2012 IEEE/RSJ International Conference on
Intelligent Robots and Systems
October 7-12, 2012. Vilamoura, Algarve, Portugal

A Parameterized Control Methodology for a Modular Flying Vehicle

Raymond Oung, Miguel Picallo Cruz, and Raffaello D'Andrea

Abstract— **Designing a controller that is scalable, robust, and that can adapt to an arbitrary configuration is one of the major challenges of modular robotics. This paper takes one step forward in addressing this challenge by presenting a methodology for controlling *any* flight-feasible configuration of a modular flying vehicle, in this case the Distributed Flight Array (DFA). In this work we present a well-structured, parameterized controller and describe a method for optimizing its parameters in order to achieve the best possible performance subject to the system's physical constraints. We then show how the configuration space of the DFA can be parameterized by only a few variables and propose a straightforward approach for mapping this configuration space to its control parameter space.**

I. Introduction

One of the key challenges in modular robotics is control design. The majority of past research has focused on hardware and software challenges [1] [2] [3]; a summary of such work can be found in [4] [5]. Difficulties in controlling such systems are rooted in the fact that they are distributed by nature – each module has limited information, yet all must coordinate their actions in order to achieve a desired behaviour. Such a control system should be scalable, robust to single point failures, and adaptable to arbitrary configurations. Research that has taken a step in this direction includes [6] [7], as well as the work that is presented here.

This paper presents a generalized tractable methodology for controlling any flight-feasible configuration of a modular flying vehicle, namely the Distributed Flight Array (DFA) [8], see Fig. 1. We build upon previous work and describe a parameterized, decentralized controller for an arbitrary configuration of the DFA that enables it to hover. By decentralized, we mean that the controller relies only on its local sensors for control feedback and that no information is shared between modules during flight, which results in a scalable and straightforward implementation.

The tuning parameters for this controller can then be optimized for a given performance metric while taking into account physical constraints of the system. This decentralized controller is compared against two idealized cases: (1) a centralized architecture, where each module has direct sensor feedback from all modules in the vehicle, and (2) the centralized $\mathcal{H}_2$ optimal controller, which is used as a benchmark for performance. We then show that it is possible to parameterize the configuration space of the DFA using only a few variables and demonstrate that a relatively simple function can be used to map a vehicle's configuration to an appropriate set of control parameters, which enables the DFA to hover.

R. Oung and R. D'Andrea are with the Institute for Dynamic Systems and Control, ETH Zurich, 8092 Zurich, Switzerland {roung, rdandrea}@ethz.ch; M. Picallo is with Escola Técnica Superior d'Enginyers Industrials de Barcelona, UPC BarcelonaTech, 08028 Barcelona, Spain miguel.picallo@estudiant.upc.edu

Fig. 1. The Distributed Flight Array is a modular multi-propeller vehicle that can fly in ad-hoc configurations. Configurations may range in size, shape, and sparsity. Designing a controller by hand for every flight-feasible configuration would be impractical.

Beginning with Section II, we describe the nominal flight dynamics considered for the DFA. In Section III, details of the control architecture and its design parameters are given. We then briefly discuss the disturbances that our controller is designed to handle in Section IV. In Section V an automated procedure for computing the control parameters that optimizes a particular performance metric is discussed. Results regarding controller performance are presented in Section VI, and in Section VII we propose a simple function that maps the configuration space of the DFA to its control parameter space.

II. Nominal Flight Dynamics

The nominal operating conditions considered for the DFA's control design occur while hovering. Hover is defined here as the process of maintaining a constant position (x, y, z) in the air and a constant yaw angle α with respect to an inertial coordinate frame I. Although this approach may seem to limit the DFA's maneuverability, flying vehicles with controllers designed around the hovering equilibrium have demonstrated to be quite agile [9]. The DFA is modeled as a rigid-body and it is controlled by varying the thrusts f_i produced by each module i positioned at (x_i, y_i) with respect to the body coordinate frame B. The body coordinate frame B is positioned at the vehicle's center of mass and is oriented along the principal axes of rotation, see Fig. 2. The

978-1-4673-1736-8/12/S31.00 ©2012 IEEE

Authorized licensed use limited to: ETH BIBLIOTHEK ZURICH. Downloaded on May 15,2024 at 21:25:47 UTC from IEEE Xplore. Restrictions apply.

orientation of the body coordinate frame B with respect to the inertial frame I is described by ZYX-Euler angles: yaw α, pitch β, roll γ acting along the z-, y-, x-axis, respectively.

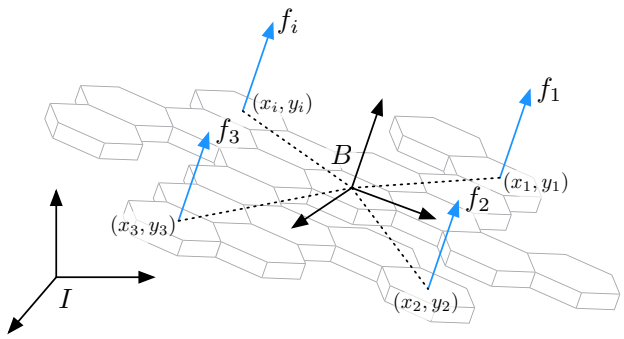

Fig. 2. The DFA's body coordinate frame B is located at its center of mass and oriented along the principal axes of rotation. The orientation of the vehicle with respect to the inertial frame I can be described by a set of ZYX-Euler angles. The DFA can be controlled by varying the thrust f_i produced by module i.

An approximately equal number of clockwise (CW) and counterclockwise (CCW) rotating propellers are required in the vehicle in order to cancel out the torques in flight. In this case, the nominal thrust required of each module counteracts the acceleration due to gravity g. As such, the control effort during nominal flight of each module i is given by $a_i = f_i - mg$, where m is the mass of a module.

During nominal flight the translational dynamics of the DFA with respect to an inertial coordinate frame I is a function of its total thrust and its orientation. To first-order approximation at hover, the translational accelerations $\ddot{x}$ and $\ddot{y}$ are respectively a consequence of the DFA's pitch β and roll γ angles:

$$\ddot{x} = \beta g, \tag{1}$$

$$\ddot{y} = -\gamma g. \tag{2}$$

The total control effort that is produced by N modules is the sum of all individual control efforts $\sum_{i=1}^{N} a_i$. The rolling and pitching torques of the DFA are a function of the control efforts respectively acting along the appropriate moment arm (y, x), $\sum_{i=1}^{N} y_i a_i$ and $-\sum_{i=1}^{N} x_i a_i$. The yawing torque around its vertical axis is the sum of all reaction torques, $\sum_{i=1}^{N} c_i a_i$, where c_i is a constant that maps force to torque and can be a negative or positive value depending on the handedness of the propeller [8]. These equations of motion can be more compactly expressed as

$$\mathbf{M} \begin{bmatrix} \ddot{z} \\ \ddot{\gamma} \\ \ddot{\beta} \\ \ddot{\alpha} \end{bmatrix} = \mathbf{P}^T \mathbf{a}, \tag{3}$$

where the vector $\mathbf{a} = (a_1, \ldots, a_N)$ contains the control efforts of each module; the diagonal matrix $\mathbf{M} = \text{diag}(Nm, I_x, I_y, I_z)$ contains the total mass Nm and the principal mass moments of inertia (I_x, I_y, I_z); and the matrix $\mathbf{P} \in \mathbb{R}^{N \times 4}$ contains information pertaining to the configuration of the vehicle:

$$\mathbf{P} = \begin{bmatrix} 1 & y_1 & -x_1 & c_1 \\ \vdots & \vdots & \vdots & \vdots \\ 1 & y_N & -x_N & c_N \end{bmatrix}. \tag{4}$$

Each DFA module uses an off-the-shelf brushless DC motor, motor speed controller, and propeller, which together function as an actuator with a particular bandwidth ω_m. This actuator has been modeled and shown to exhibit first-order dynamics around hover:

$$\dot{a}_i = -\omega_m (a_i - a_{d,i}), \tag{5}$$

where $a_{d,i}$ is the commanded (or desired) thrust.

III. Control Architecture

Previous work developed a control strategy of the form [8]:

$$\mathbf{a}_d = \mathbf{Q}\mathbf{f}, \tag{6}$$

where the set of functions $\mathbf{f} = (f_z, f_\gamma, f_\beta, f_\alpha)$ are mapped to the desired control effort $\mathbf{a}_d \in \mathbb{R}^{N \times 1}$ by the matrix $\mathbf{Q} \in \mathbb{R}^{N \times 4}$. We can decouple the DOF of the system by designing $\mathbf{Q}$ such that $\mathbf{P}^T\mathbf{Q} = \mathbf{I}_{4\times 4}$. Note that not just any configuration of the vehicle is controllable, for example a vehicle where all modules are aligned has insufficient actuation along one of its DOF. To be more precise, our methodology works for any configuration where the DOF of the system can be decoupled.

Considering each DOF separately and assuming full state feedback, we can choose the following second-order functions, each containing two design parameters (a closed-loop natural frequency ω and damping ratio ζ):

$$f_z = -Nm(2\omega_z\zeta_z\hat{\dot{z}} + \omega_z^2(\hat{z} - z_d)), \tag{7}$$

$$f_\gamma = -I_x(2\omega_\gamma\zeta_\gamma\hat{\dot{\gamma}} + \omega_\gamma^2(\hat{\gamma} - \gamma_d)), \tag{8}$$

$$f_\beta = -I_y(2\omega_\beta\zeta_\beta\hat{\dot{\beta}} + \omega_\beta^2(\hat{\beta} - \beta_d)), \tag{9}$$

$$f_\alpha = -I_z(2\omega_\alpha\zeta_\alpha\hat{\dot{\alpha}} + \omega_\alpha^2(\hat{\alpha} - \alpha_d)). \tag{10}$$

where the hat notation denotes the estimated state and the subscript d denotes the desired state.

As previously mentioned in Section II, the translational DOF (x, y) are indirectly controlled through the roll γ and pitch β angles (1) (2); they can also be designed to exhibit a second-order response:

$$\ddot{x} := -2\omega_x\zeta_x\hat{\dot{x}} - \omega_x^2(\hat{x} - x_d), \tag{11}$$

$$\ddot{y} := -2\omega_y\zeta_y\hat{\dot{y}} - \omega_y^2(\hat{y} - y_d). \tag{12}$$

The two desired angles γ_d and β_d in (8) (9) are then

$$\gamma_d = \frac{2\omega_y\zeta_y\hat{\dot{y}} + \omega_y^2(\hat{y} - y_d)}{g}, \tag{13}$$

$$\beta_d = -\frac{2\omega_x\zeta_x\hat{\dot{x}} + \omega_x^2(\hat{x} - x_d)}{g}. \tag{14}$$

Authorized licensed use limited to: ETH BIBLIOTHEK ZURICH. Downloaded on May 15,2024 at 21:25:47 UTC from IEEE Xplore. Restrictions apply.

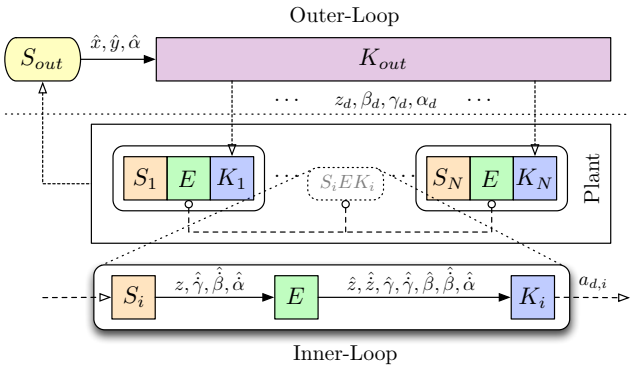

Fig. 3. A module consists of a set of sensors (S_i), a state estimator (E), and an inner-loop controller (K_i). Each module is able to communicate with its adjacent (1-hop) neighbour. The outer-loop sensor (S_{out}) measures states that are unavailable to the modules, and the outer-loop controller (K_{out}) forwards the desired states to each module.

A. Implementation

The DFA uses a cascaded control architecture, consisting of an inner control loop and an outer control loop, see Fig. 3.

Our controller design is based on continuous-time system dynamics because the sampling rate of the sensors, the rate of the control loops, and the communication rate between outer-loop and inner-loop is significantly faster than the system's dynamics. As such, performance losses are considered to be minor.

1) Inner Control Loop: The inner-loop resides on each module i and is responsible for controlling states that it can directly measure, which consists of the altitude and attitude (with the exception of yaw) DOF. The module's sensor set includes the following: a 3-axis rate-gyroscope, which measures angular body-rates[1]; and an infrared range finder used to (1) measure altitude $\hat{z}$, and (2) estimate roll $\hat{\gamma}$ and pitch $\hat{\beta}$ angles when flying over a flat surface by sharing distance measurements between modules over a communication network [11]. The DOF in z, γ, β are designed to exhibit second-order responses (7) – (9). At this stage, only yaw-rate is controlled:

$$f_\alpha = -\frac{I_z}{\tau_1}(\hat{\dot{\alpha}} - \dot{\alpha}_d), \tag{15}$$

where $\tau_1 = 1/2\omega_\alpha\zeta_\alpha$ is the closed-loop time constant of the first-order yaw-rate controller.

2) Outer Control Loop: The outer-loop resides off board and is responsible for controlling the vehicle's position (x, y) with respect to an inertial frame by communicating the desired states (γ_d, β_d) given by (13) (14) to the inner control loop. To do this, the outer-loop relies on an offboard camera-based motion capture system for state measurements [12]. Architecting the controller in this way facilitates, for example, the integration of an onboard position sensor such as a GPS in the future. Since yaw angle is also measurable, a proportional controller that acts on the measured yaw angle and outputs a desired yaw-rate can be used:

[1]Body-rates can be mapped back to Euler rates and/or vice-versa via the appropriate attitude kinematics relations [10].

$$\dot{\alpha_d} = -\frac{1}{\tau_2}(\hat{\alpha} - \alpha_d), \tag{16}$$

where $\tau_2 = 2\zeta_\alpha/\omega_\alpha$ is the closed-loop time constant.

Cascaded with the inner-loop yaw-rate controller (15), the closed-loop dynamics in yaw becomes second-order, as in (10).

IV. Expected Disturbances

A. Process Noise

The main source of disturbance is due to the turbulent flow of air generated by each propeller when spinning. These disturbances are internal to the system and are spatially uncorrelated. In an outdoor environment, one could also take into account wind gusts using a Dryden wind turbulence model [13].

Each DFA module uses an off-the-shelf brushless DC motor and connected directly to a propeller. The variation in thrust produced by the propeller and motor enclosed in the module's chassis was measured at nominal thrust using a 6-axis force-torque sensor with a sampling frequency of 7 kHz. The power spectral density (PSD) function $S_{\nu\nu}(\Omega)_k$ indicates a significant amount of power contained at $\Omega_p = 0.16$ radians per sample (or $1144\,\mathrm{rad/s}$), which as expected is twice the angular velocity of the module's 2-bladed propeller, see Fig. 4.

Since the PSD is not flat, it may not be appropriate to assume that the process noise due to the propeller is white. However, most of this power lies in a frequency band substantially higher than the closed-loop dynamics of the system. It is therefore reasonable to assume a zero-mean white noise spectrum in the frequency range of interest. In general, however, one could approximate an LTI system (or colouring filter) with the same frequency response as the noise, and absorb this system into the plant.

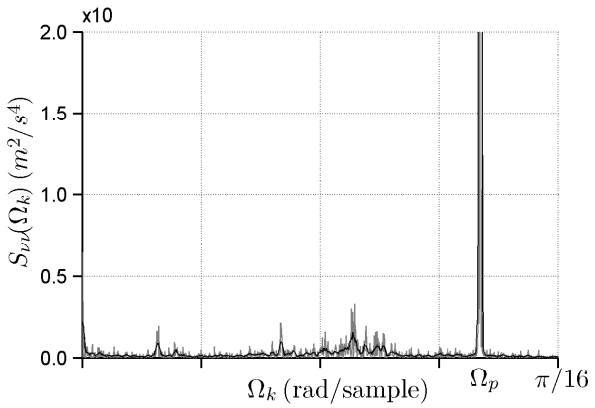

Fig. 4. The PSD function of the noise generated by a propeller at nominal thrust shows that the main portion of power in the frequency domain is contained at $\Omega_p = 0.16$ radians per sample (or $1144\,\mathrm{rad/s}$).

B. Measurement Noise

Experiments have shown that all sensors have a relatively flat PSD over the entire frequency range of interest. We can therefore model our sensors with zero-mean white noise.

Authorized licensed use limited to: ETH BIBLIOTHEK ZURICH. Downloaded on May 15,2024 at 21:25:47 UTC from IEEE Xplore. Restrictions apply.

V. Parameter Optimization

As described in Section III, our controller has two tuning parameters per DOF – a closed-loop natural frequency ω_s and damping ratio ζ_s, where $s \in \{x, y, z, \gamma, \beta, \alpha\}$. These parameters dictate the performance and stability of the DFA. The objective now is, given any configuration, to find parameters that render the system stable and produce the best performance at hover without saturating the actuators (or propeller thrust). The following section describes how the optimization is formulated in terms of the DFA, although this can be generalized to any similar system. For implementation details see extended footnote at the end of this section.

A. Performance Objective

Hover performance of the DFA is defined as a measure of its position (x, y, z) and yaw angle α response as a result of the expected input disturbances. As our metric, we propose to use the $\mathcal{H}_2$ norm $||\cdot||_2$, which quantifies the system's output (or error) power due to an exogenous disturbance with a fixed or bounded power spectrum [14]; roughly speaking, it is the RMS value of the measured response due to white noise input [15]. High performance in this context means minimizing the $\mathcal{H}_2$ norm.

The exogenous disturbance $\mathbf{w} \in \mathbb{R}^{(N+5+7N)}$ consists of process noise $\mathbf{w}_p$ and measurement noise $\mathbf{w}_m$, both of which is modeled as white noise as was described in Section IV. Measurement noise can be divided into outer- and inner-loop sensor noise, respectively denoted $\mathbf{w}_{m_{out}}$ and $\mathbf{w}_{m_{in,i}}$ for each module i. In the case of the DFA, this looks like the following:

$$\mathbf{w} = [\mathbf{w}_p, \mathbf{w}_{m_{out}}, \mathbf{w}_{m_{in,1}}, \ldots, \mathbf{w}_{m_{in,N}}]^T, \tag{17}$$

where

$$\mathbf{w}_p = [w_{p_1}, \ldots, w_{p_N}], \tag{18}$$

$$\mathbf{w}_{m_{out}} = [w_x, w_y, w_\alpha, w_{\dot{x}}, w_{\dot{y}}], \tag{19}$$

$$\mathbf{w}_{m_{in,i}} = [w_{z_i}, w_{\gamma_i}, w_{\beta_i}, w_{\alpha_i}, w_{\dot{z}_i}, w_{\dot{\gamma}_i}, w_{\dot{\beta}_i}]. \tag{20}$$

Let Z_o denote the controlled (or regulated) weighted output response that is to be minimized, and let the input $\mathbf{u} = (a_{d,1}, \ldots, a_{d,N})$ denote the desired (or commanded) control effort of each module. The continuous-time LTI state space description of our *objective* system G_o is given by:

$$\dot{X} = \mathbf{A}X + \mathbf{B}_1\mathbf{w} + \mathbf{B}_2\mathbf{u} \tag{21}$$

$$Y_i = \mathbf{C}_2 X + \mathbf{D}_i\mathbf{w}, \tag{22}$$

$$Z_o = \mathbf{C}_{1,o}X \tag{23}$$

where Y_i is the measured output of module i combined with the measured output of the offboard sensors, and X is the state vector of the DFA consisting of its six DOF $(x, y, z, \gamma, \beta, \alpha)$, their time-derivatives, and the state of each actuator i denoted by a_i,

$$X = [x, y, z, \gamma, \beta, \alpha, \dot{x}, \dot{y}, \dot{z}, \dot{\gamma}, \dot{\beta}, \dot{\alpha}, a_1, \ldots, a_N]^T. \tag{24}$$

The controlled output error signal Z_o is a weighted response of the position (x, y, z) and yaw angle α. It is sensible to set the position weights to the same value; this value is normalized to 1. We are then left with a single free design parameter: the weight of the yaw angle in units of meters per radian. Experiments have shown $\frac{1}{35}$ m/rad to demonstrate reasonable performance, with the deviation in yaw angle α small enough in comparison to the other DOF.

The measured output Y_i is fed into the controller K_i in order to obtain the desired control effort $a_{d,i}$ of module i,

$$a_{d,i} = -K_i Y_i, \tag{25}$$

where

$$K_i = \mathbf{Q}_i \mathbf{M} K_o, \tag{26}$$

and K_o is the static gain matrix of our full state feedback controller, which is a function of ω_s and ζ_s. As previously mentioned, $\mathbf{Q}$ maps the desired closed-loop response of the vehicle to the individual desired control efforts (6); the subscript i in (26) denotes the row of this matrix. Again, the elements of the matrix $\mathbf{M}$ are the vehicle's total mass and principal mass moments of inertia.

If this was a centralized architecture, where each module i used the same measured output Y_i to drive its controller K_i, we would obtain the same expression as (6).

B. Constraints

The optimization must be constrained to ensure that the actuators do not saturate their available control effort. This can be formulated as N number of inequality constraints, where the variance of the output control effort $\sigma^2_{a,i}$ for each module i must fall below a predefined threshold $\sigma^2_{a,max}$. While this does not guarantee that all control signals will not saturate, it does quantify the probability of saturation, which we want to minimize.

The standard deviation $\sigma_{a,i}$ of the output control effort a_i, can be obtained by computing the $\mathcal{H}_2$ norm of the closed-loop *constraint* system G_c consisting of (21), (22), and the weighted output response $Z_{c,i} = \mathbf{C}_{1,c_i}X$ of module i, due to an exogenous disturbance $\mathbf{w}$. It can be shown that the $\mathcal{H}_2$ norm for this single output system is in fact the standard deviation of the control effort due to white noise input.

Authorized licensed use limited to: ETH BIBLIOTHEK ZURICH. Downloaded on May 15,2024 at 21:25:47 UTC from IEEE Xplore. Restrictions apply.

The state transition matrix $\mathbf{A}$ for nominal flight dynamics including motor dynamics, as described in Section II, is given here as

$$\mathbf{A} = \begin{bmatrix} 0_{6\times 6} & I_{6\times 6} & 0_{6\times N} \\ \mathcal{A} & 0_{2\times 6} & 0_{2\times N} \\ 0_{4\times 6} & 0_{4\times 6} & \mathbf{M}^{-1}\mathbf{P}^T \\ 0_{N\times 6} & 0_{N\times 6} & -\omega_m I_{N\times N} \end{bmatrix}, \tag{27}$$

where

$$\mathcal{A} = \left[\begin{array}{c|ccc} 0_{2\times 3} & 0 & g & 0 \\ & -g & 0 & 0 \end{array}\right]. \tag{28}$$

Part of the exogenous white noise disturbance $\mathbf{w}$ is scaled in the matrix $\mathbf{B}_1$ by the standard deviation of noise caused by the propellers σ_m. This is also scaled by its mass and configuration matrix $\mathbf{M}^{-1}\mathbf{P}^T$:

$$\mathbf{B}_1 = \left[\begin{array}{c|c} 0_{8\times N} & \\ \sigma_m \mathbf{M}^{-1}\mathbf{P}^T & 0_{(12+N)\times(5+7N)} \\ 0_{N\times N} & \end{array}\right]. \tag{29}$$

The desired control effort $\mathbf{u}$ is fed into the matrix $\mathbf{B}_2$, which scales these values according to the motor dynamics,

$$\mathbf{B}_2 = \begin{bmatrix} 0_{12\times N} \\ \omega_m I_{N\times N} \end{bmatrix}. \tag{30}$$

The output response Z_o, which is to be minimized, is a function of our output states (x, y, z, α) weighted by the matrix $\mathbf{C}_{1,o}$. The position (x, y, z) weights are normalized to 1, leaving the weight on yaw angle q_α as the only free parameter:

$$\mathbf{C}_{1,o} = \left[\begin{array}{cc|c} \mathbf{I}_{3\times 3} & 0_{3\times 3} & 0_{4\times(6+N)} \\ 0_{1\times 3} & \mathbf{c} & \end{array}\right], \tag{31}$$

where $\mathbf{c} = [0, 0, q_a]$. The control effort $Z_{c,i}$ of module i is simply the unscaled value of the control effort a_i:

$$\mathbf{C}_{1,c_i} = \left[\begin{array}{c|cc} 0_{1\times 12} & \mathcal{C}_{1\times N} \end{array}\right], \tag{32}$$

where the elements c_j of $\mathcal{C}$ are:

$$c_j = \begin{cases} 1 & \text{if } j = i \\ 0 & \text{otherwise.} \end{cases} \tag{33}$$

The output Y is a measure of all the states, not including actuator states:

$$\mathbf{C}_2 = \left[\begin{array}{c|c} I_{12\times 12} & 0_{12\times N} \end{array}\right]. \tag{34}$$

The exogenous white noise disturbance $\mathbf{w}$ due to measurement noise is scaled appropriately in the matrix $\mathbf{D}_i$ of module i,

$$\mathbf{D}_i = \begin{bmatrix} 0_{12\times N} & \mathcal{D}_{out} & \mathcal{D}_{in,1} & \dots & \mathcal{D}_{in,N} \end{bmatrix}. \tag{35}$$

where $\mathcal{D}_{out}$ contains the standard deviations $\sigma_{x,y,\alpha,\dot{x},\dot{y}}$ of the states measured by the outer-loop sensors,

$$\mathcal{D}_{out} = \left[\begin{array}{cc} \mathrm{diag}(\sigma_x, \sigma_y) & 0_{2\times 3} \\ 0_{3\times 2} & 0_{3\times 3} \\ 0_{3\times 2} & \mathrm{diag}(\sigma_\alpha, \sigma_{\dot{x}}, \sigma_{\dot{y}}) \\ \hline & 0_{(4+N)\times 5} \end{array}\right], \tag{36}$$

and the elements $\mathcal{D}_{in,j}$ of $\mathbf{D}_i$ contain the standard deviations $\sigma_{z,\gamma,\beta,\dot{z},\dot{\gamma},\dot{\beta},\dot{\alpha}}$ of the states measured by the inner-loop sensors of module j

$$\mathcal{D}_{in,j} = \begin{cases} 0_{12\times 7} & \text{if } j = i \\ \mathcal{D}_{in} & \text{otherwise,} \end{cases}$$

where

$$\mathcal{D}_{in} = \begin{bmatrix} 0_{2\times 3} & 0_{2\times 4} \\ \mathrm{diag}(\sigma_z, \sigma_\gamma, \sigma_\beta) & 0_{3\times 4} \\ 0_{3\times 3} & 0_{3\times 4} \\ 0_{4\times 3} & \mathrm{diag}(\sigma_{\dot{z}}, \sigma_{\dot{\gamma}}, \sigma_{\dot{\beta}}, \sigma_{\dot{\alpha}}) \end{bmatrix}. \tag{37}$$

The static gain matrix K_o of the full state feedback controller is given by

$$\mathbf{K}_o = \begin{bmatrix} 0 & 0 & \omega_z^2 & 0 & 0 & 0 & 0 & 0 & 2\omega_z\zeta_z & 0 & 0 & 0 \\ 0 & -\omega_\gamma^2\omega_y^2 & 0 & \omega_\gamma^2 & 0 & 0 & 0 & -\omega_\gamma^2 2\omega_y\zeta_y & 0 & 2\omega_\gamma\zeta_\gamma & 0 & 0 \\ \omega_\beta^2\omega_x^2 & 0 & 0 & 0 & \omega_\beta^2 & 0 & \omega_\beta^2 2\omega_x\zeta_x & 0 & 0 & 0 & 2\omega_\beta\zeta_\beta & 0 \\ 0 & 0 & 0 & 0 & 0 & \omega_\alpha^2 & 0 & 0 & 0 & 0 & 0 & 2\omega_\alpha\zeta_\alpha \end{bmatrix}. \tag{38}$$

VI. Results

A. Performance

The performance of the controller described in the last two sections, which will be referred to from now on as the Decentralized Full State Feedback (FSF) controller, was compared to two other controllers:

1) **Centralized FSF**: This is the FSF controller used in a centralized architecture; a controller aggregates sensor measurements from all the modules in the vehicle, computes an average that is used as state estimate, and commands each module with the appropriate control effort. This is considered to be the best case scenario when using the FSF controller.
2) $\mathcal{H}_2$ **Optimal**: Using the system G_o in a centralized architecture and extending it to include control effort weights, we compute the $\mathcal{H}_2$ optimal controller [16] and adjust the weights on the control effort subject to not saturating the actuators. This is considered to be the benchmark of all control methods considered here.

Using the procedure described in Section V, the optimal control tuning parameters were computed in Matlab for DFA configurations ranging from 4 to 20 modules in size. To provide some statistics, the procedure was implemented on 50 unique and randomly chosen configurations for each configuration size.

The results in Fig. 5(a) show that as the vehicle gets larger, the performance gets better (i.e. $||G_o||_2$ decreases). This is due to the fact that the effect that process noise has on the states decreases proportionally by at least a factor of $\sqrt{N}$. This can be made mathematically precise, but the intuition is that process noise is spatially uncorrelated, smoothing out the effect it has on the measured states as the vehicle gets larger. In Fig. 5(b) one can also see that the optimal parameters vary with N. This is expected, as the effective decrease of process noise acting on the states changes the optimal performance of the system.

As expected, the $\mathcal{H}_2$ optimal controller outperforms the FSF controllers, see Fig. 5(a). In the absolute sense, the difference in performance between the $\mathcal{H}_2$ optimal controller and the FSF controllers is relatively constant for all sizes of the vehicle N. However, note that the comparison made here is not equivalent. The $\mathcal{H}_2$ optimal controller contains dynamics and grows with N, whereas the FSF controller is simply a static gain matrix containing 12 free variables for any configuration size. In spite of this, the decentralized FSF controller provides sufficiently good performance; its time response for a configuration consisting of 4-modules in a quadrotor configuration can be seen in Fig. 6. In comparison to an $\mathcal{H}_2$ optimal controller, which is essentially a black box to the designer, the tuning parameters of the structured

Authorized licensed use limited to: ETH BIBLIOTHEK ZURICH. Downloaded on May 15,2024 at 21:25:47 UTC from IEEE Xplore. Restrictions apply.

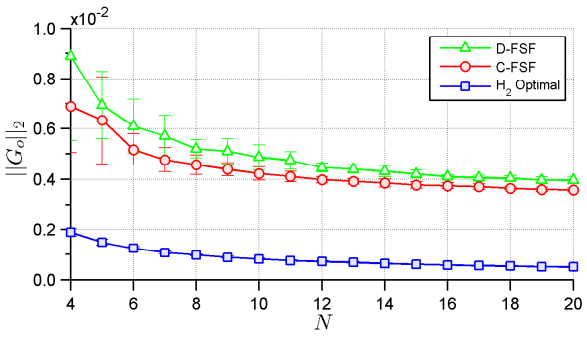

(a) The performance of the DFA improves (decreasing $||G_o||_2$) with increasing size of the vehicle N.

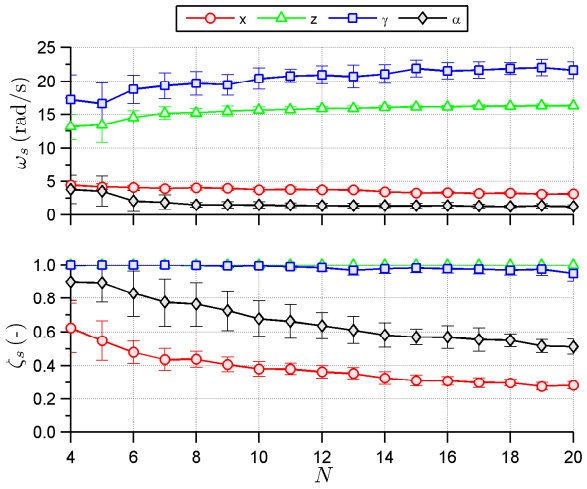

(b) Since the effect that process noise has on the states decreases with the size of the vehicle N, optimal parameters of the system are expected to vary.

Fig. 5. This data set was obtained from 50 unique and randomly chosen configurations for each vehicle size N. Error bars indicate the standard deviation in performance (a) and of the parameters (b) due to different configurations.

FSF controller are physically meaningful and thus provides some intuition about the closed-loop system. Note that the performance of the centralized FSF controller is not much better than the decentralized FSF controller, despite the incurred communication costs.

VII. Configuration to Control Parameter Map

Designing an $\mathcal{H}_2$ optimal controller for every possible configuration is intractable. The parameterized FSF controller offers a straightforward solution for creating a function that maps DFA configurations to their control parameters. These mapped control parameters are not necessarily optimal, but they should perform sufficiently.

A. Configuration Parameterization

The configuration of the DFA can be parameterized by its size N, and the parameters

$$\epsilon_x = \frac{I_x}{I_d} \quad , \quad \epsilon_y = \frac{I_y}{I_d}, \tag{39}$$

which capture the mass distribution of the vehicle, where I_d is the mass moment of inertia for a solid disk of equivalent

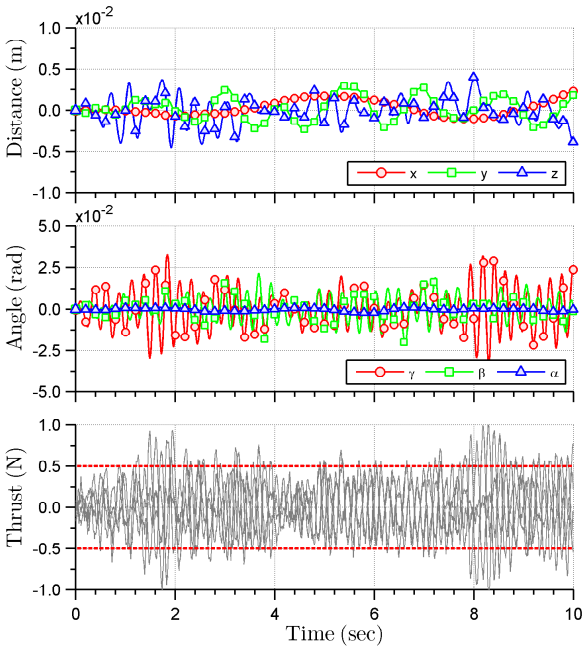

Fig. 6. The simulated time response for each DOF using the decentralized FSF controller on a 4-module vehicle in a quadrotor configuration shows good performance during nominal operating conditions while meeting the actuator constraints ($\pm$ 1 N) shown in the lower plot. The dotted lines in the lower plot denotes one standard deviation of the allowed control effort $\sigma_{a,max}$.

volume. As such, the values (ϵ_x, ϵ_y) are expected to be close to 1 for a disk-like configuration.

As expected, for a fixed sized vehicle, the optimal parameters (ω_s, ζ_s) change with its configuration, see Fig. 7. The change is most pronounced in the roll γ and pitch β DOF. Consider an example where the vehicle is elongated along its x-axis and short along its y-axis (i.e. large ϵ_x and small ϵ_y). The controller acting along the short axis will need to react more quickly than that around the long axis due to faster natural dynamics of the system.

B. Map

By fixing the size N of the DFA, we can obtain a set of functions (two for each DOF) that maps the configuration parameters (ϵ_x, ϵ_y) to each of the control design parameters (ω_s, ζ_s). One way to obtain such a map is by surface-fitting a polynomial function via least-squares to a data set of optimal control parameters.

The set of surface-fitted functions can then be tested against an independent data set. We can assess the fit by comparing the $\mathcal{H}_2$ norm of the system that uses the fitted parameters to one that uses the optimal parameters. This error can be minimized by independently adjusting the degree of each polynomial function in order to find those that fit best.

Results have shown that for a data set containing 100 unique configurations of a 10-module vehicle, the best performance was obtained by using a single degree polynomial function of the form

$$\omega_s = a_{\omega_s}\epsilon_x + b_{\omega_s}\epsilon_y + c_{\omega_s}, \tag{40}$$

$$\zeta_s = a_{\zeta_s}\epsilon_x + b_{\zeta_s}\epsilon_y + c_{\zeta_s}, \tag{41}$$

Authorized licensed use limited to: ETH BIBLIOTHEK ZURICH. Downloaded on May 15,2024 at 21:25:47 UTC from IEEE Xplore. Restrictions apply.

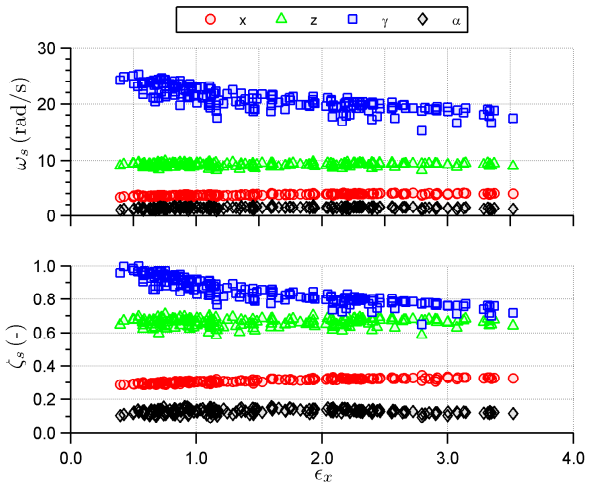

Fig. 7. The closed-loop optimal parameters for 200 unique DFA configurations of size 10, which spans the entire spectrum of configurations. Most parameters maintain a relatively constant value except in the DOF of roll γ and pitch β (not shown for clarity), which decreases and increases, respectively, for large values of ϵ_x.

where $(a_{\omega_s}, b_{\omega_s}, c_{\omega_s})$ and $(a_{\zeta_s}, b_{\zeta_s}, c_{\zeta_s})$ are all constant parameters, see Fig. 8.

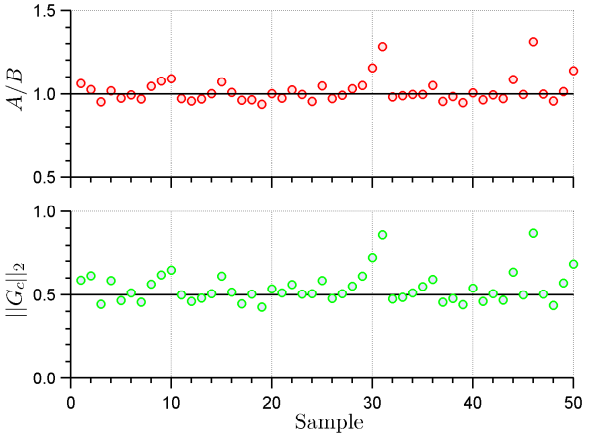

Fig. 8. The upper subplot shows the performance ratio A/B, where A is the performance of a 10-module vehicle using fitted parameters and B is the performance of the same vehicle using optimal FSF parameters. For some configurations, the performance ratio exceeds 1, indicating that the fitted parameters *perform better* than the *optimal* parameters. This is because these fitted parameters violate the actuator constraints, as seen in the lower subplot where points lie above $||G_c||_2 > 0.5$.

To obtain reasonable control parameters for flying an arbitrary configuration, all that is needed is the size N and configuration parameters (ϵ_x, ϵ_y), the latter being a function of each module's position and mass. This information can be communicated across the vehicle before taking flight.

VIII. Conclusions

We have described an approach for controlling a modular flying vehicle of arbitrary configuration, namely the Distributed Flight Array. Our methodology consists of computing the control parameters of a decentralized parameterized controller that minimizes the $\mathcal{H}_2$ norm of the system subject to its physical constraints. This controller was compared to traditional methods, and although it demonstrates slightly lower performance, it provides much more intuition into the behaviour of the system. More importantly, its simplicity enables one to map the configuration space of the vehicle to its control parameter space using a polynomial function. Finding reasonable control tuning parameters for any flight-feasible configuration is now nothing more than solving a few first-order equations.

This methodology could easily be extended to include different noise models than those already considered. It would also be straightforward to consider different performance metrics, for example the $\mathcal{H}_\infty$ norm, which would capture the worse case disturbances of a system subject to an exogenous input.

Future work will include verifying the performance of the real system with our simulated model, and to do this for various configurations.

Acknowledgements

This work is funded by the Swiss National Science Foundation (SNSF).

References

[1] M. Yim, D. Duff, and K. Roufas, "PolyBot: a Modular Reconfigurable Robot," in *IEEE International Conference on Robotics and Automation*, vol. 1. IEEE, 2000, pp. 514–520.

[2] S. Murata, E. Yoshida, A. Kamimura, H. Kurokawa, K. Tomita, and S. Kokaji, "M-TRAN: Self-Reconfigurable Modular Robotic System," *IEEE/ASME Transactions on Mechatronics*, vol. 7, no. 4, pp. 431–441, 2002.

[3] B. Salemi, M. Moll, and W. Shen, "SUPERBOT: A Deployable, Multi-Functional, and Modular Self-Reconfigurable Robotic System," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2006, pp. 3636–3641.

[4] M. Yim, W. Shen, B. Salemi, D. Rus, M. Moll, H. Lipson, E. Klavins, and G. Chirikjian, "Modular Self-Reconfigurable Robot Systems [Grand Challenges of Robotics]," *IEEE Robotics & Automation Magazine*, vol. 14, no. 1, pp. 43–52, 2007.

[5] K. Gilpin and D. Rus, "Modular Robot Systems," *IEEE Robotics & Automation Magazine*, vol. 17, no. 3, pp. 38–55, 2010.

[6] P. Varshavskaya, L. Kaelbling, and D. Rus, "Learning Distributed Control for Modular Robots," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 3. IEEE, 2004, pp. 2648–2653.

[7] R. Fitch and Z. Butler, "Million Module March: Scalable Locomotion for Large Self-Reconfiguring Robots," *The International Journal of Robotics Research*, vol. 27, no. 3-4, pp. 331–343, 2008.

[8] R. Oung and R. D'Andrea, "The Distributed Flight Array," *Mechatronics*, vol. 21, no. 6, pp. 908–917, Sept. 2011.

[9] M. Hehn and R. D'Andrea, "A Flying Inverted Pendulum," in *IEEE International Conference on Robotics and Automation*. IEEE, 2011, pp. 763–770.

[10] N. Chaturvedi, A. Sanyal, and N. McClamroch, "Rigid-Body Attitude Control," *IEEE Control Systems*, vol. 31, no. 3, pp. 30–51, 2011.

[11] M. Kriegleder, R. Oung, and R. D'Andrea, "Distributed Altitude and Attitude Estimation from Multiple Range Measurements," in *IEEE/RSJ International Conference of Intelligent Robots and Systems*, 2012.

[12] Vicon MX homepage. Last Accessed: Mar. 6, 2012. [Online]. Available: http://www.vicon.com/

[13] F. Hoblit, *Gust Loads on Aircraft: Concepts and Applications*. AIAA, 1988.

[14] B. Anderson and J. Moore, *Optimal Control: Linear Quadratic Methods*. Prentice Hall Englewood Cliffs, NJ, 1990, vol. 1.

[15] S. Skogestad and I. Postlethwaite, *Multivariable Feedback Control: Analysis and Design*. Wiley, 2007, vol. 2.

[16] J. Doyle, K. Glover, P. Khargonekar, and B. Francis, "State-Space Solutions to Standard H2 and H Control Problems," *IEEE Transactions on Automatic Control*, vol. 34, no. 8, pp. 831–847, 1989.

Authorized licensed use limited to: ETH BIBLIOTHEK ZURICH. Downloaded on May 15,2024 at 21:25:47 UTC from IEEE Xplore. Restrictions apply.