

Portals: A Showcase of Multi-Dataflow Stateful Serverless

^{*†}Jonas Spenger, [†]Chengyang Huang, [†]Philipp Haller, ^{*†}Paris Carbone

^{*}RISE Research Institutes of Sweden, Stockholm, Sweden; [†]KTH Royal Institute of Technology, Stockholm, Sweden

<https://www.portals-project.org/vldb2023demo/>

demo info



Context

- **Serverless functions** have made it easy to write and deploy distributed applications on fully managed runtimes.
- Recent developments have been on including **state management** and **compositional patterns**, enabling a wider range of applications; yet, finding the right abstractions and implementation methods remains an open problem.
- Our work is on a decentralized programming framework for **stateful serverless applications in the cloud-edge continuum**.


Summary

This demonstration presents **Portals** [1, 2], a programming framework for **stateful serverless applications**, with the following highlights:

- **Multi-dataflow applications.** Dynamically composing multiple dataflow pipelines together via **atomic streams**.
- **Inter-dataflow services / Portal services.** Exposing dataflow pipelines (operators, state) as inter-dataflow services.
- **Decentralized cloud and local execution.** Decentralized API and runtime, with end-to-end processing guarantees.

Key Features

// Programming framework for stateful serverless // **Atomic streams & exactly-once processing guarantees (WIP)** // Dynamic topology // Decentralized execution // Flexible API //




Multi-dataflow composition


- Dataflows as microservices.
- Composition using **atomic streams**.

Powered by Atomic Streams

- Enforces the **exactly-once processing guarantees**.
- Provides interface for the atomic processing contract.




- Portals Vision -



Portal services


- A Portal exposes a dataflow as a service, implemented with task operators.
- Enables **request/reply communication**.



Decentralized, dynamic topology

- Applications spanning multiple deployments.
- **Topology may change over time.**
- A runtime for cloud and edge (WIP).


Scenario 1: Shopping Cart Pipeline



Demo overview

- **Four services:** cart; orders; inventory; analytics.
- Services launched dynamically, connected.
- Portal service exposes the inventory, analytics.
- Starting new services that use the analytics, on demand.


(1) The **inventory** exposed through a **Portal**, implemented as a task with `PerKeyState`.



(2) The **cart** consumes user requests, and **interacts with the inventory through a portal** to add/remove items to the cart.

```
val inventory = Portal(...)
Dataflows("inventory").source(...)
  .taskWithReplier(inventory){...}{
    ... if State.get() > 0 then
      ... Reply(GetReply(item, true))
  }
```


```
val inventoryRef = Registry.Portal(...)
Dataflows("cart").source(...)
  .taskWithRequester(inventoryRef):
    ... Request(inventoryRef)
    (GetItem(item)).onComplete:
      case Success(rep) =>
        State.update(item, ... + 1)
```



(3) The orders app consumes the checked out orders.
(4) The analytics service produces a **Top100 list of the orders, exposed through a portal service**.
(5) New services dyn. connect to the the analytics portal.

Scenario 2: SQL to Dataflow

- **SQL API based on portal services with state managed by a dataflow** [3].
- Supports multi-table SQL queries and transactions.
- The SQL engine uses Apache Calcite.



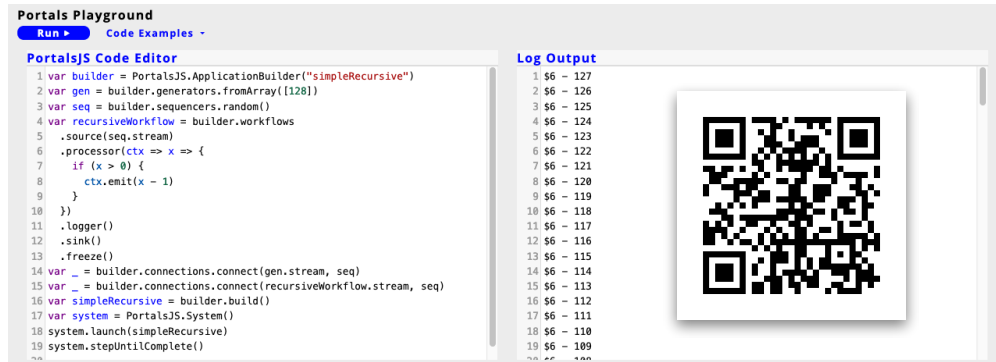
```
val table =
  TableWorkflow[Types.KV] ("KVTable")
```

Multiple querying apps connecting to the same shared KeyValue store, implemented as a library, leveraging the decentralized execution.

```
Dataflows[String, String](...)
  .source(...)
  .query(table)
```

Scenario 3: Playground

- The Portals Playground is a Javascript-based sandbox capable of **running Portals apps in the browser**. Check out the examples!
- The JS runtime can be used for web apps and edge devices.



<https://www.portals-project.org/playground/>

Ongoing / Future Work

- Impl. of the distributed serverless runtime.
- Dataflow optimiser exploiting the global view; improving the performance of cyclic dependencies across pipelines.
- Multi-dataflow ACID transactions.

References

- [1] Spenger et al., "**Portals: An extension of dataflow streaming for stateful serverless**," Onward'22.
- [2] <https://github.com/portals-project/portals>
- [3] Chengyang Huang, "Queryable Workflows: Extending Dataflow Streaming with Dynamic Request/Reply Communication." Dissertation, 2023.

Acknowledgements

We would like to thank other contributors to the Portals project: Maximilian Kurzwaski; Aleksey Veresov; Gabriele Morello; Siyao Liu. This work was partially funded by Digital Futures, the Swedish Foundation for Strategic Research under Grant No.: BD15-0006, as well as RISE AI.

