

Determination of the orbital parameters of the M51 system using a genetic algorithm

M. Wahde¹ and K.J. Donner²

¹ Div. of Mechatronics, Chalmers University of Technology, SE-412 96 Göteborg, Sweden

² Observatory, P.O. Box 14, FIN-00014 University of Helsinki, Finland

Received date / Accepted date

Abstract. Using an efficient method based on genetic algorithms, we have searched for an optimum set of parameters for the interaction between NGC 5194 (M51) and NGC 5195. The preferred model is one where the time of closest passage was 900 Myr ago. The orbit of NGC 5195 is almost perpendicular to the plane of the sky, and is slightly hyperbolic. The results from the search among non-self-gravitating models are confirmed by a fully self-gravitating simulation. Our model, which is limited to the outer regions of NGC 5194, reproduces the density distribution of the giant H I arm of NGC 5194 accurately, but is not able to find a perfect representation of the velocity field. We argue that this is caused by deficiencies in the initial disc model used for NGC 5194, rather than a failure of our search method. In general, we argue that genetic algorithms are ideally suitable for investigations of tidally interacting galaxies, where a large number of parameters need to be determined, and the goodness of fit for models contains a significant irregular component, with several local optima.

Key words. galaxies: interactions – galaxies: individual: M51

1. Introduction

Among spiral galaxies M 51 holds a special position. Apart from being the first galaxy where spiral structure was detected, it seems to be a shining example of how a spiral pattern can be generated by the tidal interaction with a passing satellite. However, the apparent early successes of Toomre & Toomre (1972) in modelling the interaction, supported by observational arguments of Tully (1974), have evaporated as new and more detailed observations (Rots et al. 1990) have become available. At present there exists no fully acceptable tidal model of the M 51 system (Barnes 1998).

Determining the orbits in interacting systems of galaxies is a problem with a large number of unknown parameters. Usually, the modelling has been based partly on guesswork, where a most likely range of orbital parameters has been chosen based on the general appearance of the tidal features. More systematic searches, combined with objective evaluations of the goodness of fit of the models, have been limited to only a few parameters.

In a recent paper (Wahde 1998) one of us has shown that a genetic algorithm (GA) provides an effective and economical method for finding the best tidal interaction model in the large parameter space of an interacting pair

of galaxies. The advantage of the GA method is that it requires minimal *a priori* restrictions on the initial assumptions. In this paper we test the method in practice by applying it to the NGC 5194/5 system.

A perturber passing a galaxy on a direct orbit in the disc plane will produce a tidal bridge with the perturber at one end. Since NGC 5195 appears to be placed at the end of one of the arms, it is a natural guess that the M 51 system should be an example of this type of interaction. The model of Howard & Byrd (1990) was based on such a picture. However, this viewpoint has become untenable with the publication of detailed VLA maps by Rots et al. (1990), revealing an extensive, apparently tidal, spiral arm at a great distance from the centre of M 51. Such a large arm requires a long time to form, indicating that the closest passage must have been at a much more remote time than had been previously presumed, and that the situation of NGC 5195 close to a spiral arm is a coincidence. The absence of any visible tidal damage to NGC 5195 (Thronson et al. 1991, Spillar et al. 1992) is also an indication that the two galaxies are not strongly interacting at present. These considerations put the problem in a whole new light, rendering all earlier models obsolete.

Attempts to incorporate the outer arm in the models have been made by Hernquist (1990), Antonioletti & Nelson (1999), and, in particular, by Engström and Athanassoula (1991; Engström 1992). None of these in-

Send offprint requests to: M. Wahde, e-mail: mwahde@me.chalmers.se

vestigations settled on any specific final description of the interaction. An additional purpose here is to produce a model that can account for the more recent observations, which is optimal according to some reasonably well-defined criteria.

Our paper thus has the dual purpose of making the case that genetic algorithms are ideally suited for orbital determination in interacting galaxy pairs, and of establishing a viable dynamical model for the M51 system. The central principles of genetic algorithms and their application to the problem at hand are presented in Section 2. The application to the specific case of M51 is discussed in Section 3, where it is shown that an optimum model reproducing the main observed features can be found. In Section 4 this optimum model is used as input for a self-gravitating simulation confirming that the non-self-gravitating results are realistic. In Section 5 we discuss the usefulness of the GA method for determining the parameters of interacting galaxy pairs, and the quality of our optimum model.

2. Method

The purpose of this section is to describe briefly an efficient method for determining the orbital parameters of pairs of interacting galaxies, given position (photometric) and velocity (spectroscopic) data. In general, the problem of determining orbital parameters is a very challenging one, due to the many unknown parameters involved in the problem. It is easy to show that an unbiased search (i.e. a search with equal spacing between successive parameter values) of the entire search space is not an option when the dimensionality of the search space is larger than five or so. Thus, without an efficient method for carrying out the search one must either settle for a very coarse-grained survey of the parameter space, or restrict the search space by fixing the values of some of the parameters. It is clear that neither of these two approaches is optimal in a general case, and that the bias introduced by either carrying out a coarse-grained search or strongly restricting the search space may prevent one from finding an acceptable set of orbital parameters.

In recent years, several new methods for search and optimization inspired by natural evolution, e.g. genetic algorithms (hereafter: GAs) and evolution strategies, have been introduced, and have been widely used in several fields (see Mitchell 1996). The use of GAs in astrophysics was pioneered by Charbonneau (1995), and Wahde (1998, 1999) first used GAs in connection with the search for orbits of interacting galaxies. Theis & Harfst (2000) applied a method similar to that used by Wahde, using a Cray supercomputer to speed up the search.

Genetic algorithms have the advantages of carrying out a parallel search through the search space, and are therefore well suited for problems with “rugged fitness landscapes”, i.e. problems for which there exists several suboptimal solutions. An additional advantage is that the amount of subjective bias is minimal. A very brief descrip-

tion of GAs, followed by an introduction to the method for orbital parameter determination will now be given. For a more complete description, see Wahde (1998).

2.1. Genetic algorithms

In a standard GA, the variables of the problem are encoded in strings of digits referred to as *chromosomes*. Initially, a *population* of *individuals*, each associated with one such string, is generated by assigning random values to all the locations (called *genes*) along the strings. Then, for each individual, the variables are decoded from the chromosome, the relevant computation (which obviously varies from case to case) is carried out, and the *fitness* of the *individual* is evaluated.

The fitness values should be such that those individuals which are close to achieving the goal determined by the user obtain higher fitnesses than those which are far from the goal. When all the individuals in the first *generation* have been evaluated, the second generation is formed by first selecting *parent strings* such that individuals with high fitness have a greater chance of being selected than those with low fitness, and then combining the genetic material contained in their chromosomes to form new strings, and, finally, allowing a small degree of *mutation* (i.e. random variation) of the newly formed chromosomes. The new chromosomes constitute the second generation, which is evaluated in the same way as the first. The process is repeated until a satisfactory solution has been found.

In any error-minimizing search, the optimum is defined only with respect to the error measure that is used to guide the search. The definition of a suitable error measure in multi-criteria optimization is a difficult issue. In this case, we wish to minimize the deviation between the observed (light) density distribution and the density distribution obtained in the simulations, while simultaneously minimizing the velocity deviation. Clearly, these two aims should not be in conflict with each other *provided that the model used in the simulation constitutes an accurate description of the undisturbed (i.e. pre-encounter) galaxies*. This, however, is unlikely to be the case, in view of the many simplifications that need to be made in order to cope with the very large number of simulations needed for the determination of an orbit. However, even if it were possible to define a model with a high degree of realism, it would probably involve the setting of a number of more or less ad hoc parameters. Here, we wish to minimize the number of such parameters, and only define our model from those properties that can be directly observed. Thus, the method is aimed at finding, in a short time, a good set of parameters which can then be used as input parameters for more advanced simulation methods involving interstellar gas, self-gravity, non-axisymmetric disturbances etc.




Fig. 1. The left panel shows the observed density distribution in the HI-arm of M51, and the right panel shows the coarser representation used in the GA simulations. The crosses represent the template which is used to block out the main body of the galaxy

2.2. Determination of orbital parameters using genetic algorithms

In the method introduced by Wahde (1998), the parameters encoded in the chromosomes are: the masses of the two galaxies, their relative position along the line of sight, their relative velocities in the plane of the sky, the direction of rotation of the two galaxies, as well as their inclinations and position angles. Thus, a total of 11 parameters need to be found, unless the smaller of the two galaxies is represented only by a point particle, in which case the total number of parameters is 7, if also the direction of rotation of the main galaxy is assumed to be known.

For each individual, the values of the unknown parameters are obtained from the chromosome, and the corresponding orbital parameters are computed. Then two point particles representing the galaxies are integrated backwards in time until they are sufficiently separated. At this point, a disc of particles is added to each of the galaxies (or only to the main galaxy, if the perturbing galaxy is represented by a point mass), and the system is integrated forward in time until the time of the observation is reached. Then, the density and velocity fields are compared with the observational data, and a fitness value is assigned based on the similarity between output from the simulation and the observational data.

For its operation, the method requires (light) density data and preferably also radial velocity data for the tidal structures caused by the interaction between the two galaxies. The data grids need to be sufficiently fine-grained to resolve the main tidal features.

In Wahde (1998) the difference between the result of a given simulation and the observational data was defined using essentially the absolute deviation between observed and simulated values. Although this definition of fitness was shown to be effective in the simulated models of Wahde (1998), it is only one among many possibilities. After trying a number of alternatives we have chosen in this paper to measure the distance between the observed and the simulated density data using a distance measure inspired by the Kullback–Leibler directed divergence

$$\delta = \sum_{i,j} (m_{i,j}^{\text{obs}} + m_{\epsilon}) \left| \ln \frac{m_{i,j} + m_{\epsilon}}{m_{i,j}^{\text{obs}} + m_{\epsilon}} \right|, \quad (1)$$

where $m_{i,j}$ denotes the mass in grid cell (i, j) obtained from the simulation, and $m_{i,j}^{\text{obs}}$ denotes the same quantity obtained from the observational data, assuming a constant mass-to-light ratio. This distance is an average of the relative density error weighted with the observed density. The logarithmic scale for the relative errors in density is motivated by the very large density variations in the system. The weighting factor m_{ϵ} is necessary in order to avoid problems in regions where the observed mass is zero. For smaller m_{ϵ} , models that place many particles in the observationally empty regions are more severely discouraged.

For the velocities we have considered the corresponding distance

$$\delta_v = \sum_{i,j} (m_{i,j}^{\text{obs}} + m_{\epsilon}) \left| \ln \frac{\bar{v}_{i,j}}{\bar{v}_{i,j}^{\text{obs}}} \right|, \quad (2)$$

where $\bar{v}_{i,j}$ and $\bar{v}_{i,j}^{\text{obs}}$ denote the average radial velocity in grid cell (i, j) from the simulation and the observation,




Fig. 2. Three projections of the best-fitting model found for the NGC 5194/95 system. The plane of the sky corresponds to the xy -plane (bottom plane in the figure). The hole in the centre of the galaxy is an artefact of the hole in our initial H I density distribution (see Sect. 3.1)

respectively. In equation (2) we rather arbitrarily use the actual observed velocities.

The total distance between the observations and simulations can be defined as a weighted combination of the expressions (1) and (2). However, our attempts with deviation measures which more strongly emphasized the optimization of the velocity field were not successful, and (perhaps predictably) tended to neglect the optimization of the density distribution. Thus, mostly we have simply used the fitness measure $1/(1 + \delta)$.

As described by Wahde (1998), the algorithm does not require data covering the entire interacting system. Instead, the central regions of the galaxies can be blocked out, and thus ignored by the genetic algorithm. Hence, one can find a model of an extended tidal tail without needing to simultaneously make a detailed model of the often very complicated central regions of the galaxies.

In Fig. 1, the grid and template are shown together with our representation of the data of Rots et al. (1990).

3. Best non-selfgravitating model

3.1. Observational constraints

In our application of the GA we aim to find a tidal model for the outer H I arm that is consistent with the observed relative positions and velocities of NGC 5194/5. An advantage of the method used here is that in principle not all observational constraints need to be used. In particular, since we are only modelling the outer features, the observational data relating to the central galaxy (mass, inclination, PA) may be incorporated into the assumptions of our models. Alternatively, these observations can be used to check for consistency once the best model has been found.

Our strict observational constraints will be the separation of the galaxies in the sky and their line-of-sight velocity difference. Thus we require that the galaxies should be separated by the observed position differences $\Delta\alpha = -6^{\circ}.2$, $\Delta\delta = 4'22''$. We use a Cartesian coordinate system with the x -axis to the west, y to the north, and z along the line




Fig. 3. A schematic illustration of the orbit of the perturbing galaxy (NGC 5195) relative to the main galaxy (NGC 5194). The interval between successive points on the orbit corresponds to 52.5 Myr. The direction towards an observer on Earth is indicated by the arrow

of sight, with the positive direction towards the observer. We use units such that the gravitational constant is 1, the unit of length is 1 kpc, the unit of time is equal to 1.05 Myr, and the unit of mass is equal to $2 \times 10^{11} M_{\odot}$. The unit of velocity then equals 931 km s^{-1} .

In order to conform to previous studies, we adopt the traditional value 9.6 Mpc of Sandage & Tammann (1975) for the distance, although a more recent estimate is 8.4 Mpc (Feldmeier et al. 1997). Using the adopted distance l' corresponds to 2.79 kpc, and the position of NGC 5195 relative to the main galaxy is $x = -2.79$, $y = 12.1$. Furthermore, following Schweizer (1977), we adopt the value 130 km s^{-1} for the line-of-sight velocity difference between the galaxies.

In order to minimize the search space for the GA, we restrict the mass of M51 to an interval consistent with the observed rotational velocities. The most extended rotation curve was determined from H I observations by Tilanus & Allen (1991). For the GA runs we extrapolate this to larger distances using a Keplerian law. This has the advantage of avoiding the vexed problem of how to keep the centre of mass at a constant position in the later non-self-gravitating simulation. Fitting a Keplerian rotation curve to the data of (Tilanus & Allen 1991) using the method of least-squares, a mass of 0.31, corresponding to 6.2×10^{10} solar masses, was obtained. Based on this estimate, we restricted the mass for NGC 5194 in the simulations to the interval [0.25, 0.50].

The mass ratio for the two galaxies in the M51 system has been estimated by Schweizer (1977) who finds that the mass of NGC 5195 is between 30 % and 50 % of that of M51. We allowed the mass ratio to vary in the range [0.20, 0.70].

In principle, the inclination and position angle of the M51 disc could also be constrained to the observed values. However, their actual values are rather uncertain, since the inclination is small. As noted by Garcia-Gomez & Athanassoula (1991), previous determinations of inclination have yielded values in the range $[15^{\circ}, 40^{\circ}]$, whereas estimates of the position angle fall into two groups, one with values around $\text{PA} = -10^{\circ}$ and one in the neighbourhood of $\text{PA} = 30^{\circ}$. Therefore we allow them to be variable in the optimization, initially only forcing them to lie in an interval around the values $i = 20^{\circ}$, $\text{PA} = -10^{\circ}$ given by Tully (1974), broad enough to include all published values. However, it was found that the best results were obtained in a more narrow interval, and in the final GA run, the inclination was constrained to the interval $[10^{\circ}, 30^{\circ}]$ and the position angle to the interval $[0^{\circ}, 25^{\circ}]$. From the fitness variation shown in Fig. 7, it is clear that the optimum falls well within these intervals.

The distance Δz along the line of sight between NGC 5194 and NGC 5195 is difficult to estimate. In the simulations, we have therefore allowed it to vary in a rather large interval, 0 to 400 length units (kpc). Finally, the relative velocities in the sky plane, Δv_x and Δv_y , were allowed to vary in the interval $[-0.3, 0.3]$ velocity units.

The simulations employed by the GA are non-self-gravitating, and therefore we are free to choose an arbitrary density distribution for the particles. We have used a distribution loosely based on the H I data of Tilanus & Allen (1991). Since we are mainly interested in the large H I arm, we leave the region $r < 3$ kpc empty. For $r > 3$ kpc, the density was taken to be constant out to $r = 7$ kpc, and outside this radius we used the density distribution

$$\sigma \propto e^{-\frac{(r-7)}{\delta}}. \quad (3)$$

We also assumed an outer cutoff of the disc at $r = 23.5$ kpc, corresponding to five scale radii of the optical disc. Although we did not include the cutoff radius as a free parameter in the GA, it is rather crucial, and a satisfactory




Fig. 4. The velocity field obtained from the best model. Darker particle colours correspond to larger velocities. Note that the strength of the counterarm is artificially enhanced by the size of the squares used to represent the velocity field.

model could only be found for values near the one that was used.

3.2. Results

In the GA run, a total of 5 000 particles were used for each simulation. Each simulation lasted for 1.05 Gyr (1 000 time units) of simulated time. A population size of 50 individuals was used. After 70 generations, no further improvements were recorded. Fig. 2 shows three projections of the appearance of NGC 5194 in the final model.

Using standard notation (see e.g. Danby 1988), the orbital parameters of the best-fitting model were $a = -20.92$, $e = 1.769$, $i = 89.98$, $\omega = 341.5$, $\Omega = 104^\circ$, and $E = 126^\circ$, where E denotes the eccentric anomaly at the end of the simulation. We made a large number of tests with simulations using up to 250 individuals and up to 200 generations, without finding any model parameters better than these.

The orbit is illustrated in Fig. 3. We checked carefully that the initial separation of the galaxies in this simulation was sufficiently large to avoid transients due to the introduction of the perturber, by carrying out another simulation in which the backward integration was extended an additional 525 Myr. With the optimal orbital parameters, NGC 5195 passed pericentre 908 Myr ago at a distance of 16.1 kpc. It crossed the disc plane 849 Myr ago at a radius of 17.0 kpc. At present the distance between the two galaxies along the line of sight is 147 kpc. Thus, according to this model, NGC 5195 is located far behind NGC 5194, and is moving away from it on a hyperbolic orbit.

The mass of NGC 5194 was found to be 0.28 and the mass of NGC 5194 0.0586, so that the ratio was equal to the rather low value 0.21. The disc inclination of NGC

5194 was found to be 29° and the position angle 5° , which is closer to the early values 35° and 0° of Burbidge & Burbidge (1964) than to those of Tully (1974).

Comparing the velocity field shown in Fig. 4 with the data of Rots et al. (1990), it is apparent that we did not arrive at a perfect model for NGC 5194/95. After attempting (and failing) to improve the model by using a modified deviation measure which weighted velocity deviations more strongly, we attempted to make a slight modification of the galaxy model used by the GA. An additional parameter was encoded in the chromosomes, namely a warp in the disc, consisting of a linear increase of the inclination angle of the disc from the stipulated value i_1 to a value $i'_i = i_i + \Delta i_1$ at radius $r = 5r_d$. However, no improvements of the model were found in any of the runs carried out with this additional parameter, despite the fact that these runs were somewhat lengthened in order to cope with the increase in the search space caused by the addition of the new parameter in the chromosomes.

4. Self-gravitating simulation

In order to reduce the computational demands in the GA simulations we had to simplify the model quite drastically. Here we give the results of a more realistic self-gravitating simulation that uses our optimum orbit.

In order to carry out this simulation, we used a new method developed by Wahde. This method attempts to combine the advantages of a grid-based code, e.g. speed of computation and an essentially linear variation of the computation time with the number of particles used in the simulation, with the advantages of a direct summation code, e.g. accuracy in the computation of forces. The code is essentially a P³M (particle-particle, particle-mesh) code, with the FFT-based mesh calculations replaced by a tree-code inspired computation of monopole contributions from the various grid cells. Since the forces between nearby particles are computed through direct summation, the resolution of the simulation is essentially determined by the applied softening length, which in our self-gravitating simulation was set to 0.15 kpc. The details of the code will be described elsewhere.

The model of NGC 5194 used in the self-gravitating simulation consisted of a halo with 32 000 particles of mass 7.00×10^{-6} , in the units used above, and a disc with 38 000 particles of mass 1.47×10^{-6} , giving a total mass of 0.28, as in our best non-self-gravitating model. In the direction perpendicular to the plane of the disc, the galaxy was modelled with a sech^2 function, with a scale height of 0.5 kpc. Typical interparticle distances in the main part of the disc were of order 0.1 kpc.

The scale radius of the (exponential) disc was equal to 4.7, and the scale radius of the (Plummer) halo was equal to 3.0. This parameter choice provided a compromise fit to the observed inner rotation curve, as determined by Tully (1974) from H α observations and by Kuno & Nakai (1997) from CO data, but consistent with our mass.




Fig. 5. A sequence of images from the self-gravitating simulation. The images are equally spaced (131 Myr) in time, which in this figure is counted backwards from the present time. The perturbing galaxy passes through the disc of the main galaxy between the second and the third frame

The setup procedure for the galaxy model started by a relaxation phase, in which the force field of the disc was allowed to grow, by slowly increasing the mass of the (stationary, during this phase) disc particles from 0 to their final values. This phase lasted for the equivalent of 735 Myr. Next, velocities were added to the disc particles. The computation of the disc particle velocities were based on the actual potential generated by the halo and disc particles at the end of the relaxation phase. The Toomre Q parameter of the undisturbed disc was set to 1.5. Finally, the perturbing galaxy, NGC 5195, was added using a simple point-particle representation with mass 0.0586.

Then, using the orbital parameters of the best non-self-gravitating model, two point particles were integrated backwards in time for a period corresponding to 1.05 Gyr in our units. The resulting positions and velocities of the two particles were taken as initial conditions for the posi-

tions and velocities of the centres-of-mass of the galaxies in the self-gravitating simulation. The forward integration was then carried out, using 10 000 time steps of length 0.1. Since the motion of the galaxies in the self-gravitating model is not exactly Keplerian, the final positions of the two galaxies differed somewhat from the observed positions. However, the difference was marginal. The results of the simulation are shown in Figs. 5 and 6. Clearly, the results of the non-self-gravitating model are confirmed.

Although the self-gravitating model shows that the results from the non-self-gravitating simulations remain valid, it also inherits the weaknesses of those models. Thus, the observed velocity variations along the tidal tail are still unexplained. Furthermore, inspection of Fig. 5 suggests that the inner parts of the tidal tail and the counter-arm have developed somewhat too far. We return to these issues in Section 5.2.




Fig. 6. The results of a fully self-gravitating simulation, using the best parameters found for the NGC 5194/95 system. In order to make the figure more realistic-looking, the light from each particle has been spread out according to a Plummer law. In order to bring out the features of the large H I-arm, the central parts of the galaxy were overexposed in this image

5. Discussion

5.1. Usefulness of the GA method

Somewhat surprisingly, despite the often enormous size of the multi-dimensional search spaces involved in the determination of galactic orbits, there have been few attempts to define a general systematic and efficient strategy for solving this search problem.

A possible systematic method consists simply of covering the search space with a grid with constant spacing between points along each axis, and carrying out an exhaustive, non-directed search through this space. However, as pointed out by Wahde (1998), the grid thus obtained will, in realistic cases, be much too coarse-grained to allow a thorough search through the entire space of orbits. For example, in a 7-dimensional search space, a search involving the testing of as many as 100 000 different orbits would only allow around 5 different values per dimension. As shown by Wahde (1998), this is far from sufficient.

Thus, with exhaustive searches excluded, the next logical step is to consider directed searches. One of the most straightforward approaches is to use a gradient-based search method which, for any given point in the search space, computes the direction of steepest descent on the error surface, and then follows this direction towards the optimum. This approach will be successful if the search space has a simple structure, i.e. if there is a single global optimum.

However, if the search space is rugged, with many local optima, a simple gradient-based method of the kind outlined above will, with high probability, become trapped in a local optimum far from the global optimum.

There are, however, search techniques which are specifically designed to cope with rugged search spaces. In this paper, we have used a GA, which conducts a parallel search through the space of possible orbits and is therefore less likely to be trapped in a local optimum. GAs have proved to be very useful in many applications involving complex search spaces with local optima¹. While GAs include some parameters that must be set in a more or less ad hoc fashion, such as e.g. the mutation probability, the performance of the method is not strongly dependent on small variations of these parameters.

Before selecting a search method one should investigate the structure of the search space. Obviously, it is very likely that the detailed structure of this space will vary from case to case. However, if the search space turns out to be complicated in any given case, it would be prudent to use, in other cases as well, methods that are able to cope with such a situation. It should be noted that the exact shape of the search space will of course depend on the error measure (or, in the case of GAs, the fitness measure) that is used.

Needless to say, it is all but impossible to visualize a 7-dimensional fitness landscape. Thus, in order to investigate the shape of this space, we have chosen to vary one parameter at a time, and to hold the remaining parameters fixed at the values corresponding to the best orbit found by the GA. The resulting curves are shown in the panels of Fig. 7, from which it is evident that the search space is very rugged indeed, even though a few dimensions show a more regular structure. Note also that the ruggedness of the curves in Fig. 7 is *not* an artefact of plotting only a few points per curve. On the contrary, each curve was plotted by connecting 100 points sampled in the range $[0.7p, 1.3p]$, where p denotes the value of the parameter in the best solution found by the GA.

The appearance of the search space, as shown in Fig. 7 strongly favours the use of a GA when searching for orbital parameters, and may also cast doubt on some orbital parameter determinations which have either used an exhaustive (and, therefore, sparse) search or simply selected a model that appeared to be similar to the observations.

5.2. Parameters of the M51 system

A strength of the GA approach is that we can search for the optimum model given well-defined simplifying assumptions. We can then be sure that any shortcomings of the final model are due to the underlying assumptions rather than to our failure to cover the search space completely.

The optimum model reached by the GA is generally similar to the models that have been previously proposed

¹ It should be noted that some authors refer to GAs as being gradient-based. This, however, is a misnomer: the *effect* of a GA is often to follow a rather smooth curve towards higher fitness (i.e. smaller error), but the search does *not* make use of explicit gradient information.

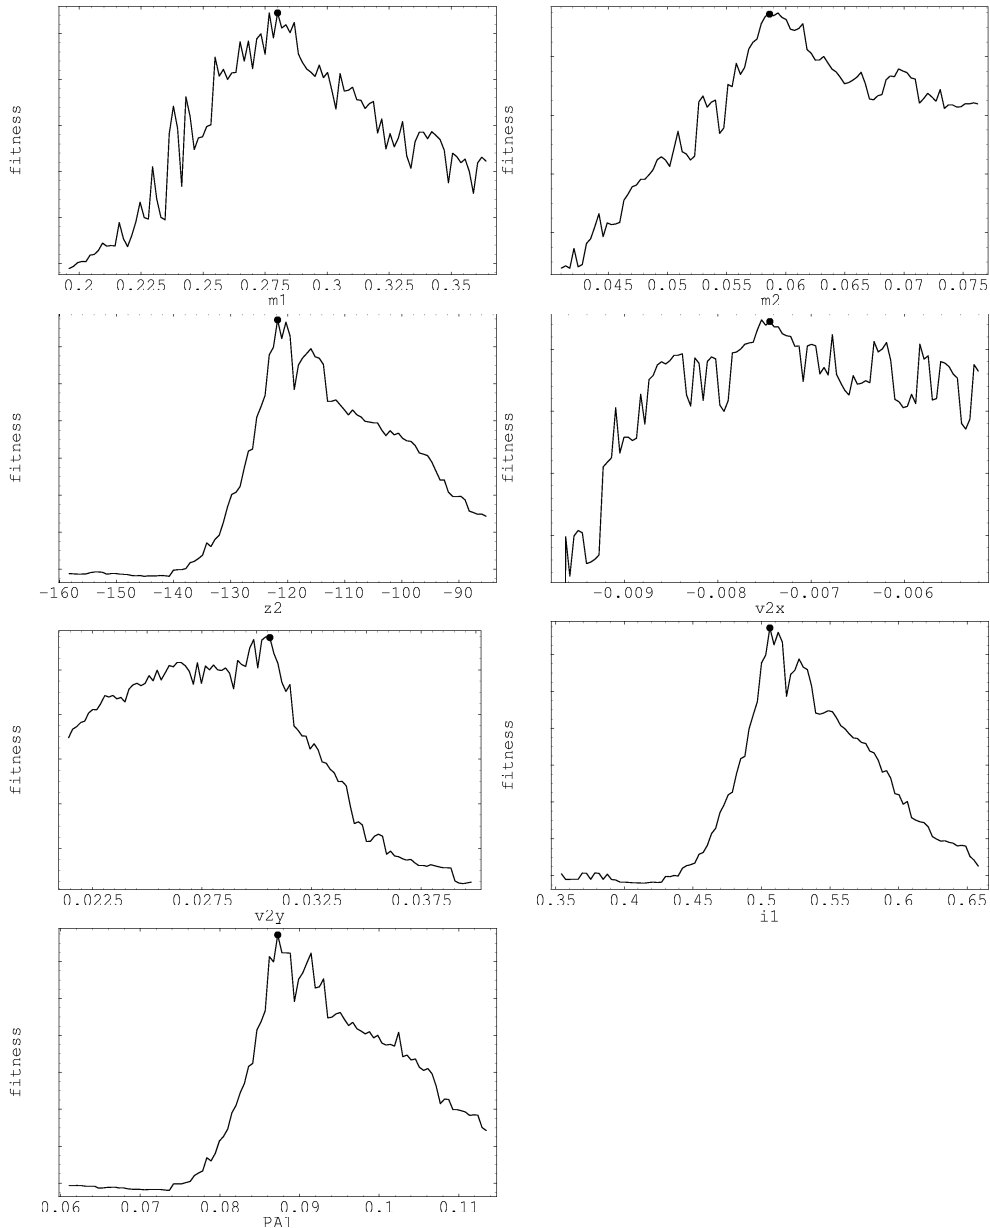


Fig. 7. The curves illustrate the fitness variation when each parameter of the interacting system is varied around the best solution (see Sect. 3.2), while keeping all other parameters fixed. m_1 and m_2 denote the masses of the two galaxies, z_2 the final z -position of NGC 5195, v_{2x} and v_{2y} the final x - and y -velocities of NGC 5195, and i_1 and PA_1 the inclination and position angle of NGC 5194 (in radians), respectively

based on trial and error (Hernquist 1990; Antonioletti & Nelson 1999; see the review of Barnes 1998). This confirms that the extended tidal tail of Rots et al. (1990) is a very robust feature of models where the perturber crosses the M51 disc from above in the south-west on a highly inclined, more or less parabolic, orbit. The fact that very simple models consistently produce this feature make more complicated models involving multiple passages, such as that of Salo & Laurikainen (2000), inherently less attractive.

Unlike previous models ours does not lead to a relatively large mass for NGC 5195. In fact, the value 0.28 we find for the mass ratio of the galaxies is near the lower

end of the observationally allowed range. Another feature of our model worth noting is the double structure of the tidal tail, which also seems to be present in the observations. Note that the double structure of the tidal tail, which is evident in the self-gravitating simulation (see the lower panels of Fig. 5, has no direct counterpart in the non-self-gravitating model shown in Fig. 2. However, as can be seen in the central view of M51 in Fig. 2, instead the densities are higher at the edges of the arm than in its centre. This structure corresponds to the double structure seen in the self-gravitating simulation.

The optimum model also produces a diffuse counter-arm to the north and west. Although this was not explic-

itly included in the optimization, it seems to have a counterpart in the extended diffuse HI detected by Haynes, Giovanelli & Burkhead (1978) and Rots et al. (1990), and in the optical emission in this region (Burkhead 1978). In other models this material has been taken to originate in NGC 5195.

In order to get an idea of how the satellite is affected by the encounter we have run a test-particle simulation of NGC 5195. We used the mass and orbit from our model, and added a disc with inclination 43° and PA 91° (Spillar et al. 1992) rotating with the east side approaching (Schweizer 1977). During the passage the galaxy develops an extended fairly symmetric two-armed tidal spiral pattern, but by the present time this pattern has largely dispersed. It is plausible that the bar in NGC 5195 could have been tidally produced, but establishing this would require a self-gravitating simulation involving uncertain model-dependent assumptions.

We have not attempted to reproduce the inner spiral structure of M51 in detail. A rather weak two-armed spiral does survive up to the present in our self-gravitating simulation. It has been shown by Donner & Thomasson (1994) that tidal spiral arms can survive for several billion years. Making the disc less stable and including gas in the simulation would produce a more pronounced spiral pattern. However, in our opinion the precise location and shape of the spiral arms are too sensitive to detailed model assumptions to allow for a proper test of the tidal model.

The most important shortcoming of our model is its failure to reproduce the observed velocities in the HI arm. We attempted to improve the velocity fit first by increasing the weight of the velocities in the fitness measure as discussed in Sect. 2.2, and then by incorporating a simple initial warp, described at the end of Sect. 3.2, in the disc model. Neither of these attempts succeeded in producing an improved fit. Although the inclusion of a simple warp did not help in this respect, a more complex initial warp is perhaps still the most plausible explanation for the velocity variation along the arm.

A second problem is that the inner part of the tidal tail and the northern counter-arm seem to have advanced too far compared to the observations. We suspect that this is related to our use of a Keplerian rotation curve. If the rotational velocities at large radii were larger, the outer part of the tidal tail would reach its present position sooner.

Another indication that our rotation curve may be declining too steeply at large radii is the rather small mass of our model, only $5.6 \times 10^{10} M_\odot$, whereas Tully (1974) finds $7.6 \times 10^{10} M_\odot$, and Kuno & Nakai (1997) $1.5 \times 10^{11} M_\odot$. In view of the asymmetry of the rotation curve at large radii (Tilanus & Allen 1991), and the possible presence of an initial warp, the extrapolation of the rotation curve is in any case highly uncertain.

Considering that our optimum model is generally similar to the models proposed in previous studies, it seems that, with the geometry of the M51 system, the determi-

nation of the orbital parameters of the interaction from the density distribution is robust in the sense that moderate variations in the initial model does not qualitatively modify the relative orbit. This is in contrast to the velocity field, where the difficulties for simple models can be traced back to the fact that the predicted velocities are sensitive to the extrapolated kinematics of the initial disc at large, unobserved radii. If we have not attempted to produce a more satisfactory model for the velocity field, it is partly because such a model might be better viewed as a way of inferring the required properties of the initial disc, given the orbit derived from the densities.

6. Conclusion

We have shown that genetic algorithms provide an efficient and reliable method for determining the orbital parameters of interacting pairs of galaxies, and we thus recommend this method as a general tool for such investigations. In addition we have shown that, at least for the interaction studied in this paper, the search space is quite rugged with several local optima, further strengthening the case for GAs in preference to e.g. gradient-based search methods.

Based on a very simple restricted three-body model the best orbit found using the GA is almost parabolic, perpendicular to the plane of the sky, and crosses the M51 disc in the south-west. In this model the mass ratio of the two galaxies is within the observational range, but the velocity variation along the extended tidal tail of M51 is not reproduced. A self-gravitating simulation using the orbital parameters of the optimum model shows that the characteristic tidal features are still present in this more realistic model.

Acknowledgements. KJD is thankful to Nordita for hospitality and support during two visits when part of this work was done.

References

- Antonioletti, M., Nelson, A.H. 1999, in IAU Symp. 186, Galaxy Interactions at Low and High Redshift, ed. J.E. Barnes, D.B. Sanders, 105
- Barnes, J.E. 1998, in Galaxies: Interactions and Induced Star Formation, ed. D. Friedli, L. Martinet, D. Pfenniger, (Springer-Verlag, Berlin Heidelberg New York), 275
- Burbidge, E.M., Burbidge, G.R. 1964, ApJ, 140, 1445
- Burkhead, M.S. 1978, ApJS, 38, 147
- Charbonneau, P. 1995, ApJS, 101, 309
- Danby, J.M.A. 1988, Fundamentals of Celestial Mechanics, 2nd ed. (Willman-Bell, Richmond VA)
- Donner, K.J., Thomasson, M. 1994, A&A, 290, 785
- Engström, S. 1992, Ph.D. Thesis, University of Göteborg and Chalmers University of Technology
- Engström, S., Athanassoula, E. 1991, in Dynamics of Disc Galaxies, ed. B. Sundelius (Göteborg University and Chalmers University of Technology, Göteborg), 215
- Feldmeier, J.J., Ciardullo, R., Jacoby, G.H. 1997, ApJ, 479, 231
- Garcia-Gomez, C., Athanassoula, E. 1991, A&AS, 89, 159

- Haynes, M.P., Giovanelli, R., Burkhead, M.S. 1978, AJ, 83, 938
- Hernquist, L. 1990, in Dynamics and Interactions of Galaxies, ed. R. Wielen (Springer-Verlag, Berlin Heidelberg New York), 108
- Howard, S., Byrd, G.G. 1990, AJ, 99, 1798
- Kuno, N., Nakai, N. 1997, PASJ, 49, 279
- Mitchell, M. 1996, An Introduction to Genetic Algorithms (MIT Press, Cambridge MA)
- Rots, A.H., Bosma, A., van der Hulst, J.M., Athanassoula, E., Crane, P.C. 1990, AJ, 100, 387
- Salo, H., Laurikainen, E. 2000, MNRAS, 319, 377
- Sandage, A., Tammann, G.A. 1975, ApJ, 196, 313
- Schweizer, F. 1977, ApJ, 211, 324
- Spillar, E.J., Oh, S.P., Johnson, P.E., Wenz, M. 1992, AJ, 103, 793
- Theis C., Harfst S. 2000, in ASP Conf. Ser. 197, Dynamics of Galaxies: from the Early Universe to the Present, ed. F. Combes, G.A. Mamon, V. Charmandaris, 357
- Thronson, H.A., Rubin, H., Ksir, A. 1991, MNRAS, 252, 550
- Tilanus, R.P.J., Allen, R.J. 1991, A&A, 244, 8
- Toomre, A., Toomre, J. 1972, ApJ, 178, 623
- Tully, R.B. 1974, ApJS, 27, 449
- Wahde, M. 1998, A&AS, 132, 417
- Wahde, M. 1999, in Numerical Astrophysics 1998, ed. S.M. Miyama, T. Hanawa, K. Tomisaka (Kluwer, Dordrecht), 401