

International Conference on Computational Science, ICCS 2013

Regularity versus Load-Balancing on GPU for treefix computations

David Defour^a, Manuel Marin^a

^aUniv. Perpignan Via Domitia, DALI F-66860, Perpignan, France
Univ. Montpellier II, LIRMM, UMR 5506, F-34095, Montpellier, France
CNRS, LIRMM, UMR 5506, F-34095, Montpellier, France

Abstract

The use of GPUs has enabled us to achieve substantial acceleration in highly regular data parallel applications. The trend is now to look at irregular applications, as it requires advanced load balancing technics. However, it is well known that the use of regular computation is preferable and more suitable when working with these architectures. An alternative to the use of load balancing is to rely on *scan* and other GPU friendly parallel primitives to build the desired result; however implying in return, the involvement of extra memory storage and computation.

This article discusses of both solutions for treefix operations, which consist of applying a certain operation while performing a tree traversal. They can be performed by traversing the tree from top to bottom or from bottom to top, applying the proper operation at each vertex. It can be accelerated using either load balancing which maintains a pool of tasks while performing only the necessary amount of computation or using a vector friendly representation that will involve twice the amount of computation than the first solution. We will explore these two approaches and compare them in terms of performance and accuracy. We will show that the vectorial approach is always faster for any category of trees, but it raises accuracy issues when working with floating-point data.

Keywords: GPU computing ; regular versus irregular algorithms ; numerical quality ;

1. Introduction

In recent years, processors such as IBM cell SPUs, FPGAs, GPUs, and ASICs were successfully considered to provide speedup on numerous classes of applications. Of these, GPUs stand out as they are produced as commodity processors and exhibiting a number of processing cores doubling every year, revealing the current architectural trend. GPUs were used to improve the performance of regular computations such as those described in [1]. On such highly regular computations, GPUs can outperform a single core CPU by a large factor on average, that could be higher than 400 in some cases [2]. These large speedups are only possible for highly regular and computationally intensive classes of application. More recently, irregular computations on graphs such as list ranking [3] and connected components [4] were also considered. However, in these cases, the observed speedup compared to single core performance is of the order of 5 or less.

Treefix operations were first introduced by Leiserson and Maggs [5] as intermediate steps in a number of higher-level graph analysis algorithms. They defined two basic operations, *Rootfix* and *Leafix*. *Rootfix* returns to

*Corresponding author. Tel.: +33-4-30-19-23-06 ; fax: +33-4-68-66-22-87.
E-mail address: david.defour@univ-perp.fr.

each vertex of the tree the result of applying a certain operation over all its ancestors; Leaffix returns to each vertex the result of applying an operation over all its descendants. Rootfix and Leaffix have application for example in the Backward-forward sweep algorithm for electrical network analysis [6] or to evaluate the parsimony score of phylogenetic trees [7, 8]. In this article, we explore the available alternatives to accelerate these computations using GPUs.

The usual implementation of Rootfix and Leaffix is based on traversing the tree, from top to bottom or from bottom to top. The vertices are updated as visited, allowing to effectively propagate the accumulated result of the operation through the whole tree as the traversal progresses. The order of visit is relevant. Starting from the root, Depth-first or Breadth-first traversals are both valid alternatives. Ultimately, Rootfix and Leaffix can be viewed as performing a complete Breadth-first or Depth-first search over a tree, updating the vertices' weights as they are visited.

Successful implementations of parallel Breadth-first search over a general graph on GPU can be found in [9, 10, 11, 12, 13]. All of them rely on *level-synchronization*, i.e. processing every level of the graph in parallel, in order of depth. This is often implemented as an iterative process that performs one iteration per level. Some versions [11, 12, 13] examine every vertex of the graph at every iteration: if the predecessor was visited during the last iteration, then the vertex is visited. These methods perform a quadratic amount of work, as the graph can have, in the worst case, as many levels as vertices. A work efficient versions [9, 10] focus on producing, at each iteration, a vertex or edge *frontier*, including only those elements to be visited or traversed during that iteration. The main advantage of these methods is to exhibit a work efficient scheme, but have to deal with the irregularity of the graph data structure, which involves load imbalance and potential underutilization of SIMD lanes. Different load balancing strategies are applied to improve the performance achieved by these methods.

An alternative for performing Rootfix and Leaffix on a GPU, is to use a parallel-friendly representation of the tree consisting of three arrays based on the Euler-tour ordering. A series of highly regular parallel operations performed over these arrays, such as *scan*, allow to compute the result of Rootfix and Leaffix for a tree with n vertices in $O(\lg n)$ parallel steps, independently of the tree topology. However this methods relies on array of size $2.n$ with two times more computations than load balancing implementations.

The purpose of this article is to determine the best solution between a work efficient scheme thanks to irregular computation or a solution with regular computation with double the amount of operation to solve the treefix problem on GPUs. It makes the following contributions in the area of parallel computing:

- **Regular vs irregular algorithm comparison.** We present two different approaches that make use of data-parallelism to perform a distinctive operation over trees. One of them leads to an application that is highly regular, the other to one that is highly irregular and compares them in terms of performance.
- **Numerical quality analysis.** We compare the numerical accuracy of both methods when dealing with floating-point data as the amount and the order of operation is different.
- **Rootfix and Leaffix OpenCL implementation.** We provide a vectorial implementation of +Rootfix and +Leaffix in OpenCL. Even if there has been some work on implementing Rootfix and Leaffix in different languages [14, 15, 16], this is, to our knowledge, the first parallel implementation that could run on a GPU.

2. Presentation of Rootfix and Leaffix

Leiserson and Maggs [5] formally defined Rootfix and Leaffix as follows: given a weighted tree and a binary operator \oplus , Rootfix assigns to each vertex the result of applying \oplus to all of the vertex's ancestors; Leaffix assigns to each vertex the result of applying \oplus to all of the vertex's descendants.

From there, we can define the +*Rootfix* and +*Leaffix* operations, where \oplus is the addition, as assigning to each vertex the sum of its ancestors and the sum of its descendants, respectively. Figure 1 shows an example. In particular, if all the vertices of the tree have weight 1, +Rootfix returns the depth of each vertex, and +Leaffix returns the size of the sub-tree rooted on every vertex.

2.1. Parallel algorithm

Regarding the type of trees considered, there are two easy cases of parallelization: balanced binary tree and linked list. For the balanced binary tree, Leiserson and Maggs [5] proposed a randomized algorithm that performs

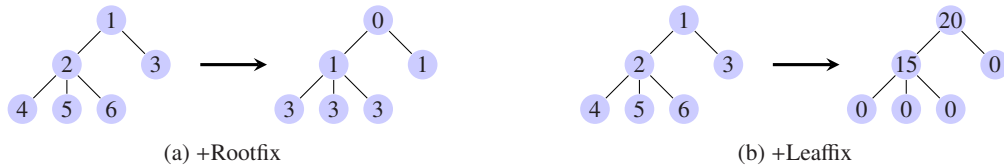


Fig. 1: Example of +Rootfix and +Leaffix.

Rootfix and Leaffix on a tree of size n in $O(\lg n)$ parallel steps, applying the contraction technique provided by Miller and Reif [17]. For the linked list, or caterpillar, there exists a $O(\lg n)$ depth algorithm based on symmetry breaking.

In this article, we consider traversing the tree using parallel Breadth-first search. The tree is expressed as a directed graph of the form $G = (V, E)$, with a set V of n vertices and a set E of $n - 1$ directed edges¹. The adjacency matrix A is defined as follows.

$$A_{ij} = \begin{cases} 1 & \text{if } (v_i, v_j) \in E \\ 0 & \text{otherwise} \end{cases}$$

We rely on compressed sparse row (CSR) format to store this matrix into two arrays. The array C contains the column indices of the non-zero elements of A arranged in row-major order. The array R contains $n + 1$ integers, and entry $R[i]$ is the index in C of the i -th row of A .

Algorithm 1 illustrates the usual way of performing +Rootfix using parallel Breadth-first search based on level-synchronization. The algorithm manipulates two queues: one input queue and one output queue. The input queue contains all the vertices to be examined during certain iteration. All these vertices are dequeued in parallel and their children are updated. As updated, the children are placed in the output queue. When all the children have been visited at a given level, the output queue is transferred in the input queue to be consumed by the next iteration. The algorithm proceeds until there are no vertices left to examine. An analog algorithm can be formulated for +Leaffix.

Algorithm 1 +Rootfix parallel algorithm

Input: Row-offsets array R , column-indices array C , weights array W . Function *LockedEnqueue*(*vertex*) safely inserts *vertex* at the end of the queue instance.

Output: Array *rootfix*[0... $n - 1$] holding the result.

```

1: rootfix[0] ←  $W[0]$ 
2: inQ ← {}
3: inQ.LockedEnqueue(0)
4: while inQ != {} do
5:   outQ ← {}
6:   for  $i$  in inQ do in parallel
7:     for offset in  $R[i] \dots R[i + 1] - 1$  do in parallel
8:        $j \leftarrow C[\textit{offset}]$ 
9:       rootfix[ $j$ ] = rootfix[ $i$ ] +  $W[i]$ 
10:      outQ.LockedEnqueue( $j$ )
11:  inQ ← outQ

```

The amount of parallel work that this algorithm can perform depends on the tree topology. The wider the level, the greater the number of parallel tasks than can be assigned for that level. This is related to the *average branching factor*, i.e. the average number of children per vertex. The worst-case scenario is when every vertex has only one child (caterpillar) and then all the vertices have to be examined sequentially.

2.2. Vectorial algorithm

The implementation of Rootfix and Leaffix for the PRAM machine model was studied by Blelloch [18], who provided a vectorial algorithm. The algorithm uses Euler-tour order, a technique first introduced by Tarjan and Vishkin [19], to compute a vector representation of the tree. The Euler-tour order is generated by replacing every

¹always directed from parent to child

edge in the tree by two directed edges, one in each sense; these edges define an Eulerian path around the tree. As they appear on this path, the edges are placed into an Euler-tour vector E . A downward edge reaching vertex v is labeled ${}_v$ and an upward edge leaving vertex v is labeled v . Figure 2 shows an example tree and the corresponding Euler-tour vector. Note that, as we doubled the number of edges, the Euler-tour vector has twice the size of the tree. The Euler-tour vector for a tree of n vertices can be found in parallel in $O(\lg^2 n)$ steps[20].

The vector tree representation consists of three arrays, ${}_V$, V_j and W . The array ${}_V$ holds, for each vertex v , the index of ${}_v$ in the Euler-tour vector E ; the array V_j , the index of v_j . The array W holds the vertices' weights.




Fig. 2: Example tree, Euler-tour ordering and vector tree representation.

These three arrays are used altogether with some regular parallel primitives to compute the result of Rootfix and Leaffix in a parallel fashion. The key primitive is the *scan* operation, that given a binary operator \oplus with identity i , takes the array

$$(x_0, x_1, \dots, x_{n-1})$$

and returns the array

$$(i, x_0, x_0 \oplus x_1 \dots, x_0 \oplus x_1 \oplus \dots \oplus x_{n-2})$$

For \oplus being the addition, the *+scan* operation takes the same input array and returns

$$(0, x_0, x_0 + x_1 \dots, x_0 + x_1 + \dots + x_{n-2})$$

There exist many GPUs implementation of this operation as it is a basic building block of many data parallel algorithms. The one in [21] operates in $O(\lg n)$ steps and $O(n)$ operations. This has been further optimized for the NVIDIA Fermi architecture in [22].

Algorithm 2 takes as input an array E of size $2n$, which is used for intermediate computation, and the three arrays ${}_V$, V_j and W of size n that hold the tree. It produces the result of *+Rootfix*. A similar algorithm is available for *+Leaffix*.

Algorithm 2 *+Rootfix* vectorial algorithm

Input: Array E of size $2n$, arrays ${}_V$, V_j and W of size n holding the tree.

Output: Array R of size n holding the result.

- 1: //Step 1: Write
 - 2: for i in $0 \dots n$ do in parallel
 - 3: $E[{}_V[i]] \leftarrow W[i]$
 - 4: $E[V_j[i]] \leftarrow -W[i]$
 - 5: //Step 2: Scan
 - 6: Run an inplace *+scan* on E
 - 7: //Step 3: Read
 - 8: for i in $0 \dots n$ do in parallel
 - 9: $R[i] \leftarrow E[{}_V[i]]$
-

Figures 3 illustrates this algorithm on an example tree. We used the sum as operation applied on integer data. It can be noticed that we could have used any set of values and with any binary operation that forms a *group*. The operation has to be associative, with an inverse and an identity value. As floating-point addition is not associative, these algorithms should not be applied in such cases. However, we will show that in this particular case the error can be bounded.




Fig. 3: +Rootfix vectorial algorithm.

3. GPU implementation

3.1. Parallel version

In section 2.1 we showed that +Rootfix and +Leaffix can be performed using parallel Breadth-first search over a tree. As Breadth-first search is a common building block for many graph analysis algorithms, there exist several GPU implementations. We used the one by Merrill *et al.* [9], written in CUDA. This version optimizes the neighbor gathering process, which corresponds to the for-loop in line 7 of algorithm 1, to balance load within the CTA. For each vertex being expanded, the row-range bounds are read from the array *R* (values $R[i]$ and $R[i + 1]$). Then, each thread uses the result of a CTA-wide parallel prefix sum over the differences $R[i + 1] - R[i]$, to perfectly pack into a buffer, which is shared by the entire CTA, the positions on the array *C* of the neighbors to be gathered (values $R[i] \dots R[i + 1] - 1$). Once the buffer has been filled, each thread in the CTA reads one position on it and gathers the corresponding neighbor from *C*, leaving no SIMD lane idle during the process. According to [9], this load balancing strategy allows to achieve a traversal rate about 5 times greater than with other parallel implementations on GPU. We did not modify the code to make it more suitable to our purposes, more details are available in [9].

3.2. Vectorial version

We have seen in section 2.2 that +Rootfix and +Leaffix can be implemented on a PRAM machine using the vector tree representation and the +scan operation. However, there was no GPU implementation available. To perform the test, we developed an OpenCL implementation of +Rootfix and +Leaffix, as this allows us to be platform independent.

The implementation for both operations follows the algorithms by Blelloch and it is built around 3 separate kernels, operating on 3 vectors of size *n* that represent the input tree ((V, V, W)). Once data allocation and data transfer are done, a first kernel *Write* is launched with *n* work items packed in workgroup sizes that maximize performance. Our test has shown that this corresponds to the maximum allowed for the selected device, which can be queried via `clGetKernelWorkGroupInfo()`. This first kernel is in charge of reading data from input vectors and placing them accordingly in the Euler-tour vector *E* located in global memory. Then the *Scan* kernel is launched to perform a prefix sum on *E*. And finally the third kernel *Read* reads the results from *E* and compute the results for each node. The execution configuration of this third kernel is identical to the first kernel.

All tree kernels are bandwidth limited. Let *idx* represent the global index of a given OpenCL work item. The Write kernel involves 3 coalesced reads ($(V[idx], V[idx]$ and $W[idx])$ and 2 uncoalesced writes in the Euler-tour vector *E* ($E[V]$ and $E[V]$) for both +Rootfix and +Leaffix. The Read kernel involves 1 coalesced read ($(V[idx])$, 1 uncoalesced read ($E[V]$) and 1 coalesced write ($R[idx]$) for both +Rootfix and +Leaffix, plus 2 coalesced reads ($V[idx]$ and $W[idx]$) and 1 uncoalesced read ($E[V]$) only for +Leaffix. Although it is possible to design an

Table 1: Suite of benchmark trees

Name	Nb. of vertices	Depth	Avg. branching factor
af_shell9	504855	490	1030.32
audikw1	943695	236	3998.71
ldoor	952203	784	1214.54
af_shell10	1508065	1098	1373.47
G3_circuit	1585478	705	2248.91
kkt_power	2063494	36	57319.28
nlpkkt120	3542400	123	28800.00
cage15	5154859	81	63640.23
nlpkkt160	8345600	163	51200.00
nlpkkt200	16240000	203	80000.00

efficient memory access pattern for the Scan kernel, it was not possible to avoid those 'uncoalesced' memory accesses for the Write and Read kernels as the scheme is highly dependent on the tree topology. This has been confirmed by the Nvidia profiler. However, we noticed that GPU with L1 and L2 cache like Fermi were benefiting of relaxed memory access pattern improving memory bandwidth. On a Fermi architecture, when performing +Rootfix on a tree of 10^7 vertices, the global memory load efficiency of the Read kernel is about 61.5 %, whereas on a pre-Fermi architecture it is about 30 %. For the Write kernel, the difference is of 42.9 % versus 22.9 %.

4. Tests and results

In this section we present the tests we carried out to measure the related performance and accuracy of different +Rootfix and +Leaffix implementations. The results are discussed in light of the different features presented in the tested implementations. Due to space limitations, we are only reporting performance results related to +Rootfix, +Leaffix implementation leading to similar conclusion.

4.1. Performance

When using Breadth-first search for performing +Rootfix over a tree, as the algorithm is completely data-driven, one can expect that the tree topology will have an impact on the performance. Moreover, if a parallel implementation is used, some types of tree will allow more parallelism than others. This is related to the average branching factor, i.e. the ratio between the number of vertices and the number of levels. The larger is this parameter, the wider the tree and thus the greater the number of parallel tasks that can be performed. On the other hand, if we use a vectorial algorithm, the impact of the tree topology over performance should be negligible. To validate this hypothesis and test the proposed implementations, we used a group of benchmarks from the University of Florida Sparse Matrix Collection [23]. This collection, maintained by Tim Davis and Yifan Hu, includes several matrices from different real-life problems on different fields. We selected ten matrices that were considered by the 10th DIMACS Implementation Challenge [24], as they are representative of the type of networks that can be found in reality in terms of size and topology. For each one of these matrices, we computed a spanning tree of the associated directed graph and used that tree as benchmark. Table 1 shows the details of the benchmarks generated, including the tree depth and the average branching factor.

For each algorithm running on each benchmark, we measured the total execution time and decoupled it into (a) data transfer time, and (b) computation time. This is motivated by the fact that these algorithms are usually included in iterative scheme where they are called alternatively until a condition is reached. In these cases, data transfer is operated only once. We compared the parallel and vectorial +Rootfix implementations to a purely sequential +Rootfix implementation running on CPU. The machine used for the tests is an Intel Xeon E645 CPU with an NVIDIA GeForce GTX670 1344 cores GPU. We used GCC 4.6.3, Cuda 4.2.1 and OpenCL 1.1.

Figure 4 shows the related performance of sequential, parallel and vectorial +Rootfix. The benchmarks are ordered from left to right by increasing number of vertices. We observe in figure 4a that the computation time for the vectorial implementation always grows with the tree size, while for the sequential and parallel implementation there are some cases where certain trees are processed in less time than other with fewer vertices. For example, the parallel implementation needs 21 milliseconds to compute the result for the *nlpkkt120* benchmark, which has

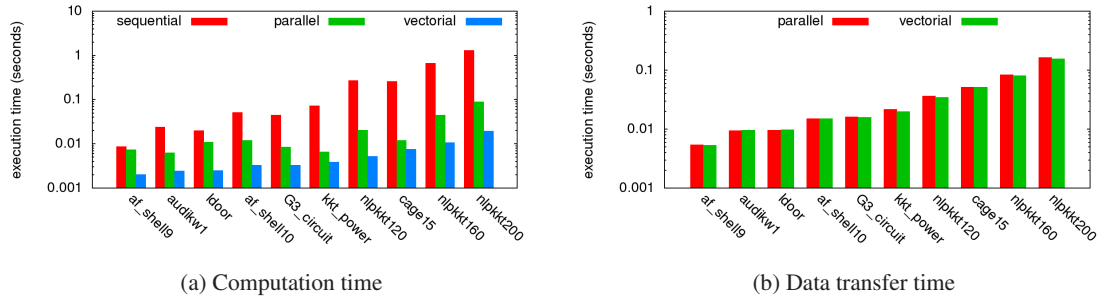


Fig. 4: Related performance of sequential, parallel and vectorial +Rootfix on the GTX670 GPU.

3.54 million vertices, and only 12 milliseconds to compute the result for the *cage15* benchmark, which has 1.5 million vertices.

The data transfer time is almost the same for both the parallel and vectorial implementations, as we can observe in figure 4b. This is consistent with the fact that the same amount of data is transferred. For a tree with n vertices, the parallel implementation transfers from host to device the CSR representation, consisting of two arrays of size respectively $n - 1$ and $n + 1$. The vectorial implementation transfers the vector tree representation, consisting of two arrays of size n . Then, both implementations transfer from the device to the host the result in the form of one array of size n .




Fig. 5: Speedup of parallel and vectorial +Rootfix over sequential +Rootfix on the GTX670 GPU.

Figure 5 shows the speedup of parallel and vectorial +Rootfix over sequential +Rootfix. We can see that, when considering only computation time, the speedup achieved by both implementations is quite substantial; it reaches more than 60x on the largest benchmarks analyzed with the vectorial implementation.




Fig. 6: Comparison of vertex distribution in the *nlpkkt120* and *cage15* benchmarks.

To measure the impact of the tree topology, we looked at the vertex distribution of pairs of benchmarks, like it is plotted in figure 6. The 5.15 million vertices of the *cage15* benchmark are concentrated in fewer levels than the 3.54 million of the *nlpkkt120* benchmark. As a consequence of this, the *nlpkkt120* benchmark takes longer to

process, even if it is smaller than the *cage15*. This explains the difference quoted in figure 4.

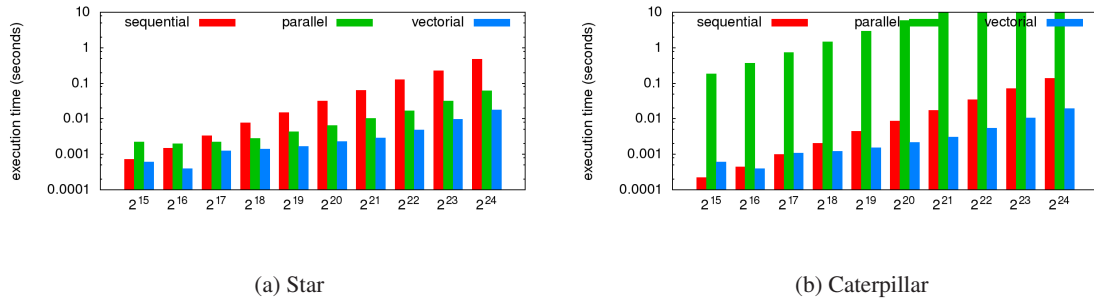


Fig. 7: Related performance of sequential, parallel and vectorial +Rootfix for star and caterpillar trees on the GTX670 GPU.

To quantify the effect of the average branching factor on the two version of the +Rootfix algorithms, we considered two extreme cases of topology: (a) the star, where the root has $n - 1$ children, and (b) the linked list, or caterpillar, where every vertex has exactly one child. In the star, the average branching factor is equal to the size of the tree; in the caterpillar, it is equal to one.

We generated a new set of benchmarks composed of stars and caterpillars of sizes varying from 2^{15} to 2^{24} vertices. Figure 7 shows the computation time of sequential, parallel and vectorial implementation of +Rootfix on these special topologies. We observe that the parallel implementation performs poorly on the caterpillar, as this algorithm finally needs to process all the vertices sequentially on the GPU. This causes a slowdown compared to the sequential implementation, as the load balancing tasks remains while bringing no benefits. For the star, all the terminal nodes are concentrated on one single level, which correspond to the perfect case for the parallel version. We can notice that, surprisingly, the vectorial implementation is faster by a factor 5 compared to the parallel implementation for the star with 2^{24} nodes. As the branching factor is decreasing, the performance of the parallel version is quickly decreasing leading to a computation time 5000 times greater than the vectorial implementation.

4.2. Accuracy

When +Rootfix and +Leaffix operate on integer both the parallel and vectorial implementations return the same result as long as no overflow occurs during intermediate computation. However, with floating-point arithmetics, rounding errors may occur for every operation. This is the case with floating-point addition that is not associative. Therefore, we could expect a variation in the result between the parallel and vectorial versions of +Rootfix and +Leaffix. For every vertex v , the +Rootfix parallel algorithm performs only as many operations as the vertex has ancestors. The +Rootfix vectorial algorithm performs as many operations as the number of elements in the Euler-tour vector before the v_j position. When using floating-point arithmetics, we can expect the +Rootfix vectorial algorithm to be less accurate than the +Rootfix parallel algorithm. Figure 8 illustrates the difference in the number of operations for both +Rootfix parallel and vectorial algorithms.

To measure the numerical quality of these algorithms, we use the *relative error*, which is a measure of how far is the observed result from the real result. If x is the real result and \hat{x} the observed result, the relative error e is given by $e = \frac{|\hat{x} - x|}{|x|}$. Given a problem and an input data, this measure is linked with the algorithm that produces the result and thus can be used to compare algorithms. The measure of the difficulty of a problem independently of the algorithm used to solve it is given by the *condition number*. The condition number is a measure of how much the result of a problem is changed by small variations in the operands. If we consider the addition of n floating-point numbers x_0, \dots, x_{n-1} , the condition number C is defined by $C = \left(\sum_{i=0}^{n-1} |x_i| \right) / \left| \sum_{i=0}^{n-1} x_i \right|$. As a rule thumb, we may lose up to $\lg(C)$ bits of accuracy.

As the result of +Rootfix and +Leaffix is a set of n values, we can use different metrics to quantify the error. We could look at each error individually, the mean error over the n results or the maximum error. In addition, the topology of the tree is impacting the computation scheme and therefore the error. For example, if we consider




Fig. 8: Different number of operations when performing Rootfix with different algorithms.

a linked list (caterpillar), then both +Rootfix and +Leaffix parallel implementations will require a recursive sum of n values with n partial sums. Whereas if we consider a tree with the root and $n - 1$ children (star) then each partial sum generated by +Rootfix will be the result of only one addition.

We choose to evaluate the numerical behavior of both the parallel and vectorial versions of +Rootfix and +Leaffix over a sum of n numbers, which corresponds to a chain of n vertices in a tree. For this set of n numbers we generated 100 random trees of 10.000 nodes with condition numbers from 10 to 10^{10} . We used the algorithm proposed by Ogita *et al.* [25] to generate series of floating-point numbers with a given condition number. We measured the relative error of the parallel and vectorial versions of +Rootfix and +Leaffix on every node using double-precision to compute the real result and single-precision to compute the observed result. With this measure we captured the numerical behavior of both algorithms on one sum among the n sums that constitute the result. By construction, this is representative of the numerical behavior in function of the condition number of the problem.




Fig. 9: Related accuracy of parallel and vectorial +Rootfix and +Leaffix algorithms.

Figure 9 shows the maximum relative error as a function of the condition number for the +Rootfix and +Leaffix parallel and vectorial algorithms. We observe that both parallel and vectorial versions of +Rootfix have similar numerical behavior. The large dispersion of points for condition number less than 10^4 may come from the difficulties we had generating vectors with such characteristics. On the other hand, the parallel version of +Leaffix seems better than the vectorial one. It seems that in this case the vectorial version is losing an extra 2 bits of accuracy compared to the parallel version.

5. Conclusion

In this paper, we have presented two different methods to solve the treefix problem on GPU and compared them. A parallel implementation, that minimizes the number of operations and intermediate storage thanks to load balancing technics and a vector friendly method that involves twice the amount of memory usage and operation than the previous one but exhibit regular computation pattern. We have shown that in terms of performance, regularity is always a better choice over reducing the amount of operations and memory usage. In addition, we have observed that depending on the tree topology, the vectorial implementation is insensitive to it which lead to speed-up factor ranging from 5 to 5000 compared to the load-balancing implementation.

When dealing with floating-point input data, we have seen that the vectorial implementation is introducing rounding error in the final result compared to the parallel implementation. These errors are the consequence of the extra operations and reordering of computations of the vectorial method, which may leads to a 2-bit lost in the worst case. Nevertheless, this accuracy impact has to be formally bounded according to the tree topology, which is planned as future work.

References

- [1] S. Sengupta, M. Harris, Y. Zhang, J. D. Owens, Scan primitives for gpu computing, in: Proceedings of the 22nd ACM SIGGRAPH/EUROGRAPHICS symposium on Graphics hardware, GH '07, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 2007, pp. 97–106.
URL <http://dl.acm.org/citation.cfm?id=1280094.1280110>
- [2] S. Collange, M. Daumas, D. Defour, Graphic processors to speed-up simulations for the design of high performance solar receptors, in: Application-specific Systems, Architectures and Processors, 2007. ASAP. IEEE International Conf. on, IEEE, 2007, pp. 377–382.
- [3] Z. Wei, J. JaJa, Optimization of linked list prefix computations on multithreaded gpus using cuda, in: Parallel Distributed Processing (IPDPS), 2010 IEEE International Symposium on, 2010, pp. 1–8. doi:10.1109/IPDPS.2010.5470455.
- [4] K. Hawick, A. Leist, D. Playne, Parallel graph component labelling with gpus and cuda, *Parallel Computing* 36 (12) (2010) 655 – 678. doi:10.1016/j.parco.2010.07.002.
URL <http://www.sciencedirect.com/science/article/pii/S0167819110001055>
- [5] C. Leiserson, B. M. Maggs, Communication-efficient parallel algorithms for distributed random-access machines, *Algorithmica* 3 (1988) 53–77.
- [6] D. Shirmohammadi, H. Hong, A. Semlyen, G. Luo, A compensation-based power flow method for weakly meshed distribution and transmission networks, *Power Systems, IEEE Transactions on* 3 (2) (1988) 753 –762. doi:10.1109/59.192932.
- [7] W. M. Fitch, Toward defining the course of evolution: Minimum change for a specific tree topology, *Syst Biol* 20 (1971) 406–416.
- [8] D. Sankoff, Minimal mutation trees of sequences, *SIAM Journal on Applied Mathematics* 28 (35–42).
- [9] D. Merrill, M. Garland, A. Grimshaw, Scalable gpu graph traversal, *SIGPLAN Not.* 47 (8) (2012) 117–128. doi:10.1145/2370036.2145832.
URL <http://doi.acm.org/10.1145/2370036.2145832>
- [10] L. Luo, M. Wong, W.-m. Hwu, An effective gpu implementation of breadth-first search, in: Proceedings of the 47th Design Automation Conference, DAC '10, ACM, New York, NY, USA, 2010, pp. 52–55. doi:10.1145/1837274.1837289.
URL <http://doi.acm.org/10.1145/1837274.1837289>
- [11] M. Hussein, A. Varshney, L. S. Davis, On implementing graph cuts on cuda, *First Workshop on General Purpose Processing on Graphics Processing Units*.
- [12] P. Harish, P. J. Narayanan, Accelerating large graph algorithms on the gpu using cuda, in: Proceedings of the 14th international conference on High performance computing, HiPC'07, Springer-Verlag, Berlin, Heidelberg, 2007, pp. 197–208.
URL <http://dl.acm.org/citation.cfm?id=1782174.1782200>
- [13] Y. S. Deng, B. D. Wang, S. Mu, Taming irregular eda applications on gpus, in: Proceedings of the 2009 International Conference on Computer-Aided Design, ICCAD '09, ACM, New York, NY, USA, 2009, pp. 539–546. doi:10.1145/1687399.1687501.
URL <http://doi.acm.org/10.1145/1687399.1687501>
- [14] M. M. T. Chakravarty, R. Leshchinskiy, S. Peyton Jones, G. Keller, S. Marlow, Data parallel haskell: a status report, in: Proceedings of the 2007 workshop on Declarative aspects of multicore programming, DAMP '07, ACM, New York, NY, USA, 2007, pp. 10–18. doi:10.1145/1248648.1248652.
URL <http://doi.acm.org/10.1145/1248648.1248652>
- [15] Scandal project home page, <http://www.cs.cmu.edu/scandal/> (2012).
- [16] The manticore project, <http://manticore.cs.uchicago.edu/> (2012).
- [17] G. L. Miller, J. H. Reif, Parallel tree contraction and its application, in: 26th Symposium on Foundations of Computer Science, IEEE, Portland, Oregon, 1985, pp. 478–489.
- [18] G. E. Blelloch, Prefix sums and their applications, Tech. rep., *Synthesis of Parallel Algorithms* (1990).
- [19] R. E. Tarjan, U. Vishkin, Finding biconnected components and computing tree functions in logarithmic parallel time, in: Proceedings of the 25th Annual Symposium on Foundations of Computer Science, 1984, SFCS '84, IEEE Computer Society, Washington, DC, USA, 1984, pp. 12–20. doi:10.1109/SFCS.1984.715896.
URL <http://dx.doi.org/10.1109/SFCS.1984.715896>
- [20] M. Atallah, U. Vishkin, Finding euler tours in parallel, *J. Comput. Syst. Sci.* 29 (3) (1984) 330–337. doi:10.1016/0022-0000(84)90003-5.
URL [http://dx.doi.org/10.1016/0022-0000\(84\)90003-5](http://dx.doi.org/10.1016/0022-0000(84)90003-5)
- [21] M. Harris, S. Sengupta, J. D. Owens, Parallel prefix sum (scan) with CUDA, in: H. Nguyen (Ed.), *GPU Gems 3*, Addison Wesley, 2007, Ch. 39, pp. 851–876.
- [22] M. Harris, M. Garland, *GPU Computing Gems Jade Edition*, 1st Edition, no. 3, MKP, 2011, Ch. Optimizing Parallel Prefix Operations for the Fermi Architecture.
- [23] The university of florida sparse matrix collection, <http://www.cise.ufl.edu/research/sparse/matrices/> (2012).
- [24] 10th dimacs implementation challenge, <http://www.cc.gatech.edu/dimacs10/index.shtml> (2012).
- [25] T. Ogita, S. M. Rump, S. Oishi, Accurate sum and dot product, *SIAM J. Sci. Comput* 26 (2005) 2005.