

Parallel Particle Swarm Optimization Algorithm Accelerated by Asynchronous Evaluations

Gerhard Venter*

Vanderplaats Research and Development, Inc., Colorado Springs, CO 80906

and

Jaroslav Sobieszczanski-Sobieski†

NASA Langley Research Center, MS 155, Hampton, VA 23681-2199

An asynchronous parallel Particle Swarm Optimization (PSO) algorithm is presented. Particle swarm optimization is a fairly recent addition to the family of non-gradient based, probabilistic search algorithms. Particle swarm optimization is based on a simplified social model and is closely tied to swarming theory. Although PSO algorithms present several attractive properties to the designer, they are plagued by high computational cost as measured by elapsed time. One approach to reduce the elapsed time is to make use of coarse-grained parallelization to evaluate the design points. Previous parallel PSO algorithms were exclusively implemented in a synchronous manner, where all design points within a design iteration are evaluated before the next iteration is started. This approach leads to poor parallel speedup in cases where a heterogeneous parallel environment is used and/or where the analysis time depends on the design point being analyzed. This paper introduces an asynchronous parallel PSO algorithm that is benchmarked on a cluster assembled of Apple Macintosh G5 desktop computers, using the multi-disciplinary optimization of a typical transport aircraft wing as an example. The asynchronous parallel PSO algorithm greatly improves the parallel efficiency for the considered example problem.

I. Introduction

PARTICLE Swarm Optimization (PSO) is a fairly recent, but rapidly growing, addition to an expanding collection of non-gradient based, probabilistic search algorithms. Some widely used algorithms that fall into this category are genetic algorithms (e.g., Michalewicz and Dugupta)¹ that model Darwin's principle of survival of the fittest and simulated annealing algorithms (e.g., Nemhauser and Wolsey)² that model the equilibrium of large numbers of atoms during an annealing process. This class of optimization algorithms provides the designer with several attractive characteristics. For example, these algorithms are generally easy to implement, can efficiently make use of large numbers of parallel processors, do not require continuity in response functions and are better suited for finding global or near global solutions. Although these non-gradient based algorithms provide the designer with several advantages, they should be applied with care. Due to their high computational cost, these algorithms should only be used when a gradient-based algorithm is not a viable alternative, such as integer/discrete and discontinuous problems and problems with multiple local minima.

Many non-gradient based search algorithms are based on some natural phenomena, and PSO is no exception. Particle swarm optimization is based on a simplified social model that is closely tied to swarming theory and was

Received 25 May 2005; revision received 16 December 2005; accepted for publication 23 December 2005. Copyright 2006 by Gerhard Venter and Jaroslav Sobieszczanski-Sobieski. Published by the American Institute of Aeronautics and Astronautics, Inc., with permission. Copies of this paper may be made for personal or internal use, on condition that the copier pay the \$10.00 per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923; include the code 1542-9423/04 \$10.00 in correspondence with the CCC.

* VisualDOC Project Manager, Senior AIAA Member, e-mail: gventer@vrand.com

† Senior Research Scientist, AIAA Fellow, e-mail: j.sobieski@larc.nasa.gov

first introduced by Kennedy and Eberhart.^{3,4} A physical analogy might be a school of fish that is adapting to its environment. In this analogy, each fish adapts to its environment by making use of its own memory as well as knowledge gained by the school as a whole.

Although PSO is a fairly recent algorithm, it is quickly gaining momentum in the engineering research community. For example Fourie and Groenwold^{5,6} applied the algorithm to structural shape and sizing and topology optimization problems. The authors^{7,8} applied the algorithm to a cantilevered beam and multi-disciplinary design optimization of a transport aircraft wing. The University of Florida has applied the algorithm to biomechanical system identification problems e.g., Schutte et al.^{9,10}

Some authors have worked on parallel PSO algorithms in an attempt to alleviate the associated high computational cost. All parallel implementations of the PSO algorithm presented to date are based on a synchronous implementation e.g., Schutte et al.^{9,10} where all design points within a design iteration are evaluated, before the next design iteration is started. One exception is a paper by Koh et al.¹¹ that also deals with an asynchronous parallel PSO algorithm. Although Koh et al.¹¹ published their paper concurrently with the present paper, the research was performed independently. The synchronous implementation can easily lead to poor parallel performance in cases where all processors are not fully utilized to evaluate all points within a design iteration, resulting in some idle processors at the end of each design iteration. This can easily happen in many cases, for example a heterogeneous parallel environment or when the analysis time depends on the design point being analyzed. The present work presents a new asynchronous parallel implementation to increase the parallel performance of the PSO algorithm in cases where the synchronous implementation would lead to idle processors at the end of each design iteration. The challenge is to keep all processors working all of the time as the algorithm moves from one design iteration to the next, without loss of numerical accuracy.

In the present work, both synchronous and asynchronous parallel versions of the PSO algorithm are presented and compared in terms of numerical accuracy and parallel efficiency. The multi-disciplinary design optimization of a typical transport aircraft wing is used as a test problem. This example makes use of a bi-level approach to perform the system level optimization of the wing for maximum range, accounting for the trade-off between the aerodynamic drag and the structural weight and is similar to that presented in Venter and Sobieszczanski-Sobieski.⁸ In this formulation, the aerodynamic optimization is performed at the system level, and the structural optimization is considered as a sub-problem at the discipline level. The motivation for using PSO in the present design problem is the presence of discrete design variables and severe numerical noise, which makes the use of a gradient-based optimizer impractical.

II. Particle Swarm Optimization Algorithm

Particle swarm optimization is based on the social behavior that a population of individuals adapts to its environment by returning to promising regions that were previously discovered, Kennedy and Spears¹². This adaptation to the environment is a stochastic process that depends on both the memory of each individual as well as knowledge gained by the population as a whole. In the numerical implementation of this simplified social model, the population is referred to as a swarm and each individual as a particle.

The numerical implementation starts with an initial swarm of randomly distributed particles, each with a random initial velocity vector. The algorithm then repeatedly updates the position of each particle over a time period to simulate the adaptation of the swarm to its environment. The position of each particle is updated using the current position, a velocity vector and a time increment. The new position of each particle at iteration $k + 1$ is calculated from Eq. (1)

$$\mathbf{x}_{k+1}^i = \mathbf{x}_k^i + \mathbf{v}_{k+1}^i \Delta t, \quad (1)$$

where: \mathbf{x}_{k+1}^i is the position of particle i at iteration $k + 1$; \mathbf{v}_{k+1}^i is the corresponding velocity vector; and Δt is the time step value. A unit time step is used throughout the present work.

The velocity vector of each particle can be obtained from many different formulations. A formulation that is widely used in the literature was introduced by Shi and Eberhart¹³ and is shown in Eq. (2)

$$\mathbf{v}_{k+1}^i = w\mathbf{v}_k^i + c_1r_1 \frac{(\mathbf{p}^i - \mathbf{x}_k^i)}{\Delta t} + c_2r_2 \frac{(\mathbf{p}_k^g - \mathbf{x}_k^i)}{\Delta t}. \quad (2)$$

In Eq. (2), r_1 and r_2 are random numbers between 0 and 1, p^i is the best position found by particle i so far, and p_k^g is the best position in the swarm during iteration k . There are three problem dependent parameters, the inertia of the particle (w), and two “trust” parameters c_1 and c_2 . The inertia controls the exploration properties of the algorithm, with larger values facilitating a more global behavior and smaller values facilitating a more local behavior. The trust parameters indicate how much confidence the particle has in itself (c_1) and how much confidence it has in the swarm (c_2).

A slightly different formulation of the velocity vector (e.g., Fourie and Groenwold)⁵ is to replace the best position of the current iteration p_k^g with the best position found so far p^g . Both formulations will be considered here.

Enhancements to the basic algorithm include dealing with constrained optimization problems, a craziness operator, and dynamically changing the inertia value. These enhancements are discussed in more detail by Venter and Sobieszcanski-Sobieski.^{7,8} Constrained optimization problems are typically handled using a penalty function approach and modifying the velocity vector by dropping the inertia term for violated design points. In the present work, an unconstrained problem is considered and the penalty function approach is not used. However, the velocity vector is modified for design points that fall outside of the design space. These design points are reset to the nearest design space bound and the inertia component is dropped from the velocity vector. The craziness operator adds additional randomness to the swarm to avoid premature convergence and is similar to the mutation operator of a genetic algorithm. A small subset of particles are selected at each design iteration, and the position and/or velocity vectors are randomly modified. As mentioned, the inertia parameter controls the exploration properties of the algorithm. The goal is to start with a larger inertia value (a more global search) that is dynamically reduced towards the end of the optimization (a more local search).

III. Parallel Particle Swarm Optimization Algorithm

The PSO algorithm is ideally suited for a coarse-grained parallel implementation on a parallel or distributed computing network. For each time step (design iteration), all particles (design points) are independent of each other and can be easily analyzed in parallel.

A. Synchronous Parallel Particle Swarm Optimization Algorithm

The most obvious PSO parallel implementation is to simply evaluate the design points within a design iteration in parallel, without changing the overall logic of the algorithm itself. In this implementation, all particles within a design iteration are sent to the parallel computing environment, and the algorithm waits for all the analyses to complete before moving to the next design iteration. This implementation is referred to as a synchronous implementation.

Although the synchronous implementation provides an easy extension to a parallel environment, it is not an ideal solution and may result in poor parallel efficiency in some cases, as shown in Schutte et al.¹⁰ The problem with this implementation is that it is nearly impossible to keep all processors working towards the end of each design iteration, before the next iteration is started. There are three common situations that would result in some of or most of the processors being idle towards the end of each design iteration, as follows.

1. Having a swarm size that is not an integer multiple of the number of processors.
2. A heterogenous distributed compute environment where processors with varying computational speed are combined into a parallel compute environment (e.g., Venter and Watson¹⁴).
3. Using a numerical simulation to evaluate each design point, where the required simulation time depends on the design point being analyzed.

The influence of having idle processors towards the end of each design iteration on the parallel efficiency is amplified when more processors are used. As a result, one can expect a reduction in the parallel efficiency of a synchronous PSO algorithm when more processors are used.

B. Asynchronous Parallel Particle Swarm Optimization Algorithm

The above limitations in the synchronous parallel PSO algorithm can be overcome by considering an asynchronous algorithm where design points in the next design iteration are analyzed before the current design iteration is completed. The goal is to have no idle processors as one moves from one design iteration to the next.

The key to implementing an asynchronous parallel PSO algorithm is to separate the update actions associated with each point and those associated with the swarm as a whole. These update actions include updating the inertia

value, the craziness operator, and the swarm and point histories. Note that the same PSO operators are used in the synchronous and asynchronous algorithms, the only difference is that they are executed at different times. For the synchronous algorithm, all the update actions are performed at the end of each design iteration. For the asynchronous algorithm, we want to perform the point update actions after each design point is analyzed, and the swarm update actions at the end of each design iteration. The parts of the algorithm that need to be considered when looking at the update actions are the velocity vector, the craziness operator and the dynamic reduction of the inertia value.

The velocity vector is the centerpoint of any PSO algorithm. For each design point, the velocity vector is updated using the following dynamic properties for that point: the previous velocity vector; the current position vector; and the best position found so far. In addition, the updated inertia value and the best position for the swarm as a whole are also required. To do the velocity update in an asynchronous fashion, we need to update the position vector and the best position found so far for each design point directly after evaluating that point. For the best position in the swarm, we have two choices: use the best position in the current iteration (p_k^g), or use the best position found so far (p^g). To keep the best position for the swarm current when moving to the next design iteration, before the current iteration is completed, it is necessary to use the best position found so far rather than the best position in the current iteration. This setup allows the algorithm to update all required dynamic properties of the velocity vector directly after evaluating each design point, except for the inertia value. The inertia value is the only iteration level update required to compute the velocity vector and is updated at the end of each design iteration. The craziness operator is the only other iteration level update and is also performed at the end of each design iteration. For the asynchronous algorithm the end of a design iteration is assumed each time a number of analyses equal to the swarm size is completed. Of course, this could result in some points of the next design iteration being analyzed before the inertia and craziness operators are applied for that design iteration. However, the influence on the overall performance of the algorithm seems to be negligible for the example problem considered here.

The asynchronous algorithm is thus very similar to the synchronous algorithm, except that we update as much information as possible after each design point is analyzed. The inertia and craziness operators are only applied when a design iteration is completed.

C. Parallel Implementation

In the present work, both synchronous and asynchronous parallel PSO algorithms are implemented and compared. The parallel scheme used here is based on the Message Passing Interface (MPI) to provide a master-worker implementation. The master-worker scheme implemented here provides dynamic load balancing between the processors. In this implementation, one processor is always used as the master processor, and all remaining processors are used as worker processors. The master processor does not do any computing and is only used to control the communication to and from the worker processors. The same master-worker scheme is used for both the synchronous and asynchronous implementations presented here.

The parallel implementation starts with the master processor sending out a design point to each worker processor. Each worker processor receives the request and generates a local temporary working directory. Within this working directory, the analysis is setup and executed for the current design point. Once the analysis is completed, the objective function and constraint values are sent back to the master processor. The master processor checks if any more analyses should be performed and if so, sends design points to the worker processors. This process is repeated until all analyses are completed. The only difference between the synchronous and asynchronous implementations are the following.

1. The synchronous algorithm performs all design point and swarm updates at the end of each design iteration, while the asynchronous algorithm updates all the design point information as soon as the point is available.
2. The synchronous algorithm waits at the end of each design iteration until all points are analyzed before starting the next design iteration, while the asynchronous algorithm directly starts the next design iteration without waiting for all points to be analyzed.
3. The synchronous algorithm creates a set of points for each design iteration and waits until all points are analyzed for the design iteration to be completed. The asynchronous continuously sends out analyses and the end of a design iteration is identified each time a number of points equal to the swarm size is completed.
4. The synchronous algorithm can use both formulations of the velocity vector, that is using the best point found so far or using the best point from the current design iteration. For the asynchronous algorithm, we always evaluate the velocity vector using only the best point found so far.

D. Parallel Hardware

The parallel hardware used is a crucial aspect for a successful parallel implementation of the PSO algorithm. In the present work, each analysis includes a finite element analysis that heavily depends on disk input/output (IO) through scratch files. In this case, it is easy to generate a severe disk IO bottleneck when sharing a common disk among the parallel processors, similar to that reported in Venter and Watson.¹⁴ To achieve the best parallel efficiency, it is vital to have a parallel setup where it is possible to access local disk storage for each processor or small subgroup of processors.

In the present work, the Virginia Tech System X computer¹⁵ was used as the parallel compute environment. System X is a collection of 1100 Apple Xserver G5 cluster nodes, where each node consists of two 2.3 GHz PowerPC 970FX processors, 4GB ECC DDR400 (PC3200) random access memory (RAM) and a 80 GB S-ATA hard disk drive.

IV. Example Problem

The two parallel PSO algorithms are applied to the multidisciplinary design of a typical long-range transport aircraft wing in the Boeing 767 class, similar to that presented in Venter and Sobieszczanski-Sobieski⁸. For this setup, a single system level analysis consists of a simplified aerodynamic analysis and a structural sub-optimization, using the commercially available GENESIS¹⁶ structural analysis and optimization code. The forces from the aerodynamic analysis are transferred to the structural analysis and optimization. The steps required to perform a single system level analysis is outlined in Fig. 1.

Two independent load cases are considered, a 3.75 G maneuver and a -1.5 G maneuver. The wing is optimized relative to a reference wing with properties summarized in Table 1. Note that the wing area (S), the take off gross weight ($TOGW$), the root chord to tip chord ratio (c_t/c_r) and the sweep angle (p) are all assumed to be constant.

A. System Level Optimization

The aerodynamic optimization is performed at the system level. The goal of the system level optimization is to maximize the aircraft flight range by changing the aspect ratio (A), the depth-to-chord ratio (h/c), the number of internal spars, the number of internal ribs, and the type of wing cover construction. The system level optimization problem is thus an unconstrained problem with five design variables, three of which are discrete variables. The aspect and depth-to-chord ratios are continuous design variables that are allowed to vary between 1.0 and 50.0 and 0.01 and

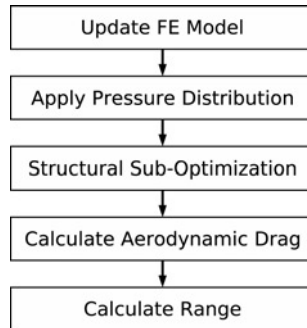


Fig. 1 System level analysis.

Table 1 Reference wing.

Parameter	Value	Parameter	Value
Span (b_{ref})	120 ft	Take off Gross Weight ($TOGW$)	300 000 lbs
Root Chord (c_{ref})	0.5 ft	Aspect Ratio (A_{ref})	6.8571
Drag (D_{ref})	4000 lbs	$(h/c)_{ref}$ Ratio	0.12
Range (R_{ref})	5000 n. mi.	c_t/c_r Ratio	0.4
Area (S)	2100 ft ²	Sweep (p)	25/120

1.0 respectively. The remaining variables are all integer variables. The number of internal spars can be either 0 or 1, the number of internal ribs can vary between 0 and 7 and the wing cover construction can be either sandwiched or hat-stiffened.

The range is calculated using the simplified Breguet formula that does not account for the required fuel reserve, as shown in Eq. (3)

$$R = C_r \frac{L}{D} \ln \left(\frac{TOGW}{W_c + W_p + W_{opt}} \right), \quad (3)$$

where R is the range; L the lift; D the total drag; W_c the non-structural weight; W_p the payload weight; W_{opt} the structural weight and C_r a constant. The take-off gross weight $TOGW$ includes all the weight components $TOGW = W_c + W_p + W_{opt} + W_f$, where W_f is the fuel weight. In the present work, a structural sub-optimization of the reference wing generates W_{opt} and the aircraft flight range of the reference wing (5000 n. mi.) is used to obtain a C_r value of 155.81. The following assumptions are maintained: $W_c = 0.48 TOGW$ and $W_p = 0.13 TOGW$. Additionally, the lift (L) must equal the $TOGW$, which is considered constant in the present work. With these assumptions, the range formula is simplified as shown in Eq. (4)

$$R = C_r \frac{TOGW}{D} \ln \left(\frac{TOGW}{0.61 TOGW + W_{opt}} \right). \quad (4)$$

Equation (4) represents the system level objective function for the aerodynamic optimization and consists of two designable components: D and W_{opt} . W_{opt} is the weight from the structural sub-optimization problem factored by 1.3 to account for structural non-optimum weight, and D is the total drag obtained from the aerodynamic analysis. The structural sub-optimization problem and the aerodynamic analysis are discussed in more detail in the following sections.

The finite element model of the wing-box for the reference wing with three spars and nine ribs is shown in Fig. 2. For the present design problem, the front and back spars and the ribs at the root and tip of the wing are always present. However, the internal spar is optional (the number of internal spars is allowed to vary between 0 and 1) and the number of internal ribs is allowed to vary between 0 and 7. The rib spacing (S_{rib}) is a constant and equal to $S_{rib} = b_{ref}/(N_{rib_max} + 1)$, where $N_{rib_max} = 7$ is the maximum number of allowable internal ribs. The internal ribs are always placed from the root out, observing the rib spacing S_{rib} , with the number of internal ribs obtained from the corresponding discrete system level design variable. Additionally, the construction of the upper and lower wing covers are selected from either a sandwich or a hat-stiffened construction.




Fig. 2 Wing-Box structural finite element model.

B. Structural Sub-Optimization

For the structural sub-optimization problem, a simplified finite element model is used to model the wing-box. The wing-box model consists only of shell elements and, for simplicity, the spar and rib caps are not modeled. It is assumed that the wing-box is manufactured from aluminum with a density of 0.1 lb/in^3 , an allowable tensile strength of 50 ksi and an allowable compression strength of 25 ksi . The reference finite element mesh is shown in Fig. 2 and consists of 72 nodes and 72 shell elements. During the system level optimization, the number of finite elements used to model the wing-box varies between 50 and 72 depending on the number of internal spars and ribs for the current configuration. The nodes shown on the leading and trailing edges of Fig. 2 are non-structural and are included only to transfer aerodynamic loads from the aerodynamic analysis to the structural analysis. The load transfer is done using rigid elements that do not add stiffness to the finite element model.

The goal of the structural sub-optimization problem is to minimize the weight of the wing-box by changing the thickness values of the shell elements, subject to allowable stress and local buckling constraints. The number of design variables in the structural sub-optimization problem depends on the number of internal spars and ribs and the upper and lower wing cover construction selected at the system level. The spars and ribs are modeled as aluminum panels with a single thickness design variable per panel. All three spars between the same two ribs share the same thickness design variable. Similarly, a single design variable is used to design the thickness of each rib. There is thus a maximum of eight design variables for designing the spar thickness and eight design variables for designing the rib thickness.

The top and bottom wing cover panels are designed on a per panel basis, where each panel is constrained by two spars and two ribs. There is a maximum of 16 panels for the top wing cover and 16 panels for the bottom wing cover of the wing-box. A typical panel with positive force directions, orientation and dimensions is shown in Fig. 3. In the model, each cover panel is represented by a single quadrilateral element if the internal spar and all seven internal ribs are present. If the number of internal spars and/or ribs is reduced, some cover panels are represented by more than one quadrilateral element.

The number of design variables used to design each panel depends on the selected construction. For the sandwich construction, there are two design variables for each panel as shown in Fig. 4. The design variables are the thickness of the wing cover t and the thickness of the core t_c . There is thus a maximum of 80 structural sub-optimization design variables when considering the sandwich construction, which corresponds to 16 spar and rib thickness design variables and 64 top and bottom wing cover design variables. The core density is assumed to be $1.736 \cdot 10^{-3} \text{ lb/in}^3$.

For the hat-stiffened construction, there are six design variables for each panel as illustrated in Fig. 5. The design variables are the number of hat stiffeners per panel n , the thickness of the wing cover t_1 , the thickness of the stiffener t_2 , and the dimensions of the stiffener d_2 , d_3 and d_4 . The number of stiffeners per panel is currently assumed to be a continuous variable that can be rounded to the next highest number in the final design. There is thus a maximum of




Fig. 3 Typical wing cover panel.




Fig. 4 Sandwich construction.

208 structural sub-optimization design variables when considering the hat-stiffened construction, which corresponds to the 16 spar and rib thickness design variables and 192 top and bottom wing cover design variables.

The structural sub-optimization problem is subject to both stress and local buckling constraints. Each element is subject to maximum Von Mises stress constraints at both the upper and lower surfaces. The local buckling constraints are applied to each upper and lower wing cover panel, and the equations are summarized in Venter and Sobieszczanski-Sobieski.⁸

The aerodynamic loads are applied as nodal forces to the bottom surface nodes only, including the non-structural leading and trailing edge nodes. The normalized span-wise and chord-wise distribution of the assumed forces are shown in Figs. 6 and 7.

The aerodynamic pressure distribution is converted to concentrated nodal forces, accounting for the areas associated with each node, using Eq. (5)

$$F_{i,j} = C_f LP_j CP_i S_{i,j}, \tag{5}$$

where i is the line number and j the chord number. Additionally, C_f is the wing loading, which is 143 lb/sf for the current wing, LP_j is the span-wise pressure distribution value based on the line number from Fig. 7, and CP_i is the chord-wise pressure distribution based on the chord number from Fig. 6. The area associated with each node is represented by $S_{i,j}$. The actual nodal forces are obtained by multiplying the forces obtained from Eq. (5) with a factor equal to the G value of the maneuver. For the $3.75 G$ maneuver, this factor is 3.75, and for the $-1.5 G$ maneuver, the factor is -1.5 . An important note regarding Eq. (5) should be made here. The aerodynamic nodal forces are influenced by the deformation of the wing. In the present work, this interaction is not accounted for. Instead, it is assumed that the wing will be built to a jig-shape that off-sets the deformation due to the aerodynamic loads.

C. Aerodynamic Analysis

The aerodynamic analysis is used to calculate the total drag required to calculate the range at the system level, using Eq. (4). A simplified drag calculation is used with the total drag (D) for the current wing calculated based on




Fig. 5 Hat-stiffened construction.




Fig. 6 Normalized span-wise pressure distribution.




Fig. 7 Normalized chord-wise pressure distribution.

the total drag of the reference wing (D_{ref}). The calculations consists of the induced drag (D_I), the wave drag (D_w), and a constant fraction of the original drag, as shown in Eq. (6)

$$D = D_I + D_w + C_{dt} D_{ref}, \quad (6)$$

where C_{dt} is a constant and is assumed to be 0.4 in the present work. The total drag of the reference wing is $D_{ref} = 40\,000\text{ lb}$. The induced drag depends on the aspect ratio of the current wing (A) relative to the aspect ratio of the reference wing (A_{ref}), as follows

$$D_I = C_{di} D_{ref} \frac{A_{ref}}{A}. \quad (7)$$

C_{di} is a constant and is assumed equal to 0.4 in the present work. The wave drag depends on the frontal area of the wing as projected on the streamline direction. The frontal area depends on the span, and the mean height of the wing as follows

$$S_h = b \frac{(h_r + h_t)}{2}, \quad (8)$$

where S_h is the frontal area, b is the span and h_r and h_t are the height at the root and the tip of the wing, respectively. The wave drag is then obtained relative to the reference wing using Eq. (9)

$$D_w = C_{dw} D_{ref} \frac{S_h}{S_{href}}, \quad (9)$$

where C_{dw} is a constant assumed equal to 0.2 in the present work, S_h is the frontal area of the current wing, and S_{href} is the frontal area of the reference wing.

V. Numerical Results

Both the synchronous and asynchronous PSO implementations were applied to the example problem. The goal was to study two effects: the numerical accuracy of the asynchronous implementation; and the parallel efficiency of both implementations. For each point shown in the following graphs, 10 independent repetitions were executed, and the results averaged over these 10 repetitions are shown. For the numerical accuracy results, the objective functions were averaged for each design iteration. For the parallel efficiency, the total elapsed time was averaged for each number of processors used.

The parameters used for the PSO algorithm are summarized in Table 2. The swarm size of 31 particles was chosen based on the maximum number of processors used in this study. A maximum of 32 processors was used, one as the master processor that does the communication and the 31 remaining processors as worker processors that do all the work. The values of the inertia weight and the two trust parameters were selected based on past results as published in Venter and Sobieszczanski-Sobieski.⁸ No attempt was made to fine tune these values for increased performance.

Table 2 PSO parameters.

Parameter	Value	Parameter	Value
Number of particles	31	Trust parameter 1, c_1	1.50
Initial inertia weight, w	1.4	Trust parameter 2, c_2	2.50

A. Baseline Accuracy Study

The synchronous parallel implementation uses exactly the same logic as the serial implementation and as a result gives the same numerical accuracy, independent of the number of processors used. However, the asynchronous algorithm can lead to different numerical accuracy as a result of starting analyses in the next iteration before the current iteration is completed. This effect could be amplified as more processors are used. As a result, the accuracy of the asynchronous algorithm needs to be validated.

To perform this study, we first solved the optimization problem using the synchronous implementation with 16 processors. Four different cases were considered: using both formulations of the velocity vector (using the best point in the current iteration versus using the best point found to date) and applying or not applying the craziness operator. The results are presented in Fig. 8. From Fig. 8, the best synchronous case seems to be using the best point found to date and applying craziness.

For the asynchronous algorithm, we always evaluate the velocity vector using only the best point found so far. As a result, only two different cases were considered: applying or not applying the craziness operator. The results are summarized in Fig. 9. From Fig. 9, it appears that the best performing asynchronous algorithm is when we make use of the craziness operator.

In addition to the average values, the variance between the 10 runs are also very important in comparing the algorithms. To compare the variance among the runs, the best synchronous and best asynchronous cases were selected and the average, minimum and maximum values obtained at each design iteration are shown in Figs. 10 and 11 below.

From Figs. 8 and 9 and Figs. 10 and 11, it is clear that both the numerical accuracy and the variance in the results obtained from the two algorithms are comparable for the example problem considered here. However, the influence of the number of processors on the numerical accuracy of the asynchronous algorithm should also be studied since the asynchronous algorithm updates the particle positions in a different fashion from that of the synchronous algorithm. The update scheme used for the asynchronous algorithm depends on the number of processors used. To check the numerical accuracy of the asynchronous algorithm as a function of the number of processors, the results for 4 (bottom curve), 8, 16 and 32 processors are compared, as shown in Fig. 12. Figure 12 also includes the results for the best performing synchronous case from Fig. 8.

Except for the 4 processor case, Fig. 12 clearly shows that the numerical accuracy is not significantly influenced when varying the number of processors for the asynchronous algorithm. Further investigation of the 4 processor



Fig. 8 Synchronous baseline study (Results averaged over 10 runs).




Fig. 9 Asynchronous baseline study (Results averaged over 10 runs).




Fig. 10 Variance in synchronous results (Results from 10 runs).

case, indicated that of the 10 optimizations performed, 9 had an optimum range of more than 5385 *n.mi.*, while one optimization resulted in an optimum range of only 4693 *n.mi.* The poor accuracy of the 4 processor case is thus due to a single optimization that did not converge on the optimum solution. If this one bad point is ignored, the 4 processor case performed as well as the others. From Figs. 8, 9 and 12, we can thus conclude that for the current




Fig. 11 Variance in asynchronous results (Results from 10 runs).

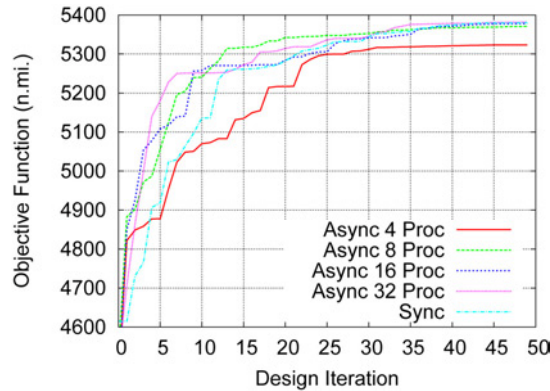


Fig. 12 Asynchronous accuracy study.




Fig. 13 Parallel speedup comparison (Averaged over 10 runs).

example problem, the numerical accuracy of the PSO algorithm is not compromised when using the asynchronous implementation, even if the number of processors used is equal to the number of particles in the swarm.

B. Parallel Efficiency Study

For the synchronous implementation, the parallel efficiency is highly problem dependent. For example, the analysis time for the current problem depends on the point being analyzed. The overall analysis time is fairly small, but there is at least an order of magnitude difference between the shortest and longest analysis times. The current example problem is thus a good test for comparing the algorithms both in terms of accuracy and efficiency. For the asynchronous algorithm, the variation in analysis times is a severe test since more design points will be analyzed in the next design iteration before the current design iteration is completed. This scenario is representative of a heterogeneous parallel environment with a big variation in processor speeds. For the synchronous algorithm, the variation in analysis times result in having idle processors at the end of each design iteration which result in poor parallel efficiency.

To test the parallel efficiency and speedup of the two algorithms, the best performing synchronous and asynchronous algorithms were executed with 1, 4, 8, 16 and 32 processors. For each case, 10 repetitions were completed and the elapsed times averaged. The resulting parallel speedup and efficiency plots are shown in Figs. 13 and 14, respectively. In each case, the ideal speedup and efficiency are also shown.

From Figs. 13 and 14, it is clear that the asynchronous algorithm significantly outperforms the synchronous algorithm for the example problem considered here. As expected, the synchronous algorithm becomes less efficient when more processors are added. The asynchronous algorithm provides close to the ideal speedup and actually becomes more efficient as more processors are used. This increase in efficiency is explained by the master-worker

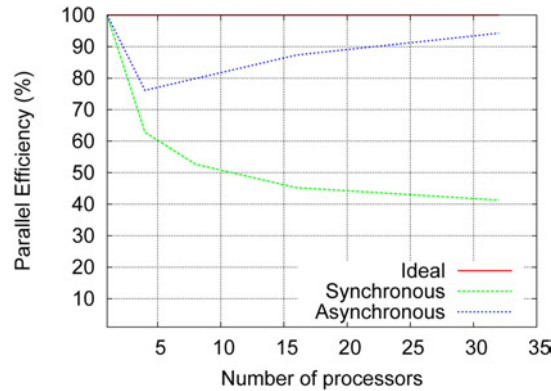


Fig. 14 Parallel efficiency comparison (Averaged over 10 runs).

implementation used. The master processor only communicates with the worker processors that do all the calculations. For small numbers of processors, the master processor is not fully utilized, and a larger fraction of the total number of processors are not working at full capacity. As the number of processors is increased, a larger fraction of the total number of processors is used to do calculations, and thus the increase in efficiency. Although the master processor is under utilized for smaller numbers of processors, it becomes very important to avoid a communications bottleneck for larger numbers of processors. For very large numbers of processors, one may even want to have multiple master processors to avoid possible communication bottlenecks between the master and worker processors.

For the synchronous implementation, there is significantly less communication within a design iteration. For example, in the best case scenario where the number of processors is equal to the swarm size, all communication is performed only at the beginning and the end of each design iteration. It could thus be argued that there is less need to use the master processor for communication purposes only. Instead the master process could also be used to perform analyses. The parallel efficiency plot can be adjusted to account for this, by assuming that the master processor can work as fast as the other processors without any loss in efficiency due to communication. Of course this is a theoretical and best case scenario only. The updated parallel efficiency plot is shown in Fig. 15 below. Note that with this adjustment the synchronous algorithm has a higher efficiency for a smaller number of processors. For larger numbers of processors there is little difference as compared to the results obtained in Fig. 14.

C. Gradient-Based Optimization

Finally, it is useful to compare the PSO results with that of a gradient-based optimization of the example problem. As illustrated in Venter and Sobieszczanski-Sobieski,⁸ the example problem considered here has significant numerical




Fig. 15 Ideal synchronous parallel efficiency comparison (Averaged over 10 runs).

Table 3 DOT vs. PSO comparison results.

Parameter	Reference wing	DOT optimum	PSO optimum
Aspect ratio	6.8571	6.9100	9.2360
Depth-to-chord ratio	0.01	0.08	0.11
Num. Internal spars	1	0	0
Num. Internal ribs	7	4	4
Construction	–	Sandwich	Sandwich
Range (n. mi.)	5000	5225	5375
Number of analyses	–	526	1500

noise associated with it. It is expected that the combination of the numerical noise and the discrete design variables would have a negative impact on the performance of a gradient-based optimizer.

For the comparison, it was decided to look at all 32 possible combinations of the 3 discrete design variables. For each combination, it is possible to solve a 2 design variable continuous problem, using a gradient-based optimizer. In our case, we use the commercially available DOT¹⁷ optimizer with the reference wing as a starting point. The reference wing, the DOT optimum results and PSO optimum results are summarized in Table 3. Although DOT found the same optimum values for the discrete variables and improved the range, it is clear that the aspect ratio could not move much from the starting point. Although DOT required only a third of the analyses to complete the optimization, one should also consider the potential for speedup in a parallel environment. In general, a gradient-based optimizer can only use as many processors as design variables to perform the finite difference gradient calculations in parallel. In contrast, the PSO algorithm can use many more processors than design variables, as illustrated in the current example problem.

VI. Concluding Remarks

An asynchronous parallel PSO algorithm is presented and is compared to a synchronous parallel PSO algorithm. The multidisciplinary design of an transport aircraft wing, where it is shown that gradient-based optimization is not a viable alternative, is used to compare the algorithms. The numerical results presented indicate that the asynchronous algorithm significantly outperforms the synchronous algorithm in terms of parallel efficiency for the example problem considered here. In addition, the numerical accuracy of the asynchronous algorithm is comparable to that of the synchronous algorithm for the current example problem.

The present work indicates that asynchronous PSO algorithm is a good candidate for efficient parallel computations in cases where the analysis time depends on the design point being analyzed, the optimization is performed on a heterogenous system of parallel computers or the swarm size that is not an integer multiple of the number of processors. Apart from the algorithm, the parallel hardware is also important and should be appropriate for the problem being solved. For example in the problem presented here, it was crucial to have local disk storage available at the processor level to avoid a disk input/output (IO) bottleneck.

The implementation presented here should scale very well to larger numbers of processors and will be the subject of future work.

VII. Acknowledgment

The authors would like to thank Prof. Cal Ribbens for providing access to the Virginia Tech System X computer. The authors would also like to acknowledge the anonymous referees appointed by the AIAA for a constructive comments.

References

- ¹Michalewicz, Z., and Dasgupta, D., *Evolutionary Algorithms in Engineering Applications*, Springer Verlag, 1997.
- ²Nemhauser, G. L., and Wolsey, L. A., *Integer and Combinatorial Optimization, Chapter 3*, John Wiley & Sons, 1988.
- ³Kennedy, J., and Eberhart, R. C., Particle Swarm Optimization, *Proceedings of the 1995 IEEE International Conference on Neural Networks (held in Perth, Australia)*, pp. 1942–1948, 1995.
- ⁴Eberhart, R. C., and Kennedy, J., A New Optimizer Using Particles Swarm Theory, *Sixth International Symposium on Micro Machine and Human Science (held in Nagoya, Japan)*, pp. 39–43, 1995.

- ⁵Fourie, P. C., and Groenwold, A. A., The Particle Swarm Optimization Algorithm in Size and Shape Optimization, *Structural and Multidisciplinary Optimization*, 23, pp. 259–267, 2002.
- ⁶Fourie, P. C., and Groenwold, A. A., Particle Swarms in Topology Optimization, *In Proceedings of the Fourth World Congress of Structural and Multidisciplinary Optimization*, Dalian, China, 2001.
- ⁷Venter, G., and Sobieszcanski-Sobieski, J., Particle Swarm Optimization, *AIAA Journal*, 41(8), pp. 1583–1589, 2003.
- ⁸Venter, G., and Sobieszcanski-Sobieski, J., Multidisciplinary Optimization of a Transport Aircraft Wing using Particle Swarm Optimization, *Structural and Multidisciplinary Optimization*, 26(1-2), pp. 121–131, 2004.
- ⁹Schutte, J. F., Fregly, B. J., Haftka, R. T., and George, A., A Parallel Particle Swarm Algorithm. *In Proceedings of the Fifth World Congress of Structural and Multidisciplinary Optimization*, Venice, Italy, 2003.
- ¹⁰Schutte, J. F., Reinbolt, J. A., Fregly, B. J., Haftka, R. T., and George, A. D., Parallel Global Optimization with the Particle Swarm Algorithm, *International Journal of Numerical Methods in Engineering (accepted)*, 2003.
- ¹¹Koh, B., George, A. D., Haftka, R. T., and Fregly, B. J., Parallel Asynchronous Particle Swarm Optimization, *International Journal for Numerical Methods in Engineering (in press)*, 2005.
- ¹²Kennedy, J., and Spears, W. M., Matching Algorithms to Problems: An Experimental Test of the Particle Swarm and Some Genetic Algorithms on the Multimodal Problem Generator, *Proceedings of the 1998 IEEE International Conference on Evolutionary Computation (held in Anchorage, AK)*, 1998.
- ¹³Shi, Y. H., and Eberhart, R. C., A Modified Particle Swarm Optimizer, *Proceedings of the 1998 IEEE International Conference on Evolutionary Computation (held in Anchorage, AK)*, 1998.
- ¹⁴Venter, G., and Watson, B., Efficient Optimization Algorithms for Parallel Applications, *Proceedings of the 8th AIAA/USAF/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization*, AIAA-2000-4819, Long Beach, CA, 2000.
- ¹⁵<http://www.tcf.vt.edu/systemX.html>
- ¹⁶*GENESIS Version 7.0 Users Manual*, Vanderplaats Research and Development, Inc., 1767 S. 8th St., Colorado Springs, CO, 2001.
- ¹⁷*DOT. Design Optimization Tools, Version 5.x, Users Manual*, Vanderplaats Research and Development, Inc., 1767 S. 8th St., Colorado Springs, CO, 2001.