

Monitor, Mitigate, Moderate: Backpressure in Stream Benchmark Generators

14th International Workshop on Load Testing and Benchmarking of Software Systems, **LTB 2026**
Florence, Italy

Iain Dixon, Matthew Forshaw, Joe Matthews

4th of May 2026

Newcastle University - School of Computing



Brief Explanation of Stream Processing Systems (SPS)




Figure 1. Logical layout of a Stream Processing System (SPS), with an upstream **source**, downstream **sink**, and pipelined operators. Each operator has a processing speed $x\rho$ where x is the multiple of the comparative processing speed of the minimum operator's processing speed ρ .

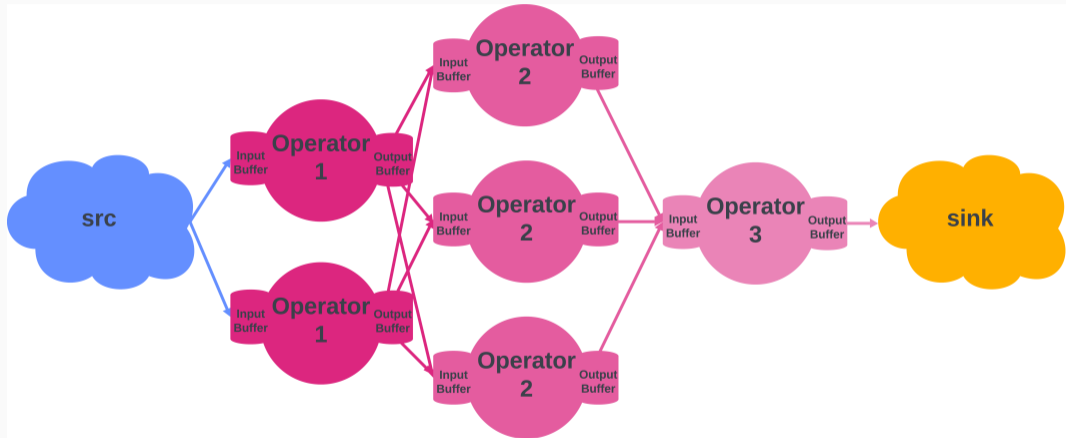


Figure 2. Physical layout of an SPS where the operators are scaled to handle processing loads between operators.




Figure 3. Load shedding maintains the external SPS rate by dropping external records received.

Backpressure halts upstream operators to throttle internal rates




Figure 4. Backpressure maintains the internal SPS rate by slowing down internal records received.




Figure 6. An open SPS benchmark operates in a similar way to the deployed case, with the generator continuing agnostic of how the SUT functions.

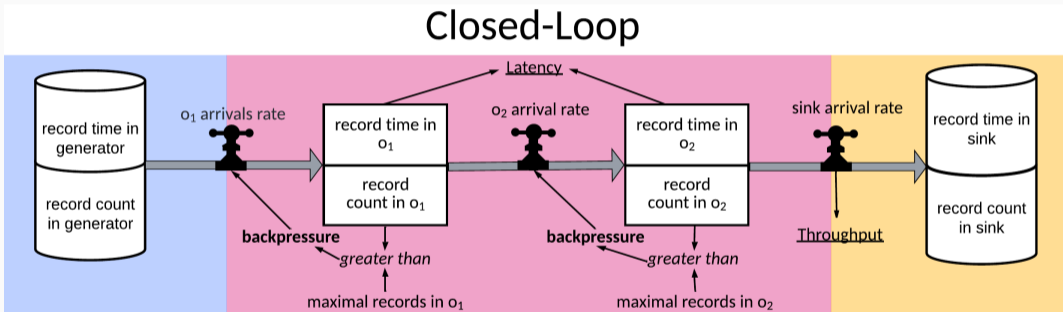


Figure 7. A closed SPS benchmark cannot ensure the level of separation between the generator and SUT, resulting in possible feedback if the SPS ingestion buffer overflows.



ALL ("backpressure" OR "back pressure" OR "back-pressure")
AND ALL ("stream" OR "streaming")
AND ALL ("benchmark" OR "benchmarking")
AND NOT ALL ("vehicle")
LIMIT-TO (SUBJAREA , "COMP")
LIMIT-TO (LANGUAGE , "ENGLISH")

Without Backpressure Limiter we received 52,620 articles, including it resulted in 63 articles as of 12th August, 2025.

	Scopus	Alt BP	Citation Only
# Papers	63	2	22
	Public Access		Final Included
# Papers	1		38

Table 1. Manual Review of Scopus Query. **Alt BP** = Excluded for being about an alternative definition of backpressure, **Citation Only** = Excluded for only having reference to backpressure in citations, **Public Access** = Excluded for not being public accessible.

0. Papers that **do not explicitly acknowledge** backpressure in the text
1. Papers that explicitly acknowledge backpressure in the text but **do not monitor** for it during benchmarking
2. Papers that explicitly acknowledge backpressure in the text and monitors for it during benchmarking but **do not mitigate** the effects
3. Papers that explicitly acknowledge backpressure in the text, monitors for it during benchmarking, and mitigates the effects to prevent it from confounding results

Level 0	Level 1	Level 2	Level 3
22	32	7	2
			One of which is our previous publication

Table 2. Categorisation of literature into levels 0, 1, 2, and 3 of backpressure acknowledgement

Generator & Workload Operators

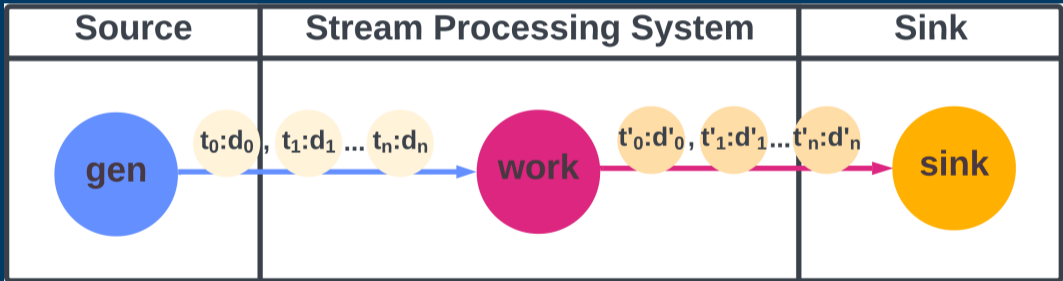


Figure 8. Experimental testbed framework.

Algorithm 1: Basic Generation Loop

Input: $W, \Lambda, \delta_L = 1000$

```
1 for  $w := 0$  to  $W - 1$  do
2    $\lambda_w := \Lambda(w)$ 
3    $n_w := \lambda_w * \delta_L / 1000$ 
4    $t_{start} := \text{TIME\_MS}()$ 
5   for  $i := 0$  to  $n_w - 1$  do
6      $\text{rec} := \text{GEN}()$ 
7      $\text{rec.gen\_time} := \text{TIME\_MS}()$ 
8      $\text{SEND}(\text{rec})$ 
9    $\delta_{transmit} := \text{TIME\_MS}() - t_{start}$ 
10   $\delta_{remain} := \delta_L - \delta_{transmit}$ 
11   $\delta_w := \delta_{transmit} + \delta_{remain}$ 
12  if  $\delta_{remain} > 0$  then
13     $\text{WAIT\_MS}(\delta_{remain})$ 
```

W	Number of windows
δ_L	Target window duration
w	Window index
Λ	Arrival process
λ_w	Target arrival rate
n_w	Target arrival load
t_{start}	Window start time
$\delta_{transmit}$	Transmit duration
δ_{remain}	Remaining duration
δ_w	Window duration
rec.gen_time	Record Generation

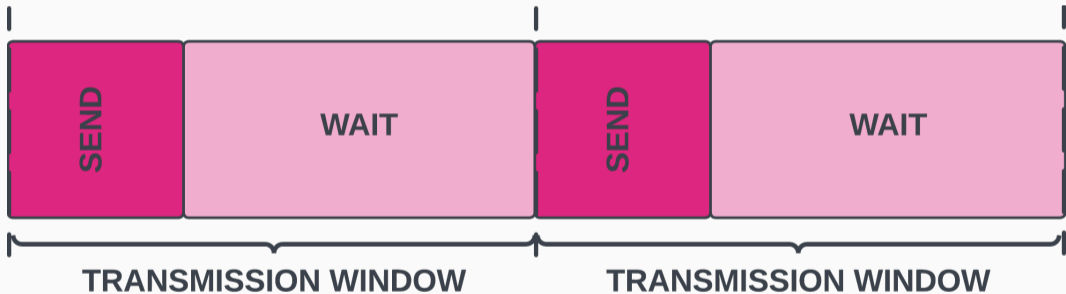


Figure 9. Algorithm 1 produces a transmission pattern where for each transmission window the generator sends all of the records, then waits until the next transmission window.

Using BGL we can generate a variety of arrival processes

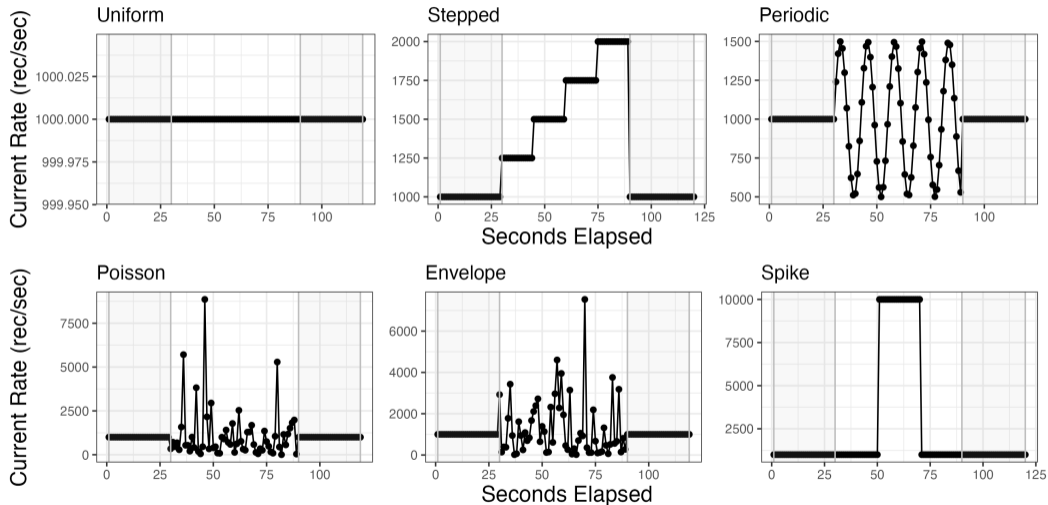


Figure 10. Implemented arrival processes using the experimental generator operator: uniform, periodic, poisson, envelope, and spike. For each process, the generator is provided with ($W = 60, \lambda = 1000$).

Algorithm 2: Workload Simulation Loop

Input: δ_{cost} , f_{work} , t_{start} , t_{end}

```
1  $i := 0$ 
2  $t_{start} := \text{TIME\_MS}()$ 
3 while  $rec = \text{GET\_NEXT}()$  do
4    $t_{current} := \text{TIME\_MS}()$ 
5    $j := (t_{current} - t_{start})/1000$ 
6   if  $t_{current} \geq t_{act} \wedge t_{current} < t_{deact}$  then
7     if  $j \bmod f_{wind} = 0 \wedge i \bmod f_{work} = 0$ 
8       then
9          $\text{WAIT\_MS}(\delta_{cost})$ 
10         $\text{SEND}(rec)$ 
11    $i := i + 1$ 
```

δ_{cost}	Per-record cost
f_{work}	Work Frequency
t_{start}	Workload start time
t_{end}	Workload end time
t_{act}	Workload activation time
t_{deact}	Workload deactivation time
f_{wind}	Window Frequency
i	Record Index
j	Window Index

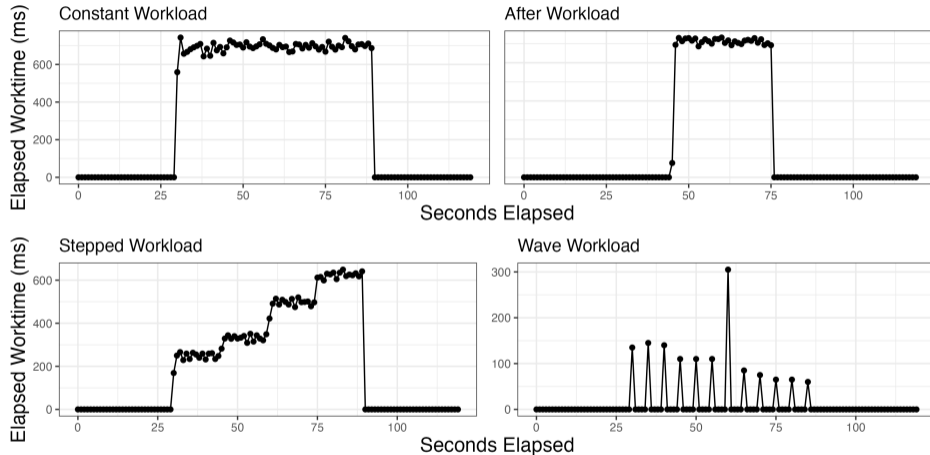


Figure 11. Implemented workload processes using the experimental workload operator: **constant workload** where δ_{cost} and f_{work} are constant, **delayed workload** where the workload does not activate immediately once the experiment starts, **stepped workload** where the workload increases in steps rather than remaining constant throughout the experiment, and **wave workload** where the work oscillates throughout the experiment.

Total workload (Δ) experienced by the workload can be thought of as a product of the per-record cost (δ_{cost}), work frequency (f_{work}), and per-second arrival rate (λ_w), such that:

$$\Delta = \frac{\delta_{cost} \cdot \lambda_w}{f_{work}}$$

By increasing total workload greater than 1000ms, we can fill the workload operator's input buffers and trigger backpressure in the generator, and by monitoring the transmission window duration detect its occurrence.

We want to run a 60 second experiment, with 30 seconds of calibration (shown in highlighted grey) time pre and post experiment for the system to stabilise and investigate metrics either side of the experiment.

Any metric captured after the experiment and calibration windows have finished will be considered an **experimental overrun** (shown in highlighted dark red).

The generator operator will provide a constant arrival process with each transmission window targetting 1000ms, and workload operator producing a total workload totalling 500ms, 1000ms¹, and 2000ms.

¹In practice this equates to slightly over 1000ms of workload due to logic overhead, producing a total workload which slightly exceeds the transmission window.

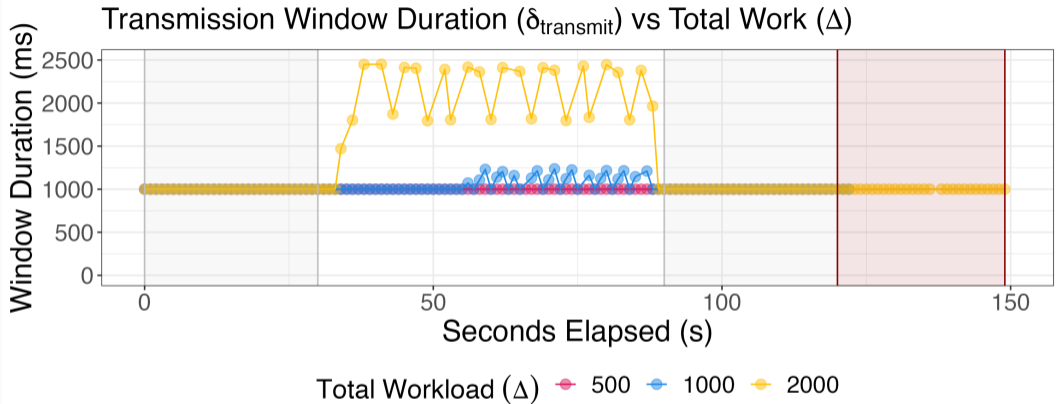


Figure 12. Transmission window durations (δ_{transmit}) for increasing total workloads ($\Delta = 500, 1000, 2000$).

msBackpressuredPerSecond

"The time (in milliseconds) this task is back pressured (soft or hard) per second. It's a sum of `softBackPressuredTimeMsPerSecond` and `hardBackPressuredTimeMsPerSecond`." Rolling average over the past 60 seconds.

isBackpressured

"Whether the task is back-pressured." Per-second boolean measure if an operator is in the backpressured state.

outPoolUsage

"An estimate of the output buffers usage. The pool usage can be > 100% if overdraft buffers are being used." Per-second measure of the output buffer utilisation across all of an operators output buffers.

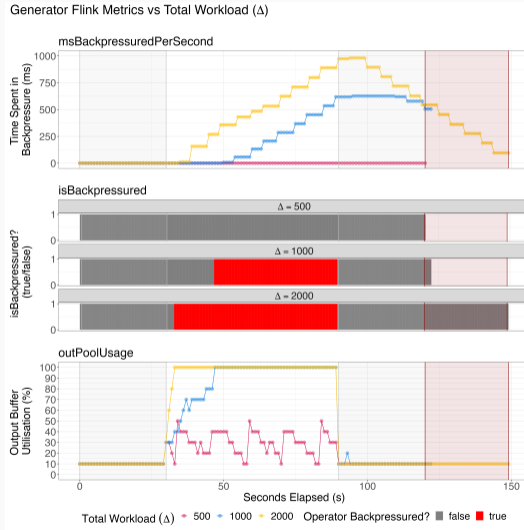
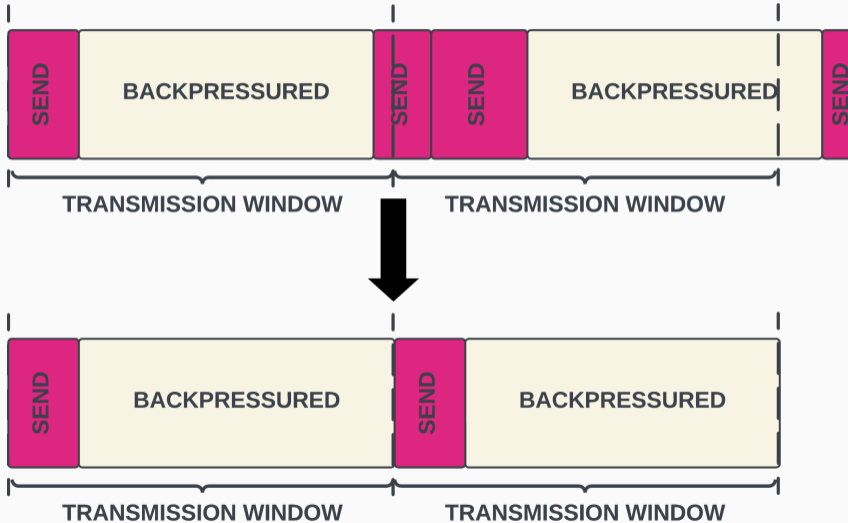


Figure 13. Flink metrics applied to escalating total workloads ($\Delta = 500, 1000, 2000$).

Cutting off transmission windows when overrun





Algorithm 3: Cutoff

Input: W, δ_L, λ

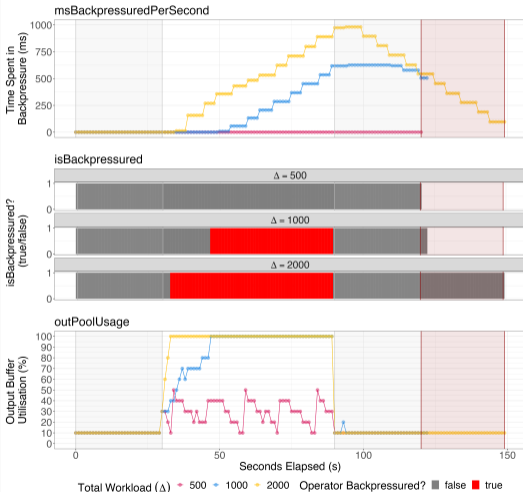
```
1  $\delta_{remain} := 0$ 
2 for  $w := 0$  to  $W - 1$  do
3    $n_w := \lambda_w * \delta_L / 1000$ 
4    $t_{start} := \text{TIME\_MS}()$ 
5   for  $i := 0$  to  $n_w - 1$  do
6     if  $\text{TIME\_MS}() - t_{start} > \delta_L$  then
7       | BREAK()
8      $\text{rec} := \text{GEN}()$ 
9     SEND( $\text{rec}$ )
10   $\delta_{transmit} := \text{TIME\_MS}() - t_{start}$ 
11   $\delta_{remain} := \delta_L - \delta_{transmit}$ 
12   $\delta_w := \delta_{transmit} + \delta_{remain}$ 
13  if  $\delta_{remain} > 0$  then
14    | WAIT\_MS( $\delta_{remain}$ )
15   $\delta_{remain} := \delta_{remain} + \delta_w$ 
```

Cutoff reactively reduces transmission duration overruns

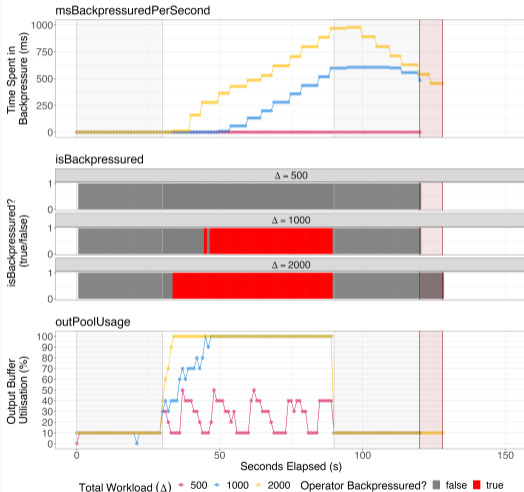


Cutoff doesn't improve Flink metrics, but it does reduce overruns

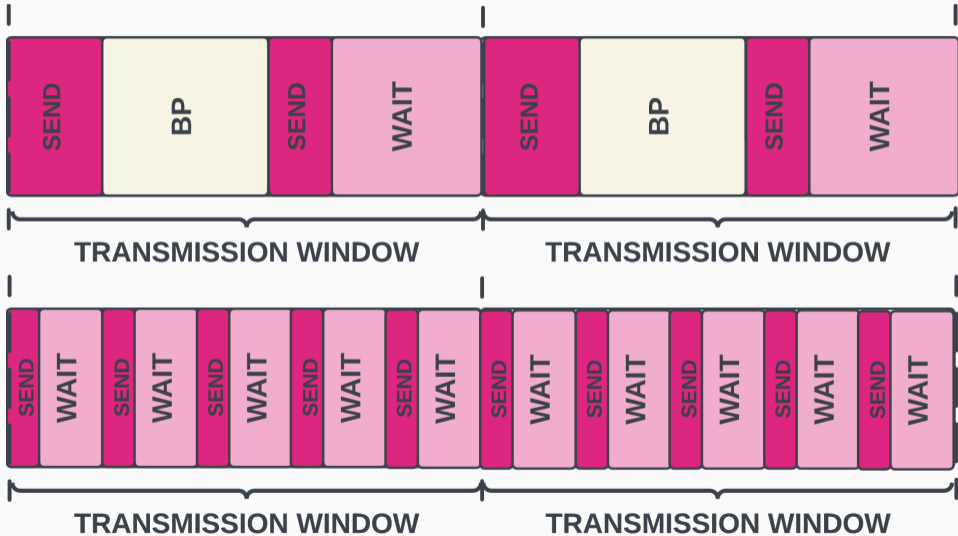
Generator Flink Metrics vs Total Workload (Δ)



Generator Flink Metrics vs Total Workload (Δ)



Spacing out records across a transmission mitigates backpressure



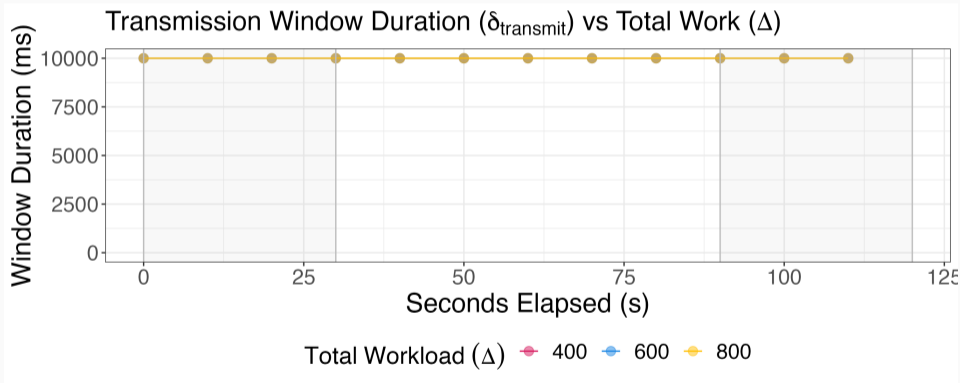
Algorithm 4: Spaced Out

Input: $W, \delta_L, \lambda, \delta_{inter}$

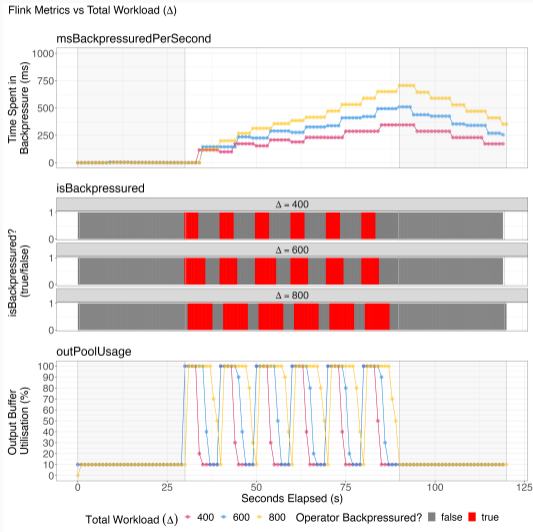
```
1  $\delta_{remain} := 0$ 
2 for  $w := 0$  to  $W - 1$  do
3    $n_w := \lambda_w * \delta_L / 1000$ 
4    $t_{start} := \text{TIME\_MS}()$ 
5   for  $i := 0$  to  $n_w - 1$  do
6      $\text{rec} := \text{GEN}()$ 
7      $\text{SEND}(\text{rec})$ 
8     if  $\delta_{remain} > 0 \wedge i \bmod \lambda_w / \delta_{remain} = 0$  then
9        $\text{WAIT\_MS}(\delta_{inter})$ 
10       $\delta_{wait} := \delta_{wait} + \delta_{inter}$ 
11    $\delta_{transmit} := \text{TIME\_MS}() - t_{start}$ 
12    $\delta_{remain} := \delta_L - \delta_{transmit}$ 
13    $\delta_w := \delta_{transmit} + \delta_{remain}$ 
14   if  $\delta_{remain} > 0$  then
15      $\text{WAIT\_MS}(\delta_{remain})$ 
16    $\delta_{remain} := \delta_{remain} + \delta_{wait}$ 
```

δ_{inter} Interspersed wait duration

δ_{wait} Total interspersed wait duration

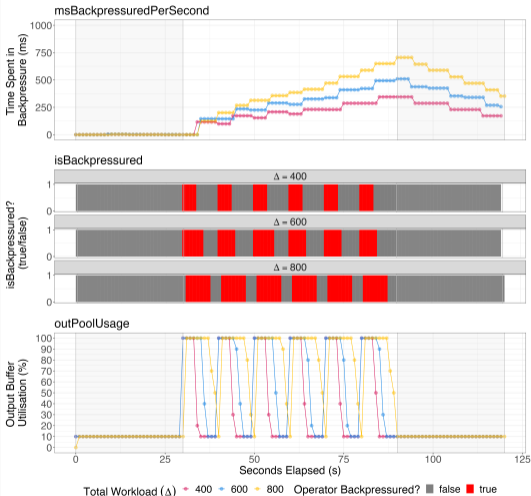


Without spacing backpressure occurs within the 10s window

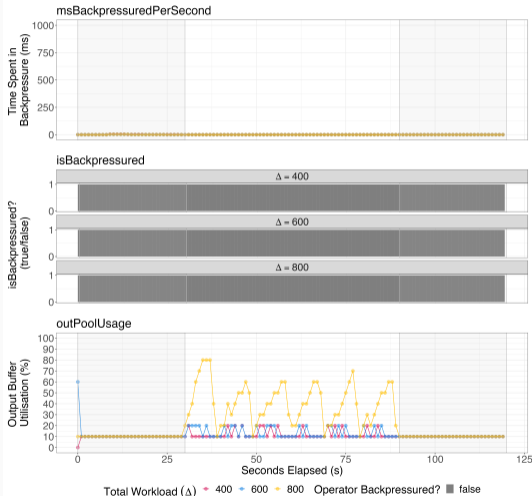


Spacing mitigates backpressure which Flink cannot capture normally

Flink Metrics vs Total Workload (Δ)



Flink Metrics vs Total Workload (Δ)



1. Is the benchmark closed-loop?
2. If so, is the generator affected by backpressure from the SUT?
3. If so, can you space records of cut off overrunning transmission windows?
4. If not, is the amount of backpressure captured in the metrics cause for concern? (Use-case specific)

Thank You!

i.g.dixon1@ncl.ac.uk



Kalavri, Vasiliki et al. **“Three steps is all you need: Fast, accurate, automatic scaling decisions for distributed streaming dataflows”**. In: *Proceedings of the 13th USENIX Symposium on Operating Systems Design and Implementation, OSDI 2018*. Cited by: 71. 2007, pp. 783–798.