

PREPRINT - GPU-based Aim Point Optimization for Solar Tower Power Plants

Laurin Oberkirsch^{a,*}, Daniel Maldonado Quinto^a, Peter Schwarzbözl^a, Bernhard Hoffschmidt^a

^aGerman Aerospace Center (DLR), Institute of Solar Research, Linder Höhe, D-51147 Köln, Germany

Abstract

In solar tower power plants, aim point optimization is suitable to find aim point distributions resulting in intercept powers close to the theoretical maximum. However, the application in real time operation often faces the problem of long optimization duration. To counteract this issue, the convergence of an existing strategy, the ant colony optimization meta-heuristic, is enhanced. The raytracing is already replaced by pre-calculated flux maps of the individual heliostats in previous works to increase the optimization speed. In this work, the optimization is merged with a grouping strategy and implemented on a GPU to achieve further time reductions. Here, a k-means clustering algorithm performs the heliostats grouping. The use of groups reduces the solution space for the optimizer and additionally the amount of pre-calculated flux maps, so that the data fits in the global memory of the GPU. Over 100 billion flux values can be evaluated per second using this adapted approach. In this way, the algorithm finds suitable aim point distributions within a few seconds up to a minute. The achieved intercepts are 1 % to 4 % higher than those found by Vant-Hull aiming for the evaluated central receiver reference power plant. Moreover, the approach has proved its applicability in clouded environments that lead to spatially fluctuating solar radiation. There, a spillage reduction compared to Vant-Hull aiming of 35 % is reached.

Keywords: Concentrating solar power, Solar tower power plant, Heliostat aiming, Aim point optimization, Cloud disturbance

1. Introduction

Solar tower power systems have to reach higher efficiencies and need to overcome durability issues to be competitive for energy generation. The highest efficiency is usually reached, when all heliostats of a solar tower power plant aim at the receiver's center. However, the receiver as one of the critical components of the system is then exposed to high peak fluxes leading to material corrosion and thermal stresses (Sánchez-González et al., 2020). Hence, well suited aiming strategies are necessary to observe temperature and stress limits and to operate at high efficiencies simultaneously (Maldonado et al., 2018). These strategies are especially challenged by high variability in the solar radiation. In particular, cloud passages complicate the safe operation, as transient incident flux distributions on the receiver increase the thermal stresses (Chu and Coimbra, 2017; García et al., 2018). Early approaches like Vant-Hull aiming (Vant-Hull, 2002) coupled with a dynamic aim point processing system applied at the Solar Two plant (Lipps and Vant-Hull, 1978) observe flux constraints during operation at varying environmental conditions. According to Vant-Hull (2002), the aim points are spread vertically from the edges of a cylindrical receiver dependent on the heliostats' beam

radii multiplied by a factor k . Moreover, the dynamic aim point processing system of Solar Two identifies flux densities exceeding the allowable flux density. It determines the heliostats causing the highest flux at this location and removes them from tracking. Finally, their fluxes are subtracted from the measured flux until it complies with the constraints to protect the receiver (Lipps and Vant-Hull, 1978). This approach is suited for the control of solar tower power plants as it is safe, but it does not optimize the intercept or receiver power.

On the one hand, Collado and Guallar (2019) extended the Vant-Hull (2002) strategy by introducing an additional factor k_3 for the heliostats located further away from the tower. Selecting $k_3 \geq k$ smooths the flux profile, which usually exhibits two peaks. Thus, the maximal flux density is reduced and simultaneously the efficiency increases slightly. On the other hand, Astolfi et al. (2017) as well as Sánchez-González et al. (2018) also expanded the Vant-Hull (2002) strategy by segmenting the field in azimuthal sectors. For each sector, an individual factor k is determined leading to lower spillage. Moreover, Sánchez-González et al. (2017) added a receiver model that calculates the *Allowable Flux Density* (AFD) based on thermal stresses and corrosion limits. This strategy finds solutions that observe the AFD within 2 min for a Gemasolar like plant.

For the same plant model, García et al. (2017) merged both extensions of the Vant-Hull (2002) strategy as they divided the field in 18 azimuthal sections, each with three groups.

*Corresponding author

Email address: laurin.oberkirsch@dlr.de (Laurin Oberkirsch)

Further, they applied a *multi-input multi-output* (MIMO) closed control loop based on *Proportional Integral Derivative* (PID) controllers. Subsequently, García et al. (2018) analyzes the dynamic performance of the method under transient conditions caused by clouds. In this case, the control of the MIMO system is supported by a *Model Predictive Control* (MPC) technique, called *Dynamic Matrix Control* (DMC). For this method, Soo Too et al. (2019) especially concerns the arising material corrosion and thermal stresses.

Transient conditions were also studied by Ashley et al. (2017) as they solved the knapsack problem of assigning heliostats to aim points by *Binary Integer Linear Programming* (BILP) in a time of 30 s using the PS10 as reference power plant. On this basis, Richter et al. (2019) developed a robust *Mixed Integer Linear Programming* (MILP) aiming strategy considering tracking errors as uncertainties by a Gamma robustness approach. Further studies to this approach are published in Kuhnke et al. (2020). Due to the increased complexity in consequence of the Gamma robustness the running time rises to 60 s.

In the recent years, several aiming strategies based on meta-heuristics are developed. Based on the HFLCAL convolution method, Salomé et al. (2013) applied a TABU algorithm that flattens the flux distribution for the THEMIS flat plate receiver, while minimizing spillage. Besarati et al. (2014) solves the same problem by replacing the TABU algorithm by a genetic one. Likewise, Yu et al. (2014) implements a TABU algorithm and Wang et al. (2017) a genetic algorithm for a cavity receiver.

Furthermore, Belhomme et al. (2013) found with the ant-colony optimization meta-heuristic (ACO) a method to maximize the intercept and comply with the limits independent of the receiver's shape. This method outperforms the Vant-Hull strategy for cylindrical receivers by 2 % according to Flesch et al. (2017). Moreover, the approach is coupled with a local search algorithm by Maldonado et al. (2018) to improve the convergence towards the global optimum. However, the ACO cannot substitute control strategies as it delivers the suited aim point distribution in periodic intervals e.g. one hour. During these intervals, the aim point distribution has to be adapted by other control techniques to react on changes in the sun position as well as on cloud passages.

In this literature review, two main drawbacks are noticeable:

- Solutions based on meta-heuristics or MILP are computationally expensive. Hence, the computational time is often too long for the application in dynamic environments with highly fluctuating solar radiation. Especially, if the plant size increases further e.g. more heliostats and more aim points, most of the algorithms reach their limits.
- Solutions based on parameter-based aiming show significantly higher spillage. Even though, the losses are

reduced by multi-parameter strategies, most of the strategies are prone to cloud passages in particular.

The downside of ACO (Belhomme et al., 2013) is clearly the first point e.g. the long optimization time. However, high intercepts are achieved according to Flesch et al. (2017). Hence, this work focuses on reducing considerably the running time of the ACO, while preserving the advantage of high intercepts.

The ACO is already based on pre-calculated flux maps as it clearly reduces the optimization time compared to raytracing. In this paper, the ACO is coupled with a grouping strategy to reduce the solution space and increase in consequence the convergence rate further. Moreover, it downsizes the amount of data, which in turn allows to shift the data into the device memory of a *Graphics Processing Unit* (GPU). As Izygon et al. (2011) demonstrated by implementing a raytracer on a GPU first, significant speed ups are achieved. Thus, GPU based raytracing was especially used for heliostat field layout design (Zhou and Zhao, 2014; Gebreiter et al., 2019). Moreover, analytical models based on elliptical Gaussian distributions became GPU-based (He et al., 2017, 2019). Due to the promised speed up, this work focuses on implementing an optimizer for heliostat-aim-point-assignment in solar tower power plants on a GPU. The ACO is especially suited as it offers a great parallelization potential. Besides, an adaption of the Vant-Hull (2002) strategy is applied as initial solution to increase robustness and reproducibility of the ACO. Finally, some test under realistic cloud conditions show the applicability of the method in reality.

2. Methods

This section introduces the ACO first. Afterwards, the pre-calculated flux maps are described before the grouping is explained in detail. Some information regarding GPUs follow and the major points of the GPU implementation are highlighted. Finally, some adaptations to the initial ACO are discussed.

2.1. Ant-colony optimization meta-heuristic

The ACO is a multi-agent method inspired by the pheromone-based communication of ants during foraging. In nature, ants randomly walk around until they find food. On the way back to their colony, they emit pheromones. If other ants discover this pheromone path, they are more likely to stay on this path and they will enhance the pheromone concentration, if they find food once again. Moreover, the pheromones evaporate with time. Hence, the concentration rises on highly frequented paths, but also on shorter paths as less ants are necessary to create the same pheromone concentration since they walk denser. In the analogy, the evaporation prevents the algorithm to converge in local optima.

According to Belhomme et al. (2013), the probability p_{ij} of an ant following path i to j is calculated as defined in

equation 1:

$$p_{ij} = \frac{\tau_{ij} \cdot \eta_{ij}^{\beta}}{\sum_{k \in N} \tau_{ik} \cdot \eta_{ik}^{\beta}}. \quad (1)$$

It is based on the local pheromone concentration τ_{ij} , the attractiveness η_{ij} and the parameter β . Here, β states the relevance of attractiveness to pheromone concentration and N the number of reachable destinations from i . When transferring the ant foraging to the optimization problem of assigning heliostats i to specific aim points j , each ant evaluates a combination of heliostat-aim-point-combinations. A combination of heliostat-aim-point-combinations represents one constellation, where each heliostat points to a certain aim point. The pheromone concentrations are stored in a pheromone matrix, which operates as a collective memory that is regularly updated during optimization. Similarly, the attractiveness e.g. the intercept factor of a heliostat assigned to an aim point is stored in an attractiveness matrix. During optimization several ants are used in parallel to form an ant generation. The paths ij of each ant or rather the combination of heliostat-aim-point-combinations is determined based on the probability and a random factor, set by a Monte-Carlo-Method, that differs for each ant. Finally, each combination of heliostat-aim-point-combinations is evaluated and the best one is selected to form the subsequent ant generation.

2.2. Pre-calculated flux maps

An advantage of the ACO is that the number of heliostats and aim points, which are usually defined on the receiver's surface in advance, is finite. Thus, the raytracing can be performed prior to the optimization for each heliostat-aim-point-combination. The emerging flux on the receiver's surface when assigning a single heliostat to a single aim point is in the following referred to as flux map. The flux distribution is different for each sun angle, heliostat and aim point. Additionally, the flux differs according to the *Direct Normal Irradiance* (DNI) onto the heliostat. However, the flux distribution remains constant and is just scaled by the DNI as a factor.

The raytracer STRAL can calculate 10 million rays per CPU core per second. Assuming a large plant with 1 million square meter heliostat surface, an accuracy of 600 rays per square meter and 400 aim points, the pre-calculation of one sun angle would require 40 min on 10 cores.

Even though the number of possible sun angles is infinite, the sun angle chart 1 shows that only sun angles between winter and summer solstice are relevant. With a grid of 2° lattice spacing in azimuth and elevation direction roughly 1830 sun angles are necessary for southern Spain requiring 50 days of computational time. This grid size leads to errors below 1 %. However, interpolation gives sufficiently accurate results even with a sparser grid.

The algorithm can use these pre-calculated flux maps and only has to superpose them during optimization.




Figure 1: Only sun angles between winter and summer solstice are relevant for the operation of a solar tower power plant. The courses are presented exemplary for southern Spain and computed by the model of Grena (2008). The graph is superimposed by a sketch of an exemplary grid with 2° grid spacing.

This significantly reduces the computation time. The attractiveness e.g. the intercept of a heliostat-aim-point-combination is constant for one sun angle. Hence, it is also part of the pre-calculation. Belhomme et al. (2013) states that compared to raytracing the computational time is reduced by a factor of $10^2 - 10^4$ depending on the complexity of the receiver model as well as the resolution of the flux on the receiver surface. At the same time, a slight accuracy loss emerges due to errors in shading and blocking calculation as the neighboring heliostats are assigned to a central aim point during pre-calculation instead of the actual one. However, if the field becomes larger, this error is negligible.

The main drawback of the ACO of Belhomme et al. (2013) was the long optimization duration leading to good optimization results only in periodic intervals of roughly one hour. This impedes the application in real plants under highly fluctuating conditions of reality like cloud shading. Thus, the interval for optimization results has to be considerably reduced. The use of pre-calculated flux maps for the optimization is already a good approach as it reduced significantly the optimization duration. Hence, they are also used in this work and generated by means of the raytracing software STRAL¹ for various sun angles to cover the application at different times of the day.

2.3. Grouping

For the optimization in the GPU, all pre-calculated flux maps of one sun angle have to fit in the global memory of the GPU. However, 40 GB would be required to store the pre-calculated flux maps as 4-byte float values assuming a large field with 10.000 heliostats and 400 aim points and 50×50 receiver bins. Hence, the amount of data needs to be reduced, which can be achieved by grouping. For example, grouping the heliostats in 500 representative groups would already result in only 2 GB that have to fit in the GPU memory.

In this work, the grouping is performed by means of the

¹STRAL is a raytracing software tool developed at the German Aerospace Center (Belhomme et al., 2009).

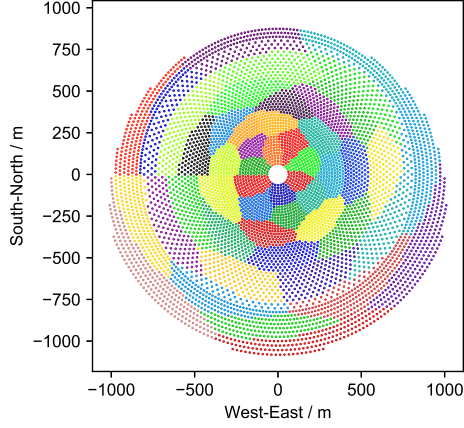


Figure 2: 6482 heliostats of a reference power plant are clustered exemplarily in 50 groups by the k-means clustering algorithm. Each group is highlighted by a different color.

k-means clustering algorithm (Lloyd, 1982). In figure 2, the 6482 heliostats of the later introduced reference power plant are exemplarily clustered in 50 groups. The clustering is performed by the parameters radius and the 2-argument arctangent $\arctan2$ of the circumferential angle. Alternatively, manual clustering as well as other algorithms are also feasible. However, this is a fast and automatic way to cluster fields with various shapes and sizes. If a group constellation is determined, the pre-calculated flux maps are superposed and stored in an overall flux map for each group.

Another advantage of the clustering is that the possible number of combinations of heliostat-aim-point-combinations, also referred to as solution space S , is reduced. Initially, the solution space for the optimizer is the number of aim points n_A to the power of the number of heliostats n_H :

$$|S| = n_A^{n_H}. \quad (2)$$

After the clustering, it decreases to the number of aim points to the power of the number of groups n_G :

$$|S| = n_A^{n_G}. \quad (3)$$

In this way, the optimizer finds better solutions in less optimization steps. Indeed, the theoretical maximum is also reduced due to the restricted solution space. Though, this drawback is acceptable for an adequate number of groups.

2.4. GPU implementation

Initially, the purpose of *Graphics Processing Units* (GPUs) was the acceleration of computer graphics and vision related tasks. By now, great computing performance has been identified in other fields and with CUDA from NVIDIA (NVIDIA Corporation, 2019) and OpenCL, the GPU programmability has been significantly improved.

In this way, the domain *General-Purpose computation on GPU* (GPGPU) emerged (Liu, 2018).

A typical CUDA program comprises kernel code, executed by the GPU device, and host code, executed by the CPU. A kernel consists of thousands of threads, which are organized in thread blocks. The kernel in turn is a grid of thread blocks. On NVIDIA hardware, groups of 32 threads within a thread block are scheduled as a unit and called warps (NVIDIA Corporation, 2019). When a kernel is called, the thread blocks are arranged to *Streaming Multiprocessors* (SM). Each SM has its own shared memory and L1 data cache. In parallel algorithms, the threads often need to share data. Therefore, the *syncthreads*-function allows to synchronize all threads of a thread block. The synchronization between threads of different thread blocks is not possible within the kernel. Thus, the different thread blocks can run parallel without waiting for the others to finish (Liu, 2018).

The ACO has to achieve good convergence in short optimization time to be applicable in reality. Hence, the ACO is implemented on the GPU as significantly higher throughput than on the CPU is promised. Due to the high degree of parallelization, the ACO is well suited. The flow chart, figure 3, sketches the workflow of the GPU implementation.

Initially, the pre-calculation maps for the current sun an-

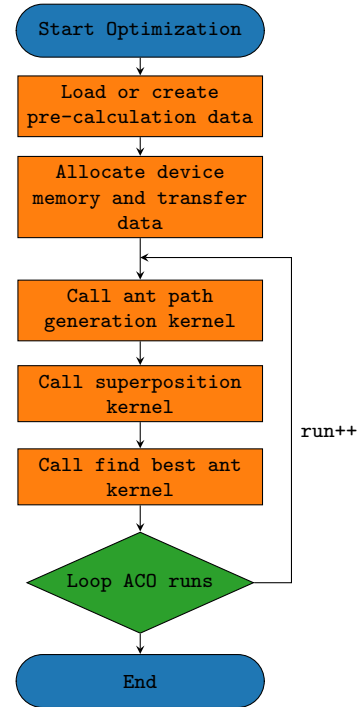


Figure 3: Flow chart illustrating the optimization workflow implemented on the GPU.

gle are only loaded to the CPU memory. At start of the optimization, the device memory is allocated and the data is shifted to the global memory of the GPU. This step has to be performed only once per sun angle.

Subsequently, the ant colony optimization can start. The combination of group-aim-point-combinations for each ant is computed similarly to the model of Belhomme et al. (2013). The pseudorandom number generator for cuda applications *Curand* replaces the Monte-Carlo-Method to create slightly different combinations of group-aim-point-combinations for each ant of one generation.

The superposition of the pre-calculated flux maps is the time-sensitive step. Per computed combination of group-aim-point-combinations, the flux maps of all groups are added to the final flux map. This means one addition per group, per ant and per receiver bin. Especially, the autonomous ants of each generation and the individual bins can be computed simultaneously and thus, provide a great potential for the parallelization in a GPU.

A memory instruction for global GPU memory requires 400-600 clock cycles. Consequently, the bottleneck of the GPU implementation of the ACO is the access of global memory and not the number of *Floating Point Operations Per Second* (FLOPS). Thus, the usage of shared memory, which requires only 2 clock cycles, is indispensable (Prisacariu et al., 2009). However, shared memory has not the size to store all pre-calculation maps. Hence, it is used to store the total flux map, on which each individual flux maps is added as well as a quality value for each thread. Hence, only the pre-calculation data is in global memory, but the receiver model is in shared memory. This allows the receiver model to be more complex without increasing the computational time.

Moreover, the speed of global memory access can be maximized by accessing it coalesced. Coalesced means that the 32 threads in a warp access consecutive 4-byte memory locations. This is the most efficient case as L1 cache always fetches 128-byte cache line even when only one 4-byte value is required. Thus, per memory instruction, which costs 400-600 clock cycles, 32 4-byte values can be maximally loaded (Liu, 2018).

In the superposition kernel, an individual thread is created for each bin of the receiver surface. The thread sums up the fluxes of all groups in this bin. The fluxes of one flux map are stored in consecutive memory. Moreover, the threads of one thread block deal with neighbored bins. Thus, the 128-byte sequence and full coalescence is achieved, which maximizes the speed of the memory access. As the memory access is the bottleneck of the kernel, it also maximizes the speed of the kernel itself. Finally, the find best ant kernel is called that evaluates the quality values determined by each parallelized ant and returns the best one. If the best ant is better than the current solution, it is applied as initial solution for the next loop e.g. the next ant generation is generated based on this ant.

2.5. Adaptions of the ACO

Further, the ACO has the drawback that it sometimes diverges, if the constraints become stricter. This can be counteracted by starting from a valid good initial guess.

In this work, a good initial solution is found by the Vant-Hull (2002) strategy combined with an Flesch et al. (2017) adaption. Vant-Hull (2002) multiplies the heliostats' beam radii with a factor. The heliostat's aim point is shifted from the edge of a cylindrical receiver towards the equator by this scaled beam radius. However, this often results in two hot spots; one on each side of the receiver's equator. Thus, Flesch et al. (2017) additionally includes the slant range of heliostat and center of the receiver r as an extension to the factor k of Vant-Hull according to equation 4:

$$k_{\sigma} = k + \epsilon \cdot r. \quad (4)$$

This adaption results in two determinable factors k and ϵ for a good initial guess. The additional factor supports the Vant-Hull strategy by providing further flexibility. In the following, this strategy is referred to as Flesch-Vant-Hull strategy.

To apply the Flesch-Vant-Hull strategy as initial guess for the ACO two adaptions are required. First, the ACO allows only specific aim points and has the condition that all heliostats of one group have to aim towards the same aim point. Hence, the mean aim point for all heliostats in a group is computed. Second, determining the flux distribution by raytracing always requires some seconds. Hence, the computational time is reduced by using the already pre-calculated flux maps. Therefore, the previously defined aim point grid is used and for each group, the closest existing aim point is selected. When referring to these strategies in this work, the word *grouped* is added. This results in the respective names grouped Vant-Hull strategy and grouped Flesch-Vant-Hull strategy.

Finally, the ACO is applied in dynamic environments, which are exposed to cloud passages. *All Sky Imager* (ASI) based nowcasting systems can determine and predict the local DNI (Nouri et al., 2018, 2019, 2020; Schroedter-Homscheidt et al., 2018). This DNI information can be mapped onto the individual heliostats or groups of heliostats. Prior to each optimization call, the pre-calculation results can be scaled by this DNI information before they are loaded into the GPU memory.

3. Results

In this section, the reference power plant used for the evaluation is introduced first. Subsequently, the speed of the GPU implementation is evaluated before a group study is performed. Then, the grouped Flesch-Vant-Hull algorithm is tested before it is used in the ACO. Finally, the ACO is compared to the Vant-Hull algorithm both for an even DNI situation as well as for a clouded case.

3.1. Reference power plant

The methods that improve the convergence rate of the ACO are tested at a reference power plant with a designed thermal power of 450 MW that is also used by Flesch et al.

(2017) and Maldonado et al. (2018). The plant has a cylindrical receiver and 6482 heliostats, each with a mirror surface of 121 m^2 . More detailed information about the virtual plant can be found in Flesch et al. (2017). In this work, the thermal model is not included to focus on the optical side of the plant. As an alternative, the *Allowable Flux Density* (AFD) is applied to limit the maximal flux density radiating onto the receiver's surface. For optimization purposes a fixed aim point grid with 470 aim points, 36 in circumferential direction times 13 in vertical direction plus 2 off receiver aim points, is employed. The 21th of March at noon is chosen as evaluation time. Initially, the DNI is 1000 W m^{-2} across the entire field.

3.2. GPU speed evaluation

The optimization on a NVIDIA Quadro P5000 performs in 98s 1000 ant generations runs, each with 16384 ants, for the reference power plant. Here, 200 groups are used and the receiver has 1500 bins, 60 in circumferential and 25 in vertical direction. The total time splits into the kernel times per run as presented in figure 4.

The time required by the superposition kernel is the




Figure 4: Kernel times in milliseconds for one run comprising 16384 ants, 200 groups and 1500 receiver bins.

longest. Here, a memory instruction and an addition operation is performed for each bin, for each group and for each ant. This results in $\frac{1500 \cdot 200 \cdot 16384}{0.044} = 112 \cdot 10^9$ flux values that are loaded and evaluated per second. The number of evaluated flux values reduces to $50 \cdot 10^9$, if the additional times for ant path generation kernel, find best ant kernel and some further CPU computations are included. However, these times scale far less than linear with the number of evaluated flux values. Thus, increasing the ants, groups or bins always results in a higher share of the superposition kernel time compared to the total time. Moreover, using a more complex receiver model only leads to a longer computational time in the superposition kernel, which is perfectly parallelized on the GPU, so that occupancies of 99 % to 100 % are achieved.

Tests with the GeForce RTX 2080 are roughly twice as fast as the Quadro P5000; tests with a low-priced GPU, the GeForce GTX 650 Ti, are roughly 0.3 times as fast. However, the issue becomes for this GPU the limited GPU memory size of only 1 GB.

3.3. Group study

The optimal group size for the ACO is determined by analyzing the convergence for different numbers of groups.

This is accomplished by figure 5 as it shows the intercept over optimization duration for various group numbers. Here, 800 kW m^{-2} is used as AFD as the results differ stronger for lower AFDs and valid solutions are harder to find.

On the one hand, low group numbers lead to strong




Figure 5: The convergence of the ACO e.g the intercept with respect to optimization time is shown for various number of groups into which the heliostats of the reference power plant are clustered.

convergence in the first seconds. Thus, the results are initially ordered contrary to the group size. On the other hand, higher group numbers result in greater intercepts if optimizing longer. 350 groups even reach an intercept beyond 0.948 in less than two minutes. The figure also exhibits that a group number of 450 is too large as the curve converges significantly slower, so that no better solution is found in an acceptable time frame.

Moreover, lower group numbers lead to better results, if the AFD restrictions are less challenging, whereas higher group numbers find better results for stronger restrictions.

3.4. Vant-Hull results

Before evaluating the ACO algorithm, the results of the Vant-Hull strategy are presented as it will be used as reference for the ACO in the following. Figure 6 shows the maximal flux density of the Flesch-Vant-Hull strategy with respect to the applied factors k and ϵ . According to equation 4, the curve for $\epsilon = 0$ represents the original Vant-Hull strategy. Furthermore, the color gradient illustrates the intercept. The figure demonstrates that raising k and ϵ results in enhanced intercepts, but also in higher maximal flux densities.

Moreover, a valley is formed in the curve of the maximal flux density. The bottom of the valley is highlighted by the black dashed line. It starts at $k = 2.6$ for the original Vant-Hull strategy with a maximal flux density of 771 kW m^{-2} and decreases further to $k = 0$ and $\epsilon = 4$ with 659 kW m^{-2} . Simultaneously, the intercept decreases from 0.91 to 0.85.

Additionally, the black dotted line describes a curve with constant maximal flux density of 771 kW m^{-2} . Initially, the intercept remains also constant, but for $\epsilon \geq 3$ a slight




Figure 6: The maximal flux density of the Flesch-Vant-Hull strategy for various factors k and ϵ is presented. The curve with $\epsilon = 0$ outlines the results of the original Vant-Hull strategy. The color gradient visualizes the intercept.




Figure 7: The maximal flux density and the intercept of the grouped Flesch-Vant-Hull strategy are shown for two different group numbers.

drop is identified.

As soon as, the grouped Flesch-Vant-Hull strategy is applied as initial solution for the ACO, the groups and the fixed grid of aim points are used. Thus, the pre-calculated flux maps can be applied and many $k - \epsilon$ - combinations can be tested in less than a second. As a result of the limited flexibility, the curves obtain a more ruffled shape as pictured in 7. However, the general shape remains unchanged, but is shifted towards higher maximal flux densities, when lowering the number of groups. Likewise, the minima increase significantly to 794 kW m^{-2} and 880 kW m^{-2} for 500 and 200 groups as outlined in figure 7a and 7b.

3.5. ACO evaluation

For a comparison, the Vant-Hull (2002) strategy is applied as this basic approach for cylindrical receivers is already introduced in this paper. Figure 8 presents the flux density distributions for an AFD of 800 kW m^{-2} for both strategies, the ACO and the Vant-Hull strategy. The intercept of the Vant-Hull solution is 93.2 %, whereas the one of the ACO is 94.5 %. This results in a performance enhancement of 1.4 % and a spillage reduction of 19 %.

The solution of the Vant-Hull strategy exhibits the typical two peak regions, one on the upper receiver half and one on the lower half, as pictured in figure 8a. Noticeable is the poor focus around bin 30 since only flux densities

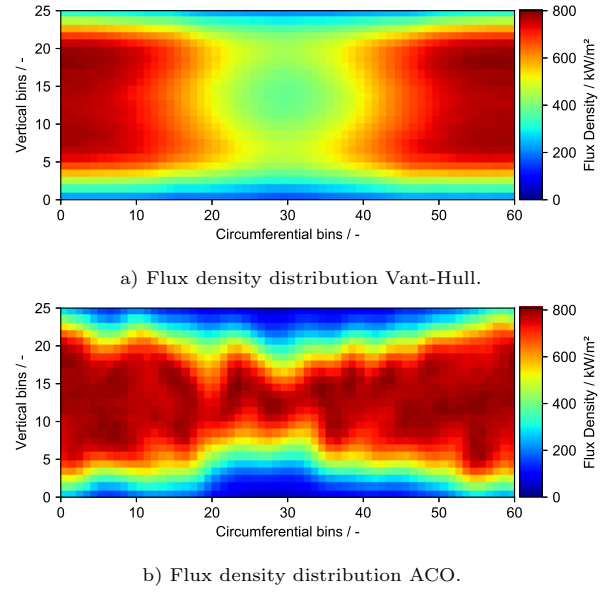


Figure 8: The flux density distribution of the results of the Vant-Hull strategy and the ACO are presented for a desired AFD of 800 kW m^{-2} .

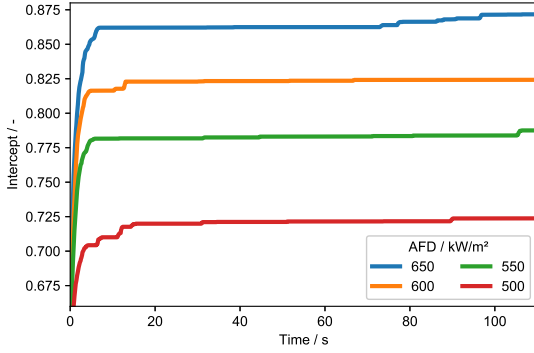
of maximal 470 kW m^{-2} are achieved. Consequently, the Vant-Hull strategy reaches lower intercepts than possible. The flux density distribution of the ACO has a better focus as illustrated in figure 8b. Here, 200 groups are used due to its strong convergence in the first seconds despite higher group numbers result in slightly better results according to section 3.3. As recognizable in the figure, the maximal flux density slightly exceeds the AFD in 7 of 1500 bins due to soft constraints in the optimizer. There, overshoots of the flux density are penalized by a factor and not strictly forbidden to support the convergence.

Figure 9 presents the convergence of the ACO for 200 groups. There, the intercept is plotted over time for AFDs between 500 kW m^{-2} to 1100 kW m^{-2} .

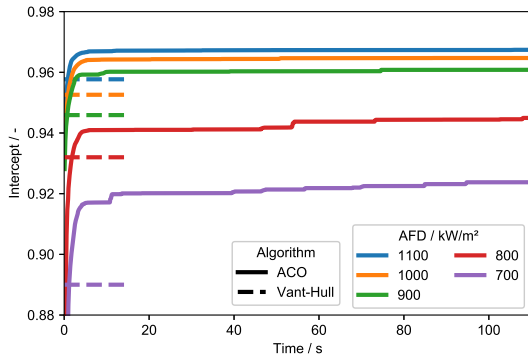
As initial solution a combination of group-aim-point-combinations found by the grouped Flesch-Vant-Hull strategy is applied. This solution is found within a second as the pre-calculated flux maps are used. The grouped Flesch-Vant-Hull strategy does not find valid solutions with AFDs below 880 kW m^{-2} as shown in figure 7b. However, even for stricter requirements regarding the AFD, this combination of group-aim-point-combinations is still a good first guess and the ACO is able to converge from there.

For the reference power plant, the ACO finds solutions starting with an AFD of 500 kW m^{-2} when using the result of the grouped Flesch-Vant-Hull strategy as initial solution. If only the result of the grouped Vant-Hull strategy is used as initial solution, results are found starting with an AFD of 650 kW m^{-2} .

Moreover, the best solutions of the original Vant-Hull strategy for the reference power plant are outlined by dashed horizontal lines in figure 9b. As for 700 kW m^{-2}



a) Low AFD.



b) High AFD.

Figure 9: The convergence of the ACO is shown for various AFDs of the reference power plant. Additionally, the results of the grouped Vant-Hull strategy or rather of the grouped Flesch-Vant-Hull strategy are marked by horizontal lines.

no solution is achieved by the original Vant-Hull strategy, the best result of the Flesch-Vant-Hull strategy is visualized. For lower AFDs, none of these strategies is able to solve the problem. Hence, only the results of the ACO are presented in figure 9a. The Flesch-Vant-Hull results are computed by raytracing for k and ϵ between 0 and 6 with a step size of 0.1.

In general, the results of the Vant-Hull strategy are presented as reference to evaluate the quality of the ACO. For all tested AFDs, the ACO outperforms the Vant-Hull strategy after 1 s to 2 s. The intercept comparison for AFDs between 800 kW m^{-2} to 1100 kW m^{-2} shows improvements of 1 % to 2 % after 60 s of optimization duration. The corresponding spillage reduction amounts to 19 % to 28 %. For 700 kW m^{-2} even an intercept improvement of 3.6 % is accomplished with respect to the Flesch-Vant-Hull strategy. Here, a spillage reduction of 31 % is achieved.

3.6. Cloud scenario

Finally, the applicability of the ACO in dynamic environments with uneven DNI is analyzed. Therefore, the heliostat field is virtually shaded by a DNI map delivered

by the ASI based nowcasting system described by Nouri et al. (2020). Here, a prediction with one minute lead time of the 19th of September 2015 at 13:00 h is used as an example. The shaded field is pictured in figure 10. The DNI of the heliostats in the southern part of the field is significantly reduced, due to a predicted cloud in this area. Here, the maximal DNI is only 882 W m^{-2} instead of 1000 W m^{-2} used previously. Likewise, the AFD is reduced to 700 kW m^{-2} for the following evaluation.

For the partially shaded heliostat field, Vant-Hull strat-

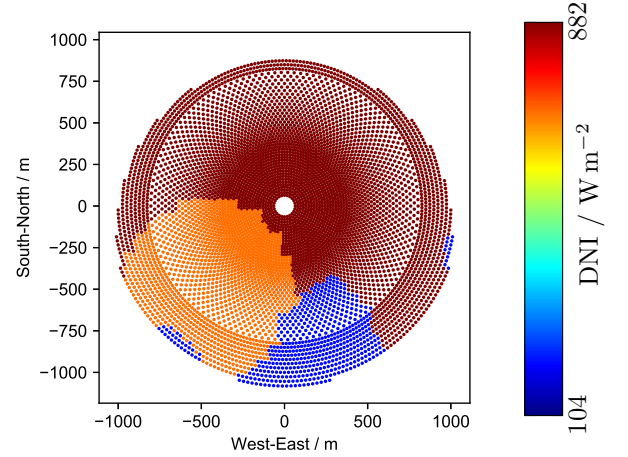


Figure 10: A real cloud passage predicted by the nowcasting system described by Nouri et al. (2020) is mapped onto the heliostat field of the reference power plant.

egy and ACO lead to flux density distributions that are presented in figure 10. The circumferential bin 0 accounts for the south side of the receiver and the vertical bin 0 for the lower receiver edge.

For this uneven distribution, the factor k of the Vant-Hull strategy is limited by the flux irradiating from the southwest part of the field. There, 700 kW m^{-2} are reached as shown in figure 11a around the circumferential bins ten to 15. All other parts of the receiver do not reach this AFD. Thus, the intercept is lower than possible and consequently, more energy than necessary is lost due to spillage. The ACO results in a flux density distribution visualized in figure 11b. Here, the flux density distribution is a bit bulky around the circumferential bins eight to twelve. However, in all parts of the receiver still the 700 kW m^{-2} and a good focus around the receiver's equator are achieved. This results in an intercept improvement of 1.9 % compared to the Vant-Hull solution. The intercept of the ACO solution is 0.967, whereas the result of the Vant-Hull strategy reaches an intercept of 0.949. Consequently, the spillage is reduced by 35 %.

4. Discussion

The main aim of this study was the reduction of the optimization time required by the ACO, while preserving

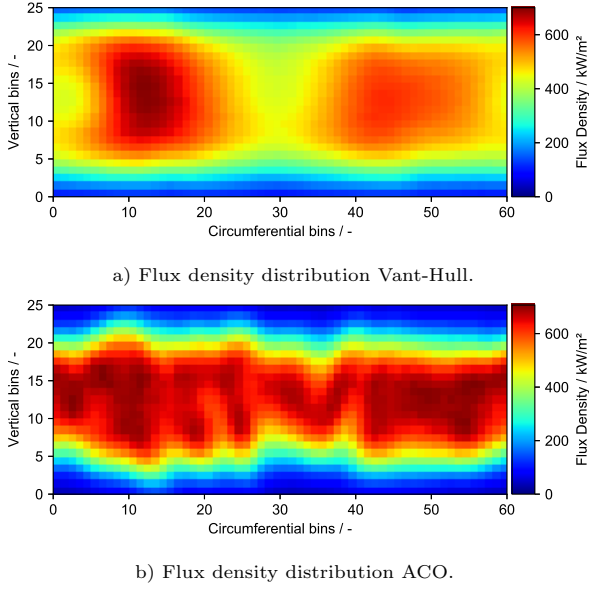


Figure 11: The flux density distribution of the results of the Vant-Hull strategy and the ACO are pictured for a desired AFD of 700 kW m^{-2} . The results are based on the DNI information illustrated in figure 10.

its ability to provide solutions close to the theoretical optimum.

The results show that the addressed drawback of the ACO, the long optimization duration, is resolved. Even for a large plant, like the reference power plant, good results are achieved after several seconds. This is accomplished by evaluating up to 112 billion flux values in a second on a GPU. According to Moore's law the computational power of GPUs will increase further in the following years (Sun et al., 2019). Moreover, multiple GPUs can be used in parallel, if a further reduction of the computational time is necessary e.g. for the application in solar multi-tower power plants. Simultaneously, the intercept improvement of 2% compared to the Vant-Hull strategy as stated by Flesch et al. (2017) for the same plant is preserved and confirmed in this work.

Moreover, the ACO becomes more robust by starting from an initial solution based on the grouped Vant-Hull strategy. Compared to the classical Vant-Hull (2002), the ACO finds solutions for AFDs that are 270 kW m^{-2} lower. Therefore, the grouped Flesch-Vant-Hull strategy is necessary as it becomes more important, if the group number decreases. This is illustrated in figure 7.

In the ungrouped case, the Flesch-Vant-Hull strategy results in similar intercepts then the original Vant-Hull strategy. However, if the AFD restrictions become more challenging, the Flesch-Vant-Hull strategy finds solutions, even if the original Vant-Hull strategy does not. For the reference power plant, the Flesch-Vant-Hull strategy solves the problem for up to 15% lower AFDs without performance losses.

The flux density distributions found by the ACO are

slightly spotted as pictured in figures 8b and 11b. The irregularities are caused by the limited number of groups and aim points. Furthermore, the mirror and tracking error in the simulation are set to 2 mrad in total. Increasing either these errors or the number of groups or aim points would lead to a smoother distribution also in the solution of the ACO. Moreover, the overshoots over the AFD caused by soft constraints in the optimizer do not enhance the risk for receiver damage as it only occurs in single bins and other errors between simulation and reality are presumably larger.

In addition, an optimal group number exists that achieves the best convergence depending on the optimization time, the plant size and the strictness of the constraints. If for instance highly fluctuating weather conditions occur, slightly reduced group numbers are favorable to reach faster convergence without forfeiting too much efficiency. Also, the ACO is able to handle uneven DNI distributions since the spillage reduction in comparison to the parameter-based technique rises clearly compared to the even DNI distribution.

In general, the ACO solves the problem for any distribution of allowable flux density and is usable for receivers of any form. Hence, all kinds of thermal receiver models or even chemical reactor models for e.g. the production of hydrogen could be connected. Besides, the k-means clustering technique that is applicable to any kind of heliostat field proved its applicability in this study.

5. Conclusion

In this paper, the ACO is adopted by a grouping strategy and supported by starting from a good initial solution found by a modified Vant-Hull strategy, which was already used by Flesch et al. (2017). Moreover, the ACO is implemented in cuda. Thus, the optimization can be performed on a GPU. As a result, the optimization duration is significantly reduced. Within a few seconds up to a minute, intercepts exceeding clearly the intercept of the Vant-Hull strategy are provided. At least, 1% to 2% greater intercepts or rather 19% to 28% less spillage are achieved for the reference power plant for all tested AFDs. Moreover, solutions are achieved even when the Vant-Hull strategy is not applicable anymore due to too strict constraints regarding the AFD.

In future, a control approach will be developed that is able to handle further measurement data like temperature and flux density distributions on the receivers surface. Using this feedback, the errors between simulation and reality should be reduced, which cannot be modeled by raytracing and consequently, cannot be represented within the ACO. In addition, the time gap during optimization should be filled by this control approach as the optimization still requires several seconds up to a minute for greater plants. Furthermore, an aim point management system should be developed that is able to evaluate nowcasting data in order to propose a certain control strategy. Finally, the system

that will include the ACO and the control approach should be tested in the solar tower power plant in Jülich.

Acknowledgement

This work was carried out in the project "HeliBo" (Grant number: PRO 0070 A) with the financial support from the Ministry of Economic Affairs, Innovation, Digitalization and Energy of the State of North Rhine-Westphalia.

References

- A. Sánchez-González, M. R. Rodríguez-Sánchez, D. Santana, Allowable solar flux densities for molten-salt receivers: In: put to the aiming strategy, Results in Engineering (2020) 100074. URL: <http://www.sciencedirect.com/science/article/pii/S259012301930074X>. doi:<https://doi.org/10.1016/j.rineng.2019.100074>.
- D. Maldonado, R. Flesch, P. Schwarzbözl, Hybridization of aim-point optimization methods for solar tower power plants, MATHMOD 2018 Extended Abstract Volume, ARGESIM Report 55 (2018). URL: <https://www.argesim.org/mathmod-vienna/mathmod2018-extended-abstracts/hybridization-of-aim-point-optimization-methods-for-solar-tower-power-plants>. doi:[10.11128/arep.55.a55230](https://doi.org/10.11128/arep.55.a55230).
- Y. Chu, C. F. Coimbra, Short-term probabilistic forecasts for direct normal irradiance, Renewable Energy (2017) 526 – 536. URL: <http://www.sciencedirect.com/science/article/pii/S0960148116308011>. doi:<https://doi.org/10.1016/j.renene.2016.09.012>.
- J. García, Y. C. Soo Too, R. Padilla, A. Beath, j.-s. Kim, M. Sanjuan, Dynamic performance of an aiming control methodology for solar central receivers due to cloud disturbances, Renewable Energy (2018) 355–367. doi:[10.1016/j.renene.2018.01.019](https://doi.org/10.1016/j.renene.2018.01.019).
- L. L. Vant-Hull, The Role of "Allowable Flux Density" in the Design and Operation of Molten-Salt Solar Central Receivers, Journal of Solar Energy Engineering (2002) 165–169. URL: <https://doi.org/10.1115/1.1464124>. doi:[10.1115/1.1464124](https://doi.org/10.1115/1.1464124).
- F. Lipps, L. Vant-Hull, A cellwise method for the optimization of large central receiver systems, Solar Energy (1978) 505 – 516. URL: <http://www.sciencedirect.com/science/article/pii/0038092X78900671>. doi:[https://doi.org/10.1016/0038-092X\(78\)90067-1](https://doi.org/10.1016/0038-092X(78)90067-1).
- F. J. Collado, J. Guallar, A two-parameter aiming strategy to reduce and flatten the flux map in solar power tower plants, Solar Energy (2019) 185 – 189. URL: <http://www.sciencedirect.com/science/article/pii/S0038092X19305663>. doi:<https://doi.org/10.1016/j.solener.2019.06.001>.
- M. Astolfi, M. Binotti, S. Mazzola, L. Zanellato, G. Manzolini, Helioast aiming point optimization for external tower receiver, Solar Energy (2017) 1114 – 1129. URL: <http://www.sciencedirect.com/science/article/pii/S0038092X16002292>. doi:<https://doi.org/10.1016/j.solener.2016.03.042>.
- A. Sánchez-González, M. R. Rodríguez-Sánchez, D. Santana, Aiming factor to flatten the flux distribution on cylindrical receivers, Energy (2018) 113 – 125. URL: <http://www.sciencedirect.com/science/article/pii/S0360544218305929>. doi:<https://doi.org/10.1016/j.energy.2018.04.002>.
- A. Sánchez-González, M. R. Rodríguez-Sánchez, D. Santana, Aiming strategy model based on allowable flux densities for molten salt central receivers, Solar Energy (2017) 1130 – 1144. URL: <http://www.sciencedirect.com/science/article/pii/S0038092X16001468>. doi:<https://doi.org/10.1016/j.solener.2015.12.055>.
- J. García, Y. C. Soo Too, R. Vasquez Padilla, R. Barraza Vicencio, A. Beath, M. Sanjuan, Heat flux distribution over a solar central receiver using an aiming strategy based on a conventional closed control loop, in: ASME 2017 11th International Conference on Energy Sustainability, Energy Sustainability, 2017. URL: <https://doi.org/10.1115/ES2017-3615>. doi:[10.1115/ES2017-3615](https://doi.org/10.1115/ES2017-3615).
- Y. C. Soo Too, J. García, R. V. Padilla, J.-S. Kim, M. Sanjuan, A transient optical-thermal model with dynamic matrix controller for solar central receivers, Applied Thermal Engineering (2019) 686 – 698. URL: <http://www.sciencedirect.com/science/article/pii/S1359431118370327>. doi:<https://doi.org/10.1016/j.applthermaleng.2019.03.086>.
- T. Ashley, E. Carrizosa, E. Fernández-Cara, Optimisation of aiming strategies in solar power tower plants, Energy (2017) 285 – 291. URL: <http://www.sciencedirect.com/science/article/pii/S0360544217311581>. doi:<https://doi.org/10.1016/j.energy.2017.06.163>.
- P. Richter, F. Kepp, C. Büsing, S. Kuhnke, Optimization of robust aiming strategies in solar tower power plants, AIP Conference Proceedings (2019) 030045. URL: <https://aip.scitation.org/doi/abs/10.1063/1.5117557>. doi:[10.1063/1.5117557](https://doi.org/10.1063/1.5117557). arXiv:<https://aip.scitation.org/doi/pdf/10.1063/1.5117557>.
- S. Kuhnke, P. Richter, F. Kepp, J. Cumpston, A. M. Koster, C. Büsing, Robust optimal aiming strategies in central receiver systems, Renewable Energy (2020) 198 – 207. URL: <http://www.sciencedirect.com/science/article/pii/S0960148119318129>. doi:<https://doi.org/10.1016/j.renene.2019.11.118>.
- A. Salomé, F. Chhel, G. Flamant, A. Ferrière, F. Thiery, Control of the flux distribution on a solar tower receiver using an optimized aiming point strategy: Application to THEMIS solar tower, Solar Energy (2013) 352 – 366. URL: <http://www.sciencedirect.com/science/article/pii/S0038092X1300090X>. doi:<https://doi.org/10.1016/j.solener.2013.02.025>.
- S. M. Besarati, D. Y. Goswami, E. K. Stefanakos, Optimal heliostat aiming strategy for uniform distribution of heat flux on the receiver of a solar power tower plant, Energy Conversion and Management (2014) 234 – 243. URL: <http://www.sciencedirect.com/science/article/pii/S0196890414003343>. doi:<https://doi.org/10.1016/j.enconman.2014.04.030>.
- Q. Yu, Z. Wang, E. Xu, Analysis and improvement of solar flux distribution inside a cavity receiver based on multi-focal points of heliostat field, Applied Energy (2014) 417 – 430. URL: <http://www.sciencedirect.com/science/article/pii/S0306261914009556>. doi:<https://doi.org/10.1016/j.apenergy.2014.09.008>.
- K. Wang, Y.-L. He, X.-D. Xue, B.-C. Du, Multi-objective optimization of the aiming strategy for the solar power tower with a cavity receiver by using the non-dominated sorting genetic algorithm, Applied Energy (2017) 399 – 416. URL: <http://www.sciencedirect.com/science/article/pii/S0306261917309698>. doi:<https://doi.org/10.1016/j.apenergy.2017.07.096>.
- B. Belhomme, R. Pitz-Paal, P. Schwarzbözl, Optimization of Heliostat Aim Point Selection for Central Receiver Systems Based on the Ant Colony Optimization Metaheuristic, Journal of Solar Energy Engineering (2013) 136 (2013). URL: <https://doi.org/10.1115/1.4024738>. doi:[10.1115/1.4024738](https://doi.org/10.1115/1.4024738).
- R. Flesch, C. Frantz, D. M. Quinto, P. Schwarzbözl, Towards an optimal aiming for molten salt power towers, Solar Energy (2017) 1273 – 1281. URL: <http://www.sciencedirect.com/science/article/pii/S0038092X17306527>. doi:<https://doi.org/10.1016/j.solener.2017.07.067>.
- D. Maldonado, R. Flesch, A. Reinholz, P. Schwarzbözl, Evaluation of aim point optimization methods, AIP Conference Proceedings (2018) 040025. URL: <https://aip.scitation.org/doi/abs/10.1063/1.5067061>. doi:[10.1063/1.5067061](https://doi.org/10.1063/1.5067061).
- M. Izzygon, P. Armstrong, C. Nilsson, N. Vu, TieSOL—a GPU-based suite of software for central receiver solar power plants, Proceed-

ings of SolarPACES (2011).

Y. Zhou, Y. Zhao, Heliostat Field Layout Design for Solar Tower Power Plant Based on GPU, IFAC Proceedings Volumes 47 (2014) 4953 – 4958. URL: <http://www.sciencedirect.com/science/article/pii/S1474667016423827>. doi:<https://doi.org/10.3182/20140824-6-ZA-1003.01581>, 19th IFAC World Congress.

D. Gebreiter, G. Weinrebe, M. Wöhrbach, F. Arbes, F. Gross, W. Landman, sbpRAY – A fast and versatile tool for the simulation of large scale CSP plants, AIP Conference Proceedings 2126 (2019) 170004. URL: <https://aip.scitation.org/doi/abs/10.1063/1.5117674>. doi:[10.1063/1.5117674](https://doi.org/10.1063/1.5117674). arXiv:<https://aip.scitation.org/doi/pdf/10.1063/1.5117674>.

C. He, J. Feng, Y. Zhao, Fast flux density distribution simulation of central receiver system on GPU, Solar Energy 144 (2017) 424 – 435. URL: <http://www.sciencedirect.com/science/article/pii/S0038092X17300440>. doi:<https://doi.org/10.1016/j.solener.2017.01.025>.

C. He, Y. Zhao, J. Feng, An improved flux density distribution model for a flat heliostat (iHFLCAL) compared with HFLCAL, Energy 189 (2019) 116239. URL: <http://www.sciencedirect.com/science/article/pii/S0360544219319346>. doi:<https://doi.org/10.1016/j.energy.2019.116239>.

R. Grena, An algorithm for the computation of the solar position, Solar Energy 82 (2008) 462 – 470. URL: <http://www.sciencedirect.com/science/article/pii/S0038092X07002162>. doi:<https://doi.org/10.1016/j.solener.2007.10.001>.

B. Belhomme, R. Pitz-Paal, P. Schwarzbözl, S. Ulmer, A New Fast Ray Tracing Tool for High-Precision Simulation of Heliostat Fields, Journal of Solar Energy Engineering 131 (2009). URL: <https://doi.org/10.1115/1.3139139>. doi:[10.1115/1.3139139](https://doi.org/10.1115/1.3139139).

S. Lloyd, Least squares quantization in PCM, IEEE Transactions on Information Theory 28 (1982) 129–137. doi:[10.1109/TIT.1982.1056489](https://doi.org/10.1109/TIT.1982.1056489).

NVIDIA Corporation, NVIDIA CUDA C++ Programming Guide, 2019. URL: <https://docs.nvidia.com/cuda/pdf/CUDA.C.Programming.Guide.pdf>.

J. Liu, Efficient Synchronization for GPGPU, Ph.D. thesis, University of Pittsburgh, 2018. URL: <http://d-scholarship.pitt.edu/34943/>.

V. Prisacariu, I. Reid, et al., fastHOG-a real-time GPU implementation of HOG, Department of Engineering Science 2310 (2009). URL: [fast{HOG}-areal-time{GPU}implementationof{HOG}](https://arxiv.org/abs/0905.3215).

B. Nouri, P. Kuhn, S. Wilbert, C. Prah, R. Pitz-Paal, P. Blanc, T. Schmidt, Z. Yasser, L. R. Santigosa, D. Heineman, Nowcasting of DNI maps for the solar field based on voxel carving and individual 3D cloud objects from all sky images, AIP Conference Proceedings 2033 (2018) 190011. doi:[10.1063/1.5067196](https://doi.org/10.1063/1.5067196).

B. Nouri, S. Wilbert, P. Kuhn, N. Hanrieder, M. Schroedter-Homscheidt, A. Kazantzidis, L. Zarzalejo, P. Blanc, S. Kumar, N. Goswami, R. Shankar, R. Affolter, R. Pitz-Paal, Real-time uncertainty specification of all sky imager derived irradiance nowcasts, Remote Sensing 11 (2019) 1059. doi:[10.3390/rs11091059](https://doi.org/10.3390/rs11091059).

B. Nouri, K. Nouredin, T. Schlichting, S. Wilbert, T. Hirsch, M. Schroedter-Homscheidt, P. Kuhn, A. Kazantzidis, L. Zarzalejo, P. Blanc, Z. Yasser, J. Fernández, R. Pitz-Paal, Optimization of parabolic trough power plant operations in variable irradiance conditions using all sky imagers, Solar Energy 198 (2020) 434–453. doi:[10.1016/j.solener.2020.01.045](https://doi.org/10.1016/j.solener.2020.01.045).

M. Schroedter-Homscheidt, M. Kosmale, S. Jung, J. Kleissl, Classifying ground-measured 1 minute temporal variability within hourly intervals for direct normal irradiances, Meteorologische Zeitschrift 27 (2018) 161–179. URL: <http://dx.doi.org/10.1127/metz/2018/0875>. doi:[10.1127/metz/2018/0875](https://doi.org/10.1127/metz/2018/0875).

Y. Sun, N. B. Agostini, S. Dong, D. Kaeli, Summarizing CPU and GPU Design Trends with Product Data, 2019. arXiv:[1911.11313](https://arxiv.org/abs/1911.11313).