An evolutionary search algorithm to generate 3D-cloud fields with measured cloud boundary statistics

Victor Venema

1 Introduction

The cloud feedback is the largest single source of uncertainty in current climate models. One of the main outstanding problems in this field is the influence of the structure of clouds on radiative transfer and exchange processes at the top of the boundary layer.

For the radiative transfer calculations though structured clouds, we need a 3-dimensional matrix of Liquid Water Content (LWC) (and maybe average droplet diameter) to represent a cloud. This report introduces the idea of using a search algorithm for this job and, as a *proof of concept*, it presents a evolutionary search algorithm which searches for a cloud with simple cloud boundary statistics. As this report is mainly intended for use inside the project, it is aimed to be more informative as comprehensive.

1.1 Cloud model (LES)

A source of clouds with realistic cloud structure can be the output of a Large Eddy Simulation (LES) model. This is a good source, especially as additional to the 3D LWC field, one gets physically consisted fields of other meteorological parameters, e.g. droplet radius, temperature and humidity. Furthermore, with LES modelling one can study the temporal development of clouds, which is very difficult to measure.

However, these LES clouds are not easily usable together with empirical data: With a LES model is almost impossible to obtain clouds that have the same statistical properties as the measured clouds, due to the non-linear, chaotic relation between initial conditions and cloud properties. In a comparison with measurements made by Brown *et al.* (2001) the LES-models had the cloud cover was up to 0.4 off and the cloud base up to 300 meter too high. Using LES modelling to get a cloud field that also has a measured power spectrum, cloud cover and LWC will be even more challenging.

1.2 Measured cloud structure

A direct measurement of a cloud field could be a potential source of structured clouds. In situ measurements provide detailed information, but are 1-dimensional and cannot provide a 3D-cloud field.

With ground-based remote sensing – (microwave) radiometry, radar, and lidar – it is nowadays possible to get high quality cloud data by combining the various sensors. To make a direct measurement of the cloud field these instruments will have to scan in all directions, which takes a lot of time. For example, if the wind direction at cloud height were known, a 2-dimensional elevation scan (varying the elevation during the measurement) with the azimuth angle perpendicular to the wind vector, would be fastest solution. If the smallest scale of interest is 100 m and the wind speed 10 m/s, then one will have about 10 seconds for this 2D elevation scan. This time is comparable to the typical integration time of the remote sensing instruments. Therefore, one will only be able to measure in a few directions. Especially, in case of the radar, a shorter integration time will be difficult, as at the moment the sensitivity is already a problem for water clouds. Another difficulty with the measurement strategy is that the measurement volume changes during such a 2D elevation scan. As many cloud properties (e.g. the correlation length, and cloud overlap) are a function of the measurement scale, elevation scans can be difficult to interpret.

It will not suffice to measure the cloud properties in some directions and interpolate in between, as the LWC-field is shows variations at all scales, is determined by nonlinear dynamics, and shows intermitted behaviour (clearly seen in the cloud cover of a cumulus field or in the vertical profile).

With optical radiometry, faster measurements are possible, but the data gives less direct information on the macro and microphysical cloud properties, making it impossible to reconstruct a full 3D-cloud field.

Concluding, in the present time a direct measurement of structured clouds is too difficult and we are limited to measuring (statistical) cloud properties.

1.3 Methods using the power spectrum

There are two much-used methods to generate LWC of LWP time series with a realistic power spectrum: the Fourier and the bounded cascade method [Marshak, 1994]. Both methods can be used in a 2-dimensional version, to generate 2D LWP fields. The cloud fields made with these methods are assumed to have plane parallel cloud boundaries and some fixed droplets radius. Gaps in the clouds are not possible.

In the Fourier method one starts with a -5/3 power spectrum (or a measured one), and calculates the LWP spectrum from this by taken the square root. These complex Fourier coefficients are then multiplied with a random phase. After a discrete Fourier transform one gets the LWP time series or field.

The bounded cascade (fractal) method starts with a homogenous cloud, divides this cloud in two equal parts and redistributes the LWP/LWC in a random direction. These two cloud parts are again divided and their LWC is again redistributed, etc. By reducing the amount of water that is redistributed in every cascade by a factor of 0.8, one gets a signal with a power spectrum of about -5/3. Interesting about this method is that the signal is stationary, i.e. every point has the same probability density function, although a typical stationarity-test on a single signal would call this signal non-stationary as for example, the mean is not constant and the power spectrum shows much variation at the largest scales.

Two realisations from these two methods are compared in figure 1. The Fourier cloud signal is clearly much smoother than the bounded cascade one. This illustrates the fact that the power spectrum is not a full description of the structure. In Fourier terms, one could say that the phases are also important. The comparison with a measured

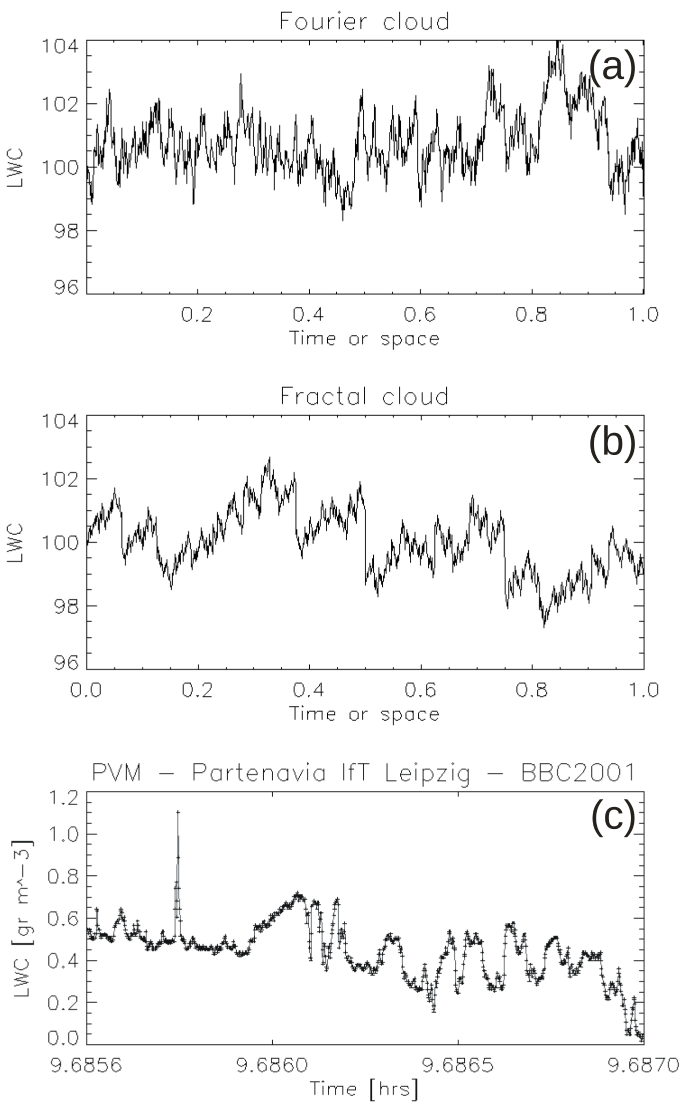

Figure 1. Comparison of the LWC/LWP time series made with the Fourier, fractal (bounded cascade) method and a real LWC time series measured with a PVM. The PVM data was measured by M. Wendisch, and S. Schmidt from the IfT, Leipzig, Germany.

LWC time series shows that the variance of the measured signal is much less smooth. The Kolomogorov theory for fully developed turbulence only claims that a passive conservative tracer has a -5/3 spectrum *on the average*. This is not supposed to hold for each realisation, which is what the above two spectral methods produce.

1.4 Search algorithm

A new approach could be to search for a cloud field: search a 3D-LWC-field that has the same statistical cloud properties as those measured, with respect to, e.g., LWP power spectrum, LWP histogram, cloud base and top height histogram. This has probably not been tried before, as the number of possibilities is very high, and one would therefore expect searching to be almost impossible. For example, if we are searching for a cloud with, $256 \times 256 \times 64 = 2^{22} = 4.2\ 10^6$ pixels, and use 256 LWC steps between zero and the maximum LWC. Then we would have: $= 2^{352}$? 10^{105} possibilities.

In the next chapter an evolutionary search algorithm will be introduced that will hopefully be able to solve this problem by including a sufficient amount of a-priory knowledge about cloud field. For example, by first solving the problem at a low resolution before going to higher resolutions make the algorithm much more efficient.

The disadvantage of a search approach is that the clouds will only have the statistical properties one puts into it. Imagine a cloudy atmosphere from a LES model that has similar feature as the measured atmosphere with regard to temperature and humidity profiles, average cloud top and base height, LWP histogram and LWP power spectrum. The other properties of this atmosphere will probably also be similar or at least realistic. These 'other properties' could be the full cloud top height histogram or some measure of intermittence. In case of a search cloud this will not happen automatically, a good statistical description of the cloud field is therefore paramount.

Another disadvantage of the searched clouds compared to LES-clouds is that one will not get physically consistent fields of humidity or temperature. These two fields are for example important in radiative transfer calculations in the infrared region. For many other applications, these fields will not be necessary. It would be possible to simultaneously search for these other fields as well and one could try to create compatible fields of humidity and temperature afterwards based on measured profiles and spatial correlations with LWC.

The most important advantage of the search approach is that one gets cloud fields whose properties are very close to the measured fields. This is, for example, important when comparing the measured radiative fluxes and cloud field properties with radiative transfer calculations. Compared to the methods using a prescribed power spectrum this method promises to be much better. Already a search cloud field based on a measured power spectrum and additionally a LWP histogram would be much more realistic than one just using a power spectrum.

2 Evolutionary search algorithm

The ideas behind evolutionary search algorithms follow the natural selection of species[†]. One takes an initial guess population of clouds and calculates how well they fit to the measured statistics. The best fitting ones are selected to produce the next generation. They are subject to small mutations to introduce diversity in the new population. This process is repeated until the quality of the fittest cloud is sufficient. As an evolutionary search is not exhaustive, it does not have to find the global optimum solution of the problem and my not converge to a sufficient solution at all.

The 3D -LWC field is the chromosome of this organism. The genotype (genes) and the phenotype (visible properties of the organism) are thus the same in this application. A set of cloud fields constitute a population. The fitness of each individual is calculated by comparing the statistical properties of the LWC -field with the measured statistics, see section 2.1. The selection pressure and the population size determine how fast algorithm converges to a solution, see section 2.2. To make potentially better clouds, various types of mutations can be introduced to the genome, e.g., point mutations, line mutation, and cross-over, see section 2.3.

The convergence of the search algorithm is accelerated enormously by first solving the problem at low resolutions before going to a higher one. This uses the scaling properties of clouds, i.e., the fact that the largest variations of clouds are at the largest scales. Without it, it would probably not be possible in practice to calculate cloud field with an interesting number of pixels. It is implemented by first solving (in 2D) the problem on a 1x1matrix, then 2x1, 2x2, 4x2, 4x4, etc. One needs a good criterion for increasing the resolution. Waiting for full convergence wastes calculation time, and may even never occur. However, failures that are still there at low resolution, take much more recourses to solve at a higher resolution.

2.1 Fitness functions and statistical parameters

The fitness function is a measure for the goodness of fit of the cloud field in the population compared to the measured cloud field. This fitness, which we are trying to maximize, is the opposite of the cost function, which one wants to minimize in optimising problems. The fitness function should not only indicate how well the solution fits, but also reward all improvements. Imagine you start with a cloud that has a cloud base at 2 km and the statistics say it should be at 3 km. In this case it would be an improvement to change the cloud base to 2.1 km, and this should be rewarded with a better fitness, although, when comparing cloud base height histograms it still does not fit at all. An example of the importance of this point is shown in figure 3.1 for the number of cloud layers.

It is furthermore important that the fitness function is computationally cheap, as this calculation accounts for most of the computation time. The fitness functions can, therefore, better be defined on aggregated quantities (e.g. 2D LWP field, or average LWC profile) than on the entire 3D LWC field.

Finding the right set of statistical parameters to describe the cloud field is one off the main challenges. Redundancy of the various statistical parameters may be a problem, because due to measurement errors, this may lead to requirements that cannot be met. On the other hand, all important measures should be used. Which parameters are important may depend on the scientific question that the cloud fields are being used for.

STATISTICAL PARAMETERS

With radiative transfer problems in mind, one can think of the following statistical cloud parameters for the various cloud properties:

[†] Most of this chapter is based on the inspiring introduction to genetic algorithms by Melanie Mittchell (1996).

? *Cloud height*. A histogram with cloud base (and top) height or alternatively cloud base height and cloud width. Additionally a measure for the spatial correlations of the cloud boundary height is needed, a power spectrum or a correlation length. Cloud boundary height data can be retrieved from collocated lidar and radar measurements.
? *Horizontal distribution.* Number of vertical cloud boundaries per km and cloud cover for every horizontal layer. Combined lidar and radar measurements can provide this.
? *Horizontally variations.* Autocorrelation function of the LWP or alternatively structure functions or LWP variance spectra. These data can be taken from the microwave radiometers and in situ (PVM) measurements (plane and balloon).
? *Vertically variations.* Similar parameters should be used as in the previous horizontal case. The problem is that there is no measured data on the vertical variations in LWC. At small scales (compared to cloud width) we can assume isotropy and use horizontal measurements, at large scales (above some hundreds of meters) we can use profiles retrieved from radar and microwave radiometer data, but we do not have any measurements for intermediate scales. This gap may be filled by looking for relationships between the variations in the larger (or smaller scales) and at the intermediate scales in LES clouds.
? *Anisotropy*. The first three measures can easily be allowed to be different in the two horizontal directions. However, if anisotrophies are present in more than one direction, their consideration will be computationally expensive.
? *LWC profile*. The average LWC profile as measured in -situ or by a retrieval using radar and microwave radiometry.

The parameters mentioned above are intentioned as an example. Many other may turn out to be valuable.

The parameters mentioned above are statistical. It may also be possible to prescribe part of the solution non -statistically. For example, one could fix the two-dimensional cloud cover based on the cloud mask of CASI, or the cloud heights of an elevation scan made by a radar. The algorithm would then function as an interpolation/extrapolation method based on the measured statistics of the same cloud field.

TOTAL FITNESS

The fitness functions for each of these parameters will have to be combined into one measure of the clouds fitness. Simply adding up the fitness functions of all statistical parameters does not work. The fitness functions are not a linear function of the progress that has been made and they are all in a different numerical range. Therefore, there will typically be one or two fitness functions that respond significantly to changes in the LWC field. Only the statistical parameters belonging to those fitness functions are optimized, the others will often even degrade. It is thus important to make the fitness functions comparable.

A common way to make fitness functions comparable is by ranking: The fitness for each of the statistical parameters is sorted and the highest/best cloud field get rank 1, the second best 2, etc. These rankings are then averaged or summed to get the total fitness of the cloud field.

2.2 Selection and population size

The fittest clouds in the population should have a higher chance to reproduce themselves. Genetic variety, on the other hand, is important for find a good solution (the survival of the species). The right balance between selection pressure and genetic variety has to be found. It might be easy to get stuck in a local maximum during the search or the peaks could be narrow or even discontinuous. In this case, one would say that the fitness landscape is difficult, and one would choose a relatively wide search with a large reproductive fraction. If the fitness landscape is relatively easy, one can go quicker toward the solution and choose a deeper kind of search with a small reproductive fraction. The evolutionary algorithm can thus be easily tuned to fit to the difficulty of a specific problem.

Selection methods can be deterministic or stochastic. A deterministic method would be to let the best half reproduce and a stochastic method to choose a random numbers between one and the number of clouds and let a cloud reproduce, if its ranking is lower than its random number. The advantage of a stochastic approach is that a hopeful monster with a low fitness may be able to produce offspring in another (hopefully higher) peak of the fitness landscape. The disadvantage is that calculation time is wasted on clouds that are not very fit.

Another distinction between selection methods is whether they are local or global. In a global selection method, one individual competes with all others for resources, i.e. its fitness is compared to all others. If one positions the individuals on a 2D grid, one can compare a cloud just with the clouds in its neighbourhood. This local selection is closer to reality and allows for speciation, i.e. that more than one solution is pursued; more peaks in the fitness landscape are explored.

If the fitness landscape is difficult, a larger population is needed to find a good solution. The size of the population is limited by the amount of RAM available. A too large population, compared to the difficulty of the problem, means that the fitness of unnecessary many clouds has to be calculated.

2.3 Reproduction and mutation

Out of the fittest clouds of the old generation, a new generation is produced. The clouds are altered to introduce new variety and find better solutions. One can make an entire new generation or maintain the best cloud(s) unaltered (which is called *elitism*). The advantage of elitism is that no quality is lost (no degeneration), however, it does reduce the variety of the population and can thus slow the convergence.

The new generation can be made using sexual or asexual reproduction. In sexual reproduction, the key process is crossover, i.e., the genome of two parents is merged to produce the offspring. This is the preferred reproductive method in the Genetic *Algorithm* (GA) community. The advantage is that good solutions in different parts of the genome can be combined and the search is thus parallelized. Next to this, the GA community uses mutations, small point changes of the genome, to make sure that selection pressure does not reduce the variety of the population too much. However, these mutations are introduced at low frequencies, e.g. once every thousand generations. In the Genetic *Strategy* community, one normally uses mutations – asexual reproduction – exclusively.

2.4 Applications

As the search method can handle all quantitative measures and measurement scales, it can be used for much more application as the 3D radiative through clouds, for which it is used in the 4D-Clouds project.

The method can easily deal with all scales for which there are sufficient measurements, it can go down to the cm scale LWC variations which are measured by the in situ particle probes and up to tenths of km measured by the microwave radiometer (depending on how stationary the situation is, as we will have assume ergodicity). Due to limitations in computational power, not all of these scales can be resolved at the same time.

Possible applications could be:

? We could test satellite retrievals methods that may depend on cloud inhomogeniety.

? These 3D-clouds can be used (together with radiative transfer calculations) to determine which measurement strategy will have the smallest error boundaries when looking for enhanced absorption (a deficit in the radiative budget of a cloud field between the incoming radiation and the transmitted and reflected radiation).

? These 3D-clouds would give a beautiful opportunity to test the Klett inversion algorithm (which calculated the extinction profile of a cloud out of a backscatter profile for a single backscatter lidar) in cases with variable clouds and maybe derive a better algorithm.
? The 3D-cloud fields could be used to study what errors are made by comparing (temporally averaged) zenith measurements of clouds with (spatially averaged) in situ measured (almost horizontal) profiles.
? The structure of cloud field in many Numerical Prediction Models is not realist. It would be interesting to look if the models are at least able to keep these more realistic 3-dimentional clouds stable [Idea Felix Ament].

The algorithm can be used to reconstruct other difficult geophysical fields as well.

3 Example algorithm for cloud boundaries

3.1 Introduction

In this chapter, the above-mentioned ideas will be applied to a simple example, to show that searching for a cloud is possible. To simplify the problem the cloud does not have various LWC level yet: the pixels can either be cloudy or clear air. Furthermore, the statistical parameters are simple one-point statistics calculated from a cloud measured by a radar and most calculations are done in 2D.

3.2 Data and resolution

The example is based on an altostratus cloud, see figure 3.1, measured by the Delft Atmospheric Research Radar (DARR) on the 18th of April 1996 during the Dutch CLARA project (CLouds And Radiation, An intensive experimental study in the Netherlands). This measurement was cleaned of measurement artefacts, converted to binary precision, and to the right a part of clear air was added to simulate a partially covered cloud field. This field was resampled to 512 pixels horizontally and 256 vertically and used as input for the search algorithm, see figure 3.2.

The search starts with a cloud field of 1x1x1, and increases the resolution in horizontal or vertical direction each time the fit is sufficient. The input cloud mentioned above, is averaged to the resolution at which the search is performed and the statistical parameters for this resolution are calculated from this low-resolution cloud. The criterion for going to a higher resolution is based on the average value of the fitness functions. It corresponds to about one in 200 wrong pixels. However, as long as the convergence still progresses at a reasonable speed, it is better not to increase the resolution yet. Therefore, the previous criterion is only used if there has not been any progress for some generations. This number of generation is chosen such that a pixel has less than 1 percent chance of not being altered in this time span (which depends on the population size and the total number of pixels). If the cloud field fits the sta-

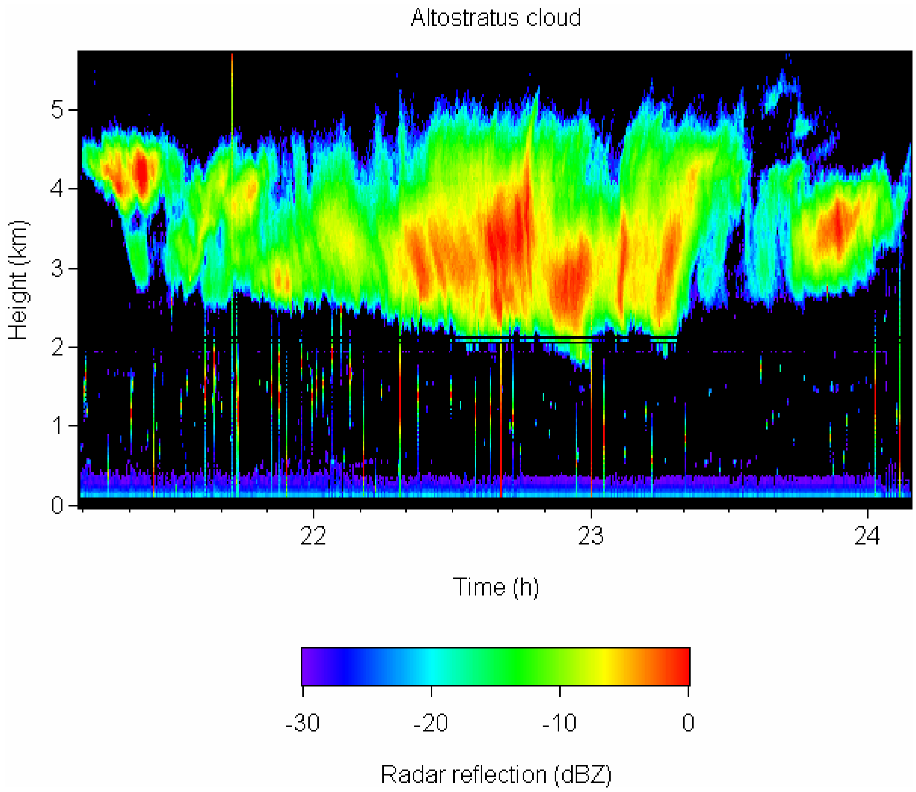

Figure 3.1. The radar measurement of an altocumulus cloud on which the example is based, see also Venema (2000).

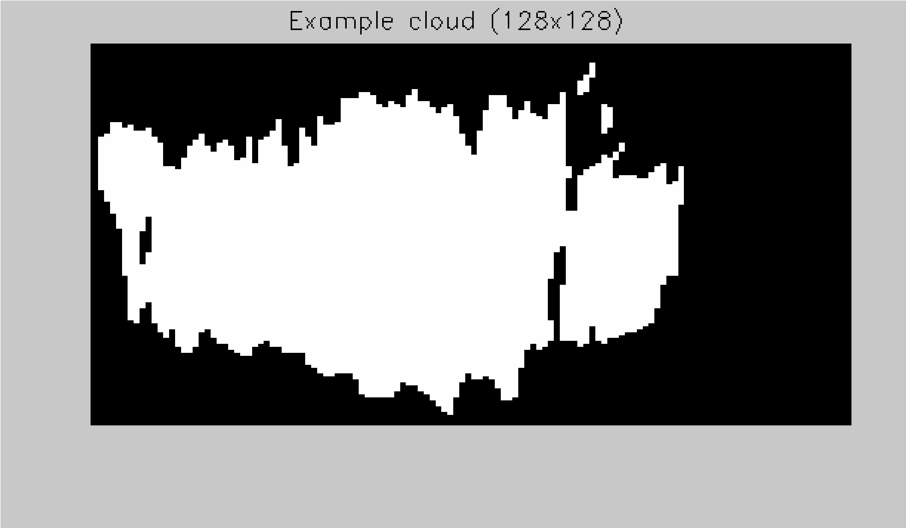

Figure 3.2. 2D cross section of a cloud field based on a radar measurement of an altocumulus cloud, see figure 3.1. This cross section is used at various resolutions to calculate the statistical cloud properties at this resolution.

tistics perfectly, the resolution is immediately increased.

Fitness functions

The fitness of the clouds depends on how well they fit to the statistical parameters. To reduce the computational effort, the statistical parameters are based on aggregated data, not directly on the 2D-cloud field itself. The aggregated data vectors are:

? Number of clouds layer for every vertical column
? Number of cloud tops for every horizontal row/height step
? Number of cloud bases for every horizontal row
? Cloud cover for every horizontal row
? Number of cloud edges (assuming periodic boundary conditions)

From the vector with the number of cloud layers, a histogram is calculated and used as statistical parameter. All other aggregated data vectors are used directly as statistical parameters. Histograms are preferred as statistic as they contain information on all moments and are very flexible. For example, compared to an average cloud height, the histogram of cloud base height also brings information on the variations in cloud height and possible other cloud layers. At the moment, all statistical parameters are one-point cloud boundary statistics. Although, one could see the number of cloud edges and cloud layers as a primitive two-point statistic.

The fitness functions of a cloud are calculated from the difference between its histogram and the one of the measured data. However, the fitness should not simply be the average difference between the two histograms. Figure 3.4 shows what can happen in such a case. In this earlier calculation, this simple difference in the histogram for the number of layers was used a fitness function. It resulted in some columns with an enormous number of layers. This happens as removing a cloudy pixel, can improve one of the other fitness functions and increasing the number of layers from, e.g., 4 to 5 is equally bad, and does not decrease the fitness of the cloud with respect to its number of layers.

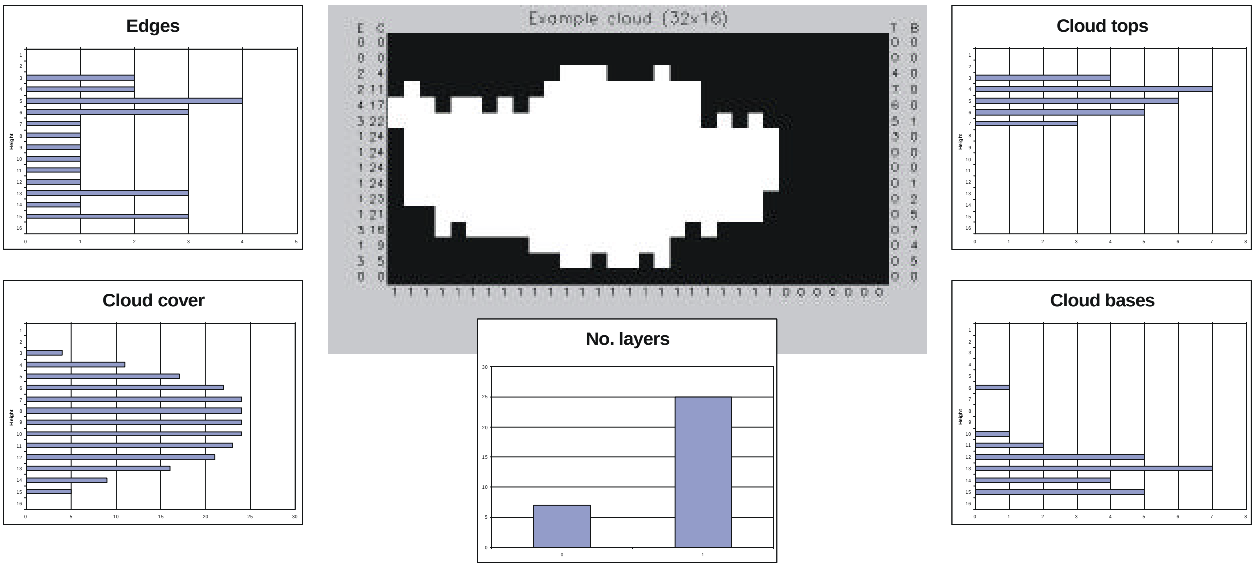

Figure 3.3. The input cloud, sampled down to 32x16 pixels and its statistical parameters describing the cloud boundaries. In this example, histograms of five statistics are used: The number of cloud layers (bottom middle), and vertical profiles of the number of edges or cloudy parts, cloud cover, cloud top height, and cloud base height.

To avoid this problem, the fitness functions are using weighted histogram comparisons, where, e.g., three cloud layers are worse than two cloud layers, when search for a single layer cloud.

3.3 *Selection and population size*

The search algorithm uses elitism, i.e., the best-ranked cloud – elite cloud – in the population is transferred to the next generation without mutations. The population size was 100 clouds. On average 8.5 clouds reproduce themselves, including the elite cloud. The chance to reproduce is a linear function of the ranking. The cloud with ranking one (elite) has 100 % chance to reproduce itself, the 17th cloud (and lower ones) 0 %.

The population size and reproductive fraction (and selection method) should still be optimized. Experience with earlier versions of the algorithm showed that the convergence speed depends strongly on these model parameters.

Figure 3.4. In this case, the fitness function for the number of layers was the average difference between the clouds histogram of number of layers and the measured histogram. This can results in a highly fragmented cloudy column, which is especially very clear at the two red arrows.

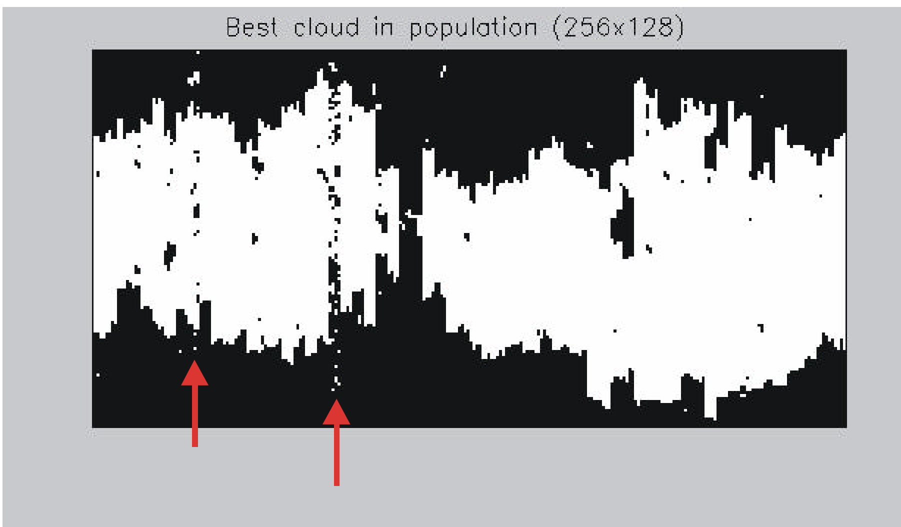

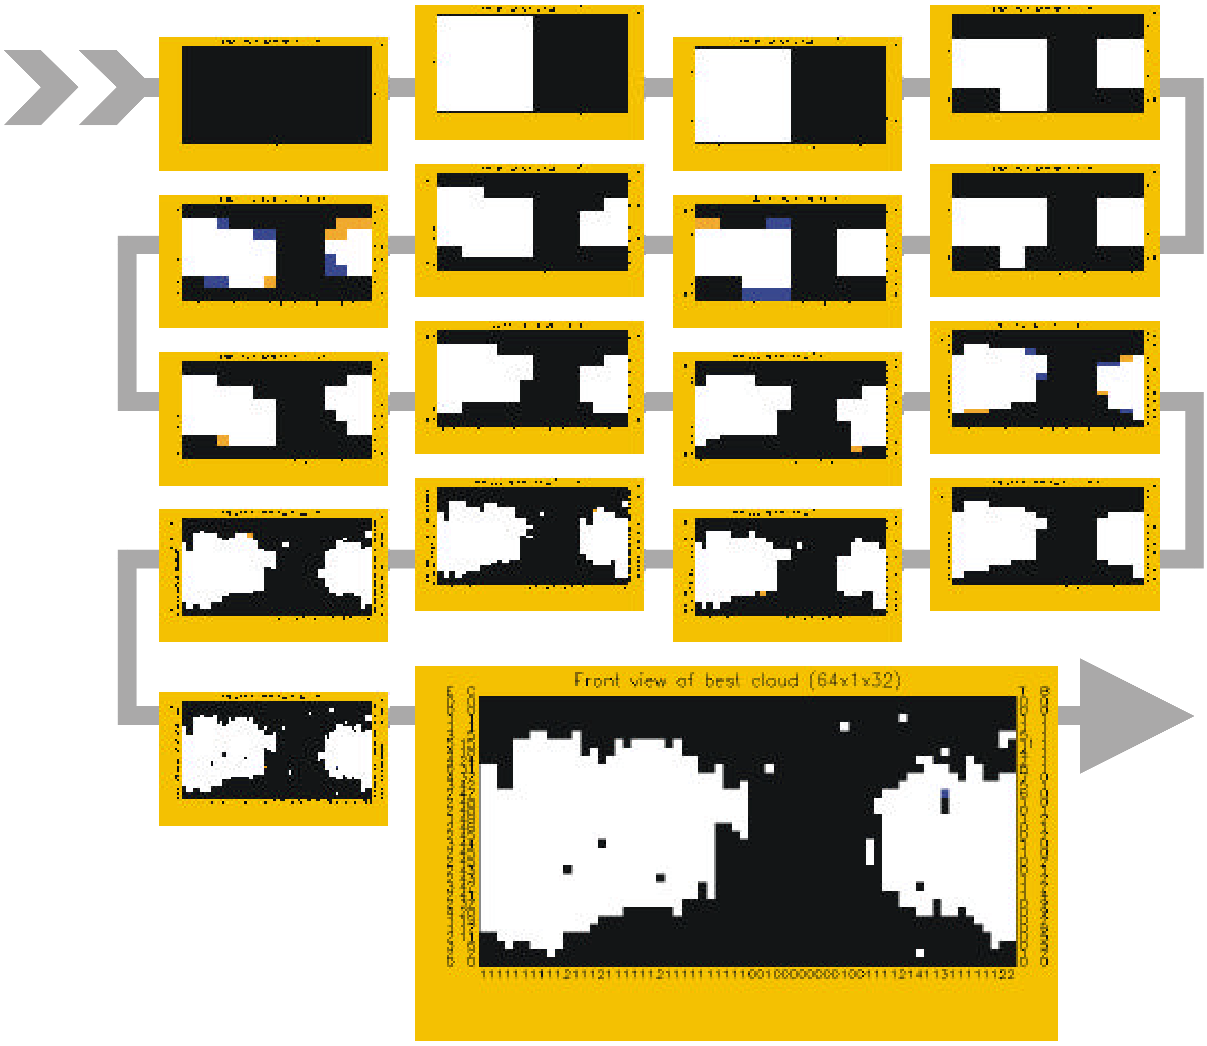

Figure 3.5. The growing of the cloud into the final solution at high resolution. At each resolution, the best cloud in one or some generations is shown. The white pixels are cloudy; the black ones clear air. The yellow and blue pixels mark the difference compared to the best cloud in the previous generation (this previous cloud is not always shown). The yellow pixels became cloudy in this generation; the blue ones just became clear air.

3.4 Reproduction and mutation

For each new cloud, one of the reproductive clouds is chosen at random as its parent. This cloud is then copied and mutated. In this search algorithm, only mutations are used; there is no cross-over. The advantage of mutations is that one can calculate from the local changes, i.e., with relatively little calculations, how the aggregated data vectors change. In case of crossover, one will have to recalculate (almost all) the aggregated data vectors using the full 3-dimensional LWC matrix.

Three ways of mutation have been implemented:

? Mutate a random pixel, i.e. change cloud into clear-air or visa versa

? Mutate a pixel next to cloud boundary and

? Exchange two columns.

Their frequency of use is adaptable, but no optimization has been made yet.

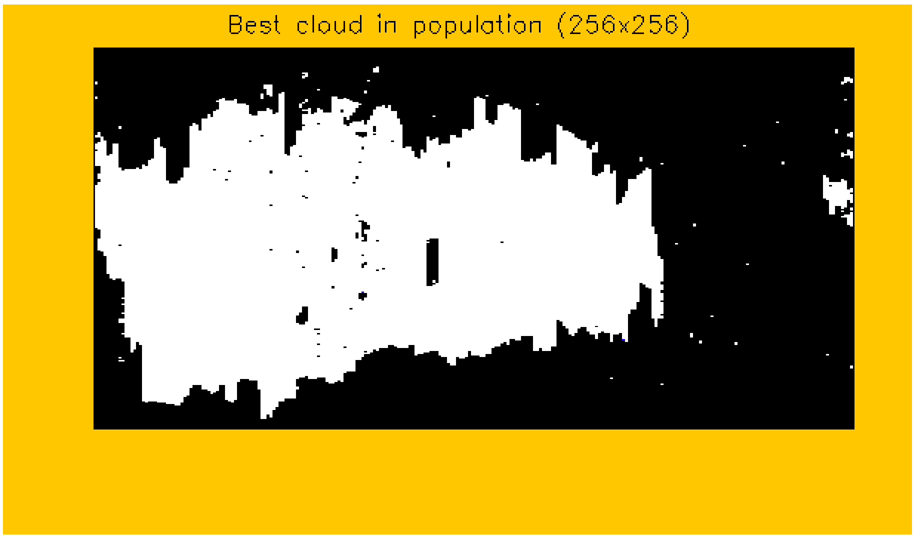

Figure 3.6. Solution of one of the searches for a cloud with the same five statistical parameters concerning its cloud boundaries as the strato-cumulus cloud measured by radar from figure 3.2.

3.5 Results and discussion

In two dimensions the search results in a cloud which is similar to the cloud the radar measured. The beginning of growing of one such cloud can be seen in figure 3.5, the final cloud in figure 3.6. This cloud should be compared to the template cloud in figure 3.2. Of course, only with respect to the five statistical cloud boundary parameters. There are two discrepancies obvious. The structure is still ragged and at pixel-level there are still many small errors.

That the structure is stilled raged is because no two-point statistics is involved yet. In fact, considering this the structure is already very good, which is probably caused by the way the

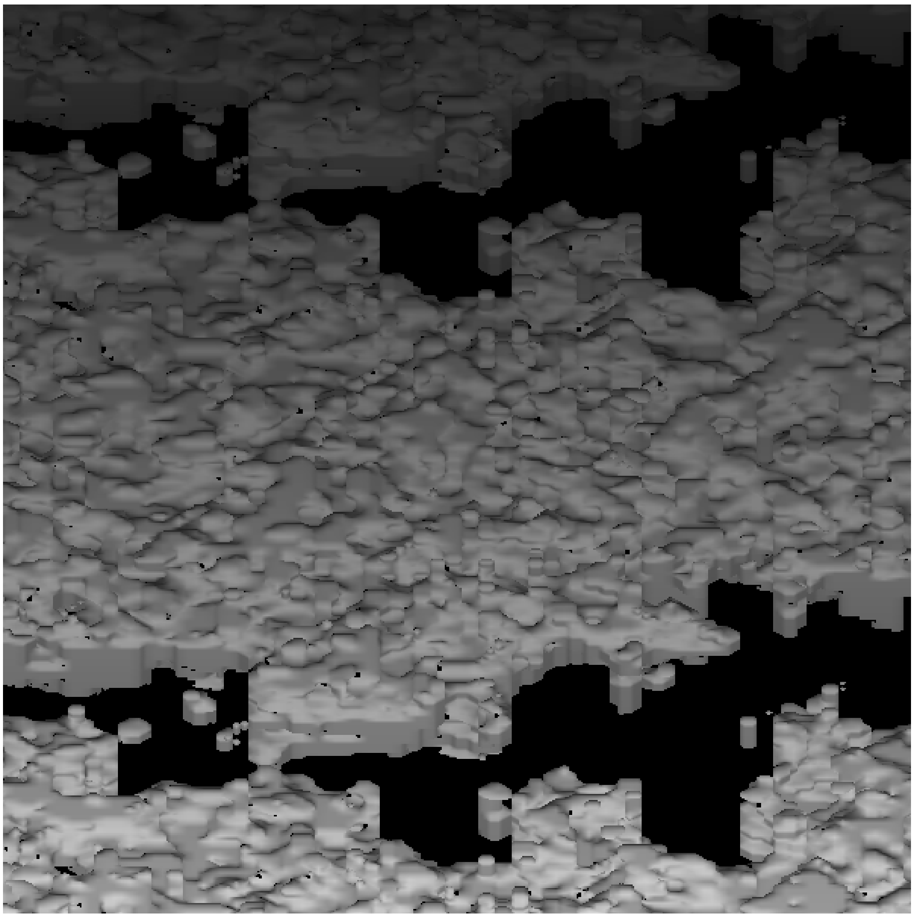

Figure 3.7. The result of for a three dimensional cloud using the statistics from the two-dimensional altostratus cloud and assuming anisotropy.

Figure 3.8. The scaling behaviour of the search algorithm with respect to the number of pixels for five different runs.

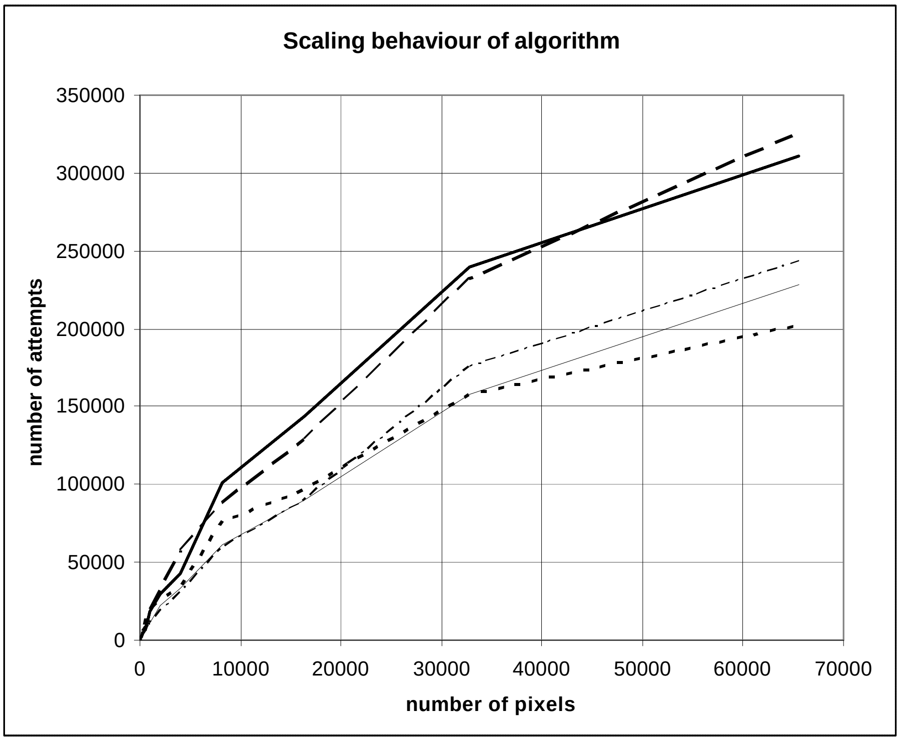

search goes from low resolution to higher ones. The structure is also partially caused by to need to satisfy the statistics in the horizontal and vertical direction simultaneously.

The pixel errors in the searched cloud are due to insufficient convergence of the solution. These errors first arrive at higher resolutions. Maybe the criterion for going to higher resolutions is not strict enough at these resolutions. Thus, it may be necessary to make a one-step deep search, on every pixel before going to the next resolution. With the current criterion for going to the next resolution, still about 1 percent of the pixels have not been tested if their change could improve the solution.

The small cloudy parts at the right of figure 3.6 are no failure; remember that the boundary conditions are periodic in the horizontal directions.

A three-dimensional calculation is shown in figure 3.7. This cloud was calculated up to a resolution of 64x64x32. At this resolution, the RAM of the PC was not sufficient anymore and the calculations became very slow. That why the calculation was stopped, the cloud median filtered and resampled back to the last useful resolution of 64x64x16 pixels.

To be able to calculate a three-dimensional cloud from two-dimensional data, anisotropy had to be assumed. Again, the horizontal boundary conditions are periodic, that is why the two gaps in the cloud field look alike. The displaying program draws shiny curves that fit to the outer parts of the clouds. This makes the clouds look somewhat like plastic. Especially the scaling behaviour of the horizontal cloud boundaries is impressive, considering that no two-point statistics is used yet.

3.6 Convergence speed

On a desktop PC, the calculation of the 2D-cloud up to the resolution of 256x1x256 pixels took about 1 hour. The program was written in IDL and was running on a 700 MHz Desktop PC with 256 Mb RAM. This research version has not been optimized for speed yet and much time is wasted on logging and plotting intermediate results. Except for this logging, the calculation of the fitness functions takes almost all of the calculation time, which is typical for search/optimization algorithms.

The search algorithm found this cloud after 2-3 x10^5 attempts, i.e. 2000 to 3000 generations with 100 clouds. The calculation of the 3D-cloud up to the resolution of 64x64x16 pixels took about 3 hours. Especially encouraging is, that the amount of attempts is linear function of the number of pixels, see figure 3.8. This means that it is relatively easy to make calculations on larger arrays, provided the computer memory is large enough. The calculation time will grow faster with the number of pixels, as the calculations on the larger arrays will take more time per attempt. Note, that with another criterion for going to a next resolution this curve may look differently.

That the search algorithm scales so well with the number of pixels is probably due to the a-priory knowledge build into the algorithm, especially the gradual increase in resolution.

4 Conclusions and outlook

The genetic algorithm is able to generate a 3D-cloud field that has statistical properties close to those measured. Its intuitive line of thinking allows for a lot of a-priory knowledge to be included in the search algorithm, and thus to speed up the search. The first experience with a simple version of the algorithm shows that the calculations can be done in a reasonable time and that the scaling behaviour with respect to the problem size is excellent.

During the development two major ideas have come forward, that one could say, have made the search approach possible: The ranking of the fitness functions to arrive at a total fitness and the solution of the problem at a low resolution before going to higher resolutions.

The next step will be to make a search algorithm with a full 3D-LWC field, instead of just cloud existence, see section 4.1. This algorithm will have to be fed with statistics from various instruments and will include two point statistics. Still a lot can and has to be done to improve the convergence speed of the algorithm, see section 4.2. Especially, many procedures and tuning variables have to be optimized and new mutation methods have to be implemented. This complete model will than have to be validated, section 4.3.

The model is not only useful for generating clouds, it also helps one to think about what kind of statistics and measurements are needed to get a full description of a structured cloud field. At the same time, it can function as a test bed for such questions, see section 4.4.

4.1 The full search algorithm

The full search algorithm should have a full 3D-LWC field. This field should have a resolution of 128x128x64 pixels and 256 LWC values, to be useful for radiative transfer calculations [Pers. comm., T. Trautmann, 2002].

This algorithm will have to be fed with statistics from various instruments. The smallest set of instruments is a lidar and a radar for the cloud boundary height statistics and for height resolved cloud cover statistics and a microwave radiometer for the LWP statistics.

Additionally, an infrared radiometer can be used to get better (i.e. more sensitive) information on total column cloud cover. The CASI, Compact Airborne Spectrographic Imager, from the University of Berlin has a 2D-cloud map as one of its products. This can be used for two-point statistics on the cloud cover. This can provide more data than the infrared radiometer and it will be possible to take anisotropies in the cloud field into account.

The in situ measurements of LWC can be used to include the smaller scale variations. In situ measurements can also provide information on the LWC profile. For some applications, it may be useful go have a corresponding 3D-field with the droplet diameter, humidity or temperature. In situ measurements are probably the best information source for these parameters.

The full search algorithm will at least have to include two-point statistics; an autocorrelation function, a structure function or a power spectrum. Most likely this will be implemented as a power spectrum comparison on the 2D-LWP field of the 3D-LWC field, to reduce the calculation time. As the power spectrum is not a complete description of the LWP structure – see section 1.3 – we may have to find additional statistical parameters.

As the final matrix contains a factor 16 more pixels and 256 LWC values instead of 2, one can expect that the calculations will take at least factor 2024 more attempts. Additionally, new fitness functions will have to be calculated and the time to calculate one fitness function increase with the matrix size. Therefore, an increase in calculation time of a factor four thousand can be expected for the full search algorithm.

A decrease in the calculation time in the same order of magnitude should still be achievable. The calculations are made at the moment on a simple desktop PC (if the problem can be efficiently parallelized, the Japanese Earth Simulator would take a few minutes to calculate one cloud). The code still needs to be optimized for speed, the current version was made to play and understand the problems. Still a lot of time is wasted in logging intermediate results. An operational version should be written in C instead of IDL. Most importantly, there are still many opportunities to increase the convergence speed of the algorithm itself, see next section.

4.2 Speeding convergence of the search algorithm

There are still many possibilities to improve the convergence speed of the algorithm itself. Improvements can be made in the mutations used, by the optimization of settings, by trying different procedures (reproduction, selection, elitism, etc.) and in the quality criterion for going to the next resolution.

MUTATIONS

New mutation methods can be implemented and tested for improved convergence:
? Mutate a pixel close to the previous one.
? Flip two pixels next to each other; exchange their values.
? Change LWC to zero.
? Change LWC to the average of the neighbours.
? Add/subtract LWC.
? Multiply LWC.
? Moving rows (horizontally) or columns (vertically).

OPTIMIZATION

Various settings in the algorithm still need to be optimized, e.g., the population size, reproductive fraction, and the frequency of the various mutation methods. Experience with an earlier version of the search algorithm showed that much can be gain by such optimizations.

In addition, various procedures are likely to be far from optimum in the present version. For example, the selection method is now implemented as partially random and partially deterministic. Maybe a full deterministic approach (let the best *x* clouds reproduce) or a full random approach (all clouds have some chance to reproduce) would be better. It still has to be examined, whether the elitism brings something. If this will be the case, it would also be possible to increase the number of clouds that go to the next generation without any alterations. Making the number of children a function of a clouds ranking may also accelerate the search. The fitness functions may still be coded differently, to reduce their calculation time and/or to improve the convergence speed.

We have to introduce a quality measure, which quantitatively says how well the fit is. With this measure, we have a better criterion for going to a next resolution level. As problems that have not been solved on a low resolution, take much more effort to repair on a higher resolution, it is important that the solution be of a very high quality before going to the next level. On the other hand, setting this quality level too high will waste a lot a of computer resources or even cause the algorithm to get stuck in an infinite loop. This quality measure can be used as a first quality check of the final solution as well.

A larger change, which could have large potential, would be to describe the LWC field with harmonic functions (like in a JPEG-picture, or in speech coding, but then 3D). One could, e.g., divide the total matrix into small matrices of 16x16x16, and describe them with 15 harmonic parameters, which would save a factor 300 in memory and in convergence speed, at the expense of more calculation time for the conversion of the harmonic function to this small LWC matrix. A problem would be how to make sure that the LWC field is continuous at the boundaries of the sub matrices.

The method used to generate fractal time series, discussed in section 1.3, can be used to generate a good quality first guess when the resolution of the LWC field is increased. This way at least one statistical parameters, the power spectrum, is already almost converged.

4.3 Validation

The quality of the clouds that are made with this search algorithm can be validated by making calculations on LES-clouds. From these LES-cloud fields the statistics can be calculated and cloud fields can be searched that have the same statistical properties. The measure that is being studied, e.g. the downward solar flux at the surface, can then be calculated for both types of cloud fields and statistically compared. Whether the searched cloud field is good enough may depend on the problem that is being studied.

After a favorable comparison wit LES clouds, a set of cloud fields can be searched with the measured micro and macrophysical cloud properties. These fields can be used for empirical studies. Following the previous example, one could compare the histograms and correlation length of the calculated and measured solar fluxes at the surface.

Interesting about this method is that, in addition, one becomes insight into the important questions: which statistical parameters are required to describe a cloud field and what is the minimal set of measurements one needs.

The subdivision of cloud structure parameters into various statistics also made obvious that we do not have any measurements on the vertical variations of the LWC. For the ramps flown with in situ probes, we can only derive an average LWC profile. The variations in LWC that are measured by the probes will be dominated by the horizontal cloud structure. Measurements of the temporal development of clouds are also failing in the measurement setup of 4D-clouds. As far as I know such measurements have only been made of cumulonimbus to study rain initiation. Such measurements are difficult as the wind drift the clouds away. Measurements with in situ probes hanging below a zeppelin or a special hot air balloon on a tether attached to a winch would be valuable.

This validation method can also be used to study which statistical cloud properties are important (for a certain question). You can specifically remove and add certain statistical parameters that are more or less independent of the rest. For example, you can fix the cloud top height or use the measured cloud top heights; you can use the LWP histogram with or without additional power spectrum; you can use an LWP power spectrum with an average LWP or a full histogram; you can remove LWP variations above or below a certain length scale.

5 References

Marshak, A. and Davis A. and Cahalan, R. and Wiscombe, W. Bounded cascade models as nonstationary multifractals, Physical Review E, **49**, pp. 55-69, 1994.

Mitchell, M. An introduction to genetic algorithms. MIT Press, cambridge, ISBN 0-262-13316-4, 1996.

Venema, V. Cloud measurements with radar. PhD thesis, 13 November 2000.