

A Parallel Approach to Object Identification in Large-scale Images

Young-Min Kang
Tongmyong University
Busan, 48520, Korea
ymkang@tu.ac.kr

Sung-Soo Kim
ETRI
Daejeon, 34129, Korea
sungsoo@etri.re.kr

Gyung-Tae Nam
GCSC Inc.
Busan, 47607, Korea
gtnam@gcsc.co.kr

ABSTRACT

As the computing power of processors is being drastically improved, the sizes of image data for various applications are also increasing. One of the most basic operations on image data is to identify objects within the image, and the connected component labeling (CCL) is the most frequently used strategy for this problem. However, CCL cannot be easily implemented in a parallel fashion because the connected pixels can be found basically only by graph traversal. In this paper, we propose a GPU-based efficient algorithm for object identification in large-scale images and the performance of the proposed method is compared with that of the most commonly used method implemented with OpenCV libraries. The method was implemented and tested on computing environments with commodity CPUs and GPUs. The experimental results show that the proposed method outperforms the reference method when the pixel density is below 0.7. Object identification in image data is the fundamental operation and rapid computation is highly requested as the sizes of the currently available image data rapidly increase. The experimental results show the proposed method can be a good solution to the object identification in large-scale image data.

KEYWORDS

GPGPU, Object Identification, CCL

1 INTRODUCTION

As the computing power of processors is being drastically improved, the sizes of image data for various applications are also increasing. Therefore, efficient algorithms for manip-

ulating the large-scale image data are required. One of the most basic operations on image data is to identify objects within the image, and the connected component labeling (CCL) is the most frequently used strategy for this problem.

Various image processing techniques can be easily implemented in parallel fashion, and GPU parallelism has been successfully exploited in this field. However, CCL cannot be easily implemented with parallel tasks because the connected pixels are represented as adjacent nodes in a graph and the adjacency among all the nodes can be investigated basically only by graph traversal.

In this paper, we propose a GPU-based efficient algorithm for object identification in large-scale images and the performance of the proposed method is compared with that of an OpenCV-based CCL algorithm.

2 RELATED WORK

Object identification is a fundamental problem in image processing. In many cases, the object identification is performed based on CCL. The most CCL algorithms are reduced to the traversal of adjacent nodes (pixels) along the edges in the graph that represents the input image. The traversal approaches are naturally sequential, and parallel implementations of typical CCL algorithms have not been very successful [10].

Even in the early stage of computer vision research, it was found that the connectivity cannot be determined by completely parallel tasks [7]. However, the rapid development of general purpose graphics processing unit (GPU) technologies made it possible for GPU-based

parallel approaches to achieve better performance than CPU-based traditional CCL methods [1, 9].

The basic approach to CCL is to use *union-find* algorithm which can determine whether two nodes in an undirected graph are connected or not [8]. However, such methods based on this approach use sequential computations. Several methods have been proposed to exploit parallel computing architectures [5, 9]. However, these methods were applied to relatively small images. Some methods utilized cluster architectures [4]. These methods split the data volume and the data segments are assigned to different computing units. Parallelism within a single GPU, therefore, cannot be exploited in these methods.

Label equivalence method is implemented on GPU [6], and this method iterates to resolve the label equivalence by finding the roots of equivalence trees. However, this approach relies on decision tables which cannot be efficiently handled on GPU [10].

Block-based labeling and efficient block processing with decision table were proposed in [2, 3]. Block equivalence method based on “scan mask” was proposed [10]. However, this method also has to iterate the equivalence resolving until it converges to the state where no label update is found.

3 PROPOSED METHOD

In this section, a GPU-based parallel approach to CCL is proposed. The method is composed of four major tasks: 1) data initialization, 2) computation of column-wise label runs, 3) label merge of connected components. Each task is explained in details in the following subsections.

3.1 Data Initialization

The pixels in an image can be classified into either ‘on’ pixels and ‘off’ pixels. The ‘on’ pixels are regarded nodes in graph representation, and it is assumed that an edge exists between two nodes of which positions neighbor each other in the image space.

The goal of CCL algorithms is to assign an identical label for linked nodes. In order to achieve this goal, each pixel is assigned unique label in the initialization stage. The simplest method is to assign sequential numbers to the pixels. Suppose we have an image with $w \times h$ pixels and the pixel at (x, y) is denoted as $p(x, y)$ where $x = [1, w], y = [1, h]$. The ‘on’ pixel at (x, y) is then labeled with the number $x + w(y - 1)$. Therefore, the labeling numbers range from 1 to wh . All the ‘off’ pixels are labeled to be -1.

The label map I_λ is an image composed of label of each pixel, and each label at (x, y) in the label map are denoted by $I_\lambda^{x,y}$. In other words, the image with initial labels can be described as follows:

$$\begin{aligned} I_\lambda &\in \mathbb{Z}^{w \times h} \\ id^{x,y} = x + w(y - 1) &\in [1, wh] \\ p(x, y) = 1 &\Rightarrow I_\lambda^{x,y} = id^{x,y} \\ p(x, y) = 0 &\Rightarrow I_\lambda^{x,y} = -1 \end{aligned} \quad (1)$$

After the initialization is done, the rest of the algorithm is to merge the positive labels in each connected component into a single label.

3.2 Computing Column-wise Label Runs

In order to merge labels, adjacent positive labels are merged. The block of contiguous object pixels in a column is a ‘run.’ The first stage of label merge is to find runs. In other words, each run is identified and labeled with a unique number. Each column is assigned to

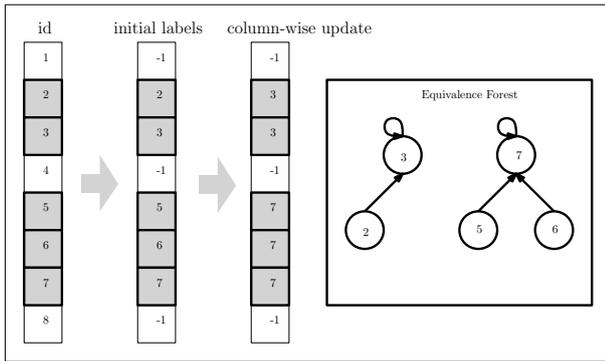


Figure 1. Label runs and equivalence forest

a CUDA thread and processed in parallel fashion. Therefore, we have w threads running separately.

The computation within a thread is to simply scan pixels and change the label of the current pixel to be that of the previously scanned pixel if the both pixels are ‘on.’

Fig. 1 shows how the label assigned to each pixel is updated through column-wise label run computation described in Algorithm 1. After the update, the each label represents the root node in the equivalence tree it belongs to as shown in Fig. 1

Algorithm 1: Column-wise label run

kernel **vertLabel**

Data: $I_\lambda \in \mathbb{Z}^{w \times h}$: In, Out

begin

```

    col = thread:[1, w]
    for row: h - 1 downto 1 do
        if  $I_\lambda^{row,col} > 0$  and  $I_\lambda^{row+1,col} > 0$  then
             $I_\lambda^{row,col} = I_\lambda^{row+1,col}$ 
    
```

3.3 Label Merge

Once the column-wise label update is finished, the vertical connectivity must be investigated. Let us suppose, for simplicity, that we have only two columns. In the previous column-wise update, the labels are merged to the largest value in the equivalence tree. If two pixels in a row are connected, the equivalence

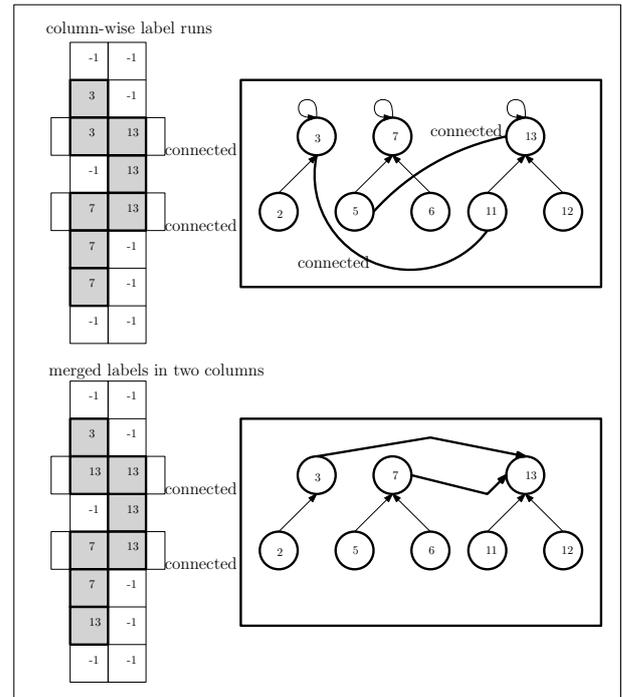


Figure 2. Label merge with two columns

trees those pixels belong to should be merged into a single tree. In order to achieve this, the root node of each tree must be found and the root node with a smaller label is relabeled to point the other root with a larger label as shown in Fig. 2.

Note that the labels of the connected pixels are not directly updated. Fig. 3 shows the merge process. As shown in Fig. 3 (a), two pixels a and b in different equivalent trees are found to be connected. The direct update of connected pixels does not successfully merge the equivalence trees as shown in Fig. 3 (b). The correct label merge can be done by comparing and update of the roots of the equivalence trees as shown in Fig. 3 (c).

In order to perform the two-column label merge for an images with w columns, $w/2$ column pairs are separately merged with $w/2$ threads. After every two adjacent column pairs are merged, we iterate the label merge. In the second merge phase, as shown in Fig. 4, we have only to consider the boundaries between the previously merged column pairs, and the computation is reduced to half the previous one. It is easily noticed that the $\lg w$ iter-

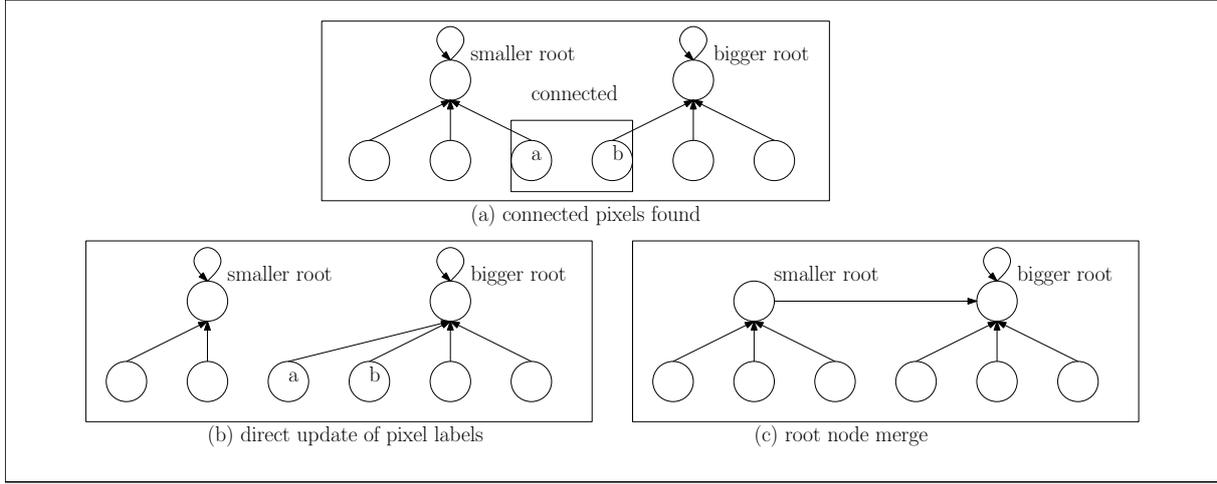


Figure 3. Label merge with two columns

Algorithm 2: Label merge algorithm

host **callMergeLabel**
Data: $I_\lambda \in \mathbb{Z}^{w \times h}$: In, Out
begin
 div = 2
 for i : 0 upto $\lg w - 1$ **do**
 merge $\langle\langle h \cdot w / \text{div} \rangle\rangle$ (div, I_λ)
 div = $2 \times \text{div}$

kernel **merge**
Data: div $\in \mathbb{Z}$: in, $I_\lambda \in \mathbb{Z}^{w \times h}$: In, Out
begin
 thread: [1, wh / div]
 nBoundary = w / div
 col = $\text{div} / 2 + (\text{thread} \% \text{nBoundary}) \times \text{div}$
 row = thread / nBoundary
 if $I_\lambda^{\text{row}, \text{col}} > 0$ and $I_\lambda^{\text{row}, \text{col}+1} > 0$ **then**
 $\text{root}_L = \mathbf{findRoot}(\text{row}, \text{col})$
 $\text{root}_R = \mathbf{findRoot}(\text{row}, \text{col}+1)$
 $I_\lambda^{\min(\text{root}_L, \text{root}_R)} = I_\lambda^{\max(\text{root}_L, \text{root}_R)}$

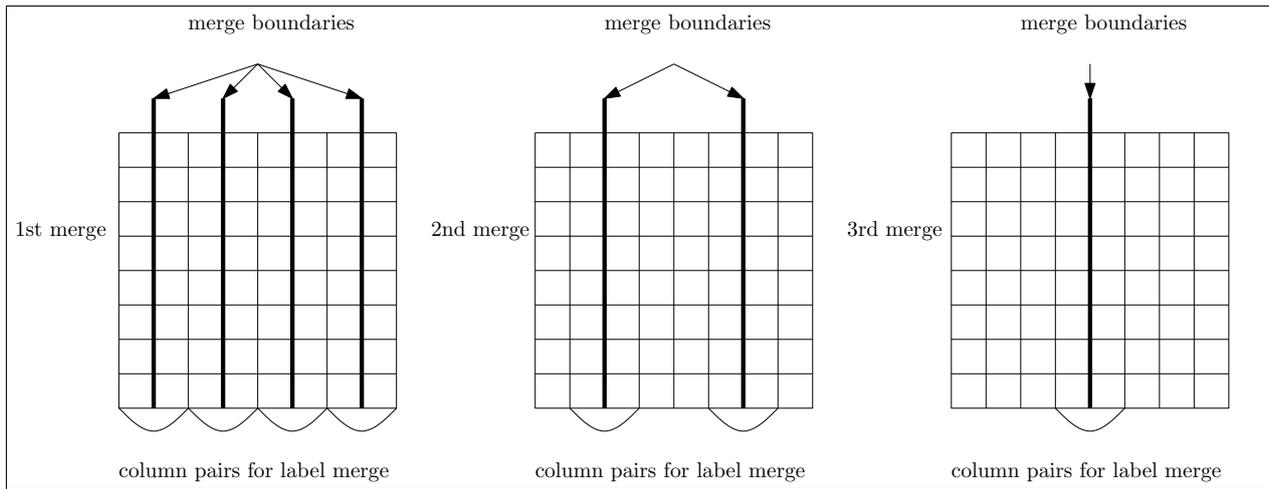
device **findRoot**
Data: row, col $\in \mathbb{Z}$: In, label $\in \mathbb{Z}$: Out
begin
 if $I_\lambda^{\text{row}, \text{col}} < 0$ **then**
 return -1
 label = $w \cdot \text{row} + \text{col}$
 while $I_\lambda^{\text{label}} \neq \text{label}$ **do**
 label := I_λ^{label}
 return label

ations are sufficient for finding the final label equivalence. Let us denote the computational cost for the first iteration to be $C(1)$. The total computational cost for the label merge is $\sum_{i=0}^{\lg w - 1} \frac{1}{2^i} C(1) = \mathbf{O}(C(1))$.

Algorithm 2 shows the implementation details of our method. The label merge is implemented with one host function, one kernel function, and one device function. In the host function callMergeLabel, we determine the number of boundaries where column-pairs are merged and call kernel function merge with the necessary number of threads. The host function iterates this call $\lg w$ times, and the number of threads to be called decreases as the iteration is repeated. In the i -th call, $w \cdot h / 2^i$ threads are required.

Every thread executes the kernel function merge. In the kernel function, every two pixels across the merge boundaries are investigated in parallel fashion. If the pixels are both ‘on’, the root nodes of equivalence trees the pixels belong to are found and compared. The label equivalence trees are merged by relabeling the root with smaller label to have the same label with the other root. The device function findRoot is called in this process to find the root of the pixel currently being investigated.

After the execution of Algorithm 2, the equivalence tree will be obtained. However, the fi-


Figure 4. Label merge iteration

nal goal of CCL is to make all the pixels in a connected component have an identical label. This can be achieved by applying the device function `findRoot` to each pixel and updating its label to be the returned value. This process can be easily performed in a parallel fashion because the label updates for any nodes in the tree do not destroy the equivalence of the nodes in the tree.

Algorithm 3 describes the operations of relabeling thread. Total $w \cdot h$ threads separately call `findRoot` for corresponding pixels and update the label in order to make the connected pixels have an identical label.

Algorithm 3: Relabeling

 kernel **relabel**
Data: $I_\lambda \in \mathbb{Z}^{w \times h}$: In, Out

begin

```

    thread:[1, w · h]
    row = thread/w, col = thread%w
     $I_\lambda^{row, col} = \text{findRoot}(row, col)$ 
    
```

4 EXPERIMENTAL RESULTS

The method proposed in this paper was implemented on computing environments with commodity CPUs and GPUs. The experimental results were collected from the tests on a system with i7-3630QM 2.4 GHz CPU and Geforce GTX 670MX GPU.

Table 1. CCL performance comparison with random density noise patterns (2048×2048 pixels).

2048×2048		
density	Grana (ms)	Proposed method (ms)
0.1	22.8	6.3
0.2	30.7	7.0
0.3	46.0	7.7
0.4	51.3	8.6
0.5	47.6	10.2
0.6	43.6	15.2
0.7	35.5	26.1
0.8	28.0	40.8
0.9	20.8	64.6

In order to verify the efficiency of our method, we compared the performance of our method with the most commonly used method. The reference method was proposed in [3], and implemented with OpenCV libraries.

In the first experiment, the performance of each method was measured by applying images with random noise. The random noise was automatically generated and the density of the noise ranges from 0.1 to 0.9. Table. 1 shows the experimental results when 2048×2048 images with random noise are applied. The first column represents the noise density, and the second column shows the measured execution time of the reference method in milliseconds. The third column shows the execution time of the proposed method. As shown in the table, the reference method (denoted by Grana) requires more execution time when the

Table 2. CCL performance comparison with random density noise patterns (4096×4096 pixels).

4096×4096		
density	Grana (<i>ms</i>)	Proposed method (<i>ms</i>)
0.1	85.7	21.4
0.2	135.3	24.3
0.3	190.4	27.1
0.4	206.1	30.8
0.5	205.4	36.9
0.6	177.6	59.0
0.7	153.1	126.7
0.8	112.1	216.7
0.9	85.6	411.5

noise density is around 0.5 while the proposed method requires more time as the density increases.

Table. 2 shows the similar experimental results except that the size of the input images is 4096×4096. As shown in the table, the computational cost is similarly changing in accordance with the density.

Table 3. CCL performance comparison with test images A and B with different image sizes.

methods	Images sizes			
	512 ²	1024 ²	2048 ²	4096 ²
Test Image A				
Grana (<i>ms</i>)	1.14	3.86	14.39	41.1
Proposed (<i>ms</i>)	0.87	2.46	7.68	26.0
Gain (%)	23.7	36.3	46.6	36.7
Test Image B				
Grana (<i>ms</i>)	1.01	3.53	12.35	40.78
Proposed (<i>ms</i>)	0.83	2.35	7.21	24.03
Gain (%)	17.8	33.4	41.6	41.1

Fig. 5 (a) visualizes the experimental results shown in Table. 1. As shown in the figure, the computational cost of the proposed method increases as the density increases. However, the proposed method is far better until the density is below 0.7.

Fig. 5 (b) shows the similar results. This figure visualizes the result shown in Table. 2. As shown in the figure, the computational cost of the proposed method increases as the density increases. However, the proposed method is far better until the density is below 0.7.

The CCL algorithms are not actually applied to noise data. In order to measure the performance of the proposed method in more feasible environments, we prepared two test images shown in Fig. 6. There are two different test images and the sizes of the images can be either 512², 1024², 2048² or 4096². Test image *A* has two components, and each of them is a long spiral curve without touching the other component. The other test image *B* has many scattered stars, and two of them are connected with a star-shaped thin line.

Table. 3 shows the execution time required for reference method and the proposed method applied to the test images with different sizes. In the last row in each data set obtained by using each test image, the performance gain is computed and shown. The performance gain is obtained by computing the ratio of the reduced computational cost by applying the proposed method to the cost of the reference method. As shown in the table, the performance gain is more noticeable as the size of the input image increases.

Fig. 7 visually compares the computational costs of the reference method and our method (lines), and the performance gain is also visualized with bars. Fig. 7 (a) shows the result when the test image *A* was used as input image. The performance gain was the largest when the size of the input image is 2048×2048.

Fig. 7 (b) shows the similar results when the other test image is used. As shown in the figure, the performance gain is again the most noticeable when the size of the image is 2048 × 2048.

5 CONCLUSION

In this paper, an efficient GPGPU implementation of connected component labeling (CCL) was proposed. The method exploits the data parallelism of GPUs to improve the performance of CCL. Object identification in image data is the fundamental operation and rapid

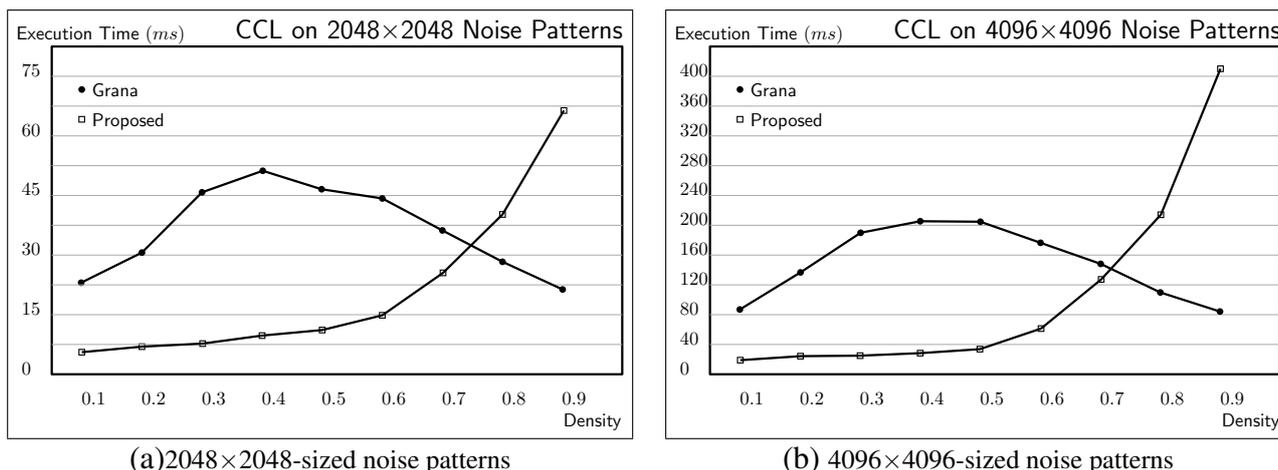


Figure 5. CCL execution time on noise patterns with different noise densities

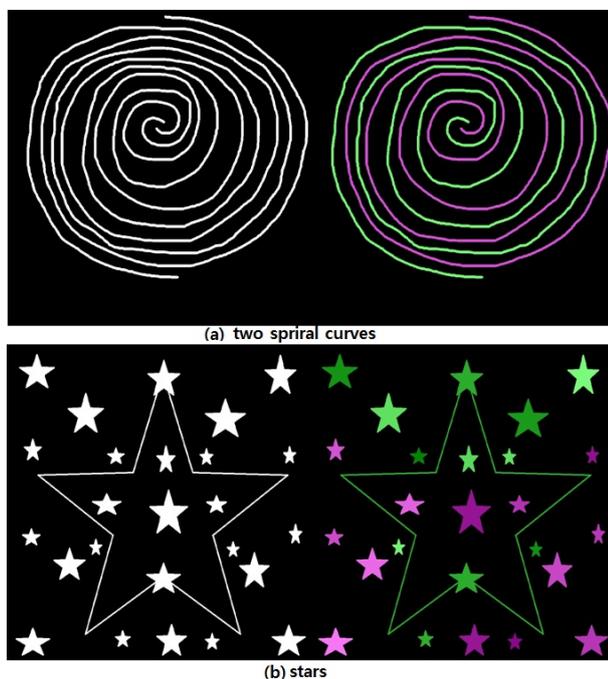


Figure 6. Test images for practical labeling: (a) two spiral curves (b) scattered stars

computation is highly requested as the sizes of the currently available image data rapidly increase. The experimental results show the proposed method can be a good solution to the object identification in large-scale image data.

ACKNOWLEDGMENT

This work was supported in part by ETRI R&D Program (Development of Big Data Platform for Dual Mode Batch-Query Analytics,

16ZS1410), and also in part by NIPA SW Convergence Technologies Enhancement Program (Development of Big Data Processing and Decision Support System for Offshore Maritime Safety, S0142-15-1014).

REFERENCES

- [1] P. Chen, H. Zhao, C. Tao, and H. Sang. Block-run-based connected component labelling algorithm for gpgpu using shared memory. *Electronics letters*, 47(24):1309–1311, 2011.
- [2] C. Grana, D. Borghesani, and R. Cucchiara. Connected component labeling techniques on modern architectures. In *International Conference on Image Analysis and Processing*, pages 816–824. Springer, 2009.
- [3] C. Grana, D. Borghesani, and R. Cucchiara. Optimized block-based connected components labeling with decision trees. *IEEE Transactions on Image Processing*, 19(6):1596–1609, 2010.
- [4] C. Harrison, H. Childs, and K. P. Gaither. Data-parallel mesh connected components labeling and analysis. In *Eurographics Parallel Graphics and Visualization Symposium, Llandudno, Wales*, 2012.
- [5] K. A. Hawick, A. Leist, and D. P. Playne. Parallel graph component labelling with gpus and cuda. *Parallel Computing*, 36(12):655–678, 2010.
- [6] O. Kalentev, A. Rai, S. Kemnitz, and R. Schneider. Connected component labeling on a 2d grid using cuda. *Journal of Parallel and Distributed Computing*, 71(4):615–620, 2011.

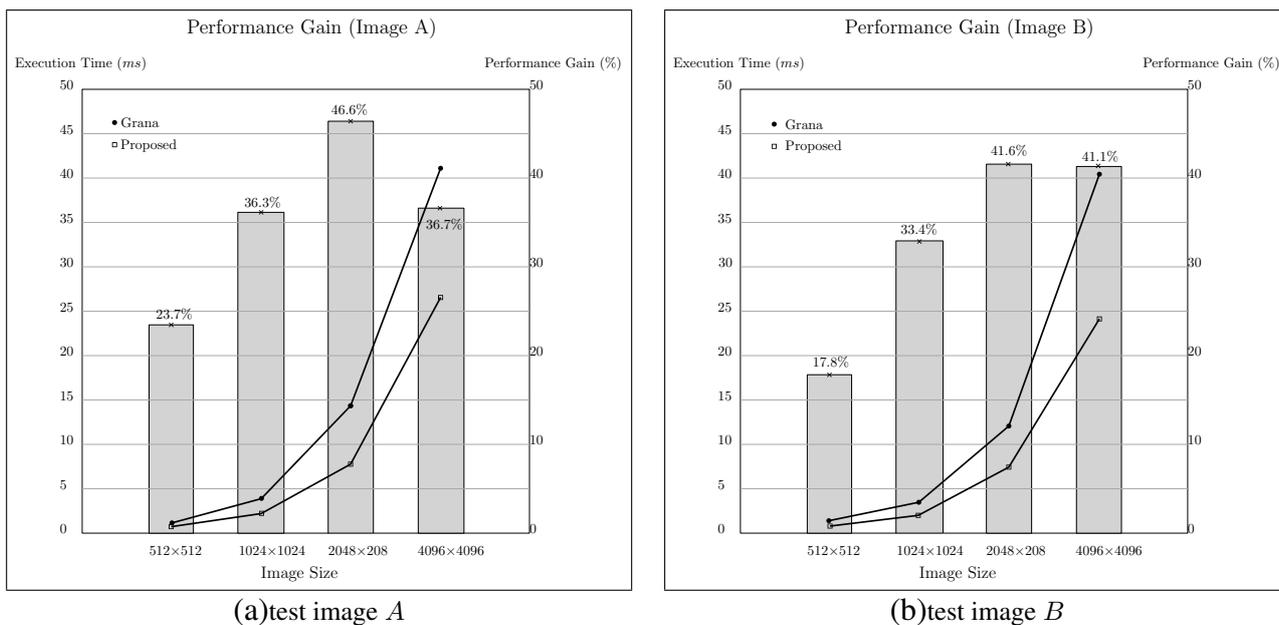


Figure 7. CCL execution time for test images and measured performance gains

- [7] M. Minsky and S. Papert. *Perceptrons*. MIT press, 1988.
- [8] B. Preto, F. Birra, A. Lopes, and P. Medeiros. Object identification in binary tomographic images using gpgpus. *International Journal of Creative Interfaces and Computer Graphics (IJ-CICG)*, 4(2):40–56, 2013.
- [9] O. Štáva and B. Beneš. Connected component labeling in cuda. *Hwu., WW (Ed.), GPU Computing Gems*, 2010.
- [10] S. Zavalishin, I. Safonov, Y. Bekhtin, and I. Kurilin. Block equivalence algorithm for labeling 2d and 3d images on gpu. *Electronic Imaging*, 2016(2):1–7, 2016.