

Energy-efficiency of software and hardware algorithms*

Maja H. Kirkeby¹, Thomas Krabben¹, Maria B. Mikkelsen¹, Mads Rosendahl¹, Mathias Larsen², Martin Sundman², Tjark Petersen³, and Martin Schoeberl³

¹ Department of People and Technology, Roskilde University,
Roskilde, Denmark

{majaht,krabben,mariabm,madsr}@ruc.dk

² IT University of Copenhagen, Copenhagen, Denmark
{mathl,sund}@itu.dk

³ DTU Compute, Technical University of Denmark,
Lyngby, Denmark
s186083@student.dtu.dk
masca@dtu.dk

Abstract. In this article, we compare the energy efficiency of hardware and software implementations of Heapsort and Dijkstra’s algorithm for route finding. The software implementations are written in C for Raspberry Pi, and the hardware implementations are crafted in Chisel for an FPGA. Our objective is to examine how we can fairly compare energy efficiency between hardware and software. These solutions are positioned to replace each other in operational contexts, necessitating a comparison of their whole-system energy consumption. This study seeks to identify circumstances where time and energy efficiency diverge, offering insights to guide hardware selection. Our findings serve as a step towards understanding the complex trade-offs in algorithm performance across different computational platforms.

Keywords: energy efficiency, performance, FPGA, CPU, algorithms

1. Introduction

Improving software’s time and energy efficiency is essential [1, 6]. Implementing the algorithms in Field-Programmable Gate Arrays (FPGAs) demonstrates substantial performance gains in highly parallelizable tasks, e.g., [18, 30]. However, only a few have considered the energy efficiency of hardware implementations [14, 17]. Previous studies on improving algorithms using FPGAs have employed FPGAs in different ways. One approach which requires specialized expertise is to implement the algorithm in hardware-specifying languages such as Verilog or VHDL [28, 30]. Alternative approaches, oriented toward software developers, use FPGAs as accelerators for, e.g., C/C++ programs [14] or, as in this study, where use the high-level programming language Chisel [3, 27].

* This work is supported by the Innovation Fund Denmark for the project DIREC (9142-00001B), by CERCIRAS Cost Action CA19135, by the Independent Research Fund Denmark Project no. 2102-00281B, and is based on work from the COST Action CERCIRAS CA19135, supported by COST (European Cooperation in Science and Technology).

We will compare two alternative solutions, i.e., a hardware and a software solution, in operational contexts. We will evaluate the energy efficiency and performance of two computing platforms (running the algorithmic implementations) that may replace each other in operational contexts. Our work extends an initial studies presented and published at the 2022 Workshop on Resource Awareness of Systems and Society [18], which is reported in Section 7. This work indicated an exciting upper bound for performance improvement when comparing a highly parallelizable program, Conway’s Game of Life [4], to software implementations. Although an upper bound is interesting, it does not provide any insight into the gain from converting ordinary software implementations into hardware implementations for ordinary programs. In this paper, we will focus on providing a more well-rounded comparison, studying algorithms where neither software nor hardware has particular advantages:

How do time and energy consumption compare for hardware and software implementations of ordinary software algorithms?

In our choice of algorithms, it is our aim that it should not provide any particular advantage for either hardware or software versions. It is possible to find highly parallelizable algorithms where the hardware implementation will have an obvious advantage. In our experiments, we use Heapsort and Dijkstra’s algorithm to find the shortest path in a graph. Both are widely used and have well-known and studied implementations. Both algorithms are not easy to parallelize [10, 13], but can still contain a fair amount of parallelism at the local level.

The main contributions of this article are:

1. *Comparative Analysis:* The paper compares the energy efficiency between software and hardware implementations of two well-known algorithms: Heapsort and Dijkstra’s algorithm. This comparison is new in that it focuses on both energy consumption and performance across software and hardware implementations and, thus, provides the first insights into the energy and performance trade-offs
2. *Methodological Innovation:* It introduces a methodological framework that ensures fair comparisons between software and hardware implementations. This includes a thorough discussion of the choice of measurement techniques.

In the following, we provide the basis for a fair comparison (Section 2), and in Section 3, we introduce the specifications and implementations of the algorithms. Section 4 describes our experimental setup, and Section 5 describes our results. Section 6 discusses the results. Section 7 compares implementations of the Game-of-Life to explore an upper bound of possible speedups and energy savings in an FPGA. Afterward, we discuss the related work (Section 8) and contextualize our results. Section 9 concludes the paper and summarizes future work.

2. A Fair Comparison

In this section, we discuss a methodological framework to ensure that comparisons between software and hardware implementations are fair and insightful. By detailing the choice of hardware and measurement protocols, we highlight the practical considerations

underlying the comparison, providing a foundation for the detailed descriptions of the specific implementations that follow.

2.1. Choice of Measurement

Previous work on energy consumption of software implementations has employed energy estimations using Intel's Running Average Power Limit (RAPL) [11], as it has been reported as having negligible overhead and providing precise results [9, 24]. e.g., [23]. However, while RAPL is precise and highly correlated (a value of 0.99) [16]) with the actual power dissipation, it does not provide accurate results, i.e., the RAPL measurements are not close to the true energy consumption. Thus, instead, we could employ external measurements that provide each systems's ground truth energy consumption. This choice also has drawbacks. For instance, it introduces more noise since it measures the energy consumption of the entire device compared to only the CPU. It also introduces an imprecision in the synchronization between the time in the measuring unit obtaining the energy consumption and the time in the measured device that executes the algorithm.

In this study, the two computing platforms running the algorithmic implementations are positioned to replace each other in operational contexts. Therefore, we employ external energy measurements to capture the energy consumption of the entire systems. This approach allows for a fair comparison across platforms by reflecting the energy impact of platform-specific overheads, such as memory and I/O subsystems, which are integral to real-world operation. While external measurements introduce noise and potential synchronization imprecision, they provide a more accurate and comprehensive evaluation of whole-system energy consumption.

2.2. Choice of Hardware

One factor that influences the energy consumption of both hardware and software implementations is the choice of hardware. Since the hardware is very different in the two cases, we will use the same requirement for choosing the hardware, namely, to use cheap and available hardware.

There are many cheap and available computers for software implementations, e.g., Orange Pi, Asus Tinker, Nvidia Jetson Nano, or Raspberry Pi. However, in this study, we have chosen a standard Raspberry Pi 4 computer Model B with 4GB RAM and a 1.5 GHz 64-bit quad-core ARM Cortex-A72 processor running the standard Raspberry Pi OS, a Linux version.

There are two types of FPGA units: pure FPGA boards and system-on-chip FPGA boards. Choosing the pure FPGA board will allow us to measure the exact energy consumption of the FPGA unit alone without interference from other processors, which would be the case with System on Chip FPGA boards.

For the hardware implementations, we have chosen a Digilent FPGA board Cmod A7 (version Cmod A7-35T) with an XC7A35T-1CPG236C FPGA unit, an MSPS On-chip ADC, 20800 Look-up Tables (LUTs), 41600 Flip-Flops, 225 KB Block RAM and 5 Clock Management Tiles.

2.3. Choice of Implementations

In addition, previous studies, e.g., [6, 8, 23], highlight that whole-systems energy consumption provides insights into the varying energy consumption across different implementations, illustrating how choices in software development impact overall energy use. Thus, for instance the choice of programming language [23], implementation style [8], and compiler flags [26] influence the energy consumption and execution time of the programs. For the hardware implementation, energy factors typically include the number of logic elements and routing resources, see, e.g., [29].

For the software solutions, we base it on standard C-implementations, e.g., following the descriptions in [5], using standard settings of optimizations in the GNU GCC compiler. The Hardware solutions are hand-written versions in Chisel with common approaches to optimize the design for improved performance. How we have realized the algorithms in hardware is discussed in the next section.

To ensure fairness, we align the implementations of the algorithms across both platforms. This alignment ensures that the differences in energy consumption and execution time reflect the platforms' inherent characteristics rather than variations in algorithmic design. The implementations will:

1. carry out Dijkstra's algorithm and Heapsort, respectively.
2. avoid external communication by including the algorithm inputs in the implementation instead of reading the input from external files.
3. return the smallest subset of the result to ensure computation of the results while reducing the read/write accesses.

This standardized approach supports meaningful comparisons while adhering to the operational requirements of each platform.

3. Algorithms

This section provides detailed descriptions of Heapsort and Dijkstra's implementations. With a focus on implementation specifics, the next Section 4 allows us to finalize the experimental design before providing the results in Section 5.

3.1. Heapsort: Software

The implementation of Heapsort uses a standard implementation from the Rosetta repository⁴, see Figure 1. It uses a max-heap data structure, storing values in a balanced tree where a node is bigger than its children. The tree is represented as an array. In a binary tree, the children at array index i can be found at array index $i * 2 + 1$ and $i * 2 + 2$. The implementation uses a k -heap with slightly better complexity measures since the tree will have a smaller depth for the same number of store values.

The k -heap structure is initially established by moving values down the tree structure if they are smaller than their children. In the second phase, the biggest value is repeatedly moved to the back of the array, and the heap structure is reestablished.

⁴ <http://www.rosettacode.org/>

```

// heapsort start
int max (int *a, int n, int parent) {
    int largest = parent;
    for (int child = (K*parent)+1; child < (K*parent)+K+1; child++)
        if (child < n && a[child] > a[largest]) largest = child;
    return largest;
}
void downheap (int *a, int n, int i) {
    while (1) {
        int j = max(a, n, i);
        if (j == i) break;
        int t = a[i];
        a[i] = a[j];
        a[j] = t;
        i = j;
    }
}
void heapsort (int *a, int n) {
    int i;
    for (i = (n - 2) / K; i >= 0; i--) downheap(a, n, i);
    for (i = 0; i < n; i++) {
        int t = a[n - i - 1];
        a[n - i - 1] = a[0];
        a[0] = t;
        downheap(a, n - i - 1, 0);
    }
} // heapsort end

```

Fig. 1. The software implementation of Heapsort in C, where a is the input array to be sorted and n is its size

3.2. Heapsort: Hardware

The hardware solution of Heapsort uses higher order k -max-heaps to increase parallelism. In the k -heap, $k+1$ elements must be compared in each step while re-establishing the heap order. In hardware, this comparison can be done in parallel, thus theoretically allowing for the heap order to be established in $\log_k(n)$ clock cycles. Practically, fetching all required values from memory, finding the largest of them, and swapping the parent and the largest child if the heap order is violated has to be spread over multiple clock cycles to allow the circuit to be operated at high clock frequencies. An architectural diagram of the heap

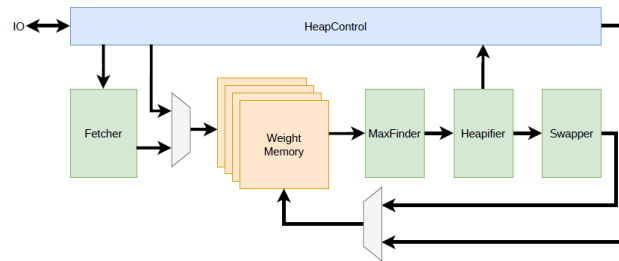


Fig. 2. The architecture of the hardware implementation for Heapsort

module is shown in Figure 2. The circuit can not easily be pipelined since deciding which child should be the next parent while traversing the heap downwards is always delayed, resulting in a pipelined implementation having to stall most of the time.

Nonetheless, parallelism can be exploited in multiple instances, giving the hardware implementation an edge over a software implementation in clock cycles per iteration. All k children of a given node can be fetched simultaneously using wide memories, which store k concatenated values per address. Thus, only two load operations are needed per iteration. The maximum between the parent and all k children nodes can be found using a tree of blocks, where the maximum of two values is selected. The total delay of this circuit is $\log_2(k)$ times the delay of a single block. For larger values of k , the tree has to be divided into multiple levels separated by registers to allow for operation at high clock frequencies. The write operations to the memory associated with a swap can be overlapped with fetching the next parent or children, depending on the direction of the traversal.

The described heap module for sorting is paired with a small circuit that inserts values from memory into the heap and then extracts the ordered sequence from the heap by continuously removing the root until the heap is empty.

3.3. Dijkstra: Software

The implementation of Dijkstra's algorithm represents the graph as an adjacency list, where all edges from a vertex are grouped together. The algorithm computes the shortest route from a given vertex (here, the vertex at index 0) to all other vertices. For each vertex, it will find the previous vertex in the shortest route from the start to that vertex and the weight of that route. The actual route can be found by following the previous vertex indices through the graph until the start vertex is reached. The implementation can be seen in Figure 3.

3.4. Dijkstra: Hardware

A priority queue-based system allows for an easy determination of the next node to visit. Using the hardware heap of the heap sort experiment, a hardware priority queue could be constructed that accepts a new value worst case every $\log_k(n)$ clock cycles, where n is the number of nodes in the graph and k is the degree of parallelism in the heap. This results in $n(n-1)d \log_k(n)$ clock cycles to execute Dijkstra's algorithm, where d is the density of the graph with $d = 1$ representing a fully connected graph. This solution is not easily parallelizable since it would require a priority queue that can insert multiple values simultaneously.

An alternative implementation uses a linear search through all nodes to determine the next node to visit. This proves to be an advantage since the route updates and search for the next node to visit can be conducted in parallel in hardware. The length of an alternative path through the currently visited node is calculated for each node. If the alternative path is shorter, it is written to the route table instead of the previously known distance from the start. At the same time, a state element holding the closest unvisited node yet to be encountered is updated if this node is closer to the start. After all nodes have been visited, the state element holds the next node to visit. This results in $n(n-1)$ clock cycles to execute Dijkstra's algorithm. This solution can easily be parallelized by working

```

// the graph is stored as an adjacency list where edges are
// ordered by start vertex input to Dijkstra's algorithm
int n;
#define n 6 // number of vertices
#define m 9 // number of edges
int node1[]={ 0, 0, 0, 1, 1, 2, 2, 3, 4}; // edges' start vertex,
int node2[]={1, 2, 5, 2, 3, 3, 5, 4, 5}; // edges' end vertex
int dist[]={7, 9, 14, 10, 15, 11, 2, 6, 9}; // edges' weight
int edge[]={0, 3, 5, 7, 8, 0}; // vertex to edge index link
// Data structures for the algorithm
int done[n], prev[n], wght[n];
const int MAX=1000000;
void main(){
    int start = 0;
    // Dijkstra start
    for(int i=0;i<n;i++){ done[i]=0; prev[i]=-1; wght[i]=MAX;}
    wght[start]=0;
    int cur=-1, w = 0;
    for(int k=0;k<n;k++){
        cur = -1; w = MAX;
        for (int i = 0; i < n; i++)
            if (done[i]==0 && wght[i] < w) { cur = i; w = wght[i];}
        if (cur<0) break;
        int j = edge[cur];
        while (j < m) {
            int n1 = node1[j], n2 = node2[j], d = dist[j];
            if (n1 != cur) break;
            int d2 = wght[n2], d3 = w + d;
            if (d2 > d3) { prev[n2] = cur; wght[n2] = d3;}
            j++;
        }
        done[cur] = 1;
    }
    cur = n-1;
    // Dijkstra end
}

```

Fig. 3. Dijkstra's shortest path algorithm; it finds the path from vertex 0




Fig. 4. The architecture of the hardware implementation of Dijkstra's algorithm

on multiple path updates simultaneously. This requires multiple read ports on the weight memory and multiple read and write ports to the route table. Furthermore, not only one route has to be compared to the closest unvisited not yet to be encountered but multiple at the same time.

The second solution is chosen over the first since it is easily parallelizable while requiring significantly fewer hardware resources. This translates into lower power dissipation and only worse performance on graphs of density below $1/\log_k(n)$. Considering the total energy of the execution of Dijkstra's algorithm, the density at which both solutions perform equally well is offset even more in favor of the second solution since its lower power consumption compensates for the longer run time.

An architectural diagram of the hardware implementation of Dijkstra's algorithm is shown in Figure 4. A group of k nodes reads their preliminary shortest path, their visited status from the route memory, and their weight to the currently visited node from the weight memory. For each node in the group, the distance of a new route through the currently visited node is calculated and compared to the old shortest path. The shorter of the two routes is selected and written to the route memory. Furthermore, the state element holding the closest unvisited node to the start is updated if the distances sent to the route memory are shorter. In the hardware implementation, the design is pipelined with the two memory modules separating the circuit into two stages.

To support single-cycle access to the weights without using $O(n^2)$ memory, a buffering solution that exploits the known direction of traversal of the weights is employed. This solution stores weights as packed, unaligned adjacency lists, thus only requiring $O(e)$ memory, where e is the number of directed edges in the graph.

4. Experimental Setup

The hardware implementations were executed on a Digilent FPGA board Cmod A7 (version Cmod A7-35T), and the software implementations were compiled with gcc (Raspbian 8.3.0-6+rpi1) 8.3.0 with flag `-O2` and executed on a Raspberry Pi 4 computer Model B 4GB RAM. There is one program per input array, i.e., one file per software and hardware implementation and input array.

4.1. Input spaces

In Heapsort, the runtime and, therefore, perhaps also the energy consumption depend on the number of input elements and their order. A previous study [17] showed that while the order matters for runtime, the comparability arises from using the same order in all experiments. In this study, we (1) always use the same order, namely ordered input, which causes the highest number of comparisons, and (2) vary the number of elements: 4096, 6144, 8192, 10240, 12288, 14336, and 16384.

For Dijkstra's shortest path algorithm, the execution time depends on the type of graph, i.e., sparse or dense and directed or undirected, and the start node for which the path is calculated. In the sorting algorithm, the focus was on growth, and for the shortest path algorithm, we varied only the form and the starting node. The implementation will calculate routes switching between all the possible start vertices.

4.2. Pre-study: Configuration of Heapsort

An exploratory study [17] found that the optimal degree of parallelism in hardware programs differs when considering energy and time. The kind of parallelism in hardware algorithms and software algorithms differs, as described in Section 3, and, thus, the optimal degree of parallelism differs for software and hardware implementations and differs for each algorithm. A fair comparison allows the optimal degree of parallelism, and therefore, we have run a pre-study to find the optimal degree of parallelism. For both hardware and software implementations, we have optimized for energy. Both Dijkstra and Heapsort can have different degrees of parallelism in the hardware implementation, while for the software implementation, this is only true for Heapsort.




Fig. 5. The current over time for Heapsort using different $k \in \{2, 4, 8, 16, 32, 64\}$ on FPGA inputs of size 4096. The k -values are not tested in order; the k -values are annotated above the associated execution. In the graph, we also see error-prone executions without annotations. The $k = 16$ provides the energy optimal execution

Hardware The test-run graph obtained directly from the Siglent can be seen in Figure 5⁵. It shows the current over time (in minutes) for Heapsort using different k -values, i.e., $k \in \{2, 4, 8, 16, 32, 64\}$ on FPGA inputs of size 4096.

The k values are not tested in order and are annotated above the associated execution. In the graph, we also see error-prone executions without annotations. The $k = 16$ provides the energy-optimal execution since it has the least area under the curve, i.e., the shortest execution time and the lowest current (the voltage is fixed).

Software The growth rate is almost linear, as expected from an algorithm that runs at $n \log_k n$, see Figure 6. Interestingly, the constant factor does matter since the fastest time is with $k = 4$. For bigger k , finding the subtree with the biggest root note becomes the main bottleneck in the execution.

⁵ The original experimental measure data was lost, and we cannot provide precise energy consumption.




Fig. 6. The average execution time in seconds for Heapsort using a k -heap on Raspberry Pi over inputs from 4096 to 16384. The $k = 4$ provides the fastest executions

4.3. Measuring and Adjustment of Repetitions

We measure the energy consumption for the entire Raspberry Pi using a programmable power supply Siglent SPD3303X-E Linear DC 3CH in connection with a Siglent SDM30-45X Digital Multimeter is a 4 1/2 digit (66,000 count) multimeter. This setup allows for the required high-precision readings of the current. The equipment spans the current of both the FPGA and the Raspberry Pi and allows us to measure the power dissipation each 100ms.

```

#define iter 500000
int main(){
    int maxl=iter/n, lng=0;
    for(int n1=0;n1<n;n1++){
        for(int k1=0;k1<maxl;k1++){
            int start=n1;
            // Dijkstra start ... end
            while(cur!=start){
                lng+=wght[cur];
                cur = prev[cur];
            }
            printf("done %d\n",lng);
        }
    }
    return 0;
}

```

(a) Dijkstra

```

#define iter = 3000
// Heapsort start ... end
int main () {
    int i, j;
    for (i = 0; i < iter; i++) {
        for (j = 0; j < N; j++){
            a[j] = j;
            heapsort(a, N);
            printf("done %d\n",a[0]);
        }
    }
    return 0;
}

```

(b) Heapsort

Fig. 7. Dijkstra's shortest path software algorithm (Figure 7a) will calculate 500.000 routes switching between the n possible start vertices and the Heapsort software algorithm will iterate its sorting 3000 times

In this experiment, we measure power continuously and the total execution time a single run for each input size. We want to calculate an average power dissipation and average execution time for one execution of the algorithm for each input size. Because

the algorithm executes very fast, we cannot measure a single execution. To ensure a large enough sample size of the current, we adjust the total execution time to be at least 3 seconds by adjusting the number of iterations. The adjustments occur within both the software and hardware implementation, see Table 1 for a precise number of iterations used in the hardware and software implementations. The adjustment setup within the software implementations differ, see Figure 7b for adjustments used for Heapsort and Figure 7a for the adjustments used for Dijkstra. This adjustment will also reduce the synchronization imprecision between the equipment and the hardware.

Table 1. Overview of the number of in-algorithm iterations that ensure a least execution time of 3 seconds

Algorithm	Mode	input size(s)	Iterations
Heapsort	Software	all	3000
	Hardware	4096	1550
		6144	1033
		8192	775
		10240	620
		12288	516
		14336	442
		16384	387
Dijkstra	Software	all	500000
	Hardware	all	9000

From the measured power, we calculate an average power dissipation, and using the number of iterations, see Table 1, and the measured total execution time, we calculate the average execution time for a single iteration.

5. Results

Our results demonstrate significant and contrasting differences in energy efficiency and performance between Heapsort and Dijkstra's software and hardware implementations. These findings may necessitate more nuanced insights into the optimal hardware selection based on the algorithmic demands.

5.1. Heapsort

Comparative data on Heapsort's time and energy consumption in an FPGA and the Raspberry Pi is shown in Figure 8 and Table 2, with a focus on the energy efficiency achieved through hardware implementation. These results demonstrate energy savings and extra time costs when employing FPGAs over traditional CPUs. It is worth noticing that while the power dissipation by software and hardware implementations are different, they seem constant for the individual implementation.

Table 2. Heapsort's energy consumption within FPGA (k=16) and Raspberry Pi (k=4)

Input size	Raspberry Pi			FPGA		
	Time pr. iter. (ms)	Power (W)	Energy pr. iter. (mJ)	Time pr. iter. (ms)	Power (W)	Energy pr. iter. (mJ)
4096	1.541	3.043	4.689	2.102	0.431	0.906
6144	2.383	3.050	7.269	3.614	0.571	2.062
8192	3.296	3.050	10.055	4.959	0.572	2.837
10240	4.259	3.032	12.913	6.513	0.573	3.731
12288	5.166	3.047	15.741	7.922	0.572	4.533
14336	6.150	3.064	18.847	9.944	0.569	5.656
16384	7.066	3.080	21.764	11.213	0.378	4.241

The energy consumption of the Heapsort implementation on the Raspberry Pi is highly correlated to the execution time. The time consumption follows the expected $n(\log n)$ time complexity, and there is no significant increase in energy consumption when larger parts of memory are used during execution.

The total execution times are greater than 3 seconds, see Table 3; the total executions times are not directly comparable because the number of iterations within each run differ, see Table 1. Table 3 also shows a 2-second difference between the Siglent and the Raspberry time measurements. This is because we have chosen to count an experiment as when the Raspberry current increases beyond a certain threshold; in this case, when the current reaches 0.55A and more. This methodology cuts the execution time short on both ends. The energy consumption is calculated based on the logged times from the Raspberry Pi and the average power dissipation. While the execution time becomes correct, this method increases the average power dissipation slightly, and thus, we report a slightly larger energy consumption for the Raspberry Pi.


**Fig. 8.** Heapsort's time and energy consumption within FPGA (k=16) and Raspberry Pi (k=4)

Table 3. Heapsort's total execution time (s) per input. They are not directly comparable because the number of repetitions differ

Input size	Raspberry Pi (Siglent)	FPGA (CPU)
4096	2.536	4.623
6144	5.017	7.15
8192	7.712	9.888
10240	11.388	12.777
12288	13.672	15.499
14336	16.495	18.450
16384	19.095	21.198

**Fig. 9.** Dijkstra: Avg. Time and Energy consumption within FPGA and Raspberry Pi

5.2. Dijkstra

The results show that Dijkstra's shortest path algorithm does not necessarily benefit from hardware implementations when considering performance or energy, see Figure 9. The FPGA platform demonstrates higher average times and energy consumption per iteration than the Raspberry Pi across various graph configurations. For example, in scenarios like dense directed graphs, it consumed 0.232 mJ per iteration compared to 0.064 mJ by the Raspberry Pi, demonstrating a significantly higher energy usage, see Table 4.

These results suggest that while FPGAs are typically considered for their potential to enhance performance through parallelism, their energy efficiency for Dijkstra's algorithm is less effective than traditional processing on a Raspberry Pi. This revelation is particularly crucial for applications where energy consumption is as critical as processing speed, such as in portable or embedded devices where power conservation is essential.

6. Discussion of Results

Our results reveal new trade-off considerations between algorithm efficiency and performance. This discussion contextualizes our findings, leading directly to Section 9 where we summarize the key insights and summarize avenues for further research.

Table 4. Dijkstra’s average time and energy consumption within FPGA and Raspberry Pi

Experiment	Raspberry P			FPGA		
	Time pr. iter. (ms)	Power (W)	Energy pr. iter. (mJ)	Time pr. iter. (ms)	Power (W)	Energy pr. iter. (mJ)
Dense Directed	0.019	3.32	0.064	0.477	0.488	0.232
Dense Undirected	0.021	3.377	0.072	0.488	0.505	0.246
Sparse Directed	0.015	3.242	0.049	0.483	0.448	0.216
Sparse Undirected	0.015	3.260	0.049	0.49	0.446	0.218

6.1. Algorithm Characteristics and Hardware Suitability

Heapsort and Dijkstra’s algorithms have different levels of inherent parallelism and complexity. Heapsort may benefit more from native parallelization in hardware due to its ability to efficiently parallelize the comparison and swapping elements, especially when sorting larger datasets. In contrast, our implementation of Dijkstra’s algorithm may not fully leverage FPGA capabilities due to its sequential dependencies, leading to less impressive gains or even inefficiencies on FPGAs.

While our results are less promising for Dijkstra, previous studies show that FPGA implementations of Dijkstra can provide considerable performance optimization compared to software due to their different growth rates: “The average execution time of the FPGA-based version grew only linearly, whereas the average execution time of the microprocessor-based version displayed quadratic growth” [28], see Table 5. Thus, it may be that increasing the graph size can improve the results, but we speculate that our result is a consequence of our implementation. Future work would include reimplementations and energy evaluations of previously successful hardware implementations.

Table 5. Dijkstra performance improvements by Tommiska et al. [28] when using an FPGA solution compared to a Microprocessor solution

Vertices	Logic Elements	Memory	FPGA time	Microprocessor time	Speedup
8	834	632 bits	10.6 μ s	250 μ s	23.58
16	1536	2116 bits	13.4 μ s	434 μ s	32.39
32	2744	8287 bits	17.2 μ s	802 μ s	46.63
64	5100	32894 bits	21.6 μ s	1456 μ s	67.41

6.2. Energy Efficiency versus Performance

Both algorithms emphasize the trade-offs between performance gains and energy efficiency. In both algorithms, the FPGAs increased execution time. Energy and time are highly correlated. However, comparing the execution time across software and hardware does not necessarily indicate a similar pattern for their energy consumption.

The benefit of the FPGA is that the power dissipation is 5 to 8 times lower than the Raspberry Pi's power dissipation. Therefore, a good rule of thumb would be that FPGA is a good choice when the execution time of the FPGA is 5 to 8 times faster than the Raspberry Pi. In our study, this is the case for Heapsort but not for Dijkstra. In addition, this factor may change with a different choice of hardware. Future work would include evaluating the energy and time trade-off factors for various FPGAs and computers.

6.3. System Architecture Design

Insights from both algorithms can guide system architects in designing more efficient systems by choosing the right combination of hardware and software based on the specific algorithms they expect to run most frequently. There is no clear case for always employing FPGAs to provide energy reductions equivalently to performance. However, in some cases, the developer may be able to exploit native parallelism and ensure that hardware implementation has a slower execution time growth rate for increasing input sizes. It would be beneficial to evaluate the energy consumption of high-performing hardware implementations and to identify trade-off trends on the architectural design level.

6.4. Hypotheses on Algorithmic Performance

Our results demonstrate that Heapsort, with its potential for parallel processing, capitalized on the FPGA's architecture to yield significant energy savings. One hypothesis for this outcome is that the FPGA's ability to conduct multiple comparisons in parallel, particularly through the efficient use of LUTs and parallel memory access, minimizes both time and energy when compared to the CPU's more serial processing. For instance, as the input size increased, the FPGA continued to scale effectively, likely due to its architectural suitability for handling concurrent operations. This suggests that algorithms with similar local parallelism would exhibit comparable performance benefits.

In contrast, Dijkstra's algorithm, with its inherent sequential dependencies, struggled to take advantage of FPGA's strengths. We hypothesize that this inefficiency stems from the algorithm's need to update path lengths iteratively, which bottlenecks performance and limits potential energy gains. Future research could investigate whether alternate graph traversal algorithms with fewer sequential dependencies, or different graph configurations (such as sparse versus dense graphs), might better exploit hardware acceleration. Additionally, hardware design optimizations targeting these sequential steps could mitigate some of these limitations.

7. The Limits of Speedup

Building on the hypothesis that algorithms with local parallelism would exhibit similar performance benefits, we explored the upper bounds of FPGA speedup using the highly parallelizable Conway's Game of Life [4]. This experiment serves to demonstrate how extreme parallelism can push the limits of performance and energy efficiency on FPGA, further supporting the idea that parallel algorithms are well-suited for hardware acceleration. Conway's Game of Life is a zero-player game defined on cellular automata. The cellular automata are 2D grids called worlds, where each grid cell has eight neighbors,

and each cell can have one of two states: dead or alive. For each step, the cell states are updated according to their state and the states of their neighboring cells in the previous step.

1. Any live cell with two or three live neighbors survives.
2. Any dead cell with three live neighbors becomes a live cell.
3. All other live cells die in the next generation. Similarly, all other dead cells stay dead.

The initial state is given as input to the program. Because each cell depends only on nearby cells, Game of Life is highly parallelizable.

Table 6. The execution time in us and speed-up of FPGA over software executions

World	Cells	Execution time per step (us)			FPGA Speedup	
		Mac	Raspberry Pi	FPGA	Mac	Raspberry Pi
10x10	100	0.10	1.783	0.0040	25	445
20x20	400	0.33	5.137	0.0040	82	1284
30x30	900	0.70	9.965	0.0041	170	2430
40x40	1600	1.21	17.212	0.0040	302	4302
50x50	2500	1.81	25.204	0.0044	411	5728
60x60	3600	2.76	37.822	0.0045	613	8404
70x70	4900	3.54	57.665	0.0040	884	14416
80x80	6400	4.81	64.396	0.0047	1023	13701
90x90	8100	6.50	81.309	0.0045	1444	18068
100x100	10000	7.51	109.964	0.0048	1564	22909

Table 6 shows the average execution time for a single time step and the speedup provided by the FPGA implementation compared to the Java software implementation executed on a MacBook Pro and the Raspberry Pi. These results show that the gain in performance increases with the measured world sizes.

Speedup factors range from 25 for a 10x10 world to 1500 for a 100x100 world. The speedup scales linearly with the problem size. While the Game of Life is an artificial workload, its parallelizable nature made it an ideal candidate for indicating an upper bound for speedups when moving algorithms from software into an FPGA.

7.1. Resource Utilization

We have implemented the Game of Life of different sizes in a Cyclon IV FPGA found on the DE2-115 evaluation board. We report the design size in logic elements and registers. A logic element represents one 4-bit lookup table. For synthesis, we used the Quartus 19.1.0 Lite Edition.

Table 7 shows the FPGA implementation's resource consumption for different world sizes. We can see that the size grows linear. The maximum frequency of the circuit is reported between 209 MHz and 250 MHz. Therefore, when we assume running it at 200 MHz we can compute one iteration in 5 ns.

Table 7. The resource utilization and minimum iteration time of different sized Game of Life worlds in an FPGA

World	Logic Elements	Registers	Minimum Clock Period
10 x 10	804	104	4.0 ns
20 x 20	3539	404	4.0 ns
30 x 30	7995	904	4.1 ns
40 x 40	14463	1604	4.0 ns
50 x 50	23439	2504	4.4 ns
60 x 60	34414	3604	4.5 ns
70 x 70	45119	4904	4.0 ns
80 x 80	59136	6404	4.7 ns
90 x 90	75102	8104	4.5 ns
100 x 100	97871	10004	4.8 ns

As expected, we use one register per cell. However, the number of logic elements per cell is surprisingly high, an average of around 9 logic elements per cell. We assume that the Chisel PopCount method has some room for improvement. However, as we aim for a technique that enables software developers to describe their algorithms in hardware, we are avoiding optimization tricks.

7.2. Estimated Energy Consumption

We did not measure power or energy consumption of the FPGA implementation. However, the DE2-115 FPGA board comes with a power supply of 24 W. Therefore, this is the upper bound of power consumption of the whole FPGA board, including peripheral devices and external memories.

If we assume 24 W as an upper bound on the power consumption and an operating frequency of 200MHz, then one iteration of a 100 x 100 Game of Life world consumes 96nJ. In comparison, the Raspberry Pi has been reported to consume an average of 6.4 W when all four cores are busy⁶ and one iteration of a 100x100 world takes 0.109964 ms. Thus, a conservative energy consumption estimate for one iteration of a 100x100 world is 0.703769 mJ. From these conservative estimates, the hardware implementation can significantly improve energy consumption compared to the Raspberry Pi 4.

8. Related Work

FPGA research spans a diverse range of applications, including machine learning, signal processing, communication systems, and big data [25]. Sorting algorithms, while less commonly studied, represent a valuable niche for benchmarking hardware capabilities due to their computational diversity and potential for parallelism. There are numerous studies on hardware implementations of algorithms (e.g., [7, 12, 14, 15, 19, 20, 22, 28, 30]); however, only a few focus on energy consumption or compare performance across platforms. This work fills these gaps by focusing on energy and runtime variability across

⁶ <https://www.pidramble.com/wiki/benchmarks/power-consumption>

multiple platforms, emphasizing the reproducibility of evaluations and the role of algorithm characteristics.

Jmaa et al. [14] study the implementation of sorting algorithms on FPGAs using high-level synthesis, comparing their performance with CPU-based implementations. They evaluate algorithms such as BubbleSort and QuickSort, focusing on execution time and its variability to demonstrate the acceleration benefits of FPGAs. Their results show that FPGAs significantly outperform CPUs in execution time for sorting tasks, particularly for larger input sizes. Their results align with ours in highlighting the benefits of FPGAs for parallelizable tasks like Heapsort, where execution time improves significantly compared to CPUs. However, their study does not consider energy efficiency, a key focus of our work, nor does it explore intra-platform variability or reproducibility. However, their work does not explore the impact of algorithm characteristics, such as parallelism, on energy efficiency, a key focus of our study.

In a separate study, Jmaa et al. [15] analyze sorting algorithms implemented as software on ARM Cortex A9 processors, part of the Zynq platform. They measure computational time, energy consumption, and stability to determine the most efficient sorting algorithm for embedded systems. Their results highlight ShellSort as the most efficient algorithm for larger input sizes, with a high correlation between energy consumption and execution time. This study is close in content to our study, but they do not compare the hardware implementations with other implementations. Instead, they compare the time and energy usage of their FPGA implementations. Their results partially align with ours by showing that energy consumption is highly correlated with execution time.

Asiatici et al. [2] investigate the use of FPGAs as accelerators in a hybrid CPU-FPGA system for parallel string sorting. Their results demonstrate that FPGAs can significantly reduce runtime and improve energy efficiency for classification tasks, complementing CPU operations. While their study highlights FPGA efficiency and energy efficiency for parallel workloads, its focus on hybrid systems differs from our evaluation of standalone FPGA and CPU performance.

Mihhailov et al. [21] propose parallel FPGA-based implementations of recursive sorting algorithms, leveraging hierarchical finite state machines to improve performance. They focus on execution time, demonstrating significant performance gains through parallelization on FPGAs. Their results align with ours in showing the benefits of FPGAs for parallelizable tasks like Heapsort, particularly in terms of execution time. However, their study does not include energy efficiency analysis or comparisons with CPUs.

Zhou et al. [30] investigate the implementation of the A* algorithm on FPGAs for real-time path planning, leveraging a custom hardware architecture to optimize performance. Their study primarily focuses on execution time and resource utilization, demonstrating significant speed improvements over CPU implementations, particularly for large-scale pathfinding tasks. Their results align with ours in showing that FPGAs can speed up computationally intensive tasks but differ in focusing exclusively on path planning rather than broader algorithmic comparisons. Unlike our study, Zhou et al. do not consider energy efficiency or intra-platform variability.

Koch et al. [19] explore FPGA-based sorting architectures designed for large datasets, leveraging run-time reconfiguration to optimize performance. They measure execution time and demonstrate the scalability of their approach, particularly for sorting tasks that require high throughput. Their results align with ours in showcasing the benefits of FPGAs

for sorting tasks like Heapsort, particularly for large input sizes. However, they do not address energy consumption, a central focus of our study.

Mueller et al. [22] investigate data processing on FPGAs, comparing their performance with modern CPUs. They focus on asynchronous sorting networks and their integration into heterogeneous computing systems. Their results indicate that FPGAs can achieve competitive performance with CPUs while offering significant advantages in power consumption. Their study aligns with ours in comparing CPUs and FPGAs but differs in its broader focus on data processing and integration strategies rather than algorithm-specific energy and runtime variability.

Tommiska and Skyttä [28] implement Dijkstra's shortest path algorithm on FPGAs, achieving notable speed advantages over traditional CPU-based implementations. Their study focuses on execution time improvements enabled by FPGA-specific optimizations, demonstrating the potential for hardware acceleration in pathfinding tasks. As discussed earlier in Section 6, their performance results do not fully align with ours. They find FPGAs to be advantageous for Dijkstra's algorithm and we speculate that result difference is due to our hardware implementation. Their study does not explore energy efficiency.

Lei et al. [20] propose an FPGA implementation of the single-source shortest path (SSSP) problem using a systolic array priority queue. They evaluate the performance of their approach in terms of execution time and scalability. Their results show that FPGAs outperform CPUs for large-scale graph problems. Their results partially align with ours, as they highlight the advantages of FPGAs for large-scale, parallelizable problems, but differs in evaluating only execution time without considering energy efficiency.

Favaro et al. [7] compare multi-core CPUs and FPGAs for numerical linear algebra tasks, such as dense matrix multiplication and sparse matrix-vector multiplication. Their findings show that multi-core CPUs outperform FPGAs in runtime for smaller or less parallelizable tasks, benefitting from efficient caching and optimized libraries like Intel MKL. In contrast, FPGAs excel in energy efficiency for larger, highly parallelizable workloads like SPMV, though their runtime advantage remains limited. This aligns with our findings for tasks like Heapsort, which leverage parallelism effectively on FPGAs, and contrasts with sequential tasks like Dijkstra's algorithm, where CPUs dominate. While their study aligns with our emphasis on task-dependent hardware suitability, it does not address how software variability or task design impacts energy and runtime efficiency within a single hardware platform, which is central to our work.

9. Conclusion

This study compared the energy efficiency and performance of software and hardware implementations of Heapsort and Dijkstra's algorithms. Our findings reveal both the advantages and limitations of using FPGAs compared to traditional software implementations on a Raspberry Pi.

For Heapsort, the hardware implementation on an FPGA demonstrated a clear advantage in terms of energy efficiency, confirming the potential of FPGAs for algorithms where some operations can be done in parallel. The results showed that the energy consumption for Heapsort on FPGA was consistently lower than on the Raspberry Pi, particularly as input sizes increased. This suggests that FPGAs can effectively reduce energy consumption for tasks where some parallel processing can be exploited.

In contrast, Dijkstra's algorithm did not exhibit the same level of energy efficiency on FPGA. Despite FPGAs' inherent capabilities for handling parallel tasks, the complex dependencies and sequential nature of Dijkstra's algorithm limited the expected gains. The study highlighted that traditional CPU implementations might still hold an advantage in terms of both performance and energy consumption for algorithms with significant sequential operations.

The comparison also underscored the importance of choosing hardware or software solutions based on the specific requirements and characteristics of the algorithm. While FPGAs offer considerable reductions in power dissipation, they are not universally superior for all computational tasks. Our findings suggest that the decision to use FPGA over CPU should be guided by a more detailed knowledge of the algorithm's structure and the potential for parallelism.

Additionally, this study contributes to the ongoing discussion about the trade-offs between computational speed and energy efficiency. It provides a first step towards a nuanced perspective that can aid system architects and developers in making informed decisions about the hardware-software configurations that best meet their performance and efficiency goals.

While this study focuses on Heapsort and Dijkstra's algorithm, both of which offer limited parallelism, future work will explore more parallelizable algorithms, such as quicksort, matrix operations or image processing, to better demonstrate FPGA's energy and performance potential. Our preliminary results with Conway's Game of Life (Section 7) show that substantial speedups can be achieved through parallel computation, and similar gains are expected from a wider range of algorithms. This suggests that algorithms with similar characteristics can significantly reduce energy consumption by distributing computational tasks evenly across FPGA's processing elements. We hypothesize that applying this approach to a wider range of highly parallelizable algorithms will result in similar improvements in both energy efficiency and performance, as distributing tasks across FPGA's processing elements can significantly reduce energy consumption. These extensions will provide further insights into how parallelism can be optimized to enhance both computational speed and energy savings across a variety of workloads.

Future work will focus on optimizing hardware implementations and expanding the range of algorithms tested to further explore their energy and performance potential. Additionally, we plan to evaluate the scalability and applicability of our findings across a wider variety of hardware platforms, including more powerful multicore processors, GPUs, and different FPGA models and configurations. This broader assessment will help establish more generalized guidelines for selecting between software and hardware implementations, particularly in terms of energy efficiency and performance metrics.

References

1. Andrae, A.S.G.: New perspectives on internet electricity use in 2030. *Engineering and Applied Science Letter* 3, 19–31 (2020), <https://doi.org/10.30538/psrp-eas12020.0038>
2. Asiatici, M., Maiorano, D., Ienne, P.: How many cpu cores is an fpga worth? lessons learned from accelerating string sorting on a cpu-fpga system. *Journal of Signal Processing Systems* 93, 1405–1417 (12 2021), <https://doi.org/10.1007/s11265-021-01686-8>

3. Bachrach, J., Vo, H., Richards, B.C., Lee, Y., Waterman, A., Avizienis, R., Wawrzynek, J., Asanovic, K.: Chisel: constructing hardware in a scala embedded language. In: Groeneveld, P., Sciuto, D., Hassoun, S. (eds.) *The 49th Annual Design Automation Conference 2012, DAC '12*, San Francisco, CA, USA, June 3-7, 2012. pp. 1216–1225. ACM (2012), <https://doi.org/10.1145/2228360.2228584>
4. Berlekamp, E.R., Conway, J.H., Guy, R.K.: *Winning ways for your mathematical plays*. Vol. 2. Academic Press Inc., London (1982), games in particular
5. Cormen, T.H., Leiserson, C.E., Rivest, R.L., Stein, C.: *Introduction to Algorithms*. MIT Press, Cambridge, MA, 3rd edn. (2009)
6. Eder, K., Gallagher, J.P., López-García, P., Muller, H., Banković, Z., Georgiou, K., Haemmerlé, R., Hermenegildo, M.V., Kafle, B., Kerrison, S., Kirkeby, M., Klemen, M., Li, X., Liqat, U., Morse, J., Rhiger, M., Rosendahl, M.: Entra: Whole-systems energy transparency. *Microprocessors and Microsystems* 47, 278–286 (2016), <https://doi.org/10.1016/j.micpro.2016.07.003>
7. Favaro, F., Dufrechou, E., Ezzatti, P., Oliver, J.P.: Energy-efficient algebra kernels in fpga for high performance computing. *Journal of Computer Science & Technology* 21(2), 80–92 (2021)
8. Field, H., Anderson, G., Eder, K.: Eacof: a framework for providing energy transparency to enable energy-aware software development. In: *Proceedings of the 29th Annual ACM Symposium on Applied Computing*. p. 1194–1199. SAC '14, Association for Computing Machinery, New York, NY, USA (2014), <https://doi.org/10.1145/2554850.2554920>
9. Hähnel, M., Döbel, B., Völz, M., Härtig, H.: Measuring energy consumption for short code paths using rapl. *SIGMETRICS Perform. Eval. Rev.* 40(3), 13–17 (Jan 2012), <https://doi.org/10.1145/2425248.2425252>
10. Hirata, H., Nunome, A.: A modified parallel heapsort algorithm. *International Journal of Software Innovation* 8(3), 1–18 (2020), <https://doi.org/10.4018/IJSI.2020070101>
11. Intel Corporation: *Intel 64 and IA-32 Architectures Software Developer's Manual* (2023), <https://www.intel.com/content/www/us/en/developer/articles/technical/intel-sdm.html>, accessed: 2024-11-04
12. Jagadeesh, G., Srikanthan, T., Lim, C.: Field programmable gate array-based acceleration of shortest-path computation. *IET Computers & Digital Techniques* 5, 231–237 (2011), <https://doi.org/10.1049/iet-cdt.2009.0072>
13. Jasika, N., Alispahic, N., Elma, A., Ilvana, K., Elma, L., Nosovic, N.: Dijkstra's shortest path algorithm serial and parallel execution performance analysis. In: *2012 Proceedings of the 35th International Convention MIPRO*. pp. 1811–1815 (2012), <https://ieeexplore.ieee.org/document/6240942>, accessed: 2024-11-04
14. Jmaa, Y.B., Atitallah, R.B., Duvivier, D., Jemaa, M.B.: A comparative study of sorting algorithms with fpga acceleration by high level synthesis. *Computacion y Sistemas* 23, 213–230 (2019), <https://doi.org/10.13053/CyS-23-1-2999>
15. Jmaa, Y.B., Duvivier, D., Abid, M.: Sorting algorithms on arm cortex a9 processor. In: Barolli, L., Woungang, I., Enokido, T. (eds.) *International Conference on Advanced Information Networking and Applications*. vol. 227, pp. 355–366. Springer International Publishing (2021), https://doi.org/10.1007/978-3-030-75078-7_36
16. Khan, K.N., Hirki, M., Niemi, T., Nurminen, J.K., Ou, Z.: RAPL in Action. *ACM Transactions on Modeling and Performance Evaluation of Computing Systems* 3(2), 1–26 (jun 2018), <https://doi.org/10.1145/3177754>
17. Kirkeby, M.H., Krabben, T., Larsen, M., Mikkelsen, M.B., Petersen, T., Rosendahl, M., Schoeberl, M., Sundman, M.: Energy consumption and performance of heapsort in hardware and software (2022), <https://arxiv.org/abs/2204.03401>
18. Kirkeby, M.H., Schoeberl, M.: Towards comparing performance of algorithms in hardware and software (2022), <https://arxiv.org/abs/2204.03394>

19. Koch, D., Torresen, J.: Fpgasort: a high performance sorting architecture exploiting run-time reconfiguration on fpgas for large problem sorting. In: *Proceedings of the 19th ACM/SIGDA International Symposium on Field Programmable Gate Arrays*. p. 45–54. FPGA '11, Association for Computing Machinery, New York, NY, USA (2011), <https://doi.org/10.1145/1950413.1950427>
20. Lei, G., Dou, Y., Li, R., Xia, F.: An fpga implementation for solving the large single-source-shortest-path problem. *IEEE Transactions on Circuits and Systems II: Express Briefs* 63(5), 473–477 (2016), <https://doi.org/10.1109/TCSII.2015.2505998>
21. Mihhailov, D., Sklyarov, V., Skliarova, I., Sudnitson, A.: Parallel fpga-based implementation of recursive sorting algorithms. In: *2010 International Conference on Reconfigurable Computing and FPGAs*. pp. 121–126 (2010)
22. Mueller, R., Teubner, J., Alonso, G.: Data processing on fpgas. *Proc. VLDB Endow.* 2(1), 910–921 (aug 2009), <https://doi.org/10.14778/1687627.1687730>
23. Pereira, R., Couto, M., Ribeiro, F., Rua, R., Cunha, J., Fernandes, J.P., Saraiva, J.: Ranking programming languages by energy efficiency. *Sci. Comput. Program.* 205, 102609 (2021), <https://doi.org/10.1016/j.scico.2021.102609>
24. Rotem, E., Naveh, A., Ananthakrishnan, A., Weissmann, E., Rajwan, D.: Power-management architecture of the intel microarchitecture code-named sandy bridge. *IEEE Micro* 32(2), 20–27 (2012), <https://doi.org/10.1109/MM.2012.12>
25. Ruiz-Rosero, J., Ramirez-Gonzalez, G., Khanna, R.: Field programmable gate array applications—a scientometric review. *Computation* 7(4) (2019), <https://www.mdpi.com/2079-3197/7/4/63>
26. Santos, B., Fernandes, J.P., Kirkeby, M.H., Pardo, A.: Compiling haskell for energy efficiency: Empirical analysis of individual transformations. In: *The 39th ACM/SIGAPP Symposium on Applied Computing (SAC '24)*, April 8–12, 2024, Avila, Spain. Association for Computing Machinery, New York, NY, United States (2023), <https://doi.org/10.1145/3605098.3635915>
27. Schoeberl, M.: *Digital Design with Chisel*. Kindle Direct Publishing (2019), available at <https://github.com/schoeberl/chisel-book>
28. Tommiska, M., Skyttä, J.: Dijkstra's shortest path routing algorithm in reconfigurable hardware. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 2147, 653–657 (2001), https://doi.org/10.1007/3-540-44687-7_73
29. Xilinx: Vivado design suite user. guide power analysis and optimization (October 2021), [https://docs.amd.com/r/2021.2-English/ug907-vivado-power-analysis-optimization, uG907 \(v2021.2\), Accessed: 2024-11-04](https://docs.amd.com/r/2021.2-English/ug907-vivado-power-analysis-optimization, uG907 (v2021.2), Accessed: 2024-11-04)
30. Zhou, Y., Jin, X., Wang, T.: FPGA implementation of a* algorithm for real-time path planning. *Int. J. Reconfigurable Comput.* 2020, 8896386:1–8896386:11 (2020), <https://doi.org/10.1155/2020/8896386>

Maja H. Kirkeby is Associate Professor at Roskilde University, Denmark, specializing in energy consumption of software and IT systems.

Thomas Krabben is Lab Manager at FlexLab, Roskilde University, Denmark, with expertise in hardware and IT systems setup and operational management.

Mathias Larsen is MSc student at IT University, Denmark, with focus on in energy consumption of software algorithms.

Maria B. Mikkelsen is Research Assistant at Roskilde University, Denmark, specialized in energy consumption and program transformation of computational systems.

Mads Rosendahl is Associate Professor at Roskilde University, Denmark, with expertise in program analysis and program transformation.

Martin Sundman is MSc student at IT University, Denmark, specialized in energy consumption of software algorithms.

Tjark Petersen is MSc student at DTU Compute, Denmark, expertise in efficient hardware algorithm design and implementation.

Martin Schoeberl is Professor at DTU Compute, Denmark, world-leading research on real-time systems and hardware design.

Received: September 14, 2024; Accepted: December 19, 2024.

