Evaluation of aim point optimization methods

Daniel Maldonado, Robert Flesch, Andreas Reinholz, and Peter Schwarzbözl

Citation: AIP Conference Proceedings **2033**, 040025 (2018); doi: 10.1063/1.5067061
View online: https://doi.org/10.1063/1.5067061
View Table of Contents: http://aip.scitation.org/toc/apc/2033/1
Published by the American Institute of Physics

Evaluation of Aim Point Optimization Methods

Daniel Maldonado[1, a)], Robert Flesch[2], Andreas Reinholz[3] and Peter Schwarzbözl[4]

[1]*Dr.-Ing.-Researcher, Institute of Solar Research, German Aerospace Center (DLR), Linder Hoehe, 51147 Cologne, Germany*
[2]*Dr.-Ing.-Researcher, Institute of Solar Research, German Aerospace Center (DLR), Prof. Rehm-Strasse 1, 52428 Juelich, Germany*
[3]*Researcher, Institute of Solar Research, German Aerospace Center (DLR), Linder Hoehe, 51147 Cologne, Germany*
[4]*Dipl.-Ing.-Group Leader, Institute of Solar Research, German Aerospace Center (DLR), Linder Hoehe, 51147 Cologne, Germany*

[a)]Corresponding author: Daniel.MaldonadoQuinto@dlr.de

Abstract. In the present work we applied the universal methodology of Belhomme et. al [3] for optimizing the aim point assignment of a solar power tower with a cylindrical molten salt receiver. The methodology uses a combinatorial formulation of the problem and solves it with the ant colony optimization metaheuristic (ACO). Additionally, the local search algorithm (LS) has been implemented for optimizing the aim point assignment. A realistic simulation scenario, including a numerical receiver model of a receiver with molten salt as heat transfer fluid, is used to analyze both optimization strategies and to compare their performances regarding the thermal output of the receiver.

INTRODUCTION

In power tower plants each heliostat has the possibility to aim on an arbitrary point of the receiver surface. For safety reasons the heliostat can also align on an off receiver point. A good aiming strategy is of great importance for an efficient operation. Most of the radiation will very likely hit the receiver if all heliostats aim to its center. Unfortunately, this aiming strategy cannot be applied as temperature and/or stress limits of the receiver, which can be expressed in a limit for the flux density, will be exceeded. Therefore, the aim points of the heliostats have to be distributed over the receiver surface in order to lower the peak flux densities. But this will presumably reduce the amount of radiation hitting the receiver. For this reason the aiming strategy can be characterized as a constraint optimization problem, in which the optical or thermal performance is optimized with respect to all limits for the receiver. Strategies for finding a suitable aim point assignment were presented by Vant-Hull [1] and García-Martín et al. [2]. However, these procedures do not optimize the receiver power. The optimal assignment of each heliostat to one of a fixed set of available aim points is a NP-hard discrete optimization problem of high dimension. The size of the solution space is equal to the number of fixed aim points to the power of the number of heliostat. Belhomme et al. [3] presented a universal procedure for the aim point optimization, giving a combinatorial formulation of the problem and solving it with the ant colony optimization metaheuristic (ACO). In addition to the ACO, the local search algorithm was implemented in this work in order to perform a comparison of both algorithms.

OPTMIZATION METHODS

A restriction-free choice of the aim points would define a continuous optimization problem. In this work we distribute a fixed number of aim points over the receiver surface and allow also off receiver alignment. Allowing only a fixed number of aim points with fixed positions formulates the optimization problem to a discrete one. Due to the NP-hard nature of the problem only heuristic methods can lead to a solution within a reasonable computation

SolarPACES 2017
AIP Conf. Proc. 2033, 040025-1–040025-8; https://doi.org/10.1063/1.5067061
Published by AIP Publishing. 978-0-7354-1757-1/$30.00

time. First we give a short review about the applied methodology of Belhomme et. al [3] and subsequently we describe the new implemented strategy using a local search algorithm. Both optimization algorithms are implemented in the raytracing environment STRAL [4].

Ant Colony Optimization Metaheuristic

The ACO method is a probabilistic technique that benefits from the principles of swarm intelligence and imitates the behavior of ant colonies during foraging. A suboptimal aim point configuration at the beginning is improved according heuristic information. The aiming optimization is transferred to a suitable problem for the ACO by defining the aim point configuration as the trail of an ant. One trail represents one possible heliostat - aim point assignment. Figure 1 shows a matrix representation of the aim point alignment of an exemplary scenario consisting of 4 heliostats and six aim points. The blue line in Fig. 1 is a representative trail where the ant selected the waypoints heliostat 1 / aim point 2, heliostat 2 / aim point 3, heliostat 3 / aim point 5 and at last the waypoint heliostat 4 / aim point 1. The superposition of the computed flux maps of each heliostat - aim point assignment (waypoints of the ant) results in an irradiation map for the receiver.

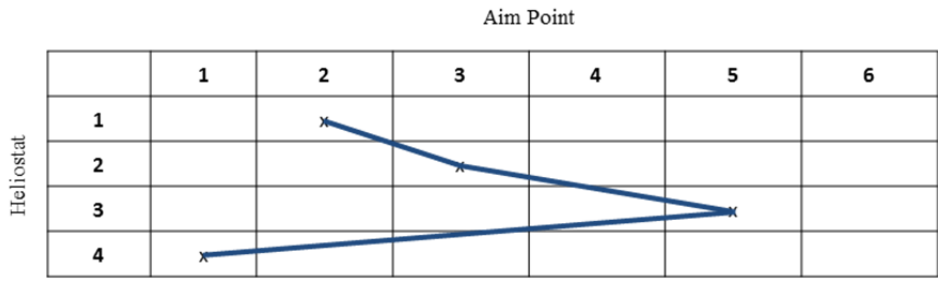

FIGURE 1. Matrix representation of the aim point alignment of an exemplary scenario with 4 heliostats and 6 aim points.

The resulting receiver performance corresponds to the "length" of the trail and represents the global quality value of the optimization problem. The intercept factor of one heliostat corresponds to the myopic information of a single ant while discovering the trail. The myopic information is called the local quality value. Additionally, a global information, the pheromone value, gives an information on the performance of a trail using that specific combination. The trail is divided in edges and the choice of the path for each ant is firstly based on a so called transition rule. This rule determines whether the edge with a high value for a pheromone concentration in combination of the local quality value, $max(\tau_{ij} \cdot \eta_{ij}^{\beta})$, is more likely to be chosen or a random edge based on their probability value p_{ij}. The probability p_{ij} is calculated with a Monte-Carlo-Method for several virtual ants. As described by Belhomme et. al [3], Equation 1 states the probability of an ant to follow the edge from i to j, with τ as the local pheromone concentration, η is the myopic attraction (in this study the intercept factor of the specific heliostat aim point combination), and β is the parameter that rules the relevance of η with respect to τ.

$$p_{ij} = \frac{\tau_{ij} \cdot \eta_{ij}^{\beta}}{\sum_{k \in N} \tau_{ik} \cdot \eta_{ik}^{\beta}} \tag{1}$$

The transition rule determines if the algorithm has a local greedy character. Otherwise, if the way is mainly chosen based on p_{ij}, the algorithm has a random character.

After one trail is chosen a numerical model calculates the receiver performance for the resulting flux distribution of this trail as the global quality value. If constraints like flux density or temperature limits are not fulfilled the global quality value is penalized. This value is used to update the pheromone concentration on all single edges. The amount of pheromone, and therefore the attractiveness of one edge, depends on the ratio between global quality value and the estimated maximum of this value. An essential characteristic of the algorithm is that it can proof constraints only at the end of the trail and an exceedance penalizes the entire track.

Local Search

Due to inherent characteristic constraints of solar receivers, e.g. flux limits, a second approach was implemented in STRAL, called local search (LS). In case the flux density limit is exceeded on a small area of the receiver the above described ACO would penalize the whole aim point configuration, even though this configuration is an optimal solution for the rest of the receiver. In those circumstances the ACO will impede itself finding an optimal solution. Therefore an algorithm with local view was implemented.

The LS starts from a solution and moves towards an improvement, equally to the ACO. In contrast, it only manipulates the aim point configuration in a local region. . The options for each heliostat changing to a new aim point are restricted. In Fig. 2 the matrix representation is shown for the exemplary scenario where the neighborhoods are marked in blue. Heliostat 1, currently aiming on aim point 2, is restricted to the options aim point 1 or aim point 3. Heliostat 2, currently aiming on aim point 3, is restricted to the option aim point 5 and so on.

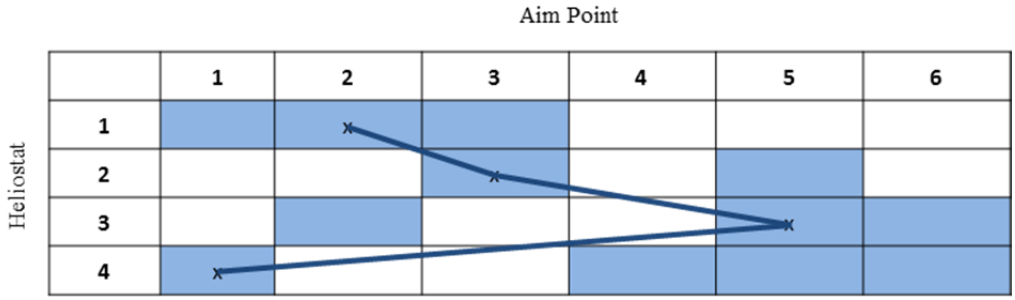

FIGURE 2. Matrix representation of the aim point alignment and marked neighborhood.

In a single LS run all heliostats' assignments are examined one by one in a certain sequence. In each examination step the assignment for a single heliostat can change by shifting to one of the neighboring aim points. In contrast to the ACO the overall receiver performance is calculated after each shift until an improved solution is found in the neighborhood.On the other hand, if no improvement is found, the initial configuration is maintained. In Fig. 3 an exemplary LS-run is demonstrated. In the first step the aim point of heliostat 1 is shifted from aim point 2 to aim point 3 and the thermal power of the receiver is calculated. Since the performance is lower than the original configuration the algorithm continues to search inside the neighborhood and the aim point changes from 3 to 1. Afterwards the thermal power is calculated once again. In this example the performance is higher than the original configuration. Therefore heliostat 1 now points to aim point 1 and the algorithm continues with heliostat 2. In the third step the aim point of heliostat 2 is realigned from aim point 3 to aim point 5, where again a better performance is determined. The procedure is executed for all heliostats. Compared to the ACO a local change of a single heliostat-aim point assignment is evaluated and local exceedances of allowable restriction can be prevented without affecting the entire aim point configuration.

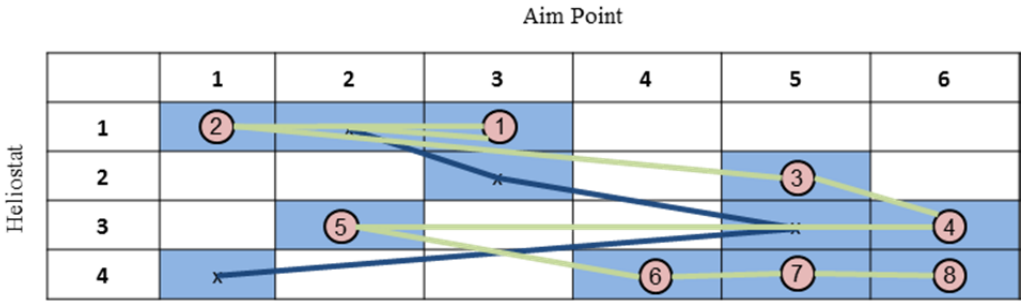

FIGURE 3. Walking through an exemplary LS-Loop.

An alternative mode is to examine the entire neighborhood and shift to the best configuration. We call this strategy the deep search (LS-DS). This means for the example in Fig. 3 all blue marked combinations will be determined.

NUMERICAL MODEL

Example Solar Tower Power Plant

The simulations in this paper are carried out for a virtual solar thermal power tower with an approximate thermal power of 450 MW. The cylindrical receiver with a circumscribing diameter of 12.8 m receiver consists of 12 panels arranged in a conventional way in two flow paths each with six panels and a crossover behind the third panel of each flow path. Each panel is built of 117 tubes, each with an outer diameter of 35 mm and a wall thickness of 2.5 mm. The tubes are radiated on a length of 20 m. The salt temperature rises inside the receiver from 290 °C to 565 °C.

The receiver is located on top of a 190 m high tower. The tower is surrounded by 6482 heliostats each with a reflective surface area of 121 m^2. All heliostats are assumed to be operational and used for optimization. This virtual plant is placed at the site of the Redstone thermal power plant (28° S and 23° E). More details of the virtual plant can be found in Flesch et al. [5].

Receiver Model

The API of STRAL delivers the possibility to implement a user defined receiver model which can be included as a dynamic link library (DLL). Here the model is implemented in C++. For a given mass flow it calculates local salt temperatures at multiple locations in the discretized tube. The user can define a number of tubes which represent all tubes in a panel. The chosen approach neglects longitudinal conduction and uses an effective radiation loss which replaces the radiation exchange between neighboring tubes. The backside of the tube is assumed to be adiabatic and conduction heat transfer from front to backside is neglected as well. These assumptions lead to a decoupled model in which the temperatures of each element can be calculated by solving a nonlinear equation. Fortunately, this nonlinear equation can be solved with an analytic approach. An additional equation is solved to obtain the fluid temperature at the tube crown as this temperature is limited to 600 °C. In the end all salt and tube temperature of the receiver can by calculated analytically for a given mass flow. Then, the mass flow is iterated until the desired outlet temperature is reached. One full evaluation of the thermal model takes approximately 0.1 s on one core. Again, a more detailed description including a validation of the model can be found in Flesch et al. [5].

Molten Salt Receiver Constraints

Receivers using molten salt are an interesting case for aim point optimization as their limitation of the flux density is quite complex. The limit for the flux density, often referred as allowable flux density (AFD), can be expressed as function of the mass flow and the local salt temperatures [1]. In case of the Solar Two Receiver the resulting AFD as function of the temperature and the design point referenced velocity is plotted in Fig. 4a. Applied to the receiver of the current study the resulting local distribution of AFD for on the cylinder surface at the design mass flow looks like Fig. 4b. In this configuration the cold salt enters centrally from the top and the hot salt leaves the receiver on top of the outer panels. Two interesting things can be observed: firstly, the AFD increases with temperature in the low temperature region close to the inlet. Secondly, the AFD drops dramatically for the hot panels close to the outlet. This drop at the outlet is caused by the second restriction arising from the limits for the film temperature. The transition between the two limits can easily been identified by the kink in Fig. 4a.

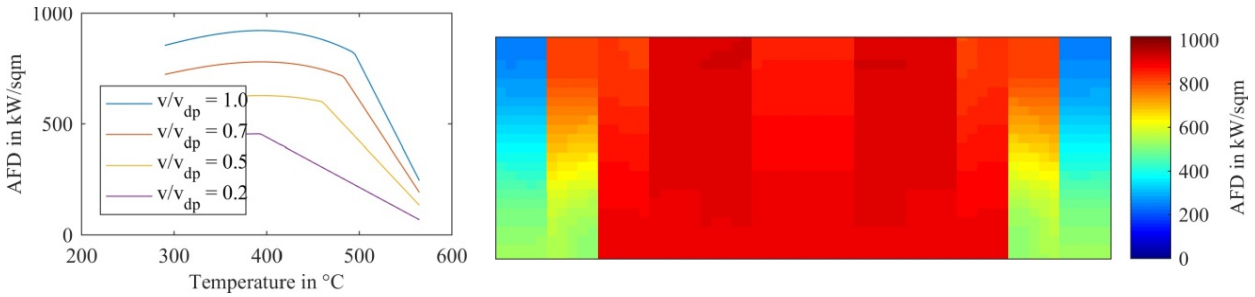

(a) (b)

FIGURE 4. (a) AFD as function of the temperature. (b) AFD applied to receiver of current study.

The thermal model of the receiver is used to calculate mass flow and local salt temperatures which are required to evaluate the AFD. The AFD is evaluated in this paper with the formulation given by Vant-Hull [1]

$$AFD = \left(842.27 - 1.5514\,\theta + 4.613\ \cdot 10^{-3}\,\theta^2 - 3.2073 \cdot 10^{-6}\,\theta^3\right) \cdot \left(0.3 + 0.7\left(\frac{v}{v_{dp}}\right)^{0.8}\right) \qquad (2)$$

as function of salt temperature in Fahrenheit θ and the ratio of the current salt velocity v and design point velocity v_{dp}. This expression gives the AFD resulting from stresses in the tube and describes the AFD for the inlet panels. In case of the Solar Two receiver the temperature restriction was reformulated as AFD expression as well as it was shown above. Because the model described above is capable of calculating the peak salt film temperature a limit of 600 °C is applied for the temperatures directly.

Simulation Case

An aim point grid with 36 points in circumferential direction and 30 points in vertical direction is used. Each heliostat is allowed to aim on all vertical aim points which are closest to the direct connection of heliostat center and axis of the cylinder. Additionally, all four neighboring stacks of vertical aim points in each circumferential direction can be chosen as well.

For the comparison of the two optimization methods four different cases are used. Table 1 gives and overview about these cases. All calculations are performed for the 21 of March. The design point (12:00) and one off-design point (7:00) are used. Each algorithm is started from two different initial aim point configuration: one of them is a random configuration; the second configuration is obtained using the parameter based strategy described by Flesch et al. [5].

TABLE 1. Case definitions

Initial Aim Point Configuration	Design-point (12:00)	Off-design-point (7:00)
Parameter based	Case 1	Case 2
Random	Case 3	Case 4

For each case both optimization methods are tested with different settings. In case of the local search the neighborhood and the way it is searched are changed. Firstly, the standard neighborhood (LS-SN) as shown in Fig. 5 is used. Secondly, a neighborhood with increased size (LS-IN) is analyzed. Thirdly, the standard neighborhood is used, but this time the deep search (LS-DS) is used. In total ten optimization runs are used for the each setting of the local search. The total number of evaluations of the thermal model depends on the actual configuration as the number of evaluation per heliostat depends on the size of the neighborhood and how it is searched.

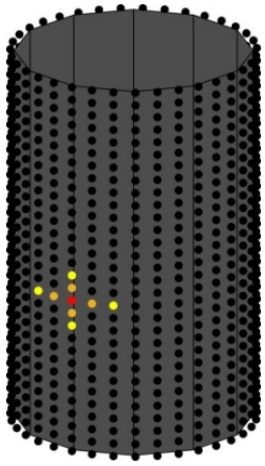

FIGURE 5. Aim point distribution on receiver surface and defined neighborhood for an exemplary aim point. The aim points belonging to the standard neighborhood of the red aim point are colored in orange (SN). The extended neighborhood contains the orange and yellow aim points (IN).

For the ant colony optimization two different settings were tested on each case. As described above the ACO-algorithm can be set greedy (ACO-GR) by focusing the ants on the myopic information rather than doing a selection randomly (ACO-RD). Firstly, for the ACO-GR settings 90% of the time the edge with $max(\tau_{ij} \cdot \eta_{ij}^{\beta})$ was chosen. Secondly, for the ACO-RD settings only 50% of the time the edge with $max(\tau_{ij} \cdot \eta_{ij}^{\beta})$ was chosen. The optimization procedure is stopped after 10^6 evaluations of the thermal model.

RESULTS AND DISCUSSION

For the comparison of the optimization methods the thermal output of the receiver during the optimization process is plottedover the number of thermal receiver evaluations. The comparison is given in Figure 6 for the four different cases defined in the previous section. For each case the results of the three different settings of the local search and of the two different settings of the ant colony algorithm are shown.The plots give an idea of the computational effort which is required to reach an aim point configuration with a certain output. The optimization progress of the LS is shown with single markers. Each marker indicates the thermal output of one completed run of the LS. In contrast, one run of the ACO requires only a few evaluations (~100) and therefore, the results of the ACO are shown with a dashed or solid line. Otherwise, the markers would be too close together to be distinguishable.

When comparing the results of the different settings of the local search one can easily observe that the LS using the deep search strategy performs best in all the cases. Additionally, it takes only a few more evaluations than the standard neighborhood without deep search. The additional evaluations are performed within the first iterations. At the end of the optimization the local search with the setting LS-SN analyzes almost the whole neighborhood as no improvement can be found for the first aim points in the neighborhood. Therefore, one step of LS-SN takes the same effort as one step of LS-DS. In contrast, the local search with extended neighborhood (LS-IN) is much more expensive as a larger neighborhood has to examined. But the optimization does not gain a significant benefit from this additional effort; in some cases this configuration performs even worse than the reference case (LS-SN).

The competition of the two different settings of the ant colony optimization has a clear winner as well. The “greedy”-setting (ACO-GR) performs much better than the more chaotic configuration (ACO-RD) in all of the analyzed cases. This result could have been expected as the usage of the myopic information gives the orientation to the ants.

A comparison of the two different optimization strategies is more complicated. Here, we will focus on the discussion of both algorithms in their best configuration. In case 1, shown in Fig, 6 (a), the local search algorithm clearly outperforms the ACO optimization algorithm. Only in the very beginning in the first $1 \cdot 10^4$ evaluations the ACO finds an aim point assignment with a higher thermal output. The LS quickly finds an improved assignment within the first 5 runs. After these 5 runs the LS improves the thermal output marginally. The ACO only finds an aim point assignment with approximately 0.2% less thermal power of the receiver. In case 2 the results are similar.

The LS quickly finds improved aim point assignments and starts saturating afterwards. But this time, when being granted more evaluations, the ACO finds a slightly better solution than the LS. It's possible that if the LS would run longer it would find a similar or even better aim point assignment. The situation is different when looking at the results of case 3 (Fig. 6 (c)). In this case the same random aim point assignment is used as initial assignment for all algorithms. At each number of evaluations the ACO gives an assignment with a better performance and the highest difference among the methods was estimated in this case. In the end after 10^6 evaluations the ACO finds an assignment with a similar performance as the initial solution used in case 1. The LS is stopped after its 10 runs with a thermal output of 444.8 MW which is approximately 2.3% less than the final value of the ACO. However, the LS has only been granted approximately $3 \cdot 10^5$ evaluations. But its progress has slowed down pretty much, so we do not expect that it would find a much better solution when being granted more runs. In case 4 (Fig. 6 (d)) both algorithm perform similarly, though the LS gives at the same number of evaluation an aim point assignment with slightly better performance. This is still the case when the LS is stopped after 10 runs. Since the LS still improves the thermal output of the receiver with a similar slope as the ACO, we expect it to perform similarly as the ACO if continued. Anyway, as it is stopped, the final solution of the ACO has a higher thermal output. The output of 147.3 MW is still little lower than the value obtained in case 2.

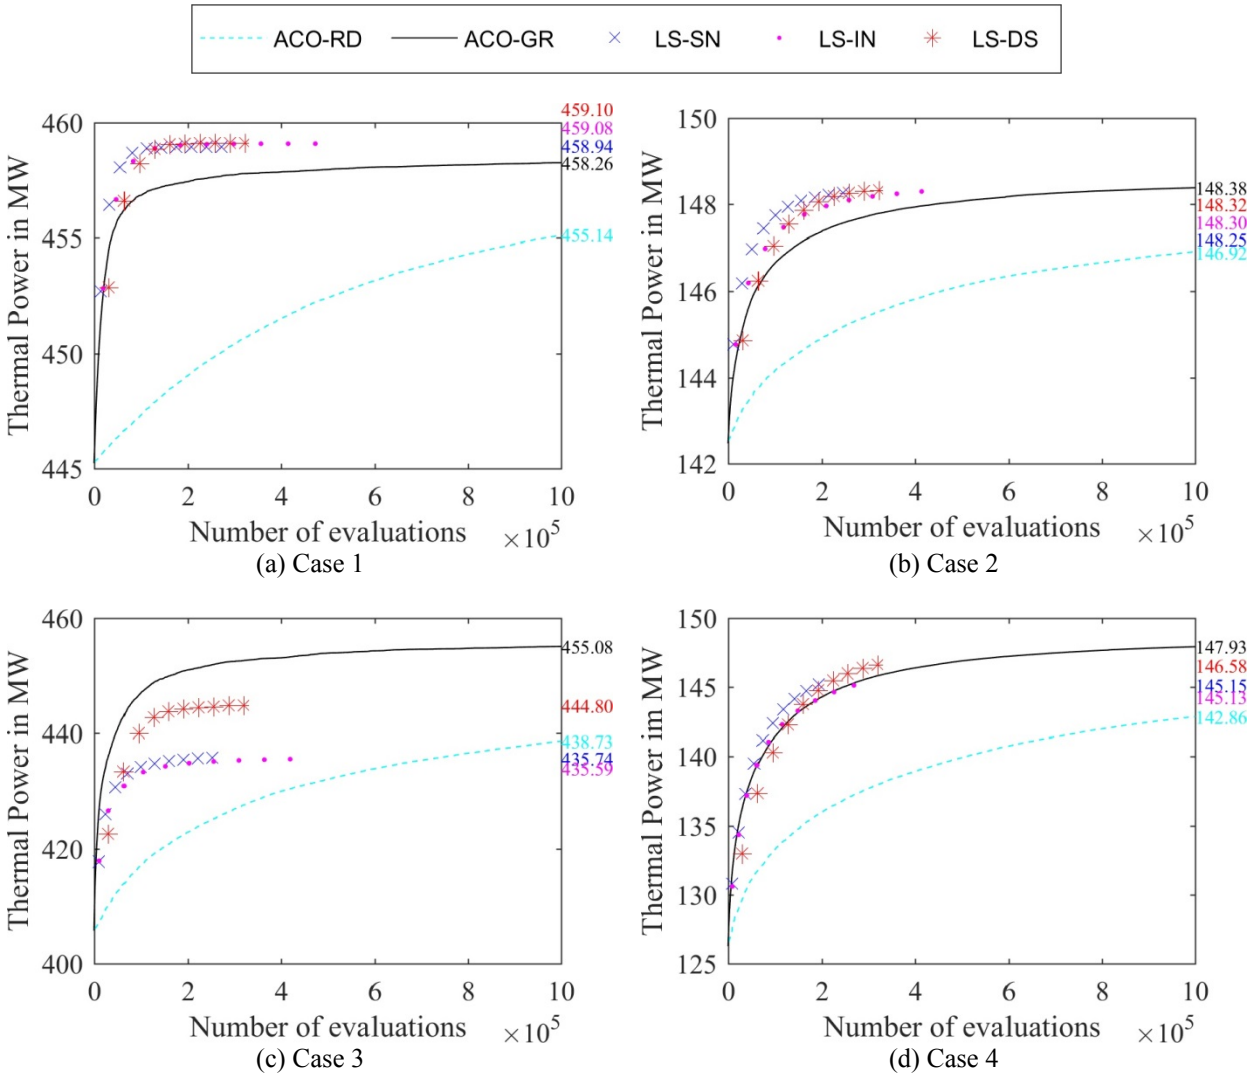

FIGURE 6. Results of the optimization: for each of the four cases the results of the two algorithms with its individual settings are compared. The numbers on the right side of each plot give the reached maximum of the thermal power during the optimization procedure of each configuration.

The main strength of the LS is that it is focused: if a maximum lies within the neighborhood of the initial solution, the LS quickly converges towards this maximum. This can be seen in the first two cases. In these cases the initial aim point assignment has a high thermal performance. The LS quickly finds an improved assignment within the neighborhood. But the focus on a direct neighborhood has also its drawbacks which emerge in case 3. If the initial solution is "far" away from the optimal assignment, convergence of the LS is very slow as it has to evaluate many combinations to get closer to the optimum. Additionally, it might get stuck in a local optimum if the neighborhood is too small to offer a better solution than in the local optimum. The ACO on the other hand has a global view and is not bound to a neighborhood. Therefore, it performs well in cases in which the initial solution is "far" away from the final solution. In the other cases it performs well, too, but it requires more evaluations as it is not as focused as the local search. Additionally, the closer the results get to the optimum, the closer the flux density will be to the limit in some regions. Therefore, many solutions chosen during the optimization process will be penalized as explained previously. The case 4 is of special interest as one could have expected that the ACO performs better, because a random initial solution was used. However, the LS even performs slightly better. A possible explanation is again the constraints of the molten receiver mentioned above. The present case at 7:00 in the morning is tightly constrained, because the mass flow rate is very low. Therefore, the limit for the film temperature is reached very quickly. In such a constrained case the local view of the LS is advantageous as it applies only changes to regions where the limits have not already been reached.

In the preceding discussion the thermal output as function of number of evaluations has been used for the comparison. We would like to point out, that a comparison of the time required for the optimization would be interesting as well, but such a comparison is hard to realize. The ant colony optimization is parallelized whereas a parallelization of the LS is hard to realize as the starting point for the aim point changes of one heliostat depend on the changes applied after evaluating the preceding heliostats (this is obviously an advantage of the ACO). Due to the parallelization the ACO has the potential to run faster. We obtained an acceleration-factor of 32 due to the parallelization on a medium-sized workstation with 32 cores and 64 Threads. About 200 evaluations took 1 second of computation time.

CONCLUSION AND OUTLOOK

Two different optimization algorithms were used to optimize the aim point configuration of a solar power tower with cylindrical molten salt receiver. Both algorithms were used with different parameter settings and their performance was compared. The LS performs well if the initial solution is within they analyzed neighborhood or somehow close or if the case is highly constrained. The ACO in contrast performs fairly well for all analyzed cases and better than the LS if the initial solution is "far away" from the global optimum and not heavily constrained. Additionally, the algorithm guarantees to find the global optimum and does not get stuck in a local optimum. Hence, a combination of both algorithms seems to be a good approach: The ACO is used to get close to the optimum from any initial solution. The final solution is then improved by a very few iterations of the LS. Such a hybrid algorithm should be implemented including some automatic switch if the convergence of the ACO slows down.
In real operation there are several factors affecting the aiming which were not modelled here, e.g. wind. It should be proved whether the optimal solutions found here are also optimal under real operation conditions.

ACKNOWLEDGMENTS

The work was carried out in the project "DynaSalt" with the financial support of the Federal Ministry of Economic Affairs and Energy of the Federal Republic of Germany under the contract 0325756.

REFERENCES

1. L. L. Vant-Hull, Journal of Solar Energy Engineering **124**, 165–169 (2002).
2. F. J. García-Martín et al., Solar Energy **66**, 355–369 (1999).
3. B. Belhomme and R. Pitz-Paal and P. Schwarzbözl, Journal of Solar Energy Engineering **136**, 011005-011005-7 (2014).
4. B. Belhomme and R. Pitz-Paal and P. Schwarzbözl and S. Ulmer, Journal of Solar Energy Engineering **131**, 031002-031002-8 (2009).
5. R. Flesch et al., Solar Energy **155**, 1273–1281 (2017)