



LARGER THAN LIFE:
This RoboBee is magnified
to show detail; turn the
page to view at actual size.

IN BRIEF

RoboBees are flying robots the size of bees. Their size presents a huge assortment of physical and computational challenges. At such small dimensions, off-the-shelf parts such as motors and bearings will prove too inefficient, so the bees must employ specially designed artificial muscles to power and control flight.

In addition, the tiny bees must think on their own, using

miniature sensors to process environmental cues and processors to make decisions on what to do next.

Like real bees, RoboBees will work best when employed as swarms of thousands of individuals, coordinating their actions without relying on a single leader. The hive must be resilient enough so that the group can complete its objectives even if many bees fail.

Robert Wood is Charles River Professor of Engineering and Applied Sciences at Harvard University and the Wyss Institute for Biologically Inspired Engineering at Harvard. The National Science Foundation recognized his work with a 2012 Alan T. Waterman Award.



Radhika Nagpal is Fred Kavli Professor of Computer Science at Harvard and the Wyss Institute. Her work on collective behavior spans artificial intelligence, robotics and biology.



Gu-Yeon Wei is Gordon McKay Professor of Electrical Engineering and Computer Science at Harvard. His research interests span a broad range of topics in energy-efficient computing systems.



TECHNOLOGY

flight of the robobees

Thousands of robotic insects will take to the skies in pursuit of a shared goal

By Robert Wood, Radhika Nagpal and Gu-Yeon Wei

NOT TOO LONG AGO A MYSTERIOUS AFFLICTION CALLED colony collapse disorder (CCD) began to wipe out honeybee hives. These bees are responsible for most commercial pollination in the U.S., and their loss provoked fears that agriculture might begin to suffer as well. In 2009 the three of us, along with colleagues at Harvard University and Northeastern University, began to seriously consider what it would take to create a robotic bee colony. We wondered if mechanical bees could replicate not just an individual's behavior but the unique behavior that emerges out of interactions among thousands of bees. We have now created the first RoboBees—flying bee-size robots—and are working on methods to make thousands of them cooperate like a real hive.


Superficially, the task appears nearly impossible. Bees have been sculpted by millions of years of evolution into incredible flying machines. Their tiny bodies can fly for hours, maintain stability during wind gusts, seek out flowers and avoid predators. Try that with a nickel-size robot.

Now consider the hive. A bee colony appears to have no supervisor and no centralized authority. Yet colonies of tens of thousands of honeybees intelligently divide their labor to accomplish tasks critical for the health of the overall hive. When the hive requires more pollen, additional bees forage; when the hive requires tending, the bees stay home. And when something goes wrong—perhaps a queen dies unexpectedly—the bees rapidly adapt to the changing circumstances. How does such a large colony make these complex decisions—without taking forever or caus-

HOW IT WORKS

Micro Flight Plan


The RoboBee project aims to create an autonomous flying robot the size of an insect. RoboBees will be able to fly under their own power, navigate toward an objective, adapt to changing circumstances and work together as a group.




RoboBee at Actual Size

On this tiny frame, a RoboBee must carry its artificial muscles, optical-flow sensors, computing processors and a power supply to keep it all running. The total package weighs around half a gram.


1 Machining



2 Lamination



3 Release



ing havoc through miscommunication—if no one is in charge?

A robot hive could do much more than just pollinate flowers (although agriculture is one potential use). Indeed, small, agile, simple, inexpensive robots could perform many tasks more effectively than a few highly capable ones. For example, consider a rescue worker with a box full of 1,000 RoboBees—a package that would weigh less than a kilogram. The RoboBees could be released at the site of a natural disaster to search for the heat, sound or exhaled carbon dioxide signature of survivors. If only three of the robots accomplish their task while the others fail, this is a success for the swarm. The same cannot be said about the current generation of \$100,000 rescue robots.

Yet a colony of robotic bees imposes a huge number of tech-

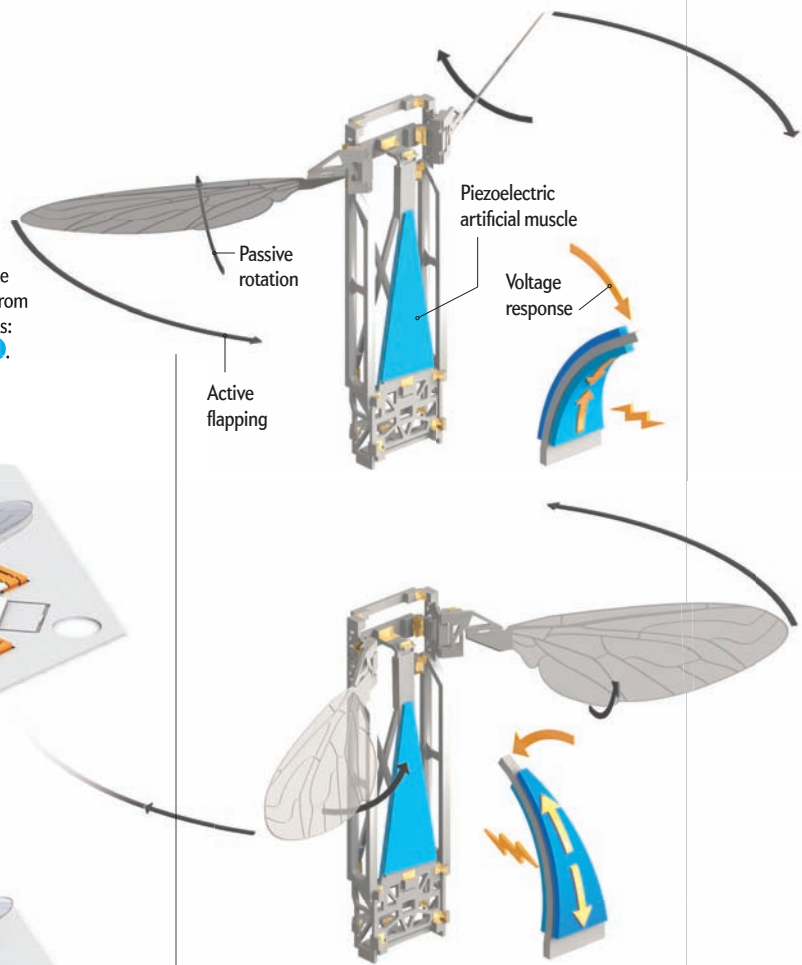

nological challenges. These tiny robots would stretch no more than a few centimeters from end to end and weigh around half a gram—about 100th the weight of the world’s lightest autonomous flying craft. That minuscule package must hold each bee’s flight system, its electronic brain and vision system, and the controls that govern how that bee interacts with other members of its hive. Recent progress in materials science, sensor technology and computing architecture are putting those goals in reach.

BODY AND FLIGHT

THE MOST OBVIOUS CHALLENGE in creating a small flying robot is figuring out a way to get it to fly. Unfortunately, the steady progress that has been made in miniaturizing robots over the past decade

Pop-up Assembly Line

Manufacturing such tiny robots carries its own challenges. The team builds each RoboBee from layers of hard material such as carbon fiber, sandwiching softer polymers. Gaps in the carbon fiber allow the polymers to flex, creating a flexure joint. The beauty of this layering method is that it lends itself to an efficient assembly line (below left). The researchers first make precise cutouts of the constituent layers using an ultraviolet laser ①. In the second step, they align all the layers of the multilayer sandwich and laminate them together with adhesive ②. They can then release individual components from the substrate ③. Finally, they borrow the final step from children's pop-up books: with one motion, they pop 3-D structures out of the two-dimensional surface ④.



The Big Flap

RoboBee flight relies on so-called artificial muscles—piezoelectric materials that contract when a voltage is applied. The wings can move in two ways—stroking back and forth and rotating their pitch. Instead of the up-and-down motion characteristic of bird flight, think of how you would tread water in a pool with your arms. Muscles control the flapping, but rotation is passive—determined by wing inertia, the interaction of the wing with the air, and the elasticity of the wing hinge.

is of little help to us because the small size of the RoboBee changes the nature of the forces at play. Surface forces such as friction begin to dominate over volume-related forces such as gravity and inertia. This scaling problem rules out most of the mechanical engineer's standard tool kit, including rotary bearings and gears and electromagnetic motors—components ubiquitous in larger robots but too inefficient for a RoboBee.

Instead of spinning motors and gears, we designed the RoboBee with an anatomy that closely mirrors an airborne insect—flapping wings powered by (in this case) artificial muscles [see box above]. Our muscle system uses separate “muscles” for power and control. Relatively large power actuators oscillate the wing-thorax mechanism to power the wing stroke while smaller


control actuators fine-tune wing motions to generate torque for control and maneuvering. Both these actuators work at the joint where the wing meets the body.

The artificial muscle is made of piezoelectric materials that contract when you apply a voltage across their thickness. Such actuators have a few drawbacks—for example, they require high voltage and are brittle—yet this is one case where the physics of scaling is on our side. The smaller these actuators are, the faster they want to move. And because the amount of work delivered per cycle (per unit mass) stays fairly constant, faster flapping leads to more power. In fact, these muscles generate an amount of power comparable to the muscles in insects of similar size.

Over the past few years we have experimented with dozens of

The Colony at Work

A colony of thousands of RoboBees will have to efficiently allocate tasks among individuals even when it lacks a full picture of the environment. In the scenario below, the hive has been assigned to find and pollinate fields of flowers. Individual RoboBees first explore different regions. As they return to the hive with information about where flowers are blooming, the new information affects where future workers will go. More robots will be allocated to areas with more work. The hive-based strategy allows bees to exhibit collective intelligence even if power constraints limit bee-to-bee communication.



different configurations of actuators and joints. One thing we look for in each of these designs is how easy they will be to build. The thousands of bees in a hive will have to be mass-produced.

The best designs we have come up with so far are created from a three-layer sandwich: hard face sheets form the top and bottom layers, and a thin polymer film rests in the middle. We create joints by carving material out of the top and bottom layers, leaving the middle-layer polymer to bend, thereby creating a flexure joint [see box on preceding page].

We have made great progress in building a bee-size robot, but we are still trying to figure out the best way to power it. To overcome the demanding energy requirements of flight at small scales, much of the bee's mass must be taken up by the main actuator and power unit (think "battery," although we are also exploring the possibility of using a solid-oxide micro fuel cell). The power question also proves to be something of a catch-22: a large power unit stores more energy but demands a larger propulsion system to handle the increased weight, which in turn requires an even bigger power source.

Although we cannot yet make a RoboBee fly under its own power, we have demonstrated a 100-milligram bee capable of producing enough thrust to take off (we kept it tethered to an external power source). The RoboBee was also able to stabilize itself using a combination of active and passive mechanisms. Given the state of the art in battery energy density and the efficiency of all the body components, our best estimate for flight time remains only a few tens of seconds. To increase flight time, we are working to minimize the mass and maximize the efficiency of each body component.

BRAIN AND NAVIGATION

POWER IS NOT THE ONLY THING keeping our RoboBee tethered. An onboard brain is another unsolved problem. A RoboBee in the wild will have to constantly take stock of its surroundings, decide on the best course of action and control its flight mechanisms. External electronics work as a makeshift solution in the laboratory, but a working RoboBee will require its own brain.

At a high level, the brain constitutes intelligence that is not only responsible for controlling an individual RoboBee but also for managing its interactions with other RoboBees in the colony. We set out to build the brain in layers—sensors to interpret the physical environment, an electronic nervous system that handles basic control functions and a programmable electronic cortex to make high-level decisions. As a first step, we sought to design a brain subsystem that enables autonomous flight. This challenge requires a tight control loop that encompasses sensors, signal processors and the movement of body parts.

To figure out what sensors to use and how to structure the brain circuitry, we once again looked to nature. Flies (and other fauna) use two broad types of sensors to make their way about the world. Proprioceptive sensors give a fly information about its internal states—how fast its wings are flapping, for example, or the charge left in the battery. Exteroceptive sensors provide information about the outside world.

Modern technology offers GPS, accelerometers and multiaxis gyroscopes, but such sensors are typically too heavy or consume too much power (or both) to be useful. Hence, we are investigating an electronic vision system that is similar to what natural bees have—one that analyzes "optical flow," the apparent motion of objects in the visual field of an image sensor. Imagine the view out the passenger window of a car: nearby objects appear to move quickly through your field of view while distant objects move slowly. A visual system that utilizes this information can create a detailed three-dimensional representation of its environment even if it is equipped with only a small, simple image sensor.

Yet the RoboBee brain must be powerful enough to process the stream of data coming out of its image sensors and make appropriate control decisions to drive body actuators. Here again, even advanced off-the-shelf components will not work for us. Consequently, we have been exploring a new class of computer architecture for the RoboBee brain that combines general-purpose computing with specialized circuits called hardware accelerators. Unlike general-purpose processors, the do-anything chips that run ordinary home computers, hardware accelerators are finely tuned circuit blocks that do only one thing but do it well. We would use hardware accelerators to make the fast, real-time computations required by the feedback control loop for stable flight while also staying within strict power budgets.

A major challenge has been to figure out what trade-offs we can get away with. For example, we would like to be able to use a high-resolution camera. High pixel counts, however, require larger image sensors and additional computational power to process the images. Where is the sweet spot?

To help answer these kinds of questions, we have developed a special test chamber. We mount a RoboBee body on a fixed multi-axis force and torque sensor and let it flap its wings in an attempt to fly. On the walls of the test chamber, we project images of the physical environment that the RoboBee would be flying through. In this way, we can explore how our prototype vision system,

brain and body work together to navigate through the world.

Flight control is just the beginning, of course. We also have parallel efforts that explore other types of sensors that will let RoboBees accomplish specific tasks—finding a person hidden in earthquake rubble, say.

Unfortunately, one capability that we do not foresee for our current bees is direct bee-to-bee communication—the power costs associated with wireless communications are just too great. Yet that does not mean that the bees will act alone.

COLONY AND COMMUNICATION

AN INDIVIDUAL ROBOBEE will be tiny and limited compared to the world we hope to operate it in, and power and weight budgets severely curtail the types of sensing and communications hardware any individual RoboBee can have. Thus, in addition to our research into the body and brain of a RoboBee, we must also figure out how to build a colony. As with honeybees, a RoboBee in isolation can accomplish little. But a hive? Group behavior will allow RoboBees to explore large areas, make sense of those areas by making many simple observations, efficiently divide labor, and thrive even when individual bees fail. Swarms of small, agile and potentially disposable robots can enable many new applications—pollination, for example, or search and rescue in a disaster scenario—that are not possible with individual robots.

Since the early 1990s computer scientists working in the research area of “swarm intelligence” have elucidated many powerful coordination algorithms inspired by social insects—from coordinated search strategies to smart division of labor. But even with these algorithms in hand, swarms of robots cannot be managed like a single robot.

For one, programming and reasoning at the level of individuals become untenable with thousands of entities. It would be like asking the average software developer to sit down and write out the instructions for each physical bit inside a computer. Instead, just as compilers take the human-readable instructions of a computer program and translate it into the 1's and 0's that control individual transistors inside a microprocessor, we need a higher-level, abstract way of programming the colony as a whole—one that would get translated from global instructions to individual behavioral programs. We need a programming language for colonies.

What is the right language that captures what honeybee colonies do and what we hope to do with RoboBee colonies? There is no simple answer yet, but we have developed two abstract languages to start with. In the Karma language, one can specify a flowchart of tasks that the colony must achieve. The flowchart contains links that represent conditions that trigger new tasks. The Karma system uses information that comes back from individuals to adjust the allocation of resources to tasks in a way that mimics the role of the hive in real honeybee colonies.

A different approach, called OptRAD (*Optimizing Reaction-Advection-Diffusion*), treats the colony of aerial robots as a fluid that diffuses through the environment. Any individual RoboBee uses a probabilistic algorithm to determine whether it will execute a task based on the current state of the environment. Treating the system as a fluid allows OptRAD to reason at a high level about the expected outcomes and adjust its behavior to adapt to new circumstances.

We also have a great deal to learn about building and operat-

ing a large colony of robots that contains not just tens or hundreds but thousands of autonomous robots that vastly outnumber their human operators. When there are thousands of entities, just operating robots at the level of individuals also becomes untenable. Imagine if every robot had an on/off switch—if it took five seconds to turn each one on, then turning on 1,000 robots would take nearly an hour and a half. Similar constraints apply to everything from cost to maintenance—every robot must be cheap, easy to make and simple to operate at the collective level. Ideally, every operation would be scalable—it would take some fixed amount of time that does not increase with the size of the collective (or at least that increases very slowly).

These challenges have motivated us to create the Kilobot system—a collective of hundreds of robots, each about the width of a quarter, that move by vibrating and that communicate with other nearby Kilobots. We can use this collective to test the efficacy of our programming languages and our mathematical models of emergent behavior. After all, without playing with real hardware, we are unlikely to understand the emergent behavior of physical systems.

The collective can be used to test many group behaviors that we would eventually want the RoboBee colony to achieve. For example, we can ask the group to search for targets in an environment, then, once an individual Kilobot finds a target, report the location back to the group. We have also made the Kilobot design open source for groups wanting to make their own. Or one can purchase premade Kilobots from K-Team, an educational robotics company. We hope such a standardized robotics kit will help generate new ideas and spur collective advances in science that individual groups cannot do—after all, we, too, rely on collective power to become more than the sum of our parts.

THE FUTURE

ALTHOUGH WE HAVE MADE a lot of progress, much work remains. We anticipate that within a few years we will have RoboBees flying under tightly controlled lab conditions. Within five to 10 years beyond that, you may see them in widespread use.

In 1989 renowned roboticist Rodney Brooks wrote a paper on the benefits of small robots for space exploration entitled “Fast, Cheap, and Out of Control: A Robot Invasion of the Solar System.” This is, of course, a play on the old engineer’s adage that consumer products can be typically characterized by any two—but not three—of the following adjectives: fast, cheap and reliable. With many individuals, the failure of one matters little.

Brooks was prophetic in his interpretation of this concept to robotics. Provided you can make many simple things that effectively work together, who cares if the individuals fail periodically? The only way to ensure the success of robotic explorers is to make it okay for them to occasionally fall out of the sky. ■

MORE TO EXPLORE

Kilobot: A Low Cost Scalable Robot System for Collective Behaviors. Michael Rubenstein, Christian Ahler and Radhika Nagpal in *2012 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3293–3298; May 14–18, 2012.

Progress on “Pico” Air Vehicles. R. J. Wood et al. in *International Journal of Robotics Research*, Vol. 31, No. 11, pages 1292–1302; September 2012.

Harvard Microrobotics Lab YouTube channel: www.youtube.com/MicroroboticsLab

SCIENTIFIC AMERICAN ONLINE

Watch RoboBees take flight at ScientificAmerican.com/mar2013/robobees