



How does Connected Components Labeling with Decision Trees perform on GPUs?



Federico Bolelli

Università degli Studi di Modena e Reggio Emilia, DIEF, Italy

Connected Components Labeling (CCL)

- Extracts connected components (objects) from a binary image.
- All pixels of the same object are given the same label.





Optimizing CCL Algorithms

- CCL has an exact solution.
- Computational complexity is proven to be O(N).
- Optimization can be achieved in different ways:
 - 1. Designing algorithms to reduce the average number of memory accesses and to take advantage of instruction cache;
 - 2. Using hardware acceleration.





8-connectivity



CPU-Based Algorithms

- Traditionally, on sequential machines a two-scan algorithm is employed:
 - **1.** *First scan,* scans the input assigning provisional labels to the output and storing equivalences in the Union-Find data structure.
 - 2. Flattening, establishes the definitive labels.
 - 3. Second scan, updates the output replacing provisional with final labels.
- The second scan can be avoided when only statistics about the objects are required.



Scan Mask





Scan Array-Based Union-Find Algorithm (SAUF)

It is based on the Rosenfeld scan-mask:



- Exploits a decision tree to avoid unnecessary read/write operations.
- Implements the Union-Find using arrays.



Block-Based with Decision Trees (BBDT)

• Introduce the concept of 2x2 block mask.





Tree to Directed Rooted Acyclic Graph (DRAG)

- Compresses the memory footprint of the source code.
- Does not impact on memory accesses.
- Increases the instruction cache hit rate.



Decision Trees on GPU

- A GPU uses the SIMT paradigm grouping threads into packets (warp).
- Branch divergence occurs when threads inside warps branch to different execution path.
- SIMT paradigm can group threads into warps flexibly *i.e.* can group threads supposed to truly execute the same instruction.



Adapting Tree-Based Algorithms to GPUs ...

First scan translates into two different kernels

Initialization

1	2	0	4	5	6	7	8	0	10	11
12	13	14	0	16	17	18	0	20	21	22
23	24	0	0	0	0	0	0	0	32	33
34	0	0	37	0	0	0	41	0	0	44
0	0	0	0	0	0	0	0	0	0	0
0	57	0	0	0	0	0	0	0	65	0
0	68	0	70	71	72	73	74	0	76	0
78	79	80	0	0	83	0	0	86	87	88

$$L[id_x] = id_x$$

Merge



DecisionTree



Adapting Tree-Based Algorithms to GPUs ...



 $L[id_{\chi}] = Find(L, id_{\chi})$

• AImage^{Lab}

55

Block Labels into Pixels

Comparative Evaluation

- We compared our proposals to state-of-the-art GPU algorithms.
- YACCLAB is employed:
 - Open source C++ benchmarking system \rightarrow github/prittt/YACCLAB;
 - Collection of real cases datasets;
 - Tests over diverse points of view;
- Our device is a Quadro K2200
 - CUDA capability 5.0;



Average Execution Time



- Optimized Label Equivalence (OLE):
 - propagates minimum label
- Block Equivalence (BE):
 - same as OLE, but works on 2x2 blocks
 - needs more data structures
- Union Find (UF)
 - uses the *union-find* data structure
- Distanceless Label Propagation (DPL)
 - merges the UF and OLE strategies
- Komura Equivalence (KE):
 - changes the initialization kernel of the UF









- Random images of 2048x2048 size with increasing density
- Tree-based solutions have the best performance when density is below ~18%
- Most real cases images are there!



Conclusion Remarks

- How does Connected Components Labeling with Decision Trees perform on GPUs?
 - Better than state-of-the-art approaches.
- "It is a capital mistake to theorize before one has data"

SHERLOCK HOLMES







How does Connected Components Labeling with Decision Trees perform on GPUs?

Thank you!

federico.bolelli@unimore.it