Genetic optimization for radio interferometer configurations

Adam Gauci,[1]★ Kristian Zarb Adami,[2] John Abela[1] and Babak E. Cohanim[3]

[1]*Department of Intelligent Computer Systems, Faculty of ICT, Computer Building, University of Malta, Msida, MSD2080, Malta*
[2]*Department of Physics (Astrophysics), University of Oxford, Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, UK*
[3]*Mission Design Group Leader, Draper Laboratory, 555 Technology Square, Cambridge, MA 02139, USA*

ABSTRACT
The large bandwidth and resolution specifications of today's telescopes require the use of different types of collectors positioned over long baselines. Innovative feeds and detectors must be designed and introduced in the initial phases of development. The required level of resolution can only be achieved through a ground-breaking configuration of dishes and antennas. This work investigates the possibility of the genetic optimization of radio interferometer layouts given constraints on cable length, required UV density distribution and point-spread function. Owing to the large collecting area and the frequency range required for the Square Kilometre Array (SKA) to deliver the promised science, the configuration of the dishes within each station is an important issue. As a proof of concept, the Phase 1 specifications of this telescope are used to test the proposed framework.

Key words: telescopes.

1 INTRODUCTION

The required levels of sensitivity and resolving power required from today's telescopes cannot be achieved solely through the use of state-of-the-art receivers: investigations to find the optimum element configuration must also be undertaken. For mega-structures such as the Square Kilometre Array (SKA), searching through the space of possible layouts can be challenging owing to the enormous spatial extent. In addition to a pioneering design to achieve the desired specifications, the solution should also be constrained to minimize the infrastructure, networking and other costs related to the cabling and trenching that connect the stations together. This research builds on work by Cohanim, Hewitt & Weck (2004), Hassan et al. (2005), Grigorescu et al. (2009) and Fridman (2001) to investigate the applicability of machine-learning techniques to determine the optimum configuration for the collecting elements within a large radio telescope such as the SKA.

Although the construction of the SKA will follow a phased approach, even from Phase 1 it will be a formidable instrument. A total of 250 parabolic dishes will be installed over a region of radius 100 km (Dewdney 2010). One hundred and twenty-five core stations (50 per cent) will be fixed in the central 500-m radius. The inner and middle regions will extend over radii of 2500 and 100 000 m and will contain 50 (20 per cent) and 75 (30 per cent) antennas, respectively.

In this study, a framework to search for the optimum configuration that maximizes the uniformity of the UV density distribution and simulates the required point spread function (PSF) while keeping the connecting wire length to a minimum is implemented and tested.

★ E-mail: adam.gauci@um.edu.mt

Unlike previous studies that consider only a limited number of elements, the best design for mega-dimensional telescopes such as the SKA is sought. The developed system repositions the dishes to search and converge towards an optimal configuration according to user-defined fitness criteria.

The next section contains details of the implemented genetic operators and optimization functions. The results are presented in Section 3, and the conclusions and future work in Section 4.

2 GENETIC OPERATORS AND FITNESS FUNCTIONS

In this study, genetic algorithms (GA) are applied and tested to determine an optimum configuration for dish arrays. Such search heuristics follow the natural process of evolution to determine the fittest hypothesis from a pool of possible solutions. Pioneered by John Holland in the mid-1970s (Holland 1975), this evolutionary programming methodology uses mating, mutation and a fitness function as described by Mitchell (1997) to create better chromosomes.

Encodings representing different configurations and that store the x and y coordinates of dish locations are created. A crossover operator that generates further offspring solutions by combining element positions in different regions is defined. This is applied on pairs of chromosomes (parents) to generate six new individuals. For instance, if the core, inner and middle regions of the first parent are represented by C1 I1 M1, and the second parent is made from C2 I2 M2, encodings with C1 I2 M1, C1 I1 M2, C1 I2 M2, C2 I1 M1, C2 I1 M2 and C2 I2 M1 are generated. Two other encodings are created by combining random parts of the two parent configurations.

A mutation operator is also implemented to alter the positions of randomly selected elements. This shifting allows the algorithm to search a broader space and prevents it from converging to a local minimum. Every encoding is processed by a gene repair function to ensure that the 250 stored locations are distinct and that there are no dishes that are closer than the minimum allowed distance.

To direct the search towards a solution that maximizes UV coverage, the density map is computed from every unique pair of dishes. The framework is intended for configurations with a large number of stations, and the computation of the distances between all baselines becomes prohibitively expensive very quickly. Computational efficiency is increased by following and adopting the approach used by Cohanim (2004) and Cohanim et al. (2004). A fixed nominal grid with the required baseline distribution is initially generated, and each UV point is mapped to the closest grid point. The percentage of non-mapped grid points gives an accurate measure of the difference between the actual and the required baseline distributions. As shown in Cohanim et al. (2004), this can be calculated using equation (1). Here, N_{total} is the total number of points in the nominal grid, and $N_{matched}$ is the total number of matched points. The numerator equates to the percentage of grid points that are not matched with any UV point. Ideally, all nominal grid points are flagged by at least one UV point.

$$f_{UV} = \frac{N_{total} - N_{matched}}{N_{total}}. \tag{1}$$

Owing to its size and planned operation, the SKA will be a log-based structure, so a log nominal grid distribution is used. The goal of the GA is set to minimize the percentage of non-matched nominal grid points.

To speed up the fitness evaluation process, the nominal grid is represented as a k-dimensional tree that is built only once and stored in the memory. The nearest point neighbours are then identified by traversing the constant binary tree data structure. Non-leaf nodes represent a perpendicular hyperplane that divides the space into two subspaces. The left subtree points to other nodes on the left, while the right subtree represents points to the right.

Apart from optimizing for UV distribution, the fitness function must also take into account the length of cable required to connect the dishes together. Various approaches to optimizing a telescope layout infrastructure have been presented in the literature. Grigorescu et al. (2009) describe a set of algorithms that take into account trenching, as well as connection, costs. In Cohanim et al. (2004), the Single Linkage algorithm is used. In this work, the shortest sequence that connects all vertices together is determined using the Kruskal minimum spanning tree (MST) algorithm (Cormen et al. 2001). A cable with unit cost per unit length is assumed. Dish locations are connected so as to create an undirected graph where the edges (connections) between each vertex (dishes) have no particular direction. The weight of every edge is taken to correspond to the Euclidian distance between the two connecting nodes.

A normalizing function that allows the computed wire length to be scaled and compared with the other fitness criteria is defined. Chromosomes with elements that can be connected by a cable length of less than 250 km are not penalized. Encodings with a minimum wire length greater than 500 km are strongly discouraged through a fitness assignment of 1. Intermediate cable lengths are given a weighting that varies linearly as described by equation (2). Results from initial tests helped to empirically determine and fine-tune such threshold values.

$$f_{wire} = \begin{cases} 0 & \text{if } 0 \leq \text{wire length} < 250; \\ 0 \to 1 & \text{if } 250 \leq \text{wire length} < 500; \\ 1 & \text{if wire length} \geq 500. \end{cases} \tag{2}$$

The residual error between the resulting PSF and a delta function is also considered to determine the fitness of each chromosome. In particular, OSKAR-2 (Mort, Dulwich & Salvini 2012) is used to carry out the beam-forming simulations at a frequency of 1.4 GHz over one time step. This software package makes possible the computation of the PSF in reasonable time by utilizing the Message Passing Interface (MPI) and Graphics Processing Units (GPUs). A sky model with a single source at zenith is assumed, and the positions of 250 stations are defined in each run. The simulator is set to consider a field of view (FOV) of 0.01 and to produce images of 256×256 pixels, making each element equal to 0.1406 arcsec. The average difference between the resulting PSF and the ideal scenario of one bright central pixel is computed and considered as the third fitness function (f_{psf}).

It is important to keep in mind that the implemented GA solver assumes a minimization optimization function. All fitness evaluation criteria are defined in such a way as to give smaller values for better encodings. The goal is to minimize the overall fitness with subsequent generations and to stop iterating when no improvement is detected.

3 RESULTS

As a proof of concept, the GA framework is tested on the proposed SKA dish array specifications. The core and inner regions are taken to have 125 and 50 randomly positioned receivers respectively, while elements within the middle region are grouped in small clusters of three to eight dishes each. Because the outermost region extends from 2.5 to 100 km, all configurations are checked to ensure that they have at least three groups of dishes between 2.5 and 20 km, 20 and 40 km, 40 and 60 km, 60 and 80 km, and 80 and 100 km. This restricts the algorithm from moving all dishes towards the core in subsequent iterations.

The UV distributions for small groups of receivers positioned in a straight line, in a triangle, as a snowflake, in a circular pattern and in a Reuleaux triangle orientation are initially compared to determine the configuration that produces the best results.

The GA is set to have an initial population of 512 random chromosomes. The sum of the UV distribution fitness as defined by equation (1), the wire-length fitness as defined by equation (2) and the PSF fitness discussed in Section 2 are considered for performance ranking. Crossover is achieved by swapping the core, inner and middle regions between parents. This generates valid offspring as the number of dishes will remain constant. The implementation of the mutation function is slightly more complicated, as different encodings can have different numbers of groups with different numbers of dishes. When the position of a dish within a group is to be changed, the entire cluster is repositioned. In every generation, the highest-ranking (lowest-fitness) individuals are selected from: (i) all 512 chromosomes of the previous population, (ii) 1024 chromosomes created by crossover, (iii) 512 chromosomes created by mutation, and (iv) 512 new random chromosomes.

Half of the chromosomes in the initial population are given random dish positions. The rest are jittered versions (± 50 m) of the generic configuration as defined by the SKA Configurations Task Force (CTF) which is shown in Fig. 1. The ultimate goal of this experiment is to determine whether the GA manages to improve on the fitness of the CTF configuration.
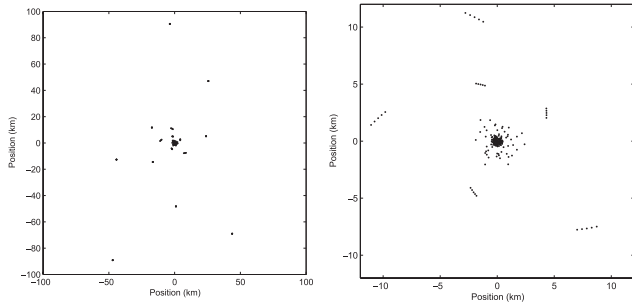
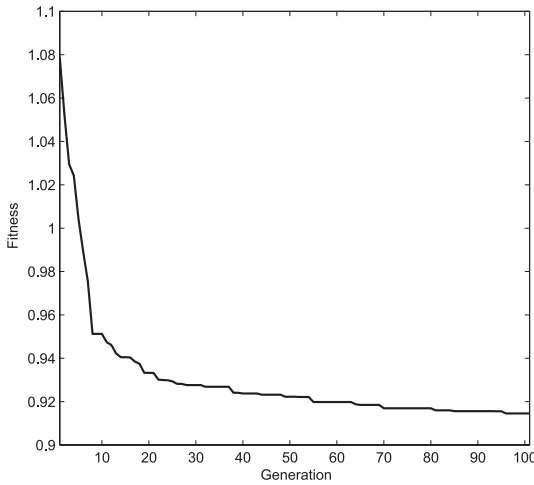**Figure 1.** Full (left) and inner (right) regions of the generic SKA CTF dish configuration.

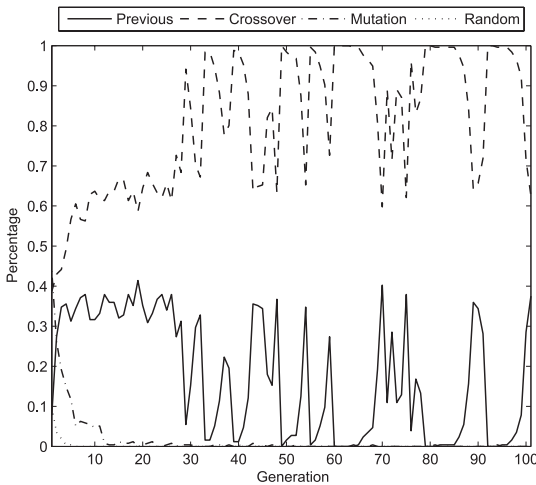**Figure 2.** Fitness calculated at every generation.

Figure 3. Percentage of chromosomes selected from the previous generation (solid line), created by crossover (dashed line), created through mutation (dash–dotted line), and randomly generated (dotted line), which are selected at every generation.

Fig. 2 presents the minimum computed fitness at every generation, which decreases with subsequent iterations, indicating convergence towards better encodings. Fig. 3 shows the percentage of chromosomes from the previous generation, the percentage of chromosomes created by the genetic operators, and the percentage of new random chromosomes that are selected at every generation.
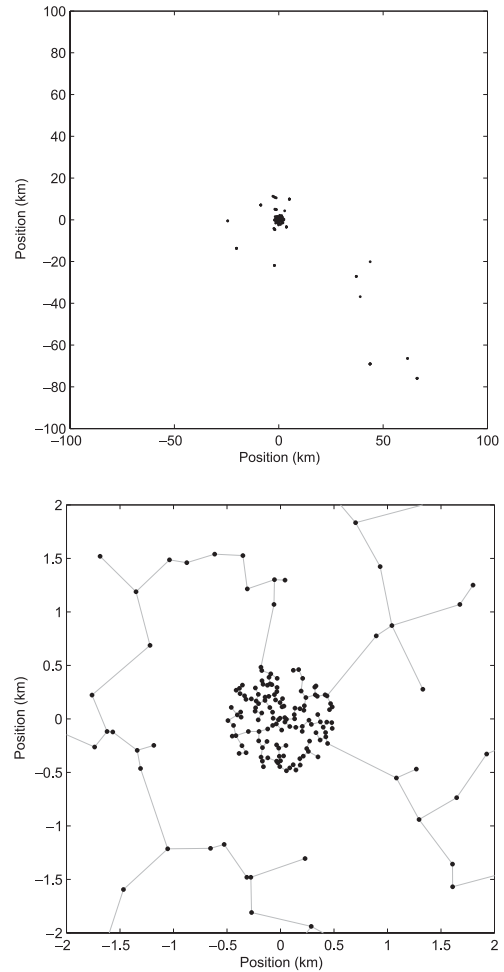

Figure 4. Full (top) and inner regions with connecting wire (bottom) of the genetic algorithm configuration.

Random configurations stop being the fittest and are not selected to procreate after the first few iterations. This indicates that the developed genetic operators are producing better offspring. As the algorithm progresses, less chromosomes created by mutation are preferred, signifying that random changes are not generating fitter individuals. The last few populations are found to contain either individuals created in previous iterations or encodings generated by the crossover operator. This indicates that the algorithm is converging and an optimized solution is found.

The resulting configuration after 101 iterations (including a wiring plan for the central region) is shown in Fig. 4. Fig. 5 presents the corresponding UV distribution and PSF. The total cable length is found to be 229.5584 km, which is less than the minimum penalty value, while f_{UV} and f_{PSF} are 0.7214 and 0.1931, respectively, giving a total fitness of 0.9145.

To be able to compare and evaluate the results, the same fitness functions are also computed for the CTF configuration. The UV fitness (f_{UV}) is found to be 0.7127. The MST algorithm gives a wire length of 384.67 km and hence an f_{wire} of 0.5387. OSKAR-2 (Mort et al. 2012) is again used for PSF estimation, and an f_{PSF} of 0.2507 is determined. Fig. 6 presents the UV distribution and the output from the imager. The total fitness of the CTF configuration is 1.5021.

A configuration based on Releaux triangles, as defined by Keto (1997), is also tested. The dishes in the core are positioned over two slightly rotated triangles. Similarly, dishes in the inner region
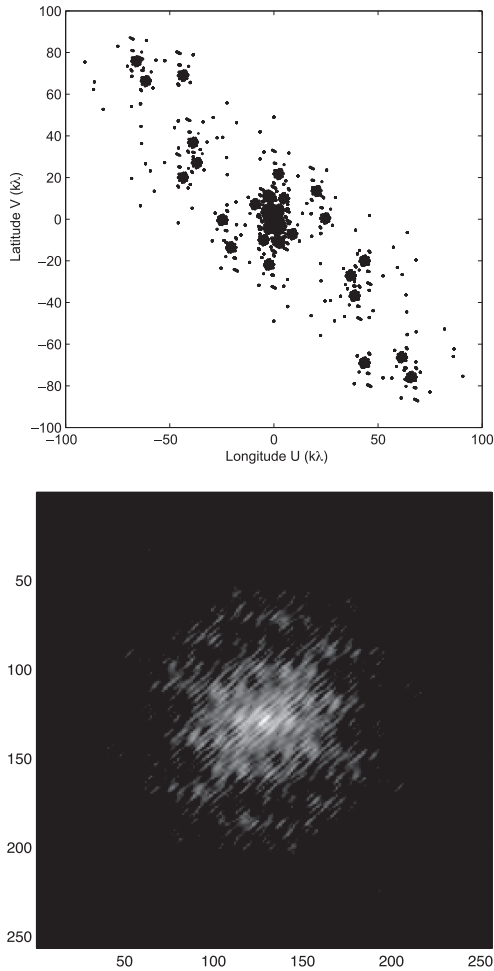
Figure 5. Resulting UV coverage (top) and PSF (bottom) of the genetic algorithm configuration.
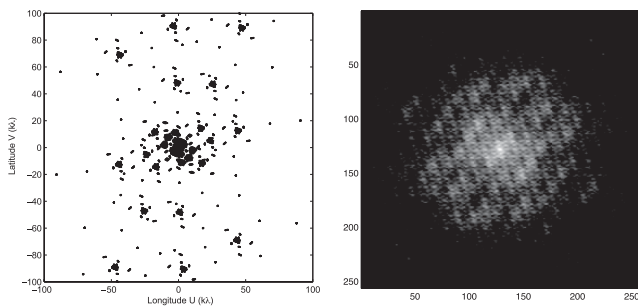

Figure 6. UV coverage (left) and PSF (right) of the CTF configuration.

are placed according to a similar but larger shape. Receivers in the middle region are grouped but positioned on randomly selected points from a predefined triangle. The fitness criteria f_{UV}, f_{wire} and f_{PSF} are found to be 0.8794, 1.000 and 0.2134 respectively.

4 CONCLUSIONS

In this work, a genetic optimization framework for the determination of interferometer layouts based on a desired UV distribution, PSF and wire constraints is developed. The 250 dishes planned in Phase 1 of the SKA telescope are used as a proof of concept. Although

a log-based nominal grid is used, this can be changed depending on the required distribution. The assumed ideal PSF can also be changed accordingly.

Even though a large number of elements and long baselines are involved, large genetic population sizes are used as much as possible. The computation time required is reduced through parallelization. Table 1 summarizes the results and compares the fitness values obtained by the GA, SKA CTF and Realuax triangle configurations. When using the defined fitness functions, encodings that have a good UV coverage, minimum cable length and a well-defined main central peak in the PSF are obtained. Although equal weighting is given to the three criteria, this can easily be adapted to accommodate user-specific requirements, budget and scientific goals. For instance, giving less importance to the wire-length criterion will improve performance but may significantly increase the infrastructure costs.

Seeding the initial population with chromosomes based on the CTF configuration helps the GA to start off with a good dish layout. The random individuals are the least fit encodings and are immediately filtered out. In subsequent iterations, the algorithm tries to reduce the wire length by moving all dishes towards the centre of the telescope and positioning the clustered elements in the same quadrant. However, the UV fitness directs convergence towards an evenly spaced distribution. The output is proof of the algorithm's attempt to direct the search towards a solution that tries to satisfy all of the defined constraints.

Although the UV fitness decreases slightly by 0.02 between the CTF and GA configurations, improvements are noted in the wire-length and PSF residuals. An overall fitness increase of 0.5876 is observed. The final wire-length fitness of the GA configuration is zero because of how equation (2) is defined. As discussed in Section 2, this is related to the total minimum cable length required to connect all nodes together and does not relate to the number and length of baselines involved.

Simulations that take into account the Earth's rotation and that consider the UV projection over a number of hours were also performed to investigate the improvements in resolution. As expected, longer exposure times cause more of the nominal grid points to pair, giving better UV fitness values.

Through this and similar work, the potential of machine-learning techniques to aid in identifying optimal dish configurations is demonstrated. Promising results are obtained, and further analysis can be undertaken once more detailed specifications on the SKA are made available. The work done by Bounova & deWeck (2005), which describes an optimized framework to model robust and scalable networks, may also be considered to derive the best configuration for the full SKA. Planned future work includes an analysis of various normalizing functions for the wire length, a detailed analysis of the effects of the power spectrum, and other fitness measures.

Table 1. Results of the defined UV, wire-length and PSF fitness functions as well as the total chromosome fitness. Smaller values indicate better encodings.

Test case	f_{UV}	Wire length (km)	f_{wire}	f_{PSF}	f
Genetic algorithm (seeded with random and CTF configurations)	0.7214	229.5584	0.0000	0.1931	0.9145
SKA CTF configuration	0.7127	384.6702	0.5387	0.2507	1.5021
Reuleaux triangles configuration	0.8794	608.8108	1.0000	0.2134	2.0928

REFERENCES

Bounova G. A., 2005, MSc thesis, MIT

Cohanim B. E., 2004, MSc thesis, MIT

Cohanim B. E., Hewitt J. N., Weck O., 2004, ApJ, 154, 705

Cormen T. H., Leiserson C. E., Rivest R. L., Stein C., 2001, Introduction to Algorithms. McGraw-Hill, New York

Dewdney P., 2010, Ska-phase 1: Preliminary system description and cost estimate. SKA Memo 130. Available online at: http://www.skatelescope.org/uploaded/21705_130_Memo_Dewdney.pdf (last accessed on 09/02/2013)

Fridman P. A., 2001, Proc. SPIE 4472, Applications of Digital Image Processing XXIV, p. 518, doi:10.1117/12.449782

Grigorescu G., Alexander P., Bolton R., McCool R., 2009, Exp. Astron., 27, 95

Hassan R., Cohanim B., De Weck O., Venter G., 2005, 46th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, Austin, Texas, pp. 1–13

Holland J. H., 1975, Adaptation in Natural and Artificial Systems. University of Michigan Press, Ann Arbor, MI

Keto E., 1997, ApJ, 475, 843

Mitchell T. M., 1997, Machine Learning. McGraw-Hill, New York

Mort B., Dulwich F., Salvini S., 2012, Oskar-2: Beamforming simulation for the ska aperture array, 2012. Available online at: http://www.oerc.ox.ac.uk/research/oskar1 (last accessed on 09/02/2012)

This paper has been typeset from a TEX/LATEX file prepared by the author.