

MIT Open Access Articles

Robot Co-design: Beyond the Monotone Case

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Carlone, Luca and Pincioli, Carlo. 2019. "Robot Co-design: Beyond the Monotone Case." Proceedings - IEEE International Conference on Robotics and Automation, 2019-May.

As Published: <http://dx.doi.org/10.1109/icra.2019.8793926>

Publisher: Institute of Electrical and Electronics Engineers (IEEE)

Persistent URL: <https://hdl.handle.net/1721.1/137260>

Version: Original manuscript: author's manuscript prior to formal peer review

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



Robot Co-design: Beyond the Monotone Case

Luca Carlone Carlo Pincirolì

Abstract—Recent advances in 3D printing and manufacturing of miniaturized robotic hardware and computing are paving the way to build inexpensive and disposable robots. This will have a large impact on several applications including scientific discovery (e.g., hurricane monitoring), search-and-rescue (e.g., operation in confined spaces), and entertainment (e.g., nano drones). The need for inexpensive and task-specific robots clashes with the current practice, where human experts are in charge of designing hardware and software aspects of the robotic platform. This makes the robot design process expensive and time consuming, and ultimately unsuitable for small-volumes low-cost applications. This paper considers the *computational robot co-design* problem, which aims to create an automatic algorithm that selects the best robotic modules (sensing, actuation, computing) in order to maximize the performance on a task, while satisfying given specifications (e.g., maximum cost of the resulting design). We propose a binary optimization formulation of the co-design problem and show that such formulation generalizes previous work based on strong modeling assumptions. We show that the proposed formulation can solve relatively large co-design problems in seconds and with minimal human intervention. We demonstrate the proposed approach in two applications: the co-design of an autonomous drone racing platform and the co-design of a multi-robot system.

Index Terms—Mechanism Design, Multi-Robot Systems, Aerial Systems: Perception and Autonomy.

I. INTRODUCTION

Recent advances in sensor design and rapid prototyping are enabling the manufacturing of low-cost robots with advanced sensing and perception capabilities. For instance, one can implement high-precision visual-inertial navigation with inexpensive camera and MEMS inertial measurement units [1]. Similarly, modern embedded CPU-GPU [2] offer high-performance computing in a compact form-factor and at a relatively affordable cost. These trends, together with the availability of low-cost micro-motors are enabling fast and cheap production of robotic platforms. In these cases, the most expensive resource becomes the effort of the expert human designers who are in charge of designing all the aspects of the robotic platform, including hardware and software. While this solution is still acceptable in the case where low-cost robots must be produced in volumes (e.g., vacuum cleaning robots), it may not be desirable when only few robots need to be deployed. Consider, for instance, the design of a drone for hurricane monitoring [3]: the drone must be disposable, hence inexpensive, and it is not typically produced in volumes. In other contexts, one may need to devise a design quickly, in order to create a robot to be deployed in a time-sensitive




Fig. 1. (a) We apply the proposed co-design framework to the design of the modules (motors, frame, computing, camera, battery pack) of an autonomous racing drone, (ii) Force diagram during forward motion (side view).

mission. For instance, one may need to design a search-and-rescue robot tailored to a specific mission (e.g., search for survivors in a narrow cave with a given size of the entry point). Finally, human design does not necessarily lead to optimal solutions. Designers typically consider different robotics modules in isolation in order to tame the design complexity, and such decoupling usually leads to suboptimal performance.

These reasons motivate us to investigate *computational robot co-design*, which aims to create an automatic algorithm that selects the best robotic modules (sensing, actuation, computing) in order to maximize the performance on a task, while satisfying given specifications (e.g., maximum cost of the resulting design). Here the term “computational” refers to the fact that the design techniques can be implemented on a machine, and require minimal human intervention. Moreover, the term “co-design” refers to the attempt to consider the robotic system as a whole, rather than (sub-optimally) decoupling the design of each module.

Related Work. The problem of co-design in robotics touches a wide span of research topics. In the most general sense, it can be considered as the problem of designing both the mind (software) and the body (hardware) of the robot at the same time. A few seminal works tackle this problem from an evolutionary standpoint [4], [5], [6], with approaches that aim to either optimize specific behaviors (e.g., walking) or to explore catalogs of possible solutions to a general problem (e.g., locomotion). While related to the problem of computational co-design as discussed above, these approaches typically abstract away several practical aspects of hardware design, such as power consumption and material selection.

At the opposite side of the spectrum, a large body of work on software/hardware design exists in the embedded systems community, e.g. FPGA and ASIC design [1]. In this

L. Carlone is with the Laboratory for Information & Decision Systems (LIDS), Massachusetts Institute of Technology, USA, lcarlone@mit.edu

C. Pincirolì is with Robotics Engineering and the Department of Computer Science, Worcester Polytechnic Institute, USA, cpincirolì@wpi.edu

This work was partially funded by ARL DCIST CRA W911NF-17-2-0181.

domain, the focus is on the synthesis of practical solutions that explicitly consider low-level aspects such as power optimization and real-time scheduling in electronic devices. Since Chinook [7], a highly influential early approach, the field of software/hardware co-design has flourished, reaching relevant applications in safety-critical fields such as the aerospace [8] and automotive industry. The inception of domain-specific languages such as the *Architecture Analysis and Design Language* (AADL) [9] has enabled a systems-of-systems approach in which co-design can be coupled with a wide array of tools for property verification, both at the software and hardware level. These tools, however, are thought to support *human* design.

The field of modular robotics attempts to bridge low-level embedded system design with high-level functionality specifications. In modular robotics, the challenge is to create a diverse set of composable, programmable modules that can be used to form robots with different capabilities. The recent, rapid development of 3D printing is offering increasingly powerful tools to streamline the automatic design of robotic platforms, and will arguably foster the development of modular robotics. In a seminal paper, Mehta *et al.* [10] demonstrate the feasibility of a procedure that takes a high-level human-defined specification of a robot and outputs the 3D description files of the components, along with instructions to assemble the robot, manufacture the electronics, and automatically generate the control firmware. To the best of our knowledge, no work exists in the automatic generation of specifications that can be input to this system.

A core issue in computational co-design is the conception of a suitable formalism to express the design problem. Ideally, an effective formalism for automatic design should combine adequate granularity along with guarantees of correctness. Evolutionary methods [4], [5], [6] combine a low-level parametric representation, such as rigid bodies connected by joints or voxel-based artifacts, with neural network-based control. The main issue of this representation is the amount of detail involved in the optimization process, which limits scalability, and the absence of correctness guarantees. Other formalisms closer to the human design process, such as the one from Mehta *et al.* [10] and AADL [9], achieve expressive power by naturally incorporating modularity. Both languages represent a co-design problem as a tree composed of heterogeneous nodes which represent hardware and software aspects of a robot. AADL, in particular, offers a set of tools specifically designed for early analysis of candidate solutions, but without support for automatic exploration of the solution space. Correct-by-design compositional approaches have also been the focus of recent studies in control theory. For instance, Hussien *et al.* [11] proposed an automatic method to decompose a large control problem into a cascade of simpler problem, under the assumption that the latter are feedback linearizable, while Kim *et al.* [12] devised a method to arbitrarily construct system abstractions from simpler, well-posed components. Lastly, Censi [13], [14] proposed a powerful co-design approach that enables the automatic generation of solutions from a specification that includes recursive constraints. Censi's approach assumes that a monotone mapping between resources and performance exists, thus casting computational co-design as

a fixed-point problem in the solution space.

Contribution. This paper builds on recent work by Censi [13], [14] and pushes the boundary of computational co-design by relaxing its monotonicity assumption. Monotonicity implies that investing more resources leads to an increase in the performance. However, it is easy to find interesting examples for which this property does not hold. For instance, in the design of a drone frame, the fact that we increase the size of the frame (resource) does not imply that the drone will fly faster (performance). Similarly, in a collaborative transportation problem, deploying a larger number of robots (resources) does not necessarily lead to an increase in the overall performance of the system [15].

This paper presents a more general computational co-design approach. Our problem formulation is presented in Sec. II, where we discuss the characterization of the design space and classify the design specifications in terms of *system-level performance*, *system-level constraints*, and *intrinsic constraints*. Sec. II-C rephrases the resulting co-design problem in terms of a binary optimization problem, where each binary variable indicates whether a given component is chosen as part of the design or not. Sec. III discusses in which case we can expect to be able to solve the binary optimization problem using off-the-shelf optimization tools; in general, binary optimization is intractable, but there exist several algorithms and implementations that are able to solve moderate-sized problems efficiently as long as some property (e.g., linearity) holds. We show that our approach allows recasting several functions that are nonlinear in the properties (or *features*) of the robotic modules as linear functions. The proposed optimization framework also makes it possible to model other constraints found in practice, such as *compatibility* constraints (e.g., we cannot use a LIDAR-based algorithm to process data from a monocular camera), which was beyond reach for existing methods [13].

We conclude the paper by presenting two applications of the proposed co-design approach to the design of an autonomous drone racing platform and a multi-robot system. These examples are presented in Sec. IV. Drone design is an interesting topic per se, and has already received attention in the literature [16], [1], which however still lacks a computational design approach. New and untapped challenges arise in the co-design of multi-robot systems. We are not aware of approaches that specifically tackle this problem; however, a vast literature exist on the automatic generation of control systems for robots swarms, including evolutionary methods [17] and automatic generation of composable control structures [18].

II. GENERAL ROBOT CO-DESIGN

The co-design problem consists in jointly designing robot software and hardware components to maximize task-dependent performance metrics (e.g., endurance, agility, payload) subject to constraints on the available resources (e.g., cost). The complexity of the problems stems from the fact that a robotic system involves an intertwining of modules. Each module contributes to the overall performance of the system while potentially consuming resources. In this paper, we consider the realistic case in which we have to choose

each module (e.g., motor, embedded computer, camera) in the robotic system from a given catalog, and we formulate the co-design as the combinatorial selection of modules that maximize performance while satisfying given system-level and module-level constraints. In Sec. II-A we introduce our abstraction of the modules forming the overall robotic system. In Sec. II-B we describe the interactions among the modules and how they contribute to the overall system performance and resource utilization. In Sec. II-C we state the co-design problem as a binary optimization problem.

A. Modules, Catalogs, Features

We consider the case in which the robotic system comprises a given set of *modules* \mathcal{M} . The modules may include, for instance, the actuators, the sensors, the computational board, the perception algorithms, the control algorithms, the planning algorithms, etc. For each module $i \in \mathcal{M}$, we have a *catalog* \mathcal{C}_i of potential choices: for instance, we can purchase different motor models, or we can utilize different approaches and implementations of a planning algorithm.

Design vector. The goal of the co-design is to select an element $j \in \mathcal{C}_i$ (catalog for module i), for each module $i \in \mathcal{M}$. We can represent this selection using a binary vector \mathbf{x}_i for each module, where the j -th entry of \mathbf{x}_i is 1 if we select the j -th element in the catalog \mathcal{C}_i or zero otherwise. Clearly, $\mathbf{x}_i \in \{0, 1\}^{|\mathcal{C}_i|}$, where $|\mathcal{C}_i|$ is the cardinality (number of elements) of the catalog. Therefore, the design is fully specified by the *design vector* \mathbf{x} obtained by stacking \mathbf{x}_i for each $i \in \mathcal{M}$. The design vector has size $N \doteq \sum_{i \in \mathcal{M}} |\mathcal{C}_i|$, which we refer to as the *dimension of the design space*.

Feature matrix. Each module has a number of *features* describing the technical specifications of the module. For example, the features of a motor may include the cost, torque, weight, maximum speed, power consumption, size of the motor, among other technical data. In general, the features are a list of properties one would find in the datasheet of a component. Similarly, for an algorithm, the set of features may include information about the expected performance and computational cost of the algorithm.

Clearly, each element in the catalog of module i (e.g., different motor models) will have different values of each feature. We can thus succinctly describe the list of features for each element in the catalog of module i as a *feature matrix* \mathbf{F}_i , where each row correspond to a given feature, and different columns correspond to different elements in the catalog. For instance, in a toy problem, we can have the following feature matrix for the motor module:

$$\mathbf{F}_m = \begin{bmatrix} W^{(1)} & W^{(2)} & W^{(3)} \\ V^{(1)} & V^{(2)} & V^{(3)} \\ A^{(1)} & A^{(2)} & A^{(3)} \\ C^{(1)} & C^{(2)} & C^{(3)} \\ T^{(1)} & T^{(2)} & T^{(3)} \end{bmatrix} \begin{array}{l} \leftarrow \text{weight} \\ \leftarrow \text{voltage} \\ \leftarrow \text{current} \\ \leftarrow \text{cost} \\ \leftarrow \text{torque} \end{array} \quad (1)$$

where the j -th column describes the features of the j element in the motor catalog. In practice, the values in the matrix are known from datasheet or from prior experiments.

Remark 1 (Beyond monotonicity): A way to relate our approach (describing modules via catalogs and feature matrices)

to the one of Censi [13] is as follows. Censi splits what we call *features* into resources and functionalities (imagine for instance, that the first n_r rows or (1) are labeled as “resources” and the last n_p as “functionalities”). Then, Censi assumes that the columns of (1) satisfy the monotonicity property, i.e., choices of components (columns) leading to better functionalities require more resources. This is not necessarily true in practice: if we design a drone system, choosing a larger frame (i.e., consuming more resources) does not necessarily imply that our drone will fly faster (i.e., better functionality). We go beyond [13] with two main innovations. First, the entries in our feature matrix (1) are completely arbitrary, hence relaxing the monotonicity assumption. We will further highlight the importance of relaxing monotonicity for multi-robot design in Sec. IV-B. Second, while resources and functionalities of each module may be problem-dependent, feature matrices are *design-agnostic*, as discussed in Remark 2 below.

Remark 2 (Feature matrices are design-agnostic): The notion of “module” and “feature matrix” are agnostic of the co-design problem, which is only defined in Sec. II-C. Indeed, the notion of “features” is general—the same modules can be used in any co-design problem involving that module. This is in sharp contrast with [13], where the definition of “resources” and “functionalities” (or “performance”) for each module depends on the interactions between modules, hence it is problem-dependent. For instance: Fig. 21 in [13] classifies the battery capacity as a performance metric, but in other problems the battery capacity can be a resource. Our choice to create an intermediate abstraction, the *feature matrix*, resolves this dependence, enabling re-usability of modules across problems.

B. Design Specifications: Performance and Constraints

In a robotic system, the different modules interact to contribute to the overall performance of the system and potentially consume resources. In particular, both the overall system as well as each module require a minimum amount of resources to operate properly, thus imposing *constraints* on the design.

In this context, we distinguish three main *design specifications*: *system-level performance*, *system-level constraints*, and *module-level constraints*. Intuitively, system-level performance defines a set of metrics the co-design has to maximize, while system-level (resp. module-level) constraints specify a set of constraints that need to be satisfied for the overall system (resp. each module) to operate correctly. These design specifications, which we discuss in more detail below, together with the catalogs of the modules we want to design, fully specify the co-design problem (Section II-C).

Explicit specifications: system-level performance and constraints. As an input to the co-design process, the user provides a set of performance metrics the design must maximize, as well as a set of constraints the *overall system* must satisfy. These specifications, which we refer to as *explicit specifications*, are at the system level, in the sense that they describe the task that the system must perform.

System-level Performance. The *system-level performance* is a vector-value function of a choice of components. Recalling

that a design is fully characterized by the design vector \mathbf{x} , the system-level performance is a function:

$$\text{SP}(\mathbf{x}) : \{0, 1\}^N \mapsto \mathbb{R}^{N_p}, \quad (2)$$

where N is the dimension of the design space and N_p is the number of performance metrics the user specifies. To clarify (2), let us consider a simple example, in which the design must maximize the torque of the wheel motors of an autonomous race car. In this case, the system-level performance is described by:

$$\text{SP}(\mathbf{x}) = 4 [\mathbf{F}_m]_T \mathbf{x}_m \quad (3)$$

In (3), \mathbf{F}_m is the motor feature matrix in (1), \mathbf{x}_m is the design vector for the motor module, and $[\mathbf{F}_m]_T$ extracts the row of \mathbf{F}_m corresponding to the motor torque. The factor “4” captures the fact that, for simplicity, we assumed four wheels mounting identical motors. We remark that the (linear) operator $[\cdot]_T$ selects the row corresponding to a specific feature from the feature matrix, while the multiplication by \mathbf{x}_m has the effect of selecting a single column (i.e., choosing a motor model) from \mathbf{F}_m , due to the fact that \mathbf{x}_m has a single non-zero element.

System-level Constraints. The *system-level constraints* are (scalar) equality or inequality constraints, describing hard requirements on the desired behavior of the system or constraints on the resources that can be used for the design. We express system-level constraints involving the design vector \mathbf{x} :

$$\text{SC}(\mathbf{x})_k \leq 0 \quad \text{for } k \in \mathcal{SC}_{\leq} \quad (4)$$

$$\text{SC}(\mathbf{x})_{k'} = 0 \quad \text{for } k' \in \mathcal{SC}_{=} \quad (5)$$

where \mathcal{SC}_{\leq} and $\mathcal{SC}_{=}$ are the sets of inequality and equality system-level constraints, respectively. For instance, we can have an upper bound on the overall cost of the design:

$$\sum_{i \in \mathcal{M}} [\mathbf{F}_i]_C \mathbf{x}_i \leq \text{budget} \Rightarrow \text{SC}(\mathbf{x})_k = \sum_{i \in \mathcal{M}} [\mathbf{F}_i]_C \mathbf{x}_i - \text{budget} \quad (6)$$

where, as before, the linear operator $[\cdot]_C$ extracts the cost from the feature matrix of each module i . Note that the “budget” must be provided by the user and it is specific to the design instance, hence (6) is a system-level constraint.

An example of system-level equality constraint is the case in which the user wants to consider only a subset of elements in the catalog (e.g., within a catalog of motors, only two motors are available in-house). For instance, for the choice of module i to be restricted to the subset $\mathcal{S}_i \subset \mathcal{C}_i$, the user can add the following system-level equality constraint:

$$\sum_{j \in \mathcal{S}_i} [\mathbf{x}_i]_j = 1 \quad \Rightarrow \quad \text{SC}(\mathbf{x})_{k'} = \sum_{j \in \mathcal{S}_i} [\mathbf{x}_i]_j - 1 \quad (7)$$

where again $[\mathbf{x}_i]_j$ selects the j -th element of \mathbf{x}_i ; the constraint (7) enforces one element in \mathcal{S}_i to be chosen, due to the binary nature of the vector \mathbf{x}_i .

Implicit specifications: module-level constraints. The user must provide explicit specifications to describe *what* the robotic system is required to do within *which* operational constraints. On the other hand, *implicit* specifications are transparent (or uninteresting) to the user and are only needed to

guarantee that each module has sufficient resources to function as expected.

Similarly to the system-level performance, we express module-level (implicit) constraints as inequality or equality constraints involving the design vector \mathbf{x} :

$$\text{IC}(\mathbf{x})_k \leq 0 \quad \text{for } k \in \mathcal{IC}_{\leq} \quad (8)$$

$$\text{IC}(\mathbf{x})_{k'} = 0 \quad \text{for } k' \in \mathcal{IC}_{=} \quad (9)$$

where \mathcal{IC}_{\leq} and $\mathcal{IC}_{=}$ are the sets of inequality and equality module-level constraints, respectively. While the mathematical nature of the system-level and module-level constraints is similar, we believe it makes sense to distinguish them, since the user has control over system-level constraints (e.g., to increase the design budget in (6)), while he/she typically does not have control over the implicit constraints.

For instance, the user cannot change the fact that, for the system to function properly, the onboard battery “b” has to provide enough power for all the active modules (say, motors “m”, sensors “s”, and computing “c”), which indeed translates into an implicit inequality constraint:

$$\begin{aligned} A_b V_b &\geq A_m V_m + A_s V_s + A_c V_c \quad \Rightarrow \\ \text{IC}(\mathbf{x})_k &= \sum_{i \in \{m, s, c\}} ([\mathbf{F}_i]_A \mathbf{x}_i)([\mathbf{F}_i]_V \mathbf{x}_i) - ([\mathbf{F}_b]_A \mathbf{x}_b)([\mathbf{F}_b]_V \mathbf{x}_b) \end{aligned} \quad (10)$$

where A_i and V_i are the current and voltage at module i , and \mathbf{F}_i is again the feature matrix for module i .

Implicit constraints are also useful to model *compatibility constraints*, which, again, the user is not typically free to alter. For instance, we may want to model the fact that we cannot run a certain algorithm (e.g., designed for FPGA) on a certain hardware (e.g., CPU), or we cannot use a LIDAR-based signal processing front-end to process data from a monocular camera. Compatibility constraints can be expressed as follows, for a pair of modules “a” and “b”:

$$\begin{aligned} [\mathbf{x}_a]_j &\leq \sum_{j' \in \mathcal{S}_j} [\mathbf{x}_b]_{j'} \quad (\text{compatibility}) \\ [\mathbf{x}_a]_j &\leq 1 - \sum_{j' \in \mathcal{S}_j} [\mathbf{x}_b]_{j'} \quad (\text{incompatibility}) \end{aligned} \quad (11)$$

The first inequality in (11) imposes that, when the j -th element in the catalog of module “a” is selected, we can only choose the module b from the subset \mathcal{S}_j (subset of *compatible* modules in the catalog of b). The second inequality in (11) imposes that, when the j -th element in the catalog of module “a” is selected, then we cannot choose the module b from the subset \mathcal{S}_j (subset of *incompatible* modules in the catalog).

C. Co-design

The co-design problem can be now stated as follows.

Definition 3 (Robot Co-design): Given the catalogs (and the corresponding feature matrices \mathbf{F}_i) for each module $i \in \mathcal{M}$ to be designed, given the system-level performance function $\text{SP}(\mathbf{x})$ and the set of system-level inequality and equality constraints ($\mathcal{SC}_{\leq}, \mathcal{SC}_{=}$), as well as the set of module-level

inequality and equality constraints ($\mathcal{IC}_{\leq}, \mathcal{IC}_{=}$), robot co-design searches for the choice of modules \mathcal{M} that maximizes the system-level performance, while satisfying the constraints:

$$\begin{aligned} & \max_{\mathbf{x} \in \mathcal{X}} \quad \text{SP}(\mathbf{x}) \\ & \text{subject to} \quad \text{SC}(\mathbf{x})_k \leq 0 \quad \text{for } k \in \mathcal{SC}_{\leq} \\ & \quad \text{SC}(\mathbf{x})_{k'} = 0 \quad \text{for } k' \in \mathcal{SC}_{=} \\ & \quad \text{IC}(\mathbf{x})_k \leq 0 \quad \text{for } k \in \mathcal{IC}_{\leq} \\ & \quad \text{IC}(\mathbf{x})_{k'} = 0 \quad \text{for } k' \in \mathcal{IC}_{=} \end{aligned} \quad (12)$$

where \mathcal{X} is the set of binary vectors that correspond to unique choices of each module (mathematically: $\mathcal{X} \doteq \{0, 1\}^N : \sum_{j=1}^{|\mathcal{C}_i|} [\mathbf{x}_i]_j = 1 \text{ for each module } i \in \mathcal{M}\}$).

Problem (12) is a *binary* optimization problem, since the vector-variable \mathbf{x} has binary entries. The formulation does not take any assumption on the nature of the functions involved in the objective function and the constraints. Our formulation is this *general* and it *does not assume monotonicity*: indeed, by introducing the notion of “feature matrix”, we circumvented the problem of reasoning in terms of resources and functionality of each module. As we will see, this framework allows one to solve non-monotonic problems (see Sec. IV).

Remark 4 (Total Order): In general, the optimization problem (12) is a multi-objective maximization, since the objective is vector-valued. While a major concern in [13] was how to deal with partially ordered sets (e.g., vectors), here we take a more pragmatic approach. In the formulation (12) we assume a total order for the vector in the objective, while we restricted the constraints to be scalar equalities and inequalities, hence working on the totally ordered set of reals. In particular, we use the *lexicographical order* to enforce a total order on a vector space of performance vectors. In the lexicographical order two vectors $\mathbf{y} = [y_1 \ y_2 \ \dots \ y_n]$ and $\mathbf{z} = [z_1 \ z_2 \ \dots \ z_n]$ satisfy $\mathbf{y} \leq \mathbf{z}$ if and only if $y_1 \leq z_1$, or $y_1 = z_1$ and $y_2 \leq z_2$, or $y_1 = z_1$, $y_2 = z_2$ and $y_3 \leq z_3$ etc. This order implies that the entries of the vector are sorted by “importance”. For instance, if we *minimize* a performance vector that includes [cost, size], then we search for the design that minimizes cost, and if two designs have the same cost, we prefer the design with smaller size. Note that we can also use the lexicographical order to generalize the constraints to be vector-valued functions.

III. LINEAR CO-DESIGN SOLVERS

We now discuss several cases in which we can expect to solve Problem (12) in reasonable time using off-the-shelf optimization tools. While Problem (12) is fairly general (we did not take any assumptions on the functions involved in the objective and the constraints), we do not expect to be able to solve (12) globally and efficiently in general. Indeed, binary optimization is NP-hard and the computational cost of solving a problem grows exponentially with its size [19].

Despite the intrinsic intractability of binary optimization, integer and binary programming algorithms keep improving and modern implementations (e.g., IBM CPLEX [20]) are already able to solve *linear and quadratic* binary optimization problems involving thousands of variables in reasonable time (i.e., seconds to few minutes). In our co-design problem, this

means that we can expect to solve problems with $N = 2,000$ (dimension of \mathbf{x}) in seconds, which would be the case if we have to design 10 modules, where each module catalog has 200 potential choices (remember $N \doteq \sum_{i \in \mathcal{M}} |\mathcal{C}_i|$).

The possibility of solving linear and quadratic binary optimization problems of interesting size in reasonable time motivates use to investigate when we can expect to rephrase (12) as a linear or quadratic optimization problem (note: the answer is not as trivial as it might seem). In the rest of this section we focus on the cases where we can rephrase (12) as a binary *linear* program (BLP), since this already includes several cases of practical interest. Since the linearity of Problem (12) relies on the capability of expressing both the objective and the constraints are linear functions, in the following we discuss which type of functions we can expect to rephrase as linear.

(a) Linear functions. Choosing linear functions in the objective and constraints would make (12) a BLP. Therefore, if the objective and the constraints in (12) have the following form, then (12) is a BLP:

$$f(\mathbf{x}) = \sum_{i \in \mathcal{M}} \mathbf{v}_i^T \mathbf{x}_i + \text{constant} \quad (13)$$

where \mathbf{v}_i is a known vector. We remark that most of the examples in Section II-B, including (3), (6), (7), and (11), can be directly expressed in this form.

(b) Sum of nonlinear functions of a module. Under the setup of Section II, we can express the sum of any nonlinear function involving a single module as a linear function. Consider for instance the following function:

$$f(\mathbf{x}) = \sum_{i \in \mathcal{M}} f_i(\mathbf{F}_i \mathbf{x}_i) \quad (14)$$

where each function f_i depends nonlinearly on the features of module i . We already observed that $\mathbf{F}_i \mathbf{x}_i$ picks the j -th column of the feature table \mathbf{F}_i , where j is the only non-zero entry in \mathbf{x}_i and is equal to 1. It is easy to show that:

$$f_i(\mathbf{F}_i \mathbf{x}_i) = \hat{f}_i(\mathbf{F}_i) \mathbf{x}_i \quad (15)$$

where $\hat{f}_i(\mathbf{F}_i)$ denotes the (known) row vector obtained by applying $f_i(\cdot)$ column-wise to the matrix \mathbf{F}_i :

$$\hat{f}_i(\mathbf{F}_i) \doteq \begin{bmatrix} f_i(\mathbf{F}_i^{(1)}) & f_i(\mathbf{F}_i^{(2)}) & \dots & f_i(\mathbf{F}_i^{(|\mathcal{C}_i|)}) \end{bmatrix} \quad (16)$$

Using (15), we can rewrite (14) as:

$$f(\mathbf{x}) = \sum_{i \in \mathcal{M}} \hat{f}_i(\mathbf{F}_i) \mathbf{x}_i \quad (17)$$

where $\hat{f}_i(\mathbf{F}_i)$ is a known vector. The expression (17) is now linear and has the same form as (13). Note that the “trick” (15) only holds since in our co-design problem each \mathbf{x}_i has a single non-zero entry equal to 1. As an example, (10) involves a sum of nonlinear functions of each module and so, according to our discussion, it can be expressed as a linear constraint.

(c) Rational functions. We now show that we can express a family of rational functions of the modules as linear functions. Let us start by considering the case in which the objective function in (12) is a rational function in the following form:

$$f(\mathbf{x}) = \frac{\prod_{i \in \mathcal{N}} f_i(\mathbf{F}_i \mathbf{x}_i)}{\prod_{i \in \mathcal{D}} f_i(\mathbf{F}_i \mathbf{x}_i)} \quad (18)$$

where $\mathcal{N} \subseteq \mathcal{M}$ and $\mathcal{D} \subseteq \mathcal{M}$ are arbitrary subsets of modules. Interestingly, it is possible to transform (18) into a linear function. For this purpose, we note that maximizing a quantity is the same as maximizing its logarithm, since the logarithm is a non-decreasing function. Therefore, we can replace the objective function (18) with the equivalent objective:

$$f'(\mathbf{x}) = \log \frac{\prod_{i \in \mathcal{N}} f_i(\mathbf{F}_i \mathbf{x}_i)}{\prod_{i \in \mathcal{D}} f_i(\mathbf{F}_i \mathbf{x}_i)} = \sum_{i \in \mathcal{N}} \log f_i(\mathbf{F}_i \mathbf{x}_i) - \sum_{i \in \mathcal{D}} \log f_i(\mathbf{F}_i \mathbf{x}_i) \quad (19)$$

Now, we note that (19) is a sum of nonlinear functions involving a single module, hence it can be simplified to the following linear function:

$$f'(\mathbf{x}) = \sum_{i \in \mathcal{N}} \log \hat{f}_i(\mathbf{F}_i) \mathbf{x}_i - \sum_{i \in \mathcal{D}} \log \hat{f}_i(\mathbf{F}_i) \mathbf{x}_i \quad (20)$$

The same approach can be applied to inequality (as well as equality) constraints in the following form (\bar{r} is a given scalar):

$$\frac{\prod_{i \in \mathcal{N}} f_i(\mathbf{F}_i \mathbf{x}_i)}{\prod_{i \in \mathcal{D}} f_i(\mathbf{F}_i \mathbf{x}_i)} \leq \bar{r} \quad (21)$$

which can be reformulated as equivalent linear constraints.

(d) Nonlinear functions of multiple modules. We conclude this section by considering the more general case, in which we have a non-linear function involving multiple modules. For simplicity of exposition, let us consider the case of a generic non-linear function involving two modules “a” and “b”:

$$f(\mathbf{x}) = f(\mathbf{x}_a, \mathbf{x}_b) \quad (22)$$

When $f(\cdot)$ does not have a specific structure (as the cases discussed above), it is still possible to obtain a linear expression for (22), but, as we will see, we will pay a price for this lack of structure. To express (22) as a linear function, we introduce an extra variable, a matrix \mathbf{Z}_{ab} which has size $|\mathcal{C}_a| \times |\mathcal{C}_b|$ and that represents the joint choice of the modules “a” and “b”. In other words, \mathbf{Z}_{ab} is zero everywhere and has a single entry equal to 1 in row j_a and column j_b , when we choose element j_a for module “a” and element j_b for module “b”. Clearly, the variables \mathbf{Z}_{ab} , \mathbf{x}_a , and \mathbf{x}_b are not independent and they have to satisfy the following linear constraints:

$$\mathbf{Z}_{ab} \mathbf{1} = \mathbf{x}_b, \quad \mathbf{1}^T \mathbf{Z}_{ab} = \mathbf{x}_a \quad (23)$$

where $\mathbf{1}$ is a vector of ones of suitable dimension ($\mathbf{Z}_{ab} \mathbf{1}$ returns the row-wise sum of the entries of \mathbf{Z}_{ab} , while $\mathbf{1}^T \mathbf{Z}_{ab}$ returns the column-wise sum). Intuitively, the constraints make sure that the matrix \mathbf{Z}_{ab} and the vectors \mathbf{x}_a and \mathbf{x}_b encode the same choice of components. By introducing the variable \mathbf{Z}_{ab} , we can write (22) as a linear function of \mathbf{Z}_{ab} following the same ideas of case (b) discussed above. The price to pay is an increase in the size of the optimization problem. In the general case, in which more than 2 variables are involved in a generic non-linear function, \mathbf{Z} becomes a (sparse) tensor with a number of entries equal to $\prod_{i \in \mathcal{N}} |\mathcal{C}_i|$, where \mathcal{N} is the set of modules involved in the (generic) non-linear transformation.

When increasing the size of the optimization problem is not an option, we can still use (12) and substitute non-linear

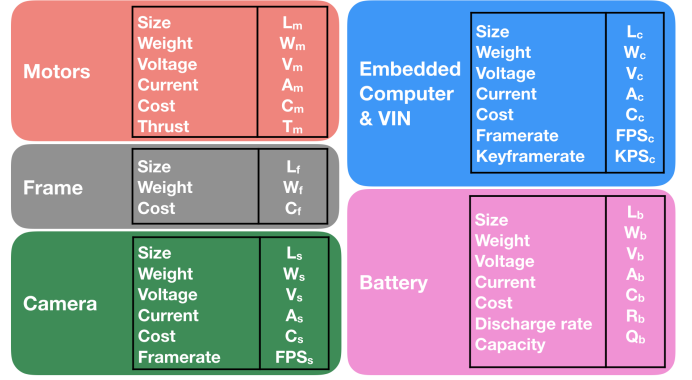


Fig. 2. Drone co-design overview: modules and features. functions with linear (or *linearizable*, as the cases above) surrogates. We show an example of this approach in Sec. IV-A.

IV. CODESIGN EXPERIMENTS

This section presents two examples of applications of the proposed co-design approach: Sec. IV-A focuses on the co-design of an autonomous racing drone, while Sec. IV-B considers the co-design of a team of robots for collaborative transportation. In both examples we use IBM CPLEX [20] to solve the binary optimization program (12).

A. Autonomous Drone Co-design

This section applies the proposed co-design framework to the problem of designing an autonomous micro aerial vehicle for drone racing. In particular, we answer the question: *what is the best autonomous drone design we can obtain on a \$1000 budget, using components chosen from a given catalog?*

Modules. The first step of the design process is to identify the set of modules \mathcal{M} we want to design and prepare a catalog for each module. In this example, we design five key modules that form an autonomous drone: motors, frame, computation, camera, and battery. For the motors, frame, and batteries, we selected real components that are commonly used for drone racing. In particular, we considered 17 candidate motors, 6 candidate frames, and 12 candidate batteries. The websites we used to select those components (together with their cost and their specs) as well as the actual numbers we used in this example are available at <https://bitbucket.org/lucacarlone/codesigncode/>. For the camera and the computation, we took a number of simplifying assumptions. For the camera, we disregarded the choice of the lens and we defined a set of (realistic) candidate cameras, where, however, only a single one is chosen from an actual catalog. For the design space of the “computation”, we actually considered the joint selection of the computing board and the visual-inertial navigation (VIN) algorithm used for state estimation. Considering them as a single component has the advantage of allowing the use of statistics reported in other papers, e.g., stating that a VIN algorithm runs at a given framerate on a given computer [1]. Therefore our modules are $\mathcal{M} \doteq \{m, f, s, c, b\}$ (m = motors, f = frame, s = camera, c = computer, b = battery).

An overview of the modules we design and their corresponding features is given in Fig. 2. For the sake of simplicity,

in this toy example we preferred not to design other components. For instance, we disregarded other algorithms running on the board, e.g., for control, which are typically less computationally demanding. We also neglected the presence/design of voltage adapters and connectors, while we assumed that each choice of motors comes with a suitable choice of ESC (Electronic Speed Control boards) and propellers.

System-level performance. The second step of the design process is to quantitatively define the system-level performance metrics (what is the “best drone”?). Since we consider an autonomous drone racing application, the best drone is one that can complete a given track as quickly as possible, hence a system that can navigate at high speed. Therefore, in this example, the system performance metric is the top speed of the drone. We mainly consider *forward* speed, but the presentation can be extended to maximize agility and accelerations.

In order to derive an expression for the top (forward) speed, we observe Fig. 1(b) and note that at its top speed, the forward acceleration is zero, hence the horizontal component T_h of the thrust T must compensate the drag force (i) $T_h = D$, and the vertical component T_v of the thrust must compensate the force of gravity (ii) $T_v = F_g = Mg$, where M is the overall mass of the drone and g is the acceleration due to gravity. Let us focus on (i), and note that (iii) $T^2 = T_h^2 + T_v^2$ and that the drag can be modeled as (iv) $D = \frac{1}{2} \rho c_d A v^2$, where ρ is the air density, c_d is the drag coefficient (we take $\rho = 1.2 \text{ kg/m}^3$ and $c_d = 1.3$), A is the cross-sectional area, and v is the forward speed of the drone. Substituting (ii), (iii), and (iv) back into (i), we obtain:

$$\sqrt{T^2 - (Mg)^2} = \frac{1}{2} \rho c_d A v_{\max}^2 \quad (24)$$

Now we note that the cross-sectional area can be approximated as $A = \sin(\beta) \ell^2 = \frac{Mg}{T} L_f^2$, where β is the pitch angle (basic trigonometry shows $\sin(\beta) = \frac{Mg}{T}$), and L_f is the length of the frame (approximated as a square). Substituting this expression for A in (24) and rearranging the terms we obtain:

$$v_{\max} = \left(\frac{4T^2}{\rho^2 c_d^2 L_f^4} \left(\frac{T^2}{(Mg)^2} - 1 \right) \right)^{\frac{1}{4}} \quad (25)$$

Finally, assuming that we design a quadrotor, the cumulative thrust is the sum of the thrusts T_m provided by each of the four motors, while the mass is the sum of the weight W_i of all the modules $i \in \mathcal{M}$:

$$v_{\max} = \left(\frac{4(4T_m)^2}{\rho^2 c_d^2 L_f^4} \left(\frac{(4T_m)^2}{(g \sum_{i \in \mathcal{M}} \omega_i W_i)^2} - 1 \right) \right)^{\frac{1}{4}} \quad (\text{SP})$$

where $\omega_i = 4$ if $i = m$ (we have 4 motors on a quadrotor, each one weighting W_m), or $\omega_i = 1$ otherwise.

In order to make explicit that the values of T_m, W_i, L_f depend on our design, we observe that $T_m = [\mathbf{F}_m]_T \mathbf{x}_m$, $W_i = [\mathbf{F}_i]_W \mathbf{x}_i$, $L_f = [\mathbf{F}_f]_L \mathbf{x}_f$, which makes (SP) a function of our design vector. Eq. (SP) represents the system performance metric our design has to maximize. It is apparent from (SP) that the design encourages drones which are small and lightweight (L_f and W_i appear at the denominator) and with large thrust (T_m appears at the numerator).

System-level constraints. System constraints provide further specifications (in the form of hard constraints) on the task the drone is designed for. For our drone example, we consider two main constraints: monetary budget and flight time.

1) *Budget*: Given a monetary budget \bar{b} , the budget constraint enforces that the sum of the costs C_i of each module i is within the budget:

$$\sum_{i \in \mathcal{M}} \omega_i C_i \leq \bar{b} \quad (\text{SC1})$$

where again $\omega_i = 4$ if $i = m$ or $\omega_i = 1$ otherwise.

2) *Flight time*: Given a minimum flight time \bar{T} (this would be between 5-10 minutes in a real application), and calling Q_b the battery capacity, and A_i the average Ampere drawn by the i -th component (all assumed to operate at the same voltage), then the time it takes to drain the battery must be $\geq \bar{T}$:

$$\frac{\alpha Q_b}{\sum_{i \in \mathcal{M}} \omega_i A_i} \geq \bar{T} \quad (\text{SC2})$$

where $\omega_i = 4$ if $i = m$ (again, we have 4 motors on a quadrotor), or $\omega_i = 1$ otherwise; the constant $\alpha \in (0, 1]$ (typically chosen to be around 0.8) models the fact that we might not want to fully discharge our battery (e.g., LiPo batteries may be damaged when discharged below a recommended threshold).

We remark that while flight time and budget constraints already make for an interesting problem, one can come up with many more system constraints, e.g., size and weight limits to participate into a specific drone racing competition, or motor power constraints for safety or regulatory constraints.

Implicit (module-level) constraints. The implicit constraints make sure that each module can operate correctly and it is compatible with the other modules in the system. We identified five implicit constraints:

1) *Minimum thrust*: the cumulative thrust provided by the four motors has to be sufficient for flight. Calling T_m the thrust provided by each motor, and W_i the weight of the i -th module, the minimum thrust constraint can be written as:

$$4T_m \geq \bar{r} g \sum_{i \in \mathcal{M}} \omega_i W_i \quad (\text{IC1})$$

where \bar{r} is a given minimum thrust-weight ratio ($\bar{r} = 2$ in this example), to ensure that the thrust is sufficient to maneuver with agility, besides allowing the drone to hover.

2) *Power*: the battery should provide enough power to support the four motors, the camera, and the computer:

$$A_b V_b \geq 4A_m V_m + A_s V_s + A_c V_c \quad (\text{IC2})$$

where A_i, V_i are the current and voltage for the i -th module.

3) *Size*: all components should fit on the frame. Assuming we can stack components on top of each other, the size constraint enforces that the size of each component should be smaller than the size of the frame:

$$L_f \geq L_i \quad \text{with } i \in \{m, s, c, b\} \quad (\text{IC3})$$

4) *Minimum camera frame-rate*: the framerate of the camera should be fast enough to allow tracking visual features, during fast motion. This is necessary for the visual-inertial navigation (VIN) algorithms to properly estimate position and attitude




Fig. 3. (a) Estimated maximum speed and (b) cost for any potential configuration of modules in our catalog. Infeasible designs are marked in red, while feasible designs are marked with a green star.

of the drone. Assuming that our VIN front-end can track a feature moving by at most δ_u pixels between frames, and corresponding to the projection of a 3D point at distance d from the camera (we assume $\delta_u = 30$ pixels, $d = 5$ meters), the frame-rate FPS_s of the camera is bounded by [1]:

$$\text{FPS}_s \geq \frac{f v_{\max}}{\delta_u d} \quad (\text{IC4})$$

where f is the focal length of the camera, and v_{\max} is the maximum speed of the drone, given in (SP).

5) *Minimum VIN frame-rate*: the VIN algorithm operating on the embedded computer on the drone should be able to process each frame, hence the VIN frame-rate (recall that we select from a catalog of VIN+computer combinations) must be larger than the camera frame-rate:

$$\text{FPS}_c \geq \text{FPS}_s \quad (\text{IC5})$$

Implementation and results. The “optimal” drone design solves the following optimization problem:

$$\begin{aligned} & \max_{\mathbf{x} \in \mathcal{X}} && (\text{SP}) \\ & \text{subject to} && (\text{SC1}), (\text{SC2}) \\ & && (\text{IC1}), (\text{IC2}), (\text{IC3}), (\text{IC4}), (\text{IC5}) \end{aligned} \quad (26)$$

Most of the constraints are already in a form that is amenable for our approach (Sec. III). Only the objective (SP) and the constraints (SC2) and (IC4) have a more involved expression, which we further develop in the appendix, where we show how to approximate those expressions in a form that fit a linear co-design solver.

The optimal design we found in our implementation suggests that an optimal configuration of modules would include a Stormer 220 FPV Racing Quadcopter Frame Kit, EMAX1045 motors, an NVIDIA TX2 computer, a 60 frame-per-second camera, and a Tattu 5100mAh 3S 10C Battery Pack. The cost of such a drone would be \$700.84, and the drone would have a flight time of 10.42 minutes and a top speed of 35.52 meters per second, which is compatible with the performance one expects from a racing drone [21]. CPLEX was able to find an optimal design in 0.3 seconds. Since the design space is relatively small we can compute the estimated maximum speed and the cost of every potential combination of the modules in our catalog. These results are shown in Fig. 3, where we also report whether the configuration is feasible (it satisfied all system and implicit constraints) or not. The figure shows that there are indeed only four feasible designs in our catalog all attaining fairly similar performance and cost.

B. Codesign of Heterogeneous Multi-Robot Teams

This section applies the proposed framework to the co-design of a heterogeneous multi-robot team. In particular, we consider a *collective transport scenario*, in which the robots must collectively carry a heavy object while avoiding obstacles. The robots must be configured to allow for efficient carrying and for wide sensor coverage, while battery power and frame size constrain the capabilities that any single robot possesses. The end result of our design is a heterogeneous team, composed of robots that specialize in carrying and robots that specialize in sensing. As a simplifying assumption, we focus on objects and robots with circular shapes (see Fig. 4).

Modules. We consider four types of modules: *frames* (f), *sensors* (s), *motors* (m), and *batteries* (b). Frames’ features include size and weight; sensors’ features include coverage (as a percentage of the surrounding area), size, weight, and power consumption; motors’ features include weight, size, power consumption and force exerted; and batteries’ features include size, weight, and power generated.

Since we have to design multiple robots, we use the notation \mathbf{x}_i^k to denote the design vector associated to the i -th module of robot k . As before, we use $[\mathbf{x}_i^k]_j$ to denote the j -th entry of \mathbf{x}_i^k . $[\mathbf{x}_i^k]_j = 1$ if we chose the j -th element in the catalog \mathcal{C}_i of module i on robot k , or zero otherwise. We let $k \in \{1, \dots, K\}$ where we calculate the upper bound K by considering the maximum number of robots that can encircle the object when the smallest frame is used. If the radius of the object is R^{object} and the radius of the smallest frame is $R^{\text{min frame}}$, then

$$K = \left\lceil \frac{\pi(R^{\text{object}} + R^{\text{min frame}})}{R^{\text{min frame}}} \right\rceil. \quad (27)$$

System-level performance. In a multi-robot system, performance is inherently non-monotonic. In Hamann’s analysis [15], performance is expressed as the ratio of two components: cooperation $C(\mathbf{x})$ and interference $I(\mathbf{x})$. Cooperation refers to those phenomena that contribute to the task at hand; interference corresponds to the phenomena that diminish the system performance. A simple approach to capture both aspects is to cast the co-design problem as an *interference minimization* problem, while using cooperation measures as system constraints. In [15], $I(\mathbf{x})$ is expressed as an exponential that decays with the size of the team R , from which it follows

$$\min_{\mathbf{x} \in \mathcal{X}} I(\mathbf{x}) \equiv \min_{\mathbf{x} \in \mathcal{X}} e^{R(\mathbf{x})} \equiv \min_{\mathbf{x} \in \mathcal{X}} R(\mathbf{x}) \equiv \max_{\mathbf{x} \in \mathcal{X}} -R(\mathbf{x}) \quad (28)$$

where the symbol “ \equiv ” denotes that these changes of objective do not alter the solution of the optimization problem. In (28) we also observed that the team size is a function of our design ($R = R(\mathbf{x})$), i.e., the design algorithm can decide to use less robots than the upper bound (27). Eq. (28) shows that our design will attempt to use the least number of robots, in order to minimize interference. To capture $R(\mathbf{x})$, we augment \mathbf{x} with a binary element (later called the “slot”) that indicates whether the features of a robot k are being used or not (in other words: if the robot is part of the team or not). We can thus define $R(\mathbf{x}) = \sum_{k=1}^K \mathbf{x}_{\text{slot}}^k$ and write

$$\text{SP}(\mathbf{x}) = - \sum_{k=1}^K \mathbf{x}_{\text{slot}}^k. \quad (\text{SP})$$

TABLE I
IMPLICIT CONSTRAINTS FOR THE MULTI-ROBOT CO-DESIGN EXAMPLE.

Meaning	Formalization	Tag
If a slot is activated, the robot has one frame	$\forall k \quad \mathbf{x}_{\text{slot}}^k = \sum_{j \in C_f} [\mathbf{x}_f^k]_j$	(IC1)
If a slot is activated, the robot has one motor	$\forall k \quad \mathbf{x}_{\text{slot}}^k = \sum_{j \in C_m} [\mathbf{x}_m^k]_j$	(IC2)
If a slot is activated, the robot has one battery	$\forall k \quad \mathbf{x}_{\text{slot}}^k = \sum_{j \in C_b} [\mathbf{x}_b^k]_j$	(IC3)
If a slot is activated, the robot has at most one sensor	$\forall k \quad \mathbf{x}_{\text{slot}}^k \geq \sum_{j \in C_s} [\mathbf{x}_s^k]_j$	(IC4)
The power consumption of a robot is lower than the power given by the battery	$\forall k \quad \sum_{j \in C_b} P_{b,j} [\mathbf{x}_b^k]_j \geq \sum_{i \in \{m,s\}, j \in C_i} P_{i,j} [\mathbf{x}_i^k]_j$	(IC5)
The total size of the components of a robot is lower than the size of the chassis	$\forall k \quad \sum_{j \in C_f} A_{f,j} [\mathbf{x}_f^k]_j \geq \sum_{i \in \{m,s,b\}, j \in C_i} A_{i,j} [\mathbf{x}_i^k]_j$	(IC6)




Fig. 4. Collective transport example (top view), where multiple robots (gray circles) carry a large object (black circle). The figure shows two solutions found by CPLEX when the object weighs 10 kg (left) and 80 kg (right).

System-level constraints. We consider two families of system constraints concerning motors and sensing.

1) *Push for object transportation:* the team as a whole must be able to carry the object; this is essentially the cooperation measure mentioned above. For this, we need that the push P_j exerted by the choice of motor j exceeds the sum of the weights $W_{i,j}$ of each choice modules ($i \in \mathcal{M}$ and $j \in C_i$) plus the weight W of the object to transport:

$$\sum_{k \in \{1, \dots, K\}, j \in C_m} P_j [\mathbf{x}_m^k]_j \geq \sum_{k \in \{1, \dots, K\}, i \in \mathcal{M}, j \in C_i} W_{i,j} [\mathbf{x}_i^k]_j + W \quad (\text{SC1})$$

In addition, each robot's motor push must exceed the weight of the robot itself:

$$\sum_{j \in C_m} P_j [\mathbf{x}_m^k]_j \geq \sum_{i \in \mathcal{M}, j \in C_i} W_{i,j} [\mathbf{x}_i^k]_j \quad \forall k \in \{1, \dots, K\} \quad (\text{SC2})$$

2) *Sensor coverage:* the robot sensors must ensure that at least 50% of the area around the object is covered at any time:

$$\sum_{j \in C_s} S_j [\mathbf{x}_s^k]_j \geq 50\% \quad (\text{SC3})$$

where S_j is the coverage provided by the j -th sensor choice.

Implicit (module-level) constraints. Table I summarizes the implicit constraints in our example. The symbol $P_{i,j}$ denotes the power offered by battery j or used by the j -th choice of module i . $A_{i,j}$ denotes the area offered by a frame j or the area used by the j -th choice of module i .

Implementation and results. To form the catalog of possible modules, we considered 10 alternatives for each type of module and 2 alternatives for the size of the robot frame, for a total of 2,000 combinations. Using the size and weight of the object to transport, we explored the space of optimal solutions, which involved teams of up to $K = 65$ robots. Fig. 4 reports the solutions we found by solving two instances of the problem, where the object to carry weighed 10 kg and 80 kg, respectively. In the left diagram, CPLEX concluded

that two large robots are sufficient to carry a 10 kg object. The robots are equipped with identical motors and batteries, and differ only in their sensor coverage. In the right diagram, CPLEX generated a solution including two types of robots: the larger type offers a lower pushing margin, but is capable of wide sensing; the smaller type is lighter and offers a higher pushing margin, but it performs no sensing. The presence of robots that carry no sensors allows them to collectively shoulder the bulk of the pushing. The interested reader can find the complete CPLEX implementation at <https://bitbucket.org/lucacarlone/codesigncode/>.

V. CONCLUSION

We presented an approach for *computational robot co-design* that formulates the joint selection of the modules composing a robotic system in terms of mathematical programming. While our approach is rooted in the general context of binary optimization, we discussed a number of properties—specific to our co-design problem—that allow rephrasing several co-design problems including non-linear functions of the features of each module in terms of *binary linear programming* (BLP). Modern BLP algorithms and implementations can solve problems with few thousands of variables in seconds, which in turn allows attacking interesting co-design problems. We demonstrated the proposed co-design approach in two applications: the design of an autonomous drone and the design of a multi-robot team for collective transportation. Future work includes extending the set of functions for which we can solve the co-design problem in practice, and adding continuous variables (e.g., wing length, 3D-printed frame size) as part of the co-design problem.

APPENDIX

A. Drone Co-design

This section discusses how to approximate the objective (SP) and the constraints (SC2) and (IC4) using binary linear functions. The following paragraphs deal with each case.

A linear lower bound for the objective (SP). The design has to maximize the maximum forward speed, which, from (SP), has the following expression:

$$v_{\max} = \left(\frac{4(4T_m)^2}{\rho^2 c_d^2 L_f^4} \left(\frac{(4T_m)^2}{(g \sum_{i \in \mathcal{M}} \omega_i W_i)^2} - 1 \right) \right)^{\frac{1}{4}} \quad (29)$$

This expression does not fall in cases (a), (b), (c) in Section III; moreover, it involves all modules, hence taking the approach (d) of Section III would be impractical (it would simply lead to enumerating every possible design choice, implying a combinatorial explosion of the state space).

The approach we take in this section is to approximate Problem (26) by replacing its objective with a *linear lower bound*. We remark that, as shown in Section III, while the design is required to be linear in \mathbf{x} , it can be heavily nonlinear in the features \mathbf{F}_i . In the following we show how to obtain a linear lower bound for (29). For this purpose, we observe

that from (IC1), the thrust-weight ratio must be larger than \bar{r} , hence $(4T_m)/(g \sum_{i \in \mathcal{M}} \omega_i W_i) \geq \bar{r}$. Therefore, it holds:

$$v_{\max} = \left(\frac{4(4T_m)^2}{\rho^2 c_d^2 L_f^4} \left(\frac{(4T_m)^2}{(g \sum_{i \in \mathcal{M}} \omega_i W_i)^2} - 1 \right) \right)^{\frac{1}{4}} \geq \left(\frac{4(4T_m)^2}{\rho^2 c_d^2 L_f^4} (\bar{r}^2 - 1) \right)^{\frac{1}{4}} = \kappa \frac{T_m^{\frac{1}{2}}}{L_f} = \kappa \frac{([\mathbf{F}_m]_T \mathbf{x}_m)^{\frac{1}{2}}}{[\mathbf{F}_f]_L \mathbf{x}_f} \quad (30)$$

where κ is a constant, irrelevant for the maximization. The function (30) now falls in the case (c) discussed in Section III and can be made linear in \mathbf{x} by taking the logarithm.

A conservative linear approximation for (IC4). Similarly to the case discussed above, we approximate the “problematic” constraints via surrogates that are linear in \mathbf{x} .

The approach we take is to substitute a constraint in the form $\hat{f}(\mathbf{x}) \leq 0$ with a linear constraint $\tilde{f}(\mathbf{x}) \leq 0$, where $\tilde{f}(\mathbf{x}) \leq \hat{f}(\mathbf{x})$ for any $\mathbf{x} \in \mathcal{X}$. This guarantees that for any \mathbf{x} that makes $\tilde{f}(\mathbf{x}) \leq 0$, then also $\hat{f}(\mathbf{x}) \leq 0$, i.e., we still guarantee to compute feasible (but potentially more conservative) designs.

Therefore, in the following we show how to obtain a linear upper bound for (IC4). Let us start by writing (IC4) more explicitly, by substituting v_{\max} from (9):

$$\text{FPS}_s \geq \frac{f}{\delta_u d} \left(\frac{4(4T_m)^2}{\rho^2 c_d^2 L_f^4} \left(\frac{(4T_m)^2}{(g \sum_{i \in \mathcal{M}} \omega_i W_i)^2} - 1 \right) \right)^{\frac{1}{4}} \quad (31)$$

or equivalently (taking the 4-th power of each member):

$$\text{FPS}_s^4 \geq \frac{4f^4}{\delta_u^4 d^4 \rho^2 c_d^2 L_f^4} \left(\frac{(4T_m)^2}{(g \sum_{i \in \mathcal{M}} \omega_i W_i)^2} - 1 \right) \quad (32)$$

Taking the logarithm of both sides we obtain:

$$4 \log(\text{FPS}_s) \geq \beta + 2 \log(4T_m) - 4 \log(L_f) + \log \left(\frac{(4T_m)^2}{(g \sum_{i \in \mathcal{M}} \omega_i W_i)^2} - 1 \right) \quad (33)$$

where we defined the constant $\beta \doteq \log(\frac{4f^4}{\delta_u^4 d^4 \rho^2 c_d^2})$.

Rearranging the terms:

$$+ \log \left(\frac{(4T_m)^2}{(g \sum_{i \in \mathcal{M}} \omega_i W_i)^2} - 1 \right) - 4 \log(\text{FPS}_s) + \beta \leq 0 \quad (34)$$

We note that so far we did not take any approximation since each operation we applied (4-th power, logarithm, reordering) preserves the original inequality. Moreover, (34) is in the form $\hat{f}(\mathbf{x}) \leq 0$. In the rest of this paragraph, we show how to

compute a linearized upper bound for $\hat{f}(\mathbf{x})$. For this purpose we note that the following chain of inequalities holds:

$$\begin{aligned} & 2 \log(4T_m) - 4 \log(L_f) + \\ & + \log \left(\frac{(4T_m)^2}{(g \sum_{i \in \mathcal{M}} \omega_i W_i)^2} - 1 \right) - 4 \log(\text{FPS}_s) + \beta \stackrel{(i)}{\leq} \\ & 2 \log(4T_m) - 4 \log(L_f) + \\ & + \log \left(\frac{(4T_m)^2}{(g \sum_{i \in \mathcal{M}} \omega_i W_i)^2} \right) - 4 \log(\text{FPS}_s) + \beta \stackrel{(ii)}{=} \\ & 4 \log(4T_m) - 4 \log(L_f) + \\ & - 2 \log \left(\sum_{i \in \mathcal{M}} g \omega_i W_i \right) - 4 \log(\text{FPS}_s) + \beta \stackrel{(iii)}{\leq} \\ & 4 \log(4T_m) - 4 \log(L_f) + \\ & - 2 \sum_{i \in \mathcal{M}} \frac{1}{|\mathcal{M}|} \log(g \omega_i W_i) - 2 \log(|\mathcal{M}|) - 4 \log(\text{FPS}_s) + \beta \end{aligned} \quad (35)$$

where in (i) we dropped the -1 and used the fact that the logarithm is a non-decreasing function, in (ii) we simply developed the logarithm of the ratio, and in (iii) we used the Jensen’s inequality:

$$\log \left(\frac{\sum_{i=1}^n y_i}{n} \right) \geq \frac{\sum_{i=1}^n \log(y_i)}{n} \iff \quad (36)$$

$$\log \left(\sum_{i=1}^n y_i \right) \geq \sum_{i=1}^n \frac{1}{n} \log(y_i) + \log(n) \quad (37)$$

which holds for any integer n and positive scalar y_i .

Making the dependence of T_m , L_f , W_i , and FPS_s on the design vector \mathbf{x} explicit in (35), we obtain:

$$\begin{aligned} \hat{f}(\mathbf{x}) & \doteq 4 \log(4[\mathbf{F}_m]_T \mathbf{x}_m) - 4 \log([\mathbf{F}_f]_L \mathbf{x}_f) + \\ & - 2 \sum_{i \in \mathcal{M}} \frac{1}{|\mathcal{M}|} \log(g \omega_i [\mathbf{F}_i]_W \mathbf{x}_i) + \\ & - 2 \log(|\mathcal{M}|) - 4 \log([\mathbf{F}_s]_{\text{FPS}} \mathbf{x}_s) + \beta \end{aligned} \quad (38)$$

which now falls in the case (b) discussed in Section III and can be easily made linear with respect to \mathbf{x} .

A conservative linear approximation for (SC2). Similarly to the previous section, here we approximate the “problematic” constraint (SC2) via surrogates that are linear in \mathbf{x} . In particular, as before, we substitute a constraint in the form $\hat{f}(\mathbf{x}) \leq 0$ with a linear constraint $\tilde{f}(\mathbf{x}) \leq 0$, where $\tilde{f}(\mathbf{x}) \leq \hat{f}(\mathbf{x})$ for any $\mathbf{x} \in \mathcal{X}$. This guarantees that for any \mathbf{x} that makes $\tilde{f}(\mathbf{x}) \leq 0$, then also $\hat{f}(\mathbf{x}) \leq 0$, i.e., we still guarantee to compute feasible (but potentially more conservative) designs.

Let us start by reporting the original expression we want to approximate (SC2):

$$\frac{\alpha Q_b}{\sum_{i \in \mathcal{M}} \omega_i A_i} \geq \bar{T} \quad (39)$$

which, again, does not fall in cases (a), (b), (c) in Section III, and involves all modules, hence making the approach described in case (d) of Section III impractical.

In order to obtain a conservative approximation of the constraint (39), we first take the logarithm of both members:

$$\log(\alpha Q_b) - \log \left(\sum_{i \in \mathcal{M}} \omega_i A_i \right) \geq \log(\bar{T}) \quad (40)$$

which does not alter the inequality since all involved quantities are positive. Rearranging the terms:

$$\log\left(\sum_{i \in \mathcal{M}} \omega_i A_i\right) - \log(\alpha Q_b) + \log(\bar{T}) \leq 0 \quad (41)$$

We can find an upper bound for the left-hand-side of (41) as follows:

$$\log\left(\sum_{i \in \mathcal{M}} \omega_i A_i\right) - \log(\alpha Q_b) + \log(\bar{T}) \leq \log(6A_m) - \log(\alpha Q_b) + \log(\bar{T}) \quad (42)$$

where we used $\sum_{i \in \mathcal{M}} \omega_i A_i = A_s + A_c + 4A_m \leq 6A_m$ (note: frame and batteries are assumed to draw zero current) which follows from the fact the motors are the modules that draw more current. Therefore we can substitute the original constraint (39) with a conservative linear approximation:

$$\hat{f}(\mathbf{x}) \doteq \log(6[\mathbf{F}_m]_A \mathbf{x}_m) - \log(\alpha[\mathbf{F}_b]_Q \mathbf{x}_b) + \log(\bar{T}) \leq 0 \quad (43)$$

which now falls in the case (b) discussed in Section III and can be easily made linear with respect to \mathbf{x} .

REFERENCES

- [1] Z. Zhang, A. Suleiman, L. Carlone, V. Sze, and S. Karaman, "Visual-inertial odometry on chip: An algorithm-and-hardware co-design approach," in *Robotics: Science and Systems (RSS)*, 2017.
- [2] NVIDIA, "Nvidia jetson tx2 module specifications." [Online]. Available: <https://developer.nvidia.com/embedded/buy/jetson-tx2>
- [3] D. Lipinski and K. Mohseni, "Feasible area coverage of a hurricane using micro-aerial vehicles," in *AIAA Scitech*, 2014, pp. 824–829.
- [4] H. Lipson and J. B. Pollack, "Automatic design and manufacture of robotic lifeforms," *Nature*, vol. 406, no. 6799, pp. 974–978, Aug. 2000.
- [5] H. Lipson, V. Sunspirai, J. Bongard, and N. Cheney, "On the difficulty of co-optimizing morphology and control in evolved virtual creatures," in *Proceedings of the European Conference on Artificial Life 13*. MIT Press, 2016, pp. 226–233.
- [6] N. Cheney, J. Bongard, V. SunSpiral, and H. Lipson, "Scalable co-optimization of morphology and control in embodied machines," *Journal of The Royal Society Interface*, vol. 15, no. 143, p. 20170937, June 2018.
- [7] P. H. Chou, R. B. Ortega, and G. Borriello, "The Chinook Hardware/Software Co-synthesis System," in *Proceedings of the 8th International Symposium on System Synthesis*, ser. ISSS '95. New York, NY, USA: ACM, 1995, pp. 22–27.
- [8] B. C. Schafer and K. Wakabayashi, "Divide and Conquer High-level Synthesis Design Space Exploration," *ACM Trans. Des. Autom. Electron. Syst.*, vol. 17, no. 3, pp. 29:1–29:19, July 2012.
- [9] "Architecture Analysis and Design Language." [Online]. Available: <http://www.aadl.info/aadl/currentsite/>
- [10] A. Mehta, J. DelPreto, and D. Rus, "Integrated Codesign of Printable Robots," *Journal of Mechanisms and Robotics*, vol. 7, no. 2, pp. 021 015–021 015–10, May 2015.
- [11] O. Hussien, A. Ames, and P. Tabuada, "Abstracting Partially Feedback Linearizable Systems Compositionally," *IEEE Control Systems Letters*, vol. 1, no. 2, pp. 227–232, Oct. 2017.
- [12] E. S. Kim, M. Arcak, and M. Zamani, "Constructing Control System Abstractions from Modular Components," in *Proceedings of the 21st International Conference on Hybrid Systems: Computation and Control (Part of CPS Week)*, ser. HSCC '18. New York, NY, USA: ACM, 2018, pp. 137–146.
- [13] A. Censi, "A Mathematical Theory of Co-Design," *arXiv:1512.08055 [cs, math]*, Dec. 2015, arXiv: 1512.08055.
- [14] —, "A Class of Co-Design Problems With Cyclic Constraints and Their Solution," *IEEE Robotics and Automation Letters*, vol. 2, no. 1, pp. 96–103, Jan. 2017.
- [15] H. Hamann, "Towards Swarm Calculus: Universal Properties of Swarm Performance and Collective Decisions," in *Swarm Intelligence*, D. e. a. Hutchison, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, vol. 7461, pp. 168–179.
- [16] A. Kushleyev, D. Mellinger, C. Powers, and V. Kumar, "Towards a swarm of agile micro quadrotors," *Autonomous Robots*, vol. 35, no. 4, pp. 287–300, 2013.
- [17] V. Trianni, *Evolutionary Swarm Robotics - Evolving Self-Organising Behaviours in Groups of Autonomous Robots*, ser. Studies in Computational Intelligence. Springer, 2008, vol. 108.
- [18] G. Francesca, M. Brambilla, A. Brutschy, L. Garattoni, R. Miletitch, G. Podevijn, A. Reina, T. Soleymani, M. Salvaro, C. Pincirolì, F. Mascia, V. Trianni, and M. Birattari, "AutoMoDe-Chocolate: automatic design of control software for robot swarms," *Swarm Intelligence*, vol. 9, no. 2, pp. 125–152, Sept. 2015.
- [19] A. Schrijver, *Theory of Linear and Integer Programming*. New York, NY, USA: John Wiley & Sons, Inc., 1986.
- [20] IBM, "CPLEX: IBM ILOG CPLEX Optimiz Studio." [Online]. Available: <https://www.ibm.com/products/ilog-cplex-optimization-studio>
- [21] FPV Drone Reviews, "Fastest FPV drones," <https://fpvdroneviews.com/guides/fastest-racing-drones/>, 2017.