Regret-Optimal Cross-Layer Co-Design in Networked Control Systems—Part I: General Case

Mohammad H. Mamduhi, *Senior Member, IEEE*, Dipankar Maity, *Senior Member, IEEE*, Karl H. Johansson, *Fellow, IEEE*, and John Lygeros, *Fellow, IEEE*

(Invited Paper)

***Abstract*— Performance of control systems interacting over a shared communication network is tightly coupled with how the network provides services and distributes resources. Novel networking technology such as 5G is capable of providing tailored services for a variety of network demands. Stringent control requirements and their critical performance specification call for online adaptable and control-aware network services. This perspective suggests a co-design of physical and network layers aiming to ensure that the necessary quality-of-service is provided to achieve the desired quality-of-control. An optimal co-design is in general challenging due to cross-layer couplings between the physical and network layers and their layer-specific functionalities. Furthermore, the complexity of the co-design depends on the level of actionable information the layers share with each other. In this Part I of a two-letter series, we present a general co-design of physical operations and service allocation aiming to minimize a *social regret* measure for networked control systems. We introduce an optimal networked co-design scenario using the regret index as the joint quality-of-control and quality-of-service (QoC-QoS) measure, and discuss the role of cross-layer awareness in the structure of optimization problems. We mainly focus on the finite-horizon case but we briefly present the infinite-horizon case as well. In Part II, we discuss regret-optimal cross-layer policies for Gauss-Markov systems and derive the optimal solutions based on the general problems introduced in Part I.**

***Index Terms*— Regret-optimal co-design, networked systems, service allocation, QoC-QoS trade-off, awareness structure.**

I. Introduction

NETWORKED Control Systems (NCS) are generally characterized by many physical devices that use a common network infrastructure to interact and exchange information between sensors, controllers and actuators. [1]. The exchanged information can take various forms; for example, a control signal to close a feedback loop for a controlled task, some remotely measured data that needs to be transmitted to an observer, or a live data feed generated by a digital twin for a real-time localization task [2]. Depending on the purpose, end-to-end data transmission may need to be performed with specific Quality-of-Service (QoS), e.g., transmission reliability, latency, and bandwidth [3]. This is in line with the user-oriented serviceability of 5G mobile technology, which is capable of providing massive connectivity and parallel communication services, each with tailored QoS specifications, known as *network slices* [4]. Furthermore, in time and safety-critical NCS applications, the desired operations and objectives are context-dependent, specially across the physical layer. Therefore, quality of physical layer performance, referred to as Quality-of-Control (QoC), should be guaranteed over finite time intervals and not only asymptotically. The performance may also vary over time, hence, finite-horizon performance requirement might be judged by different measures [5], [6]. Since QoC is coupled with QoS, the network should be reconfigurable to adapt its services for the physical layer's needs [7], [8]; modern communication technology provides this through the virtual and software-defined networking capabilities [9].

Although 5G networks have radically improved service portfolio compared to the Long-Term Evolution era, efficient service management is still a challenge. In a resource-limited scenario where all users QoS demands may not be met, the service provider needs to decide how to allocate the limited resources. To do so, the service provider needs to have certain knowledge of the QoC needs of the users for which the communication service is being allocated. Providing coordinated network access with tailored QoS (optimal QoC-QoS trade-off) typically requires a control-communication co-design framework where the effects of QoS are explicitly quantified in terms of control performance, which is a challenging problem especially for a highly dynamic and interactive physical layer with time-varying objectives and non-stationary constraints.

In this letter, we study the optimal QoC-QoS trade-off for NCS consisting of multiple heterogeneous systems that use a common adaptable network to exchange information. The QoC-QoS trade-off is associated with the combined control and communication cost. Generally higher QoS leads to better control performance but at the expense of a higher communication cost. Therefore, a trade-off exists between using enhanced QoS to reduce control cost, or tolerate higher control cost but paying less for worse QoS. We introduce the concept of regret and define a cross-layer regret function to measure the QoC-QoS trade-off. We also discuss the role of cross-layer awareness in structural properties of the regret-optimal policies, and show that a higher level of cross-layer awareness results in an improved performance at the expense of higher

Manuscript received 2 July 2023; accepted 17 August 2023. Date of publication 29 August 2023; date of current version 9 November 2023. The work of Dipankar Maity was supported by ARL grant ARL DCIST CRA W911NF-17-2-0181. The work of Mohammad H. Mamduhi and John Lygeros was supported by the Swiss National Science Foundation, NCCR Automation Grant 180545. The work of Karl H. Johansson was supported by Swedish Research Council Distinguished Professor Grant 2017-01078, Knut and Alice Wallenberg Foundation Wallenberg Scholar Grant, and Swedish Strategic Research Foundation SUCCESS Grant FUS21-0026. The associate editor coordinating the review of this letter and approving it for publication was N. Pappas. *(Corresponding author: Mohammad H. Mamduhi.)*

Mohammad H. Mamduhi and John Lygeros are with the Automatic Control Laboratory, ETH Zürich, 8092 Zürich, Switzerland (e-mail: mmamduhi@ethz.ch; jlygeros@ethz.ch).

Dipankar Maity is with the Department of Electrical and Computer Engineering, The University of North Carolina at Charlotte, Charlotte, NC 28223 USA (e-mail: dmaity@uncc.edu).

Karl H. Johansson is with the Division of Decision and Control Systems, KTH Royal Institute of Technology, 114 28 Stockholm, Sweden, and also with Digital Futures 100 44 Stockhom, Sweden (e-mail: kallej@kth.se).

Digital Object Identifier 10.1109/LCOMM.2023.3309880

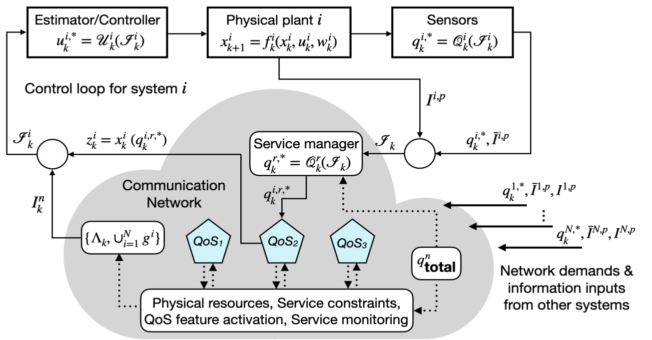

Fig. 1. NCS with N control loops and a shared-service adaptive network with three QoS-tailored channels/slices. Dotted arrows show in-network connections, and solid arrows connections that involve physical systems.

solution complexity [10]. While focus is on the finite-horizon case, we briefly present the co-design over the infinite-horizon.

II. NCS Model and Cross-Layer Performance

Consider an NCS with N physical control systems whose dynamics are expressed by the discrete-time stochastic model

$$x_{k+1}^i = f_k^i(x_k^i, u_k^i, w_k^i), \tag{1}$$

where $i \in \{1, \ldots, N\}$ is the system index, $x_k^i \in \mathbb{R}^{n_i}, u_k^i \in \mathbb{R}^{m_i}, w_k^i \in \mathbb{R}^{n_i}$, represent the states, control inputs, and stochastic exogenous disturbances of system i, respectively, and $k \in \{0, 1, 2, \ldots\}$ is the discrete time index, see Fig. 1. The initial state $x_0^i \in \mathbb{R}^{n_i}$ is randomly selected from an arbitrary distribution with finite covariance matrix $\Sigma^i \succ 0$. The initial states of different systems are mutually independent also with respect to the disturbances w_k^i, for all i and k; the disturbances are assumed to be independent for different i and k and identically distributed for the same i and different k.

Each system i has sensors to measure the states x_k^i that are transmitted to their corresponding controllers via a communication network. The measurements received by the controller at time k is denoted z_k^i, where the content of z_k^i depends on the network QoS. For example, we may set $z_k^i = x_{k-d}^i$ if the whole state is communicated but there is a delay $0 \leq d$ between the time the measurement is taken and when it arrives at the controller (For notational convenience, we set $x_{-1}^i = x_{-2}^i = \ldots = \varnothing, \forall i$), or $z_k^i = \Delta(x_k^i)$ if the state measurement is quantized using a quantization function $\Delta(\cdot)$, or $z_k^i = \{x_{k-d}^i, \Delta(x_k^i)\}$ if two state measurements simultaneously arrive at the controller at time k, one with a d-step delay and the other non-delayed but quantized, or $z_k^i = \varnothing$ if no information is received at time k due, for example, to packet dropout. Hence, the network QoS during the interval $[0, k]$ determines the available state information $\mathcal{Z}_k^i \triangleq \{z_0^i, z_1^i, \ldots, z_k^i\}$ at the controller at time k.

To measure NCS performance, various index functions can be used, each emphasizing a certain physical component or parameter, without considering the network layer parameters for QoS provision. For example, Age-of-Information optimization maximizes the temporal freshness of sample updates [11], Value-of-Information optimizes data distribution with respect to the resulting cost improvement [12], risk-sensitive cost functions (and its Linear-quadratic Gaussian and H_∞ variations) minimize the weighted values of state and control inputs, and discounted cost functions maintain a trade-off between short and long-term performance [13]. On the network side, metrics such as latency, reliability, throughput, and congestion are often indicators of QoS without appropriately quantifying their effects on the key performance measures of the physical layer. To study joint performance of control and communication in NCS, it is then essential to perform co-designs using measures that involve parameters and policies of both layers.

Regret measures performance discrepancy between *ideal* and *realized* scenarios [14]. In our setup, the *ideal scenario* refers to the case where the QoS provided by the communication network exactly matches the QoS demands of the physical systems, while the *realized scenario* pertains to the actual QoS served by the network. To maintain a satisfactory collective performance, the regret function in this letter measures the social performance discrepancy [10], averaged across all components and all time-steps over finite or infinite horizons. The performance discrepancy results from the mismatch between QoS demand and provision, which the discrepancy vanishes if network QoS matches exactly with the QoS demands.

Finite Horizon Case: Reconfigurable networks allow users to demand certain levels of QoS and the network can adapt its services to meet such demands as close as possible. Let q_k^i be the desired QoS for system i to transmit the state x_k^i from the sensor to the controller at time k over a network with multiple service channels with tailored QoS features. For a system i, the joint control–communication cost over an interval $[T_s, T_f]$ is measured by the following local cost function

$$J_{[T_s,T_f]}^i(u^i, q^i) = \sum_{k=T_s}^{T_f} \mathsf{E}\left[c_k^i(x_k^i, u_k^i, w_k^i) + g_k^i(q_k^i)\right], \tag{2}$$

where the stage cost c_k^i measures the control cost per time-step, and $g_k^i(q_k^i)$ denotes the incurred communication service usage cost at time k, which depends on the requested QoS q_k^i. Generally, function c_k^i is non-decreasing with increasing the size of the variables x_k^i, u_k^i, w_k^i, and g_k^i increases for more enhanced requested QoS q_k^i. There is a causal trade-off between the cost terms c_k^i and g_k^i, since more available data together with enhanced QoS (higher g_k^i) lead generally to improved control performance (lower c_k^i).

At time $t \in [T_s, T_f]$, computation of the cost in (2), involves two separate parts; the realized cost over $[T_s, t)$, and the cost-to-go over $[t, T_f]$. The cost-to-go is

$$J_{[t,T_f]}^i(u^i, q^i) = \sum_{k=t}^{T_f} \mathsf{E}\left[c_k^i(x_k^i, u_k^i, w_k^i) + g_k^i(q_k^i)\right], \tag{3}$$

where we use $X_{[t_1,t_2]}$ to denote $\{X_{t_1}, X_{t_1+1}, \ldots, X_{t_2}\}$. A difference between the two parts is in the realization of the random parameters and the decision variables. At $t > T_s$, the random variables $w_{[T_s,t-1]}^i$ and $x_{T_s}^i$ are known. Therefore, the expectation operator does not act on the part of the cost function (2) that corresponds to $[T_s, t)$. In (3) the expectation is with respect to $w_{[t,T_f]}^i$, and $x_{T_s}^i$, only if $t = T_s$.

From the cost-to-go (3), one can compute optimal control inputs and desired QoS demands from a current time t up to T_f, subject to cross coupling with the network layer. We denote the policies from which control inputs and QoS demands are computed by $\mathcal{U}_k^i$ and $\mathcal{Q}_k^i$, respectively, such that

$$u_k^i = \mathcal{U}_k^i(\mathcal{I}_k^i), \quad q_k^i = \mathcal{Q}_k^i(\mathcal{I}_k^i), \quad k \in [T_s, T_f], \tag{4}$$

where $\mathcal{I}_k^i$ is the available information for system i at time k, further discussed in Sec. III. Under such partial information $\mathcal{I}_k^i$, often a state observer is required to estimate state evolution. In fact, we assume that the dynamic policy $\mathcal{U}_k^i$ is given by a combined controller and observer policy generating control inputs as functions of the state beliefs computed by local observers. We discuss the derivation of observers under partial information in Part II for Gauss-Markov systems and show that a set of certainty equivalence controllers with nonlinear state estimators form the optimal control and observer policy.

The optimal control inputs and QoS demands of system i can be obtained for $k \geq t$, by solving (3) as follows:

$$\begin{aligned} &\min_{\mathcal{U}_k^i, \mathcal{Q}_k^i} \; J^i_{[t,T_f]}(u^i, q^i), \\ &\text{subject to } \; u_k^i = \mathcal{U}_k^i(\mathcal{I}_k^i), \;\; q_k^i = \mathcal{Q}_k^i(\mathcal{I}_k^i), \\ &\qquad x_{k+1}^i = f_k^i(x_k^i, u_k^i, w_k^i), \;\; k \in [t, T_f]. \end{aligned} \tag{5}$$

Problem (5) is locally computed to derive $u_k^{i,*}, q_k^{i,*}$ that jointly minimize the QoC-QoS cost $J^i_{[t,T_f]}(u^i, q^i)$ in (3). Solving (5) at all iterations leads to the optimal cost for (2).

Since network services are distributed among multiple systems with possibly diverse QoS demands, a network service manager runs service allocation to handle service constraints and resource limitations, see Fig. 1. This may lead to mismatches between QoS demands and provisions that eventually influence closed-loop performance. Indeed, (2) is the *ideal cost* over the interval $[T_s, T_f]$ from the i-th system perspective without knowing the other systems' service demands. The *realized scenario* occurs after QoS distribution has been decided at the network. We define the realized cost for system i, as

$$J^{i,r}_{[T_s,T_f]}(u^{i,r}, q^{i,r}) = \sum_{k=T_s}^{T_f} \mathsf{E}[c_k^i(x_k^i, u_k^{i,r}, w_k^i) + g_k^i(q_k^{i,r})], \tag{6}$$

where $q_k^{i,r}$ is the network service allocated to system i at time k, and $g_k^i(q_k^{i,r})$ is the associated communication cost. Similar to the policies introduced for the physical systems in (4), we assume $q_k^{i,r} = \mathcal{Q}_k^r(\mathcal{I}_k)$, with $\mathcal{I}_k$ being the available information at the network to decide on QoS allocation. Note that the control input u_k^i is coupled with the network service: if $q_k^{i,r} \neq q_k^i$, then $u_k^{i,r}$ may not equal u_k^i, which leads in general to $c_k^i(\cdot, u_k^{i,r}, \cdot) \neq c_k^i(\cdot, u_k^i, \cdot)$. Clearly, if $q_k^{i,r} = q_k^i, \forall k \in [T_s, T_f]$, then $J^{i,r}_{[T_s,T_f]} = J^i_{[T_s,T_f]}$. Note that, the system state x_{k+1}^i changes with changing the control input u_k^i, according to (1).

Similar to the cost index (2), the cost in (6) includes two parts: one during $[T_s, t)$ and one $[t, T_f]$. The realized cost $J^{i,r}_{[T_s,t)}$ is computed at time t without the expectation operator in (6). For $J^{i,r}_{[t,T_f]}$, we notice that at time t, system i computes the optimal decision variables $u^{i,*}_{[t,T_f]}, q^{i,*}_{[t,T_f]}$. In response, the network decides on the allocated QoS by computing $q^{i,r,*}_{[t,T_f]}$, for each i. Therefore, at time t, system i can predict its cost-to-go $J^{i,r}_{[t,T_f]}$ given $q^{i,r,*}_{[t,T_f]}$. Note that the future control inputs $u^{i,r}_{[t,T_f]}$ can be estimated since system i is aware of the future allocated services $q^{i,*}_{[t,T_f]}$ and the model-based estimation of system states $x^i_{[t,T_f]}$ using the system model in (1). The expectation in (6) is thus with respect to the random variables $w^i_{[t,T_f]}$ and $u^{i,r}_{[t,T_f]}$. Finally, note that the decision variables $q^{i,r,*}_{[t,T_f]}$ are generated in response to $q^{i,*}_{[t,T_f]}$, and they may change in the future, if system i decides to change its QoS demands due to $q^{i,*}_{[t,T_f]}$ and $q^{i,r,*}_{[t,T_f]}$ mismatches. Therefore, the cost-to-go needs to be evaluated in real-time.

We introduce local and social regret functions, the latter forms the basis for the network to perform socially-fair resource allocation, while with the former physical systems may adapt their future QoC expectations and demands. The local and social regrets over the interval $[T_s, T_f]$ are defined as

$$R^i_{[T_s,T_f]} = J^{i,r}_{[T_s,T_f]} - J^{i,*}_{[T_s,T_f]}, \tag{7}$$

$$R^s_{[T_s,T_f]} = \frac{1}{N} \sum_{i=1}^{N} v^i(R^i_{[T_s,T_f]}), \tag{8}$$

where $v^i(\cdot)$ is a known function that determines how the local regrets affect the social regret, for example, $v^i(R^i) = e^{R^i}$ or $\beta^i R^i$, with $\beta^i \in (0,1)$ a known weight assigned to system i.

III. Regret Optimization and Awareness Structure

The index functions (2), (3), (6) and the associated regret functions (7)–(8) all involve cross-layer coupling. Hence, optimizing joint performance across physical and network layers depends on the cross-layer awareness model. Technically, the optimal control and QoS demands computed at local system i depend on $\mathcal{I}_k^i$, and the optimal network service allocation depends on $\mathcal{I}_k$. Hence, the information the two layers incorporate for the co-design influence the optimal QoC-QoS tradeoff. Generally, higher awareness leads to improved decision making and performance of corresponding functionalities, but entails higher communication and computation cost to extract and analyze models and policies from data.

A. Awareness Structure and Optimal Co-Design

We define three layer-specific information sets that the corresponding layer may share with the other layer. At the physical layer for system i, the set of parametric information ($I^{i,p}$), the set of dynamic information ($\bar{I}_t^{i,p}$) at time t, and, at the network layer, the set of dynamic network information at time t (I_t^n), are defined, respectively, as

$$I^{i,p} = \{T_s, T_f, W^i, \Sigma^i, v^i, \epsilon_q^i, \epsilon_u^i\}, \tag{9}$$

$$\bar{I}_t^{i,p} = \{q^i_{[t,T_f]}, J^i_{[t,T_f]}, J^{i,r}_{[t,T_f]}\}, \tag{10}$$

$$I_t^n = \{q^{i,r}_{[t,T_f]}, g^i_{[t,T_f]}\}, \tag{11}$$

where $W^i \succ 0$ and $\Sigma^i \succ 0$ are, respectively, the covariance matrices of distributions for the uncertainty w_k^i and the initial condition x_0^i, for all k, ϵ_u^i denotes the constraint sets for control inputs; for example, it implies a saturation constraint to bound the size of the control signals, and ϵ_q^i denotes the constraint sets for QoS requirements; for example, only medium or low-latency services should be requested for system i. The minimum information received from the network layer at time t incorporated into the physical layer co-design is $g^i_{[t,T_f]}$. The knowledge about the physical layer received at the network at time t includes at least the QoS demands $q^i_{[t,T_f]}$, for all i. We assume that the communications of information $q^i_{[t,T_f]}$ from any system i to the network manager renders error-free.

1) Online Awareness — Reactive Physical Layer: Here we consider online cross-layer awareness with respect to the information sets (9)–(11). The network shares I_t^n with the physical layer at time t, therefore, the information set $\mathcal{I}_t^i$ is updated as $\mathcal{I}_t^i = \mathcal{I}_{t-1}^i \cup \{z_t^i, u_{t-1}^i, \bar{I}_t^{i,p}, I_t^n\}$, and $\mathcal{I}_{T_s}^i = \{I^{i,p}, z_{T_s}^i, g_{T_s}^i\}$. Note that $\bar{I}_t^{i,p}$ and I_t^n contain information on future variables that will over-write their realizations in $\mathcal{I}_{t-1}^i$, for example, the information $q^i_{[t,T_f]}$ exists in $\mathcal{I}_{t-1}^i$ (see (10)) will be replaced by the update from $\bar{I}_t^{i,p}$. Since the physical systems are aware of the QoS mismatches resulting from the service allocation, they may decide to adjust their QoS demands accordingly; we refer to this as the *reactive* physical layer. Similarly, the information set available at the network layer is updated as $\mathcal{I}_t = \mathcal{I}_{t-1} \cup_{i=1}^{N} \{I^{i,p}, \bar{I}_t^{i,p}, I_t^n\}$. This scenario is relevant for private networks with services located at the edge/local clouds, as sharing sensitive data with the network entails low privacy or security concerns. The optimal control and QoS demands for system i are computed online at every time t, from (5), as

$$\begin{aligned}(u^{i,*}_{[t,T_f]}, q^{i,*}_{[t,T_f]}) = \operatorname*{argmin}_{\mathcal{U}_k^i, \mathcal{Q}_k^i} & \sum\nolimits_{k=t}^{T_f} \mathsf{E}\left[c_k^i(x_k^i, u_k^i, w_k^i) + g_k^i(q_k^i)\right] \\ \text{subject to} \quad & x_{k+1}^i = f_k^i(x_k^i, u_k^i, w_k^i), \\ & u_k^i = \mathcal{U}_k^i(\mathcal{I}_k^i), \;\; q_k^i = \mathcal{Q}_k^i(\mathcal{I}_k^i) \\ & u_k^{i,*} \in \epsilon_u^i, \;\; q_k^{i,*} \in \epsilon_q^i, \\ & \forall k \in [t, T_f]. \end{aligned} \tag{12}$$

The regret-optimal resource allocation is computed as

$$\begin{aligned} & q^{r,*}_{[t,T_f]} \\ & = \operatorname*{argmin}_{\mathcal{Q}_k^r} \frac{1}{N} \sum\nolimits_{i=1}^{N} v^i \left(\mathsf{E}\left[J^{i,r}_{[t,T_f]}(u^{i,r}, q^{i,r}) - J^i_{[t,T_f]}(u^i, q^i)\right]\right) \\ & \text{subject to} \;\; q_k^{i,r} = \mathcal{Q}_k^r(\mathcal{I}_k), \\ & \qquad q_k^{i,r,*} \in \epsilon_q^i, \; \forall i \in \{1, \ldots, N\}, \; \forall k \in [t, T_f], \\ & \qquad \sum\nolimits_{i=1}^{N} q_k^{i,r,*} \le q^n_{\text{total}}, \; \forall k \in [t, T_f], \; t \in [T_s, T_f) \end{aligned} \tag{13}$$

where $q_k^{r,*} = [q_k^{1,r,*}, \ldots, q_k^{N,r,*}]^\top$, $J^i_{[t,T_f]}(u^i, q^i)$ is the optimal cost-to-go of system i computed based on its desired QoS $q_k^{i,*}$ obtained from (5), and $J^{i,r}_{[t,T_f]}(u^{i,r}, q^{i,r})$ is the cost based on the network-allocated QoS $q^{i,r}$. The constant q^n_{total} denotes the total service constraints at the network. Note that the network does not have the knowledge of u^i and $u^{i,r}$, but they are implicitly encoded in the cost values reported to the network.

The optimization procedure is summarized as follows.

Algorithm 1 `LocalReactiveOptimization`

for $t = T_s : T$ **do**
  $(u^{i,*}_{[t,T_f]}, q^{i,*}_{[t,T_f]}) \leftarrow$ by solving (12)
  Request $q^{i,*}_{[t,T_f]}$ from the network
  $q^{i,r,*}_{[t,T_f]} \leftarrow$ `NetworkwideQoSOptimization()`
  Use service $q_t^{i,r,*}$ to receive measurement and the sensor observes the next state x_{t+1}^i
  $t \leftarrow t + 1$
end for

Algorithm 2 `NetworkwideQoSOptimization()`

for $t = T_s : T$ **do**
  for i = 1 : N **do**
    Receive QoS request $q^{i,*}_{[t,T_f]}$
  end for
  Obtain networkwide QoS allocation $q^{r,*}_{[t,T_f]}$ by solving (13).
  send $q^{i,r,*}_{[t,T_f]}$ to the i-th system for all i and allocate the services accordingly.
  $t \leftarrow t + 1$
end for

At the network layer resource availability and service tolerance check is performed and services are pre-assigned to each system (over the entire horizon $[T_s, T_f]$) such that the constraints of the problem (13) are satisfied. At the physical layer, each system i runs the optimization (12) at time T_s, and computes $u^{i,*}_{[T_s,T_f]}, q^{i,*}_{[T_s,T_f]}$ over $[T_s, T_f]$, given the information set $\mathcal{I}^i_{T_s}$. Based on the locally computed control and QoS decisions, and the service pre-allocation at the network layer, $J^i_{[T_s,T_f]}$ and $J^{i,r}_{[T_s,T_f]}$ are computed. QoS demands from all systems as well as the local regrets are now reported to the network for service provisioning. If $q^{i,*}_{[T_s,T_f]} = q^{i,r}_{[T_s,T_f]}$, $\forall i$, then local regrets are zero and there is no need for re-allocation; otherwise, the network runs again the optimization (12), now with non-zero social regret, to find an allocation pattern that minimizes the social regret. Systems served with different QoS compared to their demands, now optimize their performance based on the mismatch between $q^i_{[T_s,T_f]}$ and $q^{i,r}_{[T_s,T_f]}$, from which only $q^{i,r}_{T_s}$ is implemented. Therefore, the optimization (12) is run again at system i over $[T_s+1, T_f]$ to re-compute the desired QoS and control inputs given the allocation pattern by the network. The re-computed QoS demands and the local regrets are reported to the network for service allocation and this loop repeats until time $T_f - 1$. Note that after service allocation and depending on QoS capacity, certain systems may receive enhanced QoS compared to what they requested.

Remark 1: At any time t, an optimal predicted version of $q^{i,*}_{[t,T_f]}$, for all i, is available at the network to solve (13). Although the desired future QoSs may change due to mismatches between $q_t^{i,*}$ and $q_t^{i,r,*}$, any changes in $q_\tau^{i,*}$, $\tau > t$, is reported to the network via the information set $\bar{I}_\tau^{i,p}$.

2) Offline Awareness — Passive Physical Layer: We now consider the information structure under which the physical systems do not incorporate the service allocation outcome when deciding on their control and QoS demands. This means that $\mathcal{I}_t^i = \mathcal{I}_{t-1}^i \cup \{z_t^i, u_{t-1}^i, \bar{I}_t^{i,p}, g^i_{[t,T_f]}\}$. We refer to this as a *passive* physical layer, since QoS demands are not updated in response to QoS demand and provision mismatches. Available information at the network layer is similar to the reactive physical layer case. The optimal control and QoS demands are then obtained from the offline optimization problem that resembles the problems presented in (12)-(13) where the only difference is the current time t is replaced with the initial time T_s. The structural difference between the reactive optimization problems (12)–(13) and their passive counterparts is that the former ones are online while the latter are offline and thus performed only once during the time horizon $[T_s, T_f]$. This emanates from the reactive and passive natures of the problems, which are a direct consequence of the awareness models $\mathcal{I}_k^i$ and $\mathcal{I}_k$. For reactive case, the network allocation in response to the individual system's QoS demands are considered at every time-step, leading to changes in future QoS demands; while for the passive case the QoS demands of the physical systems, computed at the initial time T_s, remain unchanged over the entire horizon. The higher awareness in the reactive scenario leads to a more computationally complex solution, but also improved QoC-QoS trade-off due to the online adaptation of QoS demands and allocation.

3) Absence of Parametric & Regret Information at Network Layer: In this scenario, the parametric data of the individual systems ($I^{i,p}$) and the cost values J^i and $J^{i,r}$ are not shared with the network, leading to $\mathcal{I}_t = \mathcal{I}_{t-1} \cup_{i=1}^{N} \{q^i_{[t,T_f]}, I_t^n\}$. This scenario is relevant when resources and services are located on public servers/clouds, and excessive sharing of information is not desired due to privacy concerns. Under this information model, the optimal pair $\{u_k^{i,*}, q_k^{i,*}\}$ is obtained from (12) or its offline/passive counterpart (depending on the physical layer being reactive or passive). Service allocation, however, cannot be performed using the regret function (8) since systems' regret values are not known for the network. The network layer in this case minimizes the average QoS regret among all systems. For the passive physical layer, this leads to

$$q^{r,*}_{[T_s,T_f]} = \underset{\mathcal{Q}^r_k}{\operatorname{argmin}} \frac{1}{N}\sum_{i=1}^{N} \mathsf{E}\Big[|g^i_k(q^{i,r}_k) - g^i_k(q^i_k)|\Big]$$
$$\text{subject to} \quad q^{i,r}_k = \mathcal{Q}^r_k(\mathcal{I}_k),$$
$$\sum_{i=1}^{N} q^{i,r,*}_k \le q^n_{\text{total}}, \quad \forall k \in [T_s, T_f].$$

For the reactive physical layer, the optimal service allocation problem is obtained by replacing T_s with the current time t, in the passive service allocation above.

The drawback of such co-design is that the QoC constraints or regret tolerances of individual systems are not incorporated in the resource allocation. This may lead to control performance degradation in the passive case (where QoS demand and provision are decided once over the entire time horizon), and more frequent update of QoS demands, and consequently more frequent need for resource allocation, in the reactive case.

B. Infinite Horizon Extensions

Next we discuss some cases when the finite-horizon co-design problems in previous section can be formulated over infinite time. For the infinite-horizon problems, we make two crucial assumptions: control systems are closed-loop controllable, and second all optimal policies across both layers are asymptotically stationary. That is, we assume system in (1) is asymptotically stable for some u^i_k, and second, when time $k \to \infty$, the policies $\mathcal{U}^i$, $\mathcal{Q}^i$ become time-invariant, although the decision variables may still be dynamic. In fact, under some certain structures, if the information sets $\mathcal{I}^i_k$ and $\mathcal{I}_k$ are updated asymptotically, the decision variables u^i_k, q^i_k and $q^{i,r}_k$ will be dynamic, but the policies that generate those decisions are assumed stationary. Having this assumption, the objective of system i is to find stationary policies $\mathcal{U}^i$ and $\mathcal{Q}^i$ that solve

$$\min_{\mathcal{U}^i,\mathcal{Q}^i} \quad J^i(\mathcal{U}^i, \mathcal{Q}^i), \tag{14}$$

for the average cost function

$$J^i(\mathcal{U}^i, \mathcal{Q}^i) = \limsup_{T\to\infty} \frac{1}{T}\, \mathsf{E}\Big[\sum_{k=0}^{T-1} c^i(x^i_k, u^i_k, w^i_k) + g^i(q^i_k)\Big].$$

Let $\mathcal{Q}^{i,*}$ be the optimal stationary policy for the QoS demands obtained from problem (14). The QoS demands are then computed according to $\mathcal{Q}^{i,*}$ that may not be perfectly served by the network due to service constraints. Let $\{q^{i,r}_0, q^{i,r}_1, \cdots\}$ denote the QoS served by the network according to the stationary network service allocation policy $\mathcal{Q}^{i,r}$, that is system i is served with the QoS $q^{i,r}_k$ at time k. The realized asymptotic average performance of system i then becomes

$$J^{i,r} = \limsup_{T\to\infty} \frac{1}{T}\, \mathsf{E}\Big[\sum_{k=0}^{T-1} c^i(x^i_k, u^{i,r}_k, w^i_k) + g^i(q^{i,r}_k)\Big]. \tag{15}$$

The local and social asymptotic regret measures then become

$$R^i = J^{i,r} - J^{i,*}, \qquad R^s = \frac{1}{N}\sum_{i=1}^{N} v^i(R^i). \tag{16}$$

The stationarity assumption leads to that the jointly generated QoS demand and provision decisions obey time-invariant policies. Stability guarantees should explicitly be considered in the infinite horizon co-design, either by restricting the optimal policies to a class of stabilizing policies, or as system-level constraints in the optimization problems.

In the finite-horizon, we solve the co-design problems in an adaptive and predictive way; this approach is not applicable over an infinite-horizon. To simplify the calculation, for a reactive physical layer, we can assume that a system demands a unique QoS according to a stationary policy, and the demand may be updated to another QoS realization if the network does not serve the first demand. Note that both demands are computed from the stationary policy. For a passive physical layer, on the other hand, the unique QoS remains unchanged at stationarity. The invariant optimal control and QoS demands of system i can then be obtained from (14). At the network layer, the objective is to decide on the service allocation, governed by the policy $\mathcal{Q}^{i,r}$, for all i, to minimize the social regret

$$\min_{\cup_{i=1}^N \mathcal{Q}^{i,r}} \quad R^s,$$
$$\text{subject to} \quad \sum_{i=1}^{N} q^{i,r,*} \le q^n_{\text{total}},$$
$$\text{Stability constraints.} \tag{17}$$

For NCS where stability is coupled with the network QoS (that directly affects the QoC), system-level limits are required on the physical layer and the network layer parameters (delay, data rate, reliability). We will discuss the infinite horizon case together with stability requirements quantitatively in Part II.

IV. Conclusion

In Part I of this two part letter, we presented a general joint performance optimization problem for NCS in the form of optimal trade-offs between QoC and QoS based on a regret function. We also characterized the optimal QoC-QoS co-design based on cross-layer awareness. In Part II, we provide explicit QoC-QoS regret-optimal policies for Gauss-Markov systems over a network with multiple QoS-tailored slices.

References

[1] X.-M. Zhang et al., "Networked control systems: A survey of trends and techniques," *IEEE/CAA J. Autom. Sinica*, vol. 7, no. 1, pp. 1–17, Jan. 2020.
[2] V. Jirkovský et al., "Understanding data heterogeneity in the context of cyber-physical systems integration," *IEEE Trans. Ind. Informat.*, vol. 13, no. 2, pp. 660–667, Apr. 2017.
[3] Y. Zhan et al., "Future directions of networked control systems: A combination of cloud control and fog control approach," *Comput. Netw.*, vol. 161, pp. 235–248, Oct. 2019.
[4] F. Kurtz et al., "Network slicing for critical communications in shared 5G infrastructures—An empirical evaluation," in *Proc. 4th IEEE Conf. Netw. Softwarization Workshops (NetSoft)*, Jun. 2018, pp. 393–399.
[5] M. H. Mamduhi et al., "A cross-layer optimal co-design of control and networking in time-sensitive cyber-physical systems," *IEEE Control Syst. Lett.*, vol. 5, no. 3, pp. 917–922, Jul. 2021.
[6] N. Sahlab et al., "An approach for context-aware cyber-physical automation systems," *IFAC-PapersOnLine*, vol. 54, no. 4, pp. 171–176, 2021.
[7] E. C. Balta et al., "Controller-aware dynamic network management for industry 4.0," in *Proc. 48th Annu. Conf. IEEE Ind. Electron. Soc.*, Oct. 2022, pp. 1–6.
[8] Y.-Q. Song, "Networked control systems: From independent designs of the network QoS and the control to the co-design," *IFAC Proc. Volumes*, vol. 42, no. 3, pp. 155–162, 2009.
[9] E. Molina and E. Jacob, "Software-defined networking in cyber-physical systems: A survey," *Comput. Electr. Eng.*, vol. 66, pp. 407–419, Feb. 2018.
[10] M. H. Mamduhi et al., "Delay-sensitive joint optimal control and resource management in multiloop networked control systems," *IEEE Trans. Control Netw. Syst.*, vol. 8, no. 3, pp. 1093–1106, Sep. 2021.
[11] Y.-P. Hsu et al., "Age of information: Design and analysis of optimal scheduling algorithms," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Jun. 2017, pp. 561–565.
[12] T. Soleymani et al., "Value of information in feedback control: Quantification," *IEEE Trans. Autom. Control*, vol. 67, no. 7, pp. 3730–3737, Jul. 2022.
[13] T. Runolfsson, "Stationary risk-sensitive LQG control and its relation to LQG and H-infinity control," in *Proc. 29th IEEE Conf. Decis. Control*, Dec. 1990, pp. 1018–1023.
[14] A. Karapetyan et al., "On the regret of H_∞ control," in *Proc. 61st IEEE Conf. Decis. Control*, Mar. 2022, pp. 6181–6186.