

A novel algorithm for online inexact string matching and its FPGA implementation

Alessandro Cinti¹, Filippo Maria Bianchi², Alessio Martino¹, and Antonello Rizzi¹

¹*Department of Information Engineering, Electronics and Telecommunications, University of Rome "La Sapienza", Via Eudossiana 18, 00148 Rome, Italy*

²*Machine Learning group, Department of Physics and Technology, UiT the Arctic University of Norway, Hansine Hansens veg 18, 9019 Tromsø, Norway*

Abstract

Accelerating inexact string matching procedures is of utmost importance when dealing with practical applications where huge amount of data must be processed in real time, as usual in bioinformatics or cybersecurity. Inexact matching procedures can yield multiple shadow hits, which must be filtered, according to some criterion, to obtain a concise and meaningful list of occurrences. The filtering procedures are often computationally demanding and are performed offline in a post-processing phase. This paper introduces a novel algorithm for Online Approximate String Matching (OASM) able to filter shadow hits on the fly, according to general purpose priority rules that greedily assign priorities to overlapping hits. An FPGA hardware implementation of OASM is proposed and compared with a serial software version. Even when implemented on entry level FPGAs, the proposed procedure can reach a high degree of parallelism and superior performance in time compared to the software implementation, while keeping low the usage of logic elements. This makes the developed architecture very competitive in terms of both performance and cost of the overall computing system.

1 Introduction

Approximate String Matching (ASM) studies the problem of matching two generic strings defined on the same alphabet and, unlike in Exact String Matching, a level of inexactness is allowed: two strings match if their dissimilarity is below a given threshold [1]. ASM is a particular case of subgraph matching [2, 3] that has been applied in different fields of science and technology, including computational biology, signal processing and text retrieval. In computational biology, scientists try to detect frequent patterns into DNA sequences of nucleotides (motifs) using ASM to account for mutations and evolutionary alterations of the genome sequences [4–7]. In signal processing, ASM can be used to identify patterns which have been corrupted by noise during storage, transmission, processing or conversion stages. The main application fields are message exchanges, wireless communications, audio, image and video processing [8, 9]. ASM is extensively used for information retrieval in large text collections where, due to the large amount of data and the high variety of content, classical string matching procedures are usually not suitable [10–14].

ASM also plays a fundamental role in Granular Computing, which is a human-inspired computing and information processing paradigm that explores multiple levels of granularity in data [15]. The concept of Granular Computing arose from many branches of natural and social sciences [16–18] and it is at the basis of recently developed frameworks in computational intelligence [19–21]. In this context, an ASM technique that identifies frequent and meaningful motifs in sequences database allows to design advanced machine learning systems such as Symbolic Histograms approaches [22–24], where each pattern (a sequence of objects/events) can be represented by a histogram of motif instances.

The dissimilarity between two strings can be evaluated as the cost of the edit operations required for transforming a string into the other. The most common dissimilarities are the Levenshtein, Hamming, Episode, and Longest Common Subsequence distance; they differ on the edit cost definition and on the type of the allowed edit operations, and their computational costs range from linear to NP-complete [25]. The matching procedure can be implemented following offline or online approaches, which differ on how sequences are searched and indexed [26]. Several algorithms for online pattern matching have

been designed and they can be grouped in four main classes, namely the Dynamic Programming (DP), Automata, Bit-parallel, and Filtering approach [1]. The main advantage of online ASM is to provide detection in real time, which is essential in situations where a prompt response is required.

Cybersecurity is a prominent example of application of online ASM. In particular, Network Intrusion Detection Systems (NIDSs) are devices or software applications used to identify individuals who are using a computer system without authorisation and whoever has legitimate access to the system, but is exceeding his privileges [27]. Traditional NIDSs relied on exact pattern matching to detect an attack. However, by changing the data used in the attack even slightly, it is possible to evade detection. Therefore, a more flexible detection system is required to scan efficiently *in real-time* the whole inbound and outbound traffic to match patterns from a library of known attacks, without compromising the overall network speed.

In molecular biology, online ASM is exploited in quantitative real-time Polymerase Chain Reaction (PCR) [28]. Real-time PCR aims at amplifying a small target DNA sample during PCR, performing a quantitation step after amplification. This technique is useful when only a small amount of DNA is available, which is insufficient for performing an accurate analysis. In medicine, real-time PCR is applied to the discovery of tumour cells, to the diagnosis of infectious diseases, and to a plethora of other predictive medicine and diagnostic tasks [29]; it is widely used for studying genomes in several bacteria and protists which cannot be cultured; in legal medicine and forensics it has been demonstrated to be crucial for analysing fingerprints and stains found at crime scenes [30, 31].

To increase the time performance of string matching procedures, especially when dealing with big-data, their concurrent processes can be effectively parallelized at CPU level by executing microinstructions simultaneously, or at circuit level by re-implementing arithmetic and logic operations [32, 33]. Hardware implementation provides an alternative solution to improve the speed of ASM algorithms and several architectures have been proposed. The fastest implementations are custom architectures deploying a large number of processing elements designed *ad-hoc* to execute a specific algorithm [34]. Other architectures, instead, are based on general-purpose processors that can be adopted to implement different algorithms [35, 36]. Configurable devices, such as the Field-Programmable Gate Array (FPGA), represent flexible hardware platforms that are used to accelerate the computation, providing both a high degree of programmability and reduced design time and costs [37, 38].

None of the existing methods for performing online ASM accounts for the presence of multiple shadow hits, which are overlapping hits that appear as a consequence to the flexibility in ASM. Usually, such shadow hits are filtered afterwards by means of a post-processing phase which increases the overall computational burden. Especially when dealing with data streams, this solution turns out to be unfeasible both in terms of memory and time complexity, hence the necessity to filter out shadow hits on the fly.

This paper introduces the Online Approximate String Matching (OASM) algorithm for retrieving a set of inexact matches of a known pattern from a stream of symbols, while filtering shadow hits on the fly. The proposed method performs detection in real time and is suitable for real-time applications and big-data streams. OASM is mainly based on the evaluation of the Levenshtein distance between a target pattern and a set of sequences, obtained by shifting windows of pre-established lengths over the stream. OASM follows a DP approach that provides a remarkable speedup thanks to the massive parallelization capability deriving from the inherent simplicity and regularity of its structure. Both a software (SW-OASM) and a hardware (HW-OASM) implementation of the algorithm are proposed in this paper. HW-OASM is based on the systolic arrays principle for the computation of the Levenshtein distance that extends the work in [39]. Compared to its software counterpart, experimental results show the effectiveness of the HW-OASM especially when integrated in a Multiple Online Approximate String Matching (MOASM) system. The latter, implements multiple instances of HW-OASM to perform simultaneously parallel searches of distinct patterns on a common input stream of symbols (not necessarily finite). The performance of the proposed algorithm is first evaluated in a controlled environment using synthetic data. Successively, the algorithm is applied to a case study in bioinformatics on real genome data. The analysis consists in mining part of the human genome for RNA-binding protein sites.

The remainder of the paper is organized as follows. Sec. 2 reviews related ASM approaches, including hardware-based ones. Sec. 3 introduces basic concepts and details of the dissimilarity measure considered. Sec. 4 describes the OASM algorithm and in Sec. 5 a real-world human genome case study is presented to highlight the features of SW-OASM, compared to other ASM approaches. Next, Sec. 6 provides details of the hardware implementation and in Sec. 7 the performances of SW-OASM and HW-OASM are compared on synthetic data. Finally, in Sec. 8 conclusions are reported.

2 Related Works

In literature, there are two major approaches based on DP for performing ASM on dedicated hardware, namely FPGAs and (GP-)GPUs. In this paper we focus on FPGAs and refer the interested reader to works such as [40–43] for ASM implementations on (GP-)GPUs. The first approach relies on local/global sequence alignment [44] (e.g. the Smith-Waterman’s algorithm [45], the Needleman-Wunsch’s algorithm [46] or Myers’ fast bit-vector algorithm [47]), while the second approach exploits the Levenshtein distance (e.g. the Wagner-Fischer’s algorithm [48]).

Implementation of the Smith-Waterman’s algorithm for optimal local alignment on FPGAs can be found in [49], where the authors implement the plain Smith-Waterman’s algorithm. In [50] the Smith-Waterman’s algorithm has been implemented on FPGAs without relying on systolic arrays in order to exploit all processing units, whereas in [51] the authors used systolic arrays driven by OpenCL. In [52] the authors provide FPGA implementation for DP and BLAST routines [53]. It is noteworthy that since the Smith-Waterman’s algorithm was primarily developed for nucleotide or protein sequence alignment, all works cited so far regard bioinformatics-related applications. Conversely, in [54], the authors implement the Smith-Waterman’s algorithm on FPGA for generic text search. Finally, in [55], the authors use an FPGA in order to speedup Myers’ algorithm.

General purpose applications are more frequently relying on the second approach, based on the Wagner-Fischer’s algorithm. In [39] the authors adopt the plain Wagner-Fischer’s algorithm, along with an improved version that better exploits each cell. In [56] the Wagner-Fischer’s algorithm has been applied to multimedia information retrieval, also considering text search paradigms such as wildcards and idioms. In [57] the authors provide an FPGA implementation of the Wagner-Fischer’s algorithm particularly suited for dealing with regular expressions.

FPGA implementations of systolic architectures have been proposed for ASM [39, 54] and, in bioinformatics, for DNA sequence alignment [49, 50, 52]. In the field of Music Information Retrieval (MIR), circuits implementing ASM with dynamic programming cannot retrieve fragments with different sizes [56] and other architectures with higher flexibility only support symbol substitution [58]. Better performance was achieved by Application-Specific Integrated Circuit and FPGA with a comparable computational time on different MIR approaches [59–62]. An efficient FPGA implementation of ASM has been proposed for text mining, where a restricted class of regular expressions is used to define the patterns to search [57]. A work closely related to this paper is a successful implementation of online ASM on FPGA for NIDS applications [63].

Regardless of the particular algorithm used to implement ASM (online and/or using dedicated hardware), none of the works previously discussed filters shadow hits on the fly. In fact, those are discarded *a-posteriori*, increasing the overall computational burden. The proposed OASM jointly provides (i) the ability of performing simultaneously detection and shadow hits filtering; (ii) the ability of working with data streams, avoiding to store the entire input sequence and/or all multiple shadow hits. Contrarily to previous works, in OASM not only the plain Levenshtein distance is parallelized, but also the filtering procedure. This makes the proposed approach a perfect candidate for being implemented in FPGAs.

3 An Overview on the Levenshtein Distance

Let the pattern $\mathbf{p} = \langle p^{(0)}, p^{(1)}, \dots, p^{(l_p-1)} \rangle$ and the string $\mathbf{t} = \langle t^{(0)}, t^{(1)}, \dots, t^{(l_t-1)} \rangle$ be defined, respectively, as the concatenation of l_p and l_t symbols of a finite alphabet Σ and a generic substring $\mathbf{s} = \langle s^{(0)}, s^{(1)}, \dots, s^{(l-1)} \rangle = \langle t^{(i)}, t^{(i+1)}, \dots, t^{(i+l-1)} \rangle = \mathbf{t}[i, l]$ be a subset of l contiguous symbols in \mathbf{t} starting from position i . The dissimilarity between strings \mathbf{p} and \mathbf{s} can be measured by means of the Levenshtein distance $\text{lev}(\mathbf{p}, \mathbf{s}) \in \mathbb{N}_0$, which is the minimum number of single-character edits (insertion, deletion, substitution) necessary to transform \mathbf{p} into \mathbf{s} . Let \mathbf{C} be a matrix of size $(l_p + 1) \times (l + 1)$, whose elements are

$$c_{n,m} = \begin{cases} n, & \text{if } m = 0 \\ m, & \text{if } n = 0 \\ \min(c_{n-1,m} + 1, c_{n,m-1} + 1, & \\ c_{n-1,m-1} + \delta_{n,m}), & \text{if } n, m > 0 \end{cases} \quad (1)$$

where $\delta_{n,m} = (p^{(n)} \neq s^{(m)})$, $n = 0, \dots, l_p$, $m = 0, \dots, l$. The Levenshtein distance between \mathbf{p} and \mathbf{s} corresponds to the element $c_{l_p, l}$ of the Levenshtein matrix \mathbf{C} . Eq. 1 shows that, for $n, m > 0$, each element $c_{n,m}$ of the matrix \mathbf{C} can be computed just by knowing its upper ($c_{n-1,m}$), left ($c_{n,m-1}$), and upper-left ($c_{n-1,m-1}$) neighbours. Applying the DP method to the Levenshtein distance computation consists in building column-wise (or row-wise) the matrix \mathbf{C} , element by element, by solving iteratively the same simple problem, whose result will be used as input of one of the successive iterations. Although the algorithm is $\mathcal{O}(l_p \cdot l)$ in time, the space complexity is only $\mathcal{O}(\min(l_p, l))$ because only the previous column (or row) has to be stored to compute the new one.

Applying ASM between \mathbf{p} and \mathbf{t} consists in finding the set of all the substrings $\mathbf{s} \subset \mathbf{t}$ with length $l_s \in [l_p - K, l_p + K]$ that satisfy the condition

$$\text{lev}(\mathbf{p}, \mathbf{s}) = k \leq K, \quad (2)$$

where the threshold K represents the maximum acceptable level of inexactness [64]. Fig. 1(a) depicts the




Figure 1: Example of the matrix \mathbf{C} computation with canonical (a) and wave-front (b) approach.

matrix \mathbf{C} resulting from the computation of the Levenshtein distance between the pattern $\mathbf{p} = \langle ABCD \rangle$ and the substring $\mathbf{s} = \langle ABBDA \rangle$ from the string $\mathbf{t} = \langle ABBDABCDACDB \rangle$ defined over the alphabet $\Sigma = \{A, B, C, D\}$ for $K = 1$ and $i = 0$. The gray-shaded cells in Fig. 1(a) represent the values of the Levenshtein distances between \mathbf{p} and the three substrings $\mathbf{t}[0,3] = \langle ABB \rangle$, $\mathbf{t}[0,4] = \langle ABBD \rangle$, and $\mathbf{t}[0,5] = \langle ABBDA \rangle$. There is no need to compute three different distances since by computing only the distance between \mathbf{p} and $\mathbf{t}[0,5]$ (element $c_{4,5}$ of \mathbf{C}), the other two distances are available as intermediate computations (elements $c_{4,3}$ and $c_{4,4}$ of \mathbf{C} , respectively). By increasing i , step by step, all the substrings of \mathbf{t} are generated.

The DP method can be applied to compute subsequently the Levenshtein distances between \mathbf{p} and all the substrings obtained by shifting a fixed mask of length $l_p + K$ over \mathbf{t} [65]. The DP approach reduces the overall number of computations compared to brute force approaches or when the same computation is repeated over and over [66].

The DP approach for the Levenshtein distance computation exploits the relationship between each element $c_{n,m}$ in \mathbf{C} and its three neighbours $c_{n-1,m}$, $c_{n,m-1}$, $c_{n-1,m-1}$. All the elements on an anti-diagonal j can be computed at the same time in a wave-front processing fashion, by combining the information contained in the anti-diagonals $j - 1$ and $j - 2$, together with the information related to the symbols of the two compared strings. Starting from the element $c_{0,0}$ and ending with the element $c_{l_p, l}$, the matrix \mathbf{C} is filled in $2l_p + K - 1$ steps, as shown in Fig. 1(b). On the first step, the anti-diagonals 0 and 1 are combined to obtain the anti-diagonal 2, and so on, up to the step $2l_p + K - 1$, where the anti-diagonals $2l_p + K - 1$ and $2l_p + K$ are combined to obtain the anti-diagonal $2l_p + K + 1$.

4 An Online Search Algorithm Based on Approximate String Matching

This section presents the proposed online algorithm that combines the features of the DP method with a search criterion based on priority rules to find inexact occurrences of a known pattern \mathbf{p} of length l_p within a continuous stream of symbols \mathbf{t} . All the found occurrences are stored in a set \mathcal{S} to be used for further processing. Collecting all the substrings \mathbf{s} of length $l \in [l_p - K, l_p + K]$ from the text \mathbf{t} that verify Eq. 2 at each position i , may lead to multiple hits of the same occurrence of \mathbf{p} [67]. Due to the online setting considered, it is not possible to decide whether an occurrence can be added to \mathcal{S} at the time it is found, but it depends on the symbols of \mathbf{t} that are not yet processed. For instance, if Eq. 2 is verified for two substrings \mathbf{s}' and \mathbf{s}'' that start from two consecutive positions in \mathbf{t} , respectively i and $i + 1$, the two hits will correspond to the same occurrence of \mathbf{p} . The next subsection introduces a set of rules that determine the priorities for accepting the hits found in the text \mathbf{t} as valid occurrences.

4.1 Priority Rules

Not all the overlapping substrings verifying Eq. 2 should be considered as occurrences but, at the same time, the online nature of the problem complicates deciding on the fly which one should be kept or discarded. The proposed solution consists in assigning temporary priority values to the occurrences that verify Eq. 2. In the remainder of the paper, smaller values denote higher priorities and they depend on the degree of matching between \mathbf{p} and \mathbf{s} , and on the position of the substrings in \mathbf{t} .

Three rules define a priority scheme that allows to build a greedy online algorithm with low complexity to solve the actual NP-complete problem.

R1: Consider that the substring $\mathbf{s}' = \mathbf{t}[i', l']$ for which $\text{lev}(\mathbf{p}, \mathbf{s}') = k' \leq K$ has already been found and assigned with a priority level k' . Any other overlapping substring $\mathbf{s}'' = \mathbf{t}[i'', l'']$ for which $\text{lev}(\mathbf{p}, \mathbf{s}'') = k'' \leq K$ becomes an occurrence only if $k'' < k'$.

R2: If \mathbf{s}' and \mathbf{s}'' have same priority, the first encountered substring becomes an occurrence, the other one is discarded.

R3: If \mathbf{s}' and \mathbf{s}'' have same priority and start at the same position in \mathbf{t} , the shorter becomes an occurrence.

R1:
if $\mathbf{s}' \cap \mathbf{s}'' \neq \emptyset$ **and** $(k'' < k')$ **then**
 \mathbf{s}'' is an occurrence with priority k''
else
 \mathbf{s}'' is discarded

R2:
if $\mathbf{s}' \cap \mathbf{s}'' \neq \emptyset$ **and** $k' = k''$ **then**
 if $i' < i''$ **then**
 \mathbf{s}' is an occurrence, \mathbf{s}'' is discarded
 else
 \mathbf{s}' is discarded, \mathbf{s}'' is an occurrence

R3:
if $\mathbf{s}' \cap \mathbf{s}'' \neq \emptyset$ **and** $k' = k''$ **and** $i' = i''$ **then**
 if $l' < l''$ **then**
 \mathbf{s}' is an occurrence, \mathbf{s}'' is discarded
 else
 \mathbf{s}' is discarded, \mathbf{s}'' is an occurrence

4.2 Validation Process

It is possible to notice that the three rules do not resolve completely the prioritisation of the occurrences in presence of overlaps. For example, consider three overlapping substrings \mathbf{s}' , \mathbf{s}'' , and \mathbf{s}''' , so that

- $i' < i'' < i'''$;
- $k' > k'' > k'''$;

- $s' \cap s'' \neq 0, s'' \cap s''' \neq 0, s' \cap s''' = 0$

As \mathbf{t} is processed symbol by symbol, priority $k' = 2$ is assigned to the occurrence relative to s' . Then, the substring s'' becomes an occurrence with priority $k'' = 1$ verifying rule R1. As more symbols are processed, s''' is found and according to rule R1 it becomes an occurrence with priority $k''' = 0$. Due to the condition $s' \cap s''' = 0$, only the occurrence relative to s'' should be discarded. To ensure this behaviour it is necessary to introduce a further validation procedure, which decides when it is possible to safely discard an occurrence in presence of overlapping substrings.

The validation process proposed here uses one counter $r(k)$ per priority level:

- when an occurrence with priority $k < K$ is found, its validation begins and $r(k)$ starts counting from 0 up to a value equal to the occurrence length;
- if overlapping occurrences with higher priority $k^* < k$ are encountered, $r(k)$ keeps counting until the end of the occurrences.

When the validation process of the occurrence with the highest priority is completed (its relative counter $r(k)$ reaches the target value), it can be decided which of the validated occurrences can be added to \mathcal{S} . Starting from the occurrence with the highest priority and going down to the validated occurrence with the lowest priority, the substring relative to the highest priority is added to \mathcal{S} . Moving downwards, the validated occurrence with lower priority k is added to \mathcal{S} if the following condition is verified

$$r(k) - l(k^*) > l(k) \quad (3)$$

where $k^* < k$ represents the index (priority) of the occurrence relative to the last substring added to \mathcal{S} .

The pseudo-code in Alg. 1 summarises the whole OASM procedure. The storage variable **mem** is a matrix of size $(K + 1) \times 3$, whose content is reset by calling “reset(**mem**)”. **mem** stores in his columns three quantities for each one of the $K + 1$ possible priority levels of the occurrence: its position i (column 1), its length l (column 2), and its associated validation counter r (column 3). The variable idx contains the priority of last found occurrence and the variable ins is a flag that enables the counting check in the validation process. The counters $r(k)$ of the occurrences in **mem** are incremented by one each time a new symbol of \mathbf{t} is processed. Finally, the variable acc is used to implement Eq. 3 and accumulates all the lengths l of the validated occurrence that has just been added to \mathcal{S} . Every time substrings are added to \mathcal{S} , the variables idx , ins , acc , and **mem** are reset. The symbol “:” selects all the indexes across one dimension.

Algorithm 1 Online Approximate String Matching

Input: pattern \mathbf{p} , text \mathbf{t} , threshold K

Output: set of occurrences \mathcal{S}


```

1:  $l_p = |\mathbf{p}|, i = 0, \mathcal{S} = \emptyset, idx = K, ins = 0, acc = 0, \text{reset}(\mathbf{mem})$ 
2: while true do
3:   for  $l : l_p - K \rightarrow l_p + K$  do
4:      $\mathbf{s} = \mathbf{t}(i; l)$ 
5:      $k = \text{lev}(\mathbf{p}, \mathbf{s})$ 
6:     if R1 is TRUE then
7:        $idx = k, ins = 1, \mathbf{mem}[idx][:] = [i, l, 1]$ 
8:     else if R2 is TRUE or R3 is TRUE then
9:        $ins = 1, \mathbf{mem}[idx][:] = [i, l, 1]$ 
10:    if  $ins == 1$  and  $\mathbf{mem}[:,3] == 0$  then
11:      for  $j : idx \rightarrow K$  do
12:        if  $j == idx$  or  $\mathbf{mem}[j][3] - acc < \mathbf{mem}[j][2]$  then
13:           $\mathcal{S} = \mathcal{S} \cup \mathbf{t}(\mathbf{mem}[j][1]; \mathbf{mem}[j][2])$ 
14:           $acc = \mathbf{mem}[j][2] + acc$ 
15:         $idx = K, ins = 0, acc = 0, \text{reset}(\mathbf{mem})$ 
16:    if  $\mathbf{mem}[idx][3] \neq \mathbf{mem}[idx][2]$  then
17:      for  $j : idx \rightarrow K$  do
18:         $\mathbf{mem}[j][3] = \mathbf{mem}[j][3] + 1$ 
19:     $i = i + 1$ 

```

Example The OASM algorithm is executed on the following example, where $\mathbf{p} = \langle ACBDA \rangle$ and $\mathbf{t} = \langle CCCDACCBDAACBDAA \dots \rangle$ are defined over the alphabet $\Sigma = \{A, B, C, D\}$ and $K = 2$.

Fig. 2 illustrates the evolution through time of the variables stored within **mem**: three occurrences (one per priority) relative to different substrings are overlapping. As the counter of the occurrence



index_stream	k = 0	k = 1	k = 2
0	---	---	---
1	---	---	i = 1, l = 5, r = 1
2	---	---	i = 2, l = 4, r = 1
3	---	---	i = 3, l = 3, r = 1
4	---	---	i = 3, l = 3, r = 2
5	---	i = 5, l = 6, r = 1	i = 3, l = 3, r = 3
6	---	i = 6, l = 5, r = 1	i = 3, l = 3, r = 4
7	---	i = 7, l = 4, r = 1	i = 3, l = 3, r = 5
8	---	i = 7, l = 4, r = 2	i = 3, l = 3, r = 6
9	---	i = 7, l = 4, r = 3	i = 3, l = 3, r = 7
10	i = 10, l = 5, r = 1	i = 7, l = 4, r = 4	i = 3, l = 3, r = 8
11	i = 10, l = 5, r = 2	i = 7, l = 4, r = 5	i = 3, l = 3, r = 9
12	i = 10, l = 5, r = 3	i = 7, l = 4, r = 6	i = 3, l = 3, r = 10
13	i = 10, l = 5, r = 4	i = 7, l = 4, r = 7	i = 3, l = 3, r = 11
14	i = 10, l = 5, r = 5	i = 7, l = 4, r = 8	i = 3, l = 3, r = 12
15	---	---	---

Figure 2: Content of **mem** over time, in the proposed example

with priority 0 reaches 5 the validation process is complete. The substring $s = t[10; 5] = \langle ACBDA \rangle$ is the one with highest priority and it represents the occurrence that is validated. For the substring $s = t[7, 4] = \langle CBDA \rangle$, instead, Eq. 3 has to be evaluated. Substituting numerical values in $r(1) - l(0) > l(1)$ gives $8 - 5 > 4$, which is false, so the occurrence with priority 1 is discarded. Finally the substring $s = t[3, 3] = \langle CDA \rangle$ relative to the occurrence with priority 2 has to be added because evaluating $r(2) - l(0) > l(2)$ gives $12 - 5 > 3$, which is true. In the last evaluation $l(0)$ is present rather than $l(1)$, since the last validated occurrence is the one with priority 0, whereas the one with priority 1 was discarded.

4.3 Complexity Analysis

The price to pay for filtering out unaccepted occurrences with the proposed validation process is that the validated occurrences are not instantaneously stored in \mathcal{S} .

To analyze the complexity of the algorithm, we consider the worst case scenario for a generic occurrence with priority k . It occurs when (i) overlapping occurrences with all priority values from 0 to K are

found, (ii) all the occurrences have maximum length, and (iii) each occurrence with priority $k - 1$ begins just before the end of the occurrence with priority k . In this case, the substring corresponding to the occurrence with lowest priority k_{min} cannot be added to \mathcal{S} before $d = l_p(k_{min} + 1) + k_{min}(k_{min} - 1)/2$ symbols of \mathbf{t} are processed.

The space complexity corresponds to the size of the storage variable **mem**. According to the discussed worst case scenario, the upperbound for the space required to store **mem** is $\mathcal{O}\left(l_p(K + 1) + \frac{K(K-1)}{2}(2K + 1)\right)$.

5 Case Study

microRNAs (*miRNAs*) [68, 69] are small non-coding primary transcribed molecules (approx. 22 nucleotides long) which serve as regulators of gene expression and are essential components of normal organism development. Found in plants, animals and unicellular eukaryotes, miRNAs control diverse biological functions by promoting degradation or translation inhibition of target messenger RNAs or by carrying out post-transcriptional regulation of gene expression. miRNA genes are dispersed in various genomic locations (intronic, exonic or intergenic regions) and can be transcribed independently or as a part of other host genes.

miRNAs are encoded within the genome and are often transcribed by RNA polymerase II or, rarely, by RNA polymerase III as long precursor transcripts (several hundreds or thousands of nucleotides), named primary-microRNAs (*pri-miRNAs*) [70, 71]. Those pri-miRNAs have the characteristic hairpin (or stem-loop) structure.

Within the nucleus, pri-miRNAs are cleaved by the microprocessor complex formed by proteins DROSHA and DGCR8 (also known as PASHA¹) and the shorter hairpin structure (approx. 72 nucleotides) formed by this cleavage is referred to as precursor-microRNAs (*pre-miRNAs*).

The pre-miRNAs are transported from the nucleus to the cytoplasm by Exportin-5 together with Ran-Guanine TriPhosphatase [69]. In the cytoplasm, DICER, a cytoplasmic RNase III type protein, dices the transported pre-miRs near the hairpin loop, yielding a mature miRNA-duplex of approx. 22 nucleotides. The duplex is loaded onto the Argonaute protein (*AGO*). One strand of the duplex (the *passenger strand*) is discarded, while the other strand (the *guide strand* or *mature miRNA*) remains in *AGO* to form an RNA-Induced Silencing Complex [69].

The end positions of mature miRNAs are generally assumed as DROSHA cleavage sites. However, pre- and mature-miRNAs are often subject to end modification/processing such as trimming and tailing, which hinders the exact identification of DROSHA cleavage sites [72]. Processing of pri-miRNA stem-loops by the DROSHA/DGCR8 complex is the initial step in miRNA maturation and crucial for its function. Nonetheless, the underlying mechanism that determines the DROSHA cleavage site of primary transcripts has remained unclear: while the mechanisms of pre-miRNA recognition and cleavage by DICER are well characterized [73], an understanding of how DROSHA/DGCR8 selectively recognizes and precisely cleaves pri-miRNAs remains unclear [74]. The problem of identifying cleavage sites for DROSHA relies on the fact that failures at DROSHA processing step might lead to miRNAs down-regulation, a phenomenon observed in cancer patients [75].

5.1 Baseline Algorithms for ASM

To show the effectiveness of the proposed OASM algorithm, it is compared to two baseline ASM techniques.

The first, referred as *Fully Naïve*, works in a brute force fashion [76] according to the following three steps:

1. from the input stream, extracts all possible adjacent substrings of any possible length;
2. evaluates all pairwise Levenshtein distances between substrings and the target sequence;
3. discards all substrings whose distance with respect to the target sequence is greater than the threshold K .

¹Partner of droSHA

Whilst being the most straightforward method, Fully Naïve does not consider multiple shadow hits and all the occurrences found must be stored in order to be filtered in a post-processing stage. Further, due to its brute force nature, it is unsuitable for processing large input streams, since the number of possible adjacent substrings of any possible length grows quadratically with the length of the text. Specifically, let l be the length of the text, the number of non-empty adjacent substrings of any possible length is $l(l+1)/2$. Accordingly, the time and space complexity for evaluating the Levenshtein distance with Wagner-Fischer’s algorithm, as described in Sect. 3, is $\mathcal{O}(l \cdot l_p)$ with l and l_p being the input and target string lengths, respectively. Thus, the overall complexity in the best case is $\mathcal{O}(l_p)$, whereas in the worst case is $\mathcal{O}(l^2)$. Finally, the filtering procedure must scan all possible substrings generated in step 1, and its complexity is again $\mathcal{O}(l^2)$. By strictly following the three steps above, step 1 might lead to early out-of-memory errors due to the storage of $\mathcal{O}(l^2)$ possible substrings. To avoid this, for each extracted substring one can evaluate the distance with respect to the target and discard the substring if the distance is above the user-defined threshold.

The second procedure, referred as *Less Naïve*, tweaks the Levenshtein matrix \mathbf{C} [1, 57]. Again, let l and l_p be the length of the input stream and target string, respectively. According to Sect. 3, to keep track of the prefixes one initializes the matrix \mathbf{C} with a dummy row (0 to l) and a dummy column (0 to l_p), as shown in Fig. 1. Instead, if the dummy row is initialized with all zeros, the Wagner-Fischer’s algorithm computes the Levenshtein distance between the target and substrings in the input stream. Therefore, rather than picking the bottom-right element, one takes all positions in the last row that are below the threshold K as the ending positions of the hits.

5.2 Qualitative Analysis of the Results

After consultation with field-experts², several human pre-miRNAs of interest have been selected for the analysis. As discussed above, since the end positions of mature miRNAs cannot be considered as reliable cleavage sites, mining pre-miRNAs neighbourhood regions (both upstream and downstream) can give some further insights to biologists.

To evaluate the capability of the proposed algorithm to filter shadow-hits on the fly, five of the suggested pre-miRNAs have been considered, along with the corresponding chromosomes. All data are freely available on ad-hoc biological online databases: the human miRNAs can be found at miRBase [77], whereas the human chromosomes (assembly GRCh38) can be retrieved from GenomeBrowser [78]. The three algorithms, *Fully Naïve*, *Less Naïve*, and OASM have been evaluated on the same data using the same inexactness threshold $K = 2$. The number of hits and their positions are reported in Table 1 and Fig. 3 visually depicts the amount of overlapping occurrences found in each position. From Table 1

Table 1: Number of hits across 5 different pre-miRNAs by considering 200 nucleotides neighbourhood (both upstream and downstream). The total number of nucleotides is shown in brackets.

Input Stream	Target	Fully Naïve	Less Naïve	OASM
hsa-mir-218-1 ± 200 nt (510 nt)	aaaaaaaa	29	13	4
hsa-mir-515-1 ± 200 nt (483 nt)	gcaacc	64	39	19
hsa-mir-519a-1 ± 200 nt (485 nt)	acgttgca	8	6	4
hsa-mir-105-1 ± 200 nt (481 nt)	aaccttgg	6	6	3
hsa-mir-1-2 ± 200 nt (485 nt)	ctcattca	7	7	5

and Fig. 3 clearly emerges that OASM detects each occurrence accurately just one time since, in contrast to the other two algorithms, it is able to discard those detected sequences that are overlapping. Indeed, neither Fully Naïve nor Less Naïve can deal with shadow hits, thus requiring a mandatory additional post-processing phase to validate the occurrences. Such a post-process results in additional computation and, most importantly, contrarily to the proposed OASM algorithm, the detection cannot be done online.

Tab. 2 summarizes the time complexity of the ASM algorithms. Execution times are not reported since the post-processing filtering operation, which is required when using Fully Naïve and Less Naïve algorithms, is not univocally defined. In fact, it can be implemented in several ways and executed at different times, in accordance to the requirements of the task at hand. A detailed discussion and analysis of the offline filtering operations for the naïve algorithms is beyond the scope of this work.

²Dr. Giulia Piaggio and Dr. Aymone Gurtner. Regina Elena Institute for Cancer Research. Rome, Italy.

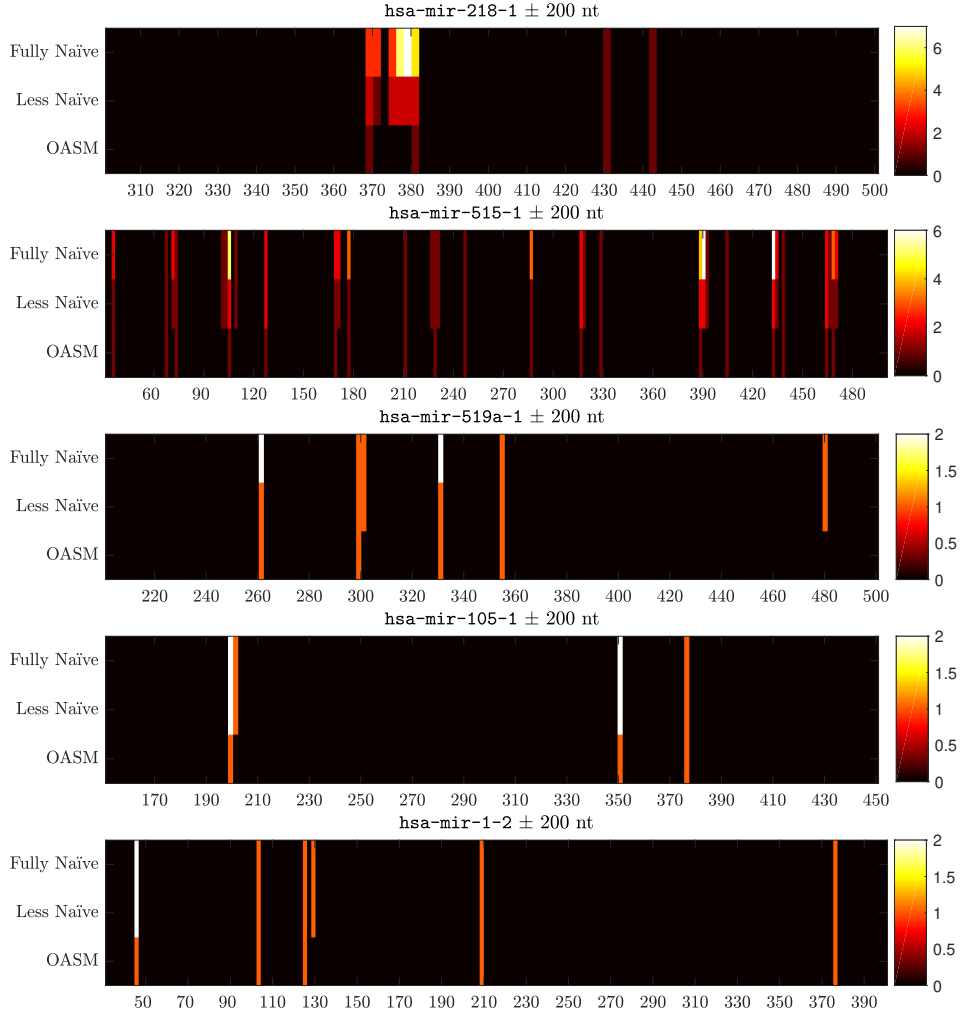


Figure 3: Hits positions for experiments in Table 1. The colour map indicates the number of overlapping occurrences found in a given position.

Table 2: Computational complexities of the ASM algorithms. Herein, let l and l_p be the length of the input stream and the length of the pattern to be search within the input stream, respectively.

Algorithm	Time Complexity
Fully Naïve	$\mathcal{O}(l^2) \cdot [\mathcal{O}(l_p) \rightarrow \mathcal{O}(l^2)] + \mathcal{O}(l^2) + \text{offline filtering cost}$
Less Naïve	$\mathcal{O}(l \cdot l_p) + \text{offline filtering cost}$
OASM	$\mathcal{O}\left(l \cdot \left(l_p(K+1) + \frac{K(K-1)}{2}\right)\right)$

6 Proposed Hardware Implementation

6.1 Top-level instance: *LEV CORE*

The LEV CORE module computes the multiple Levenshtein distances between the *pattern* \mathbf{p} and all the *substrings* \mathbf{s} extracted from a stream of symbols \mathbf{t} and returns only the list of occurrences *result* below a threshold K , according to the algorithm described in Sec. 4. Fig. 4(a) depicts the main I/O ports and the two main sub-modules of the LEV CORE:

- LEV CALC computes at every new *index_stream* value the distances between \mathbf{p} and the $2K + 1$

substrings $\mathbf{s} \in \{\mathbf{t}[i, l_p - K], \dots, \mathbf{t}[i, l_p + K]\}$.

- LEV SEARCH elaborates the received data to search for all the possible occurrences of \mathbf{p} and to validate the found occurrences that satisfy the priority scheme described in Sec. 4, discarding undesired shadow hits.

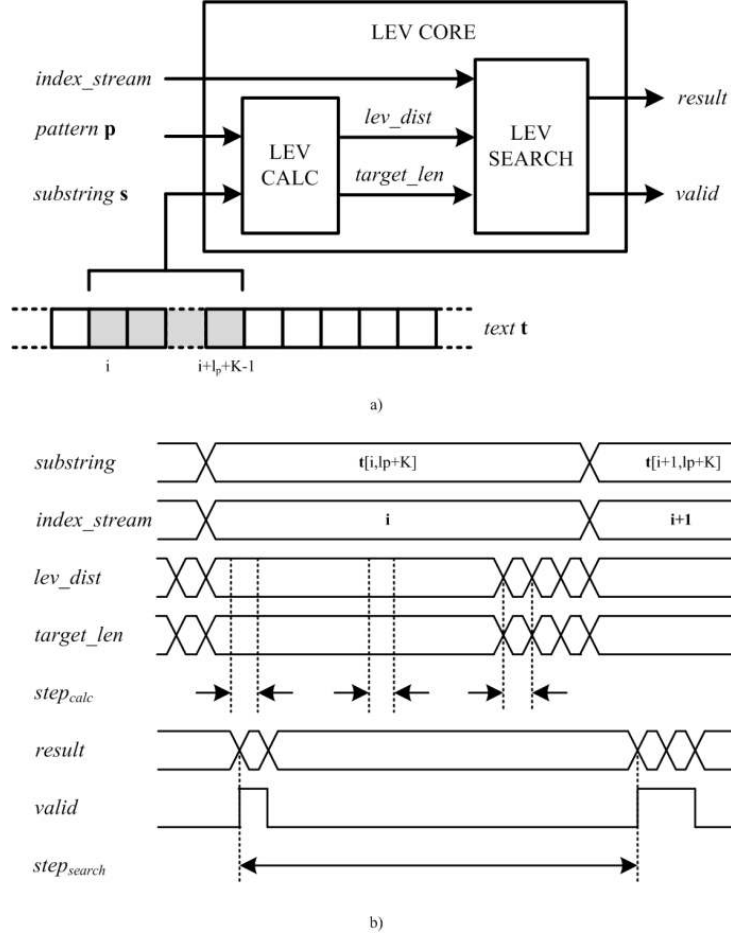


Figure 4: Conceptual block scheme for LEV CORE module (a) and relative timing for the main signals (b).

After the pattern \mathbf{p} is loaded into the proper registers, the elaboration starts as the continuous stream of data \mathbf{t} flows into the module. In Fig. 4(b) a time diagram of the main signals (port-level and internal) is shown.

6.2 Distance computation: *LEV CALC*

The systolic nature of the DP algorithm used to compute the Levenshtein distance is highly parallelizable and particularly suitable for direct hardware implementation. A systolic architecture [79] consists of identical processing elements arranged in an array that process data synchronously, executing a short invariant sequence of instructions without an intervening memory to store and exchange results across the pipelined array. The advantage of this architecture is to make both the space and time of the calculation of the Levenshtein matrix linear by parallelizing independent processes. In the ordinary systolic architectures used in string matching applications [80], two strings are shifted on each other to compare step-by-step each pair of symbols. To compute the distance, each symbol needs to be associated to an additional (stored) information, which is the number that identifies the position of the symbol within the stream \mathbf{t} . The LEV CALC sub-module here introduced implements a counter-based systolic architecture that avoids the storage of *a-priori* known data. It computes the Levenshtein

distances in a wave-front fashion returning all the elements of the current anti-diagonal of \mathbf{C} , filling the matrix in $2l_p + K - 1$ steps, each one from now on referred as $\text{step}_{\text{calc}}$ (Fig. 4(b)). Fig. 5 illustrates the conceptual architecture of two consecutive processing elements of the systolic array, on which the LEV CALC sub-module is based on.




Figure 5: Details of the systolic architecture of LEV CALC module.

The generic j -th processing element contains both sequential and combinational logic. The register $\mathbf{p_reg}(j)$ is configured with the j -th symbol $p^{(j)}$ of the pattern \mathbf{p} (if unused, filled with special symbol $\$1$). The register $\mathbf{sh_reg}(j)$ contains the shifting symbols $\text{sh}(j)$ of the substring $\mathbf{s} = \mathbf{t}[i, l_p + K]$ and at each $\text{step}_{\text{calc}}$, contains a different symbol of \mathbf{s} . The special symbol $\$2$ is used to fill the unused $\mathbf{sh_reg}$ positions during the shifting. The special symbols $\$1$ and $\$2$ belong to the alphabet Σ but cannot be used both in \mathbf{p} and in \mathbf{t} . The register $\mathbf{a_reg}(j)$ stores the result of the j -th processing element produced by $\mathbf{l_comb}(j)$ during the current $\text{step}_{\text{calc}}$. The register $\mathbf{a_reg_d}(j)$ represents a delayed version of $\mathbf{a_reg}(j)$. The combinational block $\mathbf{l_comb}(j)$ elaborates both the information contained in the above mentioned registers and the state of the upward counter cnt to yield the j -th element of the anti-diagonal of the matrix \mathbf{C} .

The counter is used both to allow the shifting process of the symbols $\text{sh}(j)$ through the j -th processing block, and to provide the suitable n or m values of Eq. 1 to $\mathbf{l_comb}(j)$ block when the corresponding processing element represents an element $c_{n,0}$ or $c_{0,m}$ of \mathbf{C} respectively. The j -th element of the anti-diagonal in the current $\text{step}_{\text{calc}}$ implementing a rearranged, yet equivalent, version of Eq. 1 is computed as

$$\mathbf{l_comb}(j) = \min \begin{bmatrix} c_{\text{left}}(j), \\ c_{\text{upper}}(j), \\ c_{\text{upper-left}}(j) \end{bmatrix} + \phi(j), \quad (4)$$

where

$$c_{\text{left}}(j) = \begin{cases} \mathbf{a_reg_d}(j), & \text{if } j < cnt \\ cnt, & \text{otherwise} \end{cases}$$

$$c_{\text{upper}}(j) = \begin{cases} cnt, & \text{if } j = 0 \\ \mathbf{a_reg}(j-1), & \text{if } j > 0 \end{cases}$$

$$c_{\text{upper-left}}(j) = \begin{cases} \mathbf{a_reg_d}(j-1), & \text{if } j < cnt \\ cnt-1, & \text{otherwise} \end{cases}$$

$$\phi(j) = \begin{cases} 1, & \text{if } c_{\text{upper-left}}(j) > \min \begin{bmatrix} c_{\text{left}}(j), \\ c_{\text{upper}}(j), \\ c_{\text{upper-left}}(j) \end{bmatrix} \\ \mathbf{p_reg}(j) \neq \mathbf{sh_reg}(j), & \text{otherwise.} \end{cases}$$

In Fig. 6 the four possible combinations of neighbors for the computation of $\mathbf{l_comb}(j)$ are located in the \mathbf{C} matrix frame to show how Eq. 1 turns into Eq. 4.

	0	1	2		m	$m+1$
0	$cnt-1$	cnt			$cnt-1$	cnt
1	cnt	$\mathbf{l_comb}(j)$			$\mathbf{a_reg_d}(j)$	$\mathbf{l_comb}(j)$
2						
n	$cnt-1$	$\mathbf{a_reg}(j-1)$			$\mathbf{a_reg_d}(j-1)$	$\mathbf{a_reg}(j-1)$
$n+1$	cnt	$\mathbf{l_comb}(j)$			$\mathbf{a_reg_d}(j)$	$\mathbf{l_comb}(j)$

Figure 6: Representation of the neighbors in the \mathbf{C} matrix to compute $\mathbf{l_comb}(j)$.

Finally, the MUX selector sel_calc is constant and uniquely determined with the length of the pattern \mathbf{p} and the value of the threshold K . Fig. 7 shows the content of $\mathbf{p_reg}$ and $\mathbf{sh_reg}$ in the inherent shifting mechanism for the example discussed in Sec. 3.

6.3 Search and validation: *LEV SEARCH*

The outputs generated by LEV CALC, together with $index_stream$, represent the inputs of the LEV SEARCH sub-module whose architecture is schematized in Fig. 8(a). This module implements the online search algorithm described in Sec. 4 and formalised in Alg. 1.

In an infinite loop cycling on i (corresponding to a new $index_stream$) the Levenshtein distances between \mathbf{p} and all the possible $2K + 1$ substrings $\mathbf{s} \in \{\mathbf{t}[i, l_p - K], \dots, \mathbf{t}[i, l_p + K]\}$ returned by LEV CALC, which satisfy both the described priority rules (R1, R2, and R3) and validation process, are stored




Figure 7: Evolution of the content of **sh_reg** step by step to implement the string shifting.

in \mathcal{S} . The logic relative to the various checks performed by the algorithm is located into the ENA GEN sub-block. There, $[i, l, k] = [index_stream, target_len, lev_dist]$ are processed by the priority rules (see Sec. 4), and Eq. 3 is evaluated on the values of the counters $r(k)$ during the validation process. Within the same block, the variables idx , ins , and acc defined in Alg. 1 are implemented as a status register, a flag and an accumulation register, respectively, and are used to generate the signals ena , $preset$, ena_acc , and $valid$. Fig. 8(b) depicts a detail of the sequential section which represents the state of the algorithm (SEARCH ELEMENTS). The two-dimensional array **o_reg** of size $(K + 1) \times 3$ implements the storage variable **mem** in Alg. 1 and it allows to retrieve information relative to the current eligible occurrences. The procedure executed in LEV CALC consists in a constant number of $step_{calc}$, whereas the algorithm implemented in LEV SEARCH may have a different duration depending on the validation process that is executed only if eligible occurrences are found. A LEV SEARCH cycle is called here $step_{search}$ and its duration coincides with number of $step_{calc}$ necessary to complete on LEV CALC cycle multiplied by its duration

$$T_{step_{search}} = (2 \cdot l_p + K - 1) T_{step_{calc}}. \quad (5)$$

The proposed hardware architecture executes one $step_{calc}$ in one clock cycle (T_{clk}), so the LEV CALC block generates the first lev_dist sample after $2l_p - K - 1$ clock cycles and in the same cycle the LEV SEARCH block starts processing it:

$$\begin{aligned} lat_{calc} &= 2l_p + K - 1 \\ lat_{search} &= \begin{cases} K + 1 & \text{without validated occurrences} \\ 2K + 1 & \text{with validated occurrences} \end{cases} \end{aligned}$$

Although the latency of the LEV SEARCH is data dependent, the operations of the two modules can be pipelined. The total execution time depends on the latency of the two blocks:

$$T_{exec} \approx T_{step_{search}} l_t = (2l_p + K - 1) \cdot T_{clk} l_t \quad (6)$$

where T_{clk} is the clock period. It should be noticed that Eq. 6 is valid when $K < 2(l_p - 1)$, which is always verified since the threshold K is always lower than l_p . Fig. 9 depicts the pipelining process.

6.4 Resource usage and parallel search

The LEV CORE has been completely described in VHDL and parameterized in terms of

- l_{p_max} : maximum number of symbols per pattern;
- l_{symb} : number of bits per symbol;
- K : maximum threshold.

Tab. 3 reports synthesis results in terms of Logic Elements (LE[#]) used, when varying the design parameters. The largest synthesized LEV CORE modules ($l_{symb} = 16$, $l_{p_max} = 32$, $K = 4$) uses only

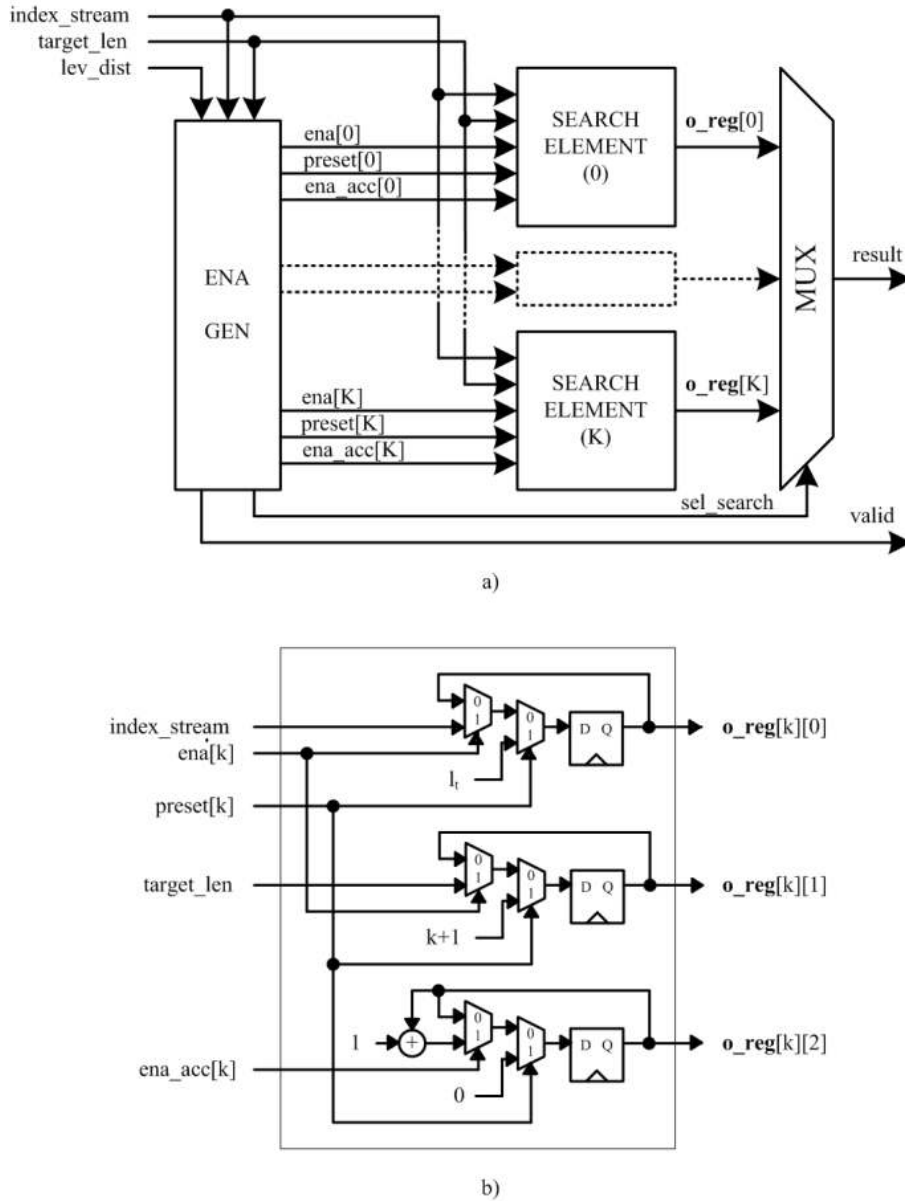


Figure 8: Conceptual scheme of the LEV SEARCH architecture (a) and detail of the k -th SEARCH ELEMENT sub-block (b)

3% of the total available resources on Altera Cyclone IV E FPGA (114800 LE). Fig. 10 depicts the results of Tab. 3, providing a clearer view of the resource usage as the design parameters vary. The planar shape of the surfaces and the lack of intersections shows the linear relation between resource usage and design parameters.

A parallel search of different patterns over the same text \mathbf{t} can be done efficiently, by deploying multiple instances of LEV CORE. This is implemented by the Multiple OASM (MOASM) system, whose architecture is schematized in Fig. 11. The synthesis of a MOASM system with 40 LEV CORE modules ($l_{\text{syms}} = 8$, $l_{\text{p_max}} = 16$, $K = 4$) and additional control logic resulted in 62936 LE, which corresponds to the 55% of the resource usage on the same target FPGA.

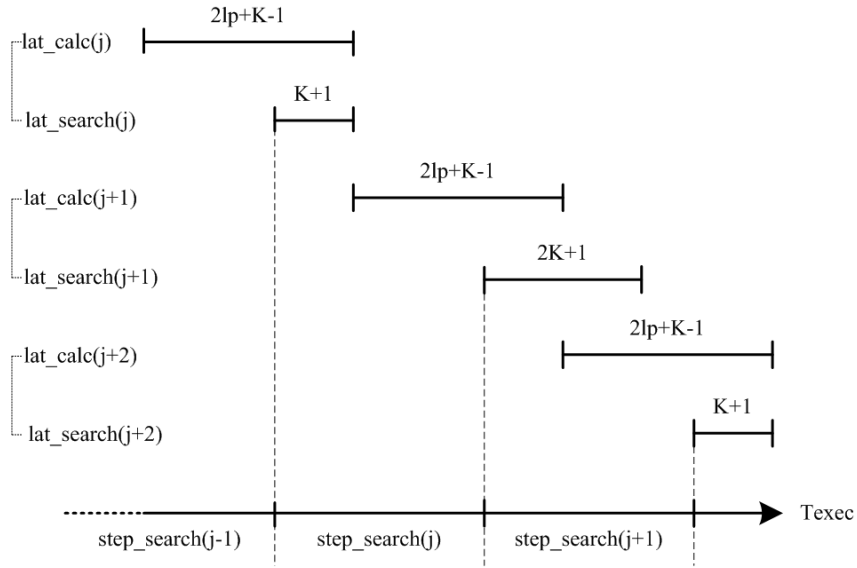


Figure 9: Time execution diagram in the pipelined structure

Table 3: LE utilization relationship with design parameters

l_{symb}	$l_{\text{p_max}}$	K	LE[#]	l_{symb}	$l_{\text{p_max}}$	K	LE[#]	
4	8	1	557	8	8	1	626	
		2	732			2	799	
		3	822			3	895	
		4	880			4	953	
	16	1	1122		16	1	1266	
		2	1286			2	1443	
		3	1391			3	1536	
		4	1472			4	1610	
	24	24	1		1687	24	1	1936
			2		1840		2	2025
			3		1960		3	2151
			4		2064		4	2217
	32	32	1		2252	32	1	2530
			2		2394		2	2869
			3		2529		3	3002
			4		2656		4	3068
l_{symb}	$l_{\text{p_max}}$	K	LE[#]	l_{symb}	$l_{\text{p_max}}$	K	LE[#]	
12	8	1	701	16	8	1	765	
		2	867			2	940	
		3	962			3	1031	
		4	1019			4	1090	
	16	16	1		1403	16	1	1543
			2		1579		2	1706
			3		1670		3	1810
			4		1738		4	1865
	24	24	1		2165	24	1	2349
			2		2233		2	2438
			3		2352		3	2547
			4		2402		4	2616
	32	32	1		2790	32	1	3067
			2		3163		2	3428
			3		3269		3	3537
			4		3359		4	3616




Figure 10: Linear relationship between design parameters and number of LEs

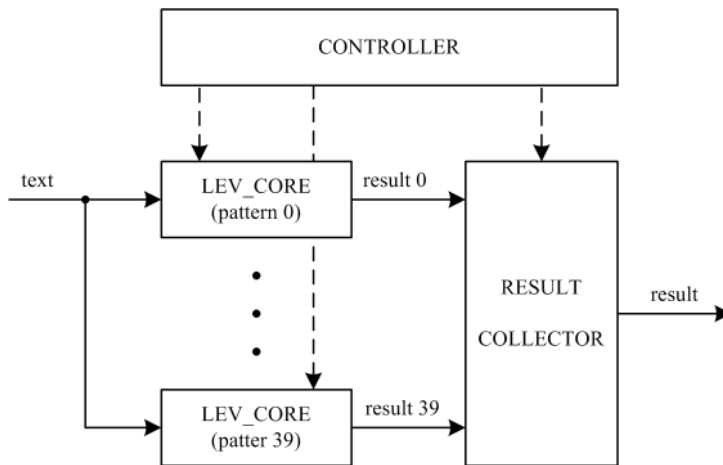


Figure 11: Conceptual scheme of a MOASM system

7 Experiments

In this section are reported experiments on synthetic data in order to compare the speed performances of the software (SW-OASM) and hardware (HW-OASM) implementation of the proposed OASM algorithm. SW-OASM has been implemented in C++ and executed on a CPU Intel Core i7-4700MQ @ 2.40 GHz. HW-OASM has been described in VHDL-1993 and implemented on the FPGA Altera Cyclone IV E mounted on the board Terasic DE2-115. Two tests are performed using the same setup with randomly generated sequences \mathbf{t} of $l_t = 3104$ symbols defined over the alphabet $\Sigma = \{A, B, C, D\}$, and randomly generated patterns \mathbf{p} that are processed by both the implementations, for different values of patterns length l_p and threshold K .

Recent FPGAs commonly host a high speed link (e.g. PCIe 2.0 with 5 GT/s and the per-lane throughput 500 MB/s) so, to compensate for the lack of a high speed link on the available FPGA, a simple link emulator for the online input data-transfer has been designed. This solution can be considered acceptable since the purpose of the tests is measuring the processing time of the LEV CORE block. Since, two special symbols $\$_1$ and $\$_2$ are added to the alphabet Σ , each symbol in \mathbf{p} and \mathbf{t} is coded with $l_{\text{symb}} = 3$ bits. The text \mathbf{t} is stored into the FPGA in 1 Mbit embedded memory and is sent, symbol by symbol, to the processing logic that operates at 100 MHz (300 Mbit/s). The resulting data flow is stored into another 1 Mbit embedded memory element and is sent to a PC by a low-speed USB link, to be analyzed once the test is completed. Fig. 12 schematises the HW-OASM architecture implemented for the tests, whose details are the following.

- LEV CORE: processing block with hardwired pattern \mathbf{p} , maximum pattern length $l_p = 15$, alphabet size = 3 bit (extended due to special symbols), maximum threshold $K = 5$.
- LINK EMULATOR: 1 Mbit ROM, containing the text \mathbf{t} (5 symbols per memory location) plus control logic. It sends a new symbol to the LEV CORE block at every new incoming request pulse.
- RAM: one 1 Mbit RAM to dump the results (1 Mbit), 16 bit index (sufficient to recognise an occurrence), 3 bit threshold, 5 bit length (max. detected length is $l_p + K = 20$).
- USB I/F: USB device interface to transfer the dumped results to an external PC.
- SYS FSM: finite state machine to coordinate the I/O operations. Main control signals are *start_elab* (start of elaboration) and *result_return* (command to start returning the data stored in RAM via USB I/F).




Figure 12: System implemented for executing the experiments

The wall-clock time of the SW-OASM execution is not deterministic, since there are usually other processes running in the background that consume computational resources. Therefore, each test run with SW-OASM for a specific configuration has been repeated 100 times with independently randomly generated datasets and the results are reported as mean and standard deviation. Regarding the tests run with HW-OASM, the results prove the effectiveness of the linear deterministic model for the execution time computation described in Eq. 6 with $T_{\text{clk}} = 10\text{ns}$. The experimental results slightly depart from the theoretical ones because of implementation details that prevent $T_{\text{step_search}}$ in Eq. 5 to be just dependent on l_p and K . In Fig. 13 the speed performance are reported, when different thresholds K are used.

Two different tests are performed. In the first, whose results are in Tab. 4, random patterns with lengths $l_p = \{5, 7, 10, 15\}$ are generated and a threshold $K = 3$ is used.

Table 4: Speed performance comparison between SW-OASM and HW-OASM implementation for the first test

l_p	SW-OASM [s]	HW-OASM [s]
5	46.51±0.74	0.0039
7	59.991±1.42	0.0051
10	78.973±3.54	0.0069
15	97.004±1.76	0.0099

In the second test, with results reported in Tab. 5, random patterns with fixed length $l_p = 5$ are generated and different thresholds $K = \{2, 3, 4, 5\}$ are used.

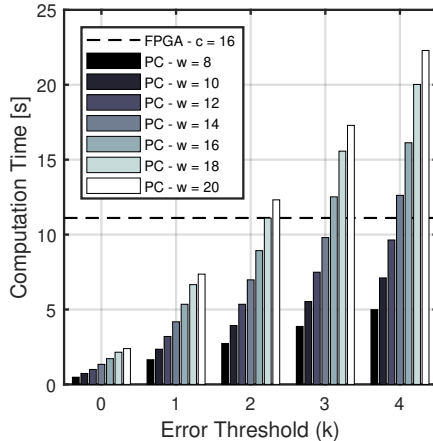


Figure 13: SW-OASM computation time for different word length (w) and error threshold (k) pairs. HW-OASM computation time (dashed) is constant on the same set of parameters used for SW-OASM

Table 5: Speed performance comparison between SW-OASM and HW-OASM implementation for the second test

K	SW-OASM [s]	HW-OASM [s]
2	9.731 ± 0.50	0.0096
3	15.179 ± 0.45	0.0099
4	30.858 ± 2.90	0.0102
5	67.953 ± 3.14	0.0105

8 Conclusions

Approximate String Matching (ASM) is a particular case of subgraph matching and is a fundamental tool in many practical applications, such as bioinformatics, cybersecurity, diagnostic systems and financial trading. However, when using a tolerance threshold, the ASM can produce shadow hits, i.e. it can identify close multiple hits that, in turn, introduce unwanted uncertainty in the overall template matching procedure and thus in the knowledge discovery task. Of course, the shadow hits could be filtered out offline afterwards in a post-processing step. However, when dealing with online applications an alternative approach must be followed.

This paper introduces the online ASM (OASM) algorithm to filter out shadow hits online, which is based on a validation procedure that takes advantage of side information provided during Levenshtein distance computations. OASM is designed in plain dynamic programming and is characterized by a high degree of parallelism, which can be leveraged to design an efficient FPGA implementation. The proposed architecture allows to perform such a complex retrieval procedure relying on inexpensive hardware, such as the adopted entry level Altera Cyclone IV E FPGA, where only 3% of available logic elements are used. This envisage further opportunities to increase the parallelism of the whole procedure, by allocating on the FPGA more LEV CORE units to concurrently search for instances of multiple motifs on the same sequence, or even on different sequences. The execution time of the hardware implementation of such an enhanced implementation (MOASM system) is independent of the number of the instantiated LEV CORE modules, whereas the computational time in the software counterpart scales linearly with the number of pattern searched. Especially when the MOASM system is implemented on high-end FPGA chips that allow the instantiation of many LEV CORE modules, the I/O interface (usually a PCI-Express bus) may easily become congested, due to the possible huge number of results conveyed on a single communication bus. This underlines the importance of the proposed approach to filter out unwanted shadow hits directly on the FPGA, which can greatly reduce the overall I/O throughput. An important future work regards the implementation of a complete MOASM system and the transfer of the Intellectual Property on a hi-performance FPGA, such as the Stratix V GX.

References

- [1] Gonzalo Navarro. A guided tour to approximate string matching. *ACM computing surveys (CSUR)*, 33(1):31–88, 2001.
- [2] Lorenzo Livi and Antonello Rizzi. The graph matching problem. *Pattern Analysis and Applications*, 16(3):253–283, 2013.
- [3] Ha-Nguyen Tran, Erik Cambria, and Amir Hussain. Towards gpu-based common-sense reasoning: Using fast subgraph matching. *Cognitive Computation*, 8(6):1074–1086, Dec 2016. ISSN 1866-9964. doi: 10.1007/s12559-016-9418-4.
- [4] Jeremy Buhler and Martin Tompa. Finding motifs using random projections. *Journal of computational biology*, 9(2):225–242, 2002.
- [5] Eleazar Eskin and Pavel A Pevzner. Finding composite regulatory patterns in dna sequences. *Bioinformatics*, 18(suppl 1):S354–S363, 2002.
- [6] Giulio Pavesi, Paolo Mereghetti, Giancarlo Mauri, and Graziano Pesole. Weeder web: discovery of transcription factor binding sites in a set of sequences from co-regulated genes. *Nucleic acids research*, 32(suppl 2):W199–W203, 2004.
- [7] Saurabh Sinha and Martin Tompa. Ymf: a program for discovery of novel transcription factor binding sites by statistical overrepresentation. *Nucleic acids research*, 31(13):3586–3588, 2003.
- [8] Rainer Typke, Frans Wiering, Remco C Veltkamp, et al. A survey of music information retrieval systems. In *ISMIR*, pages 153–160, 2005.
- [9] Marco Bertini, Alberto Del Bimbo, and Walter Nunziati. Video clip matching using mpeg-7 descriptors and edit distance. In *Image and video retrieval*, pages 133–142. Springer, 2006.
- [10] Nivio Ziviani, Edleno Silva De Moura, Gonzalo Navarro, and Ricardo Baeza-Yates. Compression: A key for next-generation text retrieval systems. *Computer*, 33(11):37–44, 2000.
- [11] Abdelhak Boukharouba and Abdelhak Bennis. Recognition of handwritten arabic literal amounts using a hybrid approach. *Cognitive Computation*, 3(2):382–393, Jun 2011. ISSN 1866-9964. doi: 10.1007/s12559-010-9088-6.
- [12] Mansi Sahi and Vishal Gupta. A novel technique for detecting plagiarism in documents exploiting information sources. *Cognitive Computation*, 9(6):852–867, Dec 2017. ISSN 1866-9964. doi: 10.1007/s12559-017-9502-4.
- [13] Luis Gravano, Panagiotis G Ipeirotis, Nick Koudas, and Divesh Srivastava. Text joins in an rdbms for web data integration. In *Proceedings of the 12th international conference on World Wide Web*, pages 90–101. ACM, 2003.
- [14] Enrico Maiorino, Francesca Possemato, Valerio Modugno, and Antonello Rizzi. Noise sensitivity of an information granules filtering procedure by genetic optimization for inexact sequential pattern mining. In *Computational Intelligence*, pages 131–150. Springer, 2016.
- [15] Yi-Yu Yao. The rise of granular computing. *Journal of Chongqing University of Posts and Telecommunications (Natural Science Edition)*, 20(3):299–308, 2008.
- [16] Newton Howard and Henry Lieberman. Brainspace: Relating neuroscience to knowledge about everyday life. *Cognitive Computation*, 6(1):35–44, Mar 2014.
- [17] Andrzej Bargiela and Witold Pedrycz. Granular computing. In *HANDBOOK ON COMPUTATIONAL INTELLIGENCE: Volume 1: Fuzzy Logic, Systems, Artificial Neural Networks, and Learning Systems*, pages 43–66. World Scientific, 2016.
- [18] Yiyu Yao. A triarchic theory of granular computing. *Granular Computing*, 1(2):145–157, 2016.

- [19] Prem Kumar Singh. Similar vague concepts selection using their euclidean distance at different granulation. *Cognitive Computation*, 10(2):228–241, Apr 2018.
- [20] Tsau Young Lin, Yiyu Y Yao, and Lotfi A Zadeh. *Data mining, rough sets and granular computing*, volume 95. Physica, 2013.
- [21] Filippo Maria Bianchi, Lorenzo Livi, Antonello Rizzi, and Alireza Sadeghian. A granular computing approach to the design of optimized graph classification systems. *Soft Computing*, 18(2):393–412, 2014.
- [22] Antonello Rizzi, Guido Del Vescovo, Lorenzo Livi, and Fabio Massimo Frattale Mascioli. A new granular computing approach for sequences representation and classification. In *Neural Networks (IJCNN), The 2012 International Joint Conference on*, pages 1–8. IEEE, 2012.
- [23] Filippo Maria Bianchi, Simone Scardapane, Antonello Rizzi, Aurelio Uncini, and Alireza Sadeghian. Granular computing techniques for classification and semantic characterization of structured data. *Cognitive Computation*, 8(3):442–461, Jun 2016. ISSN 1866-9964. doi: 10.1007/s12559-015-9369-1.
- [24] Alessio Martino, Alessandro Giuliani, and Antonello Rizzi. Granular computing techniques for bioinformatics pattern recognition problems in non-metric spaces. In Witold Pedrycz and Shyi-Ming Chen, editors, *Computational Intelligence for Pattern Recognition*, pages 53–81. Springer International Publishing, Cham, 2018. ISBN 978-3-319-89629-8. doi: 10.1007/978-3-319-89629-8_3.
- [25] Alexandr Andoni, Robert Krauthgamer, and Krzysztof Onak. Polylogarithmic approximation for edit distance and the asymmetric query complexity. In *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on*, pages 377–386. IEEE, 2010.
- [26] Leonid Boytsov. Indexing methods for approximate dictionary searching: Comparative analysis. *Journal of Experimental Algorithmics (JEA)*, 16:1–1, 2011.
- [27] Roberto Di Pietro and Luigi V Mancini. *Intrusion detection systems*, volume 38. Springer Science & Business Media, 2008.
- [28] Christian A Heid, Junko Stevens, Kenneth J Livak, and P Mickey Williams. Real time quantitative pcr. *Genome research*, 6(10):986–994, 1996.
- [29] MJ Espy, JR Uhl, LM Sloan, SP Buckwalter, MF Jones, EA Vetter, JDC Yao, NL Wengenack, JE Rosenblatt, FR 3 Cockerill, et al. Real-time pcr in clinical microbiology: applications for routine laboratory testing. *Clinical microbiology reviews*, 19(1):165–256, 2006.
- [30] Maria-Bernadette Madel, Harald Niederstätter, and Walther Parson. Trixy—homogeneous genetic sexing of highly degraded forensic samples including hair shafts. *Forensic Science International: Genetics*, 25:166 – 174, 2016. ISSN 1872-4973. doi: 10.1016/j.fsigen.2016.09.001.
- [31] H. Niederstätter, M.D. Coble, T.J. Parsons, and W. Parson. Characterization of mtdna snp typing using quantitative real-time pcr with special emphasis on heteroplasmy detection and mixture ratio assessment. *International Congress Series*, 1288:1 – 3, 2006. ISSN 0531-5131. doi: 10.1016/j.ics.2005.09.021. Progress in Forensic Genetics 11.
- [32] Akhtar Rasool and Nilay Khare. Parallelization of kmp string matching algorithm on different simd architectures: Multi-core and gpgpu’s. *International Journal of Computer Applications*, 49(11): 26–28, 2012.
- [33] Cheng Zhong and Guo-Liang Chen. A fast determinate string matching algorithm for the network intrusion detection systems. In *Machine Learning and Cybernetics, 2007 International Conference on*, volume 6, pages 3173–3177. IEEE, 2007.
- [34] Maxime Crochemore, Costas S Iliopoulos, Yoan J Pinzon, and James F Reid. A fast and practical bit-vector algorithm for the longest common subsequence problem. *Information Processing Letters*, 80(6):279–285, 2001.

- [35] F Thomson Leighton. *Introduction to parallel algorithms and architectures: Arrays, trees, hypercubes*. Elsevier, 2014.
- [36] Panagiotis D Michailidis and Konstantinos G Margaritis. A programmable array processor architecture for flexible approximate string matching algorithms. In *Parallel Processing, 2005. ICPP 2005 Workshops. International Conference Workshops on*, pages 201–209. IEEE, 2005.
- [37] Piotr Antonik, Marc Haelterman, and Serge Massar. Online training for high-performance analogue readout layers in photonic reservoir computers. *Cognitive Computation*, 9(3):297–306, Jun 2017. ISSN 1866-9964. doi: 10.1007/s12559-017-9459-3.
- [38] José L. Vásquez, Santiago T. Pérez, Carlos M. Travieso, and Jesús B. Alonso. Meteorological prediction implemented on field-programmable gate array. *Cognitive Computation*, 5(4):551–557, Dec 2013. ISSN 1866-9964. doi: 10.1007/s12559-012-9158-z.
- [39] Sadatoshi Mikami, Yosuke Kawanaka, Shin’ichi Wakabayashi, and Shinobu Nagayama. Efficient fpga-based hardware algorithms for approximate string matching. *ITC-CSCC: 2008*, pages 201–204, 2008.
- [40] Y. Mitani, F. Ino, and K. Hagihara. Parallelizing exact and approximate string matching via inclusive scan on a gpu. *IEEE Transactions on Parallel and Distributed Systems*, 28(7):1989–2002, July 2017. ISSN 1045-9219. doi: 10.1109/TPDS.2016.2645222.
- [41] Kefu Xu, Wenke Cui, Yue Hu, and Li Guo. Bit-parallel multiple approximate string matching based on gpu. *Procedia Computer Science*, 17:523 – 529, 2013. ISSN 1877-0509. doi: 10.1016/j.procs.2013.05.067. First International Conference on Information Technology and Quantitative Management.
- [42] L. S. N. Nunes, J. L. Bordim, K. Nakano, and Y. Ito. A fast approximate string matching algorithm on gpu. In *2015 Third International Symposium on Computing and Networking (CANDAR)*, pages 188–192, Dec 2015. doi: 10.1109/CANDAR.2015.29.
- [43] ThienLuan Ho, Seung-Rohk Oh, and HyunJin Kim. A parallel approximate string matching under levenshtein distance on graphics processing units using warp-shuffle operations. *PLOS ONE*, 12(10): 1–15, 10 2017. doi: 10.1371/journal.pone.0186251.
- [44] Tom Van Court and Martin C Herbordt. Families of fpga-based algorithms for approximate string matching. In *Application-Specific Systems, Architectures and Processors, 2004. Proceedings. 15th IEEE International Conference on*, pages 354–364. IEEE, 2004.
- [45] T.F. Smith and M.S. Waterman. Identification of common molecular subsequences. *Journal of Molecular Biology*, 147(1):195 – 197, 1981. ISSN 0022-2836. doi: 10.1016/0022-2836(81)90087-5.
- [46] Saul B Needleman and Christian D Wunsch. A general method applicable to the search for similarities in the amino acid sequence of two proteins. *Journal of molecular biology*, 48(3):443–453, 1970.
- [47] Gene Myers. A fast bit-vector algorithm for approximate string matching based on dynamic programming. *Journal of the ACM (JACM)*, 46(3):395–415, 1999.
- [48] Robert A Wagner and Michael J Fischer. The string-to-string correction problem. *Journal of the ACM (JACM)*, 21(1):168–173, 1974.
- [49] Chi Wai Yu, KH Kwong, Kin-Hong Lee, and Philip Heng Wai Leong. A smith-waterman systolic cell. In *New Algorithms, Architectures and Applications for Reconfigurable Computing*, pages 291–300. Springer, 2005.
- [50] Stefan Dydel and Piotr Bała. Large scale protein sequence alignment using fpga reprogrammable logic devices. In *Field Programmable Logic and Application*, pages 23–32. Springer, 2004.
- [51] Ashish Sirasao, Elliott Delaye, Ravi Sunkavalli, and Stephen Neuendorffer. Fpga based opencl acceleration of genome sequencing software. *System*, 128(8.7):11, 2015.

- [52] Martin C Herbordt, Yongfeng Gu, Bharat Sukhwani, and Tom VanCourt. Single pass, blast-like, approximate string matching on fpgas. In *Field-Programmable Custom Computing Machines, 2006. FCCM'06. 14th Annual IEEE Symposium on*, pages 217–226. IEEE, 2006.
- [53] Stephen F. Altschul, Warren Gish, Webb Miller, Eugene W. Myers, and David J. Lipman. Basic local alignment search tool. *Journal of Molecular Biology*, 215(3):403 – 410, 1990. ISSN 0022-2836. doi: [https://doi.org/10.1016/S0022-2836\(05\)80360-2](https://doi.org/10.1016/S0022-2836(05)80360-2). URL <http://www.sciencedirect.com/science/article/pii/S0022283605803602>.
- [54] Benjamin West, Roger D Chamberlain, Ronald S Indeck, and Qiong Zhang. An fpga-based search engine for unstructured database. In *Proc. of 2nd Workshop on Application Specific Processors*, volume 12, pages 25–32, 2003.
- [55] Jörn Hoffmann, Dirk Zeckzer, and Martin Bogdan. Using fpgas to accelerate myers bit-vector algorithm. In Efthyvoulos Kyriacou, Stelios Christofides, and Constantinos S. Pattichis, editors, *XIV Mediterranean Conference on Medical and Biological Engineering and Computing 2016*, pages 535–541, Cham, 2016. Springer International Publishing. ISBN 978-3-319-32703-7.
- [56] Hans-Martin Blüthgen and Tobias G Noll. A programmable processor for approximate string matching with high throughput rate. In *Application-Specific Systems, Architectures, and Processors, 2000. Proceedings. IEEE International Conference on*, pages 309–316. IEEE, 2000.
- [57] Yuichiro Utan, Shin Ichi Wakabayashi, and Shinobu Nagayama. An fpga-based text search engine for approximate regular expression matching. In *Field-Programmable Technology (FPT), 2010 International Conference on*, pages 184–191. IEEE, 2010.
- [58] Jin Hwan Park and KM George. Parallel string matching algorithms based on dataflow. In *Systems Sciences, 1999. HICSS-32. Proceedings of the 32nd Annual Hawaii International Conference on*, pages 10–pp. IEEE, 1999.
- [59] Chien-Min Ou, Che-Yu Yeh, Yong-Long Su, Wen-Jyi Hwang, and Jing-Fung Chen. Fpga implementation of content-based music retrieval systems. In *Embedded Software and Systems Symposia, 2008. ICESS Symposia'08. International Conference on*, pages 96–103. IEEE, 2008.
- [60] Michael John Sebastian Smith. *Application-specific integrated circuits*. Addison-Wesley Professional, 2008.
- [61] Stephen Brown. Fpga architectural research: a survey. *Design & Test of Computers, IEEE*, 13(4): 9–15, 1996.
- [62] Kiran Bondalapati and Viktor K Prasanna. Reconfigurable computing systems. *Proceedings of the IEEE*, 90(7):1201–1217, 2002.
- [63] Yosuke Kawanaka, Shin'ichi Wakabayashi, and Shinobu Nagayama. A systolic regular expression pattern matching engine and its application to network intrusion detection. In *FPT*, pages 297–300, 2008.
- [64] V Levenstein. Binary codes capable of correcting spurious insertions and deletions of ones. *Problems of Information Transmission*, 1(1):8–17, 1965.
- [65] Peter H Sellers. The theory and computation of evolutionary distances: Pattern recognition. *Journal of Algorithms*, 1(4):359 – 373, 1980. ISSN 0196-6774. doi: 10.1016/0196-6774(80)90016-4.
- [66] Esko Ukkonen. Algorithms for approximate string matching. *Information and Control*, 64(1):100 – 118, 1985. ISSN 0019-9958. doi: 10.1016/S0019-9958(85)80046-2. International Conference on Foundations of Computation Theory.
- [67] Tetsushi Matsui, Takeaki Uno, Juzoh Umemori, and Tsuyoshi Koide. A new approach to string pattern mining with approximate match. In *Discovery Science*, pages 110–125. Springer, 2013.
- [68] Yoontae Lee, Kipyoun Jeon, Jun-Tae Lee, Sunyoung Kim, and V Narry Kim. MicroRNA maturation: stepwise processing and subcellular localization. *The EMBO journal*, 21(17):4663–4670, 2002.

- [69] Julia Winter, Stephanie Jung, Sarina Keller, Richard I Gregory, and Sven Diederichs. Many roads to maturity: microrna biogenesis pathways and their regulation. *Nature cell biology*, 11(3):228, 2009.
- [70] Glen M Borchert, William Lanier, and Beverly L Davidson. Rna polymerase iii transcribes human micrnas. *Nature Structural and Molecular Biology*, 13(12):1097, 2006.
- [71] Yoontae Lee, Minju Kim, Jinju Han, Kyu-Hyun Yeom, Sanghyuk Lee, Sung Hee Baek, and V Narry Kim. Microrna genes are transcribed by rna polymerase ii. *The EMBO journal*, 23(20):4051–4060, 2004.
- [72] Baekgyu Kim, Kyowon Jeong, and V Narry Kim. Genome-wide mapping of drosha cleavage sites on primary micrnas and noncanonical substrates. *Molecular cell*, 66(2):258–269, 2017.
- [73] Jong-Eun Park, Inha Heo, Yuan Tian, Dharendra K Simanshu, Hyesik Chang, David Jee, Dinshaw J Patel, and V Narry Kim. Dicer recognizes the 5' end of rna for efficient and accurate processing. *Nature*, 475(7355):201, 2011.
- [74] Markus Landthaler, Abdullah Yalcin, and Thomas Tuschl. The human digeorge syndrome critical region gene 8 and its d. melanogaster homolog are required for mirna biogenesis. *Current biology*, 14(23):2162–2167, 2004.
- [75] J Michael Thomson, Martin Newman, Joel S Parker, Elizabeth M Morin-Kensicki, Tricia Wright, and Scott M Hammond. Extensive post-transcriptional regulation of micrnas and its implications for cancer. *Genes & development*, 20(16):2202–2207, 2006.
- [76] Syeda Shabnam Hasan, Fareal Ahmed, and Rosina Surovi Khan. Approximate string matching algorithms: A brief survey and comparison. *International Journal of Computer Applications*, 120(8), 2015.
- [77] Ana Kozomara and Sam Griffiths-Jones. mirbase: annotating high confidence micrnas using deep sequencing data. *Nucleic acids research*, 42(D1):D68–D73, 2013.
- [78] W James Kent, Charles W Sugnet, Terrence S Furey, Krishna M Roskin, Tom H Pringle, Alan M Zahler, and David Haussler. The human genome browser at ucsc. *Genome research*, 12(6):996–1006, 2002.
- [79] Hsiang Tsung Kung and Charles E Leiserson. Algorithms for vlsi processor arrays. *Introduction to VLSI systems*, pages 271–292, 1980.
- [80] Richard J Lipton and Daniel Lopresti. A systolic array for rapid string comparison. In *Proceedings of the Chapel Hill Conference on VLSI*, pages 363–376, 1985.