

# Towards Evolvable Hovering Flight on a Physical Ornithopter

Will Regan, Floris van Breugel, Hod Lipson

Computational Synthesis Lab  
Cornell University, Ithaca NY 14853, USA  
1 (607) 255 1686  
hod.lipson@cornell.edu

## Abstract

By presenting the first combination of evolution of simulation and hardware implementation of a hovering flapping ornithopter, this paper paves the way towards achieving artificial flight on par with nature's elegance. Our preliminary results show the feasibility of this form of flight. With the successful evolution of simulated models, we developed a hardware model to embody these findings and show physical progress towards sustained untethered hovering flight.

## Introduction and Related Work

Humans have long been fascinated by flight, but have still not yet succeeded in creating a device capable of sustained untethered hovering solely by flapping (Augustsson, Wolff, and Nordin 2002). There are, however, many such designs in nature, ranging from hummingbirds to small insects. Hitherto man has always relied on the use of an airfoil and forward motion to create enough lift to fly; for examples, see DeLaurier, 1993. Interesting work on forward flying ornithopters has also been done by Hunt, 2005 and Michelson, 2003. The Mentor robot, built by SRI and the University of Toronto, has demonstrated hovering and forward flight but still seems to require tethered control (Larijani 2001). The challenge our project addresses is the creation of an ornithopter that flies only by means of the drag forces on a flat wing. This is a difficult problem and would be nearly impossible to solve manually. Thus it provides an excellent challenge for evolutionary design (Nolfi and Floreano 2000).

Much work has been done with evolving flight in simulation, but few attempts have been made to actually build the resultant machines or even consider their production or practicality in terms of weight to lift ratio (Shim, Kim, and Kim 2004) and (Wu and Popovic 2003). Most work on evolutionary ornithopters is concerned with the flapping motion itself, without regard to the kinematic mechanism that might generate it in practice. Those attempts that have been made have shown promise (Augustsson, Wolff, and Nordin 2002), but we are thus far unaware of any successful evolved ornithopters that fly without the aid of forward motion.

There are several advantages to developing an ornithopter that does not require forward motion to fly. The primary benefit of such flight is that it would allow the machine to hover. Currently the only such devices are helicopters, which are able to create limited amounts of lift due to the rotor design. Flapping flight confers many advantages such as increased agility, independence of body orientation, and high adaptability to different situations (Smith 1997). Such mechanisms are of great interest to aerospace engineers, military operations, search and rescue, and surveillance technologies (Smith 1997).

Other projects are attempting to solve this problem using the benefit of scale (Yan et al. 1997) (Motazed 1998). It has been shown that the fluid dynamics important to flight vary with the size of the machine. On a small scale the air acts more as a viscous fluid through which insects must swim, and an airfoil gives no benefit.

The goal of our project is to use evolution to create a large-scale ornithopter (on the order of a one meter wingspan) with evolved control patterns. The larger scale could allow sizable payloads and makes adjustments easier. The first stage was to create simulations of a realistic ornithopter and show that successful control patterns could be found using genetic algorithms. Once this was accomplished, a physical model based on the simulations was built and evolved.

## Use of Evolutionary Algorithm

The physics of controlled flapping flight is a complicated problem with many linked parameters. The chances of a programmer manually designing a successful flapping pattern as well as the associated mechanical design to achieve that motion are slim. Yet there are thousands of species of insects, birds, and bats that are all well adapted to flapping flight. Modeled on the natural processes that have produced such remarkable mechanisms, simulated evolution has been shown to be a powerful tool for solving such difficult problems and producing effective behaviors (Nolfi and Floreano 2000).

## Dynamics Simulation

The first stage of this project consisted of computer simulations of realistic designs in order to verify the feasibility of actually creating a flying machine with the specified dimensions and evolutionary process. The physics was modeled using a rigid body dynamics simulator (ODE), with an emphasis on speed rather than physical accuracy (Smith 2000). ODE does not include any aerodynamic equations. Aerodynamics were modeled using the following simplified drag equation (Fox 1976):

$$Drag = \frac{1}{2} \rho C_d A_{\perp} V_{\perp}^2$$

Here  $\rho$  is the fluid density and  $C_d$  is the drag coefficient of a flat rectangular plate. In order to keep the simulation as close to the real physics our physical model would experience, full gravity and reasonable weights, sizes, torques, and flapping frequencies were used.

## Simulations

Various simulation arrangements and evolutionary techniques were tried. The following is a brief overview of the relevant experiments performed, but for a more in depth discussion see van Breugel, 2005. In order to simplify control issues, four wings were chosen rather than two, ensuring that all horizontal forces and torques cancelled, leaving only the vertical component of force. Each wing, fixed to the body by a ball-and-socket joint, was controlled by three servo motors attached to the wings through a symmetrical cable system, such that when the servo turned in one direction, the cable would give slack on one side of the wing and pull on the other side. In order to simplify the problem, an arbitrary fixed mechanical setup was used and Bezier control patterns were evolved to control the speed with which the servos pulled on the cables. Using a standard elitism genetic algorithm, successful and realistic flight was achieved.

Using a population size of 100, a weight of 310 grams, and a maximum linear force of 50N on each of the cables, the fittest member after 73 generations flew upwards with a maximum total lifting force of 190 N and a maximum upwards velocity of 12 m/s. The flapping frequency was slightly fast at approximately 6Hz, and there were some acceleration peaks that were a bit high for the scale. Overall, however, the setup seemed likely to be capable of producing lift with a real model.

## Transition to Hardware

Given the encouraging results from ODE simulation, we set upon designing and constructing a physical ornithopter based on the four-winged cable-controlled simulations. Ideally the physical model would match the simulated model as much as possible, so the simulations could be

used as a tool to help evolve successful control patterns. Unfortunately, however, the simulations are only a very rough approximation of how the physical model would behave. Aerodynamics are more complicated than the simplified drag equation used; wing bending, resonance, and mechanical play in real joints cause the simulations to be too different to use for direct evolution. In order to more accurately model the system, the number of evaluations needed would become prohibitive. Thus the physical ornithopter was designed as close to the simulations as mechanically feasible, though evolution was started from scratch. Furthermore, design choices were made to allow some variability in physical parameters such as positions of the servos, the attachment points of the wing actuators, the point of attachment of the wings to the body, and the wings themselves. It was thought that these physical parameters might be evolved by observing the rough behaviors of the ODE model.

## Physical Model

The first step in this process was to find components that were strong, quick, and light enough to perform at the level of the simulated model. The main considerations for the hardware were the motor torque and speed, wing area, and overall mass of the craft. The overall mass is the combined weight of the servos, the wings, the battery and onboard control circuitry, and any “connective tissue” or systems used in actuation of the wings.

It soon became apparent that perhaps the most important factors were the speed and power to weight ratio of the servo motors. For an approximately 300g craft, the servos would comprise upwards of  $\frac{1}{4}$  to  $\frac{1}{2}$  of the weight, and would have to provide proportionally greater lifting power. Further, the servos would have to be quick enough to provide flapping frequencies up to a few Hz.

The sheer number of servo manufacturers makes it hard to do a comprehensive comparison of all available on the market, but a good list of resources has been compiled by Ryan Carr, 2005. The best servos we found had the following properties (at 6V DC):

Torque (oz-in)	Speed (sec/60deg)	Weight (grams)
89	0.14	28
100	0.08	45

The speed of the heavier servos could allow flapping at upwards of a few Hz, if cleverly actuated. As a rough measure of the torque needs, we assumed an exact hovering situation: the work done by the torque in a wing down stroke (a half period) would have to equal the loss in gravitational potential for a free fall of one period, a

fall of distance  $h(T)$ . (Note: It was assumed that two of the heavy servos and one light servo were to be used.) Analytically, this yields:

$$N \theta = m g h(T) = \frac{1}{2} m g^2 T^2$$

Assuming  $\theta$  (angle swept out in  $\frac{1}{2}$  cycle)  $\sim \pi/2$  for a full down stroke,  $T$  (period)  $\sim 0.3$  seconds, and  $N$  (torque)  $\sim 290$  oz-in, this suggests that the servos are capable of supporting a max mass of about  $m = 740$  g, greater than the projected total mass. Granted, this estimate takes both the maximum torque and speed (and assumes we can transfer servo power to flapping power quite efficiently), but even a reduction by  $\frac{1}{2}$  in the torque would allow hovering for  $m = 370$  g, still greater than our target mass.

The wings, selected for their lightness ( $\sim 25$  g each) and large area ( $\sim 1400$  cm<sup>2</sup> each), were taken from a biplane kit. Though they were not designed for strength (loading of a few oz per square foot), we assumed they could be reinforced if needed.

The battery was a 5C 7.4 V 850mA Lithium polymer (LiPo) cell, chosen for its extremely high energy density. The battery added an additional 40g.

The control circuit, weighing in at 25g, was an assembly of a Basic Stamp 2 (BS2) and PWMPAL, allowing easy RS-232 control of the three onboard servos.

Finally, the main body was to be constructed from a material with as great a strength/lightness ratio as possible. Candidates were balsa wood, carbon-carbon fiber sheets, and light acrylic polymers. This construction was left variable until we settled on a final design, and so body/actuation system mass estimates at this point were left variable and assumed to be about 80g.

With these reasonable assumptions, the estimated weight was 313 g, which according to the simulation and estimates of servo power made flight seem feasible.

## Design

Next began the phase of designing a lightweight and sturdy body and actuation system.

Beyond being strong enough to withstand the forces of flapping and having a small footprint relative to the wing area, the body morphology was somewhat irrelevant. Thus, the main consideration was the method of actuation.


Also, to simplify the construction we made only one half of the simulated ornithopter (we kept one pair of wings, arranged to resemble a bird) with the intention of building and connecting the other half if needed for stability. Thus, the physical model was comprised of three servos (motors whose output rotation is restricted to 180 degrees) symmetrically controlling the three DoF of two wings, with the wings arranged about the body's center like a bird. The wings themselves were attached to the

body by a ball-and-socket joint, which would permit each wing to move freely in its own hemisphere.

**Cable Actuation.** Initially, we were tempted to try an actuation system similar to the cable actuation simulation to glean some of the flapping properties of its fitter individuals.

In cable actuation, each servo horn would be connected by a loop of string to a point at the bottom and top of the wing. An example of such a system is shown in Figure 1. The servos have large circular horns to which the string is fixed. One end of the string attaches to the wing bottom, whereas the other end passes through a pulley on a center tower and then attaches to the wing top.

This type of design is at first appealing, since string loops would add very little weight.




**Figure 1: CAD model of cable-actuation design. Two large servos (circular horns) provide main flapping power. Small servo controls yaw.**

However, a geometrical analysis with just two servos reveals that the allowed wing motion is quite limited. Each complete string loop has a constant total length and two invariant points (where it attaches to the servo and pulley) and so constrains the wing to an ellipsoidal surface. Therefore, the only allowed wing positions lie in the intersection of two

ellipsoids (for two servos) and one sphere (wing fixed to the body with a ball-and-socket joint), which is at best a line. Even with room for play in the string length (giving some thickness to these surfaces), the range of positions is at best two warped planes, as shown below in Figure 2. A better arrangement would have allowed a full hemisphere of wing motion.

A simple physical prototype was also built to confirm this behavior,




**Figure 2: Permitted motion of point on wing in realistic cable actuated design**

was also built to confirm this behavior, with springs spliced into each string loop to give much play in string length. A large range of motion was permitted, but with the drawback that the wing position was so free as to become underdefined.

**Pushrod Actuation.** So instead of cable actuation, we tried a more direct form of actuation still faithful to the simulated machines. The simplest solution found for well-defined actuation was simply rigid pushrod connections from the servo horns to points on the wing frame.

Although pushrod actuation makes the motion of the three servos more coupled, it allows a full range of motion without the worry of the connectors going slack, as happens with cable actuation. Furthermore, rigid metal pushrods would be less likely to break during the taxing evolution process than would string loops. We produced the following design.



**Figure 3: CAD model of the ornithopter**

Figure 3 shows a view of the craft from the side and from above. To deliver maximal vertical forces for flapping, we placed the two strong servos symmetrically about the center of the craft, directly under the wings. As these two servos rotated upwards, the pushrods pushed the wings upwards and hence yielded an up-flap; rotating downwards similarly caused a down-flap. The third servo (seen in Figure 3 in the upper right corner of top image and top center of bottom image), was responsible for yaw and was set back just far enough to be out of the path of the wings. The wings were connected to a central tower and elevated so that the lifting forces were applied well above the craft's center of mass for stability reasons.

For the structure material, we chose a combination of carbon fiber structural rods and an ABS body, the body being printed with an FDM 3000 rapid prototyping machine. Each servo horn is connected symmetrically to the two wings, secured to a carbon fiber frame to prevent wing damage during strong flapping.

This design allowed several parameters to vary, in case testing indicated a fitter arrangement was possible. Among these parameters were servo position along the length of the body, wing type (wings could be swapped in and out of the frames), the wing frame (which defines the attachment points of the servos), and wing attachment height. To clarify these points, if one were to make an analogy between the craft and a bird (with the wings arranged similarly), the servos could be moved up and down along the body's spine and the wing frame could be modified to move the wing's muscles.

As construction and assembly of the craft proceeded, extra weight added by structural reinforcements and false servo specifications (weight underestimated) brought the total craft's weight up to 440 g, a fair amount greater than the projected 313 g but still within the realm of possible flightworthiness. To get a sense of the possible wing motions, we moved the servos manually to their extreme positions; we observed plenty of yaw and roll freedom ( $\pm 50$  degrees for each), though pitch (the angle of twist of the wings) was somewhat limited at a max of about  $\pm 35$  degrees.

### Interfacing

Previous attempts at evolutionary ornithopter design, such as (Augustsson, Wolff, and Nordin 2002), relied on external power and external servo controllers. However, the long term goal for our physical model is untethered flight, which explains our choice of a lightweight power source (LiPo battery) and control circuit with onboard memory (for storing servo commands) in our design.

During the process of evolution, however, it was necessary to have an external interface in order to reprogram the control circuit and take system measurements for the fitness metric.

The first link made between the craft and computer was serial (RS-232) communication with the Basic Stamp. Short strings of servo commands, corresponding to angular positions, were sent to the onboard stamp with the Basic Stamp Editor v2.2. Once we confirmed that we could manually command the wings to flap, we moved on to system data acquisition. As a simple measure of fitness, we measured only the vertical component of force on the craft with a one dimensional LCEB-10 Omega Load cell (capable of measuring  $\pm 10$  pounds of force). Though more might be learned about the system with a three dimensional multi-axis force and torque cell, or perhaps tilt sensors in all directions, we reasoned that the best hovering individuals must maximize their lift and minimize all other forces and vibrations. The craft was mounted directly onto the load approximately through the craft's center of mass.



**Figure 4: Ornithopter machine**

Next, measurements from the load cell (a few mV) were amplified and offset to a range of 0 to 5 Volts DC with a simple operational amplifier circuit. The output of this circuit was fed into a Hytek Automation U120816 12-bit USB data acquisition module, which streamed data directly into our evolution code, thus completing the full loop required to evolve.

**Evolution of Control Patterns.** Given the success of the simulations, a similar evolutionary approach was used in order to evolve the hardware. Again, the physiology of the device was fixed and only the control patterns (the sequence of servo positions over time in a flapping cycle) for each of the three servos were evolved. Several variations on control patterns were tried.

In accordance with the simulations, the first type of servo control tried was Bezier control. A set of 10 “control points” was used to generate a Bezier curve, from which points were sampled every 0.2 sec and then used to generate a pattern of servo position commands, which was sent to the Basic Stamp. Converting these sampled Bezier points into servo positions was simply a linear transformation from the Bezier domain [-1.0, 1.0] to an output range of [30, 85], which corresponds to the allowed range of servo positions.

After testing many Bezier control patterns, it became apparent that smooth servo motion (and hence flapping) was only achievable by drastically decreasing the Bezier periods (our Bezier curves were periodic) or increasing the sampling rate, since the servos generally reached their targets quickly and had to pause until the next command. When we decreased the Bezier periods, we were still faced with jerky servo motion, perhaps due to aliasing (getting an effectively larger Bezier period since we were

still sampling slowly). As an alternative to Bezier control patterns, we instead chose randomly generated control patterns.

Two to four second cycles of random servo commands were created, which were then looped several times in the Basic code. We discovered that loop overhead caused the Basic Stamp to pause between loop iterations, proportional to the number of servo commands in the loop, so we restricted the maximum number of servo commands per loop to 8 to limit this delay and achieve smooth flapping. Thus, the variability in genomes is limited to 8 commands for each servo, which can be cycled indefinitely. These limitations still allow a large range of resultant behaviors, though we could have escaped this limitation altogether by streaming the motor commands to the craft through a control circuit such as the SV203B servo controller.

**Evolutionary Algorithm.** As in the simulations, a standard elitism algorithm (Mitchell 1998) was used to evolve the three servo control patterns. For each individual of the initially random (and later evolved) population, the control pattern was extracted from the genome, converted to servo positions, and transformed to a PBasic source file of servo commands, using the SEROUT command to talk to the servos. This source was then written, downloaded, and run on the craft’s onboard Stamp. (For consistency, all individuals were first initialized to the same servo position, the equilibrium of the servos and hence that of the wings.) During the subsequent flapping of each individual, the lift was continuously measured by the load cell and recorded. After each individual run, lift was averaged and converted to a fitness scale of 0 to 100, corresponding to the average percentage of the craft’s weight lifted. After each generation, the population was ranked by fitness, and the top 40% of the individuals were chosen randomly and evenly as parents to produce offspring to replace the bottom 40%. The middle 20% remained unchanged.

Matings were all between two different parents and produced two new offspring. Mutations were additively applied at an average rate of one per genome with a low-end weighted exponential range between 0 and  $\pm 1$ . The gene range was limited to  $\pm 1$  and scaled for appropriate range later, with mutations causing genes to exceed this limit capped. This ensured occasional large mutations and more commonly minor adjustments for fine-tuning. Two-point crossover was implemented at a rate of 80% in order to allow mixing of genes and increased variation. This method preserves possible functional groups of genes by moving large sections together.

Since the craft and load cell were limited in motion to up/down, the only forces registered in the fitness function were vertical and there were no penalties for control. This removed both the added complexity of having four wings,

as well as the need to deal with control issues as described in the two winged simulations.

The effect of noise on the load cell data was tested by running the same set of 10 random individuals 5 times and comparing the fitness from each run. The measured fitness error for these individuals was under 15%, low enough to warrant single evaluations for each individual.

## Results

In the best set of tests, five control points (spaced over a flapping period between 0.5 and 1.0 seconds) for each of the three servos were evolved. An initially random population of 40 random individuals was used. As seen below, the fitness improved dramatically in the first few




Figure 5: Best Maximum Fitness v. Generation

generations. Then after stagnating at around 10 generations, the evolution again picked up and brought the best fitness to nearly 21% lifting ability. This initial stagnation is best explained by the fact that hovering flight by flapping is a difficult problem, and with a small population size it can be difficult for evolution to find a good solution.

As noted before, the fitness was scaled such that zero fitness is equivalent to zero lift and 100 is equivalent to hovering – lift equal to the machine’s own weight. Interestingly no individuals appeared much below a fitness of 2, even in the random population. This is attributed to the fact that the wings are asymmetric due to the curvature, and thus always seem to present a larger area on a down-stroke than on the upstroke.

Qualitatively, the fittest individuals tended to show out-of-phase motion in the two primary servos responsible for

up and down power strokes. This causes the wings to present the full surface area on the down stroke and an angled, and thus smaller, area on the upstroke, yielding a net lifting force. This type of behavior has appeared in previous work on several occasions, including our ODE simulations of both four and two winged models, as well as in Augustsson’s work, 2002.




Figure 6: Pictures of evolved flapping cycle

This pattern can best be described as a figure-8 motion akin to the flapping patterns exhibited by both hummingbirds and dragonflies. Additionally, the servos seemed to work together more in later generations, causing both a higher wing amplitude and increased speed. Another high fitness pattern which emerged resembled a swimmer’s frog kick with a slow up stroke. This pattern was similar to the fittest control evolved for the Bezier cable control simulations.

In an effort to improve lift, several physical variables were modified; the wing frame was modified to produce a larger up/down flapping angle (by moving the pushrod attachment closer to the craft’s center), the servo positions were varied (along the “spine”), and alternate wings (larger, flat balsa wing) were tested. However, with none of these changes did we achieve beyond 21% lift.

## Discussion & Future Work

Although the evolved patterns were only capable of lifting just over 21% of the machine’s total weight, we saw a clear increase in both maximum and average fitness over 75 generations. This upper limit of performance is most

likely due to two factors, the weight and the highly coupled nature of the servo actuation. Considering the craft's bulk, 440 g, 21% lift means that an average of 93 g could be hovered, still a sizable amount. Several factors could reduce the weight on future models and make flight feasible. Alternate actuators with better power to weight ratios than servos could make the most impact; servo motors have rather low power to weight ratios, though ours were rated on the high end. Also, weight could be saved by removing many of the reinforcements, as flapping was not as destructive as expected. Next, the coupling of the servo motion meant much of the servos' power was wasted fighting among themselves. Due to the position-correction circuitry in the servos, movements in any one servo would cause a cascade of reactions in all servos. This fighting diminished the strength and amplitude of the flapping. Combining the mechanical limitations of the design (+/- 50 degrees in roll, yaw and +/- 35 degrees in pitch) with this diminished amplitude prevented the craft from achieving many flapping patterns which may have proved optimal, such as a sharp pitch angle before upstrokes (to reduce the drag on upstrokes).

The ultimate goal of this project is to evolve a hovering ornithopter capable of independent and untethered flight. In order to make the problem easier we plan to begin working on a lighter model, as well as utilize more concepts inspired by biology. Insects, for example, exploit passive dynamics of the system such as resonance, storing potential energy in elastic muscles and fibers. Using these ideas, we hope to construct a smaller model weighing closer to 40 grams, which can be evolved using the same successful setup shown in this paper. Hopefully, however, we will be able to use more insight from the simulation process so as to evolve both morphology and control patterns. In order to reduce the effects of the discrepancies between the simulation and hardware, we will begin implementing the Exploration-Estimation Algorithm (Bongard 2005).

**Acknowledgments.** This work was supported in part by a NASA/NY Space Grant Consortium Fellowship and by the Cornell University Presidential Research Scholars program.

## References

- Atmar, W. 1994. Notes on The Simulation of Evolution. *IEEE Trans. Neural Networks* 5(1): 130-148.
- Augustsson, P., Wolff, K., and Nordin, P. 2002. Creation of a Learning, Flying Robot by Means of Evolution. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 02)*.
- Bongard, J. and Lipson, H. 2005. Nonlinear System Identification Using Coevolution of Models and Tests. *IEEE Transactions on Evolutionary Computation* 9(4): 361-384.
- Carr, R. 2005. The Gigantic Servo Chart. Accessed 05/10/05. <http://www.fatlion.com/sailplanes/servochart.html>.
- DeLaurier, J. D. 1993. An Aerodynamic Model for Flapping Wing Flight. *Aeronautical Journal* 97(964): 125-30.
- Fox, R. W. and McDonald, A. T. 1976. *Introduction to Fluid Mechanics*. 5<sup>th</sup> ed. New York: Wiley.
- Hunt, R., Hornby, G., and Lohn, J. 2005. Toward Evolved Flight. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 05)*.
- Larijani, R. F. and DeLaurier, J. D. 2001. A Nonlinear Model for the Study of Flapping Wing Flight. Fixed and Flapping Aerodynamics for Micro Air Vehicle Applications, Part II: Flapping and Rotary Wing Aerodynamics. *Progress in Astronautics and Aeronautics AIAA* 195: 399-428.
- Michelson, R. and Naqvi, M. 2003. Beyond Biologically-Inspired Insect Flight. In *Proceedings of von Karman Institute for Fluid Dynamics RTO/AVT Lecture Series on Low Reynolds Number Aerodynamics*. Brussels, Belgium.
- Mitchell, Melanie. 1998. *An Introduction to Genetic Algorithms*. Cambridge, Mass.: MIT Press.
- Motazed, D. Vos, and Drela, M. 1998. Aerodynamics and Flight Control Design for Hovering MAVs. In *Proceedings of American Control Conference*, 681-683. Philadelphia, PA.
- Mueller, T. J. 2001. Fixed and Flapping Wing Aerodynamics for Micro Air Vehicle Applications. *AIAA Progress in Astronautics and Aeronautics* 195.
- Nolfi, S. and Floreano, D. 2000. *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. Cambridge, Mass.: The MIT Press.
- Ramakrishnananda, B. and Wong, K. C. 1999. Animating Bird Flight Using Aerodynamics. *The Visual Computer* 15(10): 494-508.
- Shim, Y., Kim, S. and Kim, C. 2004. Evolving Flying Creatures with Path Following Behavior. In *Proceedings of the 9<sup>th</sup> International Conference on the Simulation and Synthesis of Living Systems (ALIFE IX)*. Boston, Mass.
- Smith, M. J. C. 1997. Reinstating Inquiry Into Mechanized Flapping-Wing Flight: Realizing the Integrity of the Ornithopter. In *Proceedings of the 35th Aerospace Sciences Meeting and Exhibition*, 1-11. Reston, VA: AIAA.
- Smith R. 2000. Open Dynamics Engine (ODE), <http://ode.org/>.
- van Breugel, F. 2005. Evolving Buildable Flapping Ornithopters. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 05)*.
- Wu, X. C. and Popovic, Z. 2003. Realistic Modeling of Bird Flight Animations. In *Proceedings of SIGGRAPH '87*, 888-895.
- Yan, J., Wood, R., Avandhanula, S., Sitti, M. and Fearing, R. 1997. Towards Flapping Wing Control for a Micromechanical Flying Insect. In *Proceedings of the IEEE International Conference on Robotics and Automation*.