

This is the peer reviewed version of the following article:

A State-of-the-Art Review with Code about Connected Components Labeling on GPUs / Bolelli, Federico; Allegretti, Stefano; Lumetti, Luca; Grana, Costantino. - In: IEEE TRANSACTIONS ON PARALLEL AND DISTRIBUTED SYSTEMS. - ISSN 1045-9219. - (2024), pp. 1-20. [10.1109/TPDS.2024.3434357]

Terms of use:

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

23/03/2025 22:15

(Article begins on next page)

A State-of-the-Art Review with Code about Connected Components Labeling on GPUs

Federico Bolelli, *Member, IEEE*, Stefano Allegretti, Luca Lumetti, and Costantino Grana, *Member, IEEE*

Abstract—This article is about Connected Components Labeling (CCL) algorithms developed for GPU accelerators. The task itself is employed in many modern image-processing pipelines and represents a fundamental step in different scenarios, whenever object recognition is required. For this reason, a strong effort in the development of many different proposals devoted to improving algorithm performance using different kinds of hardware accelerators has been made. This paper focuses on GPU-based algorithmic solutions published in the last two decades, highlighting their distinctive traits and the improvements they leverage. The state-of-the-art review proposed is equipped with the source code, which allows to straightforwardly reproduce all the algorithms in different experimental settings. A comprehensive evaluation on multiple environments is also provided, including different operating systems, compilers, and GPUs. Our assessments are performed by means of several tests, including real-case images and synthetically generated ones, highlighting the strengths and weaknesses of each proposal. Overall, the experimental results revealed that block-based oriented algorithms outperform all the other algorithmic solutions on both 2D images and 3D volumes, regardless of the selected environment.

Index Terms—Parallel Image Processing, Connected Components Labeling, State-of-the-Art Review, GPU, CUDA.

1 INTRODUCTION

CONNECTED Components Labeling (CCL) is an image processing algorithm that plays a central role in machine vision, whenever object recognition and measurement are required. The task itself can be easily defined as the procedure of assigning to each pixel of a connected component (object) a unique identifier, typically an integer number. Starting from a binary input, a CCL algorithm generates the output symbolic image where pixels belonging to an object are given the same label. Once computed, the labeled image can be used to extract object(s) and further calculate its features such as area, perimeter, circularity, centroids, bounding boxes, etc. The process of extracting features is usually referred to as Connected Components Analysis, or CCA in short, and has been addressed by different authors in the past, both for CPU and GPU architectures [1], [2].

Multiple modern computer vision pipelines employ CCL as a pre- or post-processing algorithm when tracking [3], [4], segmenting [5], [6], [7], localizing [8], generating [9], or counting [10] objects inside images and videos, in different medical imaging applications ranging from skin lesion segmentation and classification [11], [12], [13] to Whole-Slide Image (WSI) analysis [14], [15], for document restoration [16], [17], and graph analysis [18], [19]. All the aforementioned applications benefit from efficient CCL implementations and this is why the algorithmic proposals of the last decades have focused on performance optimization [20], [21], [22], [23], [24], [25], [26], [27], [28], [29], [30].

It is worth mentioning that given the input and selected pixel connectivity, the final result of CCL and CCA is uniquely defined and the main differences among existing algorithmic solutions are the execution time and memory

requirements. Moreover, given that the object geometry and shape may be extremely complex, the task is known to be more time-consuming than any other binary image operator, e.g., filtering, thresholding, interpolation, edge detection [31]. Most importantly, the labeling procedure cannot be completed by mere parallel local operations, requiring intrinsically sequential procedures [32].

However, the fast advance of Graphics Processing Units (GPUs) in the last years encouraged the development of algorithms specifically designed to work in data parallel environments. Indeed, along with sequential solutions [31], [33], [34], many algorithms exploiting hardware parallelism have been proposed [1], [25], [26], [35], [36], [37], [38], [39].

Most of the improvements introduced for sequential implementation have been sooner or later ported to GPU algorithms, or inspired for additional optimizations on such architectures. An example is the forest-based implementation of the *union-find* data structure, originally applied to CCL in [40] to solve equivalences between labels, and later employed for implementing GPU algorithms [36], [41]. The *block-based* approach is another noticeable example: introduced in [42] and subsequently employed in different CPU-based proposals [43], [44], [45] has been recently ported to GPU-based algorithms [36], [46], [47].

This article aims to provide a state-of-the-art review of GPU-based algorithms, highlighting innovative and common traits of each proposal, and discussing their strengths and weaknesses. Whenever possible, the union-find technique will be used as the common denominator. All the algorithms are evaluated under different perspectives (e.g., total and step-level execution time using both real and synthetically generated images and memory usage) using different experimental settings (e.g., operating systems, compilers, and GPUs). Moreover, the implementations used for the evaluation are publicly released in the YACCLAB benchmark [48], allowing other research to reproduce the

The authors are with the Dipartimento di Ingegneria "Enzo Ferrari," Università degli Studi di Modena e Reggio Emilia, Italy. E-mail: {name.surname}@unimore.it

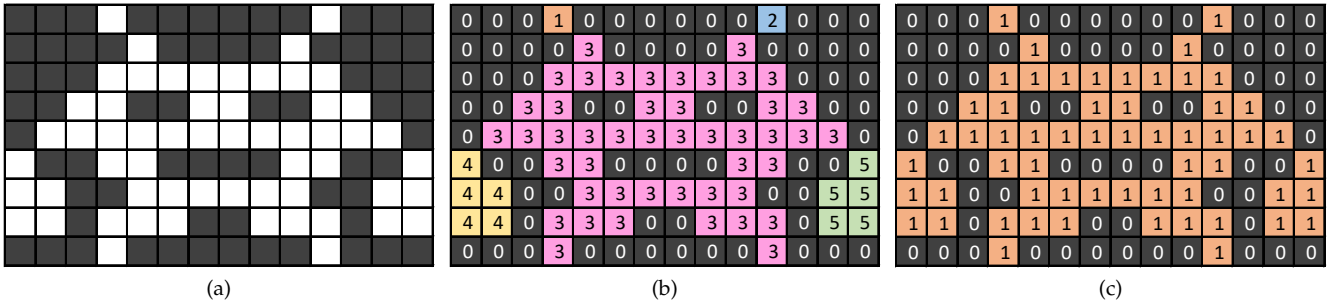


Fig. 1. Example of the labeling procedure applied on the input binary image (a), which depicts a fantasy character inspired by Space Invaders. (b) and (c) are the symbolic images produced by the labeling procedure when using 4- and 8-connectivity respectively.

experiments and verify our claims. All the code is taken from the original publication or implemented carefully following the description provided in the corresponding paper. We strongly believe in reproducible research and our effort is heading in this direction, letting everybody compare with existing solutions in the fairest possible way.

The rest of the article is organized as follows. Sec. 2 introduces a formal definition of the task, defining all the concepts that will be later employed in the discussion, including the union-find data structure. All the major GPU-based CCL proposals will be categorized and later described in Sec. 3. To better understand algorithms' details, the pseudo-code is often included. In Sec. 4, the benchmark used for the evaluation is detailed, providing a specific description of the datasets and evaluation criteria employed. Performance analysis and discussion are then provided in Sec. 5. Finally, Sec. 6 draws conclusive remarks and future research directions.

2 PROBLEM FORMALIZATION

This Section formalizes all the core elements related to the connected components labeling task, laying out the foundations for the rest of the discussion.

Let us start from the concept of image. A binary image I_d is a function defined over a d -dimensional rectangular lattice \mathcal{L}_d , where $I_d(x)$ is the value of pixel (or voxel) $x \in \mathcal{L}_d$, with x identified by its coordinates in the lattice, (x_1, \dots, x_d) . We are specifically interested in 2 and 3-dimensional images, i.e., volumes.

The distance between pixels x and y can be defined in two different ways, D_1 and D_∞ , also known as Manhattan and chessboard or Chebyshev distance, respectively:

$$D_1(x, y) = \sum_{i=1}^d |x_i - y_i| \quad (1)$$

$$D_\infty(x, y) = \max_{i=1..d} |x_i - y_i| \quad (2)$$

Each distance has an associated neighborhood:

$$V_1^i(x) = \{y \mid D_1(x, y) \leq i\} \quad (3)$$

$$V_\infty^i(x) = \{y \mid D_\infty(x, y) \leq i\} \quad (4)$$

We can now use these generic definitions to describe the different connectivities employed for 2D and 3D. In 2D images, there are two possibilities:

- *4-connectivity.* Two pixels of the lattice, p and q , are said to be 4-connected if they share a side. Formally, $\mathcal{N}_4(p) = V_1^1(p)$. We can say that p and q are 4-connected if $q \in \mathcal{N}_4(p)$, which implies $p \in \mathcal{N}_4(q)$. The name of the connectivity is intuitively derived from the number of sides a square has: four (Fig. 2a).
- *8-connectivity.* If pixels are considered connected also when they share only a vertex, four more neighbors appear, obtaining the 8-connectivity. In other words, two pixels are said to be 8-connected if they share at least one vertex. We can define $\mathcal{N}_8(p) = V_\infty^1(p)$ and say that p and q are 8-connected if $q \in \mathcal{N}_8(p)$, implying that $p \in \mathcal{N}_8(q)$ (Fig. 2b).

Similar definitions can be provided for 3D volumes. In this context, pixels are often called voxels, and can be visualized as cubes; since two of them can share a face (with at most 6 different cubes), an edge (with at most 12 cubes), or a vertex (with at most 8 cubes), three different kinds of connectivity can be defined, i.e., 6-, 18-, and 26-connectivity, whose formal definition is reported below:

$$\mathcal{N}_6(v) = V_1^1(v) \quad (5)$$

$$\mathcal{N}_{18}(v) = V_1^2(v) \cap V_\infty^1(v) \quad (6)$$

$$\mathcal{N}_{26}(v) = V_\infty^1(v) \quad (7)$$

If we want to define a mapping between 2D and 3D worlds, we can say that 4- and 8-connectivity on 2D images correspond respectively to 6- and 26-connectivity on 3D volumes. To simplify the discussion and ease the reading, in the following the word *image* is used to refer to 3D volumes also. The generic symbols \mathcal{L} and I will be used, when suitable, for both dimensionalities.




Fig. 2. Graphical representation of (a) 4-connectivity and (b) 8-connectivity. The "current" pixel is depicted in dark yellow and its neighborhood (connected pixels) is represented in yellow.

As mentioned, the CCL task aims at identifying objects within a binary image, so it is mandatory to distinguish meaningful regions from the rest of the image. The formers are usually called *foreground*, \mathcal{F} , while the latter is identified as *background*, \mathcal{B} . It is a common convention to assign value 1 (or 255) to foreground pixels, and value 0 to background ones:

$$\mathcal{F} = \{p \in \mathcal{L} \mid I(p) = 1\} \quad (8)$$

$$\mathcal{B} = \{p \in \mathcal{L} \mid I(p) = 0\} \quad (9)$$

Connected components labeling aims at identifying disjoint objects composed of foreground pixels. Given a neighborhood definition \mathcal{N} and two foreground pixels $p, q \in \mathcal{F}$, \diamond is the *connectivity* relation defined as:

$$p \diamond q \Leftrightarrow \exists \{s_i \in \mathcal{F} \mid s_1=p, s_{n+1}=q, s_{i+1} \in \mathcal{N}(s_i), i=1, \dots, n\} \quad (10)$$

Two pixels p and q are said to be *connected* when condition $p \diamond q$ is satisfied, meaning that a path of connected pixels from p to q exists or, in other words, you can move from p to q crossing only foreground pixels connected two by two. The proposed formalism does not consider background pixels, which are excluded from the concept of connectivity. \diamond is an equivalence relation since *reflexivity*, *symmetry* and *transitivity* are satisfied by pixel connectivity. For this reason, we can define $[p]$, the equivalence class of a pixel p , as:

$$[p] = \{q \in \mathcal{F} \mid p \diamond q\} \quad (11)$$

Equivalence classes based on \diamond relationship are called *Connected Components* (CCs). Every pair of connected components $[p]$ and $[q]$ is either equal or disjoint, meaning that the set of all connected components is a partition of \mathcal{F} .

Connected components labeling algorithms aim at assigning a different label to every connected component. When applied to I , the output of such an algorithm is a symbolic image L where, for every $p \in \mathcal{F}$, $L(p)$ is the label of the connected component that p belongs to ($[p]$), and for every $q \in \mathcal{B}$, $L(q) = 0$. Fig. 1 depicts an example of the labeling procedure applied to a fantasy character inspired by Space Invaders (Fig. 1a). When using 4-connectivity, the output will be that of Fig. 1b, while using 8-connectivity will produce the symbolic image reported in Fig. 1c.

2.1 The Union-Find Data Structure

During the labeling procedure, pixels belonging to the same connected component may be assigned different provisional labels; this requires a mechanism to keep track of possible equivalences and eventually solve them. One of the most effective strategies to achieve this goal is the forest-based implementation of the union-find data structure, which provides a quasi-linear solution to the disjoint-set union problem. This technique has been applied to sequential CCL for the first time by Dillencourt *et al.* [49], and it represents the basis of most modern approaches for label resolution, including those involving GPUs. As the name suggests, it consists of two basic operations —*find* and *union*— that are performed on a forest of *anti-arborescences*, i.e., directed rooted trees oriented towards the root. This forest, \mathcal{P} , represents a partition of the set \mathcal{S} that corresponds to the lattice \mathcal{L} in our scenario. The two operations can be defined as:

- *find* takes a node of a tree ($a \in \mathcal{P}$) as input and returns the corresponding root;
- *union* joins together the trees two given nodes ($a, b \in \mathcal{P}$) belong to, by setting the root of one tree (usually the one with the smallest value) as the father of the other one.

An implementation intended for GPU architectures is shown in Alg. 1. To avoid concurrency issues, the union makes use of atomic operations.

When referring to CCL, all the pixels of the image are nodes in the data structure. Initially, each foreground pixel is a standalone tree; then, by means of union operations, trees of connected pixels are merged together, until a single tree for each connected component is obtained.

From an implementation point of view, the forest can be stored in an array P , with one index per pixel: at $P[i]$ we have the parent of node i . When $P[i] = i$, the node i is the root of the tree. In Fig. 3, an example of union and find operations executed on a P array is reported.

In sequential CCL algorithms, the union-find array is usually a stand-alone data structure, while most of the GPU implementations available in the literature store it directly in the output image, avoiding expensive system calls for GPU memory allocation and improving performance. Once a single tree per connected component is obtained, the *flattening* operation can be performed to complete the CCL task. The flattening procedure links each node of the array directly to the root of its tree, meaning that all the pixels of each connected component are associated with the same identifier. A possible GPU implementation of the flattening operation is represented by the *Compress* kernel (Alg. 1). As firstly proposed in [50], the *Compress* implementation can be optimized for data-parallel environments and converted into *InlineCompression* (IC). In this case, the father of a label a is updated at every step of the tree traversal, thus making the update available to any possible concurrent thread that would save memory accesses.

To distinguish between background and foreground pixels a common approach is to shift labels, and start counting from 1 instead of 0. Such an approach ensures that every foreground pixel has a positive label, while 0 is reserved for the background.

In the following, we will simply use “union-find” when referring to its forest-based implementation.

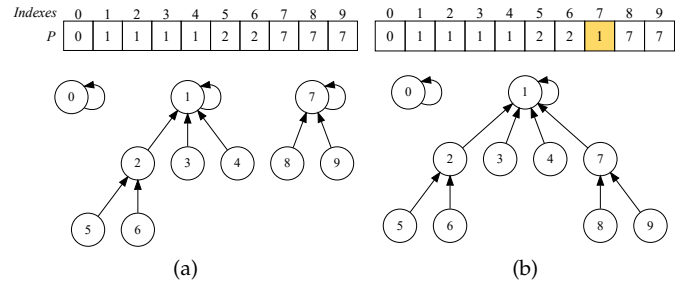


Fig. 3. Visualization of P represented as a forest of trees and stored in memory as an array. A *union*(9,4) performed on (a) will produce (b). Indeed, the root nodes of 9 and 4 are found through the *find* that respectively return 7 and 1. The tree with higher root value is then connect to the root of the other tree by updating $P[7]$ with 1.

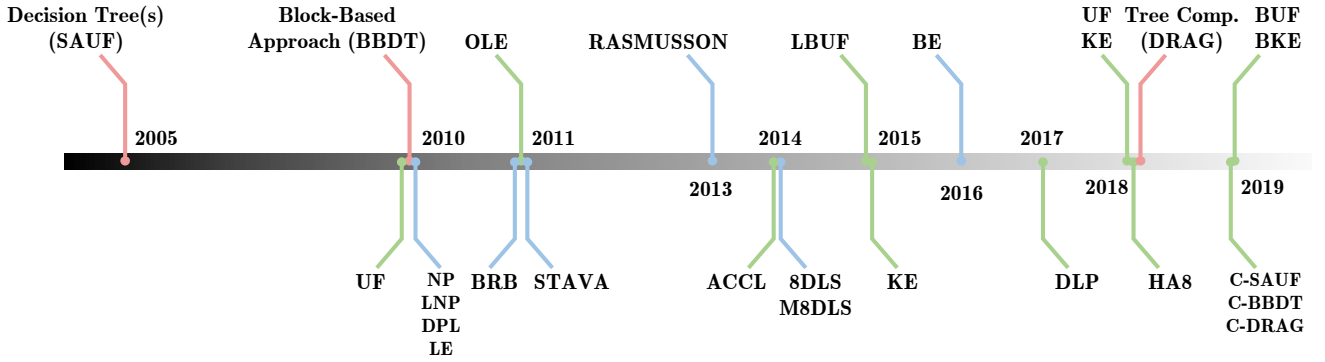


Fig. 4. Timeline of algorithm publications. Red lines represent CPU algorithms that introduced breakthrough changes in the labeling procedures, later included in GPU implementations. Blue and green lines are respectively used to identify iterative and direct solutions. Few algorithms appear twice in the timeline, meaning that the algorithm was initially released for a single connectivity type (e.g. 4-connectivity), and later extended to work with other types of connectivity (e.g. 8-connectivity) or with 3D volumes.

Algorithm 1 Pseudo-code for `Union` procedure and `Find` function. L is both the union-find array and the output labeled image, a and b are both array indexes and pixel identifiers. `Compress` and `InlineCompress` are two variations of the procedure implementing the *flattening* operation.

```

1: function FIND( $L, a$ )
2:   while  $L[a] \neq a$  do
3:      $a \leftarrow L[a]$ 
4:   return  $a$ 

5: procedure COMPRESS( $L, a$ )
6:    $L[a] \leftarrow \text{Find}(L, a)$ 

7: function INLINECOMPRESS( $L, a$ )
8:    $id \leftarrow a$ 
9:   while  $L[a] \neq a$  do
10:     $a \leftarrow L[a]$ 
11:     $L[id] \leftarrow a$ 
12:   return  $a$ 

13: procedure UNION( $L, a, b$ )
14:    $done \leftarrow false$ 
15:   while  $done = false$  do
16:      $a \leftarrow \text{Find}(L, a)$ 
17:      $b \leftarrow \text{Find}(L, b)$ 
18:     if  $a < b$  then
19:        $old \leftarrow \text{atomicMin}(\&L[b], a)$ 
20:        $done \leftarrow (old = b)$ 
21:        $b \leftarrow old$ 
22:     else if  $b < a$  then
23:        $old \leftarrow \text{atomicMin}(\&L[a], b)$ 
24:        $done \leftarrow (old = a)$ 
25:        $a \leftarrow old$ 
26:     else
27:        $done \leftarrow true$ 

```

3 ALGORITHMS - CLASSIFICATION AND DESCRIPTION

Now that a common terminology has been defined, we can dig into different proposals available in the literature. There are many different ways to categorize GPU CCL algorithms.

Broadly speaking, one of the most common categorizations distinguishes between *iterative* and *direct* solutions. The former repeats one or more kernels a data-dependent amount of times. Usually, they iterate until no more changes in data are detected. These algorithms tend to tolerate race conditions, demanding the correction of wrong results to the next iteration. Direct algorithms, on the other hand, perform a fixed number of kernel executions. Differently from iterative algorithms, each kernel must produce an exact result: therefore, most direct algorithms employ hard concurrency control to avoid race conditions, usually in the form of atomic operations.

Another possible classification may be based on the minimal set of pixels (or voxels) that are processed together. In this sense, pixel-based, run-based, and block-based algorithms can be identified:

- *Pixel-based* algorithms are the simplest: each foreground pixel has its own label;
- *Run-based* algorithms work on *runs*, i.e., chunks of consecutive foreground pixels in the same row, and assign labels to these runs instead of single pixels. This approach is beneficial in some scenarios, but performance significantly drops when foreground objects within images have elongated shapes on the y-direction;
- *Block-based* algorithms, originally proposed in [42], label 2×2 blocks of pixels at once. This is possible because, when considering 8-connectivity, any two foreground pixels in a 2×2 block are connected to each other. When moving to 3D the concept of blocks can remain at slice level or can be extended to the third dimension, considering $2 \times 2 \times 2$ blocks.

In the following of this Section, we will describe state-of-the-art solutions using the first classification. The timeline reported in Fig. 4 depicts the temporal release of each algorithm, providing the reference context. Later in this Section, a table summarizing all the characteristics of the algorithms (i.e., publication year, authors, the minimal set of pixels analyzed, scanning procedure, connectivity, input supported, data structure(s) required, and space complexity) is also provided (Tab. 1).

Algorithm 2 Neighbour Propagation and Local Neighbour Propagation kernels. I is the input image, L is the output label image, id is the thread identifier, corresponding to a pixel raster index. r and c represent the row and column index. $changed$ is an output variable to communicate to the host if any change occurs. If so, the kernel must be re-run.

```

1: kernel NP( $I, L, r, c$ )
2:    $label \leftarrow L(r, c)$ 
3:   for all  $(n_r, n_c) \in \mathcal{N}(r, c)$  do
4:     if  $I(n_r, n_c) = I(r, c) \wedge L(n_r, n_c) < label$  then
5:        $label \leftarrow L(n_r, n_c)$ 
6:        $changed \leftarrow \text{true}$ 
7:    $L(r, c) \leftarrow label$ 
8: kernel LNP( $I, L, r, c$ )
9:    $label \leftarrow L(r, c)$ 
10:  for all  $(n_r, n_c) \in \mathcal{N}(r, c)$  do
11:    if  $I(n_r, n_c) = I(r, c) \wedge L(n_r, n_c) < label$  then
12:       $label \leftarrow L(n_r, n_c)$ 
13:       $changed \leftarrow \text{true}$ 
14:   $\_\_shared\_\_ sL[blockDim.x \times blockDim.y]$ 
15:   $r_s \leftarrow threadIdx.y$ 
16:   $c_s \leftarrow threadIdx.x$ 
17:  loop
18:     $sL(r_s, c_s) \leftarrow label$ 
19:     $\_\_syncthreads()$ 
20:     $sChanged \leftarrow \text{false}$ 
21:    for all  $(n_r, n_c) \in \mathcal{N}(r, c)$  do
22:      if  $I(n_r, n_c) = I(r, c) \wedge sL(n_r, n_c) < label$  then
23:         $label \leftarrow sL(n_r, n_c)$ 
24:         $sChanged \leftarrow \text{true}$ 
25:     $\_\_syncthreads()$ 
26:    if  $sChanged = \text{false}$  then
27:      break
28:   $L(r, c) \leftarrow label$ 

```

3.1 Iterative Algorithms

3.1.1 Neighbor Propagation (NP)

The first work that addresses the GPU CCL problem is due to Hawick *et al.*, and is dated back to 2010 [51]. The authors deal with the general problem of identifying CCs in graphs and then propose four specific algorithms for d-dimensional hypercubic meshes, which include binary images. The first and simplest of these proposals is Neighbor Propagation. All foreground pixels are initialized with sequential values, each equal to the pixel raster index. Then, the provisional label of each pixel is updated to the minimum of its neighbor labels. This operation must be repeated until convergence, i.e., until there are no more changes in the output image. For this reason, the algorithm is inserted in the iterative group. The pseudocode for this algorithm is provided in Alg. 2.

3.1.2 Local Neighbor Propagation (LNP)

Local Neighbor Propagation is an optimization of NP that initializes the output image, splits it into blocks (16×16 in the original paper), and performs the update operations described above on each block separately. The advantage over the previous solution is that, in this case, the shared memory inside the block is exploited. As NP, each thread

updates a different pixel. The update kernel is divided into two steps: during the first step, each thread loads into shared memory the minimum neighbor label of its pixel, reading from device memory; then, the second step updates each label with the minimum of the neighbors, same as NP, but this time performed locally, in shared memory. This last operation is repeated until block convergence. Then, the content of the shared memory is copied into device memory, and the procedure restarts from the beginning until global convergence is reached. Alg. 2 details the kernels of the algorithm.

3.1.3 Directional Propagation Labeling (DPL)

Directional Propagation Labeling (DPL), again proposed in the same paper as NP and LNP [51], tries to overcome the limit of the neighborhood-confined propagation. In order to accomplish this, it assigns to each thread the task of propagating the minimum label through a row or a column. This operation is performed alternatively in four directions: left to right, bottom to top, right to left, and top to bottom. Unlike the two previous algorithms, this strategy is only compatible with 4-connectivity.

3.1.4 Label Equivalence (LE)

Label Equivalence is the last and certainly the most interesting proposal of the paper from Hawick *et al.* [51]. It is also one of the first iterative algorithms to record and solve equivalences using an auxiliary data structure.

The paper does not explicitly refer to the union-find, but the operations performed are the same and can be described by referring to them. In the first kernel, *Init*, the output image L is initialized as usual, with each pixel being given a provisional label equal to its raster index. Moreover, an additional data structure with the same size as the output image, H , is allocated and initialized in the same way. The algorithm then consists of three more kernels that are repeated in sequence until convergence: *Scan*, *Analysis*, and *Labeling*. The first kernel, *Scan*, is responsible for updating equivalences between pixels. Each thread works on a different pixel, finding the smallest of the neighbor labels and updating the corresponding value in H with the minimum label found. This operation does not touch L and only modifies H . The second kernel, *Analysis*, performs the compression of the equivalence tree, i.e., each thread is responsible for replacing a pixel of H with the root of the tree. This operation corresponds to *find*. The third and last kernel, *Labeling*, updates the labels in the output image L with the values contained in H . *Scan* requires atomic operations to avoid lost updates on H , in the case that two threads happen to concurrently update the same pixel.

After its appearance, LE has been improved by many authors. Kalentev *et al.* [52] noticed that the need for a separate data structure to store label equivalences could be removed through the direct use of the output image L . In this optimization, only two kernels are iterated after the initialization. *Analysis* kernel performs the compression of trees directly on the output image, updating the temporary label of each pixel. Thus, there is no need for a separate kernel devoted to copy the labels. In *Scan* kernel, every thread finds the smallest neighbor label, nl , and assigns it to the father (with label fl) iff $nl < fl$. Kalentev *et al.*

Algorithm 3 Label Equivalence kernels. I is the input image, L is both union-find array and output label image, id is the thread identifier, corresponding to a pixel raster index. r and c represent the row and column index, respectively.

```

1: kernel INIT( $I, L, r, c$ )
2:   if  $I(r, c) = 1$  then
3:      $L(r, c) \leftarrow r * w + c + 1$ 
4:   else
5:      $L(r, c) \leftarrow 0$ 

6: kernel SCAN( $L, r, c, changes$ )
7:    $l \leftarrow L(r, c)$ 
8:   if  $l > 0$  then
9:      $min\_l \leftarrow \text{FindMinNeighborLabel}(r, c)$ 
10:    if  $min\_l < l$  then
11:       $L[l - 1] \leftarrow \min(L[l - 1], min\_l)$ 
12:       $changes \leftarrow true$ 

13: kernel ANALYSIS( $L, r, c$ )
14:    $l \leftarrow L(r, c)$ 
15:   if  $l > 0$  then
16:      $L(r, c) \leftarrow \text{Find}(L, l)$ 

```

also removed the atomic operations by simply increasing the number of iterations. The authors claim that atomic operations slow down computation more than additional iterations caused by collisions do. Kernels of this optimized version of Label Equivalence are detailed in Alg. 3. The optimized version is identified as OLE, which stands for Optimized Label Equivalence.

3.1.5 Stava and Benes (STAVA)

Stava and Benes [53] proposed a solution that can be seen as an improvement of LE, obtained with the use of Tiles Merging. The algorithm is composed of four kernels: LocalLabeling, GlobalLabeling, BorderCompression, and PathCompression. In the first kernel, LocalLabeling, 16×16 tiles are labeled with a light version of LE, which does not require any additional union-find structure to store label equivalences, similar to the one proposed by Kalentev *et al.*. This kernel is detailed in Alg. 4.

The second kernel, GlobalLabeling, merges the labeled tiles, in a hierarchical order: in step n , macro-tiles are obtained by merging groups of 4×4 tiles, the output of step $n - 1$. Each group of tiles is processed by a three-dimensional thread block, where x and y components of the thread-id locally identify a tile. A union is performed between each pixel of the South and East border and its neighbors outside the tile. For each tile, the merging process is repeated until no more changes are detected. For optimization purposes, GlobalLabeling is alternated with BorderCompression, which performs find on border pixels only: the aim is to shorten union-find trees and consequently speed up the next round of global merge.

Finally, the last kernel, PathCompression, performs a find on all pixels of the output image, thus completing the CCL process.

Algorithm 4 Stava and Benes LocalLabeling kernel. I is the input image, L is both union-find array and output label image, id is the thread identifier, corresponding to a pixel raster index. r and c represent the row and column index, respectively.

```

1: kernel LOCALLABELING( $I, L, r, c$ )
2:    $\_shared\_sI[16 \times 16]$ 
3:    $\_shared\_sL[16 \times 16]$ 
4:    $\_shared\_changes[1]$ 
5:    $loc\_r \leftarrow \text{threadIdx.y}$ 
6:    $loc\_c \leftarrow \text{threadIdx.x}$ 
7:    $sI[loc\_r, loc\_c] \leftarrow I(r, c)$   $\triangleright$  Load input patch
8:    $\_syncthreads()$ 
9:    $l \leftarrow loc\_r * \text{blockDim.x} + loc\_c$ 
10:  loop  $\triangleright$  Pass 1 of the CCL algorithm
11:     $sL[loc\_r, loc\_c] \leftarrow l$ 
12:    if  $\text{threadIdx.x} = 0 \wedge \text{threadIdx.y} = 0$  then
13:       $changes[0] \leftarrow false$ 
14:       $\_syncthreads()$ 
15:       $new\_l \leftarrow l$   $\triangleright$  Find the minimal neighbor label
16:      for  $(n\_r, n\_c) \in \mathcal{N}(r, c)$  do
17:        if  $sI[loc\_r, loc\_c] = sI(n\_r, n\_c)$  then
18:           $new\_l \leftarrow \min(new\_l, sL(n\_r, n\_c))$ 
19:         $\_syncthreads()$ 
20:      if  $new\_l < l$  then  $\triangleright$  Merge the equivalence trees
21:         $\text{atomicMin}(sL[l], new\_l)$ 
22:         $changes[0] \leftarrow true$ 
23:       $\_syncthreads()$ 
24:      if  $changes[0] = false$  then
25:        break  $\triangleright$  Local solution has been found
26:       $l \leftarrow \text{Find}(sL, l)$   $\triangleright$  Pass 2
27:       $\_syncthreads()$ 
28:       $g\_l \leftarrow \text{convertLabelToGlobal}(l)$ 
29:       $L(r, c) \leftarrow g\_l$   $\triangleright$  Store result to device memory

```

3.1.6 Block-Run-Based (BRB)

Block-Run-Based connected components labeling [54], as the name suggests, combines block-based and run-based strategies in order to simplify the equivalence label-solving process. The whole algorithm is implemented in GPU shared memory, thus minimizing global memory bandwidth consumption. The image, of width W (and half-width HW), is divided into overlapping Block Rows (BR), composed of two consecutive pixel rows, where the upper pixel row is shared with the previous block row. Each thread works on a 2×2 block. The algorithm is composed of two scans, and the first scan is divided into three steps: BlockRunExtraction, PreviousBRUpdate, and CurrentBRUpdate.

In the first step, the current BR is compressed into one-pixel row by applying OR operation of its upper and lower row pixel value. Then, a run of contiguous 1s is detected from the compressed pixel row and the column index of the block containing the run starting pixel is assigned to each block in this run, by means of $_ballot_sync^1$ and $_clz$

1. $_ballot_sync$ is a warp vote function: it evaluates a predicate for all the threads in the warp and returns an integer whose n-th bit is set iff the predicate evaluates to non-zero for the n-th thread.

(Count Leading Zeros) instructions. From another perspective, each block can be considered as a node of a union-find tree, while the label value represents a linkage address pointing to the column location of its parent node. The purpose of the following two iterative steps is to merge union-find trees of current and previous BRs according to the connection relations. The second step, `PreviousBRUpdate`, merges union-find trees of pairs of connected blocks in the current and previous row, by linking one root to the other one. Then, `CurrentBRUpdate` is similar to the previous step, but only updates root nodes in the current BR. These two steps are iterated until convergence.

The second scan, `LabelAssignment`, proceeds row-wise in a reversed direction, and assigns a unique global label to each block recognized as a tree root. Then, child blocks get their label from the parent node.

Unfortunately, the implementation available only deals with specific image dimensions (e.g., 512×512 and 1024×1024). Any different size would require a specific implementation or multiple conditional checks that would significantly degrade the performance.

3.1.7 8-Directional Label Selection (8DLS)

8-Directional Label Selection [55] is another iterative algorithm in which, at every step, each pixel is assigned the minimum label found scanning each of the 8 directions radiating from it until a background pixel is encountered. Modified 8-Directional Label Solver (M8DLS) is an improved version of 8DLS, that avoids processing pixels that already have a minimum label. After two iterations of 8DLS, pixels that still have their original label are considered permanent. The specific amount of iterations, 2, is empirically determined based on exhaustive tests conducted by the authors. Unfortunately, when using such a kind of “optimization”, the correctness of the output may not be guaranteed.

3.1.8 Rasmusson

Rasmusson *et al.* [56] proposed a method based on early calculation of label propagation sizes, from the connectivity between pixels extracted in a pre-processing step, and reuse of established label propagation routes. As usual, the algorithm starts by initializing the output with provisional labels equal to pixel raster indices. The four most significant bits of each label store connection information about half the neighborhood; because of this trick, the algorithm only works on images with at most $2^{28} \approx 256\text{M}$ pixels.

The core of the algorithm is the `Propagate` kernel, which is repeated until convergence. The image is divided into 32×32 tiles, which are copied into shared memory, and a thread is created for each pixel. Then, for each of the 8 possible directions starting from a pixel, the maximum propagation is computed. These propagation sizes are iteratively calculated in a parallel manner: each thread stores the provisional propagation size in a local variable, p , and increments it at each step by the value of the provisional propagation size of the pixel distant p positions. In this way, a maximum propagation of 32 is reached after $\log_2 32 = 5$ iterations. Different threads share the respective values of p by means of shuffle operations, which efficiently exchange a variable between threads within a warp. The pixel-thread

Algorithm 5 Block Equivalence `Init` kernel. I is the input image, L is both union-find array and output label image, id is the thread identifier, corresponding to a pixel raster index. r and c represent the row and column index, respectively.

```

1: kernel INIT( $I, bL, bConn, r, c$ )
2:    $P \leftarrow 0$ 
3:    $P0 \leftarrow 777_{16}$ 
4:   if  $I[2r, 2c] > 0$  then
5:      $P \leftarrow P \mid P0$ 
6:   if  $I[2r, 2c+1] > 0$  then
7:      $P \leftarrow P \mid (P0 \ll 1)$ 
8:   if  $I[2r+1, 2c] > 0$  then
9:      $P \leftarrow P \mid (P0 \ll 4)$ 
10:  if  $I[2r+1, 2c+1] > 0$  then
11:     $P \leftarrow P \mid (P0 \ll 5)$ 
12:  if  $P > 0$  then
13:     $bL[r, c] \leftarrow r*w+c+1$ 
14:    if  $\text{HasBit}(P, 0) \wedge I[2r-1, 2c-1] > 0$  then
15:       $\text{SetBit}(bConn[r, c], 0)$ 
16:    if  $(\text{HasBit}(P, 1) \wedge I[2r-1, 2c] > 0) \vee$ 
17:       $(\text{HasBit}(P, 2) \wedge I[2r-1, 2c+1] > 0)$  then
18:       $\text{SetBit}(bConn[r, c], 1)$ 
19:    ...

```

association must change for each direction, in order for pixels aligned in a certain orientation to always be in the same warp. The label of each pixel is updated to the minimum of the 8 labels in the maximum propagation sizes, and also a propagation size of 1 is considered for the correctness of corner cases. In order to extend the label propagation to neighbor tiles, actually 33×33 blocks are copied into shared memory.

3.1.9 Block Equivalence (BE)

Proposed by Zavalishin *et al.* [46], Block Equivalence (BE in short) is a block-based version of Label Equivalence (Sec. 3.1.4), enriched with a novel method for checking block connectivity in a parallel fashion. The algorithm is composed of four kernels: `Init`, `Scan`, `Analysis`, and `FinalLabeling`, each requiring a thread per block. The `Init` kernel is responsible for finding which pairs of neighbor blocks are indeed connected. Decision trees, used by the original block-based CCL algorithm [42], tend to cause thread divergence and inefficiency in a GPU environment; therefore, a different method is proposed, that avoids nested conditional jumps:

- 1) Each internal pixel of block X is read, in order to find out which neighbor blocks could possibly be connected to X ;
- 2) Internal pixels of the neighbor blocks selected by the previous step are read, and the hypothesis of connectivity with X is confirmed or refused.

At the end of `Init`, block adjacency information is stored in an ad-hoc matrix BA , to be read again in the subsequent phases. Each block has 8 adjacent blocks, so one bitmapped byte per block is sufficient to record the neighbors (adjacent blocks sharing the same value). Kernels `Scan` and `Analysis` behave in the same way as the kernels of Label

Algorithm 6 LocalMerge kernel of UF. I is input image, L is both union-find array and output label image, id is the thread identifier, corresponding to a pixel raster index. r and c represent the row and column index, respectively.

```

1: kernel LOCAL MERGE( $I, L, r, c$ )
2:   __shared__  $sI[]$ 
3:   __shared__  $sL[]$ 
4:    $r_s \leftarrow \text{threadIdx.y}$ 
5:    $c_s \leftarrow \text{threadIdx.x}$ 
6:    $sI(r_s, c_s) \leftarrow L(r, c)$ 
7:   __syncthreads()
8:   for all  $(n_r, n_c) \in \mathcal{N}(r, c)$  do
9:     if  $I(r, c) = I(n_r, n_c)$  then
10:      Union( $sL, id_s, id_{sn}$ )
11:    $l \leftarrow \text{Find}(sL, id_s)$ 
12:    $L[id] \leftarrow \text{convertLabelToGlobal}(l)$ 

```

Equivalence, in the optimized version of Kalentev *et al.*; the only difference is that they work on blocks instead of pixels, storing block labels in a specific matrix, BL . Scan and Analysis are iterated until convergence, and then FinalLabeling copies block labels into pixel labels, thus composing the output. The pseudo-code is provided in Alg. 5.

3.2 Direct Algorithms

3.2.1 Union-Find (UF)

Union-Find, by Oliveira and Lotufo [41], is chronologically the first non-iterative proposal, and also one of the first GPU CCL algorithms overall. It consists of a parallel version of the common union-find algorithm, largely used in sequential CCL solutions. UF is based on three kernels: Initialization, Merge, and Compression. Each kernel is launched on a number of threads equal to the image size, and each thread is assigned a pixel, x . In order to save avoidable and time-consuming memory allocations, no additional memory is reserved for the union-find data structure. Instead, the equivalences between labels are stored, during the procedure, in the output image itself, in the sense that the provisional label assigned to a pixel corresponds to the linear index of its father node in the union-find forest. In the Initialization kernel, all foreground pixels in the output image are initialized with a provisional label equal to their linear index. From the union-find point of view, this procedure corresponds to the creation of a separate tree for every pixel. In kernel Merge, each thread working on a foreground pixel x checks its neighborhood mask, and for every foreground neighbor it performs a union with x . Since union is a symmetric operation, the neighborhood mask only contains half of the neighbors of a pixel. The union operation can involve reading and writing operations on other pixels than x . The possibility that two or more threads read or modify the same pixel is taken into account through the use of atomic operations in union, in order not to lose updates, and avoid the necessity of multiple iterations. After Merge, every connection between pixels in the image is reflected in the union-find structure, in the sense that every separate tree represents a connected component. Finally, the Compression kernel performs the compression of trees.

This operation makes sure that every pixel is assigned a label corresponding to the linear index of the root of its tree. Thus, every pixel in the same CC ends up sharing the same label, and the labeling task is completed. The aforementioned basic structure of the algorithm is enhanced with the addition of another common technique known as *tile merging*. This means that the entire algorithm is first performed on rectangular blocks the image is divided into in a preliminary phase called Local Merge. This allows to take full advantage of shared memory. This local phase is detailed in Alg. 6. Then, the Merge kernel is performed on border pixels only, and a final Compression is executed over the entire image. The original algorithm uses 4-connectivity, but it was extended to 8-connectivity in [57], by adding in the Merge kernel the diagonal directions to the neighborhood of a pixel.

3.2.2 Line-Based Union-Find (LBUF)

Line-Based Union-Find [50] is a variation of UF that employs single lines as tiles for the tile merging strategy. This choice reduces the neighborhood of every pixel to a subset containing only the two neighbors belonging to the same row, and consequently allows to simplify the logic of Local Merge. In fact, because only half the neighborhood needs to be checked (see Sec. 3.2.1), each thread only has to link its pixel to the one on the left, in the case that both are foreground. The remaining kernels of UF are left unchanged.

3.2.3 Komura Equivalence (KE)

Komura Equivalence [58] can be seen as a variation of UF, which involves a more complex initialization in order to remove the need for one union per pixel later. It consists of four steps: Initialization, Compression, Reduction, and Compression again. The main difference between KE and UF lies in the first kernel: differently from UF, KE Initialization does not create single-node trees. Instead, every thread checks the neighborhood of x in increasing raster index order, and the smallest foreground neighbor is set as the father node of x . This first phase aims at creating partial union-find trees right from the beginning, that are flattened by the subsequent Compression kernel. The Reduction kernel is a variation of Merge. It only performs a union between x and foreground neighbors that were not chosen during Initialization. Analogously to UF, a final Compression is required for flattening the forest trees. Same as UF, the original KE is a 4-connectivity algorithm. It is extended to 8-connectivity in [57], by adding the supplementary neighbors in both Initialization and Reduction.

3.2.4 Accelerated CCL (ACCL)

The CUDA compute capability 3.5 introduced several features, among which dynamic parallelism and Hyper-Q are the ones exploited by the Accelerated CCL (ACCL) algorithm, proposed by Paravecino and Kaeli in [59]. ACCL is configured as a run-based algorithm and is composed of two phases. In the first phase, a thread is run for each row of the image, with the aim of finding runs and recording the first and last indices. During this step, each run is assigned

a provisional label L , recorded in a separate matrix. During the second scan, connected runs from contiguous rows are joined. This particular operation involves spawning a child kernel for each merge operation, with one thread per provisional label in L . These threads update the respective labels of the added segment. The algorithm makes use of texture memory for read-only memory accesses, which improves the performance when multiple threads from the same block access contiguous memory locations. Finally, the Hyper-Q feature is exploited to process multiple images at the same time with different CUDA streams, thus increasing the total throughput.

3.2.5 Distanceless Label Propagation (DLP)

Proposed by Cabaret *et al.* [25], Distanceless Label Propagation is another algorithm based on union-find, characterized by an uncommon gather-scatter procedure used to propagate the minimum label of a CC. It includes four kernels: DLP-I, DLP-SR, DLP-R and DLP-RUF. The first kernel, DLP-I, initializes foreground pixels as usual, with a provisional label equal to the raster index. The second kernel is DLP-SR (SetRoot), and is responsible for a gather-scatter label propagation procedure. The minimum label is found in a mask selecting a reduced neighborhood of x , containing 2×2 pixels (x , right, down, and down-right); then, differently from most algorithms based on minimum label gathering, the minimum is scattered back to the whole mask with the SetRoot procedure: for each involved pixel, the root of its union-find tree is updated with the minimum label just found, by means of the usual atomic operation that characterizes direct algorithms. The third kernel, DLP-R, performs the classical compression of union-find trees, replacing provisional labels with the root of their trees. The last kernel is DLP-RUF, named after its distinctive algorithm, *recursive* union-find. It is the same as DLP-SR, but replaces the SetRoot procedure with a recursive implementation of the union. These four different kernels are used to build three sequential steps that compose the CCL algorithm: (i) *Tile Labeling* - the labeling process is performed locally on tiles, by running in sequence DLP-I, DLP-SR, DLP-R, DLP-RUF and finally another DLP-R; (ii) *Border Merging* - DLP-RUF is applied to border pixels only, with the aim of merging union-find trees of different tiles; (iii) *Relabeling* - DLP-R is applied to the whole image.

3.2.6 Hardware Accelerated 4/8 (HA4/HA8)

In 2018, Hennequin *et al.* [24] proposed HA4, a 4-connected run-based algorithm heavily based on CUDA intrinsics. The algorithm can be used to perform either connected components labeling or analysis; the CCL variation is composed of three kernels: StripLabeling, BorderMerging, and FinalLabeling.

The first kernel, which is also the most complex, performs partial labeling on horizontal strips. The block width is 32, the same as the warp size, and each thread is assigned a pixel X so that a line of 32 pixels perfectly matches a warp. The kernel begins with threads reading the values of their pixels and sharing them with the rest of the line (warp) with a `__ballot_sync` instruction, which produces a 32-bit bitmask of the foreground pixels; then, each thread checks the *start* and *end* position of the run X belongs to,

by means of efficient `__ffs` (Find First Set) and `__clz` intrinsics. The first pixel of each run is given a temporary label equal to its raster index, while the other pixels are left uninitialized. Then, in order to merge lines vertically, the first thread of the line copies the bitmask generated earlier with `__ballot_sync` in shared memory, so that it can be read from the lower line. Each thread performs a union if X or the pixel above is the start of a run. After this partial labeling of the block, the same threads shift 32 pixels to the right and continue labeling the stripe with the same method, extending runs that overcome the block border. This technique avoids the need to merge vertical borders and minimizes thread creation/destruction.

The second kernel merges the horizontal border, performing a union only on the start of runs, same as in StripLabeling. Finally, the last kernel performs the final labeling. Only the run-starting-pixel thread retrieves the final label with a *find*, and then it propagates it to the other threads with a shuffle operation; finally, each thread updates the label image with this value.

Since the original algorithm is 4-connected, we propose an 8-connected variation, which performance can be compared to the other algorithms in this survey. In order to make the algorithm 8-connected, it is sufficient to add the diagonal direction to the horizontal border merging and perform a union if the rightmost of the two adjacent pixels is a start. Specifically, each thread must merge X to the upper-left pixel if X is a start, and to the upper-right pixel if the latter is a start.

Checking and merging in a diagonal direction gets tricky at warp borders since the diagonal neighbor pixels can fall into different warps. To tackle such cases, blocks of the BorderMerging kernel encompass 1024 pixel-long lines, where each warp shares the value of the rightmost pixel to the subsequent warp using shared memory.

3.2.7 Decision Tree-Based Algorithms

C_SAUF, C_BBBDT, and C_DRAG, published in [47], are CUDA adaptations of successful CCL algorithms originally designed for running sequentially on a single CPU thread. The general structure of these algorithms is the following:

- *First scan*: scans the input image using a mask of already visited pixels, and assigns a temporary label to the current pixel(s), recording any equivalence between those found in the mask;
- *Flattening*: analyzes the registered equivalences and establishes the definitive labels to replace the provisional ones;
- *Second scan*: generates the output image by replacing provisional labels with final ones.

The first scan is the most complex and time-consuming, and it has been targeted by several improvements.

Scan Array-based Union-Find (SAUF), proposed by Wu *et al.* in [60], reduces the number of neighbors visited during the first scan using a *decision tree*. The idea is that if two already visited pixels are connected, their labels have already been marked as equivalent in the union-find data structure, so it is not necessary to check their values.

This algorithm was extended in [42] with the introduction of 2×2 blocks, in which all foreground pixels share

the same label. Since the neighboring mask becomes larger, the authors propose an optimal strategy to automatically build the decision tree by means of a dynamic programming approach [61]. The resulting algorithm is known as BBDT.

In [62], authors noticed the existence of identical and equivalent subtrees in the BBDT decision tree. *Identical* subtrees were merged together by the compiler optimizer, with the introduction of jumps in machine code, but *equivalent* ones were not. By also taking into account equivalent subtrees they converted the decision tree into a Directed Rooted Acyclic Graph, which they called DRAG. The code compression obtained does not impact neither on the memory accesses, nor on the number of comparisons, but allows a significant reduction of the machine code footprint. This heavily reduces the memory requirements, increasing the instruction cache hit rate and the overall run-time performance.

The adaptation of these algorithms to the CUDA domain, as proposed in [47], consists in translating each step into an appropriate kernel, plus an additional starting one for initializing provisional labels. Each kernel uses one thread per pixel (C_SAUF) or per block (C_BBDT, C_DRAG). The first kernel, `Initialization`, just sets the label of X to its raster index, the same as in many other algorithms. The second kernel, `Merge`, deals with the recording of equivalences between labels. During execution, thread t_x traverses a decision tree in order to decide which action needs to be performed, while minimizing the average amount of memory accesses. When no neighbors in the scanning mask are foreground, nothing needs to be done. In all other cases, the current label must be merged to those of connected pixels with a union. Then, it is easy to parallelize the flattening step: in the third kernel, `Compression`, thread t_x performs $L[id_x] \leftarrow Find(L, id_x)$ to link each provisional label to the representative of its union-find tree. This is equivalent to the homonymous kernel in UF (Sec. 3.2.1).

The last step of the sequential algorithms is the *second scan*, which updates labels in the output image L . A large part of the job of *second scan* is not necessary in the parallel implementations, because `Compression` already solves label equivalences in the output image. For C_BBDT and C_DRAG, a final processing of L is required to copy the label assigned to each block into its foreground pixels. This job is performed by the last kernel, `FinalLabeling`.

Although branches in the code are known to cause thread divergence and slow down performances when dealing with CUDA parallel architectures, the authors of [47] demonstrated that in real case scenarios, where the foreground density of input images is low ($< 10\%$) or high ($> 90\%$), the benefits introduced by the use of decision trees do compensate for the cost of thread divergence. Indeed, when foreground density satisfies such conditions the decision is usually taken in the first levels of the decision trees employed for the scanning phase, saving many memory accesses without breaking the thread execution flow.

3.2.8 Block-based Union-Find (BUF)

Block-based Union-Find (BUF) [36] is the first of the direct algorithms which employ a block-based approach. It shares the same kernels as UF (`Initialization`, `Merge`, and `Compression`), but each phase works on blocks instead

of single pixels, and there is an additional final kernel, `FinalLabeling`, to copy block labels into pixels ones. Until the end of the algorithm, block labels are stored in the output image, in the upper-left pixel of every block. In this way, it avoids the allocation of unnecessary device memory. The method used to check connectivity between 2×2 blocks is borrowed from BE (Sec. 3.1.9), and is performed in the `Merge` kernel.

3.2.9 Block-based Komura Equivalence (BKE)

Block-based Komura Equivalence (BKE), introduced along with BUF in [36], is another direct block-based algorithm, this time obtained by modifying KE (Sec. 3.2.3). BKE is quite similar to BUF in what it does; the major difference between the two is that BKE, the same as KE, starts linking together connected blocks directly in the initialization phase. This means that the task of finding what blocks are connected to the active one must be anticipated in the first phase. The method is the same one used for BUF and BE, but in this case, as in KE, only the connected neighbor block with the smallest linear index is linked to X . The remaining connectivity information is stored in a variable to be used later in `Reduce`, where the remaining neighbor blocks are merged to X . The neighbors that can be merged in `Reduce` are the ones to the North, North-East, and West. Being it the first one considered for the linking performed during `Initialization`, North-West is excluded. So, 3 bits are enough to store connectivity information for the `Reduce` phase. In order to avoid the allocation of additional memory, these bits are stored in the output image, in the top-right pixel label slot, that would otherwise be left unused until `FinalLabeling`, when connectivity information is not necessary anymore. The `Reduce` kernel then reads this connectivity information and performs the necessary union operations.

Both BKE and BUF make use of the *inline compression*, an optimization designed for the use of union-find in data parallel environments. Previously introduced and described in Sec. 2.1, it aims at reducing the size of equivalence trees created and updated during `Initialization` and `Merge/Reduction` phases. Thus, possible concurrent threads performing the find on the same node may save memory accesses.

4 BENCHMARKING

The term *reproducible research* is usually referred to as the approach of presenting scientific claims together with all information needed to reproduce the results, so that others may verify the findings and build upon them. As researchers, we strongly believe in and support reproducibility in scientific research. For this reason, in 2016 we publicly released the first version of YACCLAB [63] —Yet Another Connected Components LABELing Benchmark—, an open-source benchmarking system that allows to fairly compare CCL algorithms on different environments. YACCLAB has been originally described in [48] and later extended in [64] that included more literature algorithms and additional evaluation criteria, making it more general and flexible. In 2019, a new version of the benchmark including the possibility of evaluating 3D algorithms and GPU-based

TABLE 1

Comparison between state-of-the-art algorithms, which considers publication year, authors, minimum block of pixels processed together (i.e., single pixel, block or run), processing type (i.e., iterative or direct), connectivity, input type (2D or 3D) and data structures employed in addition to the output image L . For each additional data structure its dimension in bytes is also specified in the form $width \times height (\times depth) \times bytes$. These structures have the same name used in the algorithm description. Some of the iterative algorithms require an additional byte (Termination Byte, TB) to check whether the process is completed or not after each iteration. Please note that bytes are specified considering 32-bit labels. The space complexity of each algorithm can be easily deduced by summing the last table column (additional data structures) to common output requirements, i.e., the size in bytes of the output image L .

Name	Year	Authors	Block Size	Iterative vs Direct	Connectivity	Input	Additional Data Structure(s) [bytes]
NP [51]	2010	Hawick <i>et al.</i>	Pixel	Iterative	4, 8	2D	
LNP [51]	2010	Hawick <i>et al.</i>	Pixel	Iterative	4, 8	2D	
DPL [51]	2010	Hawick <i>et al.</i>	Run	Iterative	4	2D	
LE [51]	2010	Hawick <i>et al.</i>	Pixel	Iterative	4, 8	2D	$H = w \times h \times 4$ TB
UF [41]	2010	Oliveira and Lotufo	Pixel	Direct	4	2D	
BRB [54]	2011	Chen <i>et al.</i>	Block & Run	Iterative	8	2D	
OLE [52]	2011	Kalentev <i>et al.</i>	Pixel	Iterative	8	2D	TB
STAVA [53]	2011	Stava and Benes	Pixel	Iterative	8	2D	
RASMUSSEN [56]	2013	Rasmusson <i>et al.</i>	Pixel	Iterative	8	2D	
8DLS [55]	2014	Soh <i>et al.</i>	Pixel	Iterative	8	2D	TB
M8DLS [55]	2014	Soh <i>et al.</i>	Pixel	Iterative	8	2D	
ACCL [59]	2014	Paravecino and Kaeli	Pixel	Direct	4	2D	
LBUF [50]	2015	Yonehara and Aizawa	Pixel	Direct	4, 8	2D	
KE [58]	2015	Komura	Pixel	Direct	4	2D	
BE [46]	2016	Zavalishin <i>et al.</i>	Block	Iterative	8	2D, 3D	$BL = \frac{w}{2} \times \frac{h}{2} \times 4$ $BA = \frac{w}{2} \times \frac{h}{2} \times 1$ TB
DLP [25]	2017	Cabaret <i>et al.</i>	Pixel	Direct	8	2D	
HA8 [24]	2018	Hennequin <i>et al.</i>	Run	Direct	4, 8	2D	
KE [57]	2018	Allegretti <i>et al.</i>	Pixel	Direct	8	2D	
UF [57]	2018	Allegretti <i>et al.</i>	Pixel	Direct	8	2D, 3D	
C_SAUF [47]	2019	Allegretti <i>et al.</i>	Pixel	Direct	8	2D	
C_BBDDT [47]	2019	Allegretti <i>et al.</i>	Block	Direct	8	2D	
C_DRAG [47]	2019	Allegretti <i>et al.</i>	Block	Direct	8	2D	
BUF [36]	2019	Allegretti <i>et al.</i>	Block	Direct	8	2D, 3D	
BKE [36]	2019	Allegretti <i>et al.</i>	Block	Direct	8	2D, 3D	

implementations has been released [36]. The benchmark represents nowadays the *de-facto* standard for evaluating CCL proposals, being it employed by several authors [33], [65], [66], [67].

When measuring the performance of an algorithm, many different subtleties must be taken into account. Although datasets have been blamed for narrowing the focus of research and reducing it to a single benchmark performance number, it is clear that the ability to compare different techniques on the same data is essential to choose which algorithm suits specific needs best [68]. In order to ensure a common and standard test set, YACCLAB is provided with an open-source binary dataset, which covers most of the CCL application scenarios, including both 2D images and 3D volumes. Since the whole project is open source, anyone can contribute by increasing and improving the YACCLAB dataset with different application environments. The current set of images available in the benchmark and used to compare algorithm performance in this survey is described in Sec. 4.1. Before the publication of YACCLAB, many novel proposals have been published and almost none of them compared on the same data [27], [69] making performance comparison almost impossible.

Benchmarking may be a problem by itself, because measuring performance may not be obvious, but CCL has an



Fig. 5. Samples of the YACCLAB 2D datasets. From left to right, top to bottom: 3DPeS, Fingerprints, Medical, MIRflickr, Tobacco800, XDOCS, and Hamlet.

exact result and this reduces the burden of the evaluator to the measurement of how fast an algorithm is. However, setting the correct parameters when reproducing other people tests may be a problem, changing the final figures by orders of magnitude. This is why publishing the source code together with a scientific paper or, at the very least, an executable should be mandatory. YACCLAB tackles these aspects by including the source code of all the most significant proposals of the last twenty years, considering both




Fig. 6. Samples of the YACCLAB 3D datasets. From left to right: Hilbert space-filling curve, OASIS, and Mitochondria medical imaging data.

CPU- and GPU-based algorithms. All the implementations have been taken from the original publication, whenever available, or produced by following the paper pseudo-code and description. Being the code open source, anyone can check whether implementation details are correctly handled or not, and test algorithms on their own setting (operating system, compiler, dataset, etc.), verifying any claim found in the literature and selecting the algorithm that performs the best on their specific environment.

With this paper, YACCLAB has been further extended to include many of the above-described algorithms that were previously missing in the benchmarking suites. In the following of this Section, datasets, and evaluation criteria employed to later compare the algorithms are summarized.

4.1 Evaluation Dataset

Following a common practice in literature [25], [26], [46], the YACCLAB dataset includes both synthetic and real images, addressing both 2D and 3D scenarios. All images are provided in 1 bit per pixel PNG format, with 0 being background and 1 being foreground. The dataset can be automatically downloaded during the setup of the YACCLAB benchmark.

4.1.1 2D Datasets (Fig. 5)

MIRflickr [70]. This is the Otsu-binarized [71] version of the MIRflickr dataset. It contains a set of 25 000 natural images taken from Flickr with few connected components and an average density of 0.4459 foreground pixels.

Medical [72]. This dataset covers applications of CCL on medical data. It is composed of 343 binary histological images with an average amount of 1.21 million pixels to analyze and 484 components to label. This dataset is intended to cover applications of CCL related to medical image analysis.

Hamlet. A set of 104 images, scanned from the Hamlet provided by the Gutenberg Project [73]. Images have an average amount of 1 447 components to label and an average foreground density of 0.0789.

Tobacco800. Composed of 1 290 document images, it has been collected and scanned using a wide variety of equipment over time. Image sizes range from $1\,200 \times 1\,600$ to $2\,500 \times 3\,200$ pixels [74], [75], [76].

XDOCS. A collection of Italian civil registries recently digitalized [17], [77], [78]. This dataset is composed of 1 677 high resolution images with an average size of $4\,853 \times 3\,387$ and 15 282 components to analyze. The average foreground density is quite low: 0.0918.

Fingerprints. This dataset counts 960 fingerprint images taken from three fingerprint verification competitions




Fig. 7. Samples of the YACCLAB 2D synthetic dataset. Example images have a foreground density of 30% and, from left to right, top to bottom, granularities are 1, 2, 4, 6, 8, 12, 14, 16.

(FCV2000, FCV2002, and FCV2004) [79]. Images were collected using low-cost optical sensors or synthetically generated, binarized using an adaptive threshold [80], and finally negated. Their size varies from 240×320 up to 640×480 pixels.

3DPes. This is a set of images coming from the surveillance dataset 3DPeS (3D People Surveillance Dataset) [81]. The background models for all cameras are provided by the authors, so a very basic technique of motion segmentation has been applied to generate the foreground binary masks, i.e., background subtraction and Otsu thresholding [71]. The analysis of the foreground masks to remove small connected components and to match the nearest neighbors is a common application for CCL.

4.1.2 3D Datasets (Fig. 6)

OASIS. A MRI data collection taken from the Open Access Series of Imaging Studies (OASIS) project [82]. The 373 volumes of $128 \times 256 \times 256$ voxels were binarized with the Otsu threshold [71] calculated over the entire volume.

Mitochondria. This is the Electron Microscopy Dataset available in [83] and originally published in [84]. The original dataset contains sections taken from the CA1 hippocampus region of the brain. All the binary volumes provided by the authors have been included in YACCLAB, for a total of three volumes composed of 165 slices with a resolution of $1\,024 \times 768$ pixels.

Hilbert. The Hilbert curve is a fractal space-filling curve described by David Hilbert. This curve represents a challenging test case for the labeling algorithms, and YACCLAB 3D-dataset includes six volumes filled with the 3D Hilbert curve obtained at different iterations (1 to 6) of the construction method. Final binary volumes have a size of $128 \times 128 \times 128$.

4.1.3 Random Synthetic Images

A set of black and white random noise images to stress how the behavior of algorithms varies with respect to the *density* (percentage of foreground pixels) and *granularity* (minimum size of foreground blocks). Specifically, two different test sets are available in YACCLAB, one for 2D images and one for 3D volumes. The former contains $2\,048 \times 2\,048$ images, while the latter has $256 \times 256 \times 256$ volumes. Images and volumes have been generated with the Mersenne Twister MT19937

random number generator, included in the C++ standard library [85]. Density ranges from 0% to 100% with a step of 1%. For every density value, each integer granularity $g \in [1, 16]$ has been considered. Ten images have been generated for every couple of density-granularity values, for a total of 16 160 images, both for 2D and 3D. Samples from the YACCLAB 2D synthetic dataset are reported in Fig. 7, considering different levels of granularity. In other words, granularity corresponds to the minimum size of (non-overlapped) foreground blocks that appear in an image. When changing the granularity, it also changes the shape that foreground objects may assume, introducing additional difficulties (that vary from algorithm to algorithm) in the labeling procedure.

4.2 Assessment Strategies

The YACCLAB benchmark provides different tests that can be used to stress algorithm behavior under different perspectives. The simplest evaluation mechanism consists of a direct comparison of execution times, including the time for allocating data structures. Algorithms are run on all the available images, and the average execution time per dataset is recorded.

The benchmark also allows to separately evaluate the different phases composing the algorithms, separating the memory allocation time from the core execution of the labeling procedure. Two-stage algorithms can also distinguish between the *Local Labeling* and the *Tiles Merging* steps. The impact of each phase can thus be stressed and analyzed through such experiments, highlighting strengths and weaknesses.

Granularity is the last experiment available and consists of measuring the execution time on synthetic images when foreground pixel density and granularity change. This approach allows us to draw additional conclusions, highlighting behaviors that may not emerge from the previous tests.

All the experiments carried out with YACCLAB on GPU algorithms assume that the input image is already available in the GPU memory before the algorithm begins. The rationale behind such an evaluation choice is that the algorithms' execution time does not depend on the data transfer between the host and the GPU. Comparing performance by including the data transfer time will simply introduce a fixed equal "delay", distracting from the CCL execution time. In the same way, YACCLAB assumes that the output is required in the device memory and no additional data transfer is needed. The allocation of the output GPU image and of any other data structure potentially required by the algorithm itself is instead considered in the total execution time.

5 EVALUATION

According to the *Gestalt psychology* of perception, our senses operate the law of closure, perceiving objects as a whole even if they are loosely connected as for the 8-connectivity [42]. This is the reason why many of the latest CCL proposals focus on 8-connectivity. Based on this observation, and considering the amount of environments already covered by our analysis, we focus our experiments

on the 8-connectivity for 2D images. On 3D volumes, 8-connectivity maps to 26-connectivity, which is employed in our analysis of 3D implementations.

As previously mentioned, many of the literature solutions considered in this survey have no associated public source code, forcing us to re-implement them. All the algorithms are identified by their acronym already mentioned in Sec. 3.

5.1 Hardware and Software

The comparative evaluation is carried out on two different machines, using different experimental setups.

Setup A. This test machine is equipped with an Intel(R) Core(TM) i7-4790 CPU @ 3.60GHz (with 4×32 KB L1 cache, 4×256 KB L2 cache, 8 MB L3 cache) and two GPUs: an NVIDIA Quadro P1000 and an NVIDIA Quadro K2200, both with 4 GB of GDDR5 GPU Memory and 640 CUDA Cores. The former has a *Pascal* architecture, while the latter is based on *Maxwell* architecture. In such environments, two different operating systems have been employed: Windows 19.11.25508.2 and Ubuntu 22.04 LTS, respectively combined with MSVC 19.11.25508.2 and GCC 9.3.0 compilers. In both cases, NVCC 11.0 has been used to compile CUDA files. NVIDIA Driver versions 522.06 and 450.51 have been employed on Windows and Linux, respectively.

Setup B. The second experimental setup is a server machine equipped with 52 Intel(R) Xeon(R) Gold 5320 CPU @ 2.20GHz (with 48 KB L1d cache, 32 KB L1i cache, 1 280 KB L2 cache, and 40 MB L3 cache) and an NVIDIA A100 80GB PCIe GPU based on *Ampere* architecture with 80 GB of HBM2 with ECC GPU Memory and 6912 CUDA Cores (NVIDIA Drivers version 525.60.13). In this case, Ubuntu 22.04 LTS has been combined with GCC 8.5.0 and NVCC 12.0 compilers.

The kernel sensitivity to run-time parameters such as thread block and grid dimensions is not discussed. Although these parameters can significantly affect performance in CUDA programs, we stick to those provided in the original papers by verifying that such configurations provide the best results also on our environments.

By targeting different operating systems and different GPU architectures, our analysis aims to demonstrate how these variations affect algorithm performance.

5.2 Comparative Evaluation

Overall Comparison. To begin with, Tab. 2 and Tab. 3 provide an overall comparison based on the average total execution time and the average throughput in MPixel/ms, targeting all the considered environments. Some algorithms lack a 3D implementation, so results on 3D volumes cannot be reported.

As previously noticed in [64] for CPU-based CCL implementations, the performance of the same algorithm run on different operating systems may be drastically different, especially when memory allocation becomes important, i.e., when a dataset contains high-resolution images, or the algorithm requires larger data structures. There might be plenty of reasons behind this performance gap, including the difference in virtual page commits that Windows and Linux have, or the fact that under Windows it is almost

TABLE 2

Results obtained under Windows 10 running a Quadro K2200 and a Quadro P1000 NVIDIA GPUs (Setup A). White columns report average (per image) run-time results in ms (\downarrow) while gray columns represent the average throughput of each algorithm in MPixel/ms (\uparrow). The bold values represent the best algorithm on the specific environment/dataset. The * identifies direct algorithms.

		2D Images											3D Volumes								
		3DPeS		Fingerprints		Hamlet		Medical		MIRflickr		Tobacco800		XDOCS		Hilbert		Oasis		Mitochondria	
(a) <i>Quadro K2200</i>	M8DLS	21.757	0.019	14.245	0.010	127.800	0.021	38.200	0.032	20.180	0.009	331.300	0.014	3150.550	0.005						
	RASMUS.	5.537	0.074	4.428	0.032	117.100	0.023	9.169	0.132	5.463	0.033	249.293	0.018	2987.948	0.006						
	8DLS	3.341	0.123	7.139	0.020	37.988	0.071	18.028	0.067	34.826	0.005	185.789	0.025	4732.362	0.003						
	OLE	0.958	0.428	0.612	0.229	5.215	0.520	2.806	0.431	0.698	0.258	7.985	0.576	34.102	0.484						
	STAVA	0.829	0.495	0.620	0.226	4.077	0.665	2.682	0.451	0.916	0.197	5.532	0.832	21.853	0.755						
	LBUF*	0.602	0.681	0.306	0.457	3.214	0.843	1.978	0.612	0.293	0.613	4.423	1.040	18.105	0.911						
	DLP*	0.662	0.620	0.260	0.538	3.325	0.815	1.807	0.670	0.328	0.549	5.375	0.856	19.542	0.844						
	HAS*	0.599	0.685	0.228	0.615	2.746	0.987	1.493	0.810	0.223	0.806	4.160	1.106	14.976	1.101						
	KE*	0.565	0.726	0.254	0.551	2.929	0.925	1.735	0.697	0.260	0.693	4.207	1.093	16.954	0.973						
	C_SAUF*	0.531	0.772	0.234	0.598	2.834	0.956	1.585	0.763	0.254	0.708	3.891	1.182	15.522	1.062						
	C_BBDT*	0.514	0.798	0.224	0.624	2.424	1.118	1.318	0.918	0.240	0.749	3.461	1.329	13.087	1.260						
	C_DRAG*	0.506	0.810	0.215	0.652	2.402	1.128	1.263	0.958	0.209	0.862	3.407	1.350	12.797	1.289						
	BE*	1.020	0.402	0.694	0.202	3.781	0.717	2.029	0.596	0.924	0.195	5.283	0.871	18.765	0.879	4.786	2.622	9.540	327.982	90.104	4.320
	UF*	0.598	0.686	0.307	0.457	3.312	0.818	2.213	0.547	0.391	0.461	4.684	0.982	19.115	0.863	3.492	3.594	17.798	175.803	132.286	2.943
BUF*	0.494	0.830	0.207	0.677	2.252	1.203	1.373	0.881	0.223	0.806	3.300	1.394	12.883	1.280	2.105	5.962	6.481	482.788	67.326	5.782	
BKE*	0.487	0.842	0.196	0.714	2.127	1.274	1.216	0.995	0.178	1.010	3.214	1.431	11.688	1.411	2.714	4.624	8.259	378.853	102.297	3.805	
(b) <i>Quadro P1000</i>	M8DLS	13.109	0.031	8.583	0.016	77.058	0.035	23.067	0.052	12.163	0.015	199.626	0.023	1898.248	0.009						
	RASMUS.	3.569	0.115	2.888	0.048	67.363	0.040	5.387	0.225	3.467	0.052	140.871	0.033	1667.878	0.010						
	8DLS	2.995	0.137	5.361	0.026	30.183	0.090	13.684	0.088	23.522	0.008	149.902	0.031	3888.757	0.004						
	OLE	0.991	0.414	0.482	0.290	4.390	0.617	2.244	0.539	0.568	0.317	6.718	0.685	28.536	0.578						
	STAVA	0.689	0.595	0.405	0.346	2.881	0.941	1.966	0.615	0.676	0.266	3.749	1.227	14.483	1.139						
	LBUF*	0.515	0.796	0.179	0.782	2.386	1.136	1.532	0.790	0.192	0.938	3.077	1.495	12.718	1.297						
	DLP*	0.544	0.754	0.120	1.167	2.183	1.241	1.228	0.985	0.177	1.017	3.344	1.376	12.023	1.372						
	HAS*	0.481	0.852	0.107	1.308	1.930	1.404	1.135	1.066	0.133	1.353	2.692	1.709	9.765	1.689						
	KE*	0.511	0.802	0.137	1.022	2.320	1.168	1.408	0.859	0.169	1.065	3.182	1.446	12.886	1.280						
	C_SAUF*	0.480	0.854	0.131	1.069	2.262	1.198	1.310	0.924	0.182	0.989	2.911	1.580	11.727	1.406						
	C_BBDT*	0.468	0.876	0.106	1.321	1.927	1.406	1.091	1.109	0.145	1.241	2.591	1.775	9.748	1.692						
	C_DRAG*	0.466	0.880	0.102	1.373	1.927	1.406	1.049	1.153	0.130	1.385	2.559	1.798	9.581	1.721						
	BE*	0.775	0.529	0.465	0.301	2.952	0.918	1.703	0.711	0.481	0.374	4.151	1.108	14.324	1.151	3.459	3.628	7.023	445.529	58.013	6.710
	UF*	0.483	0.849	0.156	0.897	2.311	1.173	1.568	0.772	0.237	0.759	3.083	1.492	12.498	1.319	2.379	5.275	13.939	224.474	95.871	4.060
BUF*	0.450	0.911	0.101	1.386	1.766	1.535	1.120	1.080	0.146	1.233	2.433	1.891	9.496	1.737	1.391	9.022	4.656	672.025	43.923	8.863	
BKE*	0.442	0.928	0.089	1.573	1.672	1.621	0.997	1.214	0.104	1.731	2.394	1.921	8.595	1.919	1.753	7.159	5.605	558.243	65.939	5.904	

impossible to prevent the GUI from accessing the GPU during tests. A thorough explanation of this phenomenon would require deep experiments and goes beyond the goal of the paper. The conclusion that deserves to be raised is that the execution time ratio Linux/Windows is almost constant for different algorithms, meaning that changing the operating system does not impact on algorithm behavior.

When focusing on 2D images, the overall performance of the algorithms and their mutual ranking are not affected by dataset intrinsic characteristics. Of course, the total execution time is deeply influenced by image size and complexity (i.e., the number of connected components it contains and their geometry). XDOCS is certainly the hardest dataset to label and it holds the higher average execution time.

Experimental results also reveal that the advantages of the “M” variation proposed for 8DLS are negligible on real-case datasets, unless images are very large and the majority of connected components is small, as it happens within XDOCS. In all the other cases, the additional checks required by M8DLS ruin the overall algorithm performance.

Another conclusion that can be drawn is that (M)8DLS and RASMUSON algorithms are significantly slower than others, so they can be ignored in the rest of the analysis.

Different GPU architectures may also impact algorithm performance, changing the figures by an order of magnitudes, as it can be appreciated in Tab. 3 when comparing algorithms execution on Quadro GPUs with the newer A100 80GB. Apart from the overall performance, related to computational capabilities, memory hierarchies, and bandwidths, each GPU architecture may have specific optimizations that can be applied to improve performance.

However, none of the implementations analyzed with this paper and provided in YACCLAB explicitly leverage on GPU specific optimization and this is confirmed also by our experiments: selected a dataset, the execution time ratio of the algorithms run on different GPUs is constant. However, the performance gap between different algorithmic proposals is negligible when using newer architecture and smaller datasets. A deeper analysis is hence provided in the following of this section.

The abovementioned tables also suggest that, regardless of the operating system and the GPU architecture (at least for those considered in our experiments), BKE, currently the only GPU-based algorithm available in the OpenCV library [86], is always the best-performing algorithm on 2D datasets. For what concerns 3D volumes, BUF usually demonstrates better results, except when running on the A100 GPU where BKE is still competitive. The advantages provided by the use of a block-based approach, which allows for a considerable reduction in the number of memory accesses, are clearly demonstrated by the superior performance of BKE and BUF.

On the Role of Inline Compression (IC). The better results obtained by BUF compared to BKE on 3D datasets are explained by considering the IC optimization both the algorithms make use of (see Sec. 2.1 and Alg. 1 as a reference). This optimization is intended for reducing the number of parent nodes a thread has to traverse when searching for the root of union-find equivalence trees, saving memory accesses and time. These savings come at the expense of additional write operations, which could nullify the improvement. The net improvement of the IC optimization

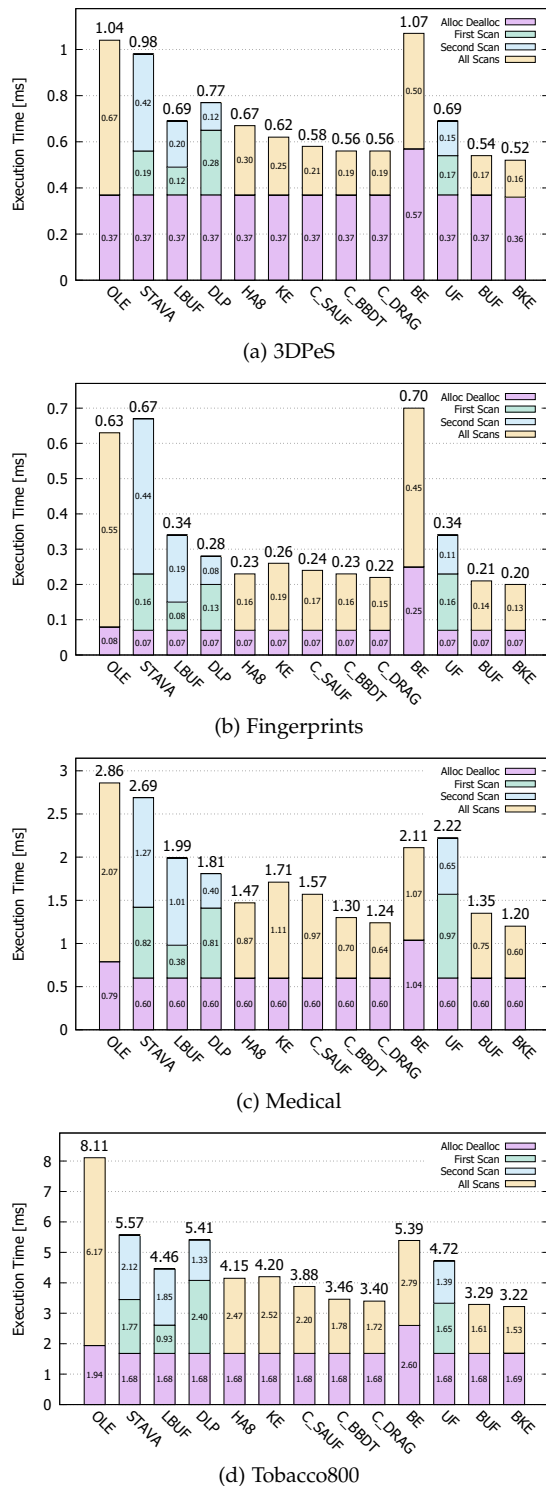


Fig. 8. Average run-time with steps test on 2D datasets under Windows using NVIDIA QUADRO K2200 (Setup A). Numbers are given in ms. Lower is better. Best viewed online.

number of connected components is small, as it happens in 3DPeS and Fingerprints, the benefit of using additional data structures is annihilated by the allocation time.

Direct algorithms, which perform a fixed number of kernel calls, tend to have better performance than iterative ones. UF is faster than OLE, STAVA, and BE on all datasets but Medical, and the same goes for its optimization, LBUF and KE, the former using better block size and the latter

with a more efficient initialization phase. DLP has a very similar approach to UF, since it also heavily relies on union-find, even implementing it slightly differently; unsurprisingly, its performance is comparable to UF.

The block-based approach proposed with BUF and BKE allows to leverage the advantages of blocks to further improve UF and KE, without introducing allocation drawbacks, making the two algorithms the fastest available in literature for both 2D and 3D scenarios. As already mentioned, BKE has been available in OpenCV since version 4.6.0 for labeling connected components on GPU devices.

How does Decision Trees Perform on GPUs? The algorithms based on decision trees, i.e., C_SAU, C_BBDT, and C_DRAG, have performance close to the state of the art, with C_DRAG always being the best among the three. Although irregular patterns of execution slow down the thread scheduling and make the decision tree the worst thing to use in GPUs, when working with “natural” images the patterns within blocks are more or less uniform, making the thread divergence not really frequent. This is also confirmed by the granularity experiments reported in Fig. 10. When granularity is 1, the computational time of both C_BBDT and C_DRAG explodes around 50% of density, in accordance with expectations related to branch divergence. These algorithms have performance close to state-of-the-art GPU-specific designs only when density is below 10% or over 90%. Indeed, when foreground densities are in such ranges the action to be performed on the current pixel is selected in the first levels of the decision trees employed by the algorithms, saving many memory accesses without ruining the thread execution flow. Since images on which CCL is usually applied are in that range of densities, the performance of C_SAU, C_BBDT, and C_DRAG highlighted in the previous charts are easily explainable taking into account granularity tests. Additionally, it can be noticed that as the granularity increases (Fig. 10b, Fig. 10c, and Fig. 10d), smaller portions of the image have irregular patterns and require to explore deeper tree levels, while the vast majority tends to be all white or all black, again minimizing thread divergence and maximizing the tree performance.

About Foreground Density. Thanks to an efficient combination of CUDA intrinsics to find runs of foreground pixels and union-find operations to merge runs of the same CC, HA8 performs close to state of the art on many datasets, especially when connected components extend in the horizontal direction (Fingerprints).

Taking into account Fig. 10, we can also notice that the performance trend of UF and BUF when varying the foreground density is similar although shifted along the y-axis since they share the same basic structure. A similar conclusion can be drawn when comparing KE and BKE.

The disadvantages of iterative approaches that do not leverage “auxiliary” data structures to record and store label equivalences are clearly visible with DLP. Independently of the granularity, the execution time of DLP (almost) exponentially increases with the number of foreground pixels.

For the sake of completeness, Fig. 11 depicts the performance of the 3D algorithms when tested over randomly generated three-dimensional volumes. The same considerations drawn for the two-dimensional cases can be applied.




Fig. 9. Average run-time with steps test on 3D datasets. The experiment ran under Windows with NVIDIA QUADRO K2200 (Setup A). Lower is better. Best viewed online.

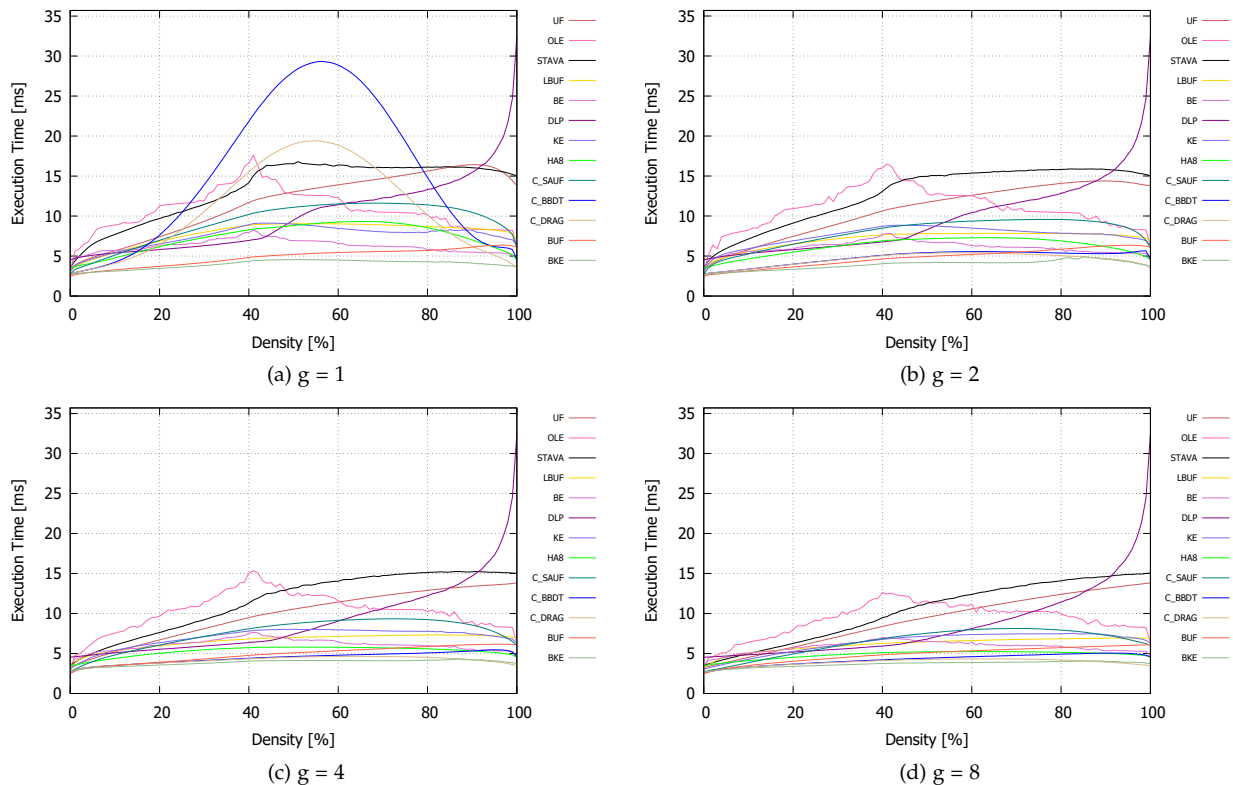


Fig. 10. Granularity tests on 2D randomly generated images. Numbers are given in ms. Lower is better. Best viewed online.

Moreover, it can be highlighted that the previously discussed negative effects of IC on BKE become less relevant as density increases. The performance of BKE and BUF are almost the same above $\sim 70\%$ and $\sim 50\%$ of foreground density, respectively with granularity 1 and 8. When granularity is 1, storing the block adjacency information in a separate data structure becomes valuable —i.e., the additional memory allocation time is compensated by the faster iterative scan approach— only if foreground density is higher than 90%, making BE slightly faster than both BUF and BKE in such scenario.

On the Effects of the Input Size. To complete our analysis, the effect of input size on execution time is depicted in Fig. 12. In this experiment, each algorithm is tested on random images with a foreground density of 0.5 and a size increasing from 10^3 to 10^7 pixels. The granularity is set to 1. In order to simplify the chart, C_DRAG has been chosen as the representative of C_BBDT and C_SAUf; the three algorithms employ the same core approach and have comparable performances. To focus on interesting algorithms,

the y-axis has been cropped.

Confirming previous conclusions, 8DLS, M8DLS, and RASMUSSESON perform much worse than other algorithms, regardless of the input dimension. These algorithms require multiple iterations that are inefficient.

When the input image contains few pixels (e.g., $< 10^5$), the execution time is dominated by the allocation time of the output image and auxiliary data structures. In this regard, we can say that all the algorithms have similar performance, except for BE and OLE, which need to allocate more memory, degrading the overall performance.

However, with larger input sizes the running time of all the algorithms grows (almost) linearly. As already noticed in Fig. 11, BE becomes more competitive as the complexity (the size in this case) of the input image increases, compensating for the additional allocation time.

In this experiment, C_DRAG is significantly slower than other direct algorithms, because the input image density of 0.5 causes the maximum branch divergence, as already discussed with granularity tests.




Fig. 11. Granularity tests on 3D randomly generated images (Setup A). Numbers are given in ms. Lower is better. Best viewed online.

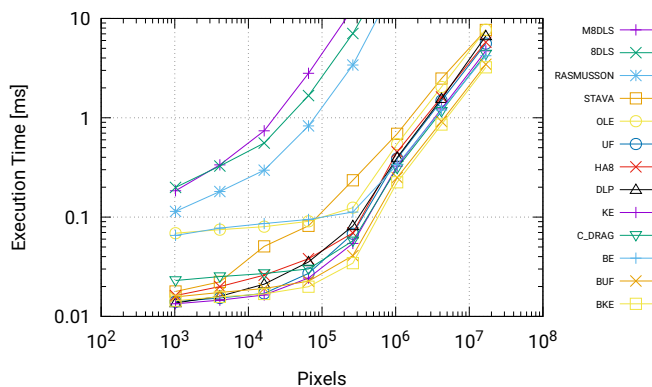


Fig. 12. Correlation between algorithm execution time and the growth of input size (Setup B). Lower is better. Best viewed online.

6 CONCLUDING REMARKS AND FUTURE WORK

This article reviewed the state-of-the-art connected components labeling algorithms designed for GPUs and published in the last decade, focusing on the main strategies adopted and analyzing their performance in different scenarios. Experimental results conducted on different operating systems, using different compilers and GPU architectures revealed that BKE is the best-performing algorithm for 2D datasets including text images, fingerprints, medical images, and surveillance datasets. When working on 3D volumes the performance of BKE is still impressive, but BUF demonstrates generally better capabilities, especially when working with hard-to-label volumes such as the Mitochondria dataset. We also showed that algorithms specifically designed for sequential architecture can be extremely effective when adapted to GPUs.

The source code of all the analyzed algorithms is collected in the YACCLAB benchmark and is provided along with this article [63]. It is worth mentioning that we have been forced to re-implement many of the algorithm solutions published in literature since the original implementation was not publicly available and some authors did not answer to our explicit requests. In any case, the code is now public so anyone can verify that our implementations are aligned with the description in the related publications.

If we consider that the allocation time is “fixed” for a given environment and that the bare minimum of any CCL algorithm is the copy of the input into the output image

(an operation that is not enough to complete the task, but strictly required) the reader can certainly spot that space for further optimization is really small, extremely so.

Future work should focus on the Connected Components Analysis (CCA) by integrating the calculation of connected components statistics in state-of-the-art solutions.

Another open research direction is the optimization of CCL on embedded systems, including FPGA. Although specific algorithmic solutions should be devised, reusing the core innovation of GPU-based solutions could benefit also such scenarios.

ACKNOWLEDGMENT

This project has received funding from the Department of Engineering “Enzo Ferrari” of the University of Modena through FARD-2023 (Fondo di Ateneo per la Ricerca 2023).

REFERENCES

- [1] L. Cabaret, L. Lacassagne, and D. Etiemble, “Parallel light speed labeling: An efficient connected component labeling algorithm for multi-core processors,” *Journal of Real-Time Image Processing*, vol. 15, no. 1, 2018.
- [2] P. Asad, R. Marroquim, and A. L. e. L. Souza, “On GPU Connected Components and Properties: A Systematic Evaluation of Connected Component Labeling Algorithms and Their Extension for Property Extraction,” *IEEE Transactions on Image Processing*, vol. 28, no. 1, 2019.
- [3] A. Dubois and F. Charpillet, “Tracking Mobile Objects with Several Kinects using HMMs and Component Labelling,” in *Workshop Assistance and Service Robotics in a human environment, International Conference on Intelligent Robots and Systems*, 2012.
- [4] C. Zhan, X. Duan, S. Xu, Z. Song, and M. Luo, “An Improved Moving Object Detection Algorithm Based on Frame Difference and Edge Detection,” in *Fourth International Conference on Image and Graphics (ICIG)*. IEEE, 2007.
- [5] A. Abramov, T. Kulvicius, F. Wörgötter, and B. Dellen, “Real-Time Image Segmentation on a GPU,” in *Facing the multicore-challenge*, 2010.
- [6] F. Pollastri, F. Bolelli, R. Paredes, and C. Grana, “Improving Skin Lesion Segmentation with Generative Adversarial Networks,” in *2018 IEEE 31st International Symposium on Computer-Based Medical Systems (CBMS)*. IEEE, 2018.
- [7] A. Körbes, G. B. Vitor, R. de Alencar Lotufo, and J. V. Ferreira, “Advances on Watershed Processing on GPU Architecture,” in *International Symposium on Mathematical Morphology and Its Applications to Signal and Image Processing*. Springer, 2011.
- [8] E. Gomel, T. Shaharabany, and L. Wolf, “Box-based Refinement for Weakly Supervised and Unsupervised Localization Tasks,” in *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023.
- [9] Y. Zhang, S. Wu, N. Snavely, and J. Wu, “Seeing a Rose in Five Thousand Ways,” in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.

- [10] I. H. Laradji, N. Rostamzadeh, P. O. Pinheiro, D. Vazquez, and M. Schmidt, "Where are the Blobs: Counting by Localization with Point Supervision," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.
- [11] A. Eklund, P. Dufort, M. Villani, and S. LaConte, "BROCCOLI: Software for fast fMRI analysis on many-core CPUs and GPUs," *Frontiers in neuroinformatics*, vol. 8, 2014.
- [12] H. V. Pham, B. Bhaduri, K. Tangella, C. Best-Popescu, and G. Popescu, "Real Time Blood Testing Using Quantitative Phase Imaging," *PLoS one*, vol. 8, no. 2, 2013.
- [13] F. Pollastri, F. Bolelli, R. Paredes, and C. Grana, "Augmenting Data with GANs to Segment Melanoma Skin Lesions," *Multimedia Tools and Applications*, vol. 79, no. 21-22, 2019.
- [14] G. Bontempo, A. Porrello, F. Bolelli, S. Calderara, and E. Ficarra, "DAS-MIL: Distilling Across Scales for MIL Classification of Histological WSIs," in *Medical Image Computing and Computer Assisted Intervention – MICCAI 2023*, Oct 2023.
- [15] B. E. Bejnordi, M. Veta, P. J. Van Diest, B. Van Ginneken, N. Karssemeijer, G. Litjens, J. A. Van Der Laak, M. Hermsen, Q. F. Manson, M. Balkenhol *et al.*, "Diagnostic Assessment of Deep Learning Algorithms for Detection of Lymph Node Metastases in Women With Breast Cancer," *Jama*, vol. 318, no. 22, 2017.
- [16] T. Lelore and F. Bouchara, "FAIR: A Fast Algorithm for Document Image Restoration," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 8, 2013.
- [17] F. Bolelli, "Indexing of Historical Document Images: Ad Hoc Dewarping Technique for Handwritten Text," in *Italian Research Conference on Digital Libraries*. Springer, 2017.
- [18] T. Berka, "The Generalized Feed-forward Loop Motif: Definition, Detection and Statistical Significance," *Procedia Computer Science*, vol. 11, 2012.
- [19] M. J. Dinneen, M. Khosravani, and A. Probert, "Using OpenCL for Implementing Simple Parallel Graph Algorithms," in *Proceedings of the International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA)*. The Steering Committee of The World Congress in Computer Science, Computer Engineering and Applied Computing (WorldComp), 2011.
- [20] C. Grana, L. Baraldi, and F. Bolelli, "Optimized Connected Components Labeling with Pixel Prediction," in *Advanced Concepts for Intelligent Vision Systems*, vol. 10016, Oct 2016.
- [21] F. Bolelli, M. Cancilla, and C. Grana, "Two More Strategies to Speed Up Connected Components Labeling Algorithms," in *Image Analysis and Processing - ICIAP 2017*, vol. 10485, Sep 2017.
- [22] W. Lee, S. Allegretti, F. Bolelli, and C. Grana, "Fast Run-Based Connected Components Labeling for Bitonal Images," in *2021 Joint 10th International Conference on Informatics, Electronics & Vision (ICIEV) and 2021 5th International Conference on Imaging, Vision & Pattern Recognition (icIVPR)*, Aug 2021.
- [23] F. Bolelli, S. Allegretti, and C. Grana, "Connected Components Labeling on Bitonal Images," in *Image Analysis and Processing - ICIAP 2021*, May 2022.
- [24] A. Hennequin, L. Lacassagne, L. Cabaret, and Q. Meunier, "A new Direct Connected Component Labeling and Analysis Algorithms for GPUs," in *2018 Conference on Design and Architectures for Signal and Image Processing (DASIP)*. IEEE, 2018.
- [25] L. Cabaret, L. Lacassagne, and D. Etienne, "Distanceless label propagation: An efficient direct connected component labeling algorithm for GPUs," in *2017 Seventh International Conference on Image Processing Theory, Tools and Applications (IPTA)*, 2017.
- [26] D. P. Playne and K. Hawick, "A New Algorithm for Parallel Connected-Component Labelling on GPUs," *IEEE Transactions on Parallel and Distributed Systems*, vol. 29, no. 6, June 2018.
- [27] L. He, X. Zhao, Y. Chao, and K. Suzuki, "Configuration-Transition-Based Connected-Component Labeling," *IEEE Transactions on Image Processing*, vol. 23, no. 2, 2014.
- [28] L. He, Y. Chao, K. Suzuki, and K. Wu, "Fast connected-component labeling," *Pattern Recognition*, vol. 42, no. 9, 2009.
- [29] D. Windisch, C. Kaever, G. Juckeland, and A. Bieberle, "Parallel Algorithm for Connected-Component Analysis Using CUDA," *Algorithms*, vol. 16, no. 2, 2023.
- [30] M. Kowalczyk, P. Ciarach, D. Przewlocka-Rus, H. Szolc, and T. Kryjak, "Real-Time FPGA Implementation of Parallel Connected Component Labelling for a 4K Video Stream," *Journal of Signal Processing Systems*, vol. 93, no. 5, 2021.
- [31] L. He, X. Ren, Q. Gao, X. Zhao, B. Yao, and Y. Chao, "The connected-component labeling problem: A review of state-of-the-art algorithms," *Pattern Recognition*, vol. 70, 2017.
- [32] A. Rosenfeld and A. Kak, *Digital Picture Processing*, ser. Computer Science and Applied Mathematics. Academic Press, 1982, no. 1.
- [33] D. Zhang, H. Ma, and L. Pan, "A gamma-signal-regulated connected components labeling algorithm," *Pattern Recognition*, vol. 91, 2019.
- [34] F. Bolelli, S. Allegretti, and C. Grana, "One DAG to Rule Them All," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 7, Jan 2021.
- [35] F. Lemaitre, A. Hennequin, and L. Lacassagne, "How to speed Connected Component Labeling up with SIMD RLE algorithms," in *Proceedings of the 2020 Sixth Workshop on Programming Models for SIMD/Vector Processing*, 2020.
- [36] S. Allegretti, F. Bolelli, and C. Grana, "Optimized Block-Based Algorithms to Label Connected Components on GPUs," *IEEE Transactions on Parallel and Distributed Systems*, vol. PP, Aug 2019.
- [37] H. Zhou, R. Dou, L. Cheng, J. Liu, and N. Wu, "A Provisional Labels-Reduced, Real-Time Connected Component Labeling Algorithm for Edge Hardware," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 69, no. 6, 2022.
- [38] Z. Li, Q. Zhang, T. Long, and B. Zhao, "A parallel pipeline connected-component labeling method for on-orbit space target monitoring," *Journal of Systems Engineering and Electronics*, vol. 33, no. 5, 2022.
- [39] K. Bok, N. Kim, D. Choi, J. Lim, and J. Yoo, "Incremental Connected Component Detection for Graph Streams on GPU," *Electronics*, vol. 12, no. 6, 2023.
- [40] H. Samet, "Connected Component Labeling Using Quadrees," *Journal of the ACM (JACM)*, vol. 28, no. 3, 1981.
- [41] V. M. Oliveira and R. A. Lotufo, "A Study on Connected Components Labeling algorithms using GPUs," in *SIBGRAPI*, vol. 3, 2010.
- [42] C. Grana, D. Borghesani, and R. Cucchiara, "Optimized Block-based Connected Components Labeling with Decision Trees," *IEEE Transactions on Image Processing*, vol. 19, no. 6, 2010.
- [43] L. He, X. Zhao, Y. Chao, and K. Suzuki, "Configuration-Transition-Based Connected-Component Labeling," *IEEE Transactions on Image Processing*, vol. 23, no. 2, 2014.
- [44] D. J. C. Santiago, T. I. Ren, G. D. Cavalcanti, and T. I. Jyh, "Efficient 2x2 block-based connected components labeling algorithms," in *2015 IEEE International Conference on Image Processing (ICIP)*, 2015.
- [45] F. Bolelli, S. Allegretti, L. Baraldi, and C. Grana, "Spaghetti Labeling: Directed Acyclic Graphs for Block-Based Connected Components Labeling," *IEEE Transactions on Image Processing*, vol. 29, no. 1, October 2019.
- [46] S. Zavalishin, I. Safonov, Y. Bekhtin, and I. Kurilin, "Block Equivalence Algorithm for Labeling 2D and 3D Images on GPU," *Electronic Imaging*, vol. 2016, no. 2, 2016.
- [47] S. Allegretti, F. Bolelli, M. Cancilla, F. Pollastri, L. Canalini, and C. Grana, "How does Connected Components Labeling with Decision Trees perform on GPUs?" in *Computer Analysis of Images and Patterns*, vol. 11678, Sep 2019.
- [48] C. Grana, F. Bolelli, L. Baraldi, and R. Vezzani, "YACCLAB - Yet Another Connected Components Labeling Benchmark," in *23rd International Conference on Pattern Recognition*. ICPR, 2016.
- [49] M. B. Dillencourt, H. Samet, and M. Tamminen, "A General Approach to Connected-Component Labeling for Arbitrary Image Representations," *Journal of the ACM*, vol. 39, no. 2, 1992.
- [50] K. Yonehara and K. Aizawa, "A Line-Based Connected Component Labeling Algorithm Using GPUs," in *2015 Third International Symposium on Computing and Networking*, 2015.
- [51] K. A. Hawick, A. Leist, and D. P. Playne, "Parallel graph component labelling with GPUs and CUDA," *Parallel Computing*, vol. 36, no. 12, 2010.
- [52] O. Kalentev, A. Rai, S. Kemnitz, and R. Schneider, "Connected component labeling on a 2D grid using CUDA," *Journal of Parallel and Distributed Computing*, vol. 71, no. 4, 2011.
- [53] O. Štáva and B. Beneš, "Chapter 35 - Connected Component Labeling in CUDA," in *GPU Computing Gems Emerald Edition*, ser. Applications of GPU Computing Series, W. mei W. Hwu, Ed., 2011.
- [54] P. Chen, H. Zhao, C. Tao, and H. Sang, "Block-run-based connected component labelling algorithm for GPGPU using shared memory," *Electronics Letters*, vol. 47, Nov. 2011.
- [55] Y. Soh, H. Raja, Y. Hae, and I. Kim, "Fast Parallel Connected Component Labeling Algorithms Using CUDA Based on 8-directional Label Selection," *International Journal of Latest Research in Science and Technology*, vol. 3, 03 2014.

- [56] A. Rasmusson, T. Sørensen, and G. Ziegler, "Connected Components Labeling on the GPU with Generalization to Voronoi Diagrams and Signed Distance Fields," in *Advances in Visual Computing*, vol. 8033, 07 2013.
- [57] S. Allegretti, F. Bolelli, M. Cancilla, and C. Grana, "Optimizing GPU-Based Connected Components Labeling Algorithms," in *2018 IEEE International Conference on Image Processing, Applications and Systems (IPAS)*, 2018.
- [58] Y. Komura, "GPU-based cluster-labeling algorithm without the use of conventional iteration: Application to the Swendsen-Wang multi-cluster spin flip algorithm," *Computer Physics Communications*, vol. 194, 2015.
- [59] F. Nina Paravecino and D. Kaeli, "Accelerated Connected Component Labeling Using CUDA Framework," in *Computer Vision and Graphics*, vol. 8671, 09 2014.
- [60] K. Wu, E. Otoo, and K. Suzuki, "Two Strategies to Speed up Connected Component Labeling Algorithms," Lawrence Berkeley National Laboratory, Tech. Rep. LBNL-59102, 2005.
- [61] C. Grana, M. Montanero, and D. Borghesani, "Optimal decision trees for local image processing algorithms," *Pattern Recognition Letters*, vol. 33, no. 16, 2012.
- [62] F. Bolelli, L. Baraldi, M. Cancilla, and C. Grana, "Connected Components Labeling on DRAGs," in *2018 24th International Conference on Pattern Recognition (ICPR)*, Aug 2018.
- [63] The YACCLAB Benchmark. Accessed on 2022-12-21. [Online]. Available: <https://github.com/prittt/YACCLAB>
- [64] F. Bolelli, M. Cancilla, L. Baraldi, and C. Grana, "Toward reliable experiments on the performance of Connected Components Labeling algorithms," *Journal of Real-Time Image Processing*, 2018.
- [65] J. Chen, K. Nonaka, H. Sankoh, R. Watanabe, H. Sabirin, and S. Naito, "Efficient Parallel Connected Component Labeling With a Coarse-to-Fine Strategy," *IEEE Access*, vol. 6, 2018.
- [66] T. Chabardès, P. Dokládál, and M. Bilodeau, "A labeling algorithm based on a forest of decision trees," *Journal of Real-Time Image Processing*, vol. 17, no. 5, 2020.
- [67] F. Diaz-del Rio, P. Sanchez-Cuevas, H. Molina-Abril, and P. Real, "Parallel connected-Component-Labeling based on homotopy trees," *Pattern Recognition Letters*, vol. 131, 2020.
- [68] A. Torralba and A. A. Efros, "Unbiased Look at Dataset Bias," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2011.
- [69] W.-Y. Chang and C.-C. Chiu, "An efficient scan algorithm for block-based connected component labeling," in *22nd Mediterranean Conference of Control and Automation (MED)*. IEEE, 2014.
- [70] M. J. Huiskes and M. S. Lew, "The MIR Flickr Retrieval Evaluation," in *MIR '08: Proceedings of the 2008 ACM International Conference on Multimedia Information Retrieval*, 2008.
- [71] N. Otsu, "A Threshold Selection Method from Gray-Level Histograms," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 9, no. 1, 1979.
- [72] F. Dong, H. Irshad, E.-Y. Oh *et al.*, "Computational Pathology to Discriminate Benign from Malignant Intraductal Proliferations of the Breast," *PLoS one*, vol. 9, no. 12, 2014.
- [73] The Hamlet Dataset. Accessed on 2019-03-21. [Online]. Available: <http://www.gutenberg.org>
- [74] G. Agam, S. Argamon, O. Frieder, D. Grossman, and D. Lewis, "The Complex Document Image Processing (CDIP) Test Collection Project," Illinois Institute of Technology, 2006.
- [75] D. Lewis, G. Agam, S. Argamon, O. Frieder, D. Grossman, and J. Heard, "Building a test collection for complex document information processing," in *Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval*. ACM, 2006.
- [76] "The Legacy Tobacco Document Library (LTDL)," University of California, San Francisco, 2007.
- [77] F. Bolelli, G. Borghi, and C. Grana, "Historical Handwritten Text Images Word Spotting Through Sliding Window Hog Features," in *19th International Conference on Image Analysis and Processing*, 2017.
- [78] —, "XDOCS: an Application to Index Historical Documents," in *Digital Libraries and Multimedia Archives*, vol. 806, Jan 2018.
- [79] D. Maltoni, D. Maio, A. Jain, and S. Prabhakar, *Handbook of Fingerprint Recognition*. Springer Science & Business Media, 2009.
- [80] J. Sauvola and M. Pietikäinen, "Adaptive document image binarization," *Pattern recognition*, vol. 33, no. 2, 2000.
- [81] D. Baltieri, R. Vezzani, and R. Cucchiara, "3DPeS: 3D People Dataset for Surveillance and Forensics," in *Proceedings of the 2011 joint ACM workshop on Human gesture and behavior understanding*. ACM, 2011.
- [82] D. S. Marcus, A. F. Fotenos, J. G. Csernansky, J. C. Morris, and R. L. Buckner, "Open Access Series of Imaging Studies (OASIS): Longitudinal MRI Data in Nondemented and Demented Older Adults," *J. Cognitive Neurosci.*, vol. 22, no. 12, 2010.
- [83] The Electron Microscopy Dataset. Accessed on 2019-03-21. [Online]. Available: <https://cvlab.epfl.ch/data/data-em/>
- [84] A. Lucchi, Y. Li, and P. Fua, "Learning for Structured Prediction Using Approximate Subgradient Descent with Working Sets," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2013.
- [85] M. Matsumoto and T. Nishimura, "Mersenne twister: a 623-dimensionally equidistributed uniform pseudo-random number generator," *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, vol. 8, no. 1, 1998.
- [86] G. Bradski, "The OpenCV Library," *Dr. Dobb's Journal of Software Tools*, 2000.



Federico Bolelli received the B.Sc. and M.Sc. degrees in Computer Engineering from Università degli Studi di Modena e Reggio Emilia, Italy. He pursued the Ph.D. degree from the same university where he is currently working as a Tenure Track Assistant Professor within the AlmageLab group at Dipartimento di Ingegneria "Enzo Ferrari." His research interests include image processing, algorithms and optimization, and medical imaging. He is currently involved in a H2020 European Project.



Stefano Allegretti received the B.Sc. and M.Sc. degree in Computer Engineering from Università degli Studi di Modena e Reggio Emilia, Italy. He is currently pursuing the Ph.D. degree at the AlmageLab Laboratory at Dipartimento di Ingegneria "Enzo Ferrari" of Università degli Studi di Modena e Reggio Emilia. His research interests include algorithm optimization and image processing.



Luca Lumetti received both B.Sc. and M.Sc. degrees in Computer Engineering from Università degli Studi di Modena e Reggio Emilia, Italy. He is currently pursuing a Ph.D. at the same University, in the AlmageLab group. His research interests regard Artificial Intelligence, Computer Vision, and Medical Imaging.



Costantino Grana graduated at Università degli Studi di Modena e Reggio Emilia, Italy in 2000 and achieved the Ph.D. in Computer Science and Engineering in 2004. He is currently Full Professor at Dipartimento di Ingegneria "Enzo Ferrari" of the same university. His research interests are mainly in medical imaging, optimization of image processing algorithms, and computer vision applications. He published 6 book chapters, 47 papers on international peer reviewed journals and more than 130 papers on

international conferences.