Adaptive Multitrack Fitting

A. Strandlie [a,1], R. Frühwirth [b,2]

[a] *Department of Technology, Gjøvik College, Gjøvik, Norway and Department of Physics, University of Oslo, Norway*

[b] *Institute for High Energy Physics of the Austrian Academy of Sciences, Vienna, Austria*

Abstract

We address the problem of estimating the parameters of very close or overlapping tracks. In order to achieve optimal performance the estimation process has to be carried out concurrently with the task of assigning the observations to the track candidates. We present a new method that is based on the Deterministic Annealing Filter (DAF) and implements a global competition. The method is studied on simulated tracks in the ATLAS Transition Radiation Tracker, both without and with mirror hits and at various noise levels. We show that the new method is superior to a sequential application of the DAF to the track candidates, and also superior to the Elastic Arms algorithm (EAA), in particular at high noise levels.

1 Introduction

Track reconstruction in the LHC detectors will be faced by the problem of finding and fitting high-momentum tracks in the presence of a large background of noise, backscattering, and low-momentum tracks. In some cases the problem is aggravated by the presence of mirror hits, for example in the Transition Radiation Tracker (TRT) of the ATLAS detector [1]. It has been shown previously [2] that for tracks that are isolated fairly well an iterated non-linear filter with annealing (the Deterministic Annealing Filter or DAF) is a powerful tool for estimating the track parameters concurrently with solving the assignment problem, i.e. deciding which hits are "signal" and which

[1] E-mail: `Are.Strandlie@hig.no`
[2] Corresponding author. Address: A-1050 Wien, Nikolsdorfer Gasse 18,
Tel: +43 1 544 73 28 33, Fax: +43 1 544 73 28 54,
E-mail: `Rudolf.Fruehwirth@oeaw.ac.at`

Preprint submitted to Elsevier Preprint *5 May 2000*

are "noise". The decision rule of the DAF is based on the assumption that in each layer of the detector several hits are competing for inclusion in the track. Note that the decision is "soft", which means that each hit participates in the estimate in concordance with its assignment probability. The DAF is, however, ill equipped to deal with situations where two (or more) tracks are so close to each other that most hits could be associated to either of them.

Another well-known method for solving the assignment problem in parallel to the estimation procedure is the Elastic Arms algorithm (EAA) [3]. The EAA has also been generalized to detectors with ambiguities [4,5], and we will make use of this generalization later. The EAA is a truly global method insofar that all tracks are treated in parallel. The energy function is formulated such that in the limit of low temperature a hit is associated to at most one track. As the tracks compete for association to the hits, the EAA can in this respect be viewed as the opposite of the DAF where the hits compete for inclusion in a track. As a consequence a hit will be associated to the closest track that is inside the compatibility limit, even if association with a different track gives a better global matching between tracks and hits. It is therefore conceivable that a global decision rule performs better than pure competition between hits (DAF) or pure competition between tracks (EAA).

In the following we present such a global competition scheme. Its performance is evaluated and compared to existing methods in a couple of simulation experiments, using the ATLAS Transition Radiation Tracker as a model detector. In order to make the results more general we run the experiment both without and with mirror hits, and for various levels of random noise superimposed on the actual tracks. Although the experiments are confined to pairs of tracks all methods can deal with an arbitrary number of arms. The paper concludes with a discussion of the results and of the limitations which are currently inherent to the multitrack fit, as well as an outline of the work to be done to overcome these limitations and make our approach operational in more complicated settings.

2 Multitrack fitting

The multitrack fit proposed here is an iterative procedure, similar to the Deterministic Annealing Filter (DAF) and to the Elastic Arms algorithm (EAA). It starts out with a certain number of track candidates or arms and a certain set of hits which are compatible with at least one of the arms. Finding the correct number of arms and reasonable initial track parameters is a non-trivial pattern recognition problem, which we assume to be solved at the moment (but see also Section 4). From the initial track parameters of the arms we compute the initial assignment probabilities of all hits to all arms. These assignment probabilities are used as the weights in the subsequent reweighted least-squares estimation of all track parameters. A hit which has a small assignment probability with respect to a certain arm thus contributes very little to the estimate of this arm's parameters. The entire procedure, which can be interpreted as an instance of the EM algorithm [6], is repeated until convergence of the assignment probabilities.

Now suppose that in a certain layer of the detector there are m arms (track candidates) and n hits which are compatible with at least one of the arms. We can describe the state (position, direction, curvature) of arm j by a state vector $\boldsymbol{x}_j$ and the observed values of hit i by an observation vector $\boldsymbol{y}_i$. A state and an observation which belong to the same track are coupled by the measurement equation

$$\boldsymbol{y} = \boldsymbol{H}\boldsymbol{x} + \boldsymbol{\varepsilon},$$

where $\boldsymbol{\varepsilon}$ is the observation error, which is assumed to be normal with zero mean and covariance matrix $\boldsymbol{V}$.

As a preparation for computing the assignment probability p_{ij} between hit i and arm j we set up the following matrix Φ, which has a row for each hit and a column for each arm:

$$(\Phi)_{ij} = \varphi_{ij} = \varphi(\boldsymbol{y}_i; \boldsymbol{H}\boldsymbol{x}_j, \boldsymbol{V}),$$

where $\varphi(\cdot; \boldsymbol{\mu}, \boldsymbol{V})$ is the multivariate normal probability density with mean vector $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{V}$. It should be noted that the matrix $\boldsymbol{V}$ depends on the current temperature, i.e. it is larger than the nominal measurement error at the outset and is then decreased in each step of the annealing procedure until the nominal size is reached. The annealing schedule will be discussed in more detail in the next section.

We now define four methods for computing the assignment probabilities.

Method 1: Competition between hits. There is competition between all hits for each track, but there is no competition between the tracks. This procedure is equivalent to the DAF [2]. The assignment probabilities are computed by dividing φ_{ij} by its associated column sum plus a constant (normalization by columns):

$$p_{ij} = \frac{\varphi_{ij}}{\sum_k \varphi_{kj} + c}.$$

Method 2: Competition between tracks. There is competition between all tracks for each hit, but there is no competition between the hits. It is equivalent to the original Elastic Arms algorithm [3]. The assignment probabilities are computed by dividing φ_{ij} by its associated row sum plus a constant (normalization by rows):

$$p_{ij} = \frac{\varphi_{ij}}{\sum_l \varphi_{il} + c}.$$

Method 3: Global competition. There is competition between all entries which are incompatible, i.e. are either in the same row or in the same column. This method is proposed here for the first time. The assignment probabilities are computed by dividing φ_{ij} by the sum of all elements in the same row and column plus a constant (normalization by columns and rows):

$$p_{ij} = \frac{\varphi_{ij}}{\sum_k \varphi_{kj} + \sum_l \varphi_{il} - \varphi_{ij} + c}.$$

Method 4: Competition between tracks and between mirror hits. This is a refinement of Method 2 developed specifically for detectors with mirror hits, e.g. drift chambers or drift tubes [4,5]. Besides competition between tracks there is also competition between a hit and its mirror hit. The assignment probabilities are computed separately for each of pair of hit and mirror hit. If (i_1, i_2) is such a pair, the assignment probabilities are computed according to

$$p_{i_k j} = \frac{\varphi_{i_k j}}{\sum_l \sum_\alpha \varphi_{i_\alpha l} + c}.$$

The normalization constant is therefore the sum of all elements in the respective submatrix plus a constant. If for some reason a hit has no mirror hit it is treated according to method 2. In a detector without mirror hits methods 2 and 4 coincide.

The constant c effectively defines a cut beyond which the assignment probability quickly drops to 0. All observations with non-zero assignment probability are combined to a single observation by a weighted mean, the weights being proportional to the respective assignment probabilities. This combined observation is then used in the least-squares estimation of the track parameters of all arms. The estimation step can be implemented in various ways, for instance as a Kalman filter, a linearized least-squares fit, or as a non-linear least-squares fit. If the track model is an exact circle, the entire multitrack fit can be carried out in a linear model, after a conformal transformation to 3-space [7].

The entire procedure is iterated until the assignment probabilities settle to their final values. As mentioned above, the iteration is carried out with annealing, in order to avoid suboptimal solutions (local minima).

3 Simulation experiments

3.1 Detector and track sample

The simulation experiments described in this section have been performed with tracks in the ATLAS Inner Detector Transition Radiation Tracker (TRT). The ATLAS TRT is a drift tube detector consisting of a barrel part and an end-cap part. The drift tubes or “straws” in the barrel are arranged in cylindrical layers, the straws being parallel to the beam axis. In the end-cap, the straws are arranged in wheels and point radially out from the beam. We will in this paper focus on simulated tracks coming from the barrel part of the TRT.

There are 75 layers of straws in the barrel, and the separation between layers is 6.8 mm. Within the same layer, the separation between the straws is also about 6.8 mm. The straw diameter is 4 mm. The barrel layers range from a radius of about 56 cm to about 106 cm, and there are about 50000 straws in total in this part of the TRT. More technical details concerning the TRT can be found in the ATLAS Inner Detector Technical Design Report [1]. A small section of the barrel TRT is shown in Figure 1.

Since we are dealing with a drift tube detector, the measurements are two-dimensional. For each measurement we know the radius, the polar angle of the sense wire of the straw, the absolute value of the drift distance and also the number of the particle creating the hit. The measurements are in principle ambiguous, but the information about which hit is the correct hit and which is the mirror hit is present. This knowledge is of course not used during the reconstruction, but we take it into account in the analysis of the results.

We have simulated a sample of 980 pairs of tracks all coming from the origin, i.e. there are 1960 tracks in total. A pair is simulated by first assigning to both tracks the same transverse momentum vector $\boldsymbol{p}_T$, the magnitude of this vector being larger than or equal to 1 GeV/c. Each component of the momentum vector is then given a random perturbation of in the order of 1 % of its magnitude. This is done independently for both tracks. The outcome of this procedure gives everything from completely separated tracks to completely overlapping tracks. If both tracks pass through the same drift tube, we keep the hits from both. This would of course not be the case in a real situation, where the signal from the straw would be some kind of sum of the contributions from each of the particles. However, the main goal of the simulations of this section is to evaluate the basic statistical properties of the different multitrack algorithms, and this is most straightforwardly done by assuming a “perfect” response from the detector.

In the LHC environment it can be expected that there is further background apart from the mirror hits, for instance electronic noise or hits from crossing or spiraling low-momentum tracks. We have therefore foreseen possibility to contaminate the tracks with noise points, at various noise levels. A noise level of n % means that for n % of the hits, a noise hit has been added in the same drift tube in a random fashion. If the mirror hits are enabled, a noise point comes together with its mirror hit. A noise level 0 means that there is no additional noise. An example of a track pair can be found in Figure 2, with noise level 0 (top) and noise level 50% (bottom), mirror hits included. Without noise it is possible to visually discern between the two tracks, but the overlap is nevertheless significant. With noise this is almost impossible.

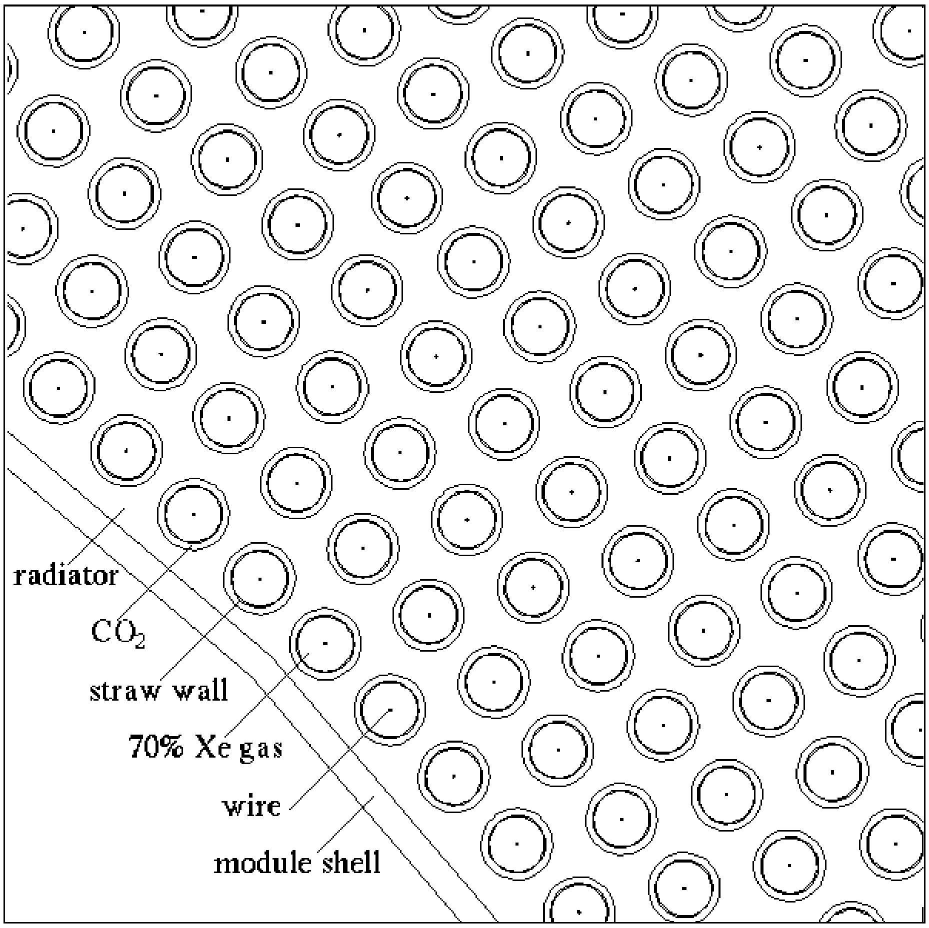

Fig. 1. A section of the barrel TRT in the ATLAS Inner Detector. Taken from the ATLAS Inner Detector Technical Design Report [1], with permission.

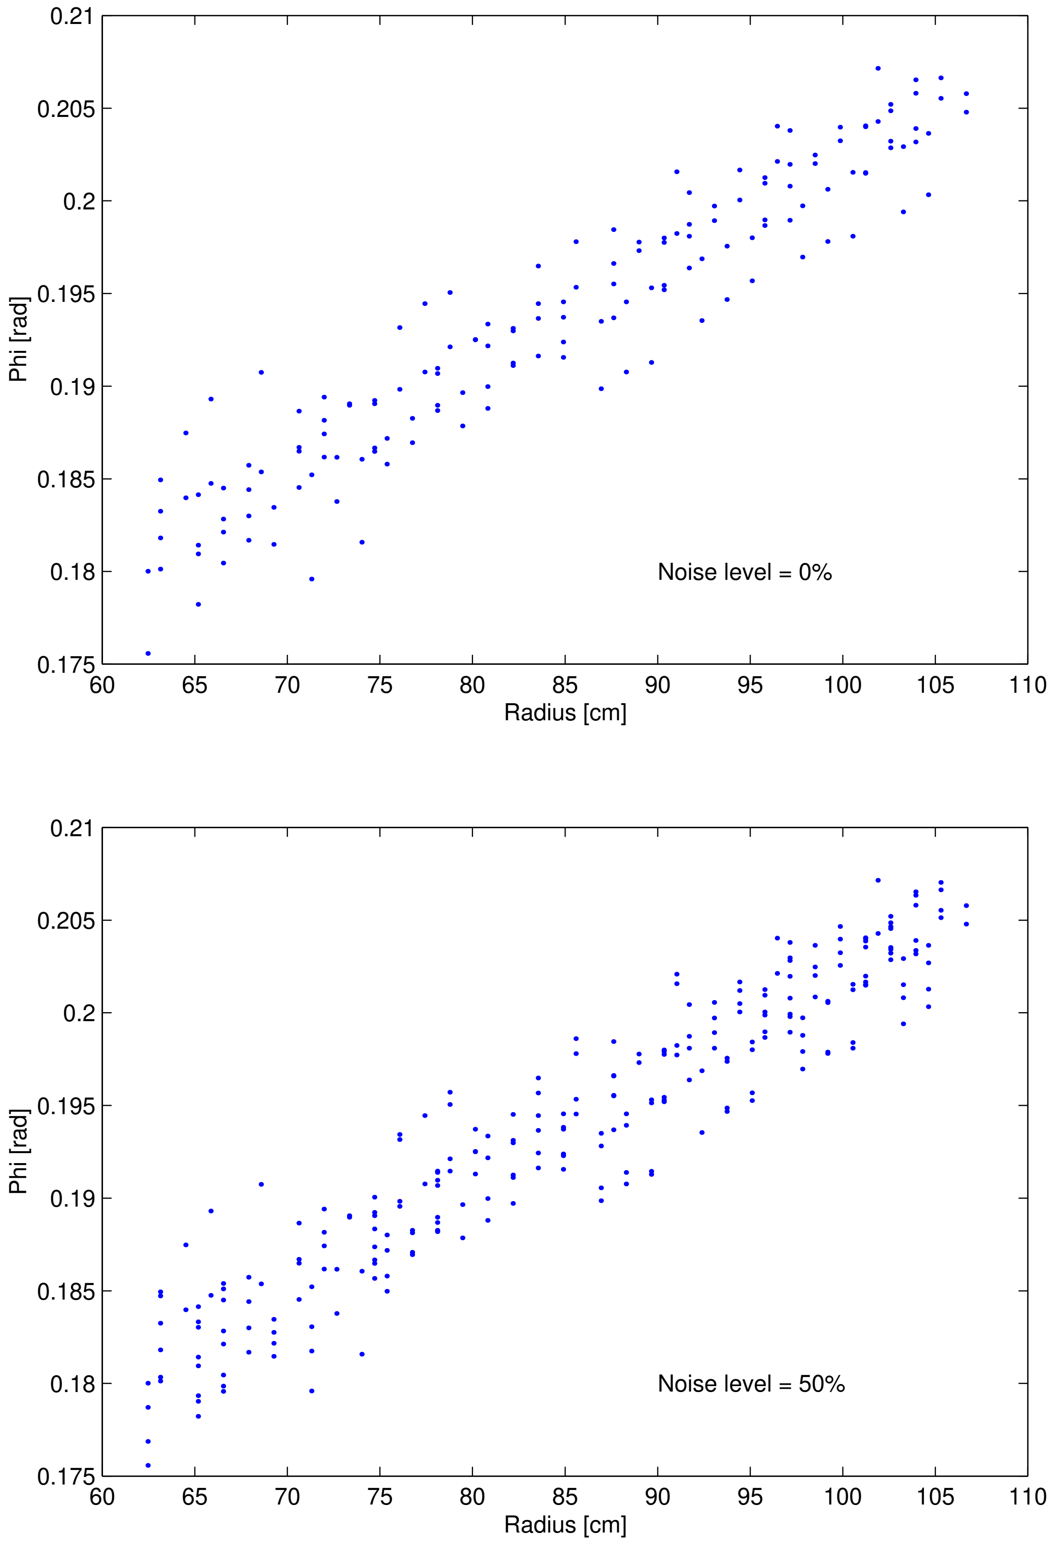

Fig. 2. A track pair in the $R\Phi$-projection. The noise level is 0 in the top figure, and 50% in the bottom figure. The mirror hits are included.

3.2 Initialization and annealing schedule

All multitrack algorithms require an initial guess of the parameters of the tracks in the pair, and we have in this work tested out two different initialization procedures. The first is simply to initialize the track parameters by their true values. By doing so we know that the initial tracks are not very far from the final tracks, as given by the output of the estimation algorithms. It is therefore hoped that the risk of ending up in local minima of the energy function is reduced and that the results mainly reflect the basic statistical properties of the estimators. On the other hand, by using the true track parameter values as the initial values there is always a risk of introducing a bias. In order to validate the results from the truth initialization, we have applied a separate pattern recognition procedure for the initialization as well. It has turned out that it is non-trivial to apply standard track finding techniques to a track pair sample as ours. Since we encounter both separated and overlapping tracks, it is not easy to tune the parameters of the pattern recognition algorithm in a satisfactory way. These parameters could for instance be the bin size and threshold value in a Hough transform [8] or the quantile in a generalized Least Median of Squares (LMS) [9] procedure. We have therefore fallen back to a simple procedure which has turned out to work rather nicely, even though it is tailored to the knowledge that we have exactly two tracks in "jets" like these. It is based on first doing a least-squares fit in the $R\Phi$-projection to all points in a jet candidate. (A jet candidate is defined as the collection of all points in a track pair.) This yields the "centre of gravity"-track (cog), given by the curvature κ and the angle ψ that the tangent to the fitted track makes with the x-axis at the origin. We then calculate the standard deviation in the differences in the angle Φ between the cog-track and all points in the jet candidate. The two initial tracks are defined by setting ψ equal to the fitted ψ plus and minus this standard deviation and keeping the same κ for both tracks. We therefore call this procedure the "cog$+/-$" initialization.

The actual annealing schedule has to be chosen with more care for these multitrack algorithms than what is the case for the single-track DAF [2]. In the single-track case, the primary point of concern is to avoid ending up in a local minimum of the energy function. This is effectively done by choosing the starting temperature not too low and by lowering the temperature not too fast. This is of course also important in the multitrack case. However, in the multitrack case it has turned out vital to avoid the situation of two tracks coming too close during the annealing, since these will have a tendency to remain close during the rest of the annealing process. This can be understood by noting that two tracks in exactly the same position never will be able to escape from one another, as a consequence of the structure of the iteratively reweighted update procedure. It is therefore plausible that two very close tracks also will be plagued by the same problem. Two tracks initialized close enough to one another will at a sufficiently high temperature be attracted to the same minimum and rapidly end up in nearly the same position. In order to avoid this effect, the starting temperature cannot be too high, and the annealing cannot be too slow. The annealing schedule therefore has to be a compromise between the two above mentioned, competing criteria, and the optimal choice will depend on the actual problem at hand.

We have also observed that the convergence of the algorithms at the final temperature is slower when treating multiple tracks. This is probably due to a more complex and ill-behaved energy landscape in the multitrack case than in the single-track case. The number of iterations at the final temperature therefore has to be increased. In our simulation experiments we have chosen the number 15 in order to obtain a satisfactory accuracy.

In the next two subsections we shall compare the results of the four different methods presented in Section 2. The only differences between them are in the structure of the weights used in the iteratively reweighted least-squares procedure. We recall that method 1 (DAF) exhibits a competition between all hits in a detector layer for each track, but no competition between the tracks. Method 2 (EAA) offers a competition between both tracks for each hit, but there is no competition between the hits. It is equivalent to the original Elastic Arms algorithm [3]. Method 3 (Global Competition) is the proposal of this paper and has a global competition between all hits and both tracks in each layer. Method 4 (Modified Elastic Arms) has weights similar to the ones given by Blankenbecler [4]. It implements a competition between both tracks for each hit, as well as a competition between a hit and its mirror hit only. Method 4 is relevant only for the experiments where the mirror hits are turned on.

The precision of the estimators will be assessed by the relative generalized variance V of the residuals of the estimated track parameters with respect to the truth values. We recall that the generalized variance is the determinant of the covariance matrix of these residuals. The baseline is a single-track fit to the measurements of each track separately. Since the "jet candidate" in this case only consists of points coming from the one track to be fitted, the precision of the single-track algorithm defines a lower limit of what can be achieved with a multitrack algorithm.

3.3 Results without mirror hits

In the first experiment the mirror hits are turned off. For the results in Table 1, the algorithms have been initialized with the true values of the track parameters. We observe that method 1 is far less precise than the other two methods. Method 1, the DAF, is the only method without any competition between the tracks, and this obviously deteriorates the performance significantly. Method 2 (Elastic Arms) is better than our proposal (method 3) up to a noise level of about 20%. For higher noise levels method 3 is significantly better than method 2. Indeed it seems to be affected very little by the addition of noise points. We ascribe this to the capability of the method to find the best global match between the tracks and the hits in each single layer.

Noise level	Method		
	1	2	3
0%	251	1.94	2.53
10%	304	2.39	2.66
20%	325	3.17	3.18
30%	233	4.00	3.19
40%	230	5.77	3.41
50%	372	8.71	3.91

Table 1
The relative generalized variance of methods 1–3 without mirror hits and with truth initialization.

Initialization with the true track parameters is of course not possible in the real world of the experiment. Table 2 shows the results with the cog+/− initialization described in Subsection 3.2. Obviously methods 1 and 3 are much more affected by the initialization than method 2, so that method 2 is now the best throughout. We can, however, rescue the superior behaviour of method 3 in the presence of noise by using the results of method 2 as the starting value of method 3. The results are in the last column (2/3) of Table 2. Method 3 is now the best throughout. We have no ready explanation for the fact that method 3 is more sensitive to a bad initialization than method 2.

Noise level	Method			
	1	2	3	2/3
0%	2.7E6	8.74	43.57	8.22
10%	2.1E6	8.90	31.78	7.39
20%	2.2E6	7.42	27.51	5.01
30%	1.9E6	8.05	22.62	4.73
40%	1.9E6	7.53	24.07	3.56
50%	1.7E6	16.49	33.83	7.14

Table 2
The relative generalized variance of methods 1–3 and method 2/3 without mirror hits and with cog+/− initialization.

3.4 Results with mirror hits

We now proceed to study our algorithms with the mirror hits turned on. In Table 3 the algorithms have been initialized with the true values of the track parameters. The results are similar to the ones obtained without mirror hits (see Table 1). Again, method 1 is consistently the least precise of all. Method 2 is now also significantly less precise than methods 3 and 4. This is due to the lack of competition between the different hits in a layer. Methods 3 and 4 are comparable but, nevertheless, not totally equivalent in behaviour. At noise levels below 20%, method 4 is better, whereas method 3 is better at higher noise levels. This can be understood by noting that method 4 exhibits a competition between a hit and its asscociated mirror hit only, while method 3 treats all hits in a layer in the same manner. Thus, if the dominant source of ambiguity is mirror hits, method 4 is the one which puts in the prior information of the structure of the data in the most realistic fashion. It is therefore the best of the two at low noise levels. However, a major disadvantage with this approach is that a noise hit/mirror hit-pair in a layer is allowed to pull on the track without any explicit competition from the correct hit/mirror hit-pair. This can be expected to deteriorate the performance of method 4 as the noise level increases. On the contrary, method 3 has been specifically designed to cope with the problem of the presence of noise hits in addition to track hits. It should therefore come as no surprise that method 3 is the best at higher levels of noise.

Noise	Method			
level	1	2	3	4
0%	281	36.2	4.52	2.84
10%	270	58.7	5.35	4.35
20%	388	100.9	6.26	7.06
30%	358	185.3	7.19	9.51
40%	409	238.9	9.50	12.44
50%	653	301.6	11.66	17.65

Table 3
The relative generalized variance of methods 1–4 with mirror hits and with truth initialization.

Finally, in Table 4 we show results with mirror hits and initialization with the cog+/− method. All methods are now less precise than with the truth initialization. It is therefore clear that the performance of the methods to some extent depends on the quality of the initialization procedure. Method 1 is again the worst, followed by method 2. Method 3 now consistently exhibits a larger spread of the track parameters than method 4, similar to what we have seen without mirror hits (Table 2). Again, we can rescue method 3 by using the result of method 4 as the initial value. The results are in the last column (4/3) of Table 4. Method 3 is now the best, with the exception of the case of no noise at all.

Noise level	Method				
	1	2	3	4	4/3
0%	7.6E5	1837	72.2	11.9	15.3
10%	7.5E5	2698	101.6	22.0	22.2
20%	9.7E5	3532	313.6	40.6	33.0
30%	9.8E5	3800	248.6	43.5	32.6
40%	9.3E5	4449	193.0	65.9	45.7
50%	1.1E6	4817	388.4	96.0	68.8

Table 4
The relative generalized variance of methods 1–4 and method 4/3 with mirror hits and with cog+/− initialization.

It is instructive to look at the final assignment probabilities in order to get a feeling of how a method works. In Figure 3 we show histograms of these assignment probabilities after the application of method 4, with mirror hits and 50% noise level.

A false track hit is here defined as a measurement originating from the other track in a track pair than the one under consideration. It can be seen that the vast majority of the true track points are correctly assigned (a). Moreover, most of the mirror hits are given a very low assignment probability (b). For the false track hits, both the hits and the mirror hits usually have a low assignment probability (c,d). The overall behaviour is thus as one would expect of an algorithm being able to correctly assign the true track points and discard the rest.

Finally, we show again the track pair of Figure 2, but this time with the fitted tracks and the true track hits being marked (Figure 4). It is remarkable how little the fitted tracks are affected by the noise, although in the bottom figure there is a noise hit (plus its mirror hit) in about half of the layers.

4 Discussion and outlook

We have compared several possible approaches to adaptive multitrack fitting. We have shown that for overlapping pairs of tracks the Elastic Arms algorithm (EAA) always works better than the Deterministic Annealing Filter (DAF). Furthermore we have shown that the performance of the EAA can be improved upon by introducing a global competition between tracks and hits

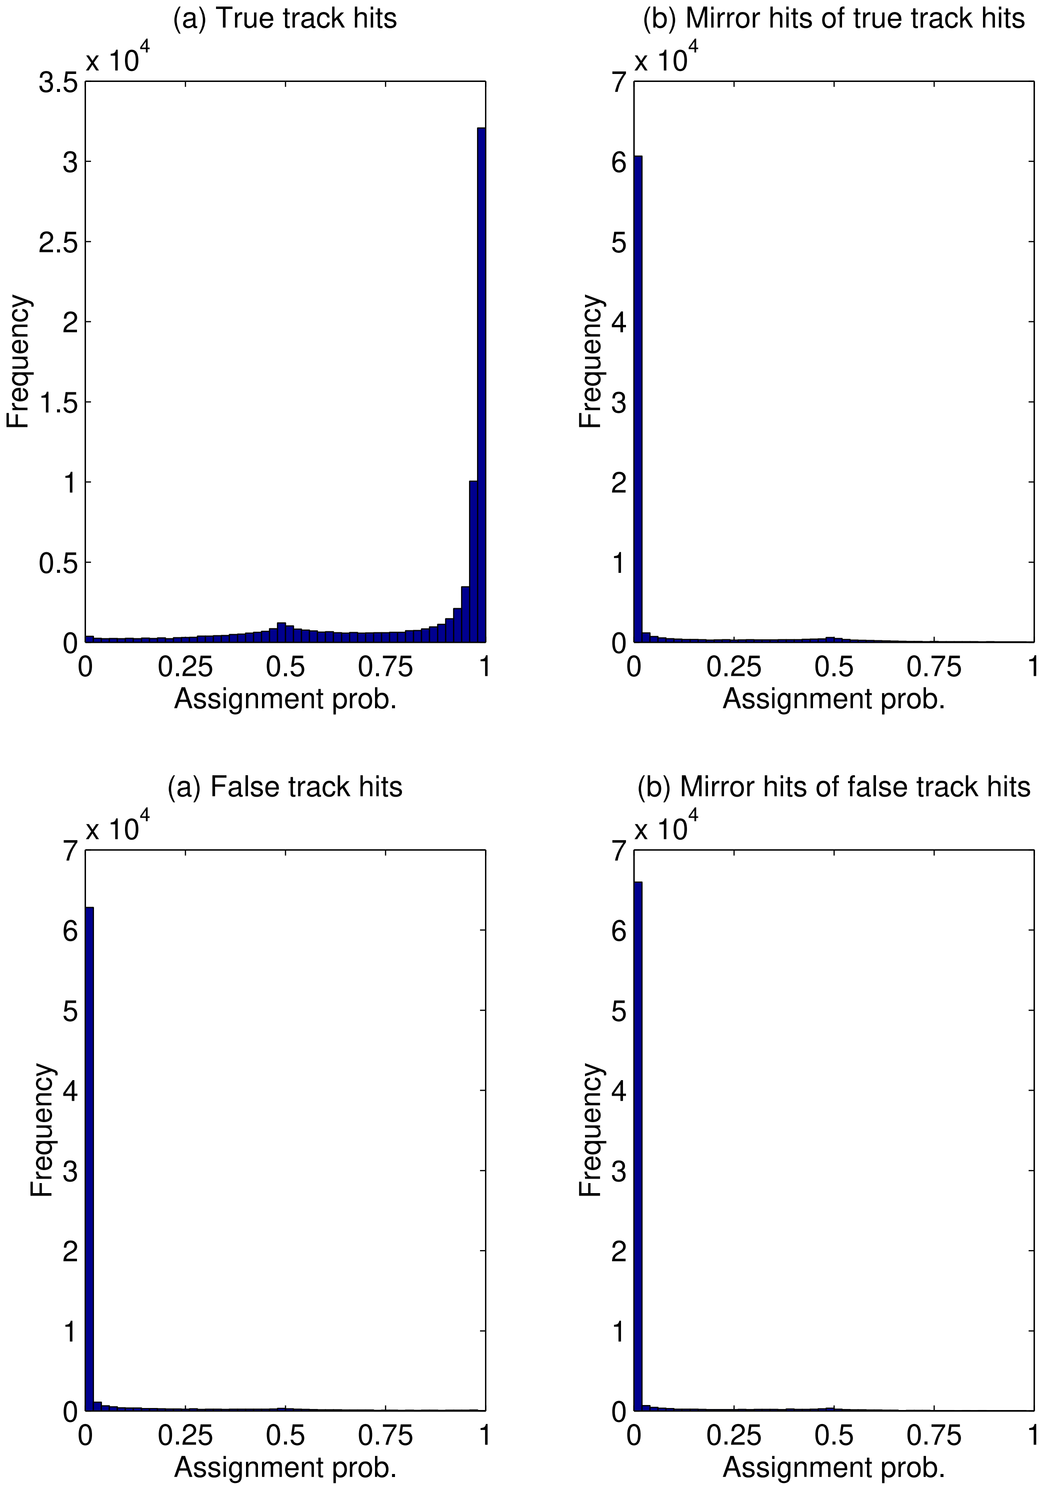

Fig. 3. Histograms of the different assignment probabilities of method 4 with mirror hits and 50% noise level.

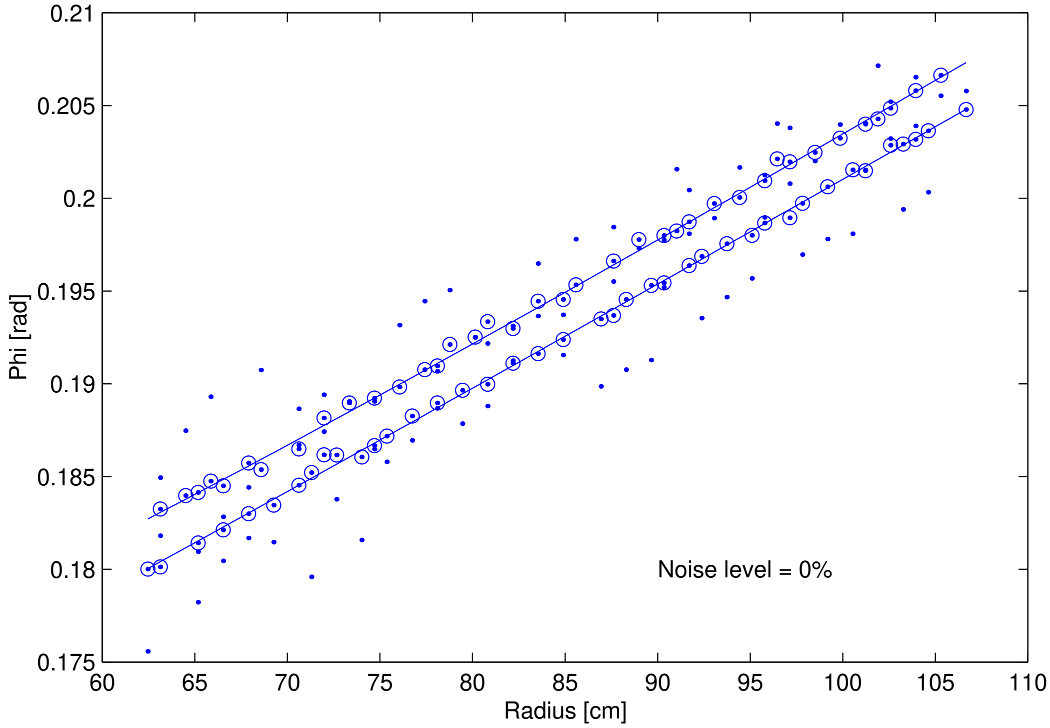

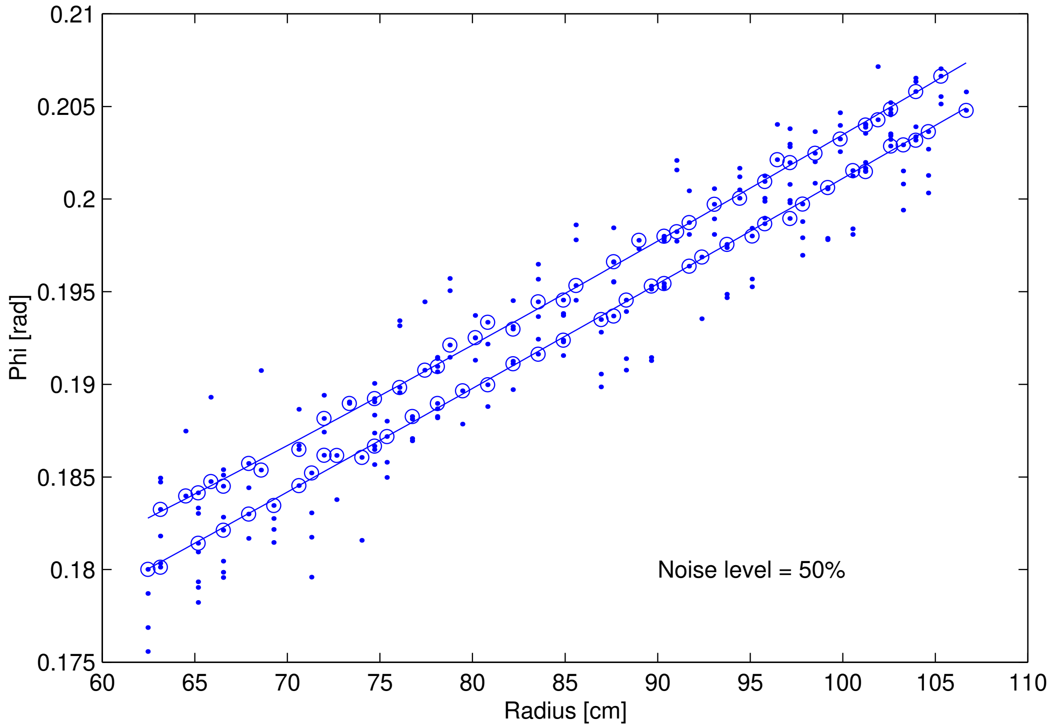

Fig. 4. A track pair in the $R\Phi$-projection. The noise level is 0 in the top figure, and 50% in the bottom figure. The mirror hits are included. The true track hits are marked by open circles. The solid lines are the fitted tracks.

at each detector layer. As can be seen from the definition of the assignment probabilities, the new scheme combines features from both the DAF and the EAA. From an inspection of the results we can draw the following conclusions which we believe to be largely independent of the actual setup:

- The improvement over the EAA is larger for higher noise levels. This is reassuring in view of the potentially very high noise levels in the LHC experiments.
- For the same noise level, the improvement over the EAA is even larger if there are mirror hits. However, in the presence of mirror hits the modified EAA (Method 4) works very well indeed, and the improvement with respect to this baseline is smaller than in the case of no mirror hits.
- Even for the best method the sensitivity to the noise is larger in the presence of mirror hits.
- All of the algorithms are affected when passing from initialization with the true parameters to an initialization which uses only the observations. We do not claim that the method employed in this study is the best one for the two-track case, but would like to stress that the search for good starting values is of paramount importance.

On the other hand, it should be clear that the present work is only a first step in the direction of a true multitrack fit. The question whether the new competition scheme introduced here is in some sense the optimal one remains open. We also do not want to conceal the fact that some non-trivial problems have to be solved before any of the methods studied in this paper can be made operational on real or at least realistic data. Although we cannot give any definite solutions to these problems we can at least offer some hopefully not entirely misguided speculations:

- How many arms or track candidates should be initialized? We have found that a standard track finding method like the Hough transform cannot be applied blindly to close or overlapping tracks. Probably a combinatorial exploration is required to solve this problem. As the multitrack fit will be activated only for a small subset of tracks this should be entirely tolerable as regards computing time.
- The computation of the initial values is closely related to the preceding question. If the number of track candidates is determined by some kind of combinatorial Kalman filter the initial track parameters would be available automatically.
- How should the annealing schedule be chosen? If the tracks are very close and the starting temperature is too high, the arms would be virtually identical after the first iteration and would never separate again. Clearly the starting temperature has to be adaptive as well, in accordance with the average hit density. This could also be found out during the initial exploration phase. Another possibility would be to introduce an explicit repulsion force

between the arms. This force would have to decrease during the annealing procedure.

It is very likely that these problems can not be solved by a generic study like the present one. The details of the initialization procedure and of the annealing schedule heavily depend on the actual detector, the actual noise level and the actual topology of the track bundle and can in all probability only be settled by dedicated simulation studies which go far beyond what has been presented here.

References

[1] *ATLAS Inner Detector Technical Design Report.* CERN/LHCC 97-16, CERN, Geneva, 1997.

[2] R. Frühwirth and A. Strandlie, Track fitting with ambiguities and noise: a study of elastic tracking and non-linear filters. *Computer Physics Communications* 120 (1999) 197.

[3] M. Ohlsson, C. Peterson and A. Yuille, Track finding with deformable templates — the elastic arms approach. *Computer Physics Communications* 71 (1992) 77.

[4] R. Blankenbecler, Deformable templates — revisited and extended — with an OOP implementation. *Computer Physics Communications* 81 (1994) 318.

[5] M. Lindström, Track reconstruction in the ATLAS detector using elastic arms. *Nuclear Instruments and Methods* A357 (1995) 129.

[6] A.P. Dempster, N.M. Laird and D.B. Rubin, Maximum Likelihood from Incomplete Data via the EM Algorithm. *Journal of the Royal Statistical Society B* 39 (1977) 1.

[7] A. Strandlie et al., Particle tracks fitted on the Riemann sphere. *Computer Physics Communications* (in press).

[8] P. V. C. Hough, Machine analysis of bubble chamber pictures. *International Conference on High Energy Accelerators and Instrumentation*, pp. 554-556. CERN, Geneva, 1958.

[9] P. Rousseeuw and A. Leroy, *Robust Regression and Outlier Detection.* John Wiley & Sons, New York, 1987.

Figure captions

Figure 1:
A section of the barrel TRT in the ATLAS Inner Detector. Taken from the ATLAS Inner Detector Technical Design Report [1], with permission.

Figure 2:
A track pair in the $R\Phi$-projection. The noise level is 0 in the top figure, and 50% in the bottom figure. The mirror hits are included.

Figure 3:
Histograms of the different assignment probabilities of method 4 with mirror hits and 50% noise level.

Figure 4:
A track pair in the $R\Phi$-projection. The noise level is 0 in the top figure, and 50% in the bottom figure. The mirror hits are included. The true track hits are marked by open circles. The solid lines are the fitted tracks.

Tables

Noise level	Method		
	1	2	3
0%	251	1.94	2.53
10%	304	2.39	2.66
20%	325	3.17	3.18
30%	233	4.00	3.19
40%	230	5.77	3.41
50%	372	8.71	3.91

Table 1
The relative generalized variance of methods 1–3 without mirror hits and with truth initialization.

Noise level	Method			
	1	2	3	2/3
0%	2.7E6	8.74	43.57	8.22
10%	2.1E6	8.90	31.78	7.39
20%	2.2E6	7.42	27.51	5.01
30%	1.9E6	8.05	22.62	4.73
40%	1.9E6	7.53	24.07	3.56
50%	1.7E6	16.49	33.83	7.14

Table 2
The relative generalized variance of methods 1–3 and method 2/3 without mirror hits and with cog+/− initialization.

Noise level	Method			
	1	2	3	4
0%	281	36.2	4.52	2.84
10%	270	58.7	5.35	4.35
20%	388	100.9	6.26	7.06
30%	358	185.3	7.19	9.51
40%	409	238.9	9.50	12.44
50%	653	301.6	11.66	17.65

Table 3
The relative generalized variance of methods 1–4 with mirror hits and with truth initialization.

Noise level	Method				
	1	2	3	4	4/3
0%	7.6E5	1837	72.2	11.9	15.3
10%	7.5E5	2698	101.6	22.0	22.2
20%	9.7E5	3532	313.6	40.6	33.0
30%	9.8E5	3800	248.6	43.5	32.6
40%	9.3E5	4449	193.0	65.9	45.7
50%	1.1E6	4817	388.4	96.0	68.8

Table 4
The relative generalized variance of methods 1–4 and method 4/3 with mirror hits and with cog+/− initialization.