

Layered Control Systems Operating on Multiple Clocks

Inigo Incer[✉], Member, IEEE, Noel Csomay-Shanklin[✉], Graduate Student Member, IEEE,
 Aaron D. Ames[✉], Fellow, IEEE, and Richard M. Murray[✉], Fellow, IEEE

Abstract—Autonomous systems typically leverage layered control architectures, created by interconnecting components that operate at multiple timescales, i.e., evolve under various clocks. To formalize this typically heuristic procedure, we introduce a new logic, Multiclock Logic (MCL), that can express the requirements of components from the point of view of their local clocks, promoting independent design and component reuse. We then use assume-guarantee contracts expressed in MCL to prove global stability properties of a system using the stability properties of its components. In particular, we consider the classic layered architecture consisting of model predictive control (MPC) layered on top of feedback linearization, and prove overall stability of the systems.

Index Terms—Hierarchical control, networked control systems, hybrid systems.

I. INTRODUCTION

OVER the years, the field of control theory has developed many tools to design control blocks in isolation—e.g., PID, feedback linearization (FBL), model predictive control (MPC) [1], [3], control Lyapunov functions (CLFs) [2], [12], etc. The design of complex control systems—legged robots, aerial robots, and autonomous vehicles, to name a few—normally involves the combination of various blocks of control functionality. It is often the case that designers working independently on each control block make assumptions on the behaviors of other blocks that are not communicated or explicitly stated—leading to a development process prone to errors. Moreover, control blocks are normally implemented at different loop rates, leading to unaccounted for timing and interfacing issues. The end result is that we interconnect our control blocks *hoping the system will work*.

This letter provides a compositional, specification-based approach to proving control properties of layered systems by

Manuscript received 8 March 2024; revised 13 May 2024; accepted 1 June 2024. Date of publication 5 June 2024; date of current version 19 June 2024. This work was supported in part by the NSF and ASEE through an eFellows Postdoctoral Fellowship; in part by the Air Force Office of Scientific Research (AFOSR) under Grant FA9550-22-1-0316 and Grant FA9550-22-1-0333; and in part by the NSF CPS Award under Grant 1932091. Recommended by Senior Editor S. Dey. (Corresponding author: Inigo Incer.)

The authors are with the Division of Engineering and Applied Science, California Institute of Technology, Pasadena, CA 91125 USA (e-mail: inigo@caltech.edu; noelcs@caltech.edu; ames@caltech.edu; murray@cds.caltech.edu).

Digital Object Identifier 10.1109/LCSYS.2024.3410150

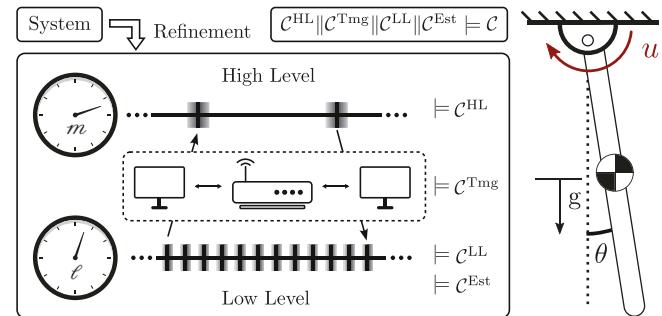


Fig. 1. A typical layered control architecture, whereby MCL contracts enable independently designed controllers to meet system-level specifications.

using the properties of each layer. We introduce a new logic, Multiclock Logic (MCL), to be able to express properties from the local point of view of a processor running on its own clock—capturing the multi-rate nature of layered control systems. We use MCL to abstract layers into assume-guarantee contracts, which are formal specifications that capture what a component guarantees and requires from its environment to be able to deliver its guarantees [6]. We use the algebra of contracts to compositionally relate the local stability properties of control algorithms with the system-level properties of their interconnection. While multi-rate control architectures were first studied in [10], the problem of compositionally writing and analyzing formal specifications for control systems made from layers that operate on distinct time bases has not been studied.

This letter is organized as follows. Section II introduces our motivating example. Section III presents the syntax and semantics of MCL. Section IV shows how to express contract specifications for a multilayer controller using MCL, and to prove system-level stability properties from the properties of the control layers. An extended version of this letter [7] contains background on behavioral modeling and contracts.

II. MULTI-RATE SYSTEM CASE STUDY

To illustrate our compositional analysis of a layered architecture, we will consider the design of a control strategy for the pendulum shown in Figure 1. Our objective will be for the system to satisfy a top-level stability objective expressed as contract \mathcal{C} . We will understand the controller for such as

system as an architecture consisting of a high-level and a low-level control layers interconnected by a network. Each control layer will be assumed to be implemented on a processor operating on its own clock, as is common in practice. We will identify contracts \mathcal{C}^{HL} for the high-level layer, \mathcal{C}^{LL} for the low-level layer, \mathcal{C}^{Est} for state estimation, and \mathcal{C}^{Tmg} for the networking between the layers. As these contracts need to state formal properties over multiple clocks, they will be expressed in MCL in Section IV. The property of these contracts is that their composition is a refinement of the desired stability contract \mathcal{C} . This means that we will be free to independently implement these aspects of the functionality, and the procedure will be guaranteed to yield a system that satisfies the system-level objective \mathcal{C} . Afterwards, we will implement each of these control blocks independently.

We now consider the details of this case study. The configuration space of the system is given by $\theta \in \mathcal{Q} \triangleq \mathbb{S}^1$, and the associated state by $x = (\theta, \dot{\theta}) \in \mathbb{T}\mathcal{Q}$. We can write the dynamics in control-affine form as $\dot{x} = f(x) + g(x)u$, with control input $u \in \mathbb{R}$, continuously differentiable drift vector $f : \mathbb{T}\mathcal{Q} \rightarrow \mathbb{R}^2$, and actuation matrix $g : \mathbb{T}\mathcal{Q} \rightarrow \mathbb{R}^2$. We will be concerned with the high level task of reaching some neighborhood of the goal state $x_g = 0$, the unstable upright equilibrium. Low-level controllers alone struggle to simultaneously enforce state and input constraints while ensuring progress to the goal is being made [3]—this motivates the use of a hierarchy for achieving this control task. Our hierarchy will consist of two layers—a high-level controller and a low-level controller—running on different clocks.

1) Low-Level Controller Design: To actuate the system, we use a feedback controller which is able to track a desired trajectory x_d . We begin by defining an output $y : \mathcal{X} \rightarrow \mathbb{R}$ as $y(x) = \theta$, whereby we can construct error coordinates $e(x, t) = [y(x) \ L_f y(x)]^\top - [y(x_d(t)) \dot{y}(x_d(t))]^\top$, with $L_f y(x)$ denoting the Lie derivative of y with respect to f . Then, a feedback control law can be created to stabilize the outputs. For example, the control law can be the feedback linearizing controller

$$k_{\text{fbl}}(x, t) = L_g L_f(x)^{-1} \left(-L_f^2 y(x) + \ddot{y}(x_d(t)) - K e(x, t) \right),$$

which exponentially stabilizes the output coordinates for elementwise positive vector $K \in \mathbb{R}^2$. Given bounded additive disturbances to the dynamics, there exists a robust invariant set $\mathcal{E} \subset \mathbb{R}^n$ such that applying k_{fbl} results in bounded tracking error, i.e., $x(t) \in x_d(t) \oplus \mathcal{E}$ [4].

2) High-Level Controller Design: This controller will be concerned with producing the pieces of trajectories x_d which make progress towards the goal x_g for the low-level controller to track. In order to do so, we set up the following MPC program:

$$\begin{aligned} \min_{x_d(t), u_d(t)} \quad & \int_0^T h(x_d(t), u_d(t)) dt \\ \text{s.t.} \quad & \dot{x}_d(t) = f^a(x_d(t)) + g^a(x_d(t))u_d(t), \\ & x_d(0) \in \hat{x} \oplus \mathcal{E}, \quad x_d(T) = 0, \\ & x(t) \in \mathcal{X}, \quad u(t) \in \mathcal{U}, \end{aligned} \quad (1)$$

with $\mathcal{X} \subset \mathbb{R}^2$ a state constraint set, $\mathcal{U} \subset \mathbb{R}$ an input constraint set, $h : \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}_{\geq 0}$ a convex stage cost, f^a and g^a linear approximations of the dynamics, $\hat{x} \in \mathcal{X}$ an estimate of the system state, $\mathcal{E} \subset \mathbb{R}^n$ a robust invariant set of the low level

controller, and $x(t)$ and $u(t)$ the continuous-time state and input of the low-level system.

Combining MPC with a low-level controller is a popular technique which is extensible to stabilizing complex nonlinear systems [5], [10]. When deploying such an architecture, however, the underlying approximations of convexification and discretization are often not explicitly reasoned about, leaving practitioners hoping rather than knowing that the closed-loop system will work. Common assumptions include: *i*) the estimate of the state \hat{x} is “close enough” to the true state x ; *ii*) the MPC runs “fast enough” and produces solutions which are “close enough” to the system dynamics; and *iii*) the low-level controller runs “fast enough” and can track “well enough.” The aim of this letter is to construct a pipeline whereby these assumptions are made explicit, with an eye towards automating the system verification process for layered control systems.

III. MULTICLOCK LOGIC

In this section, we introduce Multiclock Logic (MCL), a framework to reason about layered control operating on multiple clocks. First we define the behaviors over which the logic makes predicates, then introduce the logic. MCL will allow us to express properties from the local point of view of each control block in our design without burdening the syntax with synchronization over the multiple clocks present in the system. While some of these properties can be expressed intuitively, MCL is a formal language and so can be manipulated by a machine. This is necessary in order to automate formal system-level design.

1) System-Level Behaviors: The fundamental notion of system-modeling is the *variable*. We understand a variable as a tuple $(v, \mathcal{X}_v, \mathcal{D}_v)$, where v is the *name* for the variable, \mathcal{X}_v is the *value space* (a topological space where the variable takes values), and \mathcal{D}_v is a totally-ordered Abelian group called the *clock*, giving a notion of a progression of values. Whether a variable is discrete/continuous/constant depends on the clock that we assign to it: $\mathbb{R}/\mathbb{N}/\{\bullet\}$, respectively. We model clocks as totally-ordered Abelian groups because we need order to compare values of a clock and define intervals, and the group structure to add and subtract these values.

Suppose we have a set \mathcal{C} of clocks and that a set \mathbf{Vars} of variables shares clock $c \in \mathcal{C}$. The joint behaviors of these variables will be given by maps $c \rightarrow \prod_{v \in \mathbf{Vars}} \mathcal{X}_v$, as in the tagged signal model [9]. In order to isolate the value of a variable $u \in \mathbf{Vars}$ from those of the other variables in the same clock, we will use the projection maps that come with the definition of the product, i.e., $\pi_u : \prod_{v \in \mathbf{Vars}} \mathcal{X}_v \rightarrow \mathcal{X}_u$.

Stating certain predicates may require us to use a clock to read the values of variables belonging to another clock. Suppose $c, d \in \mathcal{C}$. To read the value of a d -variable in the clock c , we will assume the existence of a map $\tau_c^d : c \rightarrow d$. The τ maps, which we may call *synchronization maps*, determine the index of the target clock that is used when reading variables from a given clock. e.g., to read the value of a variable x clocked by d from the clock c at time $t \in c$, we would get the value of x at time $\tau_c^d(t)$. The synchronization maps capture the amount of delay incurred by reading data from a clock from which the data does not originate.

Suppose that we build a system that comprises a set of variables \mathbf{Vars} and a set of clocks \mathcal{C} . For each variable v , we

let $C(v)$ be the clock corresponding to $v \in \mathbf{Vars}$. The inverse of this map, $C^{-1}(c)$, yields the set of variables corresponding to clock $c \in \mathfrak{C}$.

Definition 1: A system-level behavior is an object $\beta = (b_d)_{d \in \mathfrak{C}} \times (\tau_c^d)_{c,d \in \mathfrak{C}}$, where the elements $b_d \in (d \rightarrow \prod_{v \in C^{-1}(d)} \mathcal{X}_v)$ carry the behaviors of the variables clocked by d , while the $\tau_c^d \in (c \rightarrow d)$ indicate how to read values of d -variables using the clock c .

Components, as usual, are defined as sets of behaviors. The following definition will be useful to give semantics to the eventual modality \diamond of the logic.

Definition 2: Let $\beta = (b_d)_{d \in \mathfrak{C}} \times (\tau_d^d)_{d,d \in \mathfrak{C}}$ be a system behavior. Given $t \in c \in \mathfrak{C}$, we define the (c, t) -execution

$$\beta_c^t \text{ as } \beta_c^t = (b_d)_{d \in \mathfrak{C}} \times (\tilde{\tau}_d^d(x) = \begin{cases} \tau_c^d(x + t) & d = c \\ \tau_d^d(x) & d \neq c \end{cases})$$

The only difference between β and β_c^t is that in β_c^t the clock c is anticipated by t units.

2) MCL Syntax: Assume we have access to a set \mathbf{Vars} of variables, a set \mathfrak{C} of clocks, and a set \mathfrak{F} of formulas of various arities. We let \mathfrak{P} be the set of symbols d^d , where d and d' denote clocks. The syntax of MCL is

$$\phi := \rho c. \Phi \mid \neg\phi \mid \phi \wedge \psi$$

$$\Phi := P(\lambda_1(t_1), \dots, \lambda_n(t_n)) \mid \neg\Phi \mid \Phi \wedge \Psi \mid \diamond_{[t_1, t_2]} \Phi \mid \diamond_{t_1} \Phi,$$

where $c \in \mathfrak{C}$, $\lambda_i \in \mathbf{Vars} \cup \mathfrak{P} \cup \mathfrak{C}$, $t_i \in c$, and $P \in \mathfrak{F}$ is an n -ary formula.

We think of the formulas ϕ as the global syntax of MCL, and of the Φ formulas as the local syntax that applies to a clock. The global syntax supports propositional logic. The syntax $\phi = \rho c. \Phi$ indicates that a global formula ϕ is created by binding a local formula Φ to the clock c . The local syntax supports propositional logic and the eventual modality. It also supports the enunciation of predicates involving multiple variables evaluated at the indicated times $t_i \in c$. The syntax also supports the enunciation of predicates on clocks $c \in \mathfrak{C}$ and on *clock pairs* $c^d \in \mathfrak{P}$. The use of clocks and clock pairs $c^d \in \mathfrak{P}$ in formulas allows us to express timing constraints in our systems. Observe that the logic leaves abstract the formulas \mathfrak{F} used to define the predicates. The formulas are functions that evaluate to Boolean values.

3) MCL Semantics: The semantics of an MCL formula apply to a system behavior. Let $\beta = (b_d)_{d \in \mathfrak{C}} \times (\tau_d^d)_{d,d \in \mathfrak{C}}$ be a system behavior. We have the following global semantics:

- $\beta \models \rho c. \Phi$ iff $\beta \models_c \Phi$
- $\beta \models \neg\phi$ iff $\beta \not\models \phi$
- $\beta \models \phi \wedge \psi$ iff $\beta \models \phi$ and $\beta \models \psi$

Formulas at the global level are either created by Boolean connectives, or they are local formulas that are assigned to a clock in the system. The symbol “ \models_c ” stands for satisfaction in the local semantics of MCL, defined as follows:

- $\beta \models_c P(\lambda_1(t_1), \dots, \lambda_n(t_n))$ iff $P(\text{Interp}(\lambda_i, t_i))_{i=1}^n$, where

$$\text{Interp}(\lambda_i, t_i) = \begin{cases} \pi_v \circ b_{C(v)} \circ \tau_c^{C(v)}(t_i) & \lambda_i = v \in \mathbf{Vars} \\ \tau_c^d(t_i) & \lambda_i = d \in \mathfrak{C} \\ \tau_d^d \circ \tau_c^d(t_i) & \lambda_i = d^d \in \mathfrak{P} \text{ and } d \neq c \end{cases}$$

- $\beta \models_c \neg\Phi$ iff $\beta \not\models_c \Phi$
- $\beta \models_c \Phi \wedge \Psi$ iff $\beta \models_c \Phi$ and $\beta \models_c \Psi$
- $\beta \models_c \diamond_{[t_1, t_2]} \Phi$ iff $\exists t. (t_1 \leq t \leq t_2) \wedge (\beta_c^t \models_c \Phi)$
- $\beta \models_c \diamond_{t_1} \Phi$ iff $\exists t. (t_1 \leq t) \wedge (\beta_c^t \models_c \Phi)$

In addition to the formation of local formulas using Boolean connectives and the eventual modality \diamond , MCL local formulas can be formed by making a predicate over a set of symbols λ_i and a set of values t_i of the clock c . In general, the meaning of the formula $P(\lambda_1(t_1), \dots, \lambda_n(t_n))$ is the evaluation of the n -ary formula P for the values that each symbol λ_i takes when the value of the clock c is $\tau_c^i(t_i)$. For example, the predicate $\rho c. \|x(0)\| > 2$ for a variable x means that we will evaluate the value of x when $c = \tau_c^i(0)$. If x is clocked by c , we simply extract the value of x at time $\tau_c^i(0)$; if it is clocked by $d \neq c$, we would return its value at time $\tau_d^i(0)$. We recall that the local semantics can anticipate the value of c by t time units—see Definition 2.

The formula $\rho c. d(0) - K > T$ for $d \in \mathfrak{C}$ and constants $T, K \in d$ has the semantics $\tau_c^d(0) - K > T$. This statement is true at ticks of clock c for which the index of clock d that c uses to read the data of d is larger than $K + T$. This type of predicate can be used as a precondition that ensures that the clock d has issued sufficient ticks. The semantics of a formula of the form $\rho c. r(0) - d^r(0) < T$ for clocks c, d, r , is $\tau_c^r(0) - \tau_d^r \circ \tau_c^d(0) < T$. It states that the difference in r time units between a tick of c and the tick of d from which c reads the data of d must be less than T . This kind of statement is useful to bound the maximum allowed time for data to have existed in its local clock d before it is read by another clock c .

The notion of satisfaction is extended from system behaviors to components as follows: we say that a component M satisfies an MCL formula ϕ if $\beta \models \phi$ for all $\beta \in M$.

We extend the given syntax and semantics in the standard way to support all Boolean connectives and the modality \square , or “globally.” Finally, when stating a formula of the form $\rho c. P(\lambda_i(t_i))_i$, we will sometimes omit the t_i arguments. In that case, the parameter should be understood as 0.

IV. VERIFYING THE MULTI-RATE CASE STUDY

We now continue the analysis of the archetypal layered architecture described in Section II. Our goal is to specify the high-level and low-level control blocks independently and carry out system-level analysis using these specifications. The objective of our system-level analysis is to show that the system reaches and remains inside a neighborhood of the goal state x_g . To do so, we assume we have a cost function \mathcal{V} that maps the state x of the system to a well-ordered set. The value of \mathcal{V} is zero in a neighborhood of x_g . The function \mathcal{V} will be further specified when we consider the implementations of component contracts in Section IV-C. As a well-order does not have infinite descending chains, the fact that the codomain of \mathcal{V} is a well-order means that any process that decreases the cost \mathcal{V} will eventually reach the minimum cost $\mathcal{V}(x) = 0$ in a finite number of iterations.

The system has three clocks: $m = \mathbb{N}$, which runs the high-level block; $\ell = \mathbb{N}$, which runs the low-level and estimation blocks; and $\tau = \mathbb{R}_{\geq 0}$, the physical time. The system state is denoted by the variable (x, \mathcal{X}, τ) , i.e., its behaviors are functions from the physical clock to the state space \mathcal{X} . The high-level block outputs a variable $(x_d, \mathcal{X}^\tau, m)$, which contains the trajectory in the state space that the system is to follow. Observe that the behaviors of this variable are functions from m to functions from τ to \mathcal{X} . Thus, at any tick of the clock m , the high-level controller will provide a function for the low-level controller. This means that x_d will be

doubly-indexed in our formulas. The first index corresponds to the clock evaluation, and the second to the time argument of the trajectory. For example, the ℓ formula $\rho\ell \cdot \|x_d(0)(T)\| < K$ evaluates to $\|x_d(\tau_\ell^m(0))(T)\| < K$. Finally, an estimation block running on clock ℓ will output a state estimate $(\hat{x}, \mathcal{X}, \ell)$. Our system's objective is $\rho\ell \cdot \Diamond\Box(\mathcal{V}(x) = 0)$, i.e., our cost will stabilize at zero.

A. Specifying the System

We consider the contracts for each layer in the system.

1) High-Level Controller: We consider the specification of the high-level controller running on a dedicated processor with clock m . The high layer has an input \hat{x} , the state estimate coming from another control layer, and a single output x_d , which is a trajectory that the low-level controller has to follow.

On a given tick of the clock m , the high-level block will assume that the state estimate is an accurate approximation of the state. The high-level controller uses this to compute a trajectory in the next tick of m . Finally, the high layer will guarantee that the trajectories it generates either decrease the cost \mathcal{V} or keep it equal to zero. We can represent the high-level controller contract $\mathcal{C}^{HL} = (A^{HL}, G^{HL})$ as follows:

$$\begin{aligned} A^{HL} : & \phi_{A_init}^{HL} \wedge \phi_{\text{timing}}^{HL} \wedge \phi_{\text{sensor}}^{HL} \wedge \phi_{\text{bound_var}}^{HL} \\ \phi_{A_init}^{HL} : & \rho m \cdot \text{Close}(x, x_i; \delta_{A_init}^{HL}) \\ \phi_{\text{timing}}^{HL} : & \rho m \cdot \Box \left(\begin{array}{l} (T_{\min}^m \leq r(1) - r \leq T_{\max}^m) \wedge \\ (r - \ell^r < T_{\text{fresh}}^m) \end{array} \right) \\ \phi_{\text{sensor}}^{HL} : & \rho m \cdot \Box \text{Close}(\hat{x}, x; \delta_{\text{sensor}}^{HL}) \\ \phi_{\text{bound_var}}^{HL} : & \rho m \cdot \Box \text{BoundedVariation}(x; D_x) \\ G^{HL} : & \phi_{G_init}^{HL} \wedge \phi_{\text{rsp_dyn}}^{HL} \wedge \phi_{\text{dynamics}}^{HL} \wedge \phi_{\text{traj_bound_var}}^{HL} \wedge \phi_{\text{progress}}^{HL} \\ \phi_{G_init}^{HL} : & \rho m \cdot \text{Close}(x_d(0)(0), x_i; \delta_{G_init}^{HL}) \\ \phi_{\text{rsp_dyn}}^{HL} : & \rho m \cdot \Box \text{RespectDynamics}(x_d) \\ \phi_{\text{dynamics}}^{HL} : & \rho m \cdot \Box \text{Close}(x_d(0)(T_{\text{avg}}^m), x_d(1)(0); \delta_{\text{dynamics}}^{HL}) \\ \phi_{\text{traj_bound_var}}^{HL} : & \rho m \cdot \Box \text{BoundedVariation}(x_d; D_d) \\ \phi_{\text{progress}}^{HL} : & \rho m \cdot \Box \left(\begin{array}{l} (\mathcal{V}(x_d(0)(0)) > \mathcal{V}(x_d(1)(0))) \vee \\ \Box(\mathcal{V}(\text{Inflate}(\text{Im}(x_d); \delta_{\text{progress}}^{HL})) = 0) \end{array} \right) \end{aligned}$$

$\Phi_{\text{sensor}}^{HL}$ captures the requirement that the sensor produces values that are close to the real state x at the time when m ticks. The predicate $\text{Close}(v, v'; \delta)$ indicates that two symbols v, v' are evaluated to quantities that are close up to a parameter δ for some notion of distance (taken to be ℓ_2 for the case study). $\Phi_{\text{bound_var}}^{HL}$ assumes that the state x has bounded variation. $\Phi_{\text{timing}}^{HL}$ requires the clock period of m to lie between T_{\min}^m and T_{\max}^m ; (we also assume that the nominal m period T_{avg}^m lies between these bounds); the assumption also requires the difference in physical time between a tick of m and the tick of ℓ from which m reads ℓ 's data to be less than T_{fresh}^m . This means that the data from ℓ that is read from m cannot be too old. $\Phi_{A_init}^{HL}$ requires x to be close to a value x_i at the beginning of the system execution.

With respect to guarantees, $\Phi_{G_init}^{HL}$ means that the first trajectory output by the high-level block will start close to x_i . $\Phi_{\text{rsp_dyn}}^{HL}$ is a predicate stating that the high-level controller will always generate trajectories x_d that should be within the competence of the low-level block to follow. In the context

of this letter, we require that the produced trajectories are dynamically feasible, i.e., there exists a feedback controller which is able to track the generated trajectories. $\Phi_{\text{dynamics}}^{HL}$ states that the beginning of every new trajectory provided by the high-level block has to be close to the value of the previous trajectory at time T_{avg}^m . $\Phi_{\text{traj_bound_var}}^{HL}$ is a promise that the trajectories provided by the high-level controller will have bounded variation.

Finally, in order to be able to promise that the system will make progress towards its objective, the high-level block will make use of the cost function \mathcal{V} . The high-level controller makes a guarantee $\Phi_{\text{progress}}^{HL}$ which says that either the cost at the beginning of a trajectory x_d is larger than the cost at the beginning of the next trajectory x_d , or that all points that are close to the trajectory have a cost of zero. The function $\text{Inflate}(A, \delta)$ takes a set A and returns the set of all points that are δ -close to any point of A . That is, each trajectory either improves the cost or stays fixed at cost equal to zero.

2) Low-Level Controller: The low-level controller takes as inputs trajectories x_d that the system's state x has to follow and promises that it can make the system follow these trajectories with a given accuracy. We have the following contract $\mathcal{C}^{LL} = (A^{LL}, G^{LL})$ for the low-level controller:

$$\begin{aligned} A^{LL} : & \phi_{\text{timing}}^{LL} \wedge \phi_{\text{rsp_dyn}}^{LL} \wedge \phi_{\text{dynamics}}^{LL} \wedge \phi_{\text{bound_var}}^{LL} \\ \phi_{\text{timing}}^{LL} : & \rho \ell \cdot \Box \left(\begin{array}{l} (T_{\min}^\ell \leq r(1) - r \leq T_{\max}^\ell) \wedge \\ (r - m^r < T_{\text{fresh}}^\ell) \end{array} \right) \\ \phi_{\text{rsp_dyn}}^{LL} : & \rho \ell \cdot \Box \text{RespectDynamics}(x_d) \\ \phi_{\text{dynamics}}^{LL} : & \rho \ell \cdot \Box \left(\begin{array}{l} (m \neq m(-1)) \wedge (m \geq 0) \\ \rightarrow \text{Close}(x_d(0)(0), x; \delta_{\text{dynamics}}^{LL}) \end{array} \right) \\ \phi_{\text{bound_var}}^{LL} : & \rho \ell \cdot \Box \text{BoundedVariation}(x; D_x) \\ G^{LL} : & \phi_{\text{upd}}^{LL} \wedge \phi_{\text{tracking}}^{LL} \\ \phi_{\text{upd}}^{LL} : & \rho \ell \cdot \Box \left(\begin{array}{l} \text{if } (m \neq m(-1)) \wedge (m \geq 0) \text{ then} \\ (\text{upd} = \ell) \text{ else } (\text{upd} = \text{upd}(-1)) \end{array} \right) \\ \phi_{\text{tracking}}^{LL} : & \rho \ell \cdot \Box \left(\begin{array}{l} (\ell - \text{upd} > 0) \rightarrow \\ \text{Close}(x, x_d(0)(T_{\text{avg}}^\ell(\ell - \text{upd})); \delta_{\text{tracking}}^{LL}) \end{array} \right) \end{aligned}$$

The assumptions of this contract are as follows. $\phi_{\text{timing}}^{LL}$ assumes that ℓ 's period is bounded below and above, and that the data read from clock m is not too old. $\phi_{\text{rsp_dyn}}^{LL}$ requires trajectories received from m to respect certain physical limits, i.e., that they satisfy the state and input constraints used in the analysis of the low level controller. $\phi_{\text{dynamics}}^{LL}$ makes sure that when ℓ detects the generation of a new trajectory x_d from m , then the starting point of that trajectory should be close to the value of the state. $\phi_{\text{bound_var}}^{LL}$ requires the state x to have bounded variation.

Regarding guarantees, ϕ_{upd}^{LL} is a helper statement that defines the variable upd . This variable contains the last value of ℓ when a new trajectory was received from m . $\phi_{\text{tracking}}^{LL}$ guarantees that the low-level controller will make the system follow the given trajectory x_d . Observe that $\phi_{\text{tracking}}^{LL}$ is enforced for all values of the trajectory x_d , except its first point. For the first point of the trajectory, the low level controller makes the assumption $\phi_{\text{dynamics}}^{LL}$ on the state. T_{avg}^ℓ is the nominal period of ℓ and respects the bounds of $\phi_{\text{timing}}^{LL}$.

3) Estimator: The estimator will guarantee that the state estimates are always accurate on the clock ticks of ℓ , yielding the contract $\mathcal{C}^{Est} = (A^{Est}, G^{Est})$, where $A^{Est} : \rho\ell$. True and $G^{Est} : \phi_{sensor}^{Est}$, with $\phi_{sensor}^{Est} : \rho\ell$. $\square \text{Close}(\hat{x}, x; \delta_{sensor}^{Est})$.

4) Timing Design: We understand the timing component of the system as “network design” in the sense that enforcing timing constraints is implemented by applying networking and clock synchronization technologies. This component will have the following contract $\mathcal{C}^{Tmg} = (A^{Tmg}, G^{Tmg})$:

$$\begin{aligned} A^{Tmg} &: \rho r \cdot \text{True} \\ G^{Tmg} &: \phi_{\ell-timing}^{Tmg} \wedge \phi_{m-timing}^{Tmg} \\ \phi_{\ell-timing}^{Tmg} &: \rho\ell \cdot \square \left(\left(T_{min}^{\ell} \leq r(1) - r \leq T_{max}^{\ell} \right) \wedge \left(r - m^r < T_{fresh}^{\ell} \right) \right) \\ \phi_{m-timing}^{Tmg} &: \rho m \cdot \square \left(\left(T_{min}^m \leq r(1) - r \leq T_{max}^m \right) \wedge \left(r - \ell^r < T_{fresh}^m \right) \right) \end{aligned}$$

B. System-Level Analysis

We have component-level contracts for all layers in our system. The following result connects these specifications with the desired top-level stability objective.

Theorem 1: Consider a two-layer control system formed by a high-level controller and a low-level controller. If each layer satisfies contracts \mathcal{C}^{HL} and \mathcal{C}^{LL} , respectively, the estimator satisfies \mathcal{C}^{Est} , and their interconnection satisfies \mathcal{C}^{Tmg} , then the top-level system will satisfy the top level contract (A, G) , where

$$A : \phi_{A_init}^{HL} \wedge \phi_{bound_var}^{HL} \quad G : \rho\ell \cdot \diamond \square(\mathcal{V}(x) = 0), \quad (2)$$

provided that the following conditions hold:

$$\delta_{sensor}^{Est} + T_{fresh}^m D_x \leq \delta_{progress}^{HL} \quad (3)$$

$$\delta_{G_init}^{HL} + \delta_{A_init}^{HL} + D_x T_{fresh}^{\ell} \leq \delta_{dynamics}^{LL} \quad (4)$$

$$\delta_{tracking}^{LL} + \delta_{dynamics}^{HL} + (T_{fresh}^m + T_{fresh}^{\ell}) D_x + D_d \Delta T^m \leq \delta_{dynamics}^{LL} \quad (5)$$

$$\delta_{dynamics}^{LL} \leq \delta_{progress}^{HL}, \quad (6)$$

where $\Delta T^m = T_{max}^m - T_{avg}^m \lfloor \frac{T_{min}^m - (T_{fresh}^m + T_{fresh}^{\ell})}{T_{max}^{\ell}} \rfloor$.

Proof: Contract composition yields the specification formed by interconnecting the components whose specifications we have available. Our system-level specification is $\mathcal{C}^{Sys} = \mathcal{C}^{HL} \parallel \mathcal{C}^{LL} \parallel \mathcal{C}^{Est} \parallel \mathcal{C}^{Tmg}$, which we compute by applying the ContractComposition routine of [8]. We now verify how this composition yields a system satisfying the assumptions of all components.

a) High-Level Controller: We consider the assumptions: ϕ_{init}^{HL} is an assumption on the initial state that we need to conserve at the system-level. $\phi_{bound_var}^{HL}$ is also a system-level assumption. ϕ_{timing}^{HL} is satisfied by $\phi_{m-timing}^{Tmg}$ of our timing design. ϕ_{sensor}^{HL} needs to be analyzed. We observe that¹

$$\frac{\phi_{sensor}^{Est} \phi_{m-timing}^{Tmg} \phi_{bound_var}^{HL}}{\rho m \cdot \text{Close}(\hat{x}, x; \delta_{sensor}^{Est} + T_{fresh}^m D_x)}$$

We need the relation (3) to hold in order to satisfy ϕ_{sensor}^{HL} .

¹This notation means that the formula below the horizontal bar is a valid deduction from the conjunction of the formulas on top of the bar.

b) Low-Level Controller: ϕ_{timing}^{HL} is satisfied by $\phi_{\ell-timing}^{Tmg}$ of the timing design. $\phi_{bound_var}^{HL}$ is an assumption on the dynamics of the physical system and should therefore be a system-level assumption. $\phi_{rsp_dyn}^{LL}$ is satisfied by the guarantee $\phi_{rsp_dyn}^{HL}$. The satisfaction of $\phi_{dynamics}^{LL}$ requires an inductive argument. For the proof that relation (4) is required for $\phi_{dynamics}^{LL}$ to hold in the initial case and that $\phi_{dynamics}^{LL}$ holds inductively when (5) is true, please see Section VI-B2 of the extended version of this letter [7].

c) Progress: Our analysis indicates that the composition of all contracts so far defined yields a situation in which all contracts have their assumptions met—provided the system parameters meet conditions (3), (4), and (5). Now we verify whether the system makes progress towards its goal. We observe that the fact that \mathcal{V} takes values in a well order means that we can carry out the following deductions:

$$\delta_{progress}^{HL}$$

$$\begin{aligned} \rho m \diamond \square(\mathcal{V}(\text{Inflate}(\text{im}(x_d); \delta_{progress}^{HL})) = 0) &\quad \phi_{progress}^{HL} \phi_{upd}^{LL} \phi_{tracking}^{LL} \phi_{dynamics}^{LL} \\ \rho\ell \cdot \diamond \square(\mathcal{V}(x) = 0) \end{aligned}$$

provided that (6) holds. ■

C. Component-Level Verifications

Theorem 1 shows that if we implement control layers adhering to the specifications introduced in Section IV-A and satisfying constraints (3)–(6), the system will satisfy the desired stability property (2). These constraints impose timing requirements in the system. For instance, (5) imposes a limit on the maximum compute time of the high-level controller. This means that constraints (3)–(6) can be used to increase the robustness of the design by ensuring that they are satisfied with margin. Now we verify that each implementation of our control layers satisfies its own contract.

1) Low-Level Controller: In order to use the methods presented in [4] which combine a low level controller with (1), we begin by showing that applying zero order held inputs, i.e., with $\rho\ell \cdot \square(\bar{u} = k_{fbl}(x, \ell - \text{upd}))$, results in bounded exogenous disturbance to the error dynamics. For the following discussion, let $x'(t)$ denote the solution to the system dynamics with continuous time control applied, and $x(t)$ the solution with zero order held inputs. Plugging in the solution to the differential equation with $x'(0) = x(0)$ and adding and subtracting $g(x'(\tau))\bar{u}$ yields

$$\begin{aligned} \|x(t) - x'(t)\| &= \int_0^t \frac{d}{dt}(x(\tau) - x'(\tau)) d\tau \\ &\leq \int_0^t (L_f + L_g \|\bar{u}\|) \|x(\tau) - x'(\tau)\| + \|g(x'(\tau))\| \|u(\tau) - \bar{u}\| d\tau, \end{aligned} \quad (7)$$

where L_f and L_g represent the local Lipschitz constants of the drift vector and actuation matrix, respectively. From the perspective of the low level controller, we assume via $\phi_{rsp_dyn}^{LL}$ that $\|u(t)\| \leq U$ and that $x(t) \in \mathcal{X}$, a compact set. These will be necessary for the analysis, and will be explicitly enforced in the proposed MPC formulation. From these constraints, we know that there exists a $G > 0$ such that $\|g(x'(\tau))\| < G$ and that all local Lipschitz constants are global over \mathcal{X} . Combining these facts and using the

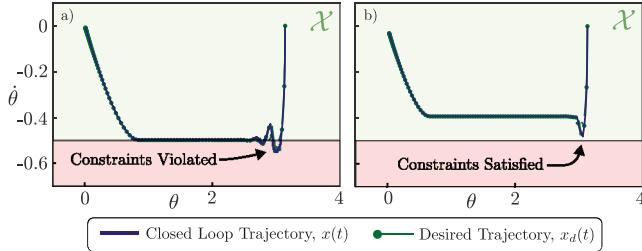


Fig. 2. a) Nominal controller whose design process did not use MCL contract verification fails to meet system objectives. b) Control blocks which adhere to MCL contracts satisfy the constraint, despite being designed in isolation.

Bellman-Gronwall Lemma [11] leads to $\|x(t) - x'(t)\| \leq UGt^2 e^{(L_f + L_g U)t} \triangleq t\rho(t)$, where $\rho \in \mathcal{K}_\infty$, a class \mathcal{K} infinity function. Next, let $e'(t) \triangleq x'(t) - x_d(t)$ denote the error dynamics with continuous time control applied, whereby the feedback linearizing controller k_{fbl} yields $\dot{e}'(t) = A_{\text{cl}}e'$ for A_{cl} a stable matrix. Taking $e(t) = x(t) - x_d(t)$ to be the error with zero order held inputs, we have $\dot{e} = w(t) + A_{\text{cl}}e$, where $w(t) = \frac{d}{dt}(x - x') + A_{\text{cl}}(x' - x)$.

Plugging in the terms developed in (7) results in $\|w(t)\| \leq \|w(T_{\max}^\ell)\| \leq T_{\max}^\ell((L_f + 2L_g U + \|A_{\text{cl}}\|)\rho(T_{\max}^\ell) + GU)$. As this bound is a composition of class- \mathcal{K} functions in time, for all $\delta_w > 0$ there exists an $\epsilon > 0$ such that $T_{\max}^\ell < \epsilon$ results in $\|w(t)\| \leq \delta_w$. Fixing an allowable δ_w and thereby upper bounding T_{\max}^ℓ , by integrating the error dynamics and using the comparison lemma we have $\|e(t)\| \leq \|x(0) - x_d(0)\|Me^{-\lambda t} + T_{\max}^\ell\delta_w$, for some $M, \lambda > 0$ as determined by the convergence rate of the low level controller. Therefore, in order to produce the guarantee of $\phi_{\text{tracking}}^{\text{LL}}$, we must enforce that $\delta_{\text{dynamics}}^{\text{LL}}Me^{-\lambda T_{\min}^\ell} + T_{\max}^\ell\delta_w \leq \delta_{\text{tracking}}^{\text{LL}}$.

2) MPC: Lemma 2 in [4] demonstrates that the trajectories produced by (1) are dynamically feasible, i.e., able to be exactly tracked via the feedback linearizing controller k_{fbl} , implying satisfaction of $\phi_{\text{rsp_dyn}}^{\text{HL}}$. As the MPC program produces solutions to a linear system with bounded state and control inputs, we have that the resulting desired trajectory x_d will have bounded variation, satisfying $\phi_{\text{traj_bound_var}}^{\text{HL}}$.

Next, we ensure that the above MPC program is recursively feasible, which requires showing that the enforced set \mathcal{E} is a robust invariant for the system. Choosing \mathcal{E} such that $B(0, \delta_{\text{tracking}}^{\text{LL}}) \subset \mathcal{E}$ results in a set that can be rendered invariant by the low level controller, even during sampling. Therefore, we choose \mathcal{E} to meet this condition as well as the design requirements imposed by the system analysis in Section IV-B. We then appeal to the results in [4] to prove recursive feasibility of the MPC algorithm used therein.

Finally, let $V: \mathcal{X} \rightarrow \mathbb{R}$ denote the sum of the running and terminal cost of MPC, often used as a Lyapunov function in stability proofs. As we have effectively transformed our nonlinear MPC program to a linear one, we can use standard MPC results [3] to state that $\rho_m \cdot (V(x_d(1)(0)) < V(x_d(0)(0)))$. Since the guarantees of the MPC contract involve \mathcal{V} , which takes values in a well-order, we can define \mathcal{V} by quantizing the function V . Then we can choose a vicinity around x_g and define $\mathcal{V}(p) = 0$ for $p \in \mathcal{E}$. In that case, \mathcal{V} takes values in a well-order, and the MPC block satisfies $\phi_{\text{progress}}^{\text{HL}}$.

3) Simulation Results: We investigate the use of MCL contracts as a design tool towards achieving stability and state constraint satisfaction for an architecture consisting of a high-level and a low-level control block designed in isolation. In both cases, the same control architecture and time delay of $T_{\text{fresh}}^\ell = T_{\max}^m = T_{\min}^m$ was used, i.e., the time delay is equal to one m clock cycle, and the m clock does not have jitter. In the first case shown in Figure 2(a), the MPC and FBL controllers were independently implemented, but their interconnection was not verified against formal specifications. Therefore, our analysis did not yield any guarantees of system behavior, and indeed we find that our system violates the desired state constraints $x(t) \in \mathcal{X}$. In Figure 2(b), the margin \mathcal{E} was modified such that (5) holds with equality. As a result of adhering to the MCL contracts, the composite control hierarchy maintains the desired system-level specification in (2) despite having independently designed components.

V. CONCLUDING REMARKS

We considered an effective way of specifying control layers running on multiple clocks using MCL. Once the specifications were expressed, we used MCL contracts to prove system-level stability properties using the local properties of each layer. Our framework is extensible to systems with additional layers. We showed in the case study how to verify that each control block satisfies its own specification. Our simulation results showed that a violation of the system-level constraints led to failure. Future work involves applications to more complex systems, such as bipedal robots.

REFERENCES

- [1] F. Allgower, R. Findeisen, and Z. K. Nagy, "Nonlinear model predictive control: From theory to application," *J. Chin. Inst. Chem. Eng.*, vol. 35, no. 3, pp. 299–316, 2004.
- [2] Z. Artstein, "Stabilization with relaxed controls," *Nonlin. Anal. Theory, Methods App.*, vol. 7, no. 11, pp. 1163–1173, 1983.
- [3] F. Borrelli, A. Bemporad, and M. Morari, *Predictive Control for Linear and Hybrid Systems*. Cambridge, U.K.: Cambridge Univ. Press, 2017.
- [4] N. Csomay-Shanklin, A. J. Taylor, U. Rosolia, and A. D. Ames, "Multi-rate planning and control of uncertain nonlinear systems: Model predictive control and control Lyapunov functions," Mar. 2022, *arXiv:2204.00152*.
- [5] R. Grandia, F. Jenelten, S. Yang, F. Farshidian, and M. Hutter, "Perceptive locomotion through nonlinear model predictive control," Aug. 2022, *arXiv:2208.08373*.
- [6] I. Incer, "The algebra of contracts," Ph.D. dissertation, Dept. Electr. Eng. Comput. Sci., Univ. California, Berkeley, CA, USA, May 2022.
- [7] I. Incer, N. Csomay-Shanklin, A. Ames, and R. M. Murray, "Specifying and analyzing networked and layered control systems operating on multiple clocks," 2024, *arXiv:2402.11666*.
- [8] I. Incer et al., "Pacti: Scaling assume-guarantee reasoning for system analysis and design," 2023, *arXiv:2303.17751*.
- [9] E. Lee and A. Sangiovanni-Vincentelli, "A framework for comparing models of computation," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 17, no. 12, pp. 1217–1229, Dec. 1998.
- [10] U. Rosolia and A. D. Ames, "Multi-rate control design leveraging control barrier functions and model predictive control policies," *IEEE Control Syst. Lett.*, vol. 5, no. 3, pp. 1007–1012, Jul. 2021.
- [11] S. Sastry, "Mathematical background," in *Nonlinear Systems: Analysis, Stability, and Control* (Interdisciplinary Applied Mathematics), S. Sastry, Ed. New York, NY, USA: Springer, 1999, pp. 76–126.
- [12] E. D. Sontag, "A 'universal' construction of Artstein's theorem on nonlinear stabilization," *Syst. Control Lett.*, vol. 13, no. 2, pp. 117–123, 1989.