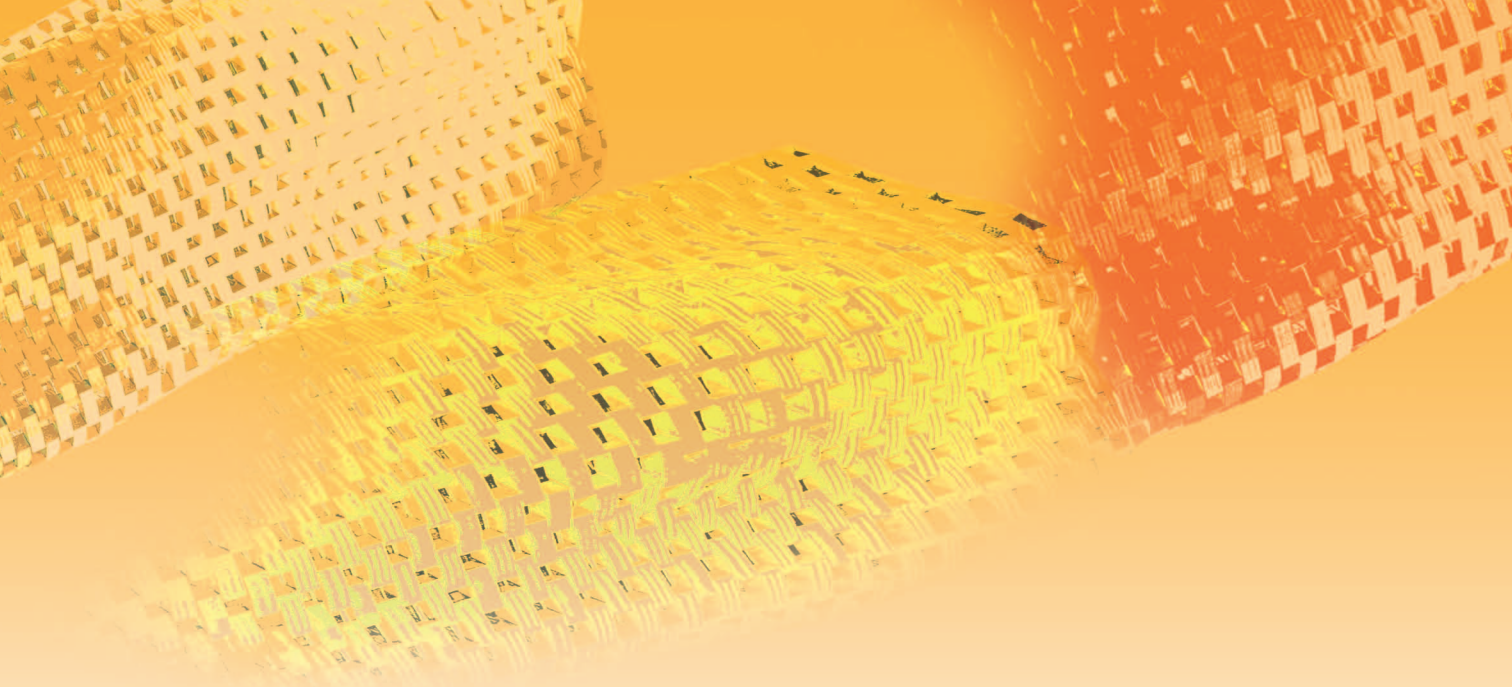

The Flight Assembled Architecture Installation

COOPERATIVE CONSTRUCTION WITH FLYING MACHINES

The art installation *Flight Assembled Architecture* [1] is one of the first structures built by flying vehicles. Culminating in a 6-m-tall tower composed of 1500 foam modules (see Figures 1 and 2), the installation was assembled by four quadrocopters in 18 hours during a four-day-long live exhibition at the Fonds Régional d'Art Contemporain (Regional Contemporary Art Fund) du Centre in Orléans, France. This article documents the design and development of specific elements of the autonomous system behind this one-of-a-kind installation and describes the process and challenges of bringing such a complex system out of the laboratory and into the public realm, where live demonstration and human-in-the-loop interaction demand high levels of robustness, dependability, and safety. The installation is a 1:100 scale model of what was originally conceived of as a 600 m-high vertical village (see "The Vertical Village" for details) and is an exploration of aerial construction in architecture. Architects have been exploring the use of digital technologies for the design and assembly of structures for some time now, and many facilities for investigating nonstandard architectural design and fabrication using industrial robots have sprung up in the past decade [2]–[4]. However, robot arms and computer numerical control (CNC) machines are limited by predefined working areas that constrain the size of the workpiece they can act upon and are thus also limited in their scale of action to a small portion or component of the overall structure, or to model-sized fabrication [5]. In contrast, flying machines are not constrained by such tight boundaries. The space that flying machines can act upon is substantially larger than the size of the machines themselves, making it feasible for the machines to work on the structure as a whole at a 1:1 scale, thus offering architects a new framework for realizing their designs.

While manned flying machines such as helicopters are commonly used to transport heavy objects to otherwise inaccessible locations, the use of autonomous unmanned aerial vehicles (UAVs) for construction tasks is still in its infancy. A first foray into autonomous UAV aerial construction was presented in [6], where quadrocopters were used to build cubic structures with the help of magnetic components. The ARCAS project focuses on aerial assembly via helicopters equipped with robotic arms [7]. The

Digital Object Identifier 10.1109/MCS.2014.2320359
Date of publication: 14 July 2014

1066-033X/14©2014IEEE

FEDERICO AUGUGLIARO, SERGEI LUPASHIN,
MICHAEL HAMER, CASON MALE, MARKUS HEHN,
MARK W. MUELLER, JAN SEBASTIAN WILLMANN,
FABIO GRAMAZIO, MATTHIAS KOHLER,
and RAFFAELLO D'ANDREA

FIGURE 1 The Flight Assembled Architecture installation. The 6-m-tall tower consisting of 1500 foam elements was assembled by four quadrocopters in France, 2011. (Photograph by François Lauginie.)

FIGURE 2 Two quadrocopters assembling the structure. During the building of the tower, each construction element was required to be at least 1.5 m from the previous element to ensure system safety and to reduce the aerodynamic interference between quadrocopters during placement. A construction module is 30-cm long. (Photograph by François Lauginie.)

prototypical assembly of tensile structures was demonstrated in [8]. Aerial manipulation research is currently addressing many of the open questions on the use of UAVs in scenarios where they must interact with their surroundings [9] and with each other to achieve a task: multiple quadrocopters cooperating to lift a payload are presented, among others, in [10]–[12]; various strategies for quadrocopters and helicopters grasping payloads are presented in [13] and [14]; and the need for aerial manipulation is also leading to the development of new concepts for flying machines such as the tritiltrotor [15] and the hex-rotor with tilted propellers [16].

Bringing aerial construction into an exhibition environment presents a host of additional challenges, many of which must also be addressed if aerial construction methods are to be implemented in practice. For example, modularity was an important design feature for the installation, which leverages a core platform of hardware and software components (such as quadrocopters, trajectory tracking controllers, vehicle state estimation, a motion capture system, or communication infrastructure; for details, see "The Flying Machine Arena" and [17]) with custom hardware and software designed for the specific task of gripping, transporting, and placing the 90-g polyurethane

The Vertical Village

Flight Assembled Architecture is not only one of the first structures built by flying machines but simultaneously represents a new architectural vision. Presented at the Fonds Régional d'Art Contemporain du Centre in Orléans, the installation addresses the next logical step of robotic fabrication and paves the way for entirely new scales of digitally fabricated architectures [25]. The installation is a model of a 600 m-tall urban structure (Figure S1) that, with 180 levels, has a total usable floor area of 1.3 million m^2, a *vertical village* [S1]. Composed of vertical core structures and horizontal module chains, the vertical village is notable on two counts: its porous arrangement not only creates living space for over 30,000 inhabitants with a great variety of programmatic and urban potential but it also enables a large degree of freedom for the spatial arrangement of the modules and their space-enclosing, self-stabilizing formation. It is not the absolute height that is decisive but rather the spatial order resulting from the structural composition. As such, the vertical village makes use of a grid-like organization. This organization does not run horizontally, as in a usual city grid, but is turned vertically and closed to form a circular entity. The village's geometry enables a varied urban program. Up to 25 individually positioned modules on each horizontal layer interact with each other. The areas in between vary and yet nevertheless form a homogeneous sequence of spaces. The modules are differentiated only internally, where they contain between one and three floors. The outer dimensions of the modules are, in contrast, unified. The modules are 30-m long, 12–15-m wide, and 10-m high. Whereas a module in architecture is traditionally defined in its function as a building component or a spatial unit, something else is apparent here: the module acquires a particular variability, freed internally from any specific functionality, and is thus versatile in its actual form while externally remaining unified and generically deployable.

FIGURE S1 The vertical village, a computer rendering of the 600-m-tall structure.

With such a network of interrelated modules, in-between spaces, and connections, the vertical village is formed by an intricate layering of private, semiprivate, and public space (see Figure S2). This design enables a decentralization that avoids not only the point-like restrictions of older urban planning and the gridlocked pathways of the modern city but also the confusing chaos that characterizes almost all unregulated urban

foam modules that compose the installation. The modularity also allows for the easy integration of charging stations that enable the system to run continuously for many hours and a navigation system that ensures collision-free trajectories for multiple vehicles. The ability to augment a core system with additional hardware and software modules that enable specialized tasks would be critical in a real-world building scenario, where the system may be required to interact with different designs, materials, and locations. Mobility is another important issue. A system that could first be tested in the lab and then transported and reassembled in the exhibition environment was required, and portability is important in real-world construction practice as well. Furthermore, in the installation, a human operator supplied the pickup station with the foam elements, triggering the quadrocopters to pick them up and carry them to the location indicated by the blueprint (see Figure 3). In addition, the audience was very close to the structure during the assembly process, and the construction space was not bounded by nets. Any time humans are in the loop (as is inevitable in a real-world construction scenario), a system must be designed with high degrees of responsiveness, robustness, and safety in mind.

This article first presents the system architecture behind the installation, explaining the various tasks performed by each component of the system, and how these interact. The system's various realization methods are then described, including strategies for accurate pickup and placement of the foam elements by the quadrocopters, as well as a navigation system for coordinating the flight of multiple vehicles. Next, the article describes the specialized components, such as the foam modules, grippers, and charging stations, that were specific to the installation. Finally, the development process and challenges presented by the live exhibition are addressed. Additional details on the hardware and software components are described in "The Flying Machine Arena" and in [1] and [17].

SYSTEM ARCHITECTURE

The autonomous system responsible for building the tower is divided into the four subsystems shown in Figure 4: the *blueprint*, which contains a list of sequential placement instructions; the *foreman*, which manages the overall

expansion today [S2]. In this sense, the question of the variety and accessibility of urban spaces and their contents becomes one of the central themes of the vertical village, where four giant continuous public double-rings (the darker-colored bands in Figure S1) with a combined length of 1 km are not located on the ground level (where public pathways are usually found) but rather spread out through the entire height of the structure, creating heterogeneous city structures [S3]. The public space thus extends over the entire height. Consequently, circulation in the vertical village can remain constrained to solely pedestrian access. Inhabitants have quick and direct access to all important functions such as schools, shops, public services, and leisure activities. As such, the vertical village offers a healthy and individual urban lifestyle characterized by short distances and a mixture of work and living; everything remains decentralized and freely accessible. Furthermore, the high-density architecture of the vertical village offers not only a high amenity value and capacity for adaptation but an enormous economic and ecological potential as well. This architecture integrates the entire construction morphology through to its detailed articulation.

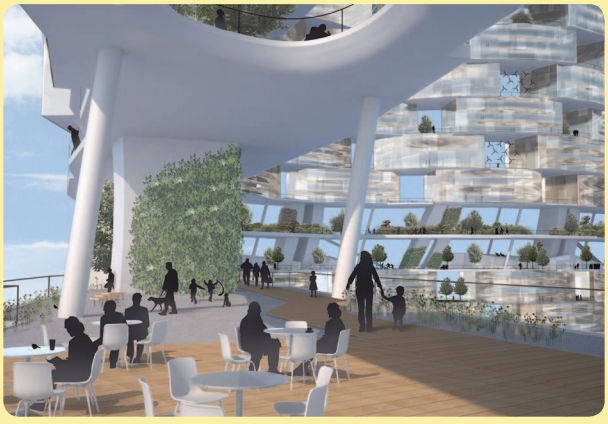

FIGURE S2 Inside the vertical village, with painted skywalks and an intricate layering of private, semiprivate, and public spaces.

The vertical village creates a new kind of urban vision in which robotic technology no longer appears as abstract but as a realistic means, not only of designing and building but also of arriving at novel architectural paradigms [S4], which allows the project to address new scales of digital fabrication and to conceptually expand the connection of architectural design and robotic technology. Here, a new urban vision can be experienced, becomes tangible, and expresses a radically new way of thinking about and materializing architecture [S5].

REFERENCES

[S1] F. Gramazio, M. Kohler, and J. Willmann, "The vertical village," in *Flight Assembled Architecture*, F. Gramazio, M. Kohler, and R. D'Andrea, Eds. Orléans, France: Editions Hyx, 2013, pp. 13–23.

[S2] B. Lootsma, "Deröffentliche raum in bewegung," *daidalos*, vol. 67, pp. 116–123, 1997.

[S3] D. Saunders, *Arrival City. How the Largest Migration in History Is Reshaping Our World*. New York: Vintage Books, 2012.

[S4] J. Willmann, S. Langenberg, F. Gramazio, and M. Kohler, "Digital by material: Envisioning an extended performative materiality in the digital age of architecture," in *Robotic Fabrication in Architecture, Art, and Design*, S. Brell-Cokcan and J. Braumann, Eds. Vienna, Austria: Springer, 2012, pp. 12–27.

[S5] A. Tönnesmann, *Die Kunst der Renaissance*. Munich: C. H. Beck, 2007.

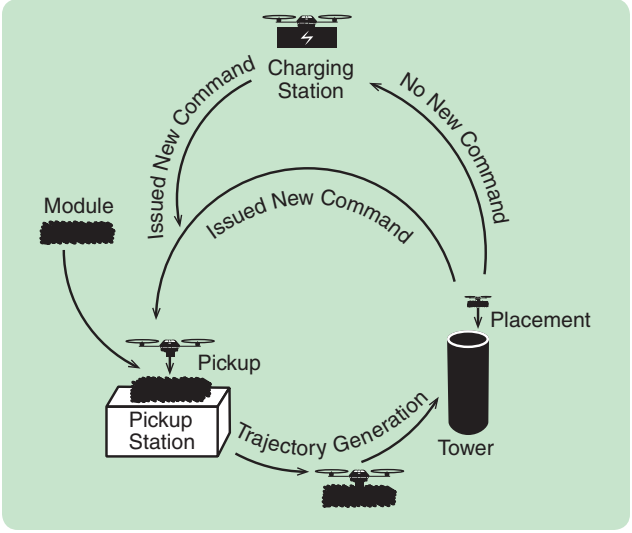

FIGURE 3 The pick-place state machine. The process of assembling a structure begins when a module is placed in the pickup station, which triggers a pick-place command to be issued to an idle quadrocopter, which then removes the module from the pickup station and places it at the desired location within the structure.

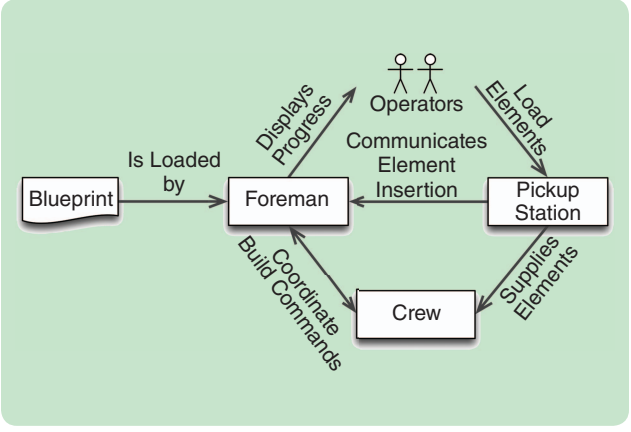

FIGURE 4 The system architecture. A block diagram shows the high-level organization and interaction of the system's components. The blueprint contains a list of sequential placement instructions. The foreman manages the overall construction process by interpreting the blueprint, issuing build orders, and tracking the construction progress. The crew is responsible for executing the build orders to fabricate the structure. The pickup station provides building elements to the crew. Operators manually put construction elements in the pickup station.

The Flying Machine Arena

The Flight Assembled Architecture project was built upon the ETH Flying Machine Arena (FMA) platform. The FMA is a research and demonstration platform for fleets of small quadrocopters that has been in development at ETH Zurich since 2008. In typical use, the FMA consists of a commercial motion capture system, a fleet of customized vehicles (based on the ascending technologies hummingbird platform described in [26]), specialized wireless and wired communication channels, and a library of building blocks and tools to create and run experiments in the system.

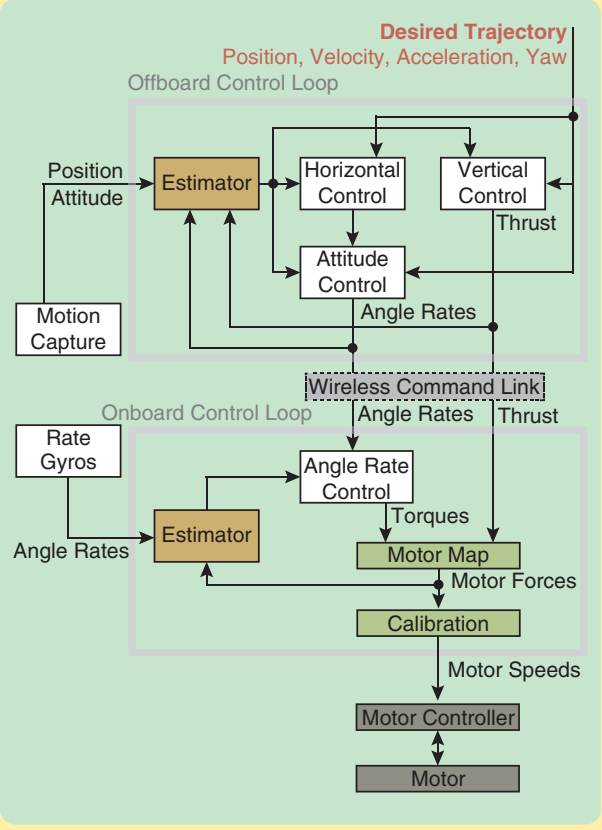

FIGURE S3 An overview of the control structure used in the flying machine arena. Some details are left out for clarity; see [17] for a more detailed description. The control strategy consists of two loops. The first control loop runs on standard computers at 50 Hz. Its inputs are the motion capture system measurements (position and attitude) and a desired trajectory (position, velocity, acceleration, and yaw). The control loop consists of an estimator and cascaded controllers. A vehicle command consisting of desired angle rates and collective thrust is generated and sent wirelessly to the vehicle. The second control loop runs on an onboard microprocessor at 800 Hz. Using the onboard rate gyroscopes, the quadrocopter tracks the received commands by controlling off-the-shelf motor controllers. The controllers are designed for tuning intuition and modularity. Each part of the control structure may be used separately, has separate calibration parameters/routines, and may be replaced on demand; for example, a more sophisticated controller may be used for attitude control instead of the in-use linear axis-separated bank-angle controller.

LOCALIZATION AND STATE ESTIMATION

The FMA uses an overhead motion capture system [19] to track the positions of marked objects in the space. For the installation, a 19-camera Vicon T-40 system was used to provide high-accuracy position and attitude information at 200 Hz for all quadrocopters in the space. The quadrocopters, pickup stations, charging stations, and the placement platform are marked using retroreflective tape. Static objects, such as the charging stations, may be calibrated once and not tracked continuously.

A predictor-corrector estimator fuses this data together with recent commands and a first-principles model of the vehicle dynamics to produce a current latency-compensated estimate of the state of each vehicle, including its current position, velocity, attitude, and rotational rates. As the dynamics model is accurate for short time durations and the total communication latency in the FMA is low (on the order of 30 ms, as detailed in [17]), a model-based prediction provides a straightforward way to improve overall system performance at a low computational cost. In a similar fashion, brief losses of position and attitude information are compensated by predicting forward the latest valid estimate based on the commands sent to the vehicle. Special care is taken to use unpredicted data in instances where the model may not be accurate, such as during module pickup or placement; during these operations, the dynamics of the flying vehicles are dominated by external contact forces due to the interaction with the environment.

CONTROL STRATEGY

An overview of the FMA control strategy is depicted in Figure S3. The strategy is composed of a cascade of controllers, where the controllers are designed with modularity and abstractability in mind. For example, from the standpoint of the position control cascade, the underlying vehicle dynamics are considered to be an ideal second-order system, which can be shown to be a reasonable assumption for appropriate tuning of the underlying control loops [17].

Similarly, calibration parameters and corresponding calibration routines are built into the various levels of the control architecture to enable automatic compensation for static nonidealities. For instance, a hover calibration step uses constraints implied by a hover vehicle (such as the balance of torques, alignment of the collective thrust vector with gravity, and other equalities) to automatically adjust compensation factors such as those for individual rotor efficiencies, overall vehicle motion capture attitude misalignment, and other factors. This calibration scheme instantly improves the performance of the system, even under severe nonidealities such as when carrying construction elements in various configurations or when adjusting for propeller wear after long-term, high-stress operation.

ROBUSTNESS FEATURES

A special software module called *Copilot* is used to help manage the vehicles, track persistent state information such as battery levels, execute common maneuvers such as takeoff and landing, manage the charge cycle, and provide a robust fallback controller for implementing emergency system-stop behavior. The structure of Copilot is further described in [17] but can be summarized as a separate, fully functional estimator and controller module, capable of safely flying the fleet of vehicles. Copilot also provides an emergency stop feature, where a physical push button may be pressed at any time to completely disable all vehicles, the last resort to shutting down the entire system in an emergency. Another robustness feature, detailed in [24] and implemented in Copilot, provides a safety blind hover behavior for each flying vehicle in case of motion capture failure, radio link failure, or other global feedback control loop failure. Each vehicle keeps an onboard estimate of its current attitude and velocity; in case of system failure, the vehicle uses this estimate to attempt to reach hover and descend in a controlled fashion.

construction process by interpreting the blueprint, issuing build orders, and tracking the construction progress; the *crew* system, which is responsible for executing the foreman-issued build orders to fabricate the structure; and the *pickup station*, which provides building elements to the crew. These processes run on an external computer.

Blueprint

The blueprint is a plain-text file containing a list of placement instructions, sequenced by placement order. A placement instruction consists of the position and orientation of each element in tower-relative coordinates, with vertical position given relative to the tower floor rather than as an absolute position. This design allows the exact vertical position of each element to be calculated at runtime based on the actual positions of the supporting modules, thus compensating for cumulative errors such as the unknown and variable thickness of the joining material (glue), which is manually applied to the bottom of the elements before being supplied to the construction system.

Static Stability and Placement Error Tolerance

From a static perspective, structural stability of a single element requires that its center of mass be within the convex hull generated by the contact surfaces between it and its supporting elements [18]. Due to inherent inaccuracies in the quadrocopter placement routine, both the shape of the supporting convex hull and the relative location of the placed module's center of mass will vary. To ensure robustness against placement errors, a stability analysis is performed on the blueprint before building the structure, taking expected placement errors into account. The stability analysis consists of generating the convex hull described above and verifying that the center of mass of the module being placed lies within it. The analysis is then repeated to account for placement errors, that is, the fact that the supporting modules and the newly placed module are shifted and rotated in different directions by the expected placement error value. Areas within the tower that are identified as being unstable are then redesigned. This stability analysis assumes that the structure is rigid and stable, which is a reasonable assumption due to the adhesive bonding between placed elements. This analysis does not take into account the adhesive bonding between the to-be-placed element and its supporting elements. In reality, the system can handle placements deemed unstable by the stability analysis due to adhesive bonding.

Placement Order Precomputation

The order of placement instructions (herein the *build order*) is precomputed to allow architectural control over the build; for example, faces of the tower were built at different rates to give the audience an ever-changing perspective of the tower. Precomputing the build order also allows additional constraints to be considered during the design of the structure. For example, during the assembly of the tower, each construction element was required to be at least 1.5 m away from the previous element to ensure system safety and to reduce the aerodynamic interference between quadrocopters during placement. Using this safety distance as a constraint, the build order was designed such that at least two quadrocopters could operate simultaneously.

Foreman

The foreman serves two functions. First, it is the graphical interface to the system, through which operators can start and stop the construction process, limit the maximum number of vehicles in flight at any given time, and monitor the build progress and subsystems. Second, the foreman listens for state feedback from both the crew and the pickup station and uses this information to coordinate the build at a high level.

Build management includes tasks such as sending the crew a new placement instruction whenever a construction element is inserted into the pickup station, responding to successful or failed module placements, and logging the build for real-time or post-construction analysis. Abstracting the task of high-level construction management from the crew system into the foreman allows the crew to focus on the execution of individual placement commands, while reducing subsystem coupling, and thus improving robustness against the failure of an individual subsystem.

Crew

The crew consists of a fleet of quadrocopters controlled by a centralized software tool. The tool communicates with the foreman, delegates tasks to members of the fleet, and

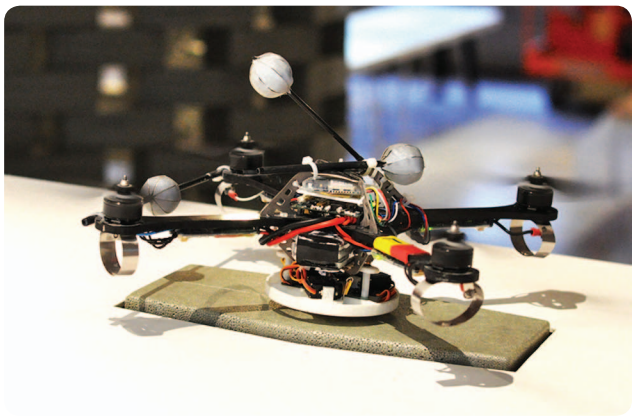

FIGURE 5 A quadrocopter picking up a construction module. To pick up the modules, vehicles must approach the construction elements from above and land precisely in their center. The modules are placed inside the pickup station by the human operator.

controls each of the quadrocopters using existing Flying Machine Arena (FMA) components, such as the state estimator, and the trajectory is loaded by tracking controller. Because of the requirement that at least two quadrocopters could operate simultaneously, the size of the fleet was set to four vehicles, which allowed two vehicles to recharge their batteries while the other two were performing the construction task. In principle, however, the crew is capable of handling a larger fleet.

The crew is responsible for the full-stack management of the quadrocopter fleet. At a low level, the crew receives the position and attitude of each quadrocopter from a motion capture system [19], runs estimation and control algorithms, and sends commands to the vehicles at 50 Hz, as discussed in "The Flying Machine Arena."

At a high level, the crew organizes the tasks of the fleet based on the battery level and state of each quadrocopter, the desired number of in-flight quadrocopters as set by the user, and based on the current placement instructions from the foreman. While the quadrocopters are in flight, the crew uses a space reservation system to ensure that the vehicles do not collide with fixed infrastructure, with the structure being built, or with each other. High-level crew management involves delegating commands to individual quadrocopters. These commands are issued by the foreman, interpreted by the crew controller, and subsequently allocated by the crew to an available quadrocopter. A build command consists of an instruction to fetch a construction element from a given pickup station and deliver it to a position in three-dimensional space.

Furthermore, the crew is responsible for reporting the state of each quadrocopter to the foreman. The feedback sent by each quadrocopter includes both the current action (for example, collecting an element from the pickup station) and confirmation of the previously completed action (for example, the placement of an element at a given location).

Pickup Station

The pickup station is the intermediate physical interface used by the operators to provide construction elements to the robotic crew. This interface allows the operators to maintain a safe distance from the autonomous operation and simplifies the crew's task of collecting construction elements. Second, the pickup station provides the operators with system feedback through a series of light-emitting diodes (LEDs). To enable simultaneous operation by two quadrocopters, two pickup stations were used during the assembly of the tower.

To initiate a cycle in the construction process, an operator manually inserts a construction element into the pickup station. If the construction element is laid flat and correctly aligned within the pickup station, its insertion is detected as successful, and the user is notified by a colored LED. Successful insertions are also communicated to the foreman, which then assigns a pick-place task to the crew. This process leverages the pickup, trajectory planning, and placement strategies documented in the next section.

REALIZATION

The successful insertion of a construction element into the pickup station triggers the construction process. First, the foreman is notified that a building element has been successfully inserted into the pickup station. The foreman then draws a placement instruction from the blueprint and delivers it to the crew subsystem. Then the centralized crew subsystem issues this instruction to an idle quadrocopter with sufficient battery power, giving preference to already in-flight vehicles. The selected quadrocopter collects the construction element from the pickup station and places it at the desired location and orientation within the tower. Once the quadrocopter has placed the element, it is free to be allocated new tasks. The pick-place state machine is shown in Figure 3. Safe navigation is achieved by means of a centralized reservation system.

The next sections present solutions that are adopted to precisely pick up and lay down construction elements. The strategy employed for safely flying multiple robots within a predesigned space is also discussed.

Picking Up an Element

Payloads carried by quadrocopters and other flying vehicles are often transported underneath the machine. The strategy used for the project is no different; the construction modules are carried by means of a gripper attached to the bottom of the quadrocopter. To pick up the modules, a vehicle must approach the construction elements from above, as shown in Figure 5. The gripper requires that the machine lands on the module at the desired gripping point before the gripper is closed. The gripper design does not include guides to assist in element positioning; thus to accurately pick up foam elements, the quadrocopter must be able to land precisely in the center of the flat-surfaced

element. Due to aerodynamic effects, such as the ground effect, this task is nontrivial. The following paragraphs present a strategy for performing accurate pickups.

To precisely land on a flat surface, some goals must be met. First, the attitude of the vehicle must be normal to the surface. Second, the lateral position error must be as small as possible. Third, good altitude tracking is required during descent to ensure that contact with the ground is made at the expected instant and at speeds low enough to prevent the vehicle from bouncing. Finally, the lateral velocity must be as small as possible during the final landing phase to prevent the vehicle from sliding after making contact with the ground. These requirements suggest that a vertical descent trajectory is preferable. The requirements are addressed by the landing strategy depicted in Figure 6 and explained below.

The quadrocopter begins its landing maneuver from a hover position above the target landing spot and outside of ground effect conditions, which allows position and heading offsets to be compensated for during the descent phase. These offsets are accentuated by any asymmetry in the vehicle configuration, such as weight distribution or propeller efficiency. A vertical trajectory that leads the vehicle to land on the construction module is planned using the trajectory generator described in the next section. Due to the intrinsic differences between vehicles, adaptation of the reference trajectory is needed to fine-tune the landing. During the descent, integral control is used along the lateral direction to compensate for position offsets. Integral action is also applied to correct the quadrocopter's heading. When the quadrocopter completes the descending trajectory, it is commanded to hover above the module, and an altitude integral controller is turned on to compensate for any residual altitude errors. At this stage, the state of the quadrocopter is constantly checked, and, when the position error, heading error, and lateral speed of the vehicle are all smaller than the given thresholds, the vehicle reduces the collective thrust below gravity to establish and maintain contact with the module, while still being able to control its rotational body rates to zero. The whole maneuver is repeated if these conditions are not satisfied after several seconds. Once the quadrocopter has landed on the module, its position is measured, and if the landing is precise enough, the gripper is closed. After a successful pickup, the extra thrust required to hover with the additional payload is estimated and taken into account during flight. Figure 7 shows the landing errors during the building of the tower, recorded by the motion capture system after the gripper has been closed. Landing errors result in an off-center pickup of the construction element. This offset is recorded in the system and compensated for during the placement; however, off-center modules might negatively affect the placement maneuver due, for example, to asymmetric weight distribution. Therefore, the target landing area was constrained to a circle of 1-cm radius, and the landing maneuver was repeated if the vehicle was not able to land inside the target area. The plot only shows successful landing attempts. Repetition of the landing maneuver occurred 3% of the time.

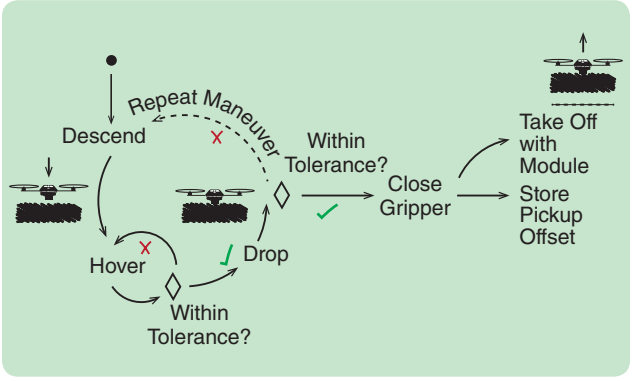

FIGURE 6 The pickup strategy. The quadrocopter begins its landing maneuver above the construction element. A vertical trajectory that leads the vehicle to land on the module is planned. The vehicle is then commanded to hover above the module. When the quadrocopter is within given tolerances, it reduces the collective thrust below gravity to establish and maintain contact with the module, while still being able to control its rotational body rates to zero. Its position is measured, and, if the landing is precise enough, the gripper is closed. A landing error results in a pickup offset, which is recorded in the system and compensated for during the placement.

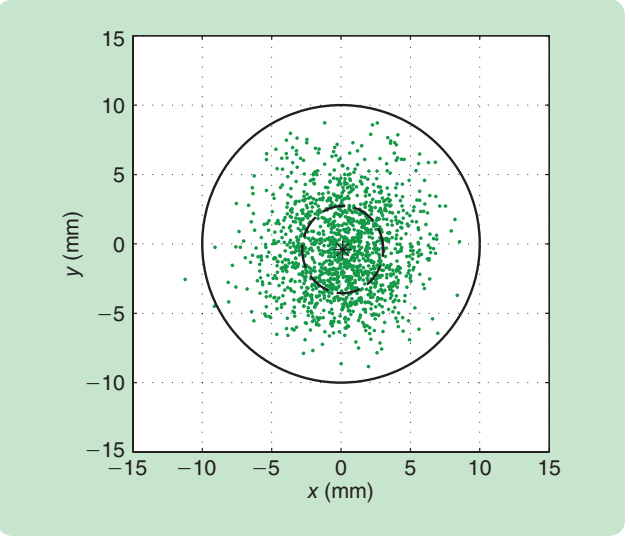

FIGURE 7 Pickup accuracy. The plot shows the pickup errors during the building of the tower, recorded after the gripper has been closed. The solid circle denotes the target landing area. Also shown are the mean (0.1 mm in the x-direction, −0.4 mm along the y-direction; indicated by the star) and the standard deviation (2.9 mm and 3.1 mm, respectively, indicated by the dashed ellipse). The plot only displays successful landing attempts: the landing maneuver is repeated if the vehicle is not able to land inside the target area. Repetition of the landing maneuver occurred 3% of the time. Landing success is measured before the gripper is closed. Closing the gripper can result in translation of the quadrocopter's position on the element, potentially moving it outside the defined landing tolerance, as indicated by the points laying outside the circle.

Trajectory Planning

Trajectory planning is crucial for performing a construction task with multiple machines in a coordinated fashion. The trajectory planner used in this project consists of three different subsystems that, together, guarantee safe trajectories. First, the space is laid out and allowable fly regions are defined. This information is used in combination with a space reservation system that allows vehicles to reserve space before flying through it. Second, waypoint-based navigation coupled with the space reservation system enables the discretization of the flyable space. Last, a trajectory planning algorithm is used to generate feasible trajectories from any initial state (given by heading, position, velocity, and acceleration) to rest (or hover, a state with zero velocity and zero acceleration), allowing quadrocopters to move between waypoints. Below, the three subsystems are presented in detail.

Freeway-Based Flight and Space Reservation System

The flight paths of the machines are controlled by a centralized space reservation system inspired by [20] and similar to the technique used by Kiva Systems [21], [22], whereby each vehicle places a request to reserve the space required for a trajectory before the trajectory is flown. The space reservation system stores all the current active space reservations and verifies if the request can be allowed. The vehicle releases the reservation as soon as it completes the trajectory. This system ensures that, while a space is reserved, only the reserving flying vehicle has access; all other vehicles must wait for the reservation to be released before flying through this space. This method guarantees collision-free navigation and is robust to communication delays [20], provided that a vehicle is able to stay within its reserved space, which is ensured through the generation of trajectories that satisfy the control inputs constraints and end at rest within the reserved space (as discussed later and in "Trajectory Generation").

Space reservation systems are prone to deadlocks, which occur, for example, when two vehicles want to swap their position by flying a straight line. None of the vehicles are able to reserve this space because it contains the current position of the other machine. This situation can last indefinitely, causing the vehicles to enter a deadlock. Deadlock situations can be solved with replanning; however, it is difficult to guarantee that the algorithm will eventually find a

Trajectory Generation

The flight paths connecting the charging stations, the pickup station, and the area above the structure where the construction elements are being placed were generated in real time. The trajectory generation used in the Flight Assembled Architecture project is based on the algorithm described in detail in [23]. An overview of the approach is given here.

For trajectory generation, the dynamics of the quadrocopter are modeled as a rigid body with a mass-normalized thrust input a and rotational body rate control inputs $\omega_x, \omega_y, \omega_z$:

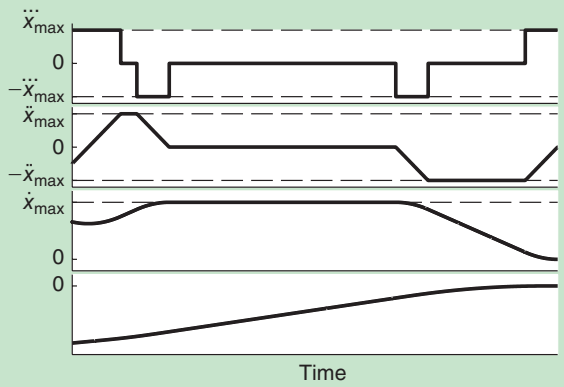

FIGURE S4 The structure of a sample trajectory. From top to bottom, the plots show the jerk, acceleration, velocity, and position profiles of a trajectory from an initial state (with nonzero velocity and acceleration) to the origin (at rest). The trajectory has a bang-singular structure and satisfies the dynamic constraints (S4)–(S6), indicated by dashed lines.

$$\dot{R} = R\begin{bmatrix} 0 & -\omega_z & \omega_y \\ \omega_z & 0 & -\omega_x \\ -\omega_y & \omega_x & 0 \end{bmatrix}, \quad \begin{bmatrix} \ddot{x} \\ \ddot{y} \\ \ddot{z} \end{bmatrix} = R\begin{bmatrix} 0 \\ 0 \\ a \end{bmatrix} - \begin{bmatrix} 0 \\ 0 \\ g \end{bmatrix}, \tag{S1}$$

where R denotes the rotation matrix representing the vehicle attitude, g is gravitational acceleration, and (x, y, z) represents the position of the quadrocopter. The control inputs are limited to be

$$a_{\min} \le a \le a_{\max}, \qquad |\omega_i| \le \omega_{\max} \text{ for } i = x, y, z. \tag{S2}$$

Model (S1) is a simplification of the true vehicle dynamics by 1) ignoring underlying dynamics (such as those associated with body rates and propeller speeds) because they are controlled by high-bandwidth control loops onboard the vehicle and 2) neglecting aerodynamic effects (such as drag acting on the vehicle) because the vehicle speed will be limited in the trajectory design, and these effects are thus not dominant.

The trajectory generation approach exploits the differential flatness of the quadrocopter dynamics to plan trajectories in the three translational degrees of freedom (DOFs) of the vehicle and by approximating the dynamics as triple integrators in each DOF:

$$\dddot{x} = u_x, \qquad \dddot{y} = u_y, \qquad \dddot{z} = u_z. \tag{S3}$$

The true control inputs a, ω_x, ω_y can then be recovered from the trajectories $x(t), y(t), z(t)$. Using the vector

$$f := \begin{bmatrix} \ddot{x} \\ \ddot{y} \\ \ddot{z} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ g \end{bmatrix}$$

suitable trajectory. An alternative solution is to adequately plan allowed paths.

To coordinate flying, the structure is encircled by two freeways that run at different heights. Downwash effect is reduced by having vehicles on the upper freeway travel in the opposite direction to vehicles on the lower freeway. The freeways are used to travel between the pickup station, the area above the structure where the construction elements are being placed, and the charging stations. These locations are physically separated and can only be accessed by one machine at a time, thus avoiding possible deadlocks. Figure 8 shows a visualization of the space reservation system in the three-dimensional environment during the actual build of the tower.

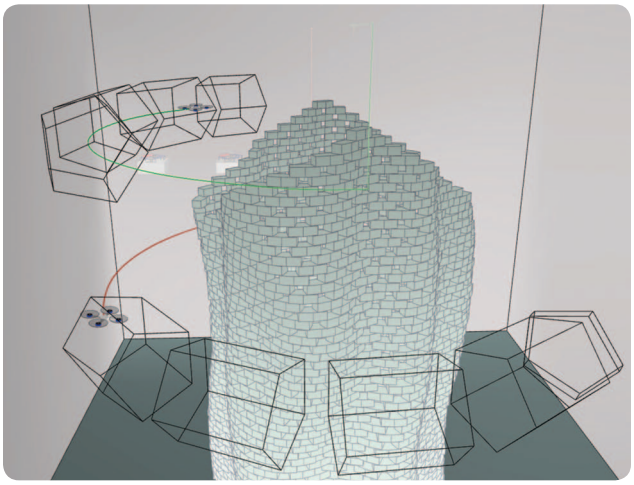

FIGURE 8 The space reservation system. The visualization of the space reservation system during the actual building of the tower in the three-dimensional environment. The two flying vehicles reserve space before flying in. The reserved space consists of cylinders with spheres at both ends and is visualized by the black prisms. In the picture, two vehicles are flying on the freeways encircling the structure after having placed a module. The red and green trails depict the flown trajectories.

Waypoint-Based Navigation

Once the allowed paths have been defined, trajectories along those paths must be generated. The trajectory design is strictly coupled with the space reservation system. For instance, when traveling from A to B, it is not convenient to reserve the space for the entire trajectory, as this would prevent other vehicles from using the space for the duration of the trajectory. Instead, segmenting the trajectory allows the

to denote the total force required to follow the trajectory, the control inputs are

$$a = \|f\|,$$

$$\begin{bmatrix} \omega_y \\ -\omega_x \\ 0 \end{bmatrix} = R^{\mathsf{T}} \left(\frac{\dot{f}}{\|f\|} - \frac{f f^{\mathsf{T}} \dot{f}}{\|f\|^3} \right).$$

The control input ω_z is not determined from the trajectories in the translational DOFs and may be determined separately. When the vehicle carries a module, ω_z is set to a constant rate for the duration of the trajectory, such that the module is rotated from the pick-up orientation to its placement orientation, as determined by the blueprint.

To satisfy the control input constraints (S2), the triple integrators (S3) are constrained in jerk and acceleration by approximating the constraints imposed by (S2) on trajectories such that feasibility remains guaranteed (as shown in [23]):

$$|\ddot{x}| \le \ddot{x}_{\max}, \quad |\ddot{y}| \le \ddot{y}_{\max}, \quad |\ddot{z}| \le \ddot{z}_{\max}, \tag{S4}$$

$$|u_x| \le \dddot{x}_{\max}, \quad |u_y| \le \dddot{y}_{\max}, \quad |u_z| \le \dddot{z}_{\max}. \tag{S5}$$

As noted above, the commonly used first-principles model of quadrocopter dynamics contains no drag term, and thus the trajectory generation algorithm described in [23] does not consider velocity constraints. However, because safety is a specific requirement of the Flight Assembled Architecture installation and to limit the influence of aerodynamic effects, it is important to have the option of limiting the maximum achievable velocity. Therefore, the trajectory generation problem was extended by a maximum allowable velocity in each DOF. The velocity in each axis is limited by

$$|\dot{x}| \le \dot{x}_{\max}, \quad |\dot{y}| \le \dot{y}_{\max}, \quad |\dot{z}| \le \dot{z}_{\max}. \tag{S6}$$

The problem—given by the dynamics (S3), the input constraints (S5), and the state constraints (S4) and (S6)—is entirely decoupled for the three DOFs. For each of the three DOFs, the time-optimal trajectory from the initial state to the final position is then computed.

Through the application of Pontryagin's minimum principle, it is straightforward to show that the time-optimal trajectory is of bang-singular structure (that is, the jerk u of each of the three axes is always minimal, maximal, or zero; see, for example, [27]), and the corresponding switching times can be computed through a bisection search. A sample trajectory for a single DOF is depicted in Figure S4. The computation of decoupled trajectories in the three DOFs using this method is on the order of tens of microseconds on a desktop computer and is therefore sufficiently low for recomputing trajectories to new waypoints.

Because of the round shape of the structure and the limited available flight space around it, it is important for the vehicles to accurately fly on the circular freeways when flying to or from a construction element placement point. To generate trajectories that follow the circular path, the trajectory generation is carried out in cylindrical coordinates when the planned flight path connects two points on a freeway. In this case, the three triple integrators (S3) are taken to represent the three cylindrical coordinates, and the constraints (S4)–(S6) are also formulated in the changed coordinates.

vehicle to reserve only a portion of the required space. During the execution of a segment, the vehicle tries to reserve the next segment. If the reservation is successful, the vehicle continues its motion without stopping. If not, the vehicle stops at the end of the segment, reaching a safe hover-state within the current reserved space. This behavior is achieved by means of waypoint-based navigation, as conceptually illustrated in Figure 9 and described in Algorithm 1. Waypoints are defined by a three-dimensional goal position in space and the vehicle's desired heading. Furthermore, a threshold in the form of a sphere can be specified; when the vehicle reaches the sphere, the vehicle will plan a trajectory that brings it from the current state to the next waypoint provided that the space required for the next trajectory has not yet been reserved by other vehicles. If instead the space has been reserved, the vehicle finishes the trajectory by coming to a rest, where it will safely hover until the required space becomes available.

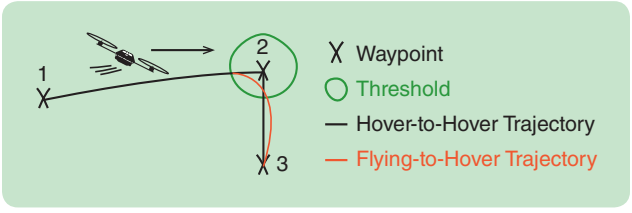

FIGURE 9 Waypoint-based navigation. Waypoints are defined by a three-dimensional hover position in space. During the execution of a segment between two waypoints, the vehicle tries to reserve the next segment: if the reservation is successful, when the vehicle crosses a predefined threshold, it directly continues its motion without stopping. If not, the vehicle stops at the end of the segment, reaching a safe hover state within the current reserved space.

ALGORITHM 1 Waypoint-Based Trajectory Navigation.

1: Waypoints: $W_0, W_1, \ldots, W_N$
2: Thresholds: $R_0, R_1, \ldots, R_N$
3: $i = 0$
4: **while** $i < N$ **do**
5: $T_i \leftarrow$ *trajectoryGeneration* (W_i, W_{i+1})
6: Submit space reservation for T_i
7: Hover at W_i **until** T_i is accepted
8: $T_{i+1} \leftarrow$ *trajectoryGeneration* $(W_{i+1} - R_{i+1}, W_{i+2})$
9: Submit space reservation for T_{i+1}
10: **while** W_{i+1} is not reached **do**
11: Fly trajectory T_i
12: **if** vehicle is at $W_{i+1} - R_{i+1}$ AND T_{i+1} is accepted **then**
13: i++, go to 8.
14: i++, go to 5.

Trajectory Generation

The waypoint navigation system relies on a trajectory generation algorithm to compute interwaypoint flight paths that satisfy the dynamic and input constraints of the vehicle. The trajectory generator accepts an initial state of the vehicle and computes a dynamically feasible trajectory to a given waypoint to be reached at rest. This trajectory is then given to the space reservation system so that it can be reserved. Due to the strategy of planning the trajectory to the next waypoint as soon as the vehicle is within a sphere from the current waypoint, it is necessary to plan from arbitrary initial states to rest, and not just from rest to rest. The trajectory generation algorithm used herein is based on [23]. An overview of the approach can be found in "Trajectory Generation."

The trajectory generation algorithm is also used for the pickup and placement of construction elements. The pickup consists of a set of waypoints guiding the vehicle to fly above the pickup station and then to the module pickup position. The placement task consists of a waypoint navigation to the hover position above the module. To then place the module, a waypoint below the actual placement is used to generate a vertical trajectory that respects the dynamic constraints and, by exploiting the structure of the maneuver, crosses the placement point at a desired velocity. This

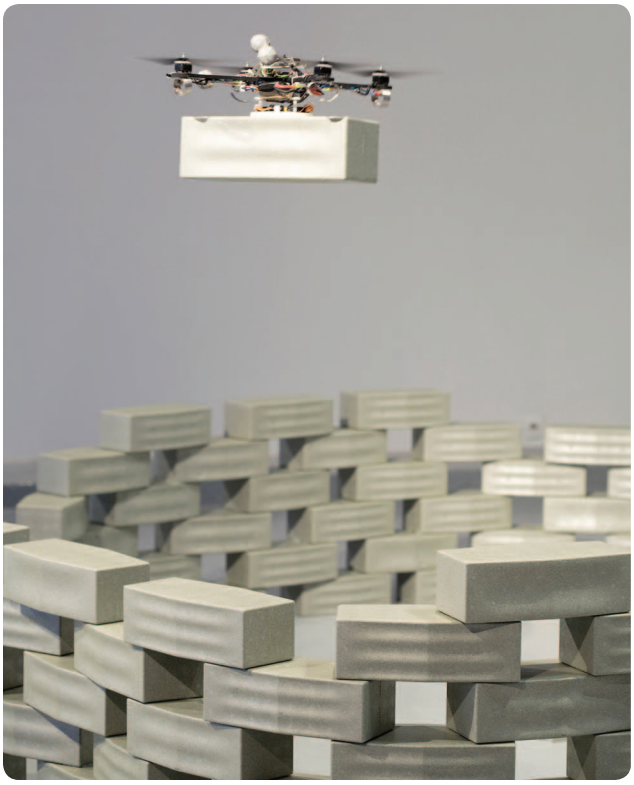

FIGURE 10 Placing a module. The vehicle hovers above the structure before placing a construction module. A vertical trajectory that results in the foam element impacting the structure at 1 m/s is executed to place the module. The actual mean impact velocity was 1.04 m/s with a standard deviation of 0.02 m/s. (Photograph by François Lauginie.)

procedure is explained in the next section and in "Trajectory Generation."

Placing an Element

The modular structure is assembled in a bottom-up manner: new elements are placed on top of already-placed elements by flying machines that descend vertically to the desired spot. The comparison between different strategies, many iterations, and fine-tuning resulted in an accurate and reliable method for placing foam elements. Starting at a specified height above the desired final location of the module (see Figure 10), the system plans a trajectory that results in the foam element impacting the structure with a desired velocity. Testing showed that low impact velocities, and thus gentler landings, are significantly affected by turbulence around the structure. For this reason, an impact velocity of 1 m/s is chosen.

During the descent maneuver, the position and heading of the quadrocopter are constantly monitored. If the tracking error is too large, the maneuver is aborted, provided there is enough time for the vehicle to recover. Integral action is used along the lateral directions to increase placement accuracy. Integral control is especially helpful for compensating for the effects of imprecise module pickups, which alter the symmetric weight distribution on the vehicle. Figure 11 shows 20 placing trajectories for two different vehicles. During the descent, zero crossing in the acceleration is used to detect the exact impact instant. At the point of impact, the vehicle sits on the module by producing thrust below gravity and controlling the vehicle rotational body rates to zero. After recording the placement position, the vehicle releases the foam element by opening the gripper and flies away.

Placing Results

The position of the placed construction element is indirectly observed through the vehicle. Once the vehicle has placed the module and is resting on it, its position is recorded. The geometry of the machine and the known pickup offset allows the position of the module to be calculated.

Given the particular assembly strategies, the vehicle has no control over the vertical location of a module, and only lateral displacement and orientation errors are of interest. The vertical error is accounted for during construction by calculating the desired vertical location of an element based on the measured vertical position of the two supporting elements, which compensates for cumulative errors such as deviations in element height and the unknown thickness of the adhesive medium used to join each layer of the structure.

Figure 12 illustrates the distribution of the lateral placement errors during construction of the 1500-module tower. The majority of the placements (91.2%) fulfilled both placement accuracy criteria (a maximum of 25-mm error in lateral displacement and 2° of error in orientation). Moreover, 98.27% of the modules satisfied lateral displacement error, which is critical to the structure's stability. The mean lateral error was 9.3 mm with a standard deviation of 6.2 mm. The median error was 8.4 mm. The mean orientation error registered was 0.89°, with standard deviation of 1.1° and a median error of 0.66°. Cumulative vertical errors at the 60th (that is, final) layer of the tower amounted to 5 cm, or 50% of an element's height, and were mostly due to the unmodeled thickness of the connective medium (glue) used to join the construction elements. Recall that the design criteria used to assess the safety of a placement relate to the aforementioned structural stability analysis, which only considers the

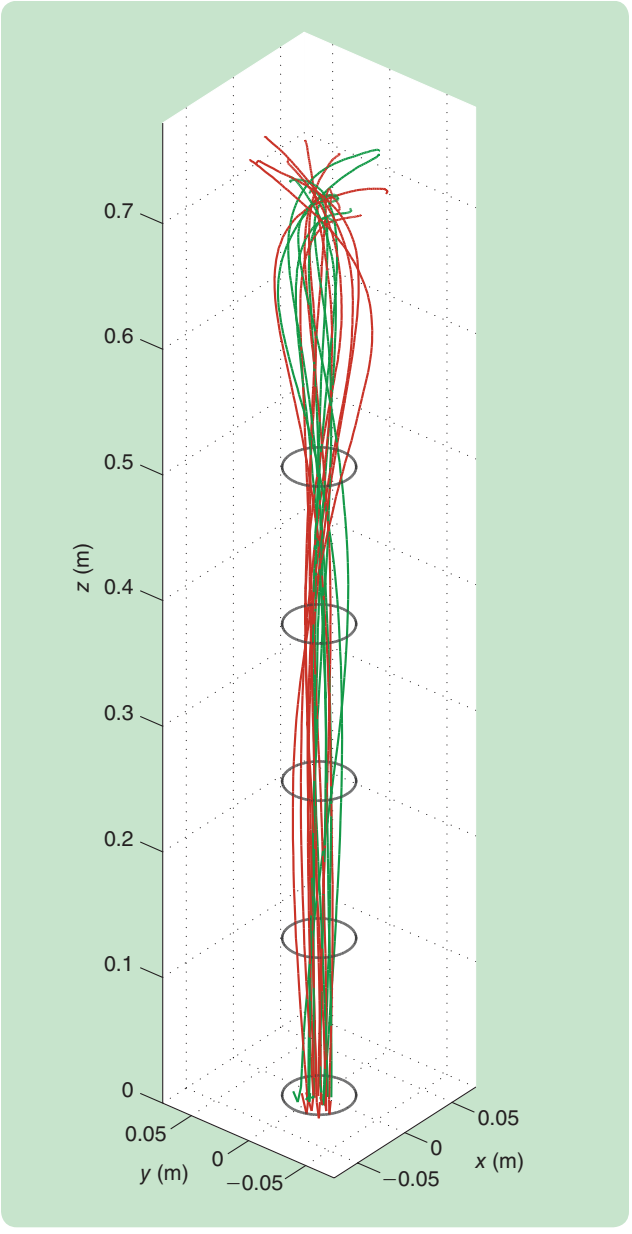

FIGURE 11 Placing trajectories. The plot shows some of the actual placing trajectories during the building of the 1500-module tower, for two different vehicles (red and green). The black circles indicate a deviation of 25 mm from the desired placing positions.

The ability to augment a core system with additional hardware and software modules that enable specialized tasks would be critical in a real-world building scenario.

worst-case scenario (connected modules being placed with maximum error in opposite directions) and does not take into account the enhanced stability obtained by the adhesive bonding created by the glue between elements. Therefore, a placing error out of tolerance does not necessarily compromise the structure's stability.

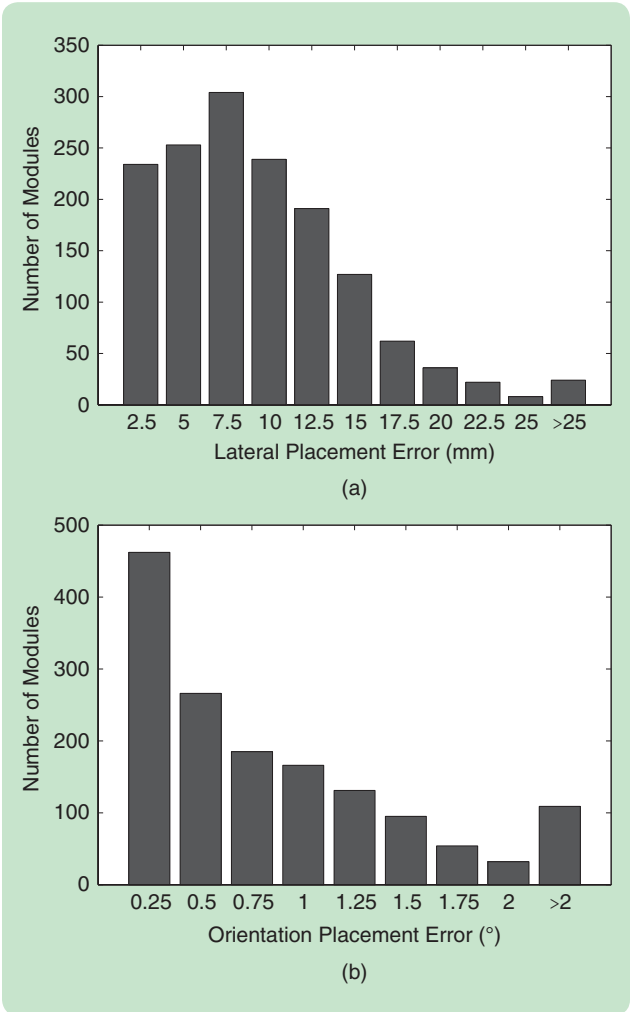

FIGURE 12 The distribution of the magnitudes of the (a) lateral placement error and (b) the orientation error during the building of the 1500-module tower. The mean lateral displacement error was 9.3 mm. The mean orientation error was 0.89°. The majority of the placements (91.2%) fulfilled both placement accuracy criteria (a maximum of 25-mm error in lateral displacement and 2° of error in orientation). Moreover, 98.27% of the modules satisfied the lateral displacement error, which is critical to the structure's stability.

Failure Mitigation

The correct autonomous functioning of the system relies upon the motion capture system measuring the state of each vehicle, with this state being correctly processed into a state estimate at the ground station from which commands are generated, the transmission of these commands over a radio link to the vehicles, and, finally, the execution of the command on board the vehicle. From this critical chain, two main fault causes that pose a significant risk to an installed system running continuously over a longer period of time were identified: 1) the motion capture not seeing a vehicle and 2) the command radio channel failing and thus the commands not arriving at the vehicle. An example of a fault of the first kind is a vehicle becoming occluded by another vehicle or by the structure and thus cannot be seen; an example of the second kind of fault is the scenario where large wireless interference is present on the same frequency as used by the radio system.

Both faults have the same effect on the system, which is that the system can no longer send the vehicle a command based on a recent state estimate. A mitigation scheme that reduces the severity of such faults was developed, which is minimally intrusive to the normal operation of the system while requiring no additional sensors. The strategy consists of periodically sending the vehicle's state to the vehicle and then using a vehicle model and the rate gyroscope measurements to predict this state forward in an open-loop fashion. Each vehicle has an onboard estimate of its own state, on which it can do short-term emergency control if the global control loop is broken.

Because the vehicles' velocity and attitude are unobservable when using only the rate gyroscopes, this estimate will diverge from the truth. This strategy thus offers only a short-term emergency solution, allowing the vehicle to remain in the air for short periods after a fault has occurred. Thus the system is able to cope with faults of short duration, while longer duration faults will still make the vehicle uncontrollable. In this case, however, the vehicle can use its internal state estimate to minimize the severity of the fault. The scheme is described in its entirety in [24].

SPECIALIZED PHYSICAL COMPONENTS

Construction Elements

Given the limited payload of flying machines, the construction elements must be lightweight. The material of choice is

polyurethane foam, which can also be gripped easily by ingressive grippers. To assemble the 1500-module tower, 90-g modules were used. Each module was trapezoidal in shape, 30-cm long, 12–15-cm wide, and 10-cm high, representing a 1:100 model of three-story modules (as described in "The Vertical Village").

The connective medium is as important as the construction element and must provide immediate adhesion to prevent bounces when the modules are flown into place. Waterborne adhesive was manually sprayed on the bottom of modules before putting them into the pickup station. The glue provides good adhesion when a module is placed by the vehicle and results in a permanent bond between elements after drying, making the structure very stable.

Gripper

The gripper, depicted in Figure 13, was designed specifically for the purpose of gripping and carrying foam elements. The gripper consists of three metal pins, each actuated by a single servo. By giving each pin its own servo, the device's mechanical complexity is minimized. The servos and pins are mounted to a three-dimensional-printed rigid gripper base, arranged in a circle with 120° of separation. Each pin end is aligned with a tapered guide in the gripper base, which leads to a small hole in the bottom through which the pin can be extended. A custom circuit board supplying power and input signal from the quadrocopter to the servos is set in the middle of the gripper base. The base itself is secured to a custom laser-cut plate that fits rigidly in the body of the quadrocopter.

The gripper servos are calibrated for two simple states: grip and release. When released, the pin ends protrude slightly from the bottom of the gripper, which helps to reduce slipping when landing on a foam module. When the base of the gripper is sitting flush on a foam module and the gripper state changes from release to grip, the pins extend through the base of the gripper and penetrate the foam module. The gripper is designed in such a way that the angle of attack of the pins decreases relative to the bottom of the gripper as they extend. This motion forces the pins to pull up on the foam module as they penetrate, creating a strong, secure connection to the module.

This design proved to be reliable and robust over the course of the live installation. Given the many cycles of use due to testing and to the installation itself, servos would occasionally expire. In these rare cases, the quadrocopter was still able to grip and lift a foam module with only two working servos, which allowed quadrocopters to transition out of the system for repairs without disturbing the overall workflow.

Charging Stations

Specialized quadrocopter charging stations were designed and built to support continuous system operation and prevent downtime due to battery changes. Two of the charging stations are depicted in Figure 14. The quadrocopter's three-cell lithium polymer battery is connected through four charging pads to an off-the-shelf commercial charger and balancer, which is housed in the lower part of the charging station. Four contacts provide independent access to each of the cells in the battery, allowing for independent voltage monitoring and charge balancing. Each contact consists of a stainless steel loop with a small magnet to ensure surface contact and to help prevent bouncing during landing. Each charger pad is a solid plate of stainless ferromagnetic steel alloy.

As the vehicle lands, the steel loop is pressed by the magnet against the charging plate, ensuring contact, and the charger is asked to begin charging. Each charger is connected through a serial link and an appropriate protocol bridge to the system, providing real-time awareness of charger status. In case of error, the vehicle is commanded to take off and land again, in an attempt to establish a better

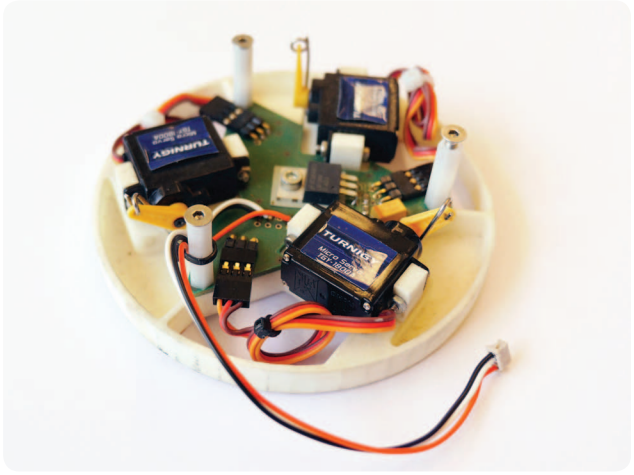

FIGURE 13 The gripper. The gripper consists of three metal pins, each actuated by a single servo. The servos and pins are mounted to a three-dimensional, printed rigid gripper base, arranged in a circle with 120° of separation. A custom circuit board supplying power and input signal from the quadrocopter to the servos is set in the middle of the gripper base.

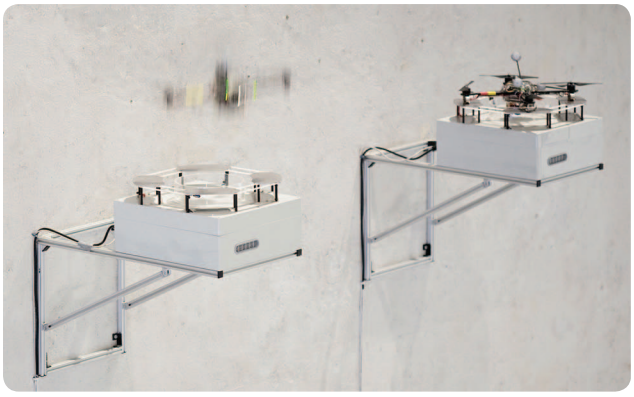

FIGURE 14 Charging stations built to support the continuous operation of the system. The quadrocopter's battery is connected through four charging pads to an off-the-shelf commercial charger and balancer, housed in the lower part of the charging station. (Photograph by François Lauginie.)

charging contact. During the exhibition, four chargers, one for each vehicle, were used continuously. As oils and dirt can easily prevent good contact, the pads on the chargers and vehicles were cleaned once at the beginning of each day of operation and a few times during the exhibit if repeated charging errors were noticed, for example, from glue dust coating the charging pads.

THE EXHIBITION

In the fall of 2011, four flying machines cooperatively assembled a 6-m-tall structure composed of 1500 polyurethane foam modules at the Fonds Régional d'Art Contemporain du Centre (FRAC) in Orléans, France. The FRAC Centre is an architectural art museum on the outskirts of Paris that promotes contemporary art, both nationally and internationally. The structure erected within the Flight Assembled Architecture installation is now part of the museum's permanent collection.

The installation was the result of intense preparation that took place during the months preceding the event. The previous sections addressed the development of the robotic systems and the various design choices that were made before and throughout the development process. However, deploying a live autonomous system outside laboratory conditions required team effort, system robustness, and extensive testing. The next sections present insights into the development process.

Development

The development of the installation began one year before the event. The team met and discussed the possibilities offered by flying machines and the constraints that they would impose on the construction process, which guided the design of a structure that, in its scale model, could be assembled by quadrocopters. At the same time, interfaces between the structure blueprint and the autonomous systems were defined to allow for the parallel development of both components, which resulted in the aforementioned system architecture, where the blueprint connects the structure design to the autonomous construction system, and the pickup station provides a physical interface between machines and humans.

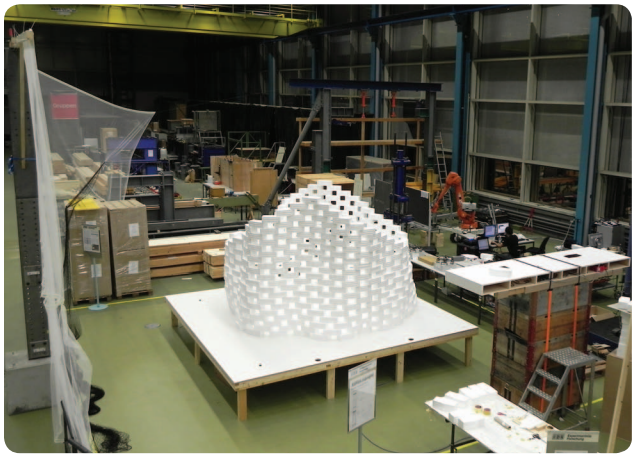

FIGURE 15 The structure during the construction test. A complete test structure was built one month before the actual event. The building test entailed deploying the motion capture system and the rest of the Flying Machine Arena infrastructure to an empty hall.

Having defined methods and goals (for example, the use of foam modules and the ability to accurately place them), the development of the different system components began. This development was performed in the FMA (for details see "The Flying Machine Arena"), a testbed for quadrocopter research. Many gripper prototypes were designed, and initial module placement tests were conducted.

The reliable pickup and placement of modules is core to the construction system, and extensive time was devoted to the development of these components. Once these components were functional, the system was complemented with the navigation system and automated charging stations. Each component was tested individually (for example, the automated charging stations were tested by continuously flying the vehicle overnight) before being integrated into the final system, which was then extensively tested in simulation. A complete test structure was built one month before the actual event. Figure 15 shows the structure during the building test. This test highlighted unexpected faults in the system (such as the sporadic crashing of third-party software that resulted in losses of vehicle pose information) and forced the team to improve the estimators and develop the aforementioned fault mitigation strategy so as to increase the system's safety and robustness. Furthermore, the building test required deploying the FMA infrastructure to an empty hall and equipping it with a motion capture system.

Deployment of the Mobile System

A mobile system for installations outside the laboratory was built in addition to the in-lab testbed. When deploying the system, special care must be taken for the placement of the motion capture system cameras. The cameras must be rigidly mounted and cannot move relative to each other. Furthermore, the cameras must adequately cover the flyable space. While nominally just two cameras are required to see a vehicle to determine its location and attitude in space, ensuring that the space is covered by three to four cameras provides redundancy against possible camera failures, temporary occlusions due to other vehicles, and the erroneous mounting of cameras. Ensuring robustness against occlusion must be considered when a large structure is being assembled in the space. To this purpose, a software tool that checks camera coverage was developed: with knowledge of the cameras' field of view, positions and orientations, and knowledge of objects in the space, the tool indicates by how many cameras a point in space is seen. The software tool was first used during the planning phase to design the motion capture system configuration. After the cameras have been placed and calibrated, their actual positions and

orientations are checked against the designed configuration and the actual camera coverage is evaluated. Figure 16 depicts the camera coverage of the installation space and the spots that are occluded by the building of the structure.

The Installation

The structure, shown in Figure 17, was assembled in 18 hours in front of exhibition attendees over four days. During the opening night, shown in Figure 18, the museum hosted 300 people. Before that night, however, the system was installed in the museum space and thoroughly tested. Part of the team arrived in Orléans almost two weeks before the event to prepare the empty museum space for the installation. The control room was set up, the cameras were mounted according to the plan, the charging stations were placed 4 m from the ground, and a first test structure was built. Two days before the opening night, the building of the final tower started, and about one-third of the tower was assembled before the first exhibition day. A video stream from an onboard camera and the on-screen visualization of a three-dimensional environment gave the audience insight into the system from the quadrocopter's point of view. Figure 19 is a photo taken from the onboard camera.

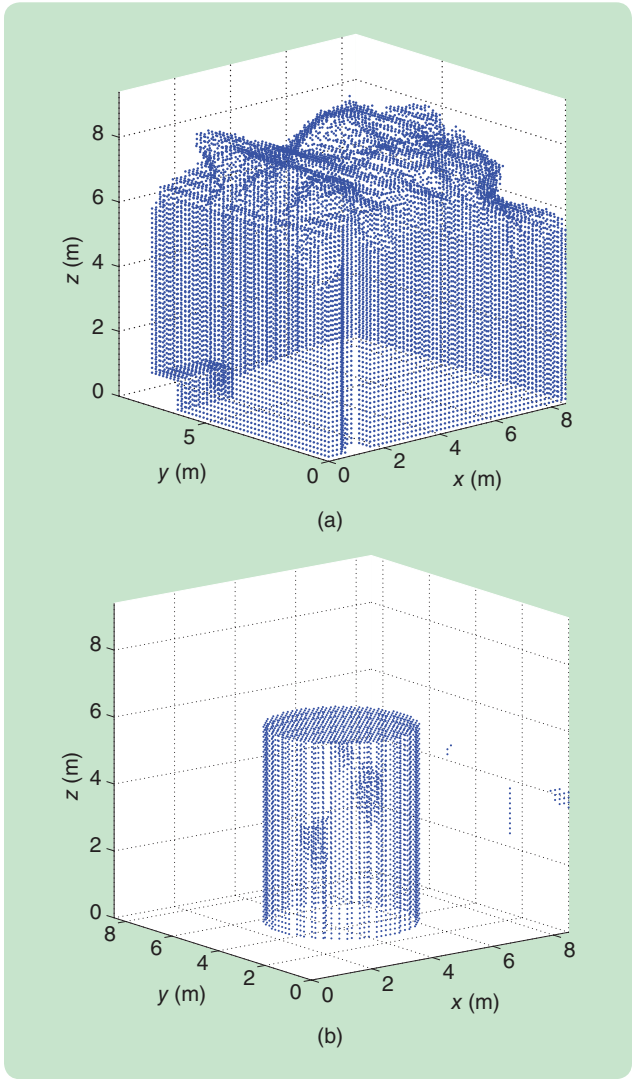

FIGURE 16 The motion capture camera coverage. A software tool uses knowledge of the cameras' field of view, positions, orientations, and knowledge of the objects in the space to indicate how many cameras can see any given point in space: (a) the blue dots represent the portion of the space seen by at least three cameras and (b) the dots represent the locations that are occluded by the building of the structure.

FIGURE 17 Time-lapse photos of the event. The structure was assembled in front of a live audience over four days. From left to right, top to bottom, the height of the structure in layers is 1, 7, 15, 18, 24, 36, 47, 52, 58, 60, 60, and 60.

FIGURE 18 People watching the construction during the opening night, which was attended by 300 people.

Culminating in a 6-m-tall tower composed of 1500 foam modules, the installation was assembled by four quadrocopters in 18 hours during a four-day-long live exhibition.

Some critical moments were faced during the building of the structure; however, these were mitigated by the robust design of the system. During the opening night, a module was placed with a lateral error greater than 6 cm. Despite being above the placement error that the tower was designed for, the module did not fall, thanks to additional connective force produced by the glue placed on the modules, which was not taken into account in the structure stability analysis. Close to the end of the installation, a module slid after being placed at a height of 5 m (the glue was not applied correctly), resulting in a very small supporting surface for the module that was to be placed on top of it. Team members could not reach the module to manually restore it in place. The situation was monitored closely, and the team decided to only fly one vehicle at a time, thus reducing the potential damages that would have resulted from a severe fault. The structure was completed without further incident. During the 18 h of flight, the system suffered a single accident when the motion capture system stopped transmitting data. The fault mitigation strategy kicked in, reducing the vehicle speed and altitude, which mitigated the effects of the fault.

These episodes show that unforeseen difficulties might occur regardless of how much planning is done. However, they also demonstrate how a robust system design and adequate backup solutions allow for smooth execution, fault mitigation, and minimal downtime. The reliability was achieved through extensive testing, development of robust submodules, and pragmatic design choices that do not add fragility to the system. Furthermore, clear interfaces and milestones were created to enable the parallel development of the different aspects of the Flight Assembled Architecture.

FIGURE 19 The quadrocopter point of view. A photo taken by the onboard camera during the opening night.

CONCLUSIONS

In the context of aerial construction, the Flight Assembled Architecture installation and the aerial construction system presented in this article should be seen as a proof of concept, demonstrating the ability of aerial vehicles to build structures. The project, however, did not preserve the real-scale spatial assembling principles of construction, that is, the methods cannot be applied one to one to real-size buildings. For aerial construction to succeed in real-world scenarios, researchers must explore strategies that combine the abilities of flying machines to reach almost any point in space and move construction elements to locations not otherwise accessible. Although the Flight Assembled Architecture installation used a motion capture system for observing vehicle position and attitude, the methods and algorithms described in this article are not reliant on such a system. However, an alternative localization system is required to achieve similar results in real-world scenarios. Researchers must also develop new material systems and novel construction processes that address the constraints imposed by these machines, such as payload and accuracy. Success in this endeavor will require researchers in a number of disciplines to work closely together [8], [25].

ACKNOWLEDGMENTS

The authors thank their teams for their efforts, particularly Luzius Brodbeck, Marc-André Corzillius, Carolina Flores, Hans Ulrich Honegger, Chanakya Hridaya, Angela P. Schoellig, and Igor Thommen of the Institute for Dynamics Systems and Control and Peter Heckeroth, Andrea Kondziela, and Marion Ott, chair of Architecture and Digital Fabrication. The authors are also grateful to the Fonds Régional d'Art Contemporain du Centre Orlêans (coproducer of the Flight Assembled Architecture installation) for both giving the opportunity of realizing such an experiment and its generous support in the overall project. The project would have not been possible without the valuable support of Dr. Lüchinger+Meyer Bauingenieure AG, Amstein+Walthert AG, and Prof. Jan Carmeliet, ETH Zurich and Empa. The authors are also grateful for the sponsorships of Région Centre, Ministry of Culture and Communication (Direction Régionale des Affaires Culturelles du Centre), Pro Helvetia Swiss

Arts Council, Centre Culturel Suisse Paris, Platform Regroupement des Fonds régionaux d'art contemporain, Vicon Motion Systems, ERCO Leuchten GmbH, and JET Schaumstoff-Formteile GmbH. Some of the photographs in this article were taken by François Lauginie, Carolina Flores, and Markus Waibel. Thanks also go to Hallie Siegel for her efforts on the editorial revision of this article. This work was supported by the Swiss National Science Foundation (SNSF).

AUTHOR INFORMATION

Federico Augugliaro (faugugliaro@ethz.ch) is currently a Ph.D. candidate at the Institute for Dynamic Systems and Control at ETH Zurich. He received the M.Sc. in robotics, systems, and control in 2011 and the B.Sc. in mechanical engineering in 2009 from ETH Zurich. During his master's thesis, he created algorithms and tools that facilitate the generation of multiquadrocopter choreographies. His research features aerial construction, human-quadrocopter physical interaction, and coordinated multivehicle flight. He received the Willi Studer Award and the Hans-Eggenberger-Prize in 2012 for his master's studies. He can be contacted at ETH Zurich, Sonneggstrasse 3, ML K38, 8092 Zurich, Switzerland.

Sergei Lupashin received the B.Sc. in electrical and computer engineering from Cornell University in 2006 and the M.Sc. and Ph.D. in mechanical engineering from ETH Zurich in 2010 and 2012, respectively. He is currently a postdoc at the Robust Perception Group at the University of Zurich, pursuing commercialization of novel aerial imaging approaches. Past activities include participation in the Robocup competition, participation in the DARPA Grand Challenge as part of the Cornell Team, and key contributions to the Flying Machine Arena at ETH Zurich.

Michael Hamer received the B.Eng. in computer engineering and the B.Sc. in computer science from Curtin University, Australia, in 2010, and the M.Sc. in robotics, systems, and control from ETH Zurich in 2013. He is currently pursuing his interests in machine learning and robot localization as a Ph.D. candidate at the Institute for Dynamic Systems and Control at ETH Zurich.

Cason Male received the B.Sc. in mechanical engineering from Carnegie Mellon University in 2009 with a minor in robotics and the M.Sc. in electrical engineering from Stanford University in 2011. His studies at Stanford were in controls, machine learning, and optimization, while his research focused on model predictive control for robotic manipulators. After graduating from Stanford, he spent seven months at ETH Zurich, working at the Institute for Dynamic Systems and Control. He is currently a robotics engineer specializing in machine learning and computer vision applied to agricultural robotics at the National Robotics Engineering Center in Pittsburgh.

Markus Hehn is a Ph.D. candidate at the Institute for Dynamic Systems and Control at ETH Zurich. He received the Diplom-Ingenieur in mechatronics from Technische Universität Darmstadt in 2009, having receiving a Robert Bosch GmbH scholarship for his graduate studies. He has worked on optimizing the operating strategy of diesel engines to reduce emissions and fuel consumption, the characterization of component load profiles for axle-split hybrid vehicle drivetrains, and the performance development of Formula One racing engines. His main research interest are the control and trajectory generation for rotorcraft during fast maneuvering, optimality-based maneuvers, multivehicle coordination, and learning algorithms.

Mark W. Mueller is a doctoral candidate at the Institute for Dynamic Systems and Control at ETH Zurich. He received the B.Eng. from the University of Pretoria in 2009 and the M.Sc. from ETH Zurich in 2011, both in mechanical engineering. He received awards for the best mechanical engineering thesis and the best aeronautical thesis for his bachelor's thesis in 2008 and the Jakob Ackeret award from the Swiss Association of Aeronautical Sciences for his master's thesis in 2011. His master's studies were supported by a scholarship from the Swiss government.

Jan Sebastian Willmann received the B.Sc. from the University of Liechtenstein in 2004, the M.Sc. from Oxford Brookes University in 2006, and the Ph.D. from the University of Innsbruck in 2010, all in architecture. He worked for numerous architectural offices as an architect and consultant; cofounded the platform Architekturtheorie.eu; and participated in a number of international research, exhibition, and publication projects, approaching architecture in the digital age as a composed theoretical, computational, and material score. He has lectured and taught at numerous universities and institutions, such as the University of Applied Arts Vienna, the Academy of Fine Arts Stuttgart, and the University of Pennsylvania. Since 2011, he has been a senior researcher of architecture and digital fabrication at ETH Zurich.

Fabio Gramazio is an architect with multidisciplinary interests ranging from computational design and digital fabrication to material innovation. In 2000, he founded, with Matthias Kohler, the architecture practice Gramazio & Kohler, where numerous award-winning designs have been realized, which include the Gantenbein vineyard façade, the Tanzhaus theatre for contemporary dance, the Christmas lights for the Bahnhofstrasse Zurich, the sWISH* Pavilion at the Swiss National Exposition Expo.02, and the Private House in Riedikon. Founding the world's first architectural robotic laboratory at ETH Zurich, the firm's academic research concentrates on additive robotic fabrication, including advanced architectural design, full-scale prototypical installations, and the design of robotically fabricated high-rise buildings. He was awarded the Swiss Art Awards, the Global Holcim Innovation Prize, and the Acadia Award for Emerging Digital Practice. His work has been published in a large number of academic journals, further contributing to numerous exhibitions around the world such as the Architecture Biennale in Venice (2008), the Storefront

Gallery for Art and Architecture (2009) New York, and the Flight Assembled Architecture at the Fonds Régional d'Art Contemporain du Centre Orléans (2011).

Matthias Kohler is an architect with multidisciplinary interests ranging from computational design and digital fabrication to material innovation. In 2000, he founded, with Fabio Gramazio, the architecture practice Gramazio & Kohler, where numerous award-winning designs have been realized, which include the Gantenbein vineyard façade, the Tanzhaus theatre for contemporary dance, the Christmas lights for the Bahnhofstrasse Zurich, the sWISH* Pavilion at the Swiss National Exposition Expo.02, and the Private House in Riedikon. Founding the world's first architectural robotic laboratory at ETH Zurich, the firm's academic research concentrates on additive robotic fabrication, including advanced architectural design, full-scale prototypical installations, and the design of robotically fabricated high-rise buildings. He was awarded the Swiss Art Awards, the Global Holcim Innovation Prize, and the Acadia Award for Emerging Digital Practice. His work has been published in a large number of academic journals, further contributing to numerous exhibitions around the world such as the Architecture Biennale in Venice (2008), the Storefront Gallery for Art and Architecture (2009) New York, and the Flight Assembled Architecture at the Fonds Régional d'Art Contemporain du Centre Orléans (2011).

Raffaello D'Andrea received the B.Sc. in engineering science from the University of Toronto in 1991 and the M.Sc. and Ph.D. in electrical engineering from the California Institute of Technology in 1992 and 1997, respectively. He was an assistant, and then an associate, professor at Cornell University from 1997 to 2007. While on leave from Cornell, from 2003 to 2007, he cofounded Kiva Systems, where he led the efforts in systems architecture, robot design, robot navigation and coordination, and control algorithms. A creator of dynamic sculpture, his work has appeared at various international venues, including the National Gallery of Canada, the Venice Biennale, Ars Electronica, the Smithsonian, the Fonds Régional d'Art Contemporain du Centre, and the Spoleto Festival. He is currently professor of dynamic systems and control at ETH Zurich.

REFERENCES

[1] F. Gramazio, M. Kohler, and R. D'Andrea, *Flight Assembled Architecture*. Orléans, France: Editions Hyx, 2013.

[2] V. Helm, S. Ercan, F. Gramazio, and M. Kohler, "Mobile robotic fabrication on construction sites: DimRob," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots Systems*, Vilamoura, Portugal, 2012, pp. 4335–4341.

[3] D. Schodek, M. Bechthold, K. Griggs, K. Kao, and M. Steinberg, *Digital Design and Manufacturing: CAD/CAM Applications in Architecture*. New York: Wiley, 2004.

[4] J. Ficca, "Inclusion of performative surfaces material and fabrication research," in *Digital Fabrications: Architectural and Material Techniques*, L. Iwamoto, Ed. New York: Princeton Architectural Press, 2009.

[5] B. Kolarevic, *Architecture in the Digital Age: Design and Manufacturing*. New York: Taylor & Francis, 2003.

[6] Q. Lindsey and V. Kumar, "Distributed construction of truss structures," in *Algorithmic Foundations of Robotics X* (Springer Tracts in Advanced Robotics, vol. 86), E. Frazzoli, T. Lozano-Perez, N. Roy, and D. Rus, Eds. Berlin Heidelberg, Germany: Springer, 2013, pp. 209–225.

[7] K. Kondak, K. Krieger, A. Albu-Schaeffer, M. Schwarzbach, M. Laiacker, I. Maza, A. Rodriguez-Castano, and A. Ollero, "Closed-loop behavior of an autonomous helicopter equipped with a robotic arm for aerial manipulation tasks," *Int. J. Adv. Robot. Syst.*, vol. 10, p. 1, Feb. 2013.

[8] F. Augugliaro, A. Mirjan, F. Gramazio, M. Kohler, and R. D'Andrea, "Building tensile structures with flying machines," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots Systems*, Tokyo, Japan, 2013, pp. 3487–3492.

[9] L. Marconi and R. Naldi, "Control of aerial robots: Hybrid force and position feedback for a ducted fan," *IEEE Control Syst.*, vol. 32, no. 4, pp. 43–65, 2012.

[10] N. Michael, J. Fink, and V. Kumar, "Cooperative manipulation and transportation with aerial robots," *Auton. Robots*, vol. 30, no. 1, pp. 73–86, 2011.

[11] M. Bernard, K. Kondak, I. Maza, and A. Ollero, "Autonomous transportation and deployment with aerial robots for search and rescue missions," *J. Field Robot.*, vol. 28, no. 6, pp. 914–931, 2011.

[12] R. Ritz and R. D'Andrea, "Carrying a flexible payload with multiple flying vehicles," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots Systems*, Tokyo, 2013, pp. 3465–3471.

[13] P. E. I. Pounds, D. R. Bersak, and A. M. Dollar, "Practical aerial grasping of unstructured objects," in *Proc. IEEE Conf. Technologies Practical Robot Applications*, Woburn, MA, 2011, pp. 99–104.

[14] D. Mellinger, M. Shomin, N. Michael, and V. Kumar, "Cooperative grasping and transport using multiple quadrotors," in *Distributed Autonomous Robotic Systems* (Springer Tracts in Advanced Robotics, vol. 83), A. Martinoli, F. Mondada, N. Correll, G. Mermoud, M. Egerstedt, M. A. Hsieh, L. E. Parker, and K. Sty, Eds. Berlin Heidelberg, Germany: Springer, 2013, pp. 545–558.

[15] C. Papachristos, K. Alexis, and A. Tzes, "Towards a high-end unmanned tri-TiltRotor: Design, modeling and hover control," in *Proc. IEEE Mediterranean Conf. Control Automation*, Barcelona, Spain, 2012, pp. 1579–1584.

[16] R. Voyles and G. Jiang, "Hexrotor UAV platform enabling dextrous interaction with structures—Preliminary work," in *Proc. IEEE Int. Symp. Safety, Security, Rescue Robotics*, College Station, TX, 2012, pp. 1–7.

[17] S. Lupashin, M. Hehn, M. W. Mueller, A. P. Schoellig, M. Sherback, and R. D'Andrea, "A platform for aerial robotics research and demonstration: The flying machine arena," *Mechatronics*, vol. 24, no. 1, pp. 41–54, 2014.

[18] G. Rumpel and H. Sondershausen, "Mechanics," in *Dubbel Handbook of Mechanical Engineering*, W. Beitz and K.-H. Küttner, Eds. London, United Kingdom: Springer, 1994, pp. A1–A72.

[19] Tracker. (2014). Vicon Motion Systems Ltd. [Online]. Available: http://www.vicon.com/Software/Tracker

[20] O. Purwin, R. D'Andrea, and J.-W. Lee, "Theory and implementation of path planning by negotiation for decentralized agents," *Robot. Auton. Syst.*, vol. 56, no. 5, pp. 422–436, 2008.

[21] R. D'Andrea, P. R. Wurman, M. T. Barbehenn, A. E. Hoffman, and M. C. Mountz, "System and method for managing mobile drive units," U.S. Patent US7 873 469 B2, Jan. 18, 2011.

[22] E. Guizzo, "Three engineers, hundreds of robots, one warehouse," *IEEE Spectr.*, vol. 45, no. 7, pp. 26–34, 2008.

[23] M. Hehn and R. D'Andrea, "Quadrocopter trajectory generation and control," in *Proc. IFAC World Congr.*, Milan, Italy, 2011, pp. 1485–1491.

[24] M. W. Muller and R. D'Andrea, "Critical subsystem failure mitigation in an indoor UAV testbed," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots Systems*, Vilamoura, Portugal, 2012, pp. 780–785.

[25] J. Willmann, F. Augugliaro, T. Cadalbert, R. D'Andrea, F. Gramazio, and M. Kohler, "Aerial robotic construction towards a new field of architectural research," *Int. J. Architectural Comput.*, vol. 10, no. 3, pp. 439–460, 2012.

[26] D. Gurdan, J. Stumpf, M. Achtelik, K.-M. Doth, G. Hirzinger, and D. Rus, "Energy-efficient autonomous four-rotor flying robot controlled at 1 kHz," in *Proc. IEEE Int. Conf. Robotics Automation*, Rome, Italy, 2007, pp. 361–366.

[27] A. E. Bryson and Y.-C. Ho, *Applied Optimal Control*. New York: Taylor & Francis, 1975.