Toward Evolved Flight

Rusty Hunt
Intelligent Systems Division
NASA Ames Research Center
Moffett Field, CA
650-604-1717
Rusty.Hunt@nasa.gov

Gregory S. Hornby
QSS Group, Inc.
Intelligent Systems Division
NASA Ames Research Center,
Moffett Field, CA
650-604-3373
hornby@email.arc.nasa.gov

Jason D. Lohn
Intelligent Systems Division
NASA Ames Research Center
Moffett Field, CA 95138
650-604-5138
Jason.D.Lohn@nasa.gov

ABSTRACT

We present the first hardware-in-the-loop evolutionary optimization on an ornithopter. Our experiments demonstrate the feasibility of evolving flight through genetic algorithms and adaptable hardware, without the requirement for a thorough knowledge of the aerodynamics of flapping flight. In this research we successfully optimized forward velocity and basic efficiency on an actual hardware ornithopter. The ornithopter was flown integrated to a "whirling-arm" test apparatus, allowing lengthy experimental flights without the risk of crashing. Flapping rate and tail position were controlled by an evolutionary algorithm with feedback of forward velocity and motor power. The system evolved an unexpected optimal configuration. This paper discusses the development of the test apparatus and experimental results from the initial phase of research.

Categories and Subject Descriptors

I.2.9 [**Artificial Intelligence**]: Robotics – *autonomous vehicles, commercial robots and applications, propelling mechanisms.*
I.2.8 [**Artificial Intelligence**]: Problem Solving, Control Methods, and Search – *control theory, dynamic programming.*
J.2 [**Physical Sciences and Engineering**]: Language Constructs and Features – *aerospace, engineering.*

General Terms

Algorithms, Performance, Design, Experimentation, Theory.

Keywords

Ornithopter, flapping flight, evolutionary algorithm, evolvable hardware, hardware-in-the-loop evolution.

1. INTRODUCTION

Scientists, philosophers and hobbyists have studied the flight of birds for hundreds of years [1], and attempted to duplicate their seemingly effortless flight, with frequently unsatisfactory results. See Figure 1. Most research has been directed toward an analytical understanding of the aerodynamics of flapping flight, with that knowledge then used to build flying machines [2] that replicate the efficiency and maneuverability of birds.

Until recently the application of flapping flight was something of a novelty. The difficulties of flapping include numerous moving parts, which increase complexity, weight, and break-downs, and the motion of the flight platform, which is generally undesirable, especially for passengers or reconnaissance. Also, with the continuing improvements in conventional propulsion technology, ornithopters have been at a disadvantage in efficiency.

Figure 1. Frost's Ornithopter, ca 1902.

Recent interest in micro air vehicles (MAV's) has provided impetus for an improved understanding of flapping flight [3]. At centimeter scale and smaller, Reynolds Number effects cause a significant degradation in the efficiency of propellers. Flapping wing flight offers advantages both in efficiency and maneuverability at this small scale.

Evolutionary algorithms (EA's) have been used successfully to solve a variety of physical and engineering problems which are difficult or impossible to solve through analytical means. EA's have been used for evolution of hardware such as satellite antennae [4] and micro air vehicles [5]. Complex motions have been successfully evolved, such as the quadrupedal gait of the Sony "Aibo" entertainment robot [6], and the flapping motion of simple rigid wings [7].

Present capabilities in Computational Fluid Dynamics (CFD) are insufficiently accurate and too computationally intensive to evolve ornithopter flight in simulation – processing time for one simulation time step could be several hours, depending on the fidelity of the simulation. Consequently, to successfully evolve ornithopter flight real hardware must be used.

GECCO'05, June 25-29, 2005, Washington, DC, USA.
ACM 1-59593-010-8/05/0006.

We have begun experiments toward evolving flight through evolutionary algorithms and adaptable hardware, without a requirement for a detailed knowledge of the aerodynamics of flapping flight. The evolution is performed on an actual hardware ornithopter. Using this system in two separate experiments, we have successfully optimized for speed and power efficiency. This paper discusses the development of the test apparatus and these experimental results along with a discussion of what we learned about the aerodynamics of our ornithopter.

2. APPROACH

The evolution of flight is a complex problem, involving many aspects of aerodynamics, structures, controls, and power. We approach this research in a series of logical steps.

2.1 Test Bed Development

The specialized hardware and software needed for the research are developed, debugged, and tested. This includes the ornithopter itself, its supporting equipment, and the control and data acquisition systems.

2.2 Tethered Flight Optimization with Static Test Vehicle

The ornithopter is operated flying suspended from a rotating support structure. The morphology of the vehicle is fixed, including wing shape and stiffness, body size, and weight distribution. The EA controls the attitude of the tail and the flapping rate, with the simple goals of achieving maximum forward velocity and efficiency, without regard to lift.

The preceding phases are complete, and are the subject of this paper. The following phases are planned for the coming months.

2.3 Evolution of Successful Flying Parameters

The ornithopter is operated flying on its support structure. The EA controls the attitude of the tail and the flapping speed, with the constraint that the vehicle velocity and lift are such that it would be flying if it were free of its support arm. The acceptable solution combinations for this configuration are expected to be small.

2.4 Propulsion Optimization with Evolvable Test Vehicle

The ornithopter is operated flying on its support structure. Modifications to the ornithopter allow variations in drive train characteristics as well as wing shape and flexibility. These parameters will be controlled by the EA with the goal of achieving the maximum efficiency from the propulsion system.

2.5 Tightly Integrated System Optimization

A long-range goal of this research program is co-evolution of several sub-systems of the ornithopter to produce a machine optimized for maximum overall performance and efficiency. Candidates include simultaneous optimization of drive train geometry and motor drive waveform for highest efficiency, and evolved coordinated control of wing trajectory and tail position for optimal maneuverability.

2.6 Free Flight

An on-board algorithm and hardware will be developed to optimize flying characteristics during free flight. Initially the ornithopter's flight path will be controlled by human radio control with the algorithm optimizing stability and propulsion efficiency. Further goals include evolution of autonomous flight.

3. EXPERIMENTAL SETUP

One of the chief difficulties in evolving flight is preventing damage to the vehicle when non-flying individuals are evaluated. For these experiments, the bird was suspended from a counter-balanced rotating arm, allowing forward travel as well as variations in altitude. This experimental technique follows rich historical precedent, first used by English mathematician Benjamin Robins in 1746 (see Figure 2), followed later by notable aerodynamicists Sir George Caley, Otto Lillianthal, and Samuel Langley [8].

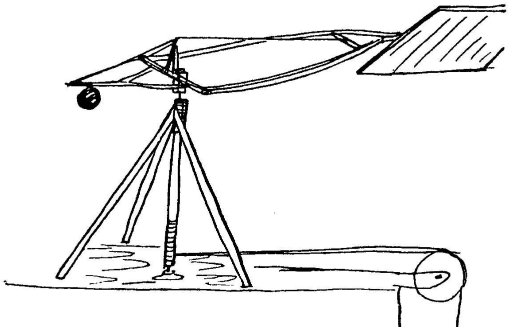

Figure 2. Robins' Whirling-Arm Apparatus.

For our experiments, the "whirling arm" setup was chosen for its low cost, portability, and multiple degrees of freedom. See Figure 3. Presently, the arm is rotated by the forward motion of the ornithopter pulling the arm along with it. Forward speed and power are measured. For the upcoming research phases, a drive motor and control system will be added, to remove the aerodynamic drag of the arm from the ornithopter characteristics. Lift measurement will be accomplished by addition of a tilt sensor on the arm.

Figure 3. Our Whirling-Arm Apparatus.

Control and data acquisition of the ornithopter is by two dedicated microprocessors. The evolutionary algorithm runs in Matlab on a desktop computer, which communicates via serial interface with the control and data acquisition processors.

Individual components are described in more detail in the following sections.

3.1 Test Vehicle

A commercially available radio control ornithopter was selected to be the test vehicle for our research. This selection gave us an easily modifiable platform at a low cost. The ornithopter is shown in Figure 4.

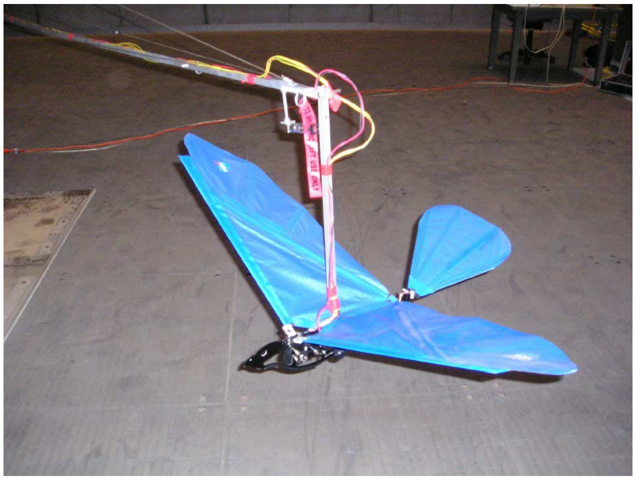

Figure 4. Ornithopter Test Vehicle.

It has a wing-span of about 4 feet, and weighs a little less than a pound. Aerodynamic characteristics are summarized in Table 1. The size and weight of the ornithopter are similar to a seagull or a large owl.

Table 1. General Aerodynamic Characteristics of Test Vehicle

Weight	4.31 N (15.5 ounces)
Span	1.17 m (46 inches)
Chord (avg)	21 cm (8.3 inches)
Length	66 cm (26 inches)
Height	7.6 cm (3.0 inches)
Airspeed (min) (*)	6.7 m/s (22.0 ft/sec)
Airspeed (max)	8.9 m/s (29.2 ft/sec)
Wing Area	0.246 m^2 (2.7 ft^2)
Wing Loading	17.6 N/ m^2 (0.37 lb/ft^2)
Aspect Ratio	5.57
Reynolds Number	1.122 x10^5
Coefficient of Lift – Cl (**)	.469
Coefficient of Drag – Cd (**)	.0385
Lift to Drag Ratio	12.2
Flapping Rate (typ)	186 beats per minute
Strouhal Number	0.192

(*) for sustained flight. (**) calculated

In 1992, Tennekes developed a plot of the weight, wing loading, and cruising velocity for flying creatures ranging in size from fruit fly to pteranodon [9], and showed that all flying creatures generally fall along a logarithmic line through those parameters. It is interesting to note that for this ornithopter the values are consistent with Tennekes's plot, falling right between the Goshawk and the Kittwake.

Also interesting is the Strouhal number, which is a dimensionless ratio of flapping rate times amplitude divided by forward velocity. Taylor and others have shown that for nearly all flying (and swimming) creatures the Strouhal number is always in the range of 0.2 to 0.4 [10]. Our bird is no exception, flying at the lower limit of this range.

The body of the vehicle is a flat, open-frame fiberglass composite structure, shown in Figure 5. It supports the wings, tail, drive train, and battery. The open frame design was especially good for this type of research. It is easily modified, and parts are visible during operation.

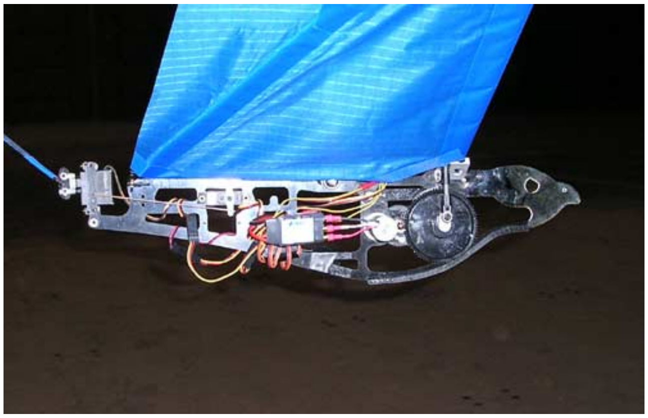

Figure 5. Ornithopter Body, Drive System, and Tail Servos.

The wings are made from a woven nylon membrane with carbon fiber spars along the leading edge, and diagonally transverse spars from the leading edge to a point on the centerline near the tail. The leading edge spars connect to the body by hinges at the centerline of the body. The arrangement of the spars in the wings creates a taught membrane over most of the wing surface, with a small amount of flexing toward the wing-tips. This flexing provides the forward thrust when the wings are flapping.

The model is powered by an electric motor driving a double-reduction gear train, which in turn drives crankshafts linked to the wing leading edge spars. This drive system imparts a simple sinusoidal motion to the wings. The flapping frequency is determined by varying the speed of the motor, accomplished by a commercially available Electronic Speed Control (ESC). This device is a standard hobby radio control device, which accepts the control signal from a receiver and adjusts the power to the motor proportionally.

Attitude control is provided by a single flat tail surface, which approximates the effect and appearance of a bird's tail. The position of the tail in relation to the rest of the bird is controlled by two servos. The tail is mounted directly to the output shaft of the rudder servo, which rotates the tail left and right. The rudder servo is attached to the ornithopter body by a hinge. The elevator servo output shaft is linked to the rudder servo body, and changes the inclination of the tail in relation to the body.

It is important to note that there is a good bit of interaction between elevator and rudder effects on this model. When the tail is in its flattest position relative to the body (down elevator), rotation of the tail has minimal effect on yaw. As the tail inclination angle is increased (up elevator), tail rotation (left or right rudder) has an increasing effect on yaw. Also, as the tail is rotated away from center, there is less effective surface area and less elevator effect on pitch. This coupling of pitch and yaw makes automatic flight control using conventional control difficult for this type of aircraft, and an interesting application for evolutionary algorithms.

For this series of experiments, only minor modifications were made to the ornithopter. The battery was removed and power supplied to the model from a DC power supply on the test cart, described below. The radio was also removed, and the pulse-width-modulation (PWM) control signal to the servos was generated by a microprocessor on the test cart. A small camera was attached to the bottom of the body, for future image processing experiments.

3.2 Support Arm and Equipment Cart

The arm is balanced so that the weight of the support arm is neutral, and the bird only feels its own weight. This arrangement allows for simple measurement of forward velocity as well as a (future) measurement of lift. The test rig is shown in Figure 6.

The arm is a truss structure constructed of aluminum channel for low mass and steel cables for strength. A counter-weight is mounted on the end opposite the bird, with adjustments available to accommodate weight changes due to modifications in the bird hardware. The arm can be used in two configurations: full-length and half-length. At full length, there are two main sections, giving about a 6 meter radius for the flight path. One section may be used alone for a half-length configuration with a 3 meter radius of travel. The half-length setup is used to allow testing and debugging to be done in a smaller space.

The arm is attached to the top of the rotating center structure on a pivot, so that the bird is free to move up and down as its lift changes. In future experiments, the angle of the arm will be measured and fed back into the EA as an indication of lift. A small color video camera is mounted on the arm to allow continuous monitoring and recording of the ornithopter.

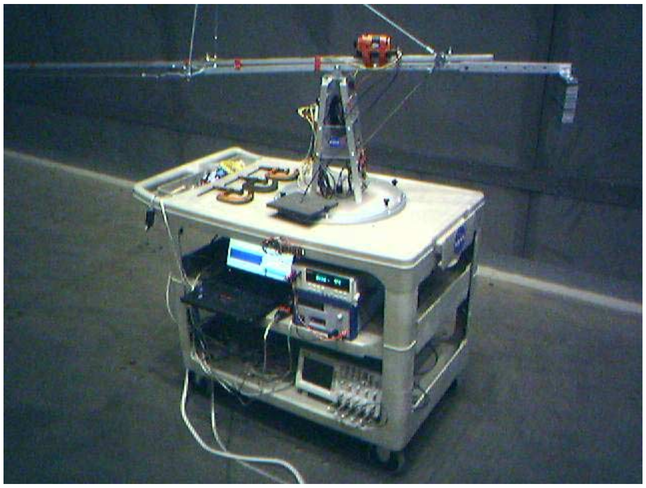

Figure 6. Support Arm on Equipment Cart.

The center structure rotates on a slewing bearing (see Figure 7), which provides low rotational friction as well as allowing for some torsional forces, such as when the bird is moving forward but not generating enough lift to fly. Two processors are mounted on the center structure to handle control, data acquisition, and serial communications.

The slewing bearing is mounted on a four wheel utility cart, shown in Figure 6. This allows portability for the entire test rig, and also provides operating locations for several pieces of support equipment.

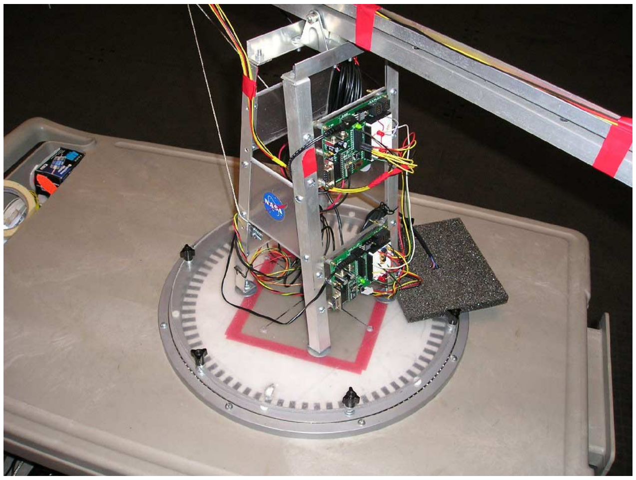

Figure 7. Center Pivot Structure on Slewing Bearing.

Power and data are passed from the rotating structure to the cart through a capsule type slip ring assembly. The slip ring has three 10 amp circuits which are used for main motor power and an isolated ground. In addition there are six 2 amp circuits which carry the serial communications from the PIC processors, video and audio from the camera, and filtered power for the electronics.

3.3 Control and Data Acquisition

The ornithopter is controlled by a Parallax BS2 processor on a Parallax development board. This processor was chosen for ease of programming and a versatile array of interface options.

The control processor provides a pulse width modulated signal to each of the servos and to the Electronic Speed Control. The PWM signal is a pulse width of 1.0 msec to 2.0 msec, representing 0-100% of travel or span. The pulses are repeated at 50 Hz. The control algorithm receives a 5 byte string via serial interface from the computer running the EA. Each byte is the commanded value for one servo.

Data acquisition is performed by a separate processor so that its operation does not interfere with the timing of the control processor. In our initial experiments the only parameter measured by this processor is the speed of rotation of the arm, which is sensed by a quadrature detector. The processor measures the period of each light-to-dark transition, internally averages 5 data points, and passes the result via serial interface to the computer for use by the EA to calculate the bird's forward velocity. A digital multimeter with a serial interface provides bird motor power measurement to the computer running the EA.

The processors are set up to provide a debugging mode. In this mode, the output to the Electronic Speed Control is disabled (so the wings won't flap). The control processor calculates a simulated bird velocity based on a simple model of the ornithopter and passes the result as a byte to the data acquisition processor. This allows for full hardware-in-the-loop debugging of the test stand, computer, algorithms, and interconnections.

4. EVOLUTIONARY ALGORITHMS

Once the hardware setup was debugged using a simple hill-climbing algorithm, optimization of flight was performed using a steady state evolutionary algorithm. An initial population of twenty individuals was generated at random and each individual was evaluated against the fitness function. Once the initial population was created, the EA created new individuals by either mutation or recombination, randomly chosen with equal probability for each evaluation. The mutation and recombination operators are:

```
function mutation(parent)
delta =
   random_integer_number(range: -10 to +10);
child_parameter[i] =
   parent_parameter[i]+delta;
```

```
function recombination(parent1, parent2)
alpha =
   random_real_number(range: -1.5 to 1.5);
child_parameter[i] =
   alpha*(parent2_parameter[i]-parent1_parameter[i])
   + parent1_parameter[i];
```

As a way of addressing the noisiness of evaluation on actual hardware the algorithm also kept track of the age of individuals in the population - the age of an individual being the number of times it has reproduced - with individuals over an age of four being replaced by new successful individuals. This EA is essentially the same as what we used for hardware-in-the-loop evolution of dynamic gaits for Sony's quadruped robots [6].

The command parameters for all runs were motor speed, tail inclination, tail rotation, and a null control variable. For all four parameters the output values are integers from 0 to 255. The genotype is represented as:

```
TrialVector = [throttle elevator rudder null]
```

The null value was passed through the entire system up to the ornithopter, identically to the other command parameters, with the only exception being that there was no actuator connected.

Initial evaluation runs were performed to optimize the ornithopter's speed without regard to lift or power. The fitness function was the ornithopter's forward velocity:

$$Fitness = Forward\ \ Velocity\ \ (feet\ \ per\ \ second\ )$$

For later runs, the fitness function was modified to represent efficiency, defined as:

$$Fitness = \frac{Forward\ \ Velocity\ \ (feet\ \ per\ \ second\ )}{Motor\ \ Power\ \ (watts\ )}$$

Note that this provides only a comparative measure of efficiency (with units of feet per second per watt) rather than a conventional dimensionless measure of absolute efficiency. Since the fitness value is only used to compare the performance of one individual to another, the comparative efficiency is adequate.

5. EXPERIMENTAL RUNS

The first runs of the system were carried out in a courtyard outside the development laboratory. This had the advantage of quick access to the lab for repairs and modifications, but the gusty winds of the outdoor location caused noisy data, sometimes to the point of being unusable. The test stand was configured with the short (3m) arm. Several hardware issues were identified and resolved during this testing, and the algorithm was tuned to provide best results. Due to inertia in the ornithopter and the support arm, long sample times were required. Each individual was evaluated for 15 seconds, with the first 1/3 of the data points of each evaluation disregarded. Runs were observed continually by the researchers, and usually terminated after 100 to 150 evaluations, due to ambient conditions, equipment problems, or no obvious movement toward convergence.

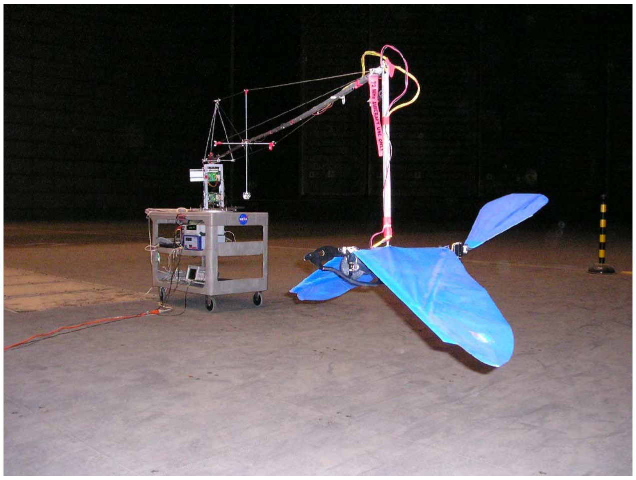

Figure 8. Test Stand with Long Arm in Wind Tunnel.

After debugging, the test rig was moved to a large enclosed area at our research site, shown in Figure 8. This area (a temporarily unused wind tunnel test section) provided protection from ambient winds, as well as a secure location for lengthy experimental runs.

Figure 9. Test Stand Setup with Laser Alignment Tool.

The larger area allowed the use of the long arm (6m) on the test stand. This produced more realistic flying conditions for the model, as the ornithopter could fly a less curved path. The long arm proved to be more flexible than desired, and runs would frequently be terminated when the arm went into a "bent over" condition. Analysis and modifications to the tension cabling on the arm corrected the problem. A laser alignment tool, shown in Figure 9, was added to ensure symmetrical stresses on the structural elements of the arm.

Once set up in the indoor test area, the hill-climber algorithm was changed to a steady state evolutionary algorithm. A population of twenty was found adequate for consistent, rapid convergence.

Individual evaluation time was kept at 15 seconds. Hardware, software, and communications problems lead to challenges in completing runs of more than 150 evaluations.

After several iterations of testing and improvements to the system, extended runs were possible. Eventually, the ornithopter could be left running unattended overnight, performing many successive runs and providing considerable data. A total of 12 satisfactory runs were completed over a period of several weeks, with run lengths of 150 to 500 evaluations, as shown in Figure 10.

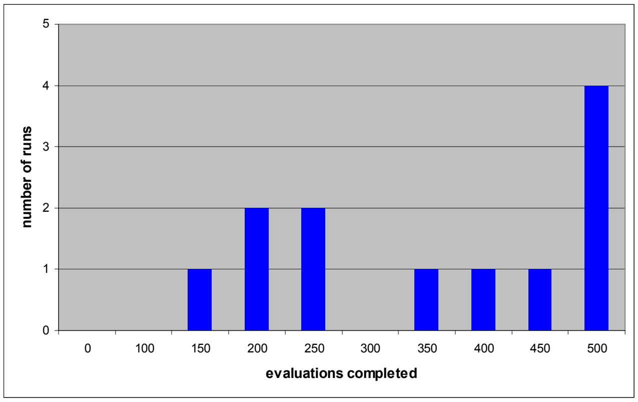

Figure 10. Histogram, Evaluations per Run.

As runs progressed, ornithopter components began to wear and perform differently. The drive gears, crankshaft linkages, and wing hinges all loosened, gradually reducing the magnitude of flapping. The primary effect was that the maximum achievable forward velocity decreased over the course of the runs, starting at about 20 ft/sec, and falling to about 17 ft/sec for the final runs.

A mechanical failure ended the first set of data runs, but provided important knowledge for our future research. At the end of a long series of overnight runs, a screw loosened at the shoulder hinges and fell out. As the bird continued to flap (and the algorithm attempted to accommodate the change), the hinge plates moved in their slots and wore away at the fiberglass body. When the screw was replaced, glue was applied to strengthen the joint and prevent any more wear. This created a small change in the position of the hinges and thus the motion of the wings. When attempting to fly the ornithopter after the repair, the bird exhibited a strong tendency to go into a severe nose-down attitude at high flapping rates. This was eventually attributed to an asymmetry in flapping motion with respect to the horizontal plane, causing a net down pitching moment that increased with flapping rate. Adjustments to the crankshaft linkages partially corrected the problem, but an inability to adequately replicate earlier results lead to the end of this phase of the research. However, knowledge of this pitching effect provides a key optimization parameter for our ongoing research.

6. RESULTS

Using the steady-state evolutionary algorithm running in the indoor test area, here we present the results of optimizing for both forward velocity and power efficiency. In these trials, the first twenty "evaluations" are randomly selected individuals, comprising the population. Initial throttle command is constrained to the range of 0-50%, requiring the algorithm to evolve the parameter out of its initial range to reach optimum fitness. After the first twenty points, the algorithm begins evolving a solution.

In our first set of experiments, we ran 12 trials with varying numbers of evaluations to optimize for forward velocity – due to hardware issues not all trials ran the full 500 evaluations. The average of all data from the 12 runs is shown in Figure 11. This graph shows that flight was consistently optimized over the course of the evolutionary runs. The values in this chart were calculated by taking the average and standard deviation of all runs for each evaluation. For example, the average and standard deviation for evaluation 125 in Figure 11 are the average and standard deviation of evaluation 125 in all 12 runs. Since some of the runs did not go to 500 evaluations (see Figure 10), the average and standard deviation of higher numbered evaluations have less data points than lower numbered evaluations. However, a comparison of runs with up to 250 evaluations versus runs with over 250 evaluations did not show a significant difference in averages or standard deviations.

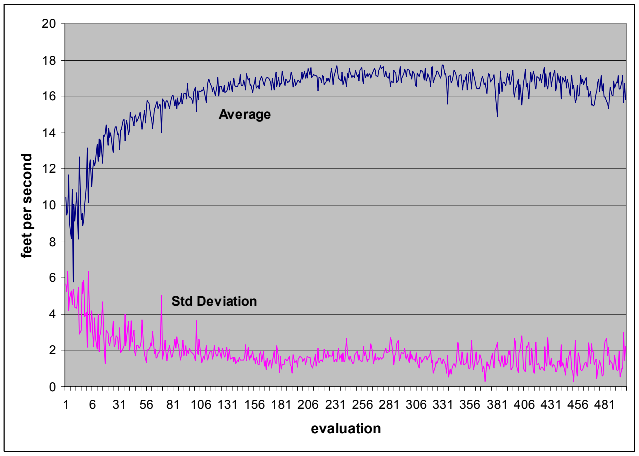

Figure 11. Velocity Optimization, Average of 12 Runs.

We also monitored the values of the three command parameters over the course of each run, as well as evolved a fourth null parameter. Scaling on the parameters is as follows. Throttle is the command parameter for the Electronic Speed Control, ranging from 0-100% of full throttle. Elevator is measured as degrees of inclination from horizontal, as referenced to the longitudinal axis of the ornithopter. Rudder angle is degree of rotation from vertical, with negative rotation to the left, or counter-clockwise as viewed from the rear of the bird. The fourth parameter is not used in controlling the ornithopter, and is included as a control variable to compare evolution on a gene that does not affect performance.

The command parameters converged over the course of each run. Averages and standard deviations at each evaluation for all data runs, for each parameter, are shown in Figs. 12a-12d. The three commands throttle, elevator, and rudder all start at the mid-point of their ranges and converge over time. As expected, the null parameter does not converge, instead its average remains at approximately the mid-range throughout the runs and its standard deviation does not decrease. This demonstrates that evolution is optimizing the three parameters that affect flight performance and is not optimizing the null parameter.

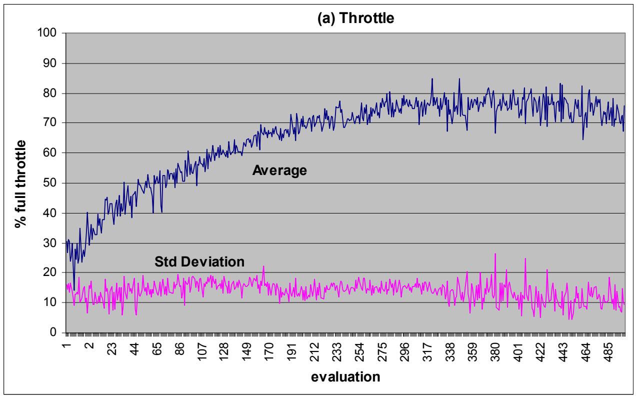

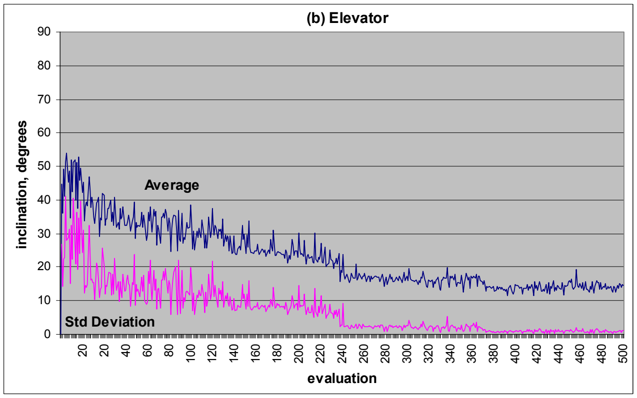

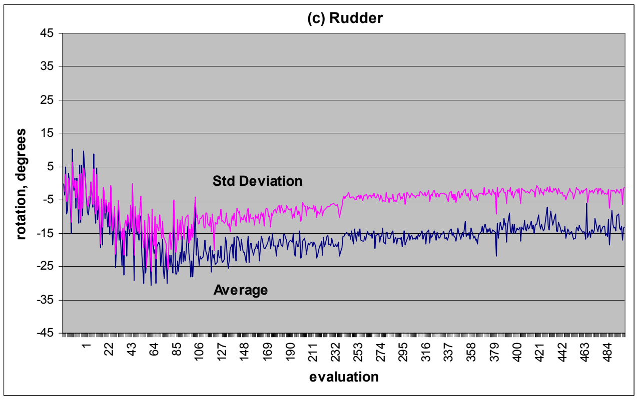

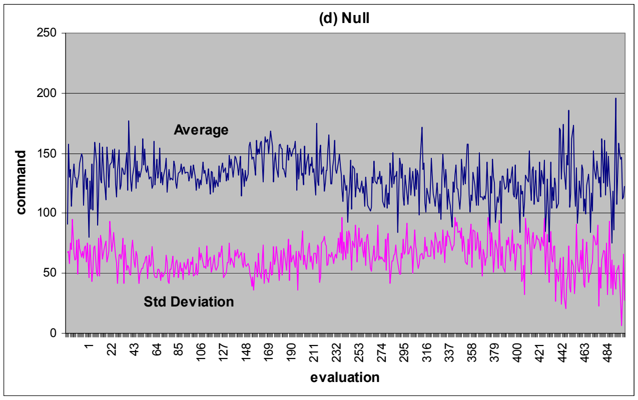

Figure 12a,b,c,d. Command Parameters, Average of 12 Runs.

The second optimization performed was efficiency, as represented by forward velocity divided by motor power. The average of 8 efficiency optimization runs is shown in Figure 13.

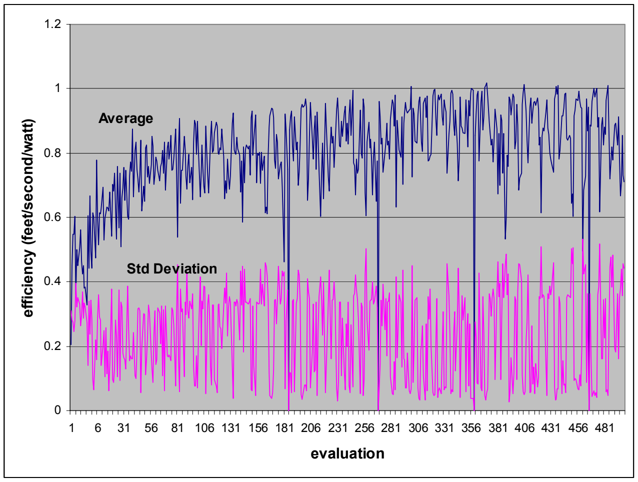

Figure 13. Efficiency Optimization, Average of 8 Runs.

Again the algorithm is able to optimize the fitness function, in this case finding solutions for optimal propulsion efficiency. Since lift is not yet a fitness criteria, the highest efficiency is achieved at a low flapping rate and low forward velocity, as seen in Figure 14.

Examination of individual runs shows numerous evaluations resulting in zero forward velocity, and thus zero efficiency. At the low flapping rate (and low power) which produce the highest efficiency, it is easy for the bird to lose all forward velocity. The algorithm is able to repeatedly recover and continue this more difficult optimization.

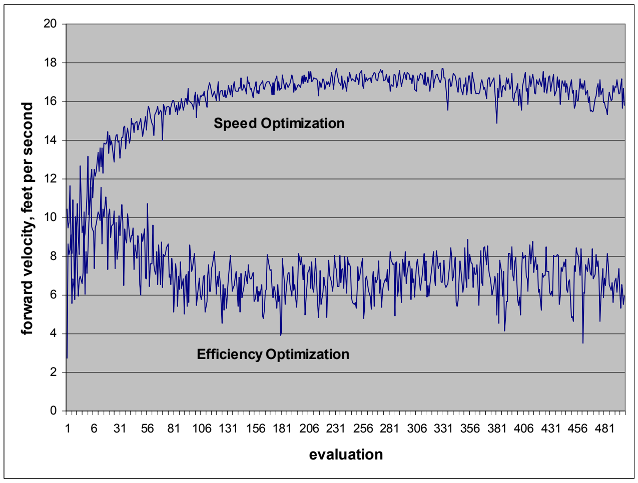

Figure 14. Forward Speed during Different Optimizations.

Figure 15 shows the bird in flight on the test rig, controlled by the evolutionary algorithm. The ornithopter is flying at the optimal configuration for best forward velocity, as found by the algorithm. The tail is nearly flat in inclination, and rotated slightly away from direction of travel, and the flapping rate is high.

Figure 15. Ornithopter in Flight, Optimized for Speed.

7. DISCUSSION

In general the ornithopter and the evolutionary algorithms performed well and behaved as expected, with a few interesting exceptions. In forward velocity optimization, the throttle command was not at its maximum and the tail position was opposite of expected.

The expected configuration for maximum forward velocity was for the wing flapping rate to be as high as possible. Results found by the algorithm and confirmed by manual radio control show that highest forward speed is actually achieved with the throttle at 75-80%. Investigation of the pitching effect seen when the shoulder hinge was repaired provides an explanation. During the speed optimization runs, the bird was flying with a slight asymmetry in flapping, with the wings above the horizontal plane slightly more than they were below the plane. Since the wings provide thrust as well as lift, they impart a downward pitching moment when above horizontal, and an upward pitching moment when below horizontal. With slightly more of their total travel above horizontal, there is a net downward pitching moment which increases with thrust derived from flapping. As flapping rate is increased the bird flies faster but the downward moment also increases. To counteract this, the tail must be inclined from horizontal to create an upward pitching moment. This tail inclination, however, also increases drag and slows the bird. The EA successfully found the optimal combination of flapping speed and tail inclination for best forward speed.

The optimal tail position was also different than expected. With the ornithopter flying in a clockwise circle (viewed from above) the bird is in a constant right turn. Intuitively one would expect the tail rotation to be to the right, pushing the tail to the left and thus the nose to the right, into the turn. In the optimization runs, however, the best forward speeds were obtained with the tail rotated about 15 degrees to the left, opposite of expectation. Manual flight via radio control seemed to confirm this as the optimal position for best forward velocity. The mechanism is not fully understood, but appears to be related to the amount of drag produced by the tail in various positions versus variations in overall drag at different body attitudes.

8. CONCLUSIONS

We have demonstrated the use of evolutionary algorithms in the optimization of flight parameters on a hardware ornithopter. The EA successfully evolved the bird's flapping speed and tail position to provide maximum forward velocity, and also maximum efficiency. By performing evolution on hardware, evolutionary optimization was able to take place without a thorough analytical understanding of the physics of the system.

As the first phase of our research, this work has demonstrated the integrated functionality of the ornithopter, test apparatus, and the computer hardware and software required to control the ornithopter. Further research toward evolved flight includes propulsion optimization by EA-controlled variations on the drive system, evolution of free-flight guidance and control, and co-evolution of tightly integrated sub-systems for optimal system performance.

9. REFERENCES

[1] Lillianthal, O. *Birdflight as the Basis of Aviation,* Translation from the second edition, Longmans, Green, and Col, London 1911.

[2] DeLaurier, J. D. *The Development of an Efficient Ornithopter Wing.* The Aeronautical Journal of the Royal Aeronautical Society, May 1993.

[3] Schenato, L., Campolo, D., Sastry, S. Controllability Issues in Flapping Flight for Biomimetic Micro Aerial Vehicles (MAVs) *IEEE International Conference on Decision and Control, Maui, Hawaii, December 2003*

[4] Lohn, Jason D., Hornby, Gregory S., Linden, Derek S. An Evolved Antenna for Deployment on NASA's Space Technology 5 Mission *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2004)* June 26-30, 2004, Seattle, Washington.

[5] Bugakska, Maddalena D., and Schultz, Alan C. Co-Evolution of Form and Function in the Design of Autonomous Agents: Micro Air Vehicle Project. *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2000)* July 8-12, 2000, Las Vegas, Nevada.

[6] Hornby, G. S., Takamura, S., Yamamoto, T., and Fujita, M. *Autonomous Evolution of Dynamic Gaits with Two Quadruped Robots.* IEEE Transactions on Robotics, Vol 21, No 2, April 2005

[7] Augustsson, P., Wolff, K., and Nordlin, P. Creation of a Learning, Flying Robot by Means of Evolution. *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2002)* July 9-13, 2002, New York, NY.

[8] http://www.centennialofflight.gov/essay/Evolution_of_Technology/first_wind_tunnels/Tech34.htm

[9] Tennekes, H. *The Simple Science of Flight – From Insects to Jumbo Jets.* MIT Press, 1992.

[10] Taylor, Graham K., Nudds, Robert L., Thomas, Adrians L.R. *Flying and Swimming Animals Cruise at a Strouhal Number Tuned for High Power Efficiency.* Nature 425, (October 16, 2003), 707-711