

Progressive Temporal Window Widening

David Tolpin

PayPal, dtolpin@paypal.com

Abstract. This paper introduces a scheme for data stream processing which is robust to batch duration. Streaming frameworks process streams in batches retrieved at fixed time intervals. In a common setting a pattern recognition algorithm is applied independently to each batch. Choosing the right time interval is tough — a pattern may not fit in an interval which is too short, but detection will be delayed and memory may be exhausted if the interval is too long. We propose here Progressive Window Widening, an algorithm for increasing the interval gradually so that patterns are caught at any pace without unnecessary delays or memory overflow.

This algorithm is relevant to computer security, system monitoring, user behavior tracking, and other applications where patterns of unknown or varying duration must be recognized online in data streams. Modern data stream processing frameworks are ubiquitously used to process high volumes of data, and adaptive memory and CPU allocation, facilitated by Progressive Window Widening, is crucial for their performance.

Keywords: temporal data streams, sliding windows, stream processing

1 Introduction

We consider here the problem of windowed data stream processing [7]. A data stream is a real-time, continuous, ordered sequence of items. In the windowed setting, the arriving data are divided into windows, either by time interval or by data size, and a pattern recognition algorithm, based on a data mining or machine learning approach, is applied to each window to discover exact or approximate patterns appearing in the window [6]. Here, we view a pattern recognition algorithm as a black box function on stream fragments. For example, a pattern can be an episode — a partially ordered sparse subsequence [10], the language of the text, or the most likely goal of the sequence of actions in the fragment.

Windowed data stream processing is frequently used in computer security [12, 11, 14], user behavior tracking [2], sensor data analysis for system monitoring [3], and other applications. The right choice of window size is crucial for efficient data processing and timely response. Data are divided either into physical windows, by time interval, or into logical, or count-based, windows, by data size or number of records in a single window [7, 6].

The choice of either physical or logical windows depends both on properties of the data stream and on the objective of the data processing algorithm. Logical

windows are more naturally handled by machine learning algorithms with inputs of fixed size [6], while physical windows allow both more efficient processing and faster online response [9, 16, 15]. This paper explores selecting a window size for physical, interval-based windows. The dilemma behind selecting a window size which inspired this research is

- whether to choose a smaller window and sacrifice context, such that no single window contains a complete pattern,
- or to increase the window size at the cost of increased consumption of computational resources and delayed response.

This dilemma is relevant to many applications of data stream processing, but in particular to security applications [12, 11, 14], where an adversary aware of the maximum window time interval can escape the detection algorithm by introducing delays between data stream entries (such as transactions or web site accesses) which exceed the interval and prevent detection. Even if the maximum duration of a pattern is known in advance, setting the window size to exceed the maximum duration means that recognition of any shorter pattern will be delayed.

To address this dilemma, we introduce an algorithm which we call *Progressive Window Widening* (PWW). PWW processes the data stream through an array of sliding windows of increasing physical size, such that shorter patterns are recognized sooner, however windows covering longer patterns are also applied to the stream. Despite employing several window sizes in parallel, PWW still remains efficient in CPU and memory consumption. The paper proceeds as follows: first, necessary preliminaries are introduced in Section 2. Then, the algorithm is described and analysed (Sections 3 and 4), as well as evaluated empirically (Section 5). Finally, related work is reviewed, and contribution and future research are discussed (Sections 6 and 7).

The full paper is available at <http://arxiv.org/abs/1604.00997>.

2 Preliminaries

2.1 Batched Stream Processing

In *batched stream processing*, which we adopt in this paper as a lower level for PWW, stream data arrives in batches — sequences of fixed duration. Several batches can be combined into a window of size equal to the total size of the batches composing the window. Along with batch size (or duration, used interchangeably here), a batch is characterized by its *length*, the number of atomic elements, or records, it contains. For example, a one-minute batch of web site log stream may contain 1000 entries — we shall say that the size, or duration of the batch is 1 minute, and the length of the batch is 1000 entries.

Further on, we extend the note of batched stream processing by stating that a data stream with batch duration t may be transformed into a data stream with batch duration kt by concatenating each k consecutive batches together.

Denoting a batch of the original stream with batch duration t by $B_{i,l}$ and a batch of the combined stream with batch duration kt by $B_{i+1,j}$ for some i, j , and l , one may write (\circ stands for batch concatenation):

$$B_{i+1,j} = B_{i,kj-k+1} \circ B_{i,kj-k+2} \circ \cdots \circ B_{i,kj} \quad \forall j \in \mathbb{N}^+ \quad (1)$$

2.2 Sliding Windows

Depending on the overlay between windows, one discerns between *tumbling* (there are gaps between windows), *jumping* (the windows are adjacent), and *sliding* (overlapping) windows [7]. PWW is based on sliding windows; the next window starts earlier than the current window terminates.

Sliding windows have several uses. We are interested in one particular case: sliding windows with a half-size overlap; the feature we are interested in is described by Lemma 1:

Lemma 1 *A sequence of sliding windows of size $2b$ with overlap b covers any interval of size at most b .*

A corollary of Lemma 1 is that if we want to recognize patterns of duration at most t , it is sufficient to use sliding windows of size $2t$ with half-size overlap.

3 Progressive Window Widening

We introduce here *Progressive Window Widening*, an algorithm for progressive widening of temporal windows. To define the algorithm efficiently, we rely on an auxiliary notion of L_{max} — the maximum length of a data sequence which may contain a pattern. For example, if a game player must complete each game round in 20 moves, then any pattern pertaining to a single round must be contained within 20 moves. Alternatively, L_{max} can be chosen such that the probability of a random occurrence of the pattern in a data sequence of length L_{max} is sufficiently low [8].

The algorithm processes the data stream in parallel, through multiple asynchronous sliding windows of different sizes.

3.1 Algorithm Outline

PWW (Algorithm 1) performs the following operations:

1. Recursively combines pairs of adjacent batches, doubling batch duration of each stream and creating a stream with batches of double duration (line 3).
2. Runs a detection algorithm in a sliding window on each stream (line 6).
3. While combining batches, discards subintervals of combined batches which cannot intersect a yet unseen pattern (see Section 3.2 for detailed explanation).

Algorithm 1 Progressive Window Widening

-
- 1: **procedure** PWW(S -- stream, t -- batch duration)
 - 2: SLEEP(t)
 - 3: Create stream S' from S with batch duration $2t$ (see Algorithm 2)
 - 4: Call PWW(S' , $2t$) *asynchronously*
 - 5: **for each** sliding window W **in** S **do**
 - 6: **if** patterns present in W **then**
 - 7: Output detected patterns
-

As the algorithm runs, multiple batched streams are created, and sliding windows move through each of the streams (Figure 1). The algorithm relies on asynchronous recursive calls to PWW (line 4). Asynchronous calls are possible because the processing of each stream is independent. Such asynchronous execution is particularly suitable for modern multi-core multi-node cluster architectures: different invocations of PWW may be executed on different cores or different nodes in the cluster.

Note that extra streams are created (lines 2–4) and processed (line 5) with exponentially increasing delays, since a window can be processed only upon termination of the window’s interval.

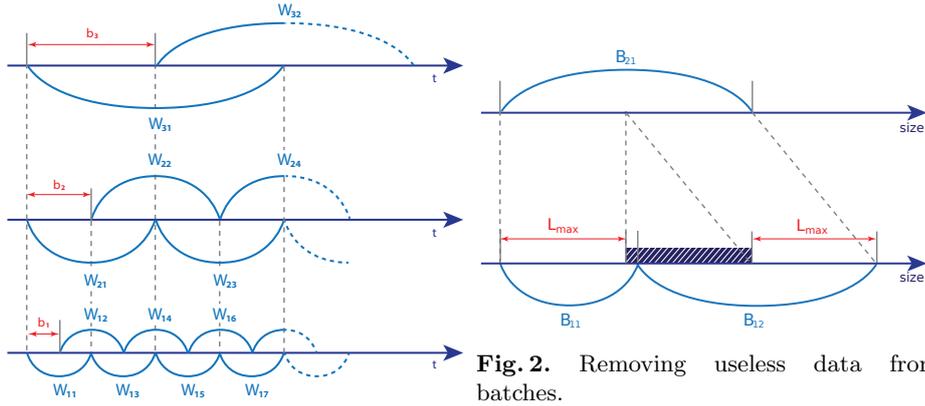


Fig. 1. Window widening.

Fig. 2. Removing useless data from batches.

3.2 Combining batches

An integral part of PWW is the optional discarding of a subinterval while combining two subsequent batches. For every stream of batches of duration t , the algorithm waits $2t$ time units for 2 batches to arrive. Then, a stream of base duration $2t$ is formed by combining the batches (Algorithm 2). PWW combines batches by concatenation (line 2). If the length of the combined batch is greater than $2L_{max}$, the middle part of the combined batch is discarded (Figure 2), leaving subsequences of length L_{max} at both ends of the batch (lines 3–4). Consequently, no batch in any stream is longer than $2L_{max}$. The subintervals may

Algorithm 2 Combining Batches

```

1: procedure COMBINE( $B_{i-1,2j-1}, B_{i-1,2j}$  -- consequent batches)
2:  $B_{i,j} \leftarrow$  CONCATENATE( $B_{i-1,2j-1}, B_{i-1,2j}$ )
3: if LENGTH( $B_{i,j}$ ) >  $2L_{max}$  then
4:   REMOVE( $B_{i,j}$ , from= $L_{max}$ , till=LENGTH( $B_{i,j}$ )- $L_{max}$ )
5: return  $B_{i,j}$ 

```

be discarded because a combined batch at the next level coincides with a sliding window at the current level, so new patterns may be discovered only between batches, rather than within a single batch.

4 Algorithm Analysis

We launch an unbounded number of parallel processes, and want to show that PWW runs in computationally bounded resources. The work that the algorithm performs is assumed to take place inside a pattern recognition algorithm run on each sliding window. Let us denote the resources (a combination of memory and amount of work) required to run a certain pattern recognition algorithm on window of length l by $\mathcal{R}(l)$. Then, the following theorem holds:

Theorem 1 *Denote by t the batch duration of the initial, uncombined stream. Assume that the maximum length of a batch of the initial stream does not exceed $2L_{max}$. Then the average resources ρ per unit time required to run PWW are bounded by a constant:*

$$\rho \leq \frac{2\mathcal{R}(4L_{max})}{t}. \quad (2)$$

In practice, the number of parallel streams may be bounded. However, even if unbounded, average resources required to run the algorithm are bounded.

5 Case Study: Detecting Remote Shells in a System Call Stream

In this case study, we monitor an online stream of system calls from a network-connected server, and want to detect possible invocations of remote shells as soon as possible. System call sequences corresponding to remote shell invocations can be interspersed with unrelated activities.

For simplicity, we limit detection to a single episode which may correspond to accepting a network connection and then launching a shell communicating with the remote user through the connection:

```

1 accept fd=x => y
2 dup fd=y => 0 | dup fd=y => 1 | dup fd=y => 2
3 execve exe=z

```

In the above pseudocode, system call name is followed by `name=value` argument pairs and then by return value preceded by `=>`. In a matching system call sequence `y` must have the same value in lines 1 and 2, three system calls in line 2 may be executed in any order, and `x`, `z` may take any value. For example, sequence

```
accept fd=5 => 6, dup fd=6 => 2, dup fd=6 => 1, dup fd=6 => 0, execve exe=sh
```

matches the episode.

For the empirical evaluation we use a sequential version of PWW which facilitates easy estimation of the amount of work. We set $L_{max} = 100$ because malicious code is often transmitted in a single packet with only a few dozens of instructions. We use a stream of 10 000 system calls recorded on a Linux machine, into which we inject episode instances with varying delays between instructions. We find that:

- The detection delay (Figure 3) is proportional to the episode duration with factor 0.5.
- The amount of work (Figure 4) approaches but stays below bound (2).

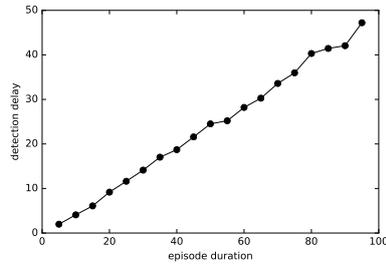


Fig. 3. Detection delay.

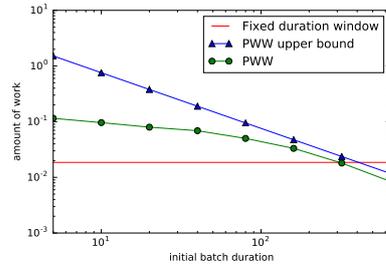


Fig. 4. Amount of work.

Both results are in accordance with the algorithm analysis. If a fixed window were used, either the average detection delay would grow, or some episodes were left undetected.

The source code, data, and results for the case study are available at <https://bitbucket.org/dtolpin/pww-paper-case-studies>. The evaluation notebook can be viewed in the browser at <http://tinyurl.com/jgknulz>.

6 Related Work

While Progressive Window Widening can be implemented from scratch on low-level data streams, the algorithm was inspired and relies in implementation on batched stream processing. Batch stream processing was introduced in Comet [9]. Apache Spark offers Spark Streaming [16, 15], a powerful implementation of programming model *discretized streams*. Discretized streams, which enable efficient batch processing in parallel architectures, is the enabling lower level for PWW.

PWW uses varying window sizes to accommodate for differences in data. Another approach in batched stream processing is to use *adaptive window size*.

Adaptive window algorithms is a field of active research [17, 4, 5, 13]. However, this research represents a different approach, in which the window size is changed sequentially and adaptively, for future windows based on earlier seen data. In PWW, several windows of fixed sizes are applied in parallel, in a parameter-free manner suitable for simple and robust implementation. Windows of doubling size were proposed for processing data streams in earlier work [1], however the approach employed in PWW is significantly different in that temporal windows of unbounded doubling durations are applied in parallel, while still ensuring efficient use of resources.

7 Contribution and Future Research

This paper introduced the Progressive Window Widening algorithm for data stream processing using temporal sliding windows. The algorithm

- solves the dilemma of smaller window size at a cost of inability to recognize longer patterns versus larger windows but slower response;
- works in parallel, in a manner suitable for modern multi-core multi-node cluster architectures;
- uses computational resources efficiently, imposing only a constant factor overhead compared to an algorithm based on a single window size.

The basic algorithm described in the paper brings a solution to the stated problem. At the same time, the algorithm design poses a number of questions and opens several research directions.

- Many adaptive window algorithms are, unlike PWW, essentially sequential. Modern data frameworks provide an opportunity to exploit the parallelism for more flexible and efficient adaptation.
- Doubling of batch durations is chosen in PWW due to simplicity of implementation and analysis. A different allocation of window sizes, either data-independent or adaptive, may bring better theoretical performance and practical results.
- PWW relies on batched stream processing, however it is only loosely coupled with the underlying computing architecture, which is both an advantage and a drawback. A tighter coupling with lower-level stream processing may be helpful.

Along with others, these directions are deemed to be important for future research.

References

1. Aggarwal, C.C., Han, J., Wang, J., Yu, P.S.: A framework for clustering evolving data streams. In: 29th International Conference on Very Large Data Bases. pp. 81–92. VLDB '03, VLDB Endowment (2003)

2. Agrawal, D., Budak, C., Abbadi, A., Georgiou, T., Yan, X.: Databases in Networked Information Systems: 9th International Workshop, chap. Big Data in Online Social Networks: User Interaction Analysis to Model User Behavior in Social Networks, pp. 1–16. DNIS '14, Springer International Publishing (2014)
3. de Aquino, A.L.L., Figueiredo, C.M.S., Nakamura, E.F., Buriol, L.S., Loureiro, A.A.F., Fernandes, A.O., Coelho, C.J.N.J.: Data stream based algorithms for wireless sensor network applications. In: IEEE 21st International Conference on Advanced Information Networking and Applications. pp. 869–876. AINA '07, IEEE (2007)
4. Bifet, A., Gavald, R.: Learning from time-changing data with adaptive windowing. In: SIAM International Conference on Data Mining. SDM '07, SIAM (2007)
5. Bifet, A., Pfahringer, B., Read, J., Holmes, G.: Efficient data stream classification via probabilistic adaptive windows. In: 28th Annual ACM Symposium on Applied Computing. pp. 801–806. SAC '13, ACM (2013)
6. Gama, J.: A survey on learning from data streams: current and future trends. *Progress in Artificial Intelligence* 1(1), 45–55 (2012)
7. Golab, L., Özsu, M.T.: Issues in data stream management. *SIGMOD Record* 32(2), 5–14 (2003)
8. Gwadera, R., Atallah, M., Szpankowski, W.: Reliable detection of episodes in event sequences. In: Third IEEE International Conference on Data Mining. pp. 67–74. ICDM '03, IEEE (2003)
9. He, B., Yang, M., Guo, Z., Chen, R., Su, B., Lin, W., Zhou, L.: Comet: Batched stream processing for data intensive distributed computing. In: 1st ACM Symposium on Cloud Computing. pp. 63–74. SoCC '10, ACM (2010)
10. Mannila, H., Toivonen, H., Inkeri Verkamo, A.: Discovery of frequent episodes in event sequences. *Data Mining and Knowledge Discovery* 1(3), 259–289 (1997)
11. Varghese, S.M., Jacob, K.P.: Anomaly detection using system call sequence sets. *Journal of Software* 2(6), 14–21 (2007)
12. Warrender, C., Forrest, S., Pearlmuter, B.: Detecting intrusions using system calls: alternative data models. In: IEEE Symposium on Security and Privacy. pp. 133–145. SP '99, IEEE (1999)
13. Yang, Y., Mao, G.: International Conference of Intelligence Computation and Evolutionary Computation, chap. A Self-Adaptive Sliding Window Technique for Mining Data Streams, pp. 689–697. ICEC '12, Springer Berlin Heidelberg (2013)
14. Yolacan, E.N., Dy, J.G., Kaeli, D.R.: System call anomaly detection using multi-HMMs. In: IEEE Eighth International Conference on Software Security and Reliability-Companion. pp. 25–30. SERE-C '14, IEEE (2014)
15. Zaharia, M., Das, T., Li, H., Hunter, T., Shenker, S., Stoica, I.: Discretized streams: Fault-tolerant streaming computation at scale. In: Twenty-Fourth ACM Symposium on Operating Systems Principles. pp. 423–438. SOSP '13, ACM (2013)
16. Zaharia, M., Das, T., Li, H., Shenker, S., Stoica, I.: Discretized streams: An efficient and fault-tolerant model for stream processing on large clusters. In: 4th USENIX Conference on Hot Topics in Cloud Computing. pp. 10–15. HotCloud'12, USENIX Association (2012)
17. Zhang, D., Li, J., Zhang, Z., Wang, W., Guo, L.: Advances in Web-Age Information Management: 5th International Conference, chap. Dynamic Adjustment of Sliding Windows over Data Streams, pp. 24–33. WAIM '04, Springer Berlin Heidelberg (2004)