

Cost-effective infrastructure for the data signal transport

G. Grigorescu¹, P. Alexander^{1,2}, R. Bolton¹, and R. McCool³

¹ Astrophysics Group, Cavendish Laboratory, J.J. Thomson Avenue, Cambridge CB3 0HE, UK, e-mail: g.grigorescu@mrao.cam.ac.uk

² Kavli Institute for Cosmology Cambridge, University of Cambridge, Madingley Road, Cambridge CB30HA, UK

³ Square Kilometre Array Program Development Office, Jodrell Bank Centre for Astrophysics, School of Physics and Astronomy, University of Manchester, Oxford Road, Manchester M13 9PL, UK

Abstract. The future Square Kilometre Array (SKA) radio telescope is an interferometer array that will use a variety of collector types, including approximately 2500 dishes distributed with separations up to a few thousand kilometres, and about 250 aperture array (AA) stations located within 200 km of the core. The data rates associated with each individual collector are vast: around 10 GBytes/s for each dish and 2 TBytes/s for an AA station. As each of these must be connected directly to a central correlator, designing a cost-effective cabling and trenching infrastructure presents a great engineering challenge. In this paper we discuss approaches to performing this optimisation.

In graph theory, the concept of a minimum spanning tree (MST) is equivalent to finding the minimum total trench length joining a set of n arbitrary points in the plane. We have developed a set of algorithms which optimise the infrastructure of any given telescope layout iteratively, taking into consideration not only trenching but also cabling and jointing costs as well. Solutions for few example configurations of telescope layout are presented. We have found that these solutions depend significantly on the collectors' output data rates. When compared to a "traditional" MST-based approach which minimises trenching costs only, our algorithms can further reduce total costs by up to 15-20%. This can influence greatly the SKA infrastructure related costs.

1. Introduction

The outline design of the SKA has been described in Schilizzi et al. (2009). The telescope will be an interferometer array, comprising approximately 2500 dishes, operating from about 700 MHz to 10 GHz. These will be distributed over continental distances (up to at least 3000 km from the centre) but will be concentrated in a core a few kilometres across. In addition to the dishes there will be around 250 aperture array stations, operating at lower frequencies, from about 70 MHz up to 1 GHz and distributed over a few hundred kilometres (Bolton et al. 2009; Alexander et al. 2007).

The design of the SKA presents many technical challenges (Hall 2004). This paper is motivated by the challenges involved in transporting data signals around the telescope. Because the SKA will be an interferometer, the data from the individual collectors must be combined in a correlator, located near to the centre of the array. Although long-distance data transport is done commercially, because of its necessarily remote location, the SKA will require a dedicated fibre-optics network. With data rates of the order of 10 Gbytes/s from each dish and possibly 2 Tbytes/s for the AA stations, the total data rate to the correlator is set to be of the order of 520 Tbytes/s, dwarfing the current (end 2008) world internet traffic rates (Minnesota Internet Traffic Studies 2009).

With so much data to transport over such large distances, it is likely that well over €100 million will need to be spent on equipment to generate, condition and receive the signals and on fibre and trenching to carry the data. Many millions of Euros can potentially be saved if the data-transport network can be optimised effectively by changing the fibre paths to trade off the amount of trenching that must be dug against the length of

fibre and the quantity of networking that is required. Individual data links cannot be treated separately in the search for a global cost minimum and so, with close to 3,000 collectors, such a complex problem cannot be solved by hand. In this paper we present algorithms that we have developed which can be used to find an optimum networking layout for a given configuration of collectors. We illustrate these with some example networks.

2. Defining the problem

In a simplified approach, the main question to be answered is the following: given a set of collectors located in a plane, where for each collector a position (x, y) and an associated data rate are known, what would be the most cost effective way to route the signal of each collector down to one (or more) central base (CB) location(s)? An optimal solution will depend on parameters such as prices per unit length of trench and cable as well as on the position of the CB. The problem is often encountered in commercial telecommunication network design where constraints such as obstacles or topographically excluded regions can arise. Other applications dealing with similar optimisations include printed circuit board design, thin soap film solutions, urban planning of pipe networks, etc.

The problem is defined by several variables. These are:

- a telescope layout: a list of collectors, each collector i being uniquely specified by an integer and having a position in space (x_i, y_i) and an output data rate D_i [Gbytes/s].
- a list of one or more central bases, acting as signal sinks.
- a 2-dimensional site mask (terrain map) of local trenching costs per unit length acting as a trenching cost metric. In

the simplified case, this value is constant over the entire telescope area.

- a list of possible cable choices and the associated bandwidths and costs per unit length. This is especially important in the context of individual collector output data rates possibly varying for each collector. We assume that more than one cable can be used in one trench.

Based on the above input, the expected solution is represented by an infrastructure layout specification of all trenching and cabling such that the total costs associated to this infrastructure are minimum.

3. Overview of relevant mathematics

From a mathematical point of view, the problem is similar to finding a minimum path length that interconnects n vertices located in the plane. In graph theory, the problem is known as the *Euclidean minimum spanning tree* (EMST). An example is given in Fig. 1, where a set of 25 points in the plane are joined by a minimum path length, defined as a summation over all links ij : $S_T = \sum s_{ij}$, with $s_{ij} = \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2}$ as the Euclidean distance between points i and j . Randomised algorithms of complexity $O(n \log n)$ are known (Cormen et al. 2001).




Fig. 1: The Euclidean minimum spanning tree associated to a set of 25 points in the plane.

Finding a shorter total path length than the EMST can be achieved by introducing new vertices. Steiner showed that the total path length of a tree interconnecting a set of n points lying in the plane can be reduced with respect to the EMST, by introducing a maximum of $n - 2$ intermediate vertices which act as hub points, also known as Steiner points (Hwang et al. 1992). The Steiner points must have a degree of 3 - i.e. connect to three other vertices - and the three edges incident to such point must form three 120° angles. For a given set of points, there may be several Steiner trees. Finding the Steiner tree of a graph is, in general, an NP-complete problem, thought to be computationally difficult. In practice, heuristics and iteration are used: namely, a large collection of points is divided into sub-sets of three or four points and their Steiner trees are individually computed and then joined together. Steiner trees for sets of three and four points respectively are shown in Fig.

2. The length of a Steiner minimum spanning tree cannot be shorter than $\sqrt{3}/2$ (about 0.86) times the length of the EMST (Du & Hwang 1992).




Fig. 2: Examples of Steiner trees for sets of three points (a) and four points (b). Steiner points are shown as squares.

4. A particular problem

Although there are obvious similarities between finding the Steiner tree of a graph and finding an optimal network infrastructure, a few important differences hint to a different optimum and therefore dictate a different approach.

First of all, not all elements in a telescope network play an equal role: the location(s) of the CB(s) will dominate the optimisation. Also, certain configurations may decouple into several sets of collectors, each set being independently connected to the CB(s) from the others.

Second, another crucial component to be taken into account, apart from the trenching, is the cable infrastructure. Depending on band-width, cables come in different varieties with different costs. This is an essential consideration since different collectors may have significantly different data-rates. Also, including the cable costs into the infrastructure cost has another effect: the problem of optimising of the SKA infrastructure layout becomes drastically different than the problem of finding the total minimum path joining a set of vertices, described in Sec. 3. Essentially, this is due to the fact that, whereas there was a direct correspondence between trenching length and the total path of a minimum spanning tree, the cable infrastructure adds another layer of complexity to the problem: routing the signal of further collectors via trenches of collectors closer to the core will demand more cable length than direct connections to the CB. This translates into higher costs for cable.

Third, variable trenching costs, depending on location, must also be allowed for, together with areas which are topographically inaccessible.

All these factors will lead to an optimal infrastructure which may differ significantly from the EMST. For example, in the extreme case of very expensive data links and low cost trenching, the solution for a sparse configuration consists of all the collectors being directly connected, through individual trenches, to the CB and thereby minimising the total cost, which is driven mainly by the cost of the cable. This solution

would actually maximise the total trench length but would still be the most cost-effective solution. An approach based on minimum spanning trees would therefore fail.

If both cable and trench costs are significant, the optimum solution for the network layout will lie between the extreme configurations of the EMST, which minimises trenching costs and the direct connecting of each collector to the CB, which minimises the cable costs.

Having considered all these aspects, it becomes clear that an EMST-based optimisation is incomplete. A different approach is therefore taken.

5. Optimisation algorithms

The solution for the optimal network layout is found by iteration. We have developed a set of algorithms which can be applied successively (Grigorescu et al. 2009). They each explore different aspects of the optimisation and several solutions may be generated by running different combinations of these algorithms. Nevertheless, a particular succession is preferred in order to improve CPU time and avoid the optimisation being stuck in local minima.

The total cost, TC , consists of three different contributions:

$$TC = K_t + K_c + K_e. \quad (1)$$

These contributions are as follows. $K_t = k_t \cdot \sum_{links} s(i, j)$ represents the trenching costs. k_t is the trenching cost per unit length and may in general be position dependent. The summation is carried out over all links. $K_c = \sum_{cable\ types} \sum_{links} k_c \cdot s(i, j)$ is the total cost of cables, k_c is the cable cost per unit length of a particular cable of type c . K_e are costs associated with non-collector elements such as Steiner points.

We consider a telescope to be a collection of elements, each element being defined by the following attributes: a unique integer identifier; a cartesian position in the plane (x, y) ; input, output and intrinsic signal data-rates, all of which can be zero; a pointer to the other element to which the current element is connected; a size and a cost; other variables related to book-keeping. In this way, no distinction is made between collectors, CB's and other passive elements of the telescope such as possible Steiner points which can be added during optimisation. The direction of the information flow is important as it is used to calculate the data rate transported through each trench. Elements are always connected such that no loops are created in the data flow.

The starting point we choose for the optimisation is to construct the minimal cable-cost network by connecting each collector directly to the nearest CB. This is called the *star configuration*. Our experiments indicated that this is the best start configuration: it converges the quickest to an optimum configuration and at the same time it is never a local minimum. The optimisation proceeds by employing a series of algorithms sequentially. Below, some of these algorithms are explained in detail.

Connect Serially: This algorithm cycles through the entire set of collectors in an orderly fashion, from the furthest to the

closest of the collectors with respect to the CB(s), and attempts to join them serially such that the total cost decreases with every attempt. The start point is the most distant collector with respect to its closest CB. Its unique identifier is added to a book-keeping list, which initially is empty. A cost-effective attempt is made to connect this to the next furthest collector. If successful, the unique identifier of the next furthest collector is also added to the book-keeping list. The next to next furthest collector is then considered as a potential connecting point. The algorithm continues in this way until the CB is reached in the chain of connections. If the book-keeping list does not contain the entire set of collectors, the algorithm is applied again, using as a start point the most distant collector which is not included in the book-keeping list. This algorithm is most powerful when applied first, directly to the star configuration because it converges only within few iterations. Also, *Connect Serially* can improve the total cost significantly within one cycle. It is very light in terms of needed CPU time per algorithm cycle.

Connect Lines to Neighbours: The purpose of this algorithm is to break longer more direct links from collectors further away from the CB into a network including more links of shorter length, making use of the elements closer to the core as hub points. It is implemented as follows: for each link in the telescope, let elements i, j be the the start, stop points respectively of the link, where start and stop are defined by the data flow direction. A search for the closest element k to the line ij is performed. If the projection of k on the line ij falls inside the segment ij , the element i is disconnected from j and reconnected to k ; a new TC is computed; if $TC_{new} < TC_{old}$ then the newly made connection is retained. It is preferred to run *Connect Lines to Neighbours* after having used *Connect Serially*: this will improve upon the possible changes missed by *Connect Serially* but will need only just few cycles to converge, as the layout had already been optimised by the first algorithm.

Add Steiner Points: This routine deals with reshaping the connections for groups of three elements by introducing Steiner points, as shown in Fig. 2. First, the algorithm searches for the following: all elements i to which at least two other elements connect: $j \rightarrow i, k \rightarrow i$; or all three elements i, j, k are connected in a chain: $k \rightarrow j \rightarrow i$. If all the angles of the ijk triangle are less than 120° , the corresponding Steiner point F is constructed. Then, both k and j elements are connected to F and F is further connected to i . TC is recomputed and the new connections are accepted if $TC_{new} < TC_{old}$. *Add Steiner Points* is used once the layout has been optimised using the algorithms presented above.

Relax Steiner Points: The Steiner points added by the previous routine were constructed such that the total trench length within the triangles was minimised. Therefore, the construction of these points does not account for minimising costs due to cables. Nevertheless, the Steiner points are good first estimates. This routine allows the Steiner vertices created previously to relax by moving iteratively in four possible orthogonal directions. The step size is taken to be a fraction of the distance to the closest vertex of the triangle. The chosen direction of movement is the one which gives the largest negative gradient in TC . The variable step size ensures a good compromise between execution time and precision. There are two possible outcomes:

either the Steiner point settles to a new position which minimises the combined cost of both trench and cable for the given triangle of collectors or the Steiner point moves such that it approaches one of existing vertices of the triangle until they coincide. In the latter case, the Steiner point is removed from the construction and new connections are made accordingly. In this way, Steiner points that are created but destroyed in the relaxation process will contribute actively to the optimisation.

6. Algorithm complexity

The scaling of algorithm complexity with the number of collectors is not trivial. First of all, the performance of all algorithms depends on the evolution of one parameter, the total cost TC. At each step within one cycle, changes in the infrastructure are performed only if TC decreases. Computing changes in TC is in itself a time-expensive task, scaling as $\sim \log(n)$, as it requires cycling through branches of collectors and summing individual costs corresponding to each encountered link in the network. The optimisation algorithms scale as $\sim n$, as they cycle through the entire set of collectors. Also, at each step within one algorithm cycle, other actions are performed, such as looking for the closest neighbour or closest link, which also may scale as $\sim n$. Indeed, the overall computation time increases in general at least as fast as $\sim n^2 \log(n)$. This scaling is inevitable and is justified by the fact that in our model, for every local change of a link, an update of the cable component of the downstream signal path is also needed. Therefore, as local changes affect the global optimum, finding a optimum solution will require significant computation time. Fortunately, the total number of collectors for the SKA telescope will not exceed a few thousand. These configurations can therefore be optimised within reasonable times using the numerical computational framework discussed above. For instance, on a modern workstation, the optimisation of layouts of around 2500 collectors takes less than 5 minutes.

7. Examples and discussion

Two different optimisations of the same layout are presented for comparison in Fig. 4 and 5. Whereas one of the examples is optimised for total trench length only, equivalent to finding the minimum spanning tree of the associated graph, the other example optimises for the combined cost of trench and cable and therefore employs the full power of our method.

Some input parameters used for the optimisation are listed in Table 1. The links were modelled using five options for cabling, namely optic fibre cables consisting of 1, 2, 6, 9 and 24 tubes. Each tube consists of 8 strands of fibre. Through each strand, we assume that 16 different colour channels can be broadcast, each channel having a maximum bandwidth of 1.25 Gbytes/s. The input values for trench and cable costs were chosen to reflect realistic commercial values (Serenate Project 2009). Cost wise, all Steiner points were treated as commercial optic fibre joints.

The two examples in Fig. 4 and 5 are two different optimisations of the same configuration of 165 collectors with a data

Name	Value
trench costs	€12k/km
cable costs (1 tube)	€0.85k/km
cable costs (2 tubes)	€1.2k/km
cable costs (6 tubes)	€2.05k/km
cable costs (9 tubes)	€2.7k/km
cable costs (24 tubes)	€5.1k/km
Steiner point costs	€0.4k

Table 1: Input parameters used for the optimisation.

rate of 2 TBytes/s each, consisting of a highly concentrated inner core of 1 km diameter and a sparse random distribution up to a 5 km diameter around the CB. This is a reasonable distribution of the central part of the AA station set. The infrastructure in Fig. 4 was optimised intentionally for total trench length only and it has an associated total cost of €898.55k, of which €560.4k are cable costs and €338.15k are trench costs. The infrastructure in Fig. 5 was optimised for total trench and cable costs combined and it has an associated total cost of €748.82k, of which €395.27k are cable costs and €353.56k are trench costs. As the individual data rate of each collector is high, about half of the total costs will be spent on cables. Therefore, this comparison shows that although the infrastructure in Fig. 4 has indeed a lower total trenching cost, the total cost is reduced in Fig. 5 by more than 16%. Namely, a trade-off of 10% in trench costs is necessary such that total costs are reduced by 16%. This is indeed a consequence of properly incorporating the cable costs into our optimisation procedure.

It is worth mentioning that, for the same layout of the 165 collectors, if one employs a star configuration, where each antenna is directly connected to the CB via its own trench and cable, the associated total cost is €1737k of which €518k are cable costs and €1219k trench costs. Our optimisation has therefore improved the total cost of the infrastructure by a factor of 230% with respect to the star configuration. We would like to emphasise that, although the total cable length used for the star configuration is less than the total cable length associated with the fully optimised infrastructure shown in Fig. 5, the costs related to the cable infrastructure are higher for the star configuration. This is due to the fact that the cable cost per unit length per unit of data rate does not increase linearly with increasing data rates, as seen in Table 1. Namely, 1-tube cables are four times more expensive than 8-tube cables, for a given data-rate. Therefore, sharing trenches between different collectors has also the intrinsic advantage of utilising larger, less expensive cables.

It was also found that for the configuration shown in Fig. 4 and 5, the addition of Steiner points (not shown in the figures) would further reduce the total cost by about 1% only. This finding is justified by the fact that the inner core is too dense to accommodate cost-effective Steiner points. The quoted 1% reduction is due to only a few Steiner points, placed outside the inner core, in the sparse section of the configuration. There, the high-data rate and therefore the expensive cables conspire against the need of more Steiner vertices, as that would require a higher total cable length. In general, Steiner points make a

greater impact on the total cost for configurations with low data rate collectors and sparse collector distributions (Bell 1999).

One more example is presented in Fig. 6. It consists of 2250 dishes, with individual data rates of 15 Gbytes/s. Most of the dishes are grouped in a highly concentrated core of 5 km across. The rest of the dishes are located along five curved arms which spiral outwards up to 150 km away from the core. The CB is placed non-symmetrically on the edge of the core. The infrastructure has been optimised for the combined cost of trench and cable costs. The total cost associated to the infrastructure of this configuration is €17469k, of which €15557k are trench costs and €1912k are cable costs. The relatively low data rate of the dishes dictate a trench dominated total cost, of which 10% are due to cable costs.

These examples show that a proper treatment of various infrastructure components and associated costs not only can reduce significantly the total cost but also can provide us with a valuable insight of how different components (Steiner vertices, cables, trenching) contribute differently to the total cost, as a function of collectors' layout and data rates.

8. Cost maps

So far, we have considered an Euclidean metric on a flat plane. A more realistic description should account for effects introduced by the Earth's curvature and varying altitudes. Also, locally variable trenching costs should be considered, which would reflect more accurately costs due to different soil types or different ways of data transport (for example, dug trenching versus over ground cable transport). One could account for the cumulated effect of all these features by employing a cost map, mostly in the form of a raster map which associates different penalty values for crossing particular regions on the map (Fig. 3). The cost map can be a sum of various weighted cost maps corresponding to independent effects. For instance, no go areas and topographic obstacles can be added to the cost map by costing very highly the forbidden regions. Costs related to access to power grids, roads, antenna deployment costs, can also be factored in. The algorithms described in Sec. 5 remain valid as long as Euclidean distances between any two collectors are replaced by "lowest penalty" path values, also called geodesic path lengths. Calculating the geodesic between two points on a cost map is not a trivial task. One may employ Dijkstra or A* algorithms for an heuristic approach. Alternatively, the fast marching algorithm, an analytical approach to finding geodesics on cost maps, may be more expensive but guarantees that the solution does not represent a local minimum but a global one (Petres et al. 2005).

9. Outlook

Future work will focus on expanding the current model to incorporate more realistic features. The option of using dark fibre (already available long distance optic fibre infrastructure) will be included. The possibility of sharing trenches between signal and power cables will also be investigated. Another realistic aspect which will be investigated is the use of cost maps, as described in Sec. 8.




Fig. 3: An example of altitude map (left) and the penalty map associated to it (right).

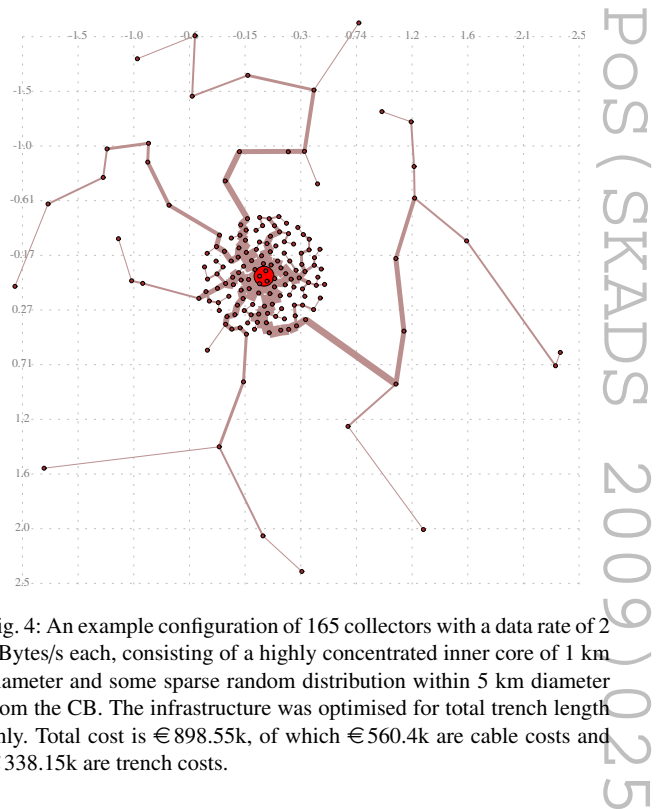


Fig. 4: An example configuration of 165 collectors with a data rate of 2 TBytes/s each, consisting of a highly concentrated inner core of 1 km diameter and some sparse random distribution within 5 km diameter from the CB. The infrastructure was optimised for total trench length only. Total cost is €898.55k, of which €560.4k are cable costs and €338.15k are trench costs.

References

- Alexander, P. et al. 2007, SKADS Benchmark Scenario Design and Costing www.skads-eu.org/PDF/SKADS_Benchmark_Scenario_D+C.v1.0.pdf
- Bell, B. 1999, <http://cse.taylor.edu/~bbell/steiner/>
- Bolton, R.C. et al. 2009, SKADS Benchmark Scenario Design and Costing 2 (SKA Phase 2 AA Scenario) www.skads-eu.org/PDF/SKADS_D&C_second_document.05112008_final.pdf
- Carilli, C. & Rawlings, S. 2004, Science with the Square Kilometre Array, *New Astronomy Reviews*, 48, 2004
- Cormen, T.H. & Leiserson, C.E. & Rivest, R.L. & Stein, C. 2001, *Introduction to Algorithms* MIT Press and McGraw-Hill, 2001
- Du, D.Z. & Hwang, F.K. 1992, A Proof of the Gilbert-Pollak Conjecture on the Steiner Ratio *Algorithmica*, 7:121–135, 1992




Fig. 5: An example configuration of 165 collectors with a data rate of 2 TBytes/s each, consisting of a highly concentrated inner core of 1 km diameter and some sparse random distribution within 5 km diameter from the CB. The infrastructure was optimised for total trench and cable cost. Total cost is € 748.82k, of which € 395.27k are cable costs and € 353.56k are trench costs.

Grigorescu, G. et al. 2009, Cost effective infrastructure in a multi-antenna telescope layout, *Experimental Astronomy*, 27, 2009, DOI:10.1007/s10686-009-9177-2

Hall, P.J. 2004, The SKA: An Engineering Perspective *Experimental Astronomy*, 17, 2004

Hwang, F.K. & Richards, D.S. & Winter, P. 1992, The Steiner Tree Problem *Annals of Discrete Mathematics*, 53, 1992

Minnesota Internet Traffic Studies 2009, www.dtc.umn.edu/mints/home.php

Petres, C. et al. 2005, Underwater path planing using fast marching algorithms, *IEEE*, 2: 814- 819, 2005

Schilizzi, R.T. et al. 2009, Preliminary Specifications for the Square Kilometre Array, www.skatelescope.org/PDF/memos/100_Memo_Schilizzi.pdf

Serenate Project 2009, www.serenate.org/publications/d14-serenate.pdf

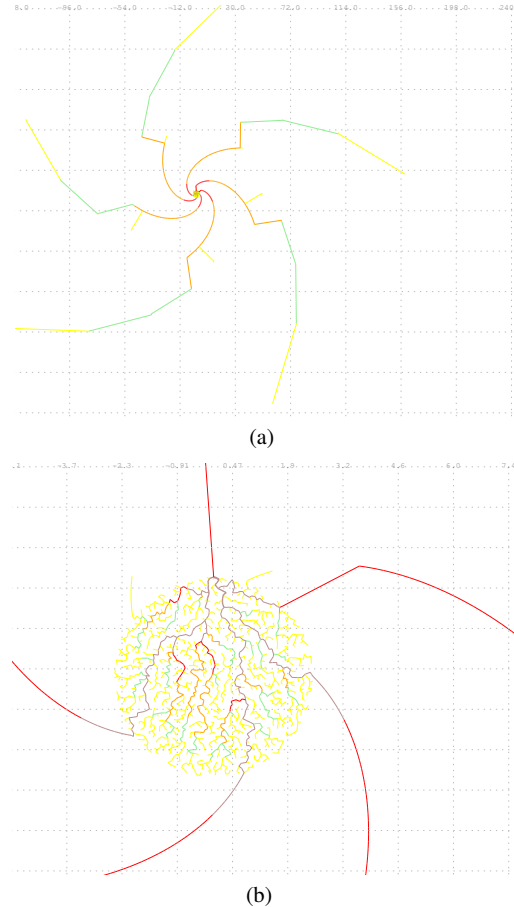


Fig. 6: A layout of 2250 dishes is shown, grouped in a highly dense core of 5 km across (enlarged plot) (b) and five spiral arms going out to 150 km radially (a). Only the links between the collectors are shown. The cables are colour-coded according to the different types of cables used. The total cost associated to the infrastructure of this configuration is € 17469k, of which € 15557k are trench costs and € 1912k are cable costs.