



ROYAL INSTITUTE
OF TECHNOLOGY



Arcon

Continuous and Deep Data Stream Analytics

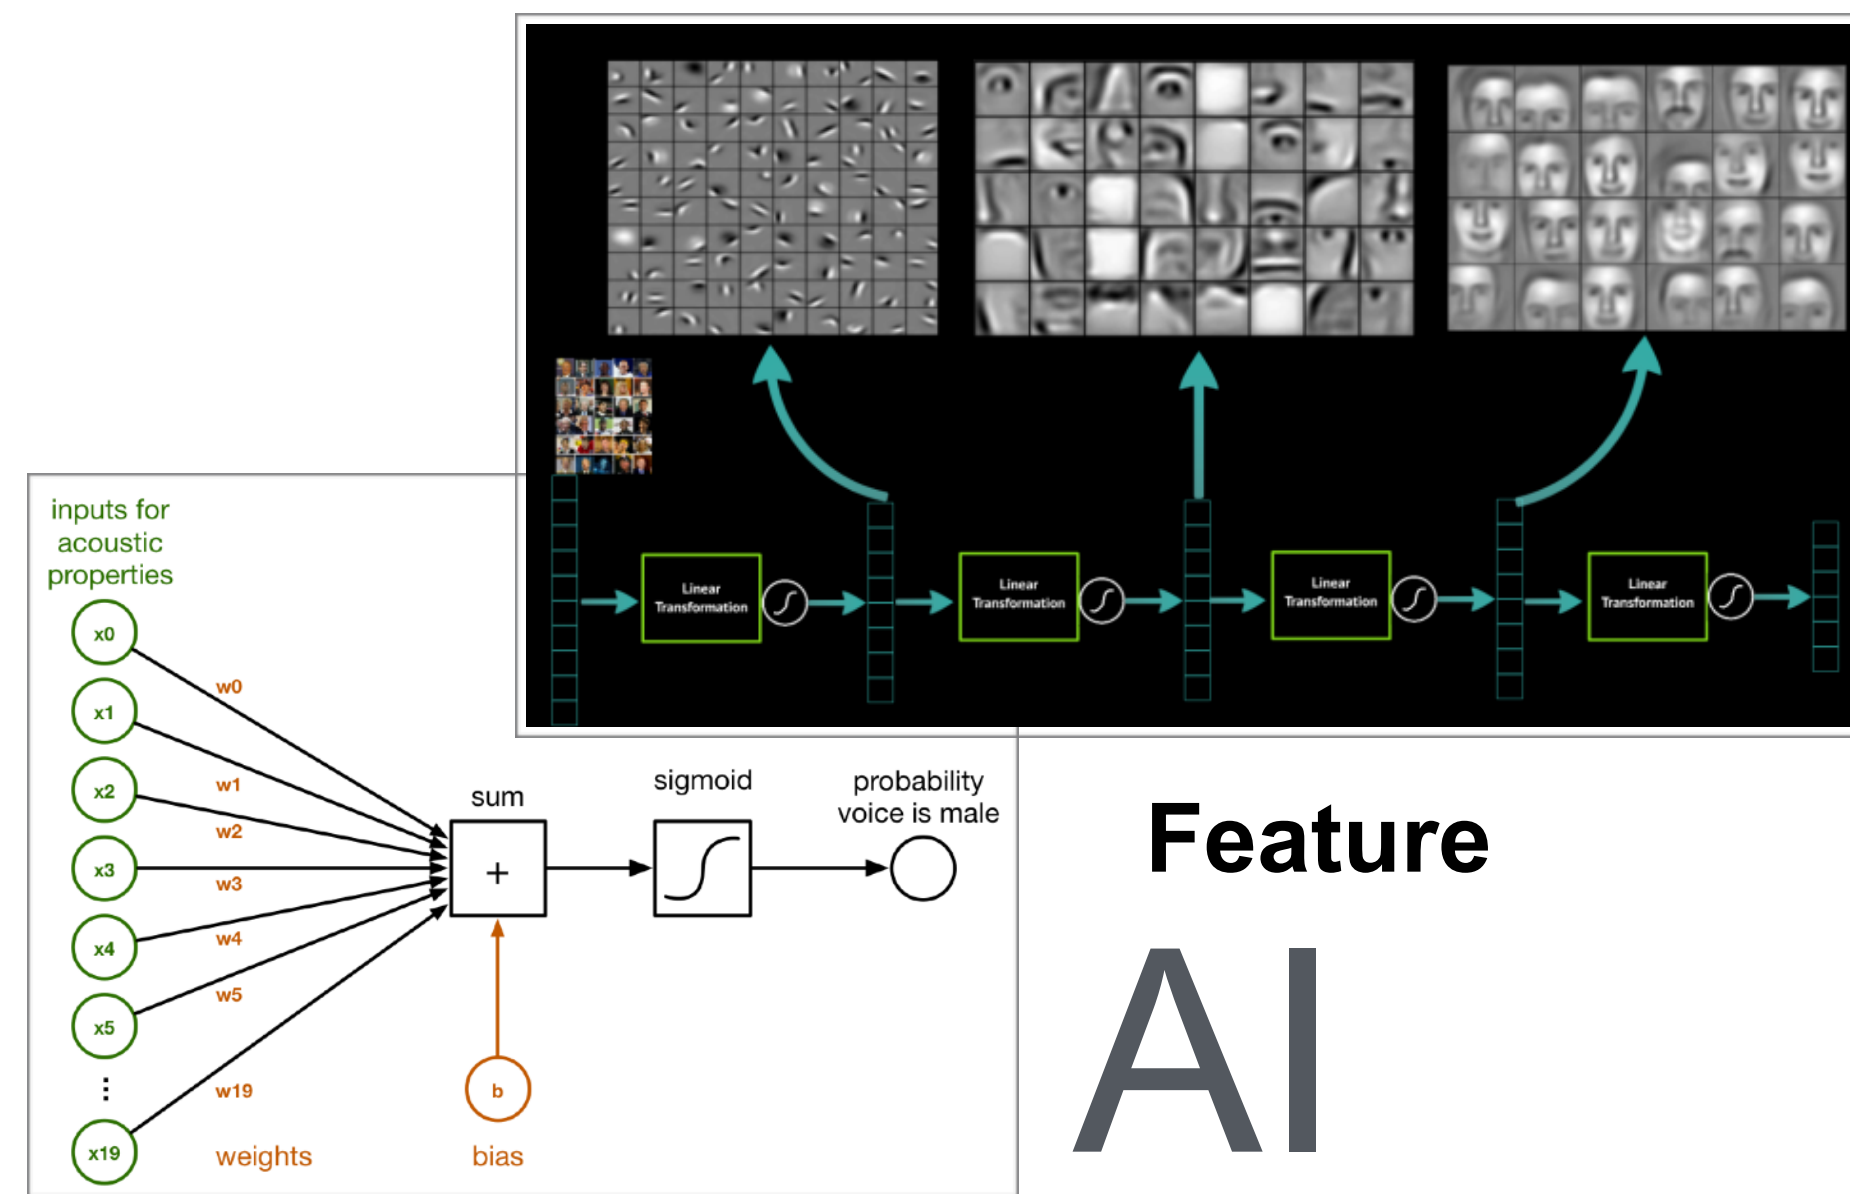
Max Meldrum, Klas Segeljakt, Lars Kroll
Paris Carbone, Christian Schulte, Seif Haridi

Outline


- Project Introduction
- Arc IR and Compilation Pipeline
- Demo (Frontend, IR, CodeGen, Execution)
- Conclusions and Future Work

Motivation

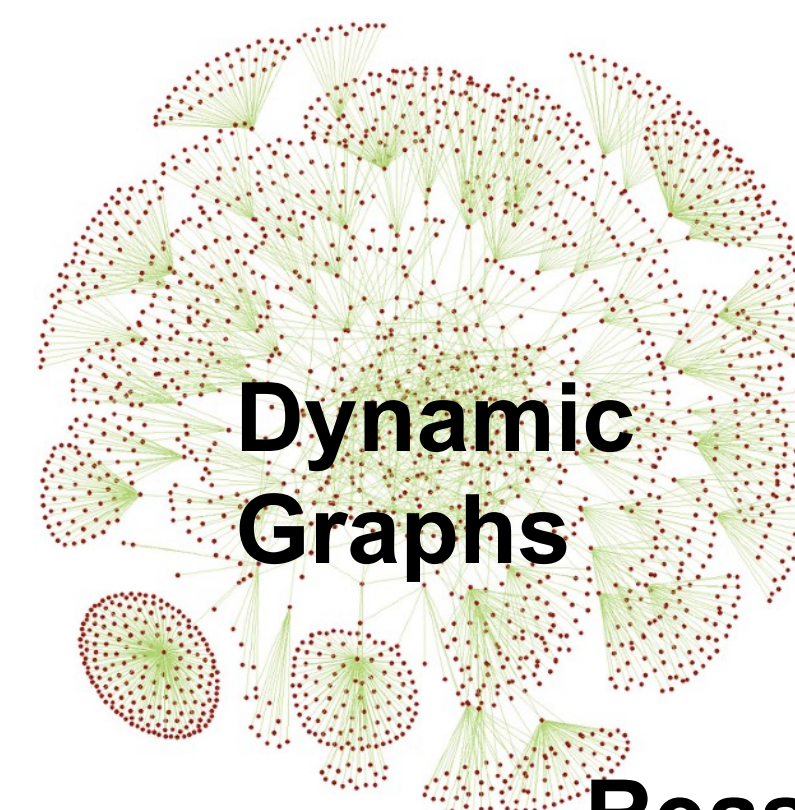
- Many Frameworks/Frontends for different needs
- (ML Training & Serving, SQL, Streams, Tensors, Graphs)



Feature
AI



Relational Data
Streams



Dynamic
Graphs

ML

Tensor

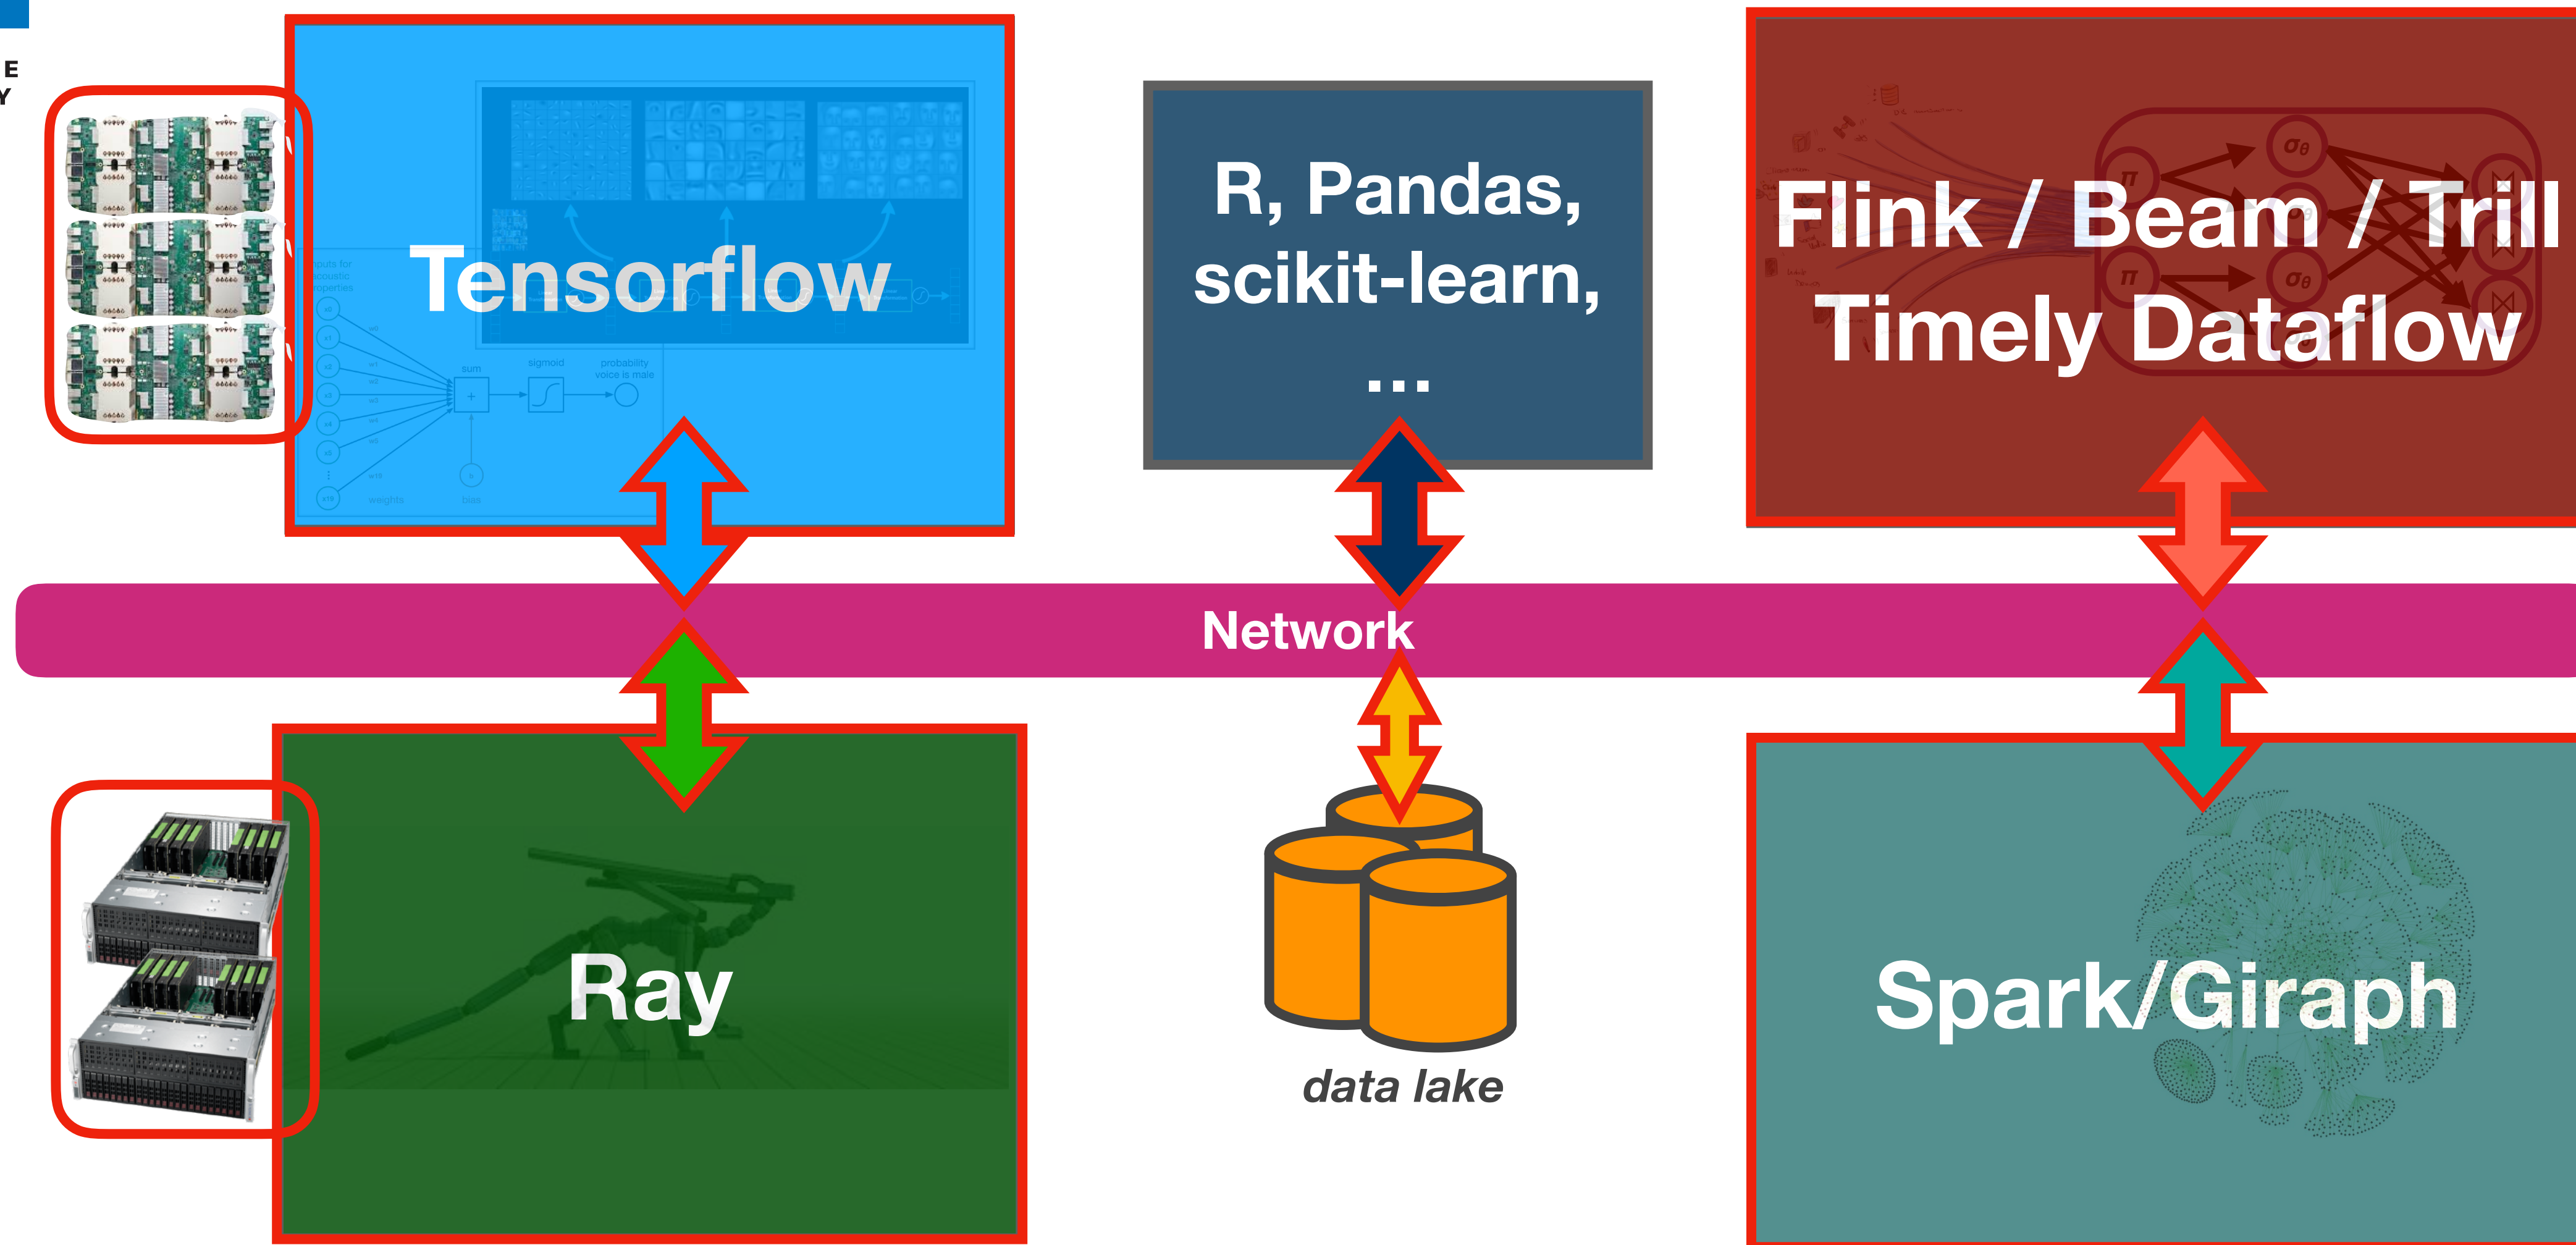


Simulation

RL

Reasoning

Silos in Data Pipelines




- Impedance Mismatch (e.g., types, guarantees, state etc.)
- Excessive IO/ Data Movement of intermediate results
- Isolated HW Execution - No cross-framework optimisation

The Arcon System

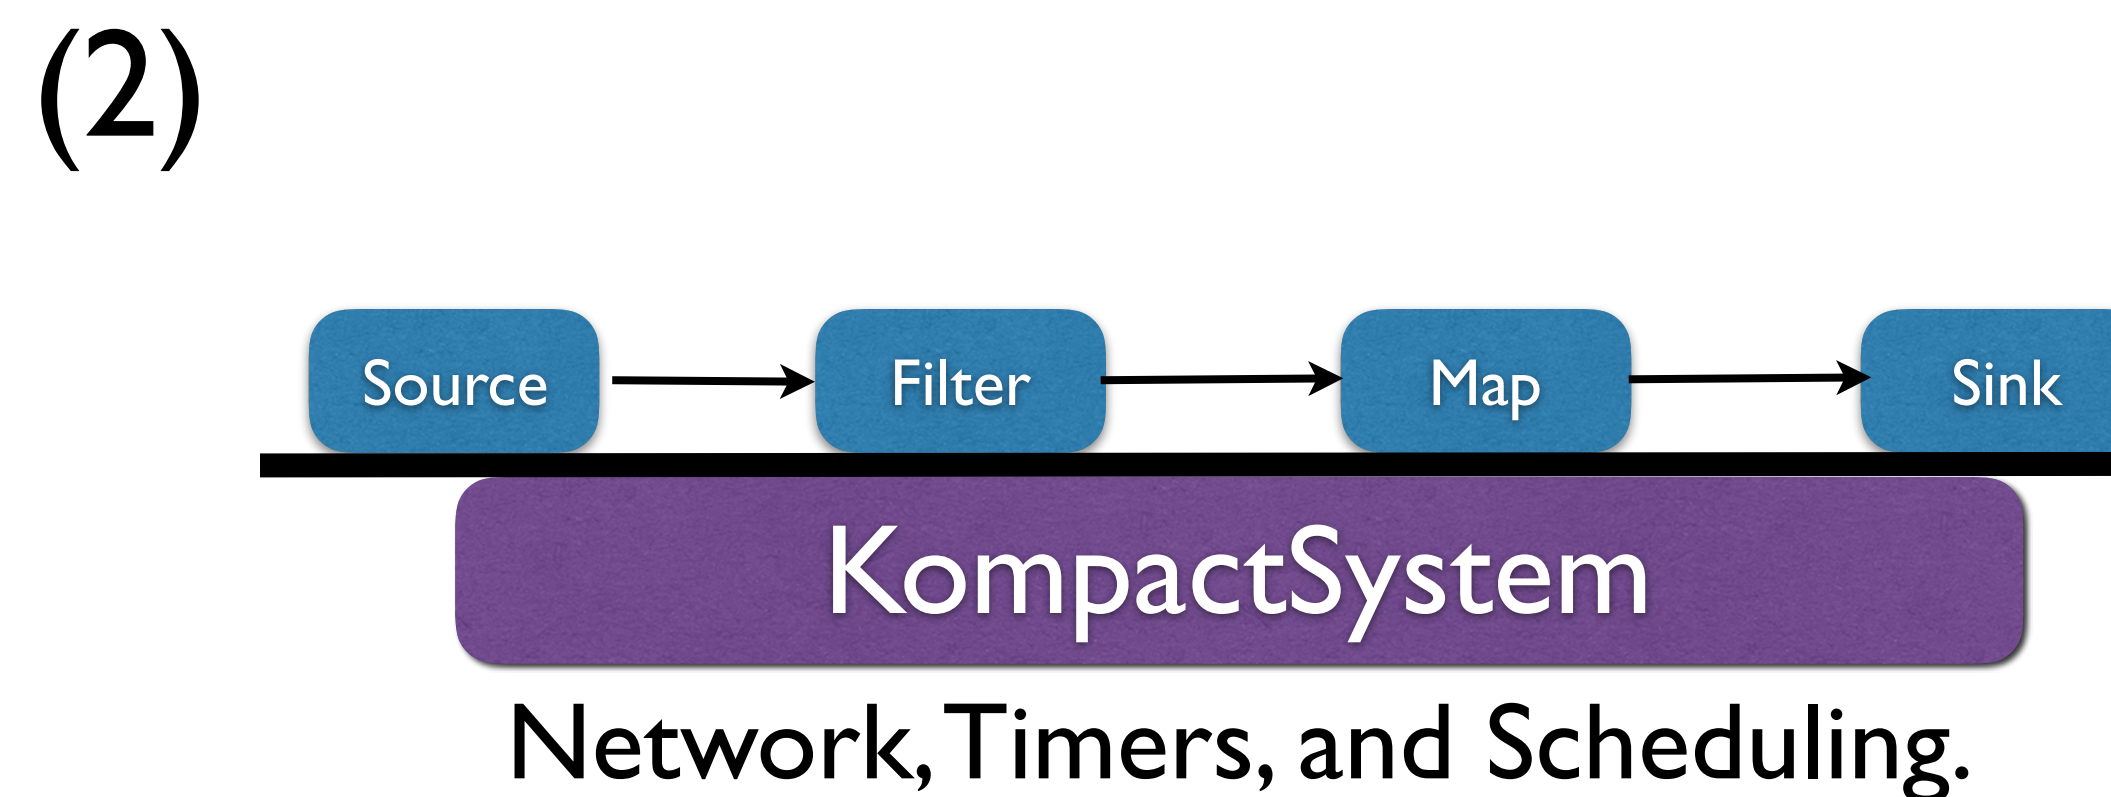
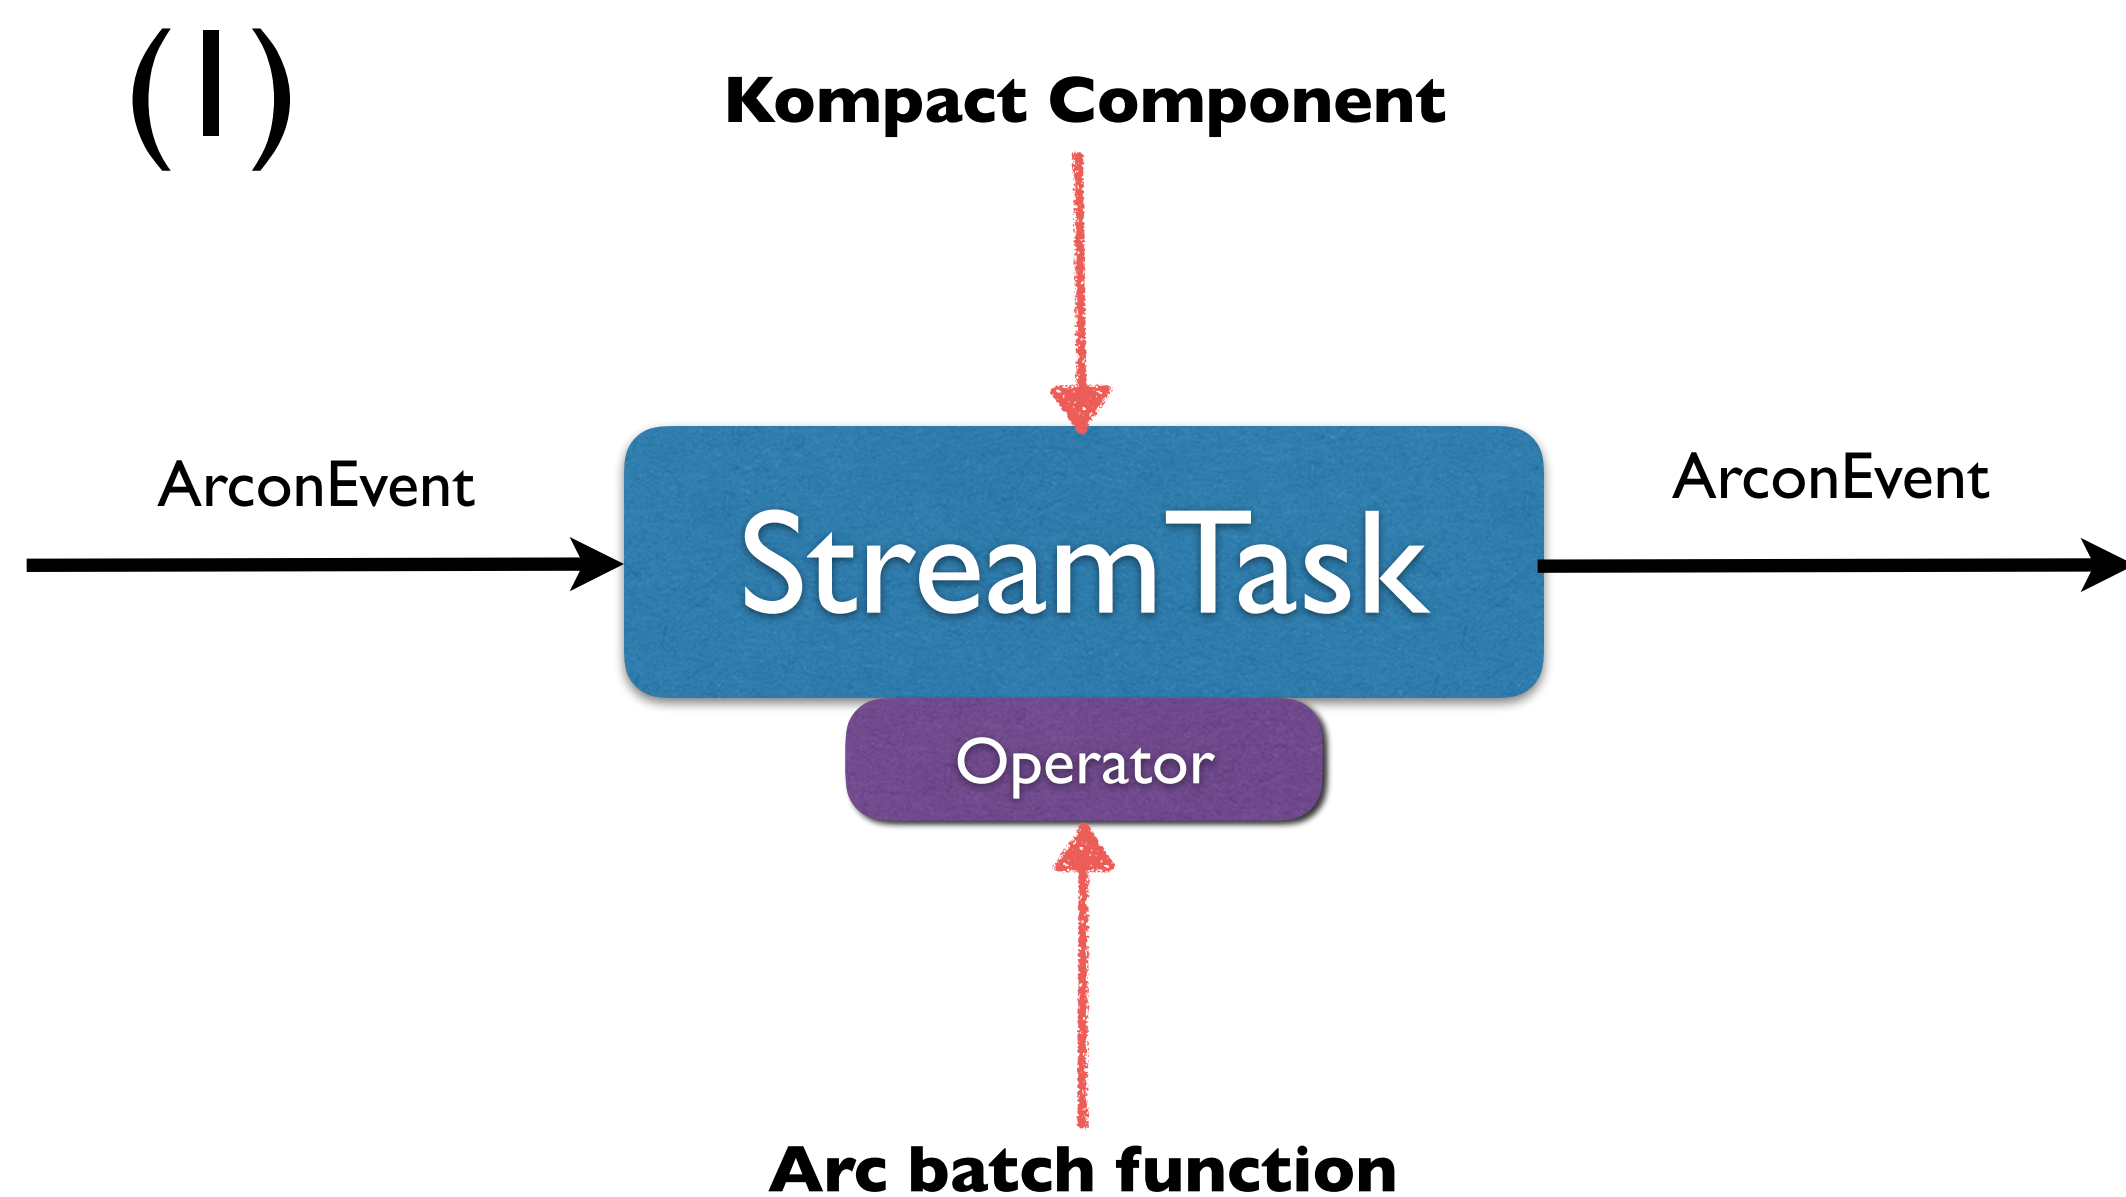
Arc IR Compiler

Arcon Runtime




Arcon Runtime

- Rust-based distributed dataflow engine
- Building Blocks:
 - Kompact: Hybrid Concurrent Component + Actor Model
 - Arc batch backend



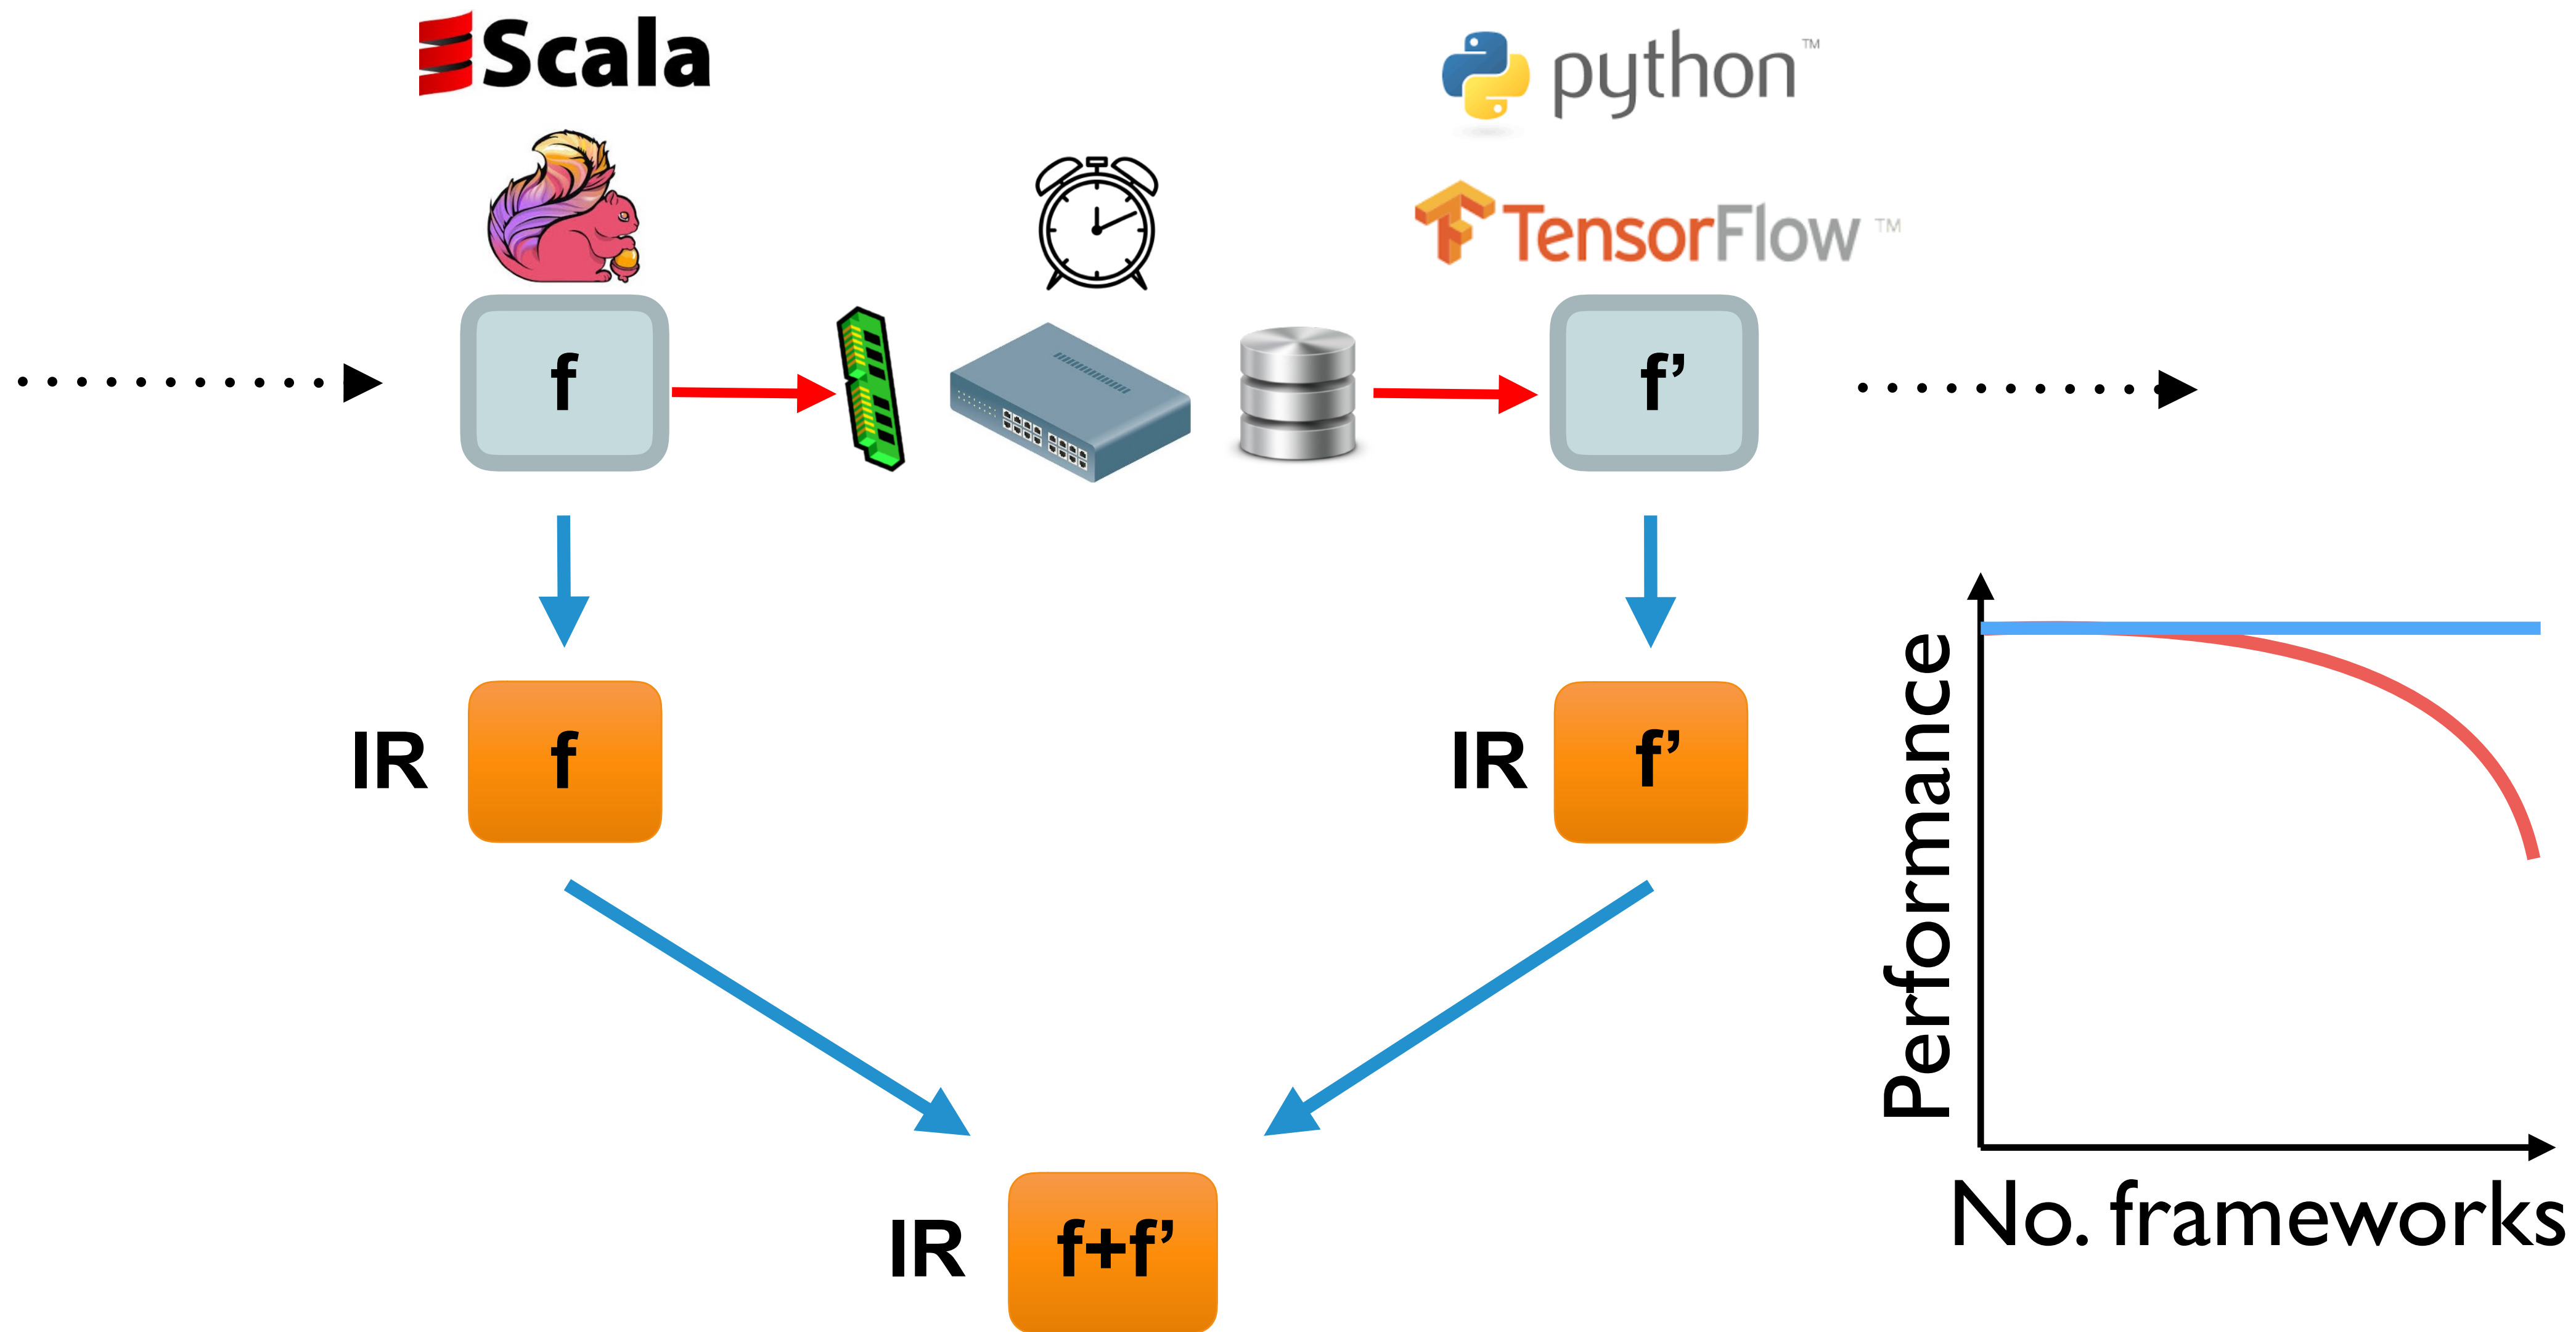
The Arcon System

Arc IR Compiler



Arcon Runtime


Intuition



Arc IR

- Streams
- Tables/Data Frames
- Vectors
- Tensors
-

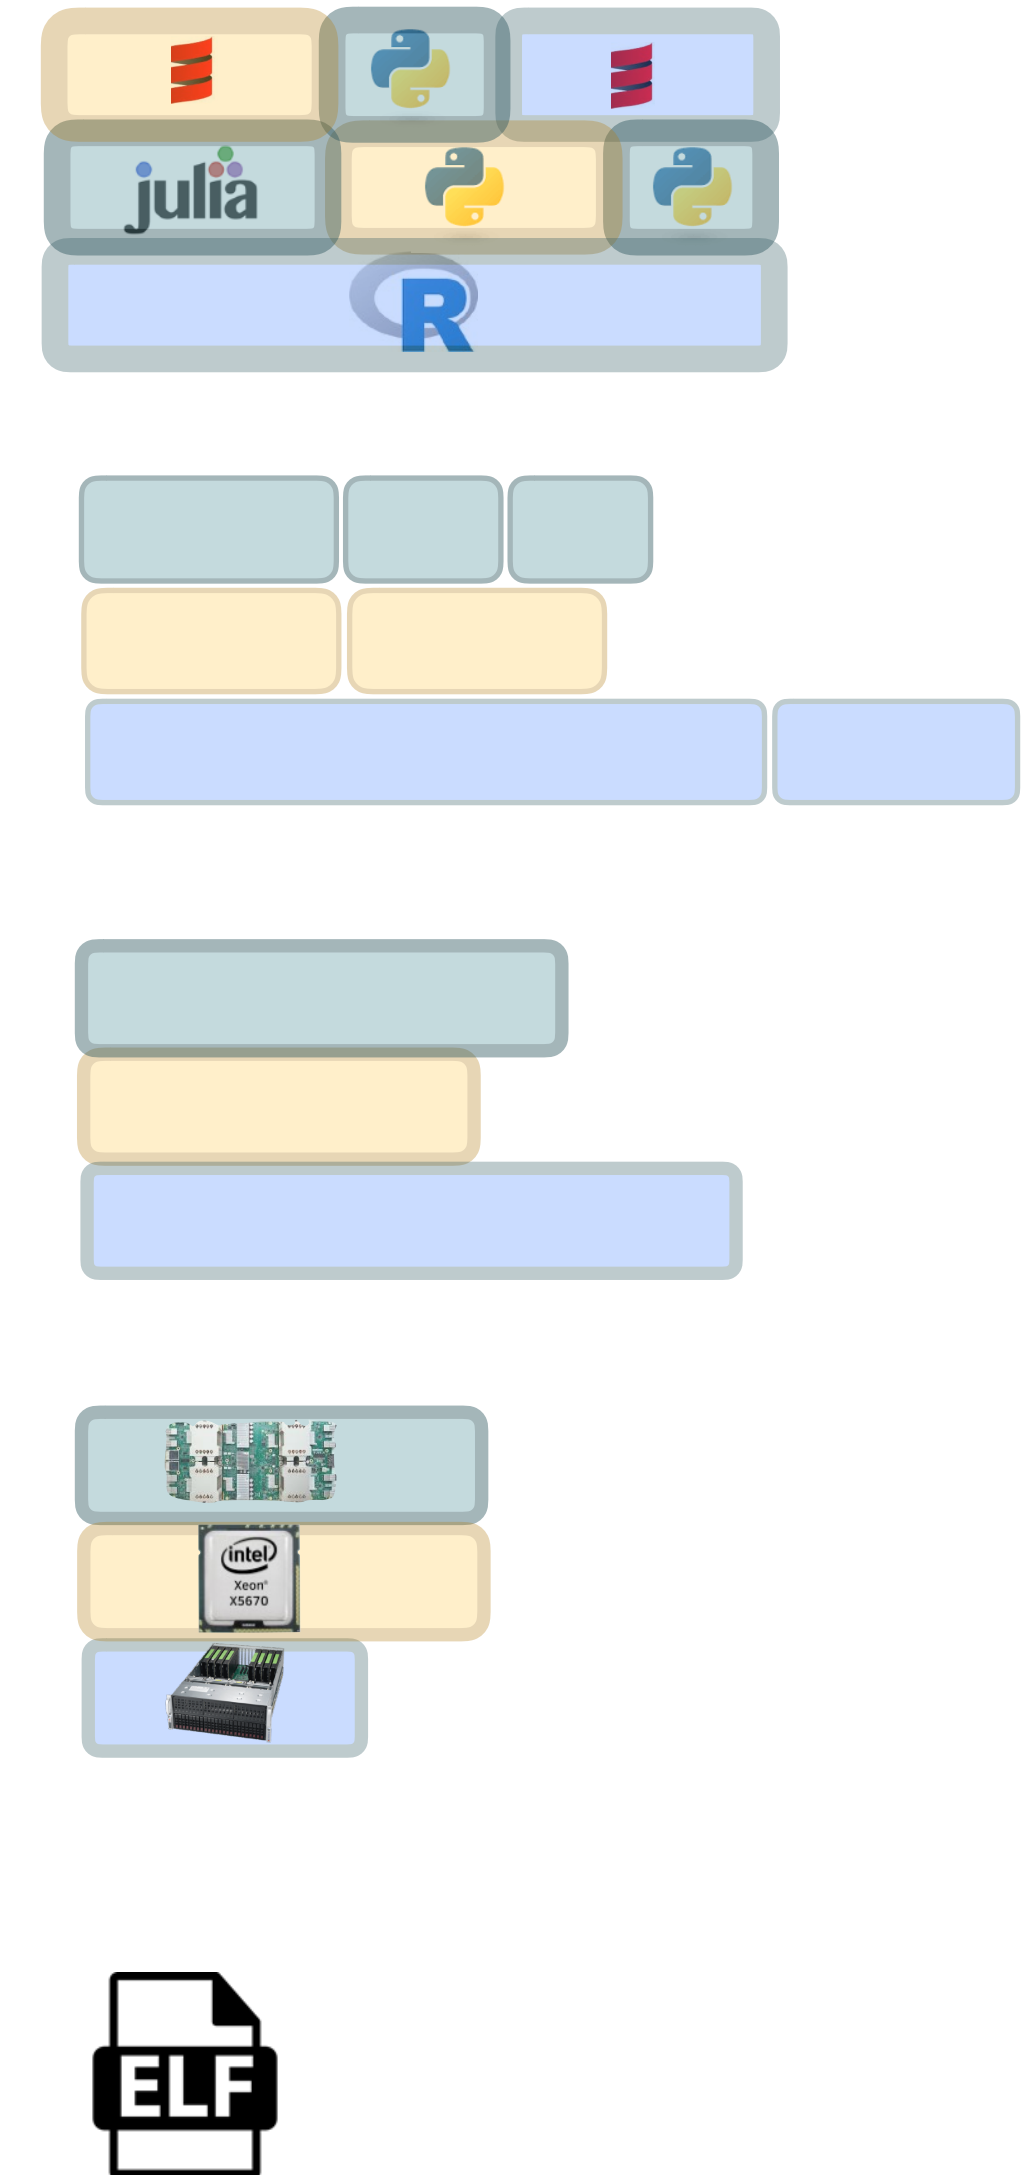
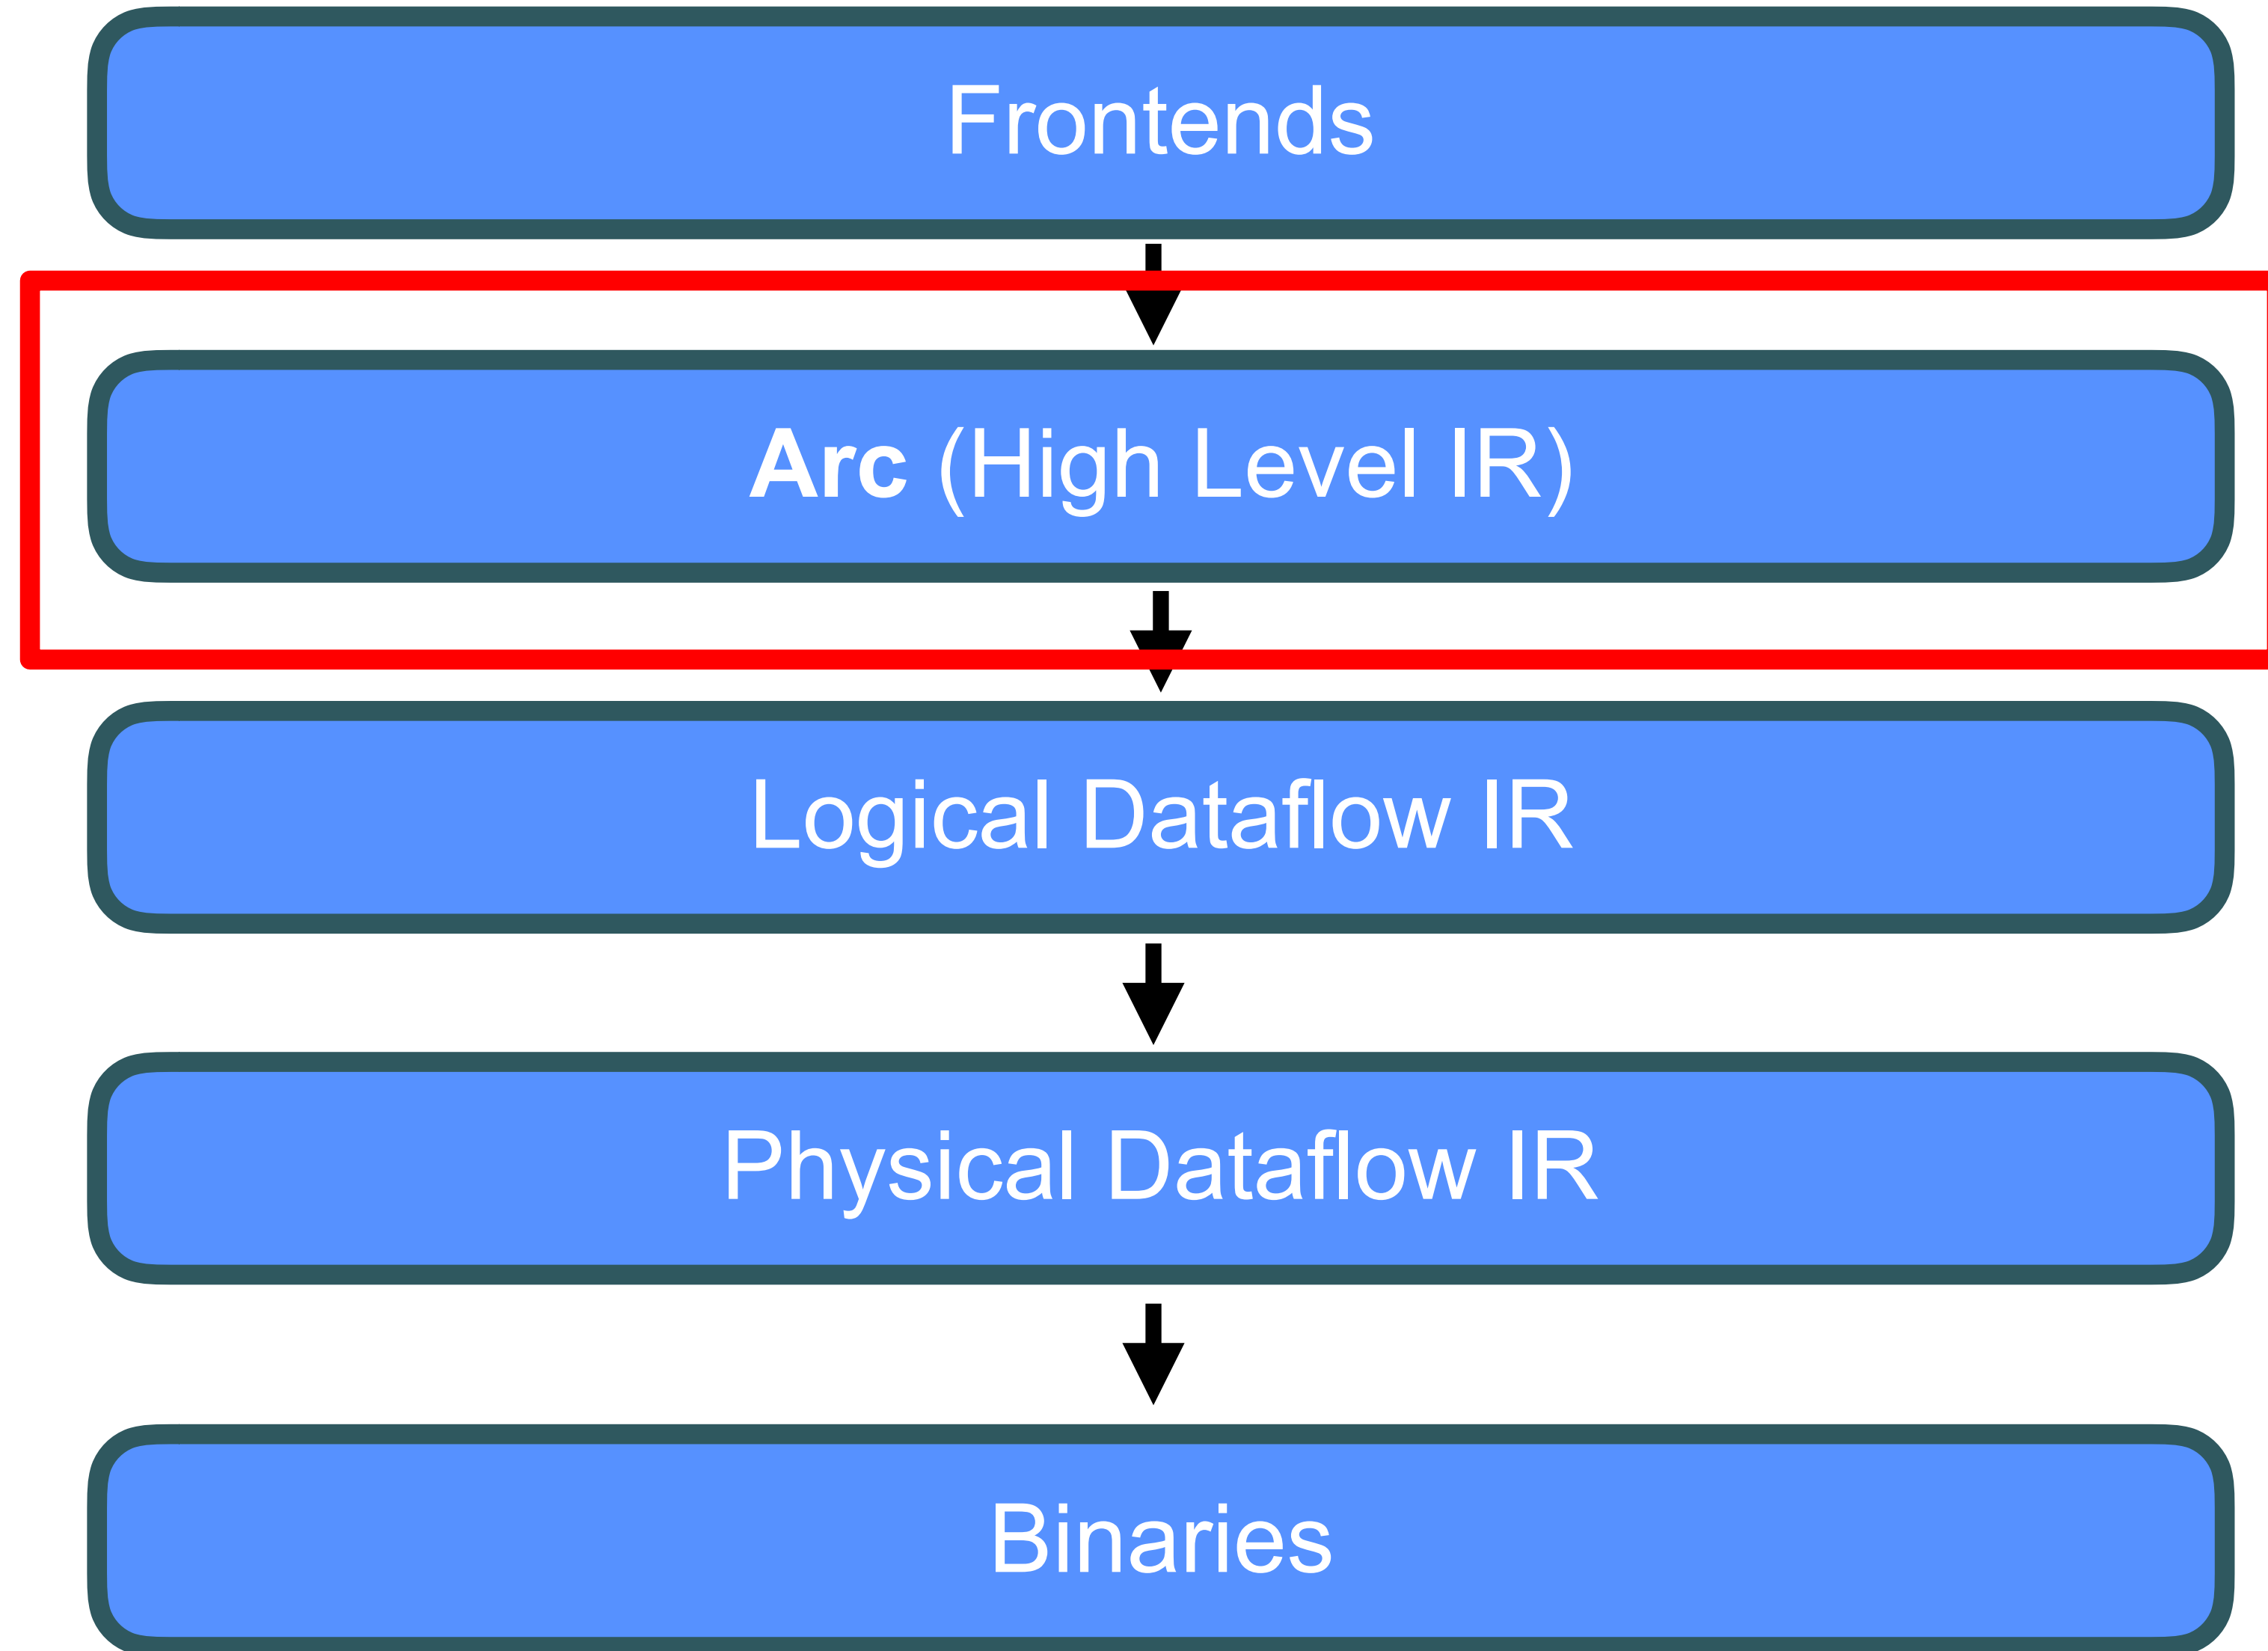
- Multicore CPUs
- GPUs
- TPUs
- FPGAs
-



- **Support both batch and streaming abstractions**
- **Sources/Sinks/Operators**
- **User-defined Windows**
- **Out of Order Processing**

Arcon Compiler Pipeline

Arcon



How does Arc work?

The **Weld IR*** is a subset of **Arc** that supports batch computations

- A restrictive language for describing data transformations
- **Values:** **Read-only** data types (e.g., `vec[T]`, `i8..i64`, `bool`, ...)
- **Builders:** **Write-only** data types (e.g., `appender[T]`)
- Calling `result` on a **builder** returns the associated **value** type

The **Arc IR**** supports both stream and batch computations

- Stream **sources** are **read-only** => **values** (i.e. `stream[T]`)
- Stream **sinks** are **write-only** => **builders** (i.e. `streamappender[T]`)
- Calling `result` on a **sink** returns a **source** and creates a **channel** between them

* Palkar, Shoumik, et al. "Weld: A common runtime for high performance data analytics." *Conference on Innovative Data Systems Research (CIDR)*. 2017.

** Kroll, Lars, et al. "Arc: an IR for batch and stream programming". In *Proceedings of the 17th ACM SIGPLAN International Symposium on Database Programming Languages*

Example: Normalisation



Normalise by dividing each element by the average

e.g. for window `[4, 2, 2, 8]`

... the average is `4`

... the output is `[1, 0.5, 0.5, 2]`

Frontend code

```
import arc_beam as beam
import arc_beam.transforms.window as window
import arc_beam.transforms.combiners as combiners
import baloo as pandas

def normalise(data):
    series = pandas.Series(data)
    avg = series.sum() / series.count()
    return series / avg

p = beam.Pipeline()

(p
 | beam.io.ReadFromText(path='input.txt').with_output_types(int)
 | beam.WindowInto(window.FixedWindows(size=5))
 | beam.CombineGlobally(normalise)
 | combiners.ToList()
 | beam.io.WriteToText(path='output.txt'))

p.run()
```

Generated Arc IR code

```

|source_0: stream[i64], sink_0: streamappender[?]|
  let operator_0 = result(for(source_0, windower[unit, appender[?], ?, vec[?]](
    |ts, windows, state| { [ts/5000L], () },
    |wm, windows, state| { result(filter(windows, |ts| ts < wm), () },
    |agg| result(agg)
  ),
  |sb, se| merge(sb, se)
));
for(operator_0, sink_0, |sb, se| merge(sb,
  let obj102 = (se);
  let obj105 = (result(
    for(obj102, merger[i64, +],
      |b: merger[i64, +], i: i64, e: i64|
        merge(b, e)
    ));
  let obj106 = (len(obj102));
  let obj107 = (obj105 / obj106);
  let obj108 = (result(
    for(obj102, appender[i64],
      |b: appender[i64], i: i64, e: i64|
        merge(b, e / obj107)
    ));
  obj108
));

```

Stream code (Beam)

Global tumbling window

Sum

Batch code (Pandas)


Count

Average

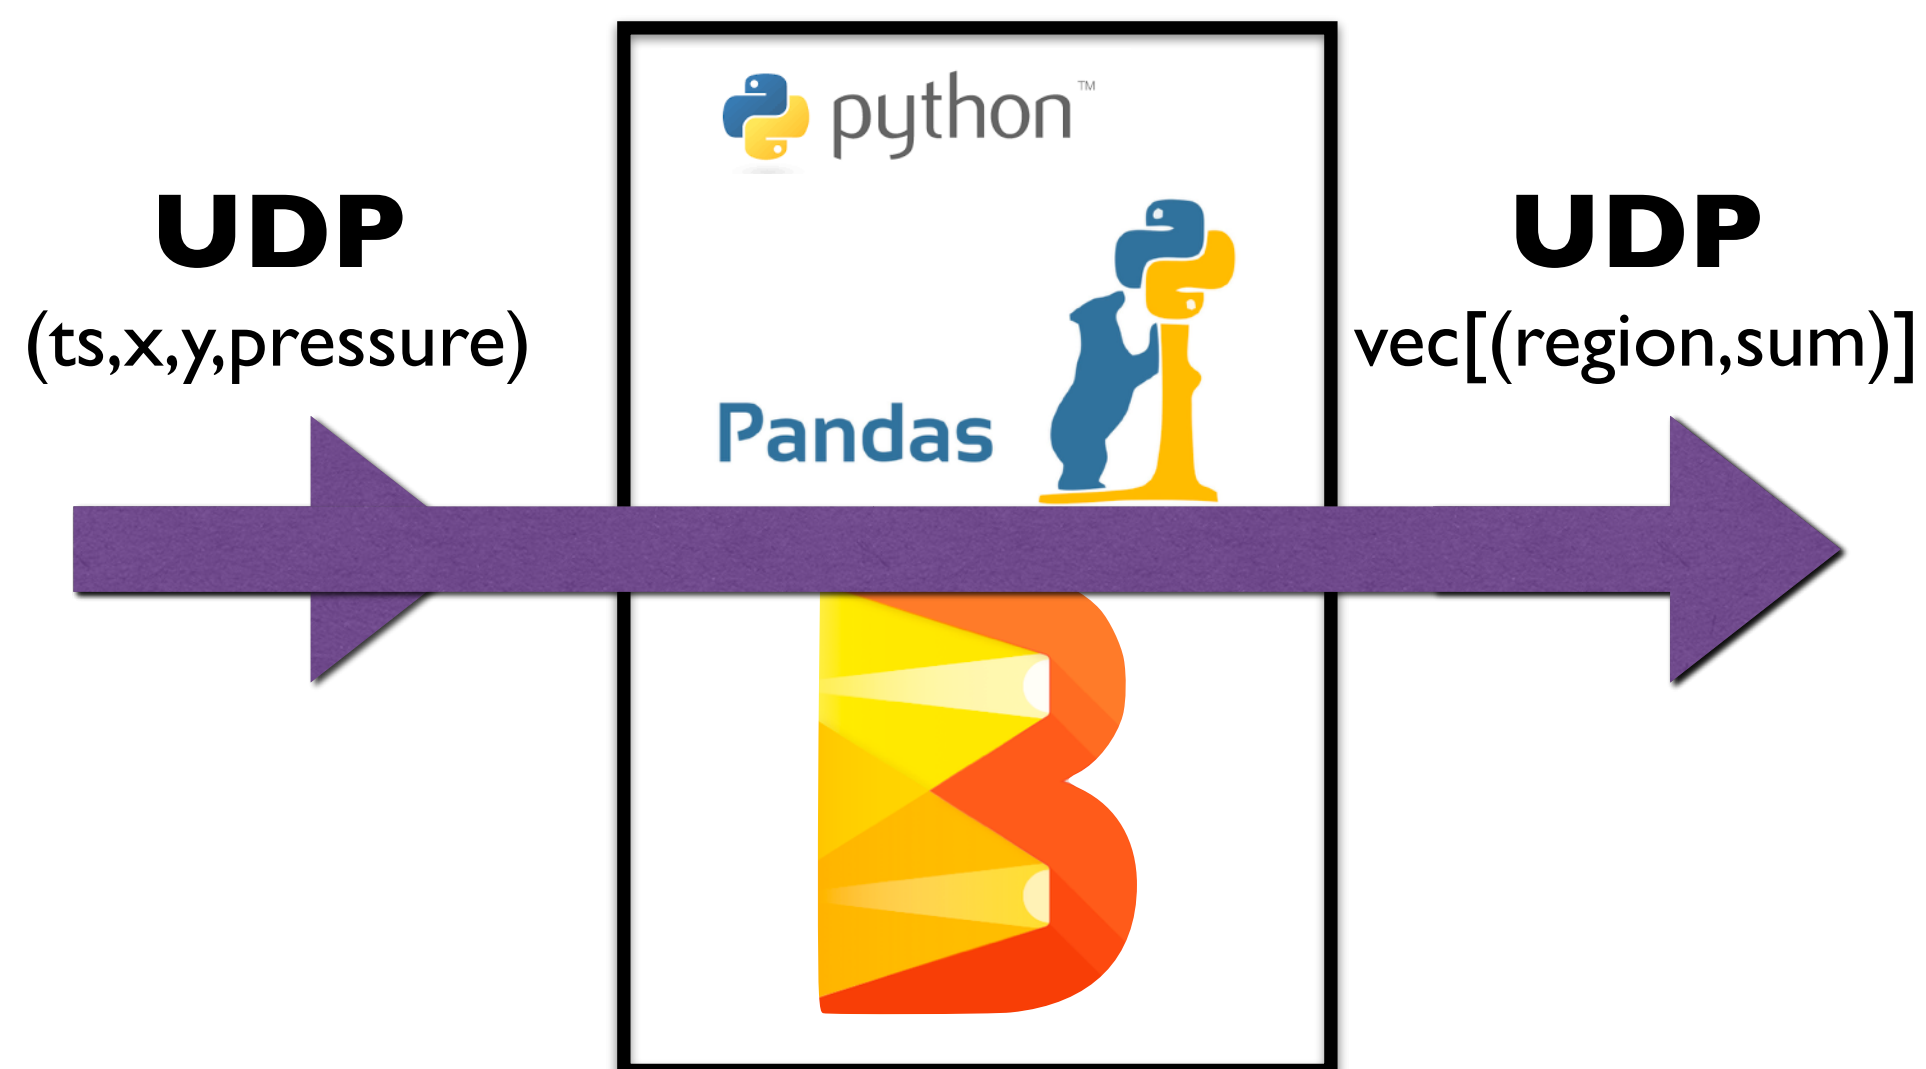
Normalisation

Demo: Touchpad Heatmap

The Input

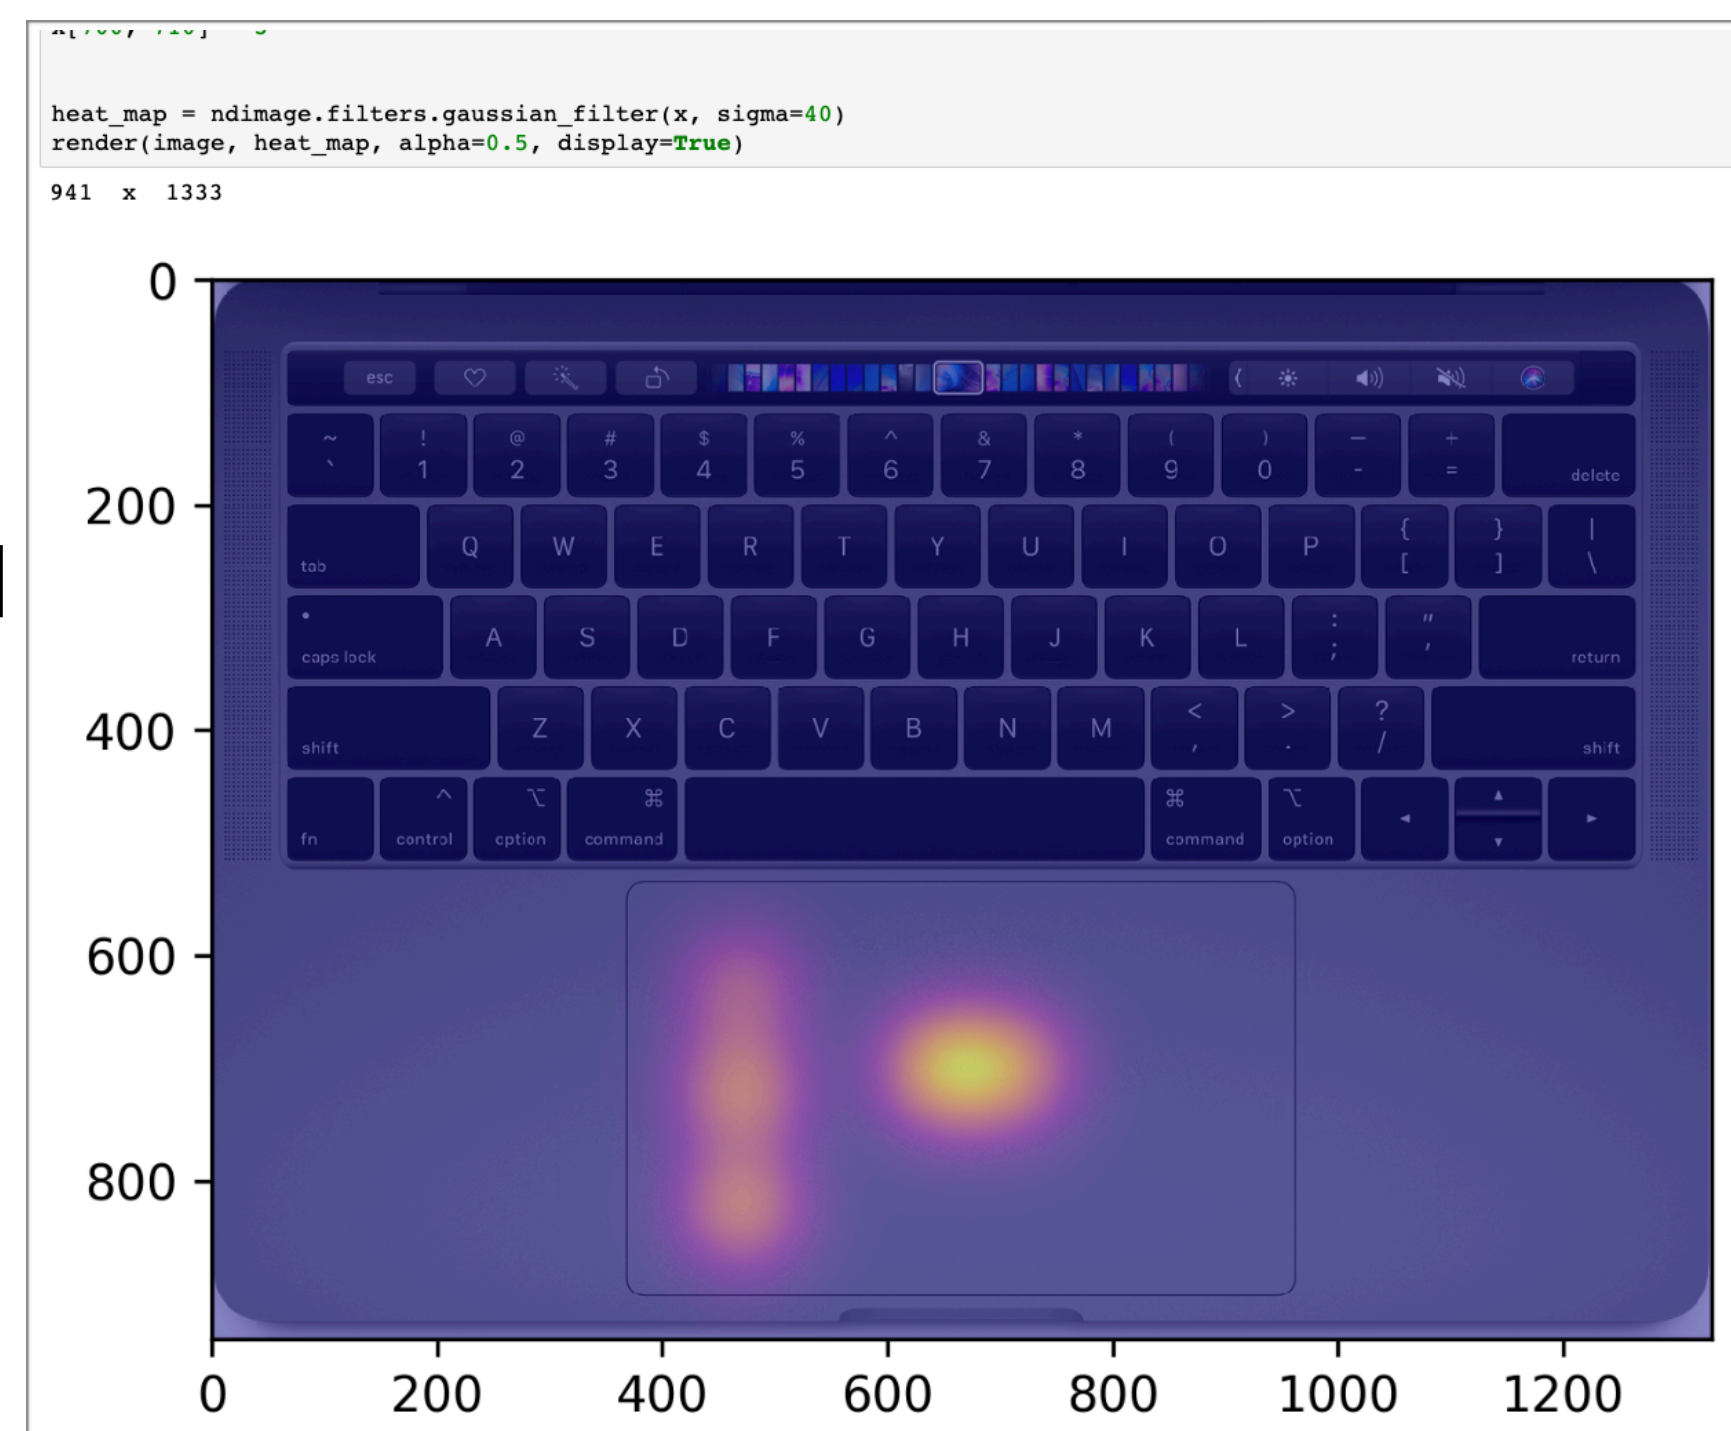


The Pipeline



- Touchpad is a grid of 5x3 regions
- Event-time window (6 seconds)
- Sum up pressure by region

The Output



Touchpad Heatmap Demo

MouseTracker

- CocoaAsyncSocket.framework
 - Headers
 - CocoaAsyncSocket.h
 - GCDAsyncSocket.h
 - GCDAsyncUdpSocket.h
- MouseTracker
 - AppDelegate.swift
 - Assets.xcassets
 - MainMenu.xib
 - Info.plist
 - MouseTracker.entitlements
 - NetSender.swift
- Products
 - MouseTracker.app

```
1 //
2 // AppDelegate.swift
3 // MouseTracker
4 //
5 // Created by Lars Kroll on 2019-08-13.
6 // Copyright © 2019 Lars Kroll. All rights reserved.
7 //
8
9 import Cocoa
10
11 @NSApplicationMain
12 class AppDelegate: NSObject, NSApplicationDelegate {
13
14     @IBOutlet weak var window: NSWindow!
15
16     var udpSender: NetSender?;
17
18     func applicationDidFinishLaunching(_ aNotification: Notification) {
19         // Insert code here to initialize your application
20         print("setting up!");
21         var frame = self.window?.frame
22         frame?.size = NSSize(width: 505, height:323)
23         self.window?.setFrame(frame!, display: true)
24         udpSender = NetSender();
25     }
26 }
```

```
setting up!
Data sent: Test!
Did send data!
594886868,11.1484375,315.83984375,0.0

Data sent: 594886868,11.1484375,315.83984375,0.0

Did send data!
```

type: Default - Swift Source

Location: Relative to Group

AppDelegate.swift

Full Path: /Users/carbone/workspace/mactouchsender/MouseTracker/AppDelegate.swift

On Demand Resource Tags

Only resources are taggable

Target Membership

- MouseTracker

Text Settings

Text Encoding: No Explicit Encoding

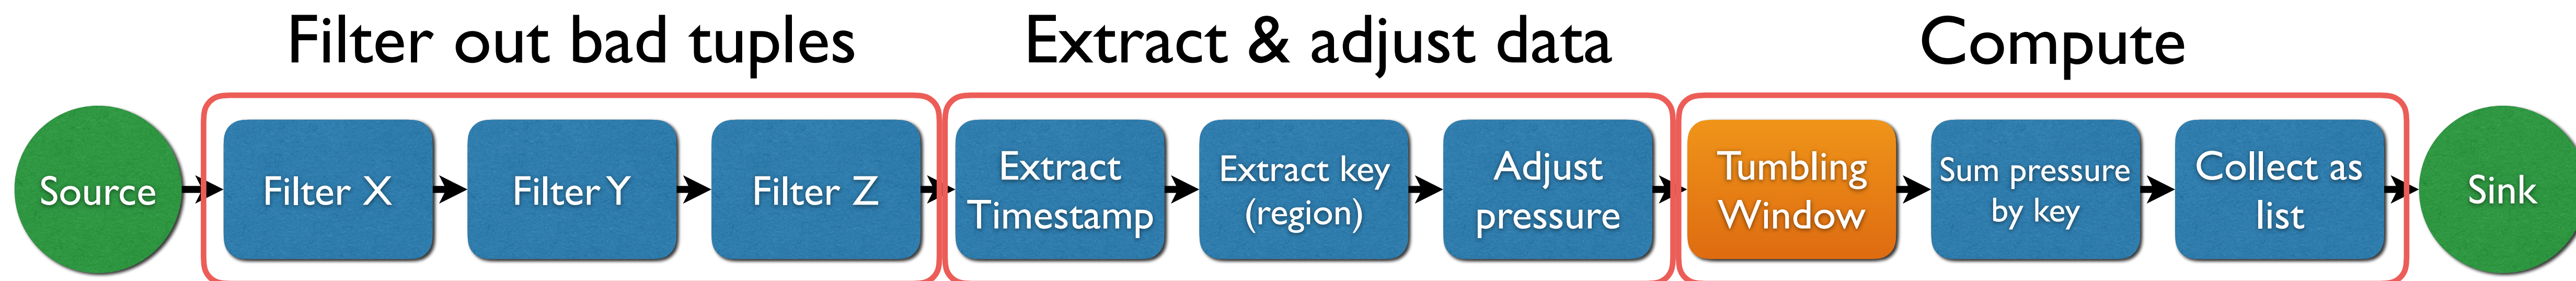
Line Endings: No Explicit Line Endings

Indent Using: Spaces

Widths: Tab: 4, Indent: 4

Wrap lines

Dataflow Graph



Touchpad Beam Code

```
p = beam.Pipeline()

(p
| beam.io.ReadFromSocket(addr=touchpad, coder=beam.coders.CSVCoder())
  .with_output_types(Tuple[ts, x, y, z])
| 'preprocess'
  >> beam.Filter(lambda e: (e[1] >= 0) & (e[1] <= width))
  beam.Filter(lambda e: (e[2] >= 0) & (e[2] <= height))
  beam.Filter(lambda e: (e[3] >= 0) & (e[3] <= max_pressure))
| 'extract timestamp'
  >> beam.Map(lambda e: window.TimestampedValue(value=e[1:4], timestamp=e[0]))
| 'extract key'
  >> beam.Map(lambda e: ((e[0] / grid_width, e[1] / grid_height), e[2]))
| 'add to pressure'
  >> beam.Map(lambda e: (e[0], e[1] + epsilon))
| 'create tumbling window'
  >> beam.WindowInto(window.FixedWindows(size=window_length))
| 'sum up pressures'
  >> beam.CombinePerKey(lambda e: pandas.Series(e).sum())
| 'collect window as list'
  >> combiners.ToList()
| beam.io.WriteToSocket(addr=display, coder=beam.coders.CSVCoder()))

p.run()
```

Code-generation steps

```

1 from typing import Tuple
2 import arc_beam as beam
3 import arc_beam.transforms.window as wind
4 import arc_beam.transforms.combiners as c
5 import baloo as pandas
6
7 width, height = 500, 300 # 500x300 point
8 num_x_grids, num_y_grids = 5, 3 # 5x3 gr
9 grid_width, grid_height = int(width / num
10 epsilon = 0.1
11 touchpad, display = '127.0.0.1:8000', '12
12 ts, x, y, z = int, float, float, float
13 max_pressure = 1.0
14 window_length = 6
15 p = beam.Pipeline()
16 (p
17 | beam.io.ReadFromSocket(addr=touchpad,
18 | .with_output_types(Tuple[ts, x,
19 | 'preprocess'
20 | >> beam.Filter(lambda e: (e[1] >= 0)
21 | beam.Filter(lambda e: (e[2] >= 0)
22 | beam.Filter(lambda e: (e[3] >= 0)
23 | 'extract timestamp'
24 | >> beam.Map(lambda e: window.Timestam
25 | 'extract key'
26 | >> beam.Map(lambda e: (((e[0] / grid_
27 | 'add to pressure'
28 | >> beam.Map(lambda e: (e[0], e[1] + e
29 | 'create tumbling window'
30 | >> beam.WindowInto(window.FixedWindow
31 | 'sum up pressures'
32 | >> beam.CombinePerKey(lambda e: panda
33 | 'collect window as list'

```

```

1 {
2   "nodes": [
3     {
4       "id": "source_0",
5       "kind": {
6         "Source": {
7           "format": "CSV",
8           "kind": {
9             "Socket": {
10              "addr": "127.0.0.1:8000"
11            }
12          }
13        }
14      },
15    },
16    {
17      "id": "sink_0",
18      "kind": {
19        "Sink": {
20          "format": "CSV",
21          "kind": {
22            "Socket": {
23              "addr": "127.0.0.1:9000"
24            }
25          }
26        }
27      },
28    }
29  ],
30  "timestamp_extractor": 0,
31  "arc_code": "|source_0: stream[{i64,
f64,f64,f64}], sink_0: streamappender[?
]|\\n# preprocess\\nlet operator_1 =

```

```

1 {
2   "id": "dfg_0",
3   "target": "x86-64-unknown-linux-gnu",
4   "nodes": [
5     {
6       "id": "source_0",
7       "parallelism": 1,
8       "kind": {
9         "Source": {
10          "source_type": {
11            "Struct": {
12              "id": "struct_source_0_tas
13              "field_tys": [
14                {
15                  "Scalar": "I64"
16                },
17                {
18                  "Scalar": "F64"
19                },
20                {
21                  "Scalar": "F64"
22                },
23                {
24                  "Scalar": "F64"
25                }
26              ]
27            }
28          },
29          "format": "CSV",
30          "channel_strategy": "Forward",
31          "successors": [
32            {
33              "Local": {

```


```

1 #![allow(dead_code)]
2 #![allow(missing_docs)]
3 #![allow(unused_imports)]
4 #![allow(unused_results)]
5 #![allow(bare_trait_objects)]
6
7 extern crate arcon;
8 use arcon::macros::*;
9 use arcon::prelude::*;
10 #![macro_use]
11 extern crate serde;
12 #![arcon_decoder(,)]
13 #![arcon]
14 pub struct struct_source_0_task_0_task_1
15     f0: i64,
16     f1: f64,
17     f2: f64,
18     f3: f64,
19 }
20 impl ::std::hash::Hash for struct_source
21     fn hash<H: ::std::hash::Hasher>(&sel
22 }
23 #![arcon_decoder(,)]
24 #![arcon]
25 pub struct struct_task_4_task_5_0 {
26     f0: struct_task_4_task_5_1,
27     f1: f64,
28 }
29 impl ::std::hash::Hash for struct_task_4
30     fn hash<H: ::std::hash::Hasher>(&sel
31 }
32 #![arcon_decoder(,)]
33 #![arcon]


```



→ Arc + Metadata



Dataflow



Dataflow Graph



Initial pipeline



Fused pipeline

Operators are fused by inlining at the instruction level

Arc Optimisations

- Arc supports **both** compiler and dataflow optimisations
 - **Compiler**: Loop unrolling, partial evaluation,
 - **Dataflow**: Operator fusion, fission, reordering, specialisation, ...
- Find optimal dataflow graph through constraint model (future work)

Conclusions & Future Work

- Arc enables cross-compiling and optimising programs from diverse libraries.
- Next steps:
 - Wider support for more frontends, Tensorflow, Flink, etc.
 - Common Pipeline DSL
 - State management for dynamic task graphs
 - Runtime Optimiser and Reconfiguration



ROYAL INSTITUTE
OF TECHNOLOGY



Extra slides

Optimisation example

