Optimal Scheduling of Downlink Communication for a Multi-Agent System with a Central Observation Post

Mario Zanon, Themistoklis Charalambous, Henk Wymeersch and Paolo Falcone

Abstract— In this paper, we consider a set of agents, which may receive an observation of their state by a central observation post via a shared wireless network. The aim of this work is to design a scheduling mechanism for the central observation post to decide how to allocate the available communication resources. The problem is tackled in two phases: (i) first, the local controllers are designed so as to stabilise the subsystems for the case of perfect communication; (ii) second, the communication schedule is decided with the aim of maximising the stability of the subsystems. To this end, we formulate an optimisation problem which explicitly minimises the Lyapunov function increase due to communication limitations. We show how the proposed optimisation can be expressed in terms of Value of Information (VoI), we prove Lyapunov stability in probability and we test our approach in simulations.

I. INTRODUCTION

The advancement of smart devices with increasing *sensing*, *computing* and *control* capabilities makes it possible for several processes to become more intelligent, energy-efficient, safe and secure. However, the components of such systems are spatially distributed and communication between smart devices (being sensors, actuators or controllers) is mainly supported by a *shared, wireless* communication network. These systems are known as Wireless Networked Control Systems (WNCSs); a thorough literature review can be found in [1], [2]. The use of wireless communications in such systems, however, introduces additional challenges due to the limitations imposed by the use of the wireless medium, e.g. packet losses, data rate constraints and interference. Therefore, it is necessary to develop Medium Access Control (MAC) mechanisms for WNCSs.

There are two types of MAC: (a) random access mechanisms, in which each agent can access the network randomly, and (b) scheduling mechanisms, in which a centralised entity decides on the allocation of the resources. Even though random access protocols can easily be implemented in a distributed fashion, it is difficult to provide any performance guarantee.

In this work, we consider a scenario in which several agents may receive an observation of their state by a central observation post, that has the ability to observe the states of all the agents in the system, via a shared wireless network.

The shared wireless network is interference-limited, thus allowing only a limited number of simultaneous transmissions. Additionally, these transmissions may not reach the agents since the communication channel can be unreliable. Such setups may arise in several realistic occasions, e.g. when a rescue operation takes place, a drone observes the scene and informs the agents by broadcasting their location with respect to the target mission, one agent at the time. In such cases on the one hand the local control schemes are required to stabilise the subsystems (agents); on the other hand, a scheduling mechanism must decide which information is sent to which subsystem.

For discrete-time linear systems, several scheduling approaches have been proposed in order to make a good use of the communication channel, see e.g. [3]–[5] and references therein. These works consider the scheduling problem for state estimation of LTI systems with Gaussian noise. With a few exceptions (e.g., [6], [7]), a finite time horizon is considered, in which the problem is a combinatorial optimisation one [8], and hence $\mathcal{NP}$-hard, making the computation of the globally optimal solution over long time horizons computationally expensive. The works considering the infinite horizon case, focus on estimation only.

In this paper, we decouple the control and communication allocation problems: (i) first, the local controllers are designed, so as to stabilise the subsystems in the perfect communication case; (ii) second, the communication schedule is decided with the aim of maximising the stability of the subsystems. To this end, we formulate an optimisation problem which explicitly minimises the Lyapunov function increases due to communication limitations. We show how the proposed optimisation can be expressed in terms of VoI which, in the considered case, resembles the concept of cost-to-go used in dynamic programming, and we prove Lyapunov stability in probability. Finally, we rely on a numerical approach for mixed-integer optimal control for cheaply solving the problem to local optimality.

The proposed approach is rather flexible, as it easily adapts to different formulations and can explicitly account for lossy communication channels, so as to (a) exploit knowledge on the probability of having a successful communication and (b) adapt the schedule whenever a message fails to be delivered. Moreover, since the choice of Lyapunov functions is not unique, there is some freedom in the definition of the cost which allows one to e.g. assign higher priority to some subsystem.

The rest of the paper is organised as follows. In Section II we give some preliminary results and the problem

formulation. In Section III we present our main theoretical contribution, i.e. the formulation of the communication allocation problem. In Section IV we evaluate our theoretical developments in a series of different scenarios. Finally, in Section V we summarise our results and discuss possible future directions.

II. PRELIMINARIES AND PROBLEM FORMULATION

We consider a network consisting of M agents with unreliable communication links from the sensors to the observer and controller, and perfect communication links from the controller to the agents, as depicted in Figure 1.

A. Definitions

Throughout the paper, we denote agents by index i and time instants by index k, such that each agent has state $x_{i,k}$ and control $u_{i,k}$. Whenever it is not necessary, we drop index i for simplicity of notation.

We consider agents described as perturbed linear systems

$$x_{k+1} = Ax_k + Bu_k + w_k, \tag{1}$$

where x, u and w denote the state, control and perturbation, respectively. We assume that x_0 and w are Gaussian random variables with zero mean and covariance $X_0 = \mathbb{E}\{x_0 x_0^\top\}$ and $W = \mathbb{E}\{ww^\top\}$. Moreover, we assume that x_0 and w are independent, such that $\mathbb{E}\{x_0 w^\top\} = 0$.

The observations of the states are given by

$$y_k = Cx_k + v_k \tag{2}$$

where v denotes the measurement noise, which we assume to be a Gaussian random variable with zero mean and covariance $\mathbb{E}\{vv^\top\} = V$. Moreover, we assume that w, x_0, v are mutually independent, i.e. $\mathbb{E}\{wv^\top\} = \mathbb{E}\{x_0 v^\top\} = 0$.

B. Kalman Filter Updates

Let $\delta_k = 1$ if the observation is sent to the agent, $\delta_k = 0$ otherwise. The *a posteriori* state estimate given by the Kalman filter [9] is

$$\hat{x}_{k+1} = A\hat{x}_k + Bu_k + \delta_{k+1}L_{k+1}(y_{k+1} - C(A\hat{x}_k + Bu_k)), \tag{3}$$

with

$$\bar{E}_{k+1} = AE_k A^\top + W, \tag{4a}$$

$$L_{k+1} = \bar{E}_{k+1} C^\top \left(V + C\bar{E}_{k+1}C^\top\right)^{-1}, \tag{4b}$$

$$E_{k+1} = (I - \delta_{k+1}L_{k+1}C)\,\bar{E}_{k+1}, \tag{4c}$$

where the estimation error is defined as $e_k = x_k - \hat{x}_k$ and its covariance as $E = \mathbb{E}\{ee^\top\}$.

C. Imperfect Communication

In the previous subsection we implicitly assumed a perfect downlink, i.e., all messages sent by the central node to the agents arrive at destination. While this assumption might be reasonable for physically connected systems, it is not when the communication is wireless. In this case, one needs to further account for the probability that the communication does not succeed.
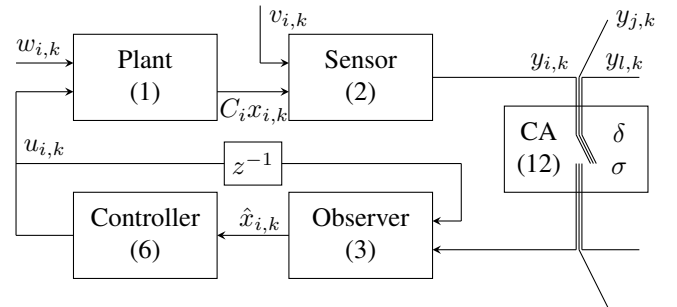

Fig. 1.   Scheme of the system: all agents share the same communication channel, which has limited resources. The Communication Allocation (CA) algorithm assigns the channel to a selected subset of the agents.

We assume in the following that the probability σ of having a successful communication from the sensor to the observer (see Figure 1) is not correlated with the estimation error e. This is in general not true, since the distance estimate is in general dependent on e. Our assumption then implicitly relies on the error e being small enough. In the case of wideband communication, this requires the position estimate error to be no larger than a few meters. Note that σ is typically also a function of the power of the transmission. In the following, we will assume that the distance and power are fixed and known a priori, such that σ_k is a known time-varying signal.

By defining s as a Bernoulli random variable with $\mathbb{E}\{s\} = \sigma$, Equation (4c) becomes

$$E_{k+1} = (I - \delta_{k+1}\sigma_k L_{k+1}C)\bar{E}_{k+1}. \tag{5}$$

Note that E defined by (5), now denotes the expected covariance. We summarise the time evolution of E as $E_{k+1} = f_k^{\mathrm{E}}(E_k, \delta_{k+1})$. This function is in general subsystem-dependent and it inherits the time-varying nature from σ_k.

D. Feedback Control

We consider the local linear feedback

$$u_k = -K\hat{x}_k. \tag{6}$$

By defining $A_K \triangleq A - BK$, the state dynamics (1) become $x_{k+1} = A_K x_k + BKe_k + w_k$.

Note that, if we consider a quadratic cost (as we will do in Section III-A), then the so-called *certainty equivalence principle* holds. It states that, the optimal solution is the same as for the corresponding deterministic problem as long as the disturbances present in the stochastic control system are zero mean [10].

III. STABILITY WITH LIMITED COMMUNICATION

In this section, we analyse the stability of the closed-loop system in the case of limited communication resources.

Note that the stability of the system in terms of expected value is obtained regardless of communication, by choosing K appropriately. The time evolution of the expected values of the state and error are given respectively by

$$\mathbb{E}\{e_k\} = 0,$$
$$\mathbb{E}\{x_{k+1}\} = A_K \mathbb{E}\{x_k\} = A_K \hat{x}_k,$$

where we assumed that $\hat{x}_0 = \mathbb{E}\{x_0\}$. Communication, however, is crucial in order to ensure that the error covariance E does not grow indefinitely large.

We define the Lyapunov function of each subsystem as $\mathcal{V}(x) = x^\top P x$. We remark here that matrix P can be easily computed by using the Lyapunov equation

$$A_K^\top P A_K - P + Q = 0, \tag{7}$$

with $Q \succ 0$. The freedom in choosing matrix Q will later be exploited in order to tune the communication allocation as desired.

In order to select how to schedule the communication between the central node and the agents, we aim at defining the cost of communicating or not with each subsystem. In this sense, our definition resembles that VoI, already used in e.g. [3], [5]. Because we find it relevant to relate the VoI to the stability of each subsystem, we first establish a framework aimed at building a cost directly related to the Lyapunov function of each subsystem. Second, we formulate the communication allocation problem in order to minimise the chosen cost function. Third, we establish a connection between our formulation and the definition of VoI. Finally, we prove Lyapunov stability in probability.

A. Stabilising Feedback and its Cost

In this subsection, we aim at defining the cost of not communicating the state measurement or estimate. We express this cost based on the Lyapunov function of each subsystem in order to explicitly account for the suboptimality resulting from control actions based on imperfect estimation. Because we consider each subsystem separately, we drop the index i for simplicity.

In order to simplify the description, we first restrict our attention to the LQR case. In a second step, we will detail how, for any stabilising feedback matrix K, one can define LQR cost matrices that yield K as optimal feedback matrix.

Consider an LQR with stage cost

$$\ell(x, u) = \begin{bmatrix} x \\ u \end{bmatrix}^\top \begin{bmatrix} Q & S^\top \\ S & R \end{bmatrix} \begin{bmatrix} x \\ u \end{bmatrix}. \tag{8}$$

The discrete-time algebraic Riccati equation (DARE) is

$$0 = Q - P + A^\top P A - (S^\top + A^\top P B)K, \tag{9a}$$
$$K = (R + B^\top P B)^{-1}(S + B^\top P A). \tag{9b}$$

We now compare the cost of applying the feedback $u = -K\hat{x}$ to the cost of applying the nominal feedback $\bar{u} = -Kx$. We first note that

$$u = -Kx + Ke = \bar{u} + Ke, \tag{10a}$$
$$x^{\bar{u}}_{k+1} = A_K x_k + w_k, \tag{10b}$$
$$x^u_{k+1} = A_K x_k + BKe_k + w_k = x^{\bar{u}}_{k+1} + BKe_k. \tag{10c}$$

The effect of the estimation error over one prediction step is given by the following difference in closed-loop cost

$$\Delta\mathcal{V}(x_k) = \ell(x_k, -K\hat{x}_k) - \ell(x_k, -Kx_k) \\ + \left(x^u_{k+1}\right)^\top P x^u_{k+1} - \left(x^{\bar{u}}_{k+1}\right)^\top P x^{\bar{u}}_{k+1}. \tag{11}$$

By taking the expected value of the cost increase and noting that $\mathbb{E}\{e\} = 0$, one obtains

$$\mathbb{E}\{\Delta\mathcal{V}(x_k)\} = \mathbb{E}\{e_k^\top K^\top R K e_k + e_k^\top K^\top B^\top P B K e_k\} \\ = \mathrm{tr}\left(K^\top \left(R + B^\top P B\right) K E_k\right).$$

Because the loop is closed using the state estimate, the expected value of the increase in the Lyapunov function due to estimation error over a given time interval $[0, N]$ is then

$$\sum_{k=0}^N \mathbb{E}\{\Delta\mathcal{V}(x_k)\} = \sum_{k=0}^N \mathrm{tr}\left(\Gamma E_k\right),$$

with $\Gamma \triangleq K^\top \left(R + B^\top P B\right) K$. Note that dynamic programming yields $\left(R + B^\top P B\right) \succ 0$ and, therefore, $\Gamma \succ 0$. Moreover, (9) yields $\Gamma = Q - P + A^\top P A$.

We now turn our attention to generalising the proposed framework to any stabilising feedback gain K.

Propositon 1: Given any linear system $x_{k+1} = Ax_k + Bu_k$ and stabilising feedback gain K, one can obtain the feedback gain K as the solution of an LQR formulated using the stage cost $\ell(x, u)$ from (8) using matrices $R = I$, $S = K$, $Q = K^\top K$.

Proof: The proposed matrices solve the DARE (9) with $P = 0$ and $R + B^\top P B \succ 0$. The proof is then obtained by noting that the DARE admits at most one stabilising solution, i.e. K, which is stabilising by assumption. ∎

Note that the LQR obtained by applying Proposition 1 has a positive semi-definite stage cost and yields $\Gamma = K^\top K$. If a positive-definite formulation is sought, the technique proposed in [11] can be applied. Alternatively, (7) can be solved and $\Gamma = Q - P + A^\top P A$ can be used.

B. Optimal Communication Allocation

After analysing each subsystem separately, we consider the aggregate cost of communication, resulting from all subsystems considered as one system. While the design of each agent's controller can be done separately, it is important to note that some degrees of freedom in the cost definition are free and should be exploited in order to define the communication cost. The simplest observation is that by multiplying the LQR tuning matrices by a factor $a_i \neq 0$, the same feedback gain is obtained, but the Lyapunov function matrix becomes $a_i P_i$. This allows one to assign higher or lower importance to some subsystems, thus indicating one possibility to further tune the cost for the aggregate system. Another approach can be to design the cost of the single systems in an aggregate way, i.e. by considering them as a single (sparse) system.

We remark here that one tempting approach could consist in applying [11, Lemma 2], so as to assign an arbitrarily chosen value $P_i \succ 0$ the cost-to-go matrix of each subsystem's controller. While this is always possible, the relative stage cost is typically indefinite. Consequently, the cost-to-go is not guaranteed to be a Lyapunov function: though positive definite, it typically fails to fulfil the decrease condition.

We formulate the communication allocation problem by summarising the previously described design choices in matrices $\Gamma_i \triangleq K_i^\top \left(R_i + B_i^\top P_i B_i \right) K_i$. The optimal allocation problem at time j can then be formulated as

$$\min_{\delta} \quad \sum_{i=1}^{M} \sum_{k=j}^{N+j} \mathrm{tr}(\Gamma_i E_{i,k}) \tag{12a}$$

$$\mathrm{s.t.} \quad E_{i,k} = f_{i,k}^{\mathrm{E}}(E_{i,k-1}, \delta_{i,k}), \quad i \in \mathbb{I}_1^M, \ k \in \mathbb{I}_j^{N+j}, \tag{12b}$$

$$\delta_{i,k} = \{0,1\}, \qquad\qquad i \in \mathbb{I}_1^M, \ k \in \mathbb{I}_j^{N+j}, \tag{12c}$$

$$\sum_{i=1}^{M} \delta_{i,k} \leq \gamma, \qquad\qquad k \in \mathbb{I}_j^{N+j}, \tag{12d}$$

where we define $\mathbb{I}_a^b \triangleq \{ n \in \mathbb{N} \mid a \leq n \leq b \}$ and $\delta = (\delta_1, \ldots, \delta_M)$ and $\delta_i = (\delta_{i,0}, \ldots, \delta_{i,N})$.

C. The Value of Information

In this subsection, we show the connection between our approach and the concept of *value of information* (VoI), used in different settings in e.g. [3], [5]. In the context of our setting, the VoI could be defined for agent j as

$$\nu_i(\delta_i, E_{i,j-1}) \triangleq \sum_{k=j}^{N+j} \mathrm{tr}(\Gamma_i E_{i,k}) \tag{13a}$$

$$\mathrm{s.t.} \quad E_{i,k} = f_{i,k}^{\mathrm{E}}(E_{i,k-1}, \delta_{i,k}), \quad k \in \mathbb{I}_j^{N+j}. \tag{13b}$$

Then, Problem (12) can be formulated in terms of VoI as

$$\min_{\delta} \quad \sum_{i=1}^{M} \nu_i(\delta_i, E_{i,j-1}) \tag{14a}$$

$$\mathrm{s.t.} \quad \delta_{i,k} = \{0,1\}, \qquad k \in \mathbb{I}_j^{N+j}, \tag{14b}$$

$$\sum_{i=1}^{M} \delta_{i,k} \leq \gamma, \qquad k \in \mathbb{I}_j^{N+j}. \tag{14c}$$

Problem (14) has the same complexity as (12). However, the problem of computing the VoI (13) is very difficult and presents strong similarities to the computation of the cost-to-go function in dynamic programming.

D. Lyapunov Stability in Probability

In this subsection, we analyse the conditions under which our approach yields closed-loop stability. The stability concept we will use is Lyapunov stability in probabilty [12]. For the sake of notational simplicity, we consider the system as a whole, such that $x_k = (x_{1,k}, \ldots, x_{M,k})$, and matrices E and P are defined consistently.

Definition 2 (Lyapunov Stability in Probability):
Lyapunov Stability in Probability (LSP) holds for a linear system if, given $P \succ 0$, $\epsilon, \bar{\epsilon}$, there exists $\rho(\epsilon, \bar{\epsilon}) > 0$ such that $|x_0| < \rho$ implies

$$\limsup_{k \to \infty} \mathrm{P}[x_k^\top P x_k \geq \epsilon] \leq \bar{\epsilon}.$$

Assumption 3: There exists a baseline schedule δ^{bs} over the time interval $[0, N]$ yielding error covariance E_k^{bs} such that $\forall \ E_0 \in \mathcal{E} \triangleq \{E \mid \mathrm{tr}(\Gamma E) \leq \mu \in \mathbb{R}\}$ it holds that $E_k^{\mathrm{bs}} \in \mathcal{E}$, $k = 1, \ldots, N$.

Theorem 4: Assume that the initial state covariance, the noise covariances W, V and the initial estimation error are finite. Suppose moreover that Assumption 3 holds. Then, the closed-loop system with communication allocation based on the solution of Problem (12) is LSP.

Proof: Since $E_N \in \mathcal{E}$, by applying the baseline schedule, $E_{N+1} \in \mathcal{E}$ and the problem is recursively feasible.

We apply Markov's inequality to obtain the upper bound

$$\mathrm{P}[x_k^\top P x_k \geq \epsilon] \leq \frac{\mathbb{E}\{x_k^\top P x_k\}}{\epsilon}. \tag{15}$$

We define $X_k \triangleq \mathbb{E}\{x_k x_k^\top\}$ and write the time evolution of $\mathbb{E}\{x_k^\top P x_k\} = \mathrm{tr}(P X_k)$ as

$$\mathbb{E}\{x_{k+1}^\top P x_{k+1}\} = \mathrm{tr}(P X_{k+1})$$
$$= \mathrm{tr}(A_K^\top P A_K X_k) + \mathrm{tr}(P W) + \mathrm{tr}(\Gamma E_k).$$

Since A_K is stable, using (7) we obtain $\mathrm{tr}(A_K^\top P A_K X_k) = \mathrm{tr}((P - Q) X_k) \leq \alpha \, \mathrm{tr}(P X_k)$, with $0 \leq \alpha < 1$.

By optimality $\sum_{j=k}^{k+N} \mathrm{tr}(\Gamma E_j) \leq \sum_{j=k}^{k+N} \mathrm{tr}(\Gamma E_j^{\mathrm{bs}})$, and we conclude that

$$\mathrm{tr}(\Gamma E_k) \leq N \sum_{j=k}^{k+N} \mathrm{tr}(\Gamma E_j^{\mathrm{bs}}) \leq N \mu.$$

This further yields the upper bound $\mathrm{tr}(P W) + \mathrm{tr}(\Gamma E_k) \leq \nu$, where $\nu \triangleq \mathrm{tr}(P W) + N \mu < \infty$, such that

$$\mathrm{tr}(P X_{k+1}) \leq \alpha \, \mathrm{tr}(P X_k) + \nu.$$

Consequently, $\limsup_{k \to \infty} \mathrm{tr}(P X_k) \leq \dfrac{\nu}{1 - \alpha} \leq \nu < \infty$. ∎

IV. NUMERICAL RESULTS

In this section, we test the theoretical developments of the previous sections in a series of different scenarios. We consider subsystems with

$$A_i = \begin{bmatrix} 1 & 0.1 \\ 0 & 1 \end{bmatrix}, \qquad B_i = \begin{bmatrix} 0.005 \\ 0.1 \end{bmatrix}. \tag{16}$$

Unless specified differently, we control each subsystem by an LQR controller designed using weighting matrices $Q_i = I$ and $R_i = 0.01$.

A. Numerical Methods

Problem (12) is a mixed-integer optimal control problem (MIOCP) formulated in discrete time. An efficient approach, called *partial outer convexification*, has been proposed in [13], [14] to handle the numerical solution of MIOCPs. Consider a generic system with state x, continuous control u, binary control ω and system dynamics $x_+ = f(x, u, \omega)$. Partial outer convexification reformulates the dynamics as $x_+ = \sum_{j=1}^{n_c} f(x, u, \omega_j) \tilde{\omega}_j$, where, vectors ω_j are binary vectors spanning the space of all possible binary combinations of inputs, n_c is the amount of such combinations and $\tilde{\omega}_j$ is a newly introduced control variable. In order to solve the MIOCP, the integer constraint is relaxed to $\tilde{\omega} \in [0, 1]^{n_c}$ and the relaxed OCP is solved using standard techniques. This relaxation can be proven to be tighter than the relaxation of the original formulation $x_+ = f(x, u, \omega)$ using $\omega \in [0, 1]^{n_\omega}$.

Often the solution of the relaxed problem is integer. In all other cases, the continuous solution can be approximated by a switching integer solution by using a rounding scheme. Note that, because the relaxed problem is nonconvex, the solver returns a local minimum. We refer to [13], [14] and references therein for more details on the topic.

We remark that the problem formulations used in this paper are already in their partial outer convexification form. We solve the optimisation problems using Casadi [15] and Ipopt [16].

B. Identical Subsystems

We simulate first the simplest case of identical subsystems with identical noise covariance. In this simple case, the optimal solution is to communicate with subsystem i only after communication has happened with all subsystems $j \neq i$.

We considered 4 identical agents with $W_i = 10^{-2}I$, $V_i = 10^{-3}I$ and a prediction horizon $N \in [1, 10]$ and $\gamma = 1$. Due to the symmetry of the problem, the solver returned non-integer solutions in the beginning of the simulation and a rounding strategy was necessary in the beginning. Afterwards, integer solutions were always obtained and the periodic solution $\delta_{i,k} = 1$ when $k \mod i = 0$, $\delta_{i,k} = 0$ otherwise, was obtained (modulo phase shifts).

C. Non-Identical Subsystem with Perfect Communication

We consider now the case of non-identical subsystems. In particular, we consider different numbers of agents M.

The closed-loop cost for $\gamma = 1$ is displayed in Figure 2 for different prediction horizons N and amounts of subsystems M. For $M < 5$, horizons $N > 1$ did not seem to improve the closed-loop cost. For $M \geq 5$, it can be seen that, while the optimiser often falls into local minima, the cost tends to decrease by using longer prediction horizons until $N = 5$. This effect is more evident when the amount M of subsystems becomes larger. In most cases, the closed-loop cost is less than 5% higher than the best value found by the optimiser. For $M \leq 9$ the solver always returned an integer solution. In all other simulations the solution needed to be rounded at most twice. Simulations with $\gamma > 1$ showed similar trends. Finally, the closed-loop cost dependence on γ is shown on the bottom graph in Figure 2.

D. Tuning the Cost of Communication

In this subsection, we briefly present how the cost of communicating to one subsystem rather than to the others can be tuned. For the sake of simplicity, we consider two identical subsystems.

We use the following tuning for the LQR: $Q_1 = I$ and $R_1 = 0.01$ and $Q_2 = a^2 Q_1$ and $R_2 = a^2 R_1$, where a is our tuning parameter. As remarked before, this choice yields the same feedback gain for both systems, i.e. $K_1 = K_2$. However, the cost-to-go matrices are $P_2 = a^2 P_1$, such that $\Gamma_2 = a^2 \Gamma_1$. As one could expect, the ratio $r = \sum_k \delta_{2,k}/\delta_{1,k}$ between the amount of communication with system 1 and 2 scales roughly linearly with increasing a. However, since the problem has an integer nature and the solvers only finds a local minimum the ratio evolves as displayed in Figure 3.
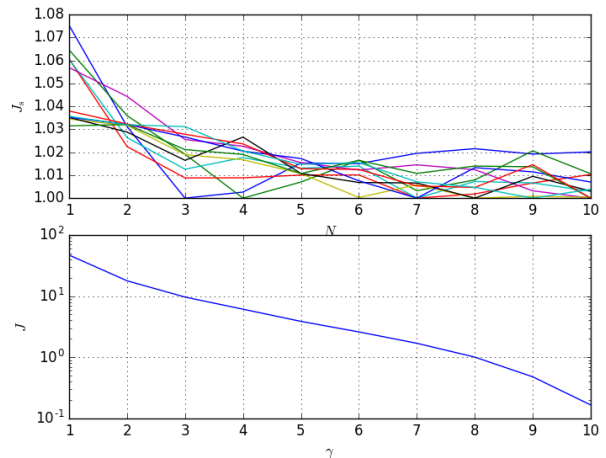

Fig. 2. Top graph: closed-loop cost for $M \in [5, 15]$, $\gamma = 1$ as a function of the prediction horizon for several scenarios considering different numbers of agents. The value J_s denotes the closed-loop cost scaled around the minimum for each scenario. Bottom graph: Closed-loop cost J for $M = 10$ as a function of the amount of available communication γ in (12d).
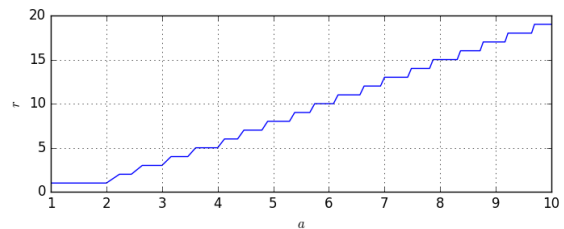

Fig. 3. Ratio $r = \sum_k \delta_{2,k}/\delta_{1,k}$ as a function of parameter a.

E. Imperfect Communication

We now consider a scenario with a lossy communication channel with probability of successful communication given by $p_{i,k} = e^{-d_{i,k}^2}$, where $d_{i,k}$ denotes the distance of agent i from the central node at time k. We use identical systems, i.e. with variances given by

$$V_1 = 10^{-3}I, \quad W_1 = 10^{-2}I, \quad V_2 = V_1, \quad W_2 = W_1.$$

We consider a scenario where $d_{i,k} = \cos\left(0.1k + i\frac{\pi}{2}\right)$, such that when one system is at the farthest point, the other system is at the closest one. We run 100 simulations over a horizon of 100 steps. We compare the formulation explicitly accounting for lossy communication to the formulation assuming perfect communication and the baseline strategy of alternating communication between the two systems at each step, which is optimal in case of no packet loss.

We consider three scenarios, with decreasing probability of successful packet reception, as displayed in Figure 4. The time evolution of the probability is periodic and the probability for the two systems is shifted by half a period. The resulting average closed-loop cost is displayed in FIgure 5. We have normalised the cost with respect to the baseline strategy. For the first two scenarios, the prediction horizon does not seem to have an important role. It is interesting to note that, even if perfect communication is assumed, the algorithm uses the information on packet loss in order to improve the closed-loop cost. While in the first scenario
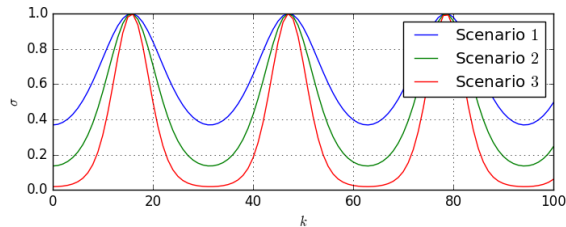
Fig. 4. Periodic probability of having a successful communication for system 1. System 2 has the same probability time evolution but shifted by half a period.
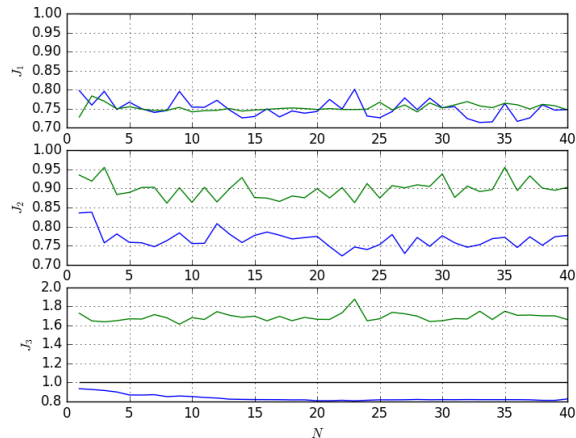

Fig. 5. Average closed-loop cost obtained by assuming perfect communication (green) and by explicitly accounting for the packet drop probability (blue) as a function of the prediction horizon N normalised with respect to the cost obtained with the baseline strategy (black). Top graph: scenario 1. Middle graph: scenario 2. Bottom graph: scenario 3.

including the information on the probability of packet loss does not improve the closed-loop cost, in the second scenario the improvement is significant. Finally, in the third and most challenging scenario, explicitly accounting for the probability of packet drop is crucial for performance, as not doing so significantly deteriorates the performance with respect to the baseline strategy. Note that in this last scenario, longer prediction horizons improve the performance. For prediction horizons $N > 15$, i.e. longer than approximately half a period of the package drop probability variation, the performance does not improve further.

V. Conclusions and Future Research

We have presented an approach for allocating communication between sensors and agents in the case of limited communication resources. The problem formulation minimises the Lyapunov function increase caused by imperfect communication and can incorporate knowledge on packet loss probability in order to improve the closed-loop cost.

In the case of perfect communication, some closed-loop solutions are periodic. This happens especially for small M, while for large amounts of subsystems this periodic behaviour does not appear. While it seems hard to prove that a periodic behaviour is optimal, the fact that the optimiser does not return a periodic solution could be explained that, due to nonconvexity, only locally optimal solutions are obtained. If periodic solutions are sought, the initial constraint in Problem (12) can be replaced by a periodicity constraint and the problem can be solved for different period lengths in order to select the solution with the lowest cost.

Moreover, in a fully distributed setting, the proposed approach can be used in combination with the approach of [4]: the offline deployment of our framework can give indications on how to select the tuning parameters of that algorithm, i.e. the quadratic cost the relative threshold.

References

[1] R. A. Gupta and M. Y. Chow, "Networked Control System: Overview and Research Trends," *IEEE Transactions on Industrial Electronics*, vol. 57, no. 7, pp. 2527–2535, July 2010.

[2] L. Zhang, H. Gao, and O. Kaynak, "Network-Induced Constraints in Networked Control Systems – A Survey," *IEEE Transactions on Industrial Informatics*, vol. 9, no. 1, pp. 403–416, Feb. 2013.

[3] A. Molin, C. Ramesh, H. Esen, and K. H. Johansson, "Innovations-based priority assignment for control over can-like networks," in *2015 54th IEEE Conference on Decision and Control (CDC)*, Dec 2015, pp. 4163–4169.

[4] M. Mamduhi, D. Tolic, and S. Hirche, "Decentralized event-based scheduling for shared-resource networked control systems," in *14th Annual European Control Conference (ECC)*, Jul 2015.

[5] T. Soleymani, S. Hirche, and J. S. Baras, "Optimal self-driven sampling for estimation based on value of information," in *2016 13th International Workshop on Discrete Event Systems (WODES)*, May 2016, pp. 183–188.

[6] L. Zhao, W. Zhang, J. Hu, A. Abate, and C. J. Tomlin, "On the Optimal Solutions of the Infinite-Horizon Linear Sensor Scheduling Problem," *IEEE Transactions on Automatic Control*, vol. 59, no. 10, pp. 2825–2830, Oct 2014.

[7] Y. Mo, E. Garone, and B. Sinopoli, "On infinite-horizon sensor scheduling," *Systems & Control Letters*, vol. 67, pp. 65–70, 2014.

[8] M. P. Vitus, W. Zhang, A. Abate, J. Hu, and C. J. Tomlin, "On efficient sensor scheduling for linear dynamical systems," *Automatica*, vol. 48, no. 10, pp. 2482–2493, 2012.

[9] B. Sinopoli, L. Schenato, M. Franceschetti, K. Poolla, M. I. Jordan, and S. S. Sastry, "Kalman filtering with intermittent observations," *IEEE Transactions on Automatic Control*, vol. 49, no. 9, pp. 1453–1464, Sept 2004.

[10] D. Bertsekas, *Dynamic programming and stochastic control*. Academic Press, 1976.

[11] M. Zanon, S. Gros, and M. Diehl, "Indefinite Linear MPC and Approximated Economic MPC for Nonlinear Systems," *Journal of Process Control*, vol. 24, pp. 1273–1281, 2014.

[12] F. Kozin, "A survey of stability of stochastic systems," *Automatica*, vol. 5, no. 1, pp. 95–112, Jan. 1969.

[13] S. Sager, G. Reinelt, and H. Bock, "Direct Methods With Maximal Lower Bound for Mixed-Integer Optimal Control Problems," *Mathematical Programming*, vol. 118, no. 1, pp. 109–149, 2009.

[14] C. Kirches, "Fast numerical methods for mixed-integer nonlinear model-predictive control," Ph.D. dissertation, Ruprecht-Karls-Universität Heidelberg, 2010.

[15] J. Andersson, "A General-Purpose Software Framework for Dynamic Optimization," PhD thesis, Arenberg Doctoral School, KU Leuven, October 2013.

[16] A. Wächter and L. Biegler, "On the Implementation of a Primal-Dual Interior Point Filter Line Search Algorithm for Large-Scale Nonlinear Programming," *Mathematical Programming*, vol. 106, no. 1, pp. 25–57, 2006.