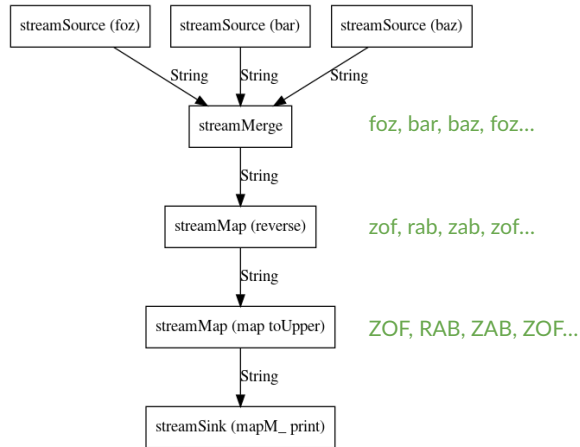


StrIoT brief introduction


Jonathan Dowland, Paul Watson
2020-10-30



User composes a stream-processing program in terms of “operators”: result is a (reverse) tree, with ≥ 1 source nodes and rooted in a single sink node.

StreamMap

Apply a function to transform each input event




We'll go through the 8 operators briefly and show how they may be represented in a Jackson network model

The operator descriptions are taken from the README at [striot.org](https://github.com/striot/striot)


StreamScan

Apply a function to transform each input event, taking into account past history



StreamFilter

Generates a stream containing only those events that meet a user-defined criteria




f is the selectivity of the operator

$\lambda_{out'}$ is the rate of the discarded events

To balance the network equations we've introduced a "phantom" output

StreamFilterAcc


Filter events, taking into account past results



f is the selectivity of the operator
 $\lambda_{out'}$ is the rate of the discarded events

StreamWindow chop


Create non-overlapping windows that are a fixed number (N) events in length



The modelling of streamWindow depends upon the user-supplied "window maker" parameter. Here we discuss four example window makers


StreamWindow chopTime

Create non-overlapping windows that are a fixed time (T) length




StreamWindow sliding

Create overlapping windows that are a fixed number of events in length




StreamWindow slidingTime

Create overlapping windows that are a fixed time length



StreamExpand

Generate an output event for each element of a list in the input event




The dual of windowing

"N" is the average number of elements in a list

Here to balance the network equations we've introduced a "phantom" input

StreamMerge


Combine ≥ 2 streams by merging their events



Here representing two inputs but could be more

StreamJoin

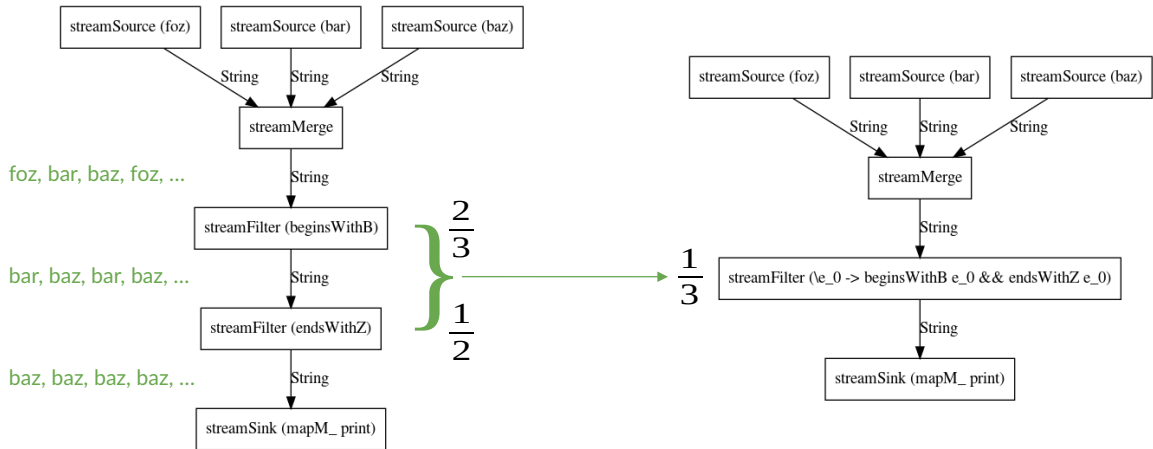
Combine two streams by creating a tuple for each pair of input events



This one is a work in progress

$$\lambda^A_{in} = \lambda^B_{in}$$


Example transformation



The design of Striot means we are able to rewrite a stream-processing program: rewrites preserve the functional behaviour but may change other aspects (e.g. performance)


Here we have replaced two serially-arranged “filter” nodes with a single “filter” which performs the work of both of the original operators.

Before fusion



This slide is busier than I'd like! The input and output arrival rates and the selectivity of each operator are known

After fusion



What are the properties of the new Server, performing the work of the two older Servers?

We are confident that we know the output rates and the filter selectivity: the product of the input selectivities

But what of the mean service rate?

Intuitively Since the new server is doing the work of the two original servers, it seems the mean average service time should be the sum of the mean average service times for the two servers. This does not model any efficiency gains (or losses)

Generalised question

“Can we represent a Jackson network as a single server and what are the properties of that server?”

I've seen the technique of applying Little's Theorem to a sub-part of a Jackson network, but Little's Theorem does not appear to help us determine the mean service rate because it is not defined in terms of service rates