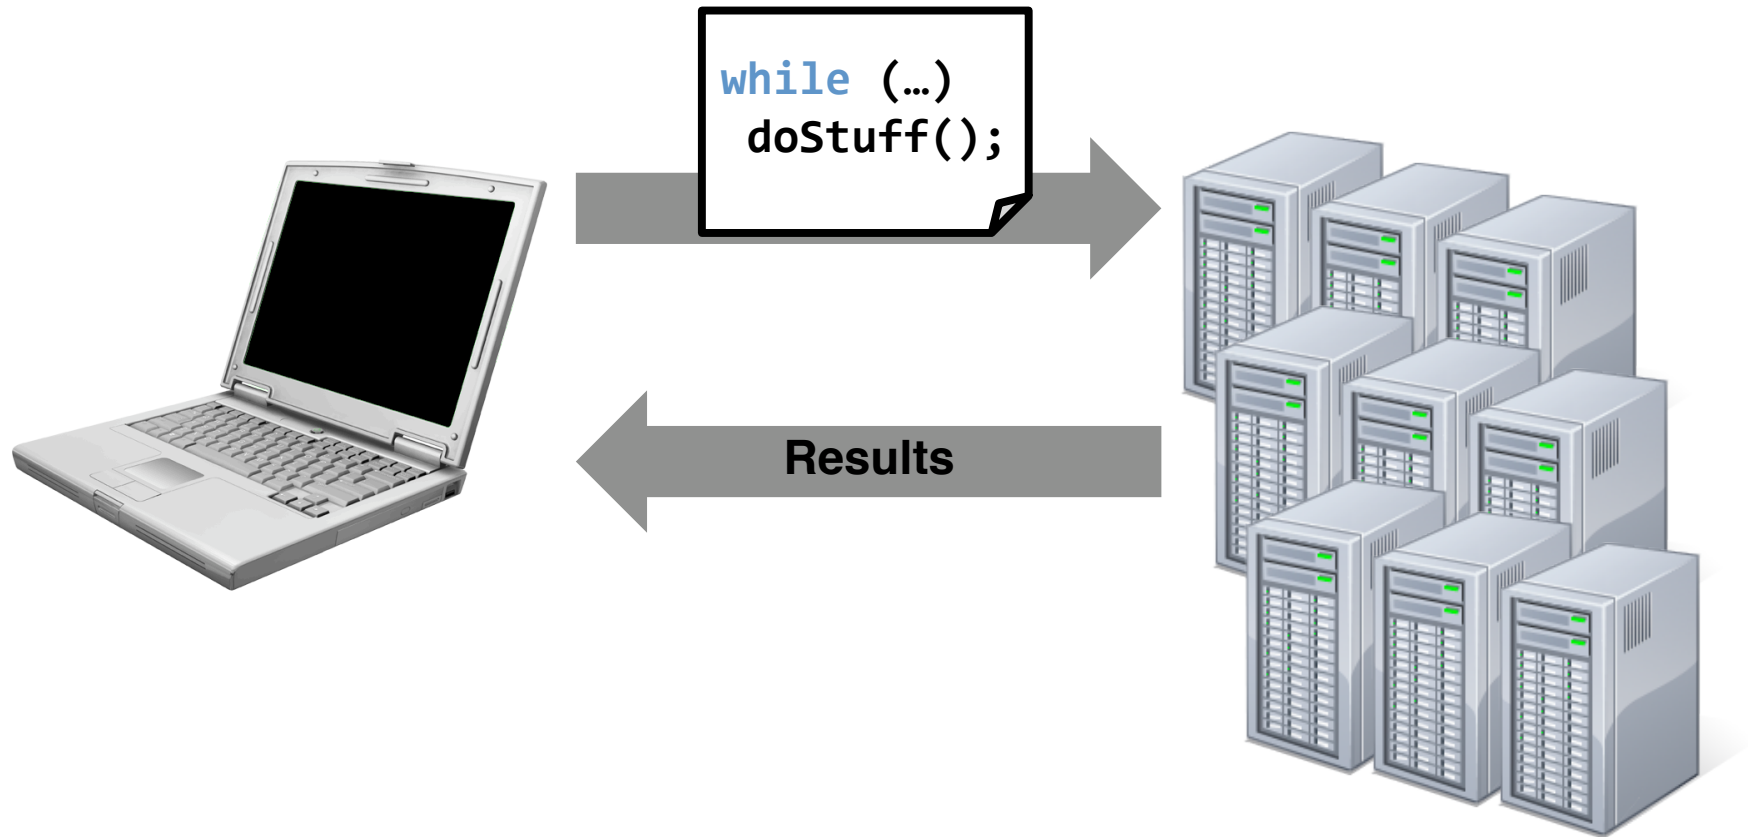
---





# Skywriting

---



# Skywriting

---

- JavaScript-like job specification language
  - Supports functional programming
  - Data-dependent control flow
- Distributed execution engine
  - Locality-based scheduling
  - Fault tolerance
  - Thread migration

# Spawning a task

---

```
function f(x) { return x + 1; }
```

```
res1 = spawn(f, [42]);
```

# Task dependencies

---

```
function f(x) { return x + 1; }  
function g(y) { ... }
```

```
res1 = spawn(f, [42]);  
res2 = spawn(g, [res1]);
```

res1 and res2 are *future references*

# Logistic regression

---

```
points = [...]; // List of partitions
w = ...;        // Random initial value
for (i in range(0, ITERATIONS)) {
    w_old = w;
    results = [];
    for (part in points) {
        results += spawn(log_reg, [part, w_old]);
    }
    w = spawn(update, [w_old, results]);
}
```

# Logistic regression

---

```
points = [...]; // List of partitions
w = ...;       // Random initial value
do {
    w_old = w;
    results = [];
    for (part in points) {
        results += spawn(log_reg, [part, w_old]);
    }
    w = spawn(update, [w_old, results]);
    done = spawn(converged, [w_old, w]);
} while (!*done);
```

# Logistic regression

---

```
points = [...]; // List of partitions
w = ...;        // Random initial value
do {
  w_old = w;
  results = [];
  for (part in points) {
    results += spawn(log_reg, [part, w_old]);
  }
}
```

\*-operator dereferences (forces) a future

```
} while (!*done);
```

# Implementation status


---

- Implemented in 4000 lines of Python
  - Also: Java, C and .NET bindings
- Many additional features
  - Native code execution
  - Introspection
  - Conditional synchronisation
- Available as open-source
  - <http://github.com/mrry/skywriting>



# Job creation overhead

---



# Future directions

---

- Multiple-scale parallel computing
  - Multiple cores, machines and clouds
- Streaming computations
  - Piping high-bandwidth data between tasks
- Better language integration
  - Hosted Skywriting on CLR or JVM

# Conclusions

---

- Turing-complete programming language for distributed computation
- Runs real jobs with low overhead
- Lots more still to do!

# Questions?

---

- Email
  - Derek.Murray@cl.cam.ac.uk
- Project website
  - <http://www.cl.cam.ac.uk/netos/skywriting/>

