

RESPONSE SURFACE MODEL EVOLUTION

Alexander I J Forrester*, Neil W Bressloff†, Andy J Keane‡
Computational Engineering and Design Centre
School of Engineering Sciences
University of Southampton SO17 1BJ
UK

ABSTRACT

Methods are presented for reducing time and effort when performing aerodynamic optimisation using response surface models. Significant time savings are made possible by monitoring the convergence of computational fluid dynamics simulations and omitting regions of poor designs. In so doing, optimal regions of the design space can be highlighted and surface refinement commenced early in the convergence of the design point set. A strategy employing surface updates with new data at points of maximum expected improvement is shown to perform more efficiently than reducing the design space to the region of the optimum. The response surface evolution methods are demonstrated through an example two parameter optimisation of a flap track fairing on a commercial airliner wing.

INTRODUCTION

Aerodynamic shape design can be optimised by simulating the flow over a number of geometries throughout a design space so as to construct a response surface model (RSM). This model is used as a surrogate instead of new simulations when searching for the optimum. Since the shape of the RSM is not known *a priori*, the same amount of computing effort is normally applied to all points when building the RSM. However, significant time savings can be achieved if more effort is directed at finding the optimum rather than modeling regions of poor designs.

A 'reasonable design space approach' has been used by Balabanov et al.¹ and Giunta et al.² to reduce the design space using a low fidelity analysis prior to evaluation with a high fidelity analysis. This approach relies on the availability of a suitable low fidelity model. Often this model will be empirically based and may not be applicable for new areas of research or radical designs.

Modern parallel computing architecture makes it possible to update the RSM at many points simultaneously. A sparsely populated design space can be updated in promising areas by using parallel updates at maximum expected improvement³ optima. Sobester et al.⁴ have shown that this method enhances the efficiency of RSMs by sampling points in areas of promising designs.

The RSM schemes mentioned above are based around finding and searching the area of the optimum whilst reducing the effort spent in areas of poor designs. The expected improvement approach is attractive because it does not exclude areas of the design space but, rather, intelligently selects promising regions. When parts of the design space are excluded there is a chance that the global optimum could be excluded at an early stage.

This paper presents the use of partially converged computational fluid dynamics (CFD) results when building RSMs. When RSMs are constructed at intervals throughout the convergence history of the CFD simulations, it is seen that all design points converge at a similar rate and information about the nature of the design space can be obtained early in the iterative process. Regions of high expected improvement can be found when just the broad shape of the RSM is known, without fully converging the CFD data.

*Graduate Research Student, School of Engineering Sciences

†Senior Research Fellow, School of Engineering Sciences

‡Professor of Computational Engineering, School of Engineering Sciences





Figure 1: FTF design variables

The following section describes the example problem used throughout the paper to demonstrate the methods used. In further sections the data and RSMs produced by partially converged CFD are examined and the techniques discussed are applied to the example problem.

FLAP TRACK FAIRING EXAMPLE OPTIMISATION

Possible schemes for monitoring the convergence of RSMs are presented through the optimisation of the Euler (inviscid) calculated lift/drag ratio ($\frac{L}{D}$) of a flap track fairing (FTF) on a commercial airliner wing (see figure 1).

The FTF geometry and flow domain is defined in a parametric computer aided design (CAD) programming language. Meshing and flow simulation are performed with the commercial CFD package, FluentTM. The complete process from design point selection using design of experiment (DoE) methods through to building the RSM is performed by an automated system. Automating the optimisation process makes searching large areas of the design space with many variables and many points feasible and is the only way to find truly optimal designs⁵.

PARALLEL COMPUTING CONSIDERATIONS

Response surface evolution methods have been investigated using the University of Southampton's 404 processor Beowulf cluster. To monitor the convergence of simulations it is necessary to have data for all the designs. Ideally every design point would be running in parallel. This is not possible for large DoEs due to the constraints of the job scheduler (a maximum of 15 jobs using a total of 120 processors can be run simultaneously). A more feasible strategy is to run the simulations a number of iterations at a time. Each design point is submitted in turn with data being collected and monitored when a full set of simulations is complete. For DoEs of 15 points or less all points can be run in parallel to each other.

EVOLUTION

A two variable optimisation is considered with the yaw angle and overhang of the FTF behind the wing trailing edge as the variables (see figure 1). Yaw varies from -15° (fairing points away from fuselage) to 15° and overhang from 10 to 40% of the fairing length. A yaw of more than 10° away from the fuselage violates the geometry description and points in this region are not simulated. An LP_τ DoE⁶ is used throughout this paper to investigate possible RSM evolution methods. Varying numbers of points are required and an LP_τ allows points to be added whilst maintaining a space filling DoE throughout. Jones et al.³ use the rule of thumb of 10 times the number of variables as the size of an initial DoE to construct a RSM. The methods investigated will be compared to a baseline RSM built from a 20 point LP_τ DoE. The force coefficients for these 20 CFD simulations converged after approximately 250 iterations. The value and position of the optimum at stages throughout the convergence of the CFD are shown in table 1. It can be seen that the approximate position of the optimum is known early in the simulation after just 50 iterations. The position seems to converge no faster than the objective function value, but it should be remembered that a far higher accuracy is required in the final objective function than that needed in the position to allow updates or reduction of the domain.

RSM approximations using the kriging method (Jones et al.³ have used kriging in global optimisation) are shown in figure 2. A RSM is plotted for

ITERATION	OVERHANG	YAW	$\frac{L}{D}$
10	0.400	15.000	4.083
50	0.354	10.386	8.612
100	0.388	9.375	10.405
150	0.400	9.375	10.429
200	0.395	9.375	10.504
250	0.398	9.375	10.526

Table 1: Successive searches of the 20pt RSM

every 10 CFD iterations. The hyper-parameters of the surface are trained by optimising the concentrated likelihood function. It is infeasible and contrary to the envisaged advantage of the evolution approach to train the hyper-parameters individually by hand for each surface. Instead a genetic algorithm followed by a dynamic hill climber are used in the same manner for all the kriging surfaces to ensure the hyper-parameters are close to the optimum. These methods are available in the Options design exploration system⁷.

Consideration of the RSM evolution shown in Figure 2 shows that the contours of the surface develop a characteristic shape very early in the process and, in fact, the region of the optimum could be found, though with little confidence, after only 10 iterations. The plots are all on the same scale so the convergence of the magnitude of the objective function can be seen as the plots tend towards white. With all the surfaces available for comparison, as in Figure 2, it is clear when the shape of the surface has stabilised, but with no prior knowledge, or if the surface has more than 2 dimensions, a means of measuring the stabilisation of the shape of the surface is required.

The correlation coefficient r^2 , which is usually used to give the quality of a least squares fitting to experimental data, is used as a measure of the quality of one approximation (\hat{f}_n) after n iterations to that of another (\hat{f}_{n-m}), m iterations earlier. The correlation coefficient is defined as follows:

$$r_m^2 = \left(\frac{N \sum \hat{f}_n \hat{f}_{n-m} - \sum \hat{f}_n \sum \hat{f}_{n-m}}{\sqrt{[N \sum \hat{f}_n^2 - (\sum \hat{f}_n)^2][N \sum \hat{f}_{n-m}^2 - (\sum \hat{f}_{n-m})^2]}} \right)^2$$

where N is the number of points to be correlated. This is either the size of the DoE if the data from a simulation is to be correlated or a number of points taken from a RSM for a surface

correlation. This measure of the convergence of the shape of an RSM can be exploited in the strategy shown in Figure 3.

The geometry is parameterised and a DoE constructed. The CFD simulations are then run, but after m iterations (this number is problem dependent and in this paper $m=10$ iterations), the data is read into a data-base and a response surface approximation of the partially converged data is constructed. The CFD is continued and after two loops of this cycle the correlation between sets of data m iterations apart can be calculated. When a high correlation is found the optimisation might take one of two routes: 1) expected improvement updates can be employed or 2) the domain can be reduced to the area around the optimum.

Method 1: The expected improvement (EI) function is searched with a gradient descent method using a number of restarts to find points of maximum EI. New simulations are run to the same number of iterations at these points and a new RSM is constructed. This process can be repeated for as many updates as necessary or time and resources permit. The simulations are then fully converged before a final search for the optimum.

Method 2: The RSM is searched and the domain reduced to the area around the optimum. This reduced domain is now updated with new CFD simulation data (run for the same number of iterations) at the optimum. After two loops through this cycle the search result can be compared with the update point. If the results match to within a certain tolerance the final loop can be entered, where the points in the reduced domain are iterated to full convergence.

Two correlations can be considered, the RSM correlation and the raw data correlation (the former compares the entire approximated surfaces, while the latter just deals with the data obtained directly from the CFD code). The kriging surface correlation for 10, 30 and 60 points are plotted up to 250 iterations in figure 4. The correlation interval is 10 iterations so the value at, for example, 50 iterations refers to the correlation between the surface at 50 iterations and 40 iterations. The kriging correlation is somewhat erratic early in the simulation, particularly for the 10 point data set. The correlation is affected by the approximation model used. The correlation falls sharply when the kriging hyper-parameters change. When more data points

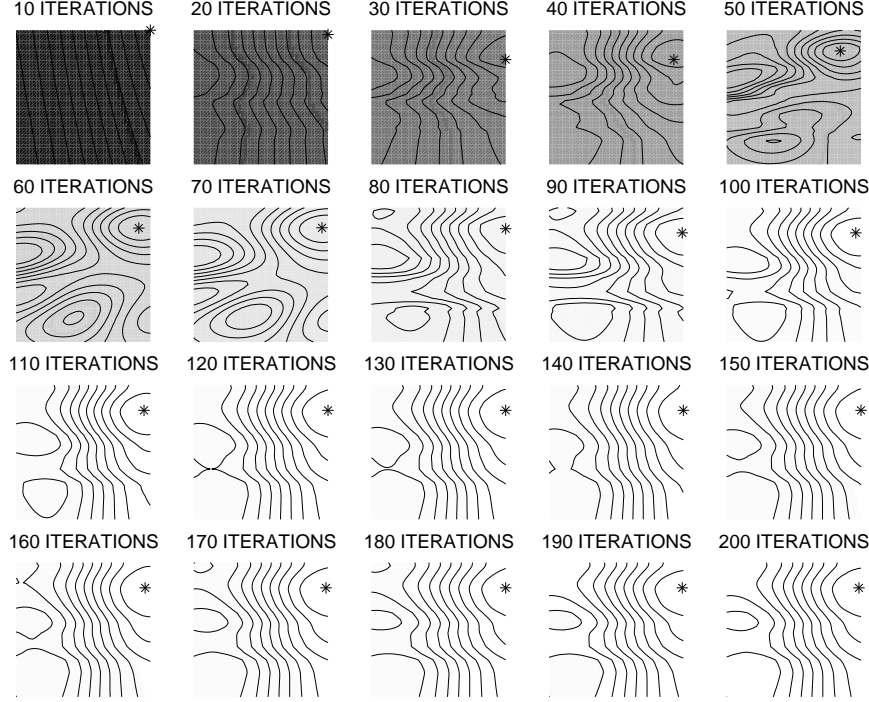


Figure 2: Evolution of the 20pt RSM

are used the hyper-parameters are more stable and the surface correlation convergence is improved. It seems that the surface correlation is an indicator of kriging's ability to represent the data - large drops in the correlation suggest a poorly trained approximation.

A polynomial regressed RSM requires no training and should therefore not be as erratic as an automatically trained kriging approximation. However the correlations in Figure 5 suggest that the polynomial surface is struggling to fit the multi-modal data in the early stages of the simulations, but the correlation does show smooth convergence for the 10 point data set after 70 iterations. However, as in the early stages of the simulations presented, a polynomial model is unable to approximate multi modal problems and is likely to perform less well for a different or multi-dimensional problem.

The data correlations shown in Figure 6 are, of course, independent of the approximation method. The drop in correlation at 50 iterations which is seen for both surface correlations is also present in the data correlation, indicating that there is an anomaly in the convergence of the data between

40 and 50 iterations. Inspection of the force coefficient history files for each design point confirm that one point shows a different convergence rate at this stage. The data correlation is higher and less erratic for the 10 point data set than either surface correlations. As more data points are added the data correlation drops slightly – naturally there is more scope for error in more points, but anomalies in the convergence of individual points have less of an impact on the overall correlation.

Because the RSM will be used when finding objective function optima, it is naturally a surface correlation that will give the better indicator of when updating or design space reduction can begin. However, for a low number of initial design points, data correlation give a better indicator of the stabilisation of the convergence of the simulations at an early stage in the iterations.

RSM evolution can be applied to a reasonable design space approach and a parallel updating scheme. The value that the correlations should reach before updating or reducing the domain depends on the size of DoE and the accuracy required. Different convergence criteria are applied to the ex-

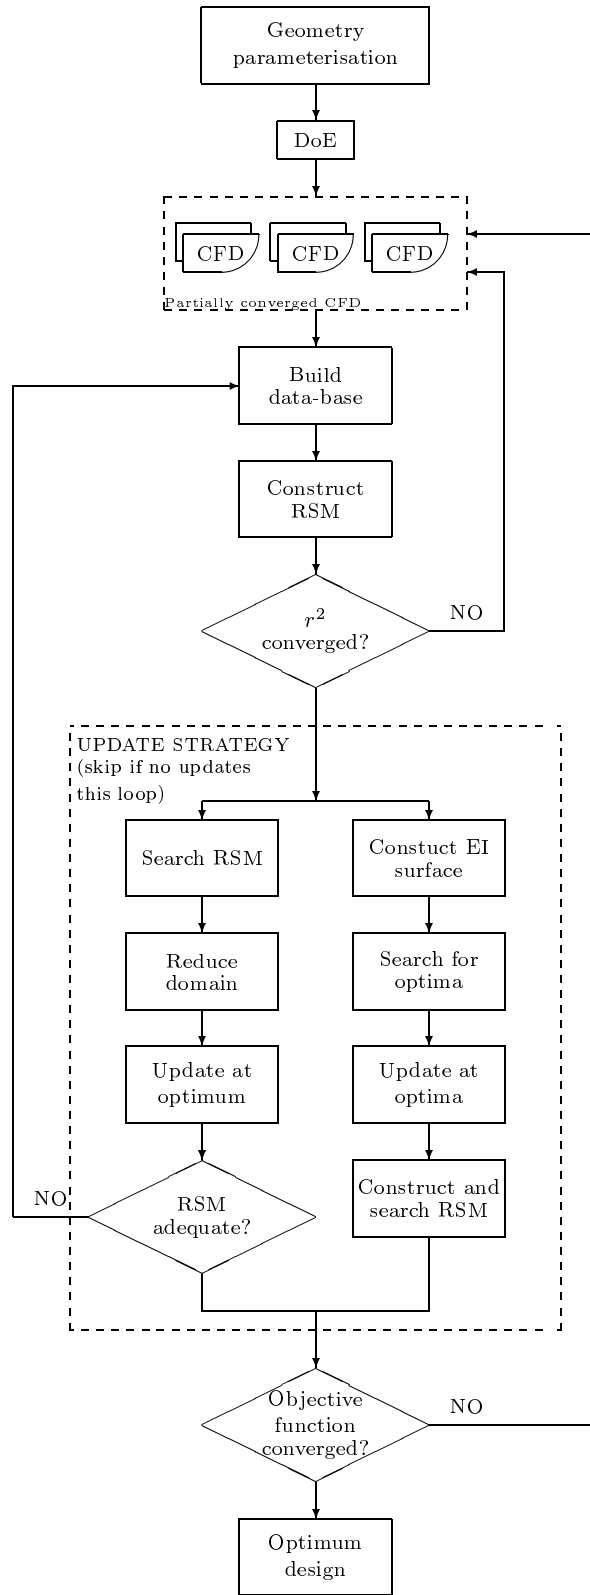


Figure 3: Search strategy




Figure 4: Kriging surface correlation




Figure 5: Polynomial surface correlation




Figure 6: Data correlation

ample optimisations later in the paper. If a parallel updating scheme is to be used then a small DoE (i.e., smaller than the 10 times number of variables rule of thumb) and the quality of the RSM through the initial data points is not crucial because the RSM gains in accuracy as more data is added. Data correlations are appropriate for use in determining when to halt the CFD before adding points. If the domain is to be reduced to a small region around the region of the optimum then a large initial DoE is needed and the surface correlation should reach a high value before clipping the domain and updating.

The number of iterations between checking the correlation should not be too small or too large. The correlation between data just one iteration apart will naturally be very high and time is wasted if early convergence is missed through letting too many iterations pass. The number chosen here is largely a tradeoff between CFD run time and restart time. If a finely meshed Navier-Stokes simulation is being computed, the time taken for restarts is negligible compared to run time and the convergence should be monitored regularly. However, a simple Euler simulation may be significantly delayed through regular stopping and restarting.

POSSIBLE OPTIMISATION METHODS

BASELINE FULLY CONVERGED 20 POINT DATA SET

Three possible methods will be compared to the baseline RSM (built from a 20 point LP_τ DoE) optimisation described above. A search on the final plot in Figure 2 gave an optimum $\frac{L}{D}$ of 10.53 at 39.8% overhang and 9.38° yaw. The final RSM, along with the optimum depicted by an asterisk, is shown in Figure 7. The total required wall clock time, assuming parallel computation to the extent possible on the computing cluster being used (as described above), is 3 hours.

METHOD 1: EXPECTED IMPROVEMENT UPDATES

The 10 point LP_τ set of designs are computed until the 10 iteration data correlation (r_{10}^2) is greater than 0.95. This criterion is satisfied after 40 iterations. A RSM of the expected improvement is constructed using this data with the kriging method. A gradient descent search started at each of the DoE points yields 5 local maximum expected improvement points. The surface is updated at these




Figure 7: 20pt RSM after 250 iterations

points. All points are now run until 250 iterations and the kriging RSM searched for the optimum. The final RSM in Figure 8 shows the optimum at 38.13% overhang and 7.5° yaw.

Significant time savings can be made with this method if sufficient processors are available because updates can be started before the initial simulations are converged. Each design point takes $1\frac{1}{2}$ hours across 6 processors. If parallel updates are performed sequentially as done by Sobester et al.⁴, i.e., the initial data set is fully converged before updates begin, the wall clock time is 3 hours. In this case the updates are started in parallel alongside the initial data set after 40 iterations. Although the total CPU time is the same, the wall clock time is now 1.74 hours – a 42% time saving.

METHOD 2: REDUCE DOMAIN AND UPDATE

The 60 point LP_τ set of designs are used to approximate the region of the optimum to allow the design space to be reduced. The correlation criterion, (r_{10}^2) is set to 0.95, as before except this time surface correlations are used since the data set is large. The surface reaches this criterion after 30 iterations (the surface at this stage is shown in Figure 9). The design space is reduced by 50% in both dimensions around the optimum, i.e., to the top right corner of Figure 9. 20 points remain in the reduced domain, which means there is no requirement for re-seeding with more points at this stage. Had a smaller initial DoE been used it would have been necessary to add points. However, the extra




Figure 8: 10pt RSM with 5 expected improvement updates

excluded points of this large DoE consumed little time and the process has been simplified considerably. There is also more confidence in the region of the optimum having used more data in the initial RSM.

The 20 points in the reduced domain are now run until 250 iterations, giving a final wall clock time of 3.36 hours. The final RSM constructed using the kriging method is shown in Figure 10 with the optimum depicted by an asterisk. It is clear by comparing Figure 10 to Figure 9 that the kriging method is not approximating the objective function well. There is some noise in the data and the variation in $\frac{L}{D}$ is small (as shown in the colour bar, there is 0.25% variation across the surface). This has led to the surface being over regressed in the yaw direction. However, the optimum of the surface is in a valid position (10.53 at 39.8% overhang and 9.38° yaw) when compared to the other methods and the process may now be continued as shown in Figure 3, with further reduction of the domain and updating at maximum expected improvement.

METHOD 3: SEARCH ON PARTIALLY CONVERGED RESULTS

Figure 4 shows high surface correlations for a large number of data points, implying that it may be possible to carry out optimisation on the partially converged results alone and only fully converge one simulation at the optimum. The 60 point LP_τ DoE




Figure 9: 60pt RSM after 30 iterations




Figure 11: 60pt RSM after 130 iterations




Figure 10: Reduced domain 20pt RSM

designs are used to construct a kriging RSM. The surface correlation criterion is set to $r_{10}^2 > 0.99$ to ensure the surface has stabilised. This criteria is met after 130 iterations and the resulting surface is shown in Figure 11. All the data points were continued to 250 iterations to allow a comparison of the partially converged surface with the corresponding converged surface. The correlation r_{120}^2 at 250 iterations of 0.996 shows that a good approximation to the position of the optimum can be made. A search on the 130 iteration surface gives the position of the optimum. This point must then be simulated to 250 iterations to give the $\frac{L}{D}$ value at that point. The final wall clock time is 3.52 hours.

The high correlation between the 130 and 250 iteration surfaces suggest that perhaps the correla-

tion criterion is too stringent. However a criteria of $r_{10}^2 > 0.98$ is met after 70 iterations and the 70 to 250 iteration r_{180}^2 correlation is 0.9272. This is low for a 60 point kriging correlation (recall Figure 4). The optimum of the 70 iteration surface is at 40% overhang and 15° yaw. This does not compare well with results from the 130 iteration surface and the other methods researched above.

COMPARISON OF PROPOSED METHODS

Table 2 is a comparison of the three RSM evolution methods investigated along with the 20 point baseline optimisation and the 60 point fully converged RSM constructed when investigating method 3. The three methods investigated all perform better than the baseline RSM insofar as a higher optimum is found. The region of the optimum of the FTF design space is very flat so there is little difference between the optimum $\frac{L}{D}$ of the different methods. Method 1 is the most attractive in terms of time saving. The calculations for the wall clock time taken assume that 15 design points are running at all times. Method 2 could be more attractive in terms of time if the design space were reduced further, but there is a chance that the optimum could be excluded, particularly if the surface had a number of local optima. Method 1 excludes none of the design space and can therefore cope well with multiple optima. Method 3 offers no time saving over the 20 point baseline optimisation, but displays that the shape of the surface is accurately predicted with partially converged simulations when

Method	O'hang	Yaw	Opt $\frac{L}{D}$	CPU hrs	Wall hrs
60 LP_τ	0.400	7.968	10.519	540	6
20 LP_τ	0.398	9.375	10.526	180.0	3
1	0.381	7.500	10.527	135	1.74
2	0.400	7.500	10.541	222	3.36
3	0.400	8.202	10.585	290	3.52

Table 2: Comparison of evolution methods

compared with the 60 point fully converged surface.

Method 1 appears to offer the best compromise of speed and accuracy, and ensures a global optimum will be found even when applied to more complex RSMs.

CONCLUSIONS

Methods are presented for reducing time and effort when performing aerodynamic optimisation using response surface models. Significant time savings are possible through monitoring the convergence of CFD simulations and reducing computational effort in regions of poor designs. Monitoring of the simulations can utilise both data and surface correlations to predict when the shape of the RSM is sufficiently stable to allow updating at optima or reduction of the domain. A strategy employing parallel expected improvement updates initiated early in the convergence of the simulations of the initial DoE is shown to perform more efficiently than domain reduction.

ACKNOWLEDGEMENTS

This work has been supported by the University Technology Partnership, a collaboration between Rolls-Royce, BAE Systems and the Universities of Sheffield, Cambridge and Southampton. The authors also thank S. J. Leary for his input.

REFERENCES

- [1] V. O. Balabanov, A. A. Giunta, O. Golovidov, B. Grossman, W. H. Mason, L. T. Watson, and R. T. Haftka. Reasonable design space approach to response surface approximation. *Journal of Aircraft*, 36(1):308–315, Jan-Feb 1999.
- [2] A. A. Giunta, R. Narducci, S. Burgee, B. Grossman, W. H. Mason, L. T. Watson, and R. T. Haftka. Variable-complexity response surface aerodynamic design of an hsct wing. *13th AIAA Applied Aerodynamics Conference San Diego California*, 19-22 June 1995.
- [3] D. R. Jones, M. Schlonlau, and W. J. Welch. Efficient global optimisation of expensive black-box functions. *Journal of Global Optimisation*, 13:455–492, 1998.
- [4] A. Sobester, S. J. Leary, and A. J. Keane. A parallel updating scheme for approximating and optimizing high fidelity computer simulations. *3rd ISSMO/AIAA Internet Conference on Approximations in Optimization*, 2002.
- [5] A. Jameson and J. C. Vassberg. Computational fluid dynamics for aerodynamic design: Its current and future impact. *39th AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV*, 8-11 January 2001.
- [6] I. M. Sobol. On the systematic search in a hypercube. *SIAM Journal of Numerical Analysis*, 16:790–793, 1979.
- [7] A. J. Keane. The options design exploration system reference manual and user guide - version b3.1. <http://www.soton.ac.uk/~ajk/options.ps>, 2002.