

Article

A Real-Time Optimization Framework for the Iterative Controller Tuning Problem

Gene A. Bunin, Grégory François and Dominique Bonvin*

Laboratoire d'Automatique, Ecole Polytechnique Fédérale de Lausanne, CH-1015, Lausanne, Switzerland

* Author to whom correspondence should be addressed; dominique.bonvin@epfl.ch, Tel.: +41 21 6933843, Fax: +41 21 6932574.

Version August 23, 2013 submitted to *Processes*. Typeset by \LaTeX using class file *mdpi.cls*


Abstract: We investigate the general iterative controller tuning (ICT) problem, where the task is to find a set of controller parameters that optimize some user-defined performance metric when the same control task is to be carried out repeatedly. Following a repeatability assumption on the system, we show that the ICT problem may be formulated as a real-time optimization (RTO) problem, thus allowing for the ICT problem to be solved in the RTO framework, which is both very flexible and comes with strong theoretical guarantees. In particular, we propose the use of a recently released RTO solver and outline a simple procedure for how this solver may be configured to solve ICT problems. The effectiveness of the proposed method is illustrated by successfully applying it to four case studies – two experimental and two simulated – that cover the tuning of model-predictive, general fixed-order, and PID controllers, as well as a system of controllers working in parallel.

Keywords: controller autotuning; real-time optimization; data-driven tuning methods

1. Introduction

The typical task of a controller consists in tracking a user-specified trajectory as closely as possible while observing certain additional specifications, such as stability, the satisfaction of safety limits, and the minimization of expensive control action when it is not needed. Mathematically, we may define such a controller by the mapping $G_c(\boldsymbol{\rho})$, where $\boldsymbol{\rho} \in \mathbb{R}^{n_\rho}$ denote the parameters that dictate the controller's behavior and represent decision variables (the “tuning parameters”) for the engineer intending to

Figure 1. Qualitative schematic of a single-input-single-output system with the controller $G_c(\rho)$. Elements such as disturbances and sensor dynamics, as well as any controller-specific requirements, are left out for simplicity. We use the notation $y(t, \rho)$ to mark the (implicit) dependence of the control output on the tuning parameters ρ (likewise for the input and the error).



19 implement the controller in practice. In the simplest scenario, this often leads to a closed-loop system
 20 that may be described by the schematic in Figure 1. No assumptions are made on the nature of G_c , which
 21 may represent such controllers as the classical PID, the general fixed-order controller, or even the more
 22 advanced MPC (model-predictive control). To be even more general, G_c may represent an entire system
 23 of such controllers – one would need, in this case, to simply replace $y_{ref}(t)$, $y(t, \rho)$, $u(t, \rho)$, and $e(t, \rho)$
 24 by their vector equivalents.

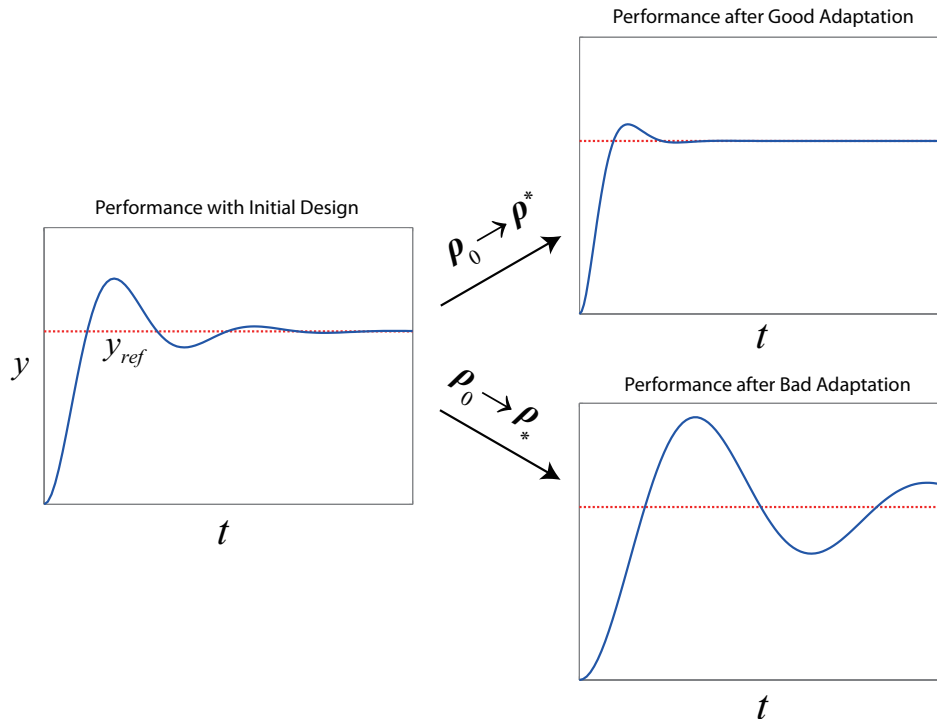
25 As with any set of decision variables, it should be clear that there are both good and bad choices of ρ ,
 26 and in every application some sort of design phase precedes the actual implementation and acts to choose
 27 a set of ρ that is expected to track the reference y_{ref} “well” while meeting any additional specifications.
 28 The classic example for PID controllers is the Ziegler-Nichols tuning method [1], with methods such
 29 as model-based direct synthesis [2] and virtual reference feedback tuning [3] acting as more advanced
 30 alternatives. Though not as developed, both theoretical and heuristic approaches exist for the design of
 31 MPC [4] and general fixed-order controllers [5,6] as well.

32 In the majority of cases, the set of controller parameters obtained by these design methods will not be
 33 the best possible with respect to control performance. There are many reasons for this, with some of the
 34 common ones being:

- 35 • assumptions on the plant, such as linearity or time invariance, that are made at the design stage,
- 36 • modeling errors and simplifications,
- 37 • conservatism in the case of a robust design,
- 38 • time constraints and/or deadlines that give preference to a simpler design over an advanced one.

39 To improve the closed-loop performance of the system, some sort of data-driven adaptation of the
 40 parameters from their initial designed values, denoted here by ρ_0 , may be done online following
 41 the acquisition of new data. These are generally classified as “indirect” and “direct” adaptations
 42 [7] depending on what is actually adapted – the model (followed by a model-based re-design of the
 43 controller) in the indirect variant, or the controller parameters directly in the direct one. This paper

Figure 2. The basic idea of iterative controller tuning. Here, a step change in the setpoint represents the repetitive control task. We use ρ_* to denote a sort of “anti-optimum” that might be achieved with a bad adaptation algorithm.



44 investigates *direct methods that attempt to optimize control performance by establishing a direct link*
 45 *between the observed closed-loop performance and the controller parameters*, with the justification
 46 that such methods may be forced to converge – at least, theoretically – to a locally optimal choice ρ^*
 47 regardless of the quality or the availability of the model, which cannot be said for indirect schemes [8].

48 Many of these schemes attempt to minimize a certain user-defined performance metric (e.g., the
 49 tracking error) for a given run or batch by playing with the controller parameters as one would in an
 50 iterative optimization scheme – i.e., by changing the parameters between two consecutive runs, trying
 51 to discover the effect that this change has on the closed-loop performance (estimating the performance
 52 derivatives), and then using the derivative estimates to adapt the parameters further in some gradient-
 53 descent manner [9,10,11,12,13]. This is essentially the *iterative controller tuning (ICT) problem*, whose
 54 goal is to bring the initial suboptimal set ρ_0 to the locally optimal ρ^* via *iterative experimentation on*
 55 *the closed-loop system*, all the while avoiding that the system become dangerously unstable from the
 56 adaptation (a qualitative sketch of this idea is given in Figure 2). A notable limitation of such methods,
 57 though rarely stated explicitly, is that the control task for which the controller is being adapted must be
 58 identical (or very similar) from one run to the next – otherwise, the concept of optimality may simply
 59 not exist since what is optimal for one control task (e.g., the tracking of a step change) need not be so
 60 for another (e.g., the tracking of a ramp). A closely-related problem where the assumption of a repeated
 61 control task is made formally is that of iterative learning control [14], although what is adapted in that
 62 case is the *open-loop* input trajectory rather than the parameters of a controller dictating the closed-loop
 63 system.

64 We observe that, as the essence of these tuning methods consists in iteratively minimizing a
 65 performance function that is *unknown* due to the lack of knowledge of the plant, the ICT problem is
 66 actually a *real-time optimization* (RTO) problem as it must be solved by iterative experimentation¹.
 67 Recent work by the authors [19,20,21] has attempted to unify different RTO approaches and to
 68 standardize the RTO problem as any problem having the following canonical form:

$$\begin{aligned} & \underset{\mathbf{v}}{\text{minimize}} && \phi_p(\mathbf{v}) \\ & \text{subject to} && \mathbf{G}_p(\mathbf{v}) \preceq \mathbf{0} \\ & && \mathbf{G}(\mathbf{v}) \preceq \mathbf{0} \\ & && \mathbf{v}^L \preceq \mathbf{v} \preceq \mathbf{v}^U \end{aligned}, \quad (1)$$

69 where $\mathbf{v} \in \mathbb{R}^{n_v}$ denote the RTO variables² (RTO inputs) forced to lie in the relevant RTO input space
 70 defined by the lower and upper limits \mathbf{v}^L and \mathbf{v}^U , ϕ_p denotes the cost function to be minimized, and \mathbf{G}_p
 71 and \mathbf{G} denote the sets of individual constraints $g_p, g : \mathbb{R}^{n_v} \rightarrow \mathbb{R}$ (i.e., safety limitations, performance
 72 specifications) to be respected. We use the symbol \preceq to denote *componentwise inequality*.

73 The subscript p (for “plant”) is used to indicate those functions that are unknown, or “uncertain”, and
 74 can only be evaluated by applying a particular \mathbf{v}_k and conducting a single experiment (with k denoting
 75 the experiment/iteration counter), from which the corresponding function values may then be measured
 76 or estimated:

$$\begin{aligned} \hat{\phi}_p(\mathbf{v}_k) &= \phi_p(\mathbf{v}_k) + w_{\phi,k} \\ \hat{g}_p(\mathbf{v}_k) &= g_p(\mathbf{v}_k) + w_{g,k} \end{aligned}, \quad (2)$$

77 with some additive stochastic error w . Conversely, the absence of p indicates that the function is easily
 78 evaluated by algebraic calculation without any error present.

79 Owing to the generality of (1), casting the ICT problem in this form is fairly straightforward and,
 80 as will be shown in this work, has numerous advantages as it allows for a fairly systematic and flexible
 81 approach to controller tuning in a framework where strong theoretical guarantees are available. The main
 82 contribution of our work is thus to make this generalization formally and to argue for its advantages while
 83 cautioning the potential user of both its apparent and hypothetical pitfalls.

84 Our second contribution lies in proposing a concrete method for solving the ICT problem in this
 85 manner. Namely, we advocate the use of the recently released open-source SCFO (“sufficient conditions
 86 for feasibility and optimality”) solver that has been designed for solving RTO problems with strong
 87 theoretical guarantees [21]. While this choice is undoubtedly biased, we put it forward as it is, to the
 88 best of our knowledge, the only solver released to date that solves the RTO problem (1) *reliably*, which
 89 is to say that it consistently converges to a local minimum without violating the safety constraints in
 90 theoretical settings and that it is fairly robust in doing the same in practical ones. Though quite simple

¹This is the definition of “real-time optimization” as it is used in the process systems engineering community (see, e.g., [15,16]) to indicate gradual improvement of an economic objective – the function ϕ_p in (1) – by iterative experimentation. We warn the reader to not confuse this with the more recent use of “real-time optimization” in the fast, on-line computation context (see, e.g., [17,18]), where an optimization problem is solved subject to real-time constraints for a purpose other than the optimization of an economic objective (e.g., for control or estimation).

²It is more common to denote RTO inputs by \mathbf{u} in the literature, but we use \mathbf{v} here so as not to create confusion with the control input $u(t)$.

91 to apply, the SCFO framework and the solver itself need to be properly configured, and so we guide the
92 potential user through how to configurate the solver for the ICT problem.

93 Finally, as the theoretical discussion alone should not be sufficient to convince the reader that there
94 is strong potential for solving the ICT problem as an RTO one, we finish the paper with a total of four
95 case studies, which are intended to cover a diverse range of experimental and simulated problems and
96 to demonstrate the general effectiveness of the proposed method, the difficulties that are likely to be
97 encountered in application, and any weak points where the methodology still needs to be improved.
98 Specifically, the four studies considered all solve the ICT problem for:

- 99 • the tracking of a temperature profile in a laboratory-scale stirred tank by an MPC controller,
- 100 • the tracking of a periodic setpoint for a laboratory-scale torsional system by a general fixed-order
101 controller with a controller stability constraint,
- 102 • the PID tracking of a setpoint change for various linear systems (previously examined in [13,22]),
- 103 • the setpoint tracking and disturbance rejection for a five-input, five-output multi-loop PI system
104 with imperfect decoupling and a hard output constraint.

105 In each case, we do our best to concretize the theory discussed earlier by showing how the resulting ICT
106 problem may be formulated in the RTO framework, followed by the application of the SCFO solver with
107 the proposed configuration.

108 2. The RTO Formulation of the Iterative Controller Tuning Problem

109 In this section, we go through the different components of the RTO problem (1) and state their ICT
110 analogues, together with any assumptions necessary to make the links between the two clean. We then
111 finish by reviewing the benefits and limitations of this approach.

112 2.1. The Cost Function $\phi_p \rightarrow$ The Control Performance Metric

113 The intrinsic driving force behind iteratively tuning a controller so that it performs “better” is the
114 somewhat natural belief that there is some sort of deterministic link between the parameters and the
115 observed performance. We qualify this via the following assumption, which was originally stated in the
116 MPC context in [23] and then extended to the general controller in [24].

117 **Assumption 1** (Repeatability). *Let $\rho \in \mathbb{R}^{n_\rho}$ denote the tuning parameters of a controller and J_k the*
118 *observed value of the user-defined performance metric at run k for a fixed control task that is identical*
119 *from run to run. The closed-loop process is repeatable with respect to performance if*

$$J_k = J(\rho_k) + \delta_k, \quad (3)$$

120 *where ρ_k are the parameters of the controller at run k , $J : \mathbb{R}^{n_\rho} \rightarrow \mathbb{R}$ is a purely deterministic relation*
121 *between the performance metric and the parameters, and δ_k is the “non-repeatability noise”, a purely*
122 *stochastic element that is independent of ρ_k .*

123 In layman's terms, the (unknown) function J is precisely the intrinsic link that we believe in, while
 124 δ_k is a representation of reality, which most often manifests itself by means of measurement noise and
 125 differs unpredictably from run to run. The discussion of the validity of such an assumption is deferred
 126 to the end of the section.

127 Comparing (2) and (3), both of which involve a deterministic function that is sampled with additive
 128 noise, we establish our first RTO \rightarrow ICT connection:

$$\underset{\mathbf{v}}{\text{minimize}} \phi_p(\mathbf{v}) \rightarrow \underset{\boldsymbol{\rho}}{\text{minimize}} J(\boldsymbol{\rho}). \quad (4)$$

129 A common general performance metric, given here in continuous form for the single-input-single-
 130 output (SISO) case, may be defined as:

$$J_k := \lambda_1 \int_0^{t_b} [y_{ref}(t) - y(t, \boldsymbol{\rho}_k)]^2 dt + \lambda_2 \int_0^{t_b} u^2(t, \boldsymbol{\rho}_k) dt + \lambda_3 \int_0^{t_b} \dot{y}^2(t, \boldsymbol{\rho}_k) dt + \lambda_4 \int_0^{t_b} \dot{u}^2(t, \boldsymbol{\rho}_k) dt, \quad (5)$$

131 where t_b denotes the total length of a single run and where the weights $\boldsymbol{\lambda} \succeq \mathbf{0}$ may be set as needed to
 132 trade off between giving preference to tracking error, the control action, the smoothness of the output
 133 response, and the aggressiveness of the controller. Modifications that include other criteria, such as
 134 frequency weighting [10], or that modify the time interval for which the performance is analyzed by
 135 adding a "mask" [22], are of course possible as well.

136 2.2. The Uncertain Inequality Constraints $\mathbf{G}_p \rightarrow$ Safety and Economic Constraints

137 Many control applications may have strict safety specifications that require a given output $y(t, \boldsymbol{\rho})$ to
 138 remain within a certain zone, defined by \underline{y} and \bar{y} , throughout the length of the run:

$$\underline{y} \leq y(t, \boldsymbol{\rho}) \leq \bar{y}, \quad \forall t \in [0, t_b]. \quad (6)$$

139 While it is not difficult to propose methods to enforce such behavior for the general controller, many
 140 of which would likely try to incorporate the constraints as setpoint objectives, such approaches remain
 141 largely *ad hoc*. This drawback has shifted particular emphasis to MPC as being the advanced controller
 142 to be able to deal with output constraints systematically [25], but even here no rigorous conditions for
 143 satisfying (6) are available for the general case where any amount of plant-model mismatch is admissible.

144 Since rigorous theoretical conditions *are* available for satisfying $\mathbf{G}_p(\mathbf{v}) \preceq \mathbf{0}$ in the RTO framework
 145 [19], we may exploit this advantage by casting the hard output constraints for the ICT problem in RTO
 146 form. To do this, we start by replacing the two semi-infinite constraints of (6) by their equivalent finite
 147 versions:

$$\underline{y} \leq \min_{t \in [0, t_b]} y(t, \boldsymbol{\rho})$$

$$\max_{t \in [0, t_b]} y(t, \boldsymbol{\rho}) \leq \bar{y}.$$

148 At this point, we need to apply a version of Assumption 1 for the constraints. For a particular run k ,
 149 we assume that *the closed-loop process is repeatable with respect to the control output range*:

$$\min_{t \in [0, t_b]} y(t, \boldsymbol{\rho}_k) = y_{min}(\boldsymbol{\rho}_k) + \delta_{min,k}$$

$$\max_{t \in [0, t_b]} y(t, \boldsymbol{\rho}_k) = y_{max}(\boldsymbol{\rho}_k) + \delta_{max,k}, \quad (7)$$

150 i.e., that the minimum and maximum values of the trajectory $y(t, \boldsymbol{\rho}_k)$ observed for a given run k are
 151 (unknown) deterministic functions (y_{min} , y_{max}) of the parameters plus a stochastic element ($\delta_{min,k}$,
 152 $\delta_{max,k}$).

153 Making the link with (2), we may now restate the hard output constraints in RTO form as:

$$\mathbf{G}_p(\mathbf{v}) \preceq \mathbf{0} \rightarrow \begin{array}{l} -y_{min}(\boldsymbol{\rho}) + \underline{y} \leq 0 \\ y_{max}(\boldsymbol{\rho}) - \bar{y} \leq 0 \end{array}, \quad (8)$$

154 where the function values y_{min} and y_{max} can be measured for a given $\boldsymbol{\rho}$ with the additive errors δ_{min} and
 155 δ_{max} .

156 Alternatively, it may occur that there are economic constraints with respect to the inputs. As an
 157 example, consider a reactor where one of the control inputs is the feed rate of a reagent. While effective
 158 for the purposes of control, the reagent may be expensive and so only a limited amount may be allotted
 159 per batch, with the constraint

$$\int_0^{t_b} u(t, \boldsymbol{\rho}) dt \leq \bar{u}_T$$

160 imposed, where \bar{u}_T is some user-defined limit. Following the same steps as above, we suppose that

$$\int_0^{t_b} u(t, \boldsymbol{\rho}_k) dt = u_T(\boldsymbol{\rho}_k) + \delta_{u,k},$$

161 with u_T the deterministic component and δ_u the non-repeatability noise, and make the connection:

$$\mathbf{G}_p(\mathbf{v}) \preceq \mathbf{0} \rightarrow u_T(\boldsymbol{\rho}) - \bar{u}_T \leq 0.$$

162 It should be clear that extension to multiple-input-multiple-output (MIMO) cases is trivial, as this
 163 only adds more elements to \mathbf{G}_p .

164 2.3. The Certain Inequality Constraints $\mathbf{G} \rightarrow$ Controller Specifications and Stability Considerations

165 In some controllers, analytically known inequality relations may need to be satisfied. One such
 166 example is the case of the MPC controller, where one may tune both the control and prediction horizons
 167 (m and n , respectively) with the built-in rule [25]:

$$m \leq n, \quad (9)$$

168 which, if we define $\rho_1 \triangleq m$ and $\rho_2 \triangleq n$, leads to the following link with (1):

$$\mathbf{G}(\mathbf{v}) \preceq \mathbf{0} \rightarrow \rho_1 - \rho_2 \leq 0.$$

169 As another example, we may want to adapt the parameters of the discrete fixed-order controller:

$$G_c(\boldsymbol{\rho}) = \frac{\rho_1 z^2 + \rho_2 z + \rho_3}{z^2 + \rho_4 z + \rho_5}, \quad (10)$$

170 but would like to limit our search to stable controllers only. Employing the Jury stability criterion [26],
 171 we generate the first four rows of the Jury table for the denominator of $G_c(\boldsymbol{\rho})$:

$$\begin{array}{l} \text{row 1 :} \\ \text{row 2 :} \\ \text{row 3 :} \\ \text{row 4 :} \end{array} \begin{array}{ccc} 1 & \rho_4 & \rho_5 \\ \rho_5 & \rho_4 & 1 \\ 1 - \rho_5^2 & \rho_4 - \rho_4 \rho_5 & 0 \\ \rho_4 - \rho_4 \rho_5 & 1 - \rho_5^2 & 0 \end{array},$$

172 from which the sufficient conditions for controller stability are obtained as:

$$\begin{aligned} |\rho_5| < 1 & \rightarrow |\rho_5| \leq 1 - \varepsilon \\ |\rho_4 - \rho_4\rho_5| < |1 - \rho_5^2| & \rightarrow |\rho_4 - \rho_4\rho_5| \leq |1 - \rho_5^2| - \varepsilon \end{aligned} \quad (11)$$

173 with the constraint set on the right representing an implementable nonstrict version with negligible
174 conservatism for $\varepsilon > 0$ small. Controller stability may now be ensured in RTO form with the
175 correspondance:

$$\mathbf{G}(\mathbf{v}) \preceq \mathbf{0} \rightarrow \begin{aligned} |\rho_5| - 1 + \varepsilon &\leq 0 \\ |\rho_4 - \rho_4\rho_5| - |1 - \rho_5^2| + \varepsilon &\leq 0 \end{aligned} \quad (12)$$

176 Finally, we note that *nominal* closed-loop stability constraints may also be incorporated in this
177 manner. As a simple example, consider the unstable plant that is modeled as

$$G(s) = \frac{1}{s-1},$$

178 and that is to be controlled by a PD controller with ρ_1 and ρ_2 the proportional and derivative gains,
179 respectively:

$$G_c(\boldsymbol{\rho}) = \rho_1 + \rho_2 s.$$

180 From the analysis of the characteristic equation $1 + GG_c = 0$, we have the stability condition, together
181 with its implementable version:

$$\frac{1 - \rho_1}{1 + \rho_2} < 0 \rightarrow \frac{1 - \rho_1}{1 + \rho_2} \leq -\varepsilon,$$

182 which, again, allows the correspondance:

$$\mathbf{G}(\mathbf{v}) \preceq \mathbf{0} \rightarrow \frac{1 - \rho_1}{1 + \rho_2} + \varepsilon \leq 0. \quad (13)$$

183 Extensions to robust nominal stability follow easily, and would simply involve a greater number of
184 constraints.

185 2.4. The Box Constraints $\mathbf{v}^L \preceq \mathbf{v} \preceq \mathbf{v}^U \rightarrow$ Controller Parameter Limits

186 Given the RTO-ICT correspondence of $\mathbf{v} \rightarrow \boldsymbol{\rho}$, the box constraints of the RTO problem are simply
187 the lower and upper limits, $\boldsymbol{\rho}^L$ and $\boldsymbol{\rho}^U$, on the adapted controller parameters. We note that certain limits
188 will be obvious for certain controllers – e.g., the integral time should be superior to 0 in a PI controller,
189 while the prediction and control horizons of an MPC controller should be equal to or greater than 1. In
190 other cases, one may have to think a little before deciding on appropriate limits. In Section 3, an easy
191 way to set parameter bounds for the general controller will be provided.

192 2.5. The ICT Problem in RTO Form: Summary

193 Having now gone through all the components of Problem (1) and having provided their ICT
194 analogues, we may make certain remarks and observations.

195 To start with the positive, almost all of the possible desired specifications in a standard ICT problem
196 are easily stated in RTO terms, although this is not surprising given the generality of (1). Of particular
197 interest with regard to this point are the constraint terms, as the flexibility of the RTO formulation has

198 allowed for us to include limits on the control outputs and inputs, as well as any controller specifications,
199 very easily. To the best of the authors' knowledge, constraints are generally avoided in the majority
200 (though not all [27]) of direct tuning formulations. This is likely because the most commonly used
201 method – the gradient descent – is not well-equipped to deal with them (apart from certain simple kinds,
202 such as the box constraints [23]). Casting the ICT problem in the RTO framework therefore allows us to
203 ignore this limitation.

204 The other big advantage is that no assumptions are needed on the nature of the controller (or
205 their number, if a system of controllers is considered) – it simply has to be something that can be
206 parametrically tuned, and so one could adapt just about anything. Likewise, the standard restricting
207 assumption of linearity is also not needed, even formally, as the black-box nature of the RTO formulation
208 does not make use of such assumptions since it ignores the actual dynamic behavior of the closed loop
209 and only considers the RTO inputs (the tuning parameters) and RTO outputs (performance, proximity
210 to constraints), both of which are static quantities with a static map between them. As such, the
211 methodology applies just as readily to nonlinear systems as it does to linear ones.

212 There are, however, points to be contested. The key linking element between RTO and ICT is
213 Assumption 1, which is, at best, only an approximation and merits justification. The driving force behind
214 this assumption is the fact that any deterministic³ controller with fixed tuning parameters, when applied
215 repeatedly to a closed-loop process to perform the same control task, should always yield the exact same
216 performance (and the exact same input/output trajectories) in the absence of non-repeatable effects such
217 as *input/measurement noise, process degradation, and disturbances*. Indeed, the absence of such effects
218 implies that the δ_k term in (3) is equal to 0, and that the repeatability assumption holds exactly. For
219 cases when these effects are minor and do not influence controller behavior significantly, we expect that
220 a given controller will yield the same performance “more or less”, with variations being lumped into δ_k
221 and the major deterministic trends being described by J . This neat way of decoupling the deterministic
222 and stochastic components may not be valid when the non-repeatable effects become large and exert a
223 significant influence on the controller behavior, however. As such, we may view this assumption as an
224 approximation of reality that tends to perfection as the magnitude of the noise/degradation/disturbances
225 in the closed-loop system tends to 0.

226 There is, as well, the issue of stability. Even with the direct incorporation of constraints in the RTO
227 problem formulation (e.g., via (12) or (13)), there is no true way to incorporate a real constraint on
228 closed-loop stability, as stability is not a real-numbered value that can be measured following a closed-
229 loop experiment (if it were, it would be trivial to include it as an uncertain constraint in the set \mathbf{G}_p).
230 Unfortunately, this is a much bigger problem that is not limited to just ICT – one cannot, for the general
231 unknown plant, *ever* guarantee stability via *any* means without making additional assumptions on the
232 nature of the plant. The bright side is that any of the standard stability-guaranteeing methods are easily
233 incorporated into the RTO formulation as certain constraints \mathbf{G} , and may be used to limit the adaptations
234 to those controllers that are at least nominally stable. Other workarounds could also be proposed – if
235 the fear of having an unstable closed-loop system stems from having some control output leave its safe

³While almost all applied controllers are deterministic, we acknowledge the possibility of stochastic controllers (e.g., those employing Monte Carlo techniques [28]), for which the methods discussed in this work may not be applicable.

operating range, then one could simply introduce an output constraint on that quantity, which, as already shown, is easily integrated into the RTO formulation as \mathbf{G}_p .

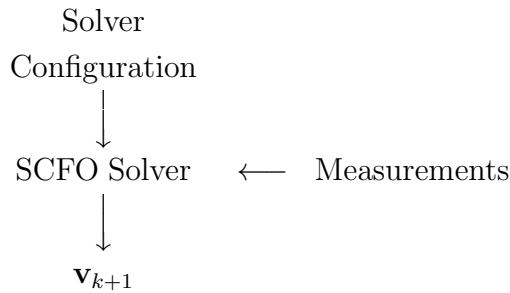
3. The SCFO Solver and Its Configuration

Having now presented the formulation of the ICT problem as an RTO one, we go on to describe how Problem (1) may be solved. Although (1) is posed like a standard optimization problem, the reader is warned that it is *experimental* in nature and must be solved by iterative closed-loop experiments on the system – i.e., one cannot simply solve (1) by numerical methods since evaluations of functions ϕ_p and \mathbf{G}_p require experiments. A variety of RTO (or “RTO-like”) methodologies, all of which are appropriate for solving (1), have been proposed over the years and may be characterized as being model-based (see, e.g., [29,30,31,32]), model-free [33,34], or as hybrids of the two [35,36]. In this work, we opt to use the SCFO solver recently proposed and released by the authors [19,20,21], as it is the only tool available to theoretically guarantee that:

- the RTO scheme converges arbitrarily close to a Karush-Kuhn-Tucker (KKT) point that is, in the vast majority of practical cases, a local minimum,
- the constraints $\mathbf{G}_p(\mathbf{v}) \preceq \mathbf{0}$ and $\mathbf{G}(\mathbf{v}) \preceq \mathbf{0}$ are *never* violated,
- the objective value is consistently improved, with $\phi_p(\mathbf{v}_{k+1}) < \phi_p(\mathbf{v}_k)$ always,

with these properties enforced approximately in practice.

The basic structure of the solver may be visualized as follows:



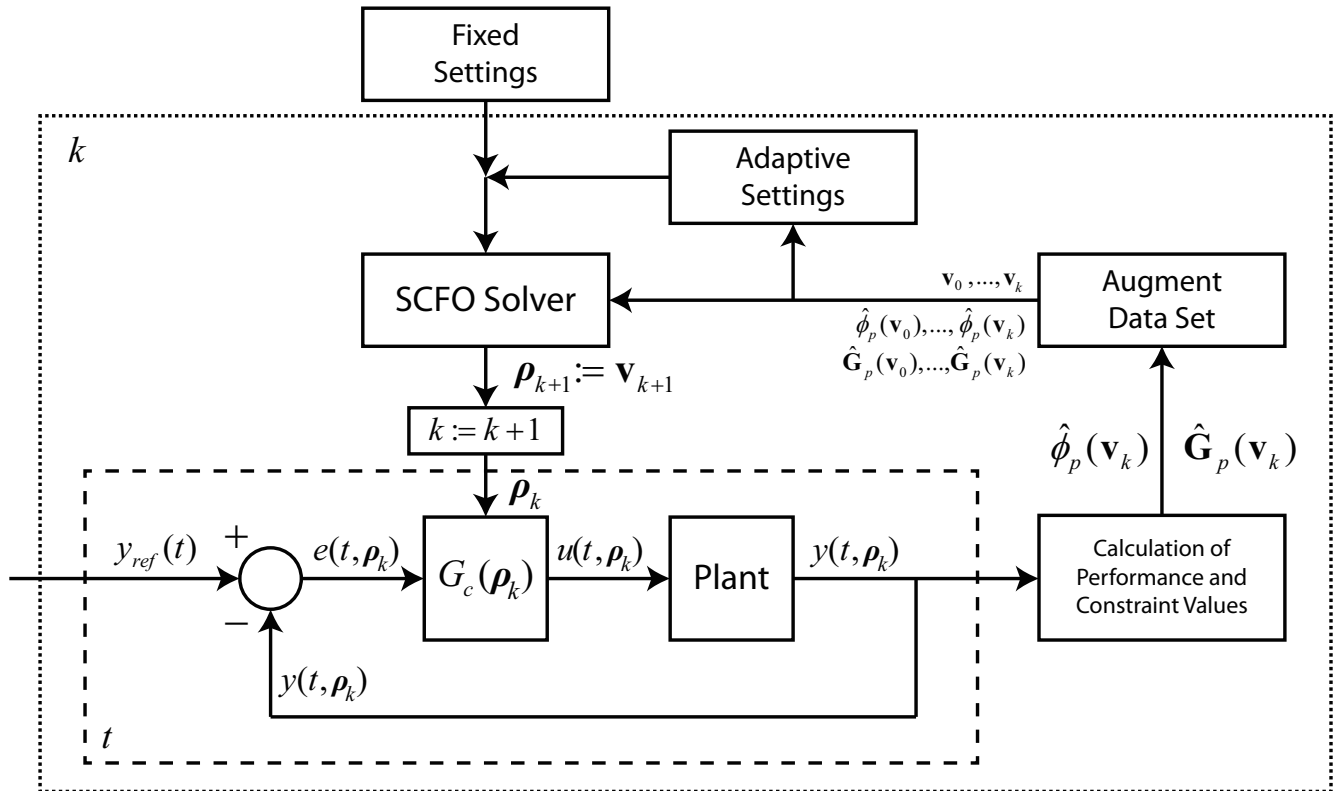
where the majority of the configuration is fixed once and for all while the measurements act as the true iterative components, with the full set of measured data being fed to the solver at each iteration, after which it does all of the necessary computations and outputs the next RTO input to be applied. This is illustrated for the ICT context in Figure 3 (as an extension of Figure 1).

The natural price to pay for such simplicity of implementation is, not surprisingly, the complexity of configuration. Table 1 provides a summary of all of the configuration components, how they are set, and the justifications for these settings. Noting that most of these settings are relatively simple and do not merit further discussion, we now turn our focus to those that do.

3.1. Solver Initialization

Prior to attempting to solve Problem (1), it is strongly recommended that the problem be well-scaled with respect to both the RTO inputs and outputs. For the former, this means that:

Figure 3. The iterative tuning scheme, where the results obtained after each closed-loop experiment on the plant (denoted by the dashed lines) are sent to the RTO loop (denoted by the dotted box), which then appends these data to previous data and uses the full data set to prompt the SCFO solver, as well as to update any data-driven adaptive settings (we refer the reader to Table 1 for which settings are fixed and which are adaptive).



$$v_1^U - v_1^L \approx v_2^U - v_2^L \approx \dots \approx v_{n_v}^U - v_{n_v}^L \approx 1,$$

265 where “ \approx ” may be read as “on the same order of magnitude as”. For the RTO outputs, it is advised that
 266 both the cost and constraint functions are such that their values vary on the magnitude of 10^0 . Once this
 267 is done, one may proceed to initialize the data set.

268 As the solver needs to compute gradient estimates directly from measured data, it is usually needed
 269 to generate the $n_v + 1$ measurements (whose corresponding RTO input values should be well-poised for
 270 linear regression) necessary for a rudimentary (linear) gradient estimate (see, e.g., [30,37]). In the case
 271 that previous measurements are already available (e.g., from experimental studies carried out prior to
 272 optimization), then one may be able to avoid this step partially or entirely.

273 We will, for generality, assume the case where no previous data are available. We will also assume that
 274 the initial point, $\mathbf{v}_0 := \boldsymbol{\rho}_0$, has been obtained by some sort of controller design technique. In addition,
 275 we require that the initial design satisfy $\mathbf{G}_p(\mathbf{v}_0) \prec \mathbf{0}$, $\mathbf{G}(\mathbf{v}_0) \prec \mathbf{0}$, and $\mathbf{v}^L \prec \mathbf{v}_0 \prec \mathbf{v}^U$ – this is expected
 276 to hold intrinsically since one would not start optimizing performance prior to having at least one design
 277 that is known to meet the required constraints with at least some safety margin. The next step is then to
 278 generate n_v additional measurements, i.e., to run n_v (n_p) closed-loop experiments on the plant.

Table 1. Summary of SCFO configuration settings for the ICT problem.

Solver Setting	Chosen As	Justification	Type
Initialization	$n_\rho + 1$ closed-loop experiments	See Section 3.1	–
Optimization target	Scaled gradient descent	See Section 3.2	Adaptive
Noise statistics	Initial experiments at ρ_0	See Section 3.3	Fixed
Constraint concavity	None assumed	No reason for assuming this property in ICT context	Fixed
Constraint relaxations	None assumed	For simplicity (should be added if some constraints are soft)	Fixed
Cost certainty	Cost function is uncertain	The performance metric is an unknown function of ρ	Fixed
Structural assumptions	Locally quadratic structure	Recommended choice for general RTO problem [21]	Fixed
Minimal-excitation radius	$0.01 (\rho_1^U - \rho_1^L)$	Recommended choice for general RTO problem [21]	Fixed
Lower and upper limits, \mathbf{v}^L and \mathbf{v}^U	Controller-dependent or set adaptively	See Section 3.4	Fixed/ Adaptive
Lipschitz and quadratic bound constants	Initial data-driven guess followed by adaptive setting	See Section 3.5	Fixed/ Adaptive
Scaling bounds	Problem-dependent; easily chosen	See [21]	Fixed
Maximal allowable adaptation step, $\Delta \mathbf{v}_{max}$	$0.1 (\rho^U - \rho^L)^T$	Recommended choice for general RTO problem [21]	Fixed

279 A simple initialization method would be to perturb each controller parameter one at a time, as this
280 would produce a well-poised data set with sufficient excitation in all input directions, thereby making the
281 task of estimating the plant gradient possible. However, such a scheme could be wasteful, especially for
282 ICT problems with many parameters to be tuned. One alternative would be to use smart, model-based
283 initializations [30], but this would require having a plant model. In the case of no model, we propose to
284 use a “smart” perturbation scheme that attempts to begin optimizing performance during the initialization
285 phase, and refer the reader to the appendix for the detailed algorithm.

286 3.2. The Optimization Target

287 The target \mathbf{v}_{k+1}^* represents a nominal optimum provided by any standard RTO algorithm that is
288 coupled with the SCFO solver, and as such actually represents the choice of algorithm. This choice
289 is important as it affects performance, with some of the results in [20] suggesting that coupling the
290 SCFO with a “strong” RTO algorithm (e.g., a model-based one) can lead to faster convergence to the
291 optimum. However, the choice is *not* crucial with respect to the reliability of the overall scheme, and so
292 one does not need to be overly particular about what RTO algorithm to use, but should prefer one that
293 generally guides the adaptations in the right direction.

294 For the sake of simplicity, the algorithm adopted in this work is the (scaled) gradient descent with a
 295 unit step size:

$$\mathbf{v}_{k+1}^* = \mathbf{v}_k - \mathbf{H}_k^\dagger \nabla \hat{\phi}_p(\mathbf{v}_k), \quad (14)$$

296 where both \mathbf{H}_k and $\nabla \hat{\phi}_p(\mathbf{v}_k)$ are data-driven estimates. We refer the reader to the appendix for how these
 297 estimates are obtained.

298 3.3. The Noise Statistics

299 Obtaining the statistics (i.e., the probability distribution function, or PDF) for the stochastic error
 300 terms δ in (3) and (7) is particularly challenging in the ICT context. One reason for this is that
 301 these terms do not have an obvious physical meaning, as both (3) and (7), which model the observed
 302 performance/constraint values as a sum of a deterministic and stochastic component, are approximations.
 303 Furthermore, even if this model were correct, the actual computation of an accurate PDF would likely
 304 require a number of closed-loop experiments on the plant that would be judged as excessive in practice.

305 As will be shown in the first two case studies of Section 4, some level of engineering approximation
 306 becomes inevitable in obtaining the PDF for an experimental system. The basic procedure advocated
 307 here is to carry out a certain (economically allowable) number of repeated experiments for $\mathbf{v}_0 := \boldsymbol{\rho}_0$
 308 *prior to the initialization step*. In the case where each experiment is expensive (or time consuming) and
 309 the total acceptable number is low, one may approximate the δ term by modeling the observed values by
 310 a zero-mean normal distribution with a standard deviation equal to that of the data. If the experiments
 311 are cheap and a fairly large number (e.g., a hundred or more) is allowed, then the observed data may be
 312 offset by its mean and then fed directly into the solver (as the solver builds an approximate PDF directly
 313 from the fed noise data).

314 3.4. Lower and Upper Input Limits

315 Providing proper lower and upper limits \mathbf{v}^L and \mathbf{v}^U can be crucial to solver performance. As already
 316 stated, for the ICT problem these are simply $\mathbf{v}^L := \boldsymbol{\rho}^L$ and $\mathbf{v}^U := \boldsymbol{\rho}^U$, but, as these values may not be
 317 obvious for certain controller designs, the user may use adaptive limits that are redefined at each iteration
 318 k :

$$\begin{aligned} \boldsymbol{\rho}_k^L &:= \boldsymbol{\rho}_k - 0.5 \\ \boldsymbol{\rho}_k^U &:= \boldsymbol{\rho}_k + 0.5 \end{aligned} \quad (15)$$

319 As the solver can never actually converge to an optimum that touches these limits, the resulting problem
 320 is essentially unconstrained with respect to them, thereby allowing us to configure the solver without
 321 affecting the optimality properties of the problem. We note that, while one could use very conservative
 322 choices and not adapt them (e.g., $\boldsymbol{\rho}^L := -1000$ and $\boldsymbol{\rho}^U := 1000$), this is not recommended as it would
 323 introduce scaling issues into the solver's subroutines.

324 3.5. Lipschitz and Quadratic Bound Constants

325 The solver requires the user to provide the Lipschitz constants (denoted by κ) for all of the functions
326 ϕ_p , \mathbf{G}_p , and \mathbf{G} . These are implicitly defined as:

$$\underline{\kappa}_{\phi,i} < \left. \frac{\partial \phi_p}{\partial v_i} \right|_{\mathbf{v}} < \overline{\kappa}_{\phi,i}, \quad \underline{\kappa}_{p,i} < \left. \frac{\partial g_p}{\partial v_i} \right|_{\mathbf{v}} < \overline{\kappa}_{p,i}, \quad \underline{\kappa}_i < \left. \frac{\partial g}{\partial v_i} \right|_{\mathbf{v}} < \overline{\kappa}_i,$$

327 for all $\mathbf{v} \in \{\mathbf{v} : \mathbf{v}^L \preceq \mathbf{v} \preceq \mathbf{v}^U\}$. Quadratic bound constants (denoted by M) on the cost function are
328 also required, and are implicitly defined as:

$$\underline{M}_{ij} < \left. \frac{\partial^2 \phi_p}{\partial v_i \partial v_j} \right|_{\mathbf{v}} < \overline{M}_{ij}, \quad \forall \mathbf{v} \in \{\mathbf{v} : \mathbf{v}^L \preceq \mathbf{v} \preceq \mathbf{v}^U\}.$$

329 For \mathbf{G} , which is easily evaluated numerically, we note that the choice is simple since one can, in many
330 cases, compute these values prior to any implementation.

331 For $\kappa_{\phi,i}$, $\kappa_{p,i}$, and M_{ij} , the choice is a very difficult one. This is especially true for the ICT problem
332 where such constants have no physical meaning, a trait that may make them easier to estimate for some
333 RTO problems [20]. When a model of the plant is available, one may proceed to compute these values
334 numerically for the modeled closed-loop behavior, and then make the estimates more conservative (e.g.,
335 by applying a safety-factor scaling) to account for plant-model mismatch.

336 For the pure model-free case, we have no choice but to resort to heuristic approaches. As a choice of
337 $\kappa_{\phi,i}$, we thus propose the following (very conservative) estimate based on the gradient estimate for the
338 initial $n_v + 1$ points (23):

$$\overline{\kappa}_{\phi,i}, \underline{\kappa}_{\phi,i} := \pm 10 \|\nabla \hat{\phi}_p\|_{\infty}, \quad i = 1, \dots, n_v,$$

339 as we expect these bounds to be valid unless $\|\nabla \hat{\phi}_p\|_{\infty}$ is small, which, however, would indicate that we
340 are probably close to a zero-gradient stationary point already, and would have little to gain by trying to
341 optimize performance further if this point were a minimum.

342 A similar rule is applied to estimate $\kappa_{p,i}$, with:

$$\overline{\kappa}_{p,i}, \underline{\kappa}_{p,i} := \pm 2 \|\nabla \hat{g}_p\|_{\infty}, \quad i = 1, \dots, n_v,$$

343 where the estimate $\nabla \hat{g}_p$ is obtained in the same manner as (23). The choice of 2, as opposed to 10,
344 is made for performance reasons, as making $\kappa_{p,i}$ too conservative can lead to very slow progress in
345 improving performance – this is expected to scale linearly, i.e., if the choice of $\pm 2 \|\nabla \hat{g}_p\|_{\infty}$ leads to a
346 realization that converges in 20 runs, the choice of $\pm 10 \|\nabla \hat{g}_p\|_{\infty}$ may lead to one that converges in 100.
347 Note, however, that this way of defining the Lipschitz constants does not have the same natural safeguard
348 as it does for the cost, and it may happen that $\|\nabla \hat{g}_p\|_{\infty} \approx 0$ at the initial point even though the gradient
349 may be quite large in the neighborhood of the optimum. When this is so, an alternate heuristic choice is
350 to set:

$$\overline{\kappa}_{p,i}, \underline{\kappa}_{p,i} := \pm 2 \frac{-g_p}{v_i^U - v_i^L}, \quad i = 1, \dots, n_v,$$

351 where \underline{g}_p denotes the smallest value that the constraint can take in practice, with $\underline{g}_p \leq g_p(\mathbf{v})$, $\forall \mathbf{v} \in \{\mathbf{v} : \mathbf{v}^L \preceq \mathbf{v} \preceq \mathbf{v}^U\}$. Combining the two, one may then use the heuristic rule:

$$\overline{\kappa}_{p,i}, \underline{\kappa}_{p,i} := \pm 2 \max \left(\|\nabla \hat{g}_p\|_{\infty}, \frac{-\underline{g}_p}{v_i^U - v_i^L} \right), \quad i = 1, \dots, n_v.$$

353 However, it may still occur that this choice is not conservative enough. This lack of conservatism may
 354 be proven if a given constraint $g_p(\mathbf{v}) \leq 0$ is violated for one of the runs, since sufficiently conservative
 355 Lipschitz estimates will usually guarantee that this is not the case (provided that the noise statistics are
 356 sufficiently accurate). As such, the following adaptive refinement of the Lipschitz constants is proposed
 357 to be done online when/if the constraint is violated with sufficient confidence:

$$g_p(\mathbf{v}_k) \geq 3\sigma_g \rightarrow \bar{\kappa}_{p,i} := 2\bar{\kappa}_{p,i}, \underline{\kappa}_{p,i} := 2\underline{\kappa}_{p,i},$$

358 where σ_g represents the estimated standard deviation of the non-repeatability noise term δ for g_p .

359 For the quadratic bound constants M , which represent lower and upper bounds on the second
 360 derivatives of ϕ_p , we propose to use the estimate of the Hessian \mathbf{H}_k as obtained in Section 3.2 (see
 361 Appendix), together with a safety factor, η , to define the bounds at each iteration k as:

$$\bar{M}_{ij}, \underline{M}_{ij} := H_{k,ij} \pm \eta |H_{k,ij}|, \quad (16)$$

362 with η initialized as 1. Since such a choice may also suffer from a lack of conservatism, an adaptive
 363 algorithm for η is put into place. Since a common indicator of choosing M values that are not
 364 conservative enough is the failure to decrease the cost between consecutive iterations, the following
 365 law is proposed for any iterations where the solver applied the SCFO conditions but increased (with
 366 sufficient confidence) the value of the cost [21]:

- 367 • If $\hat{\phi}_p(\mathbf{v}_k) - 4\sigma_\phi \geq \min_{i=0,\dots,k-1} \hat{\phi}_p(\mathbf{v}_i)$, then set $\eta := \eta + 1$,
- 368 • otherwise, set $\eta := \eta - 0.5$, with $\eta < 0 \rightarrow \eta := 0$,

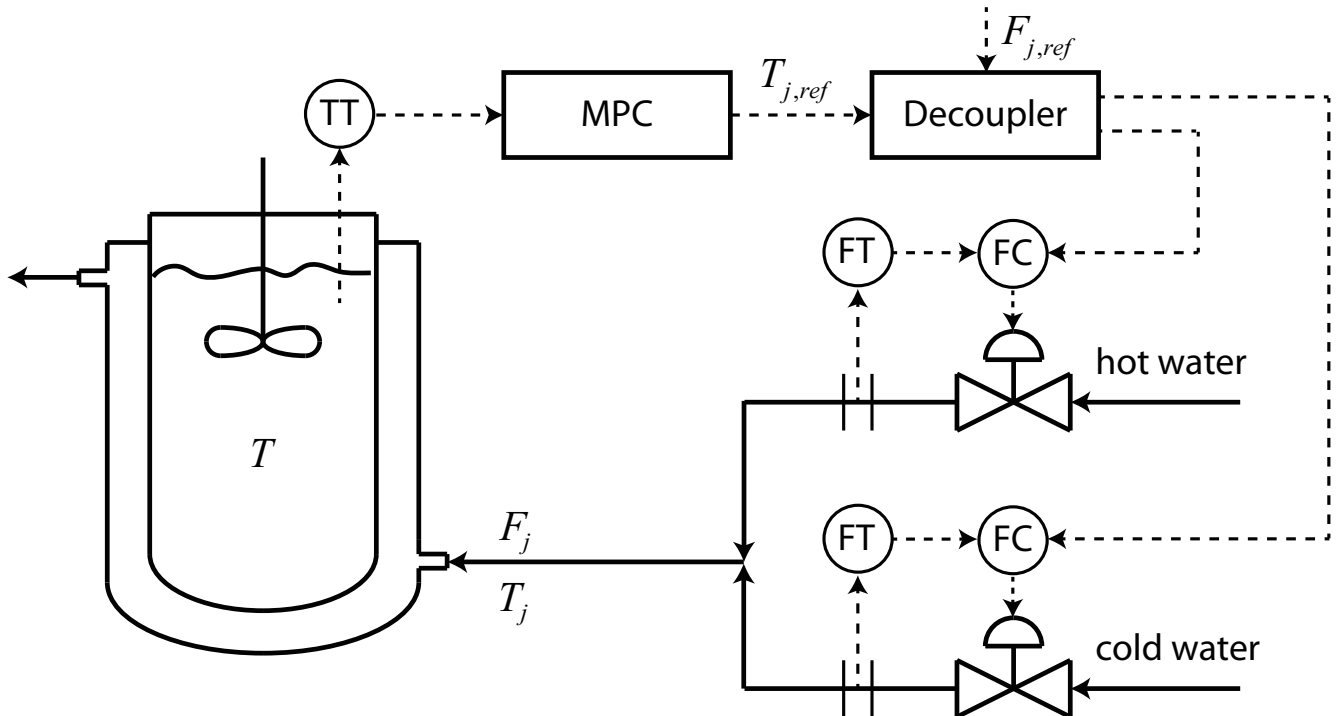
369 where σ_ϕ is the estimated standard deviation for the non-repeatability noise term in the measurement
 370 of J_k . The essence of this update law is to make M more conservative (by increasing η) whenever the
 371 performance is statistically likely to have increased in the recent adaptation, and to relax the conservatism
 372 otherwise, though at only twice the rate that it would be increased. Such a scheme essentially ensures
 373 that the M constants become conservative enough to continually guarantee improved performance with
 374 an increasing number of iterations.

375 4. Case Studies

376 The proposed method was applied to four different problems, of which the first two are of
 377 particular interest as they were carried out on experimental systems and demonstrate the reliability
 378 and effectiveness of the proposed approach when applied in settings where neither the plant nor the
 379 non-repeatability noise terms are known. Of these two, the first represents a typical batch scenario
 380 with fairly slow dynamics and time-consuming, expensive experiments for which an MPC controller
 381 is employed (Section 4.1), while the latter represents a much faster mechanical system where the
 382 optimization of the controller parameters for a general fixed-order controller must be carried out quickly
 383 due to real-time constraints, but where a single run is inexpensive (Section 4.2).

384 The last two studies, though lacking the experimental element, are nevertheless of interest as they
 385 make a link with similar work carried out by other researchers (Section 4.3) and generalize the method
 386 to *systems* of controllers with an additional challenge in the form of an output constraint (Section 4.4).

Figure 4. Schematic of the jacketed stirred tank and the cascade control system used to control the water temperature inside the tank. The reference ($F_{j,ref}$) for the water flow to the jacket (F_j) was fixed at 2 L/min.



387 In both of these cases, we have chosen to simplify things by assuming to know the noise statistics of the
 388 relevant δ terms and to let the repeatability assumption hold exactly.

389 In each of the four studies, we have used the configuration proposed in Section 3 and so will not
 390 repeat these details here. However, we will highlight those components of the configuration that are
 391 problem-dependent and will explain how we obtained them for each case.

392 4.1. Batch-to-Batch Temperature Tracking in a Stirred Tank

393 The plant in question is a jacketed stirred water tank, where a cascade system is used to control the
 394 temperature inside the tank by having an MPC controller manipulate the setpoint temperature of the
 395 jacket inlet, which is in turn tracked by a decoupled system of two PI controllers that manipulate the
 396 flow rates of the hot and cold water streams that mix to form the jacket inlet (Figure 4). As this system is
 397 essentially identical to what has been previously reported [38], we refer the reader to the previous work
 398 for all of the implementation details.

399 As the task of tracking an “optimal” temperature profile is fairly common in batch processes and the
 400 failure to do so well can lead to losses in product quality, a natural ICT problem arises in these contexts as
 401 it is desired that the temperature stay as close to the prescribed optimal setpoint trajectory as possible. In
 402 this particular case study, the controller that is tasked with this job is the MPC controller whose tunable
 403 parameters include:

- 404 • the output weight that controls the trade-off between controller aggressiveness and output tracking,
- 405 • the bias update filter gain, which acts to ensure offset-free tracking,

406 • the control and prediction horizons that dictate how far ahead the MPC attempts to look and
407 control,

408 all of which act to change the objective function at the heart of the MPC controller [38]. For this problem,
409 we decided to vary the output weight between 0.1 and 10 (i.e., covering three orders of magnitude) and
410 defined its logarithm as the first tunable variable ρ_1 . Our reason for choosing the logarithm, instead of
411 the actual value, was due to the sensitivity of the performance being more uniform with respect to the
412 magnitude difference between the priorities given to controller aggressiveness and output tracking (e.g.,
413 changing the output weight from 0.1 to 1.0 was expected to have a similar effect as changing it from
414 1.0 to 10). The bias update filter gain, defined as the second variable ρ_2 , was forced to vary between 0
415 and 1 by definition. The control and prediction horizons, m and n , were both allowed to vary anywhere
416 between 2 and 50 and, as this variance was on the magnitude of 10^2 , were divided by 100 so as to have
417 comparable scaling with the other parameters, with $\rho_3 \triangleq m/100$ and $\rho_4 \triangleq n/100$. We note as well that the
418 horizons were constrained to be integers, whereas the solver provided real numbers, and so any answer
419 provided by the solver had to be rounded to the nearest integer to accommodate these constraints.

420 As this system was fairly slow/stable and controller aggressiveness was not really an issue, and as
421 there was no strong preference between using hot or cold water, the performance metric simply consisted
422 of minimizing the tracking error (i.e., the general metric in (5) with $\lambda_1 := 1$ and $\lambda_2 := \lambda_3 := \lambda_4 := 0$)
423 over a batch time of $t_b = 40$ minutes. The setpoint trajectory to be tracked consisted of maintaining
424 the temperature at 52°C for 10 minutes, cooling by 4°C over 10 minutes, and then applying a quadratic
425 cooling profile for the remainder of the batch. Each batch was initialized by setting the jacket inlet to
426 55°C and starting the batch once the tank temperature rose to 52°C .

427 The certain inequality constraint $\rho_3 \leq \rho_4$ was enforced as this was needed by definition – see (9) –
428 thereby contributing to yield the following ICT problem in RTO form:

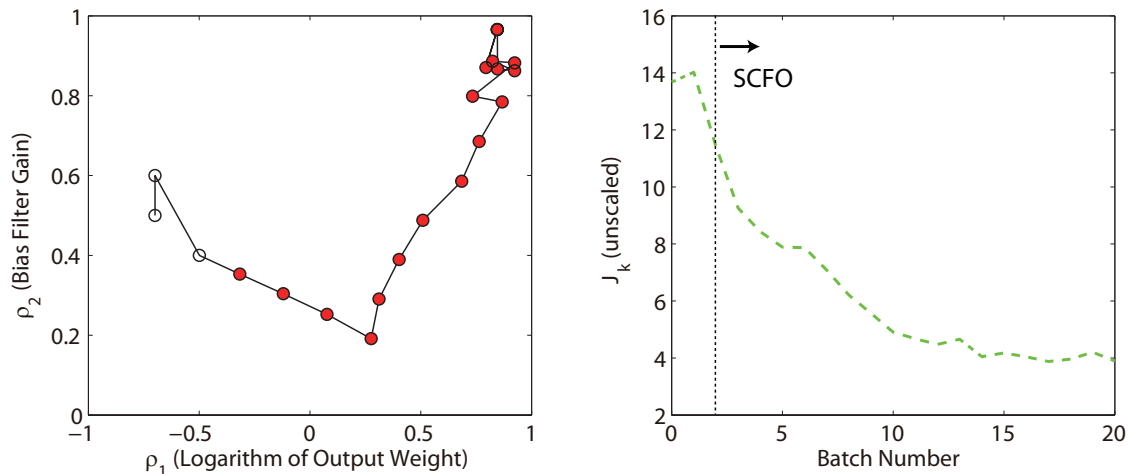
$$\begin{array}{ll}
 \underset{\boldsymbol{\rho}}{\text{minimize}} & \frac{1}{J_0} \int_0^{40} [T_{ref}(t) - T(t, \boldsymbol{\rho})]^2 dt \\
 \text{subject to} & \rho_3 - \rho_4 \leq 0 \\
 & -1 \leq \rho_1 \leq 1 \\
 & 0 \leq \rho_2 \leq 1 \\
 & 0.02 \leq \rho_3 \leq 0.50 \\
 & 0.02 \leq \rho_4 \leq 0.50
 \end{array}
 \left. \begin{array}{l}
 \right\} \phi_p(\mathbf{v}) \\
 \left. \begin{array}{l}
 \\
 \\
 \\
 \\
 \\
 \end{array} \right\} \mathbf{G}(\mathbf{v}) \leq \mathbf{0} \\
 \left. \begin{array}{l}
 \\
 \\
 \\
 \\
 \\
 \end{array} \right\} \mathbf{v}^L \preceq \mathbf{v} \preceq \mathbf{v}^U, \quad (17)
 \end{array}$$

429 where we scaled the performance metric by dividing by its initial value (thereby giving us a base
430 performance metric value of 1, which was then to be lowered). We also note that in practice
431 measurements were collected every 3 seconds, and so the integral of the squared error was evaluated
432 discretely. The initial parameter set was chosen, somewhat arbitrarily, as $\boldsymbol{\rho}_0 := [-0.7 \ 0.5 \ 0.3 \ 0.3]^T$.

433 Prior to solving (17), we first solved an easier problem where ρ_3 and ρ_4 were fixed at their initial
434 values and only ρ_1 and ρ_2 were optimized over (these two parameters being expected to be the more
435 influential of the four):

$$\begin{array}{ll}
 \underset{\rho_1, \rho_2}{\text{minimize}} & \frac{1}{J_0} \int_0^{40} [T_{ref}(t) - T(t, \rho_1, \rho_2)]^2 dt \\
 \text{subject to} & -1 \leq \rho_1 \leq 1 \\
 & 0 \leq \rho_2 \leq 1
 \end{array}
 \left. \begin{array}{l}
 \right\} \phi_p(\mathbf{v}) \\
 \left. \begin{array}{l}
 \\
 \\
 \end{array} \right\} \mathbf{v}^L \preceq \mathbf{v} \preceq \mathbf{v}^U. \quad (18)
 \end{array}$$

Figure 5. The parameter adaptation plot (left) and the measured performance metric (right) for the solution of Problem (18). Hollow circles on the left indicate batches that were carried out as part of the initialization (prior to applying the solver). Likewise, the dotted vertical line on the right shows the iteration past which the parameter adaptations were dictated by the SCFO solver.




436 In order to approximate the non-repeatability noise term for the performance, a total of 8 batches
 437 were run with the initial parameter set ρ_0 , with the (unscaled) performance metric values obtained for
 438 those experiments being: 13.45, 13.31, 13.46, 14.25, 13.80, 13.44, 13.72, 13.98 (their mean then being
 439 taken as the scaling term J_0). Rather than attempt to run more experiments, which, though it could
 440 have improved the accuracy of our approximation, would have required even more time (each batch
 441 already requiring 40 minutes, with an additional 20-30 minutes of inter-batch preparation), we chose to
 442 approximate the statistics of the non-repeatability noise term by a zero-mean normal distribution with
 443 the standard deviation of the data, i.e., 0.32.

444 The map of the parameter adaptations and the values of the measured performance metric are given
 445 in Figure 5, with a visual comparison of the tracking before and after optimization given in Figure 6.
 446 It is seen that the majority of the improvement is obtained by about the tenth batch, with only minor
 447 improvements afterwards, and that monotonic improvement of the control performance is more or less
 448 observed.

449 We also note that the solution obtained by the solver is very much in line with what an engineer
 450 would expect for a system with slow dynamics such as this one, in that one should increase both the
 451 output weight so as to have better tracking and set the bias update filter gain close to its maximal value
 452 (both of these actions could have potentially negative effects for faster, less stable systems, however).
 453 As such, the solution is not really surprising, but it is still encouraging that a method with absolutely no
 454 knowledge embedded into it has been able to find the same in a relatively low number of experiments. It
 455 is also interesting to note that the non-repeatability noise in the measured performance metric originally
 456 puts us on the wrong track, as increasing the bias update filter gain does *not* improve the observed
 457 performance for Batch 1, though it probably should, and so the solver then spends the first 6 adaptations
 458 decreasing the bias filter gain in the belief that doing so should improve performance. However, it is
 459 able to recover by Batch 7 and to go in the right direction afterwards – this is likely due to the internal

Figure 6. The visual improvement in the temperature profile tracking from Batch 1 to Batch 20. The dotted (red) lines denote the setpoint, while the solid (black) lines denote the actual measured temperature.



460 gradient estimation algorithm of the solver having considered all of the batches and thereby decoupled
 461 the effects of the two parameters.

462 Problem (17) was then solved by similar means, though we used all of the data obtained previously
 463 to help “warm start” the solver. As the results were similar to what was obtained for the two-parameter
 464 case, we only give the measured performance metric values and the temperature profile at the final batch
 465 in Figure 7. We also note that the parameter values at the final batch were $\rho_{30} = [0.89 \ 0.95 \ 0.07 \ 0.12]^T$,
 466 from which we see that, while all four variables were clearly adapted and the solver chose to lower
 467 both the control and prediction horizons, any extra performance gains from doing this (if any) appear to
 468 have been marginal when compared to the simpler two-parameter problem. This is also in line with our
 469 intuition (i.e., that the output weight and bias filter gain are more important) and reminds us of a very
 470 important RTO concept: just because one has many variables that one *can* optimize over does not mean
 471 that one *should*, as RTO problems with more optimization variables are generally expected to converge
 472 slower and, as seen here, may not be worth the effort.

473 4.2. Periodic Setpoint Tracking in a Torsional System

474 In this study, we consider the three-disk torsional system shown in Figure 8 (the technical details of
 475 which may be found in [39]). Here, the control input is defined as the voltage of the motor located near
 476 the bottom of the system, with the control output taken as the angular position of the top disk.

477 To define an ICT problem, we generalize the idea of a “run” or a “batch” as seen in the previous
 478 example and consider instead a “window” of a periodic sinusoidal trajectory defined by:

$$y_{ref}(t) = -2 \cos \frac{\pi t}{6},$$

479 with t given in seconds. As the same trajectory is repeated every 12 seconds, we can essentially consider
 480 each 12-second window as a “run” (or a “batch”) as shown in Figure 9 and adapt the relevant controller
 481 parameters in the sampling time period between two consecutive windows.

Figure 7. The measured performance metric for the solution of Problem (17), together with the tracking obtained for the final batch.






Figure 8. The ECP 205 torsional system.



482 Not surprisingly, this presents a computational challenge, as the sampling time for this system is only
 483 60 milliseconds, which is, with the current version of the solver, insufficient – the solver needing at least
 484 a few seconds to provide a new choice of parameters. While a much simpler implementation that satisfies
 485 this real-time constraint has already been successfully carried out on the same system [24], we choose
 486 to apply the methodology presented in this paper by using a wait-and-synchronize approach. Here, the
 487 solver takes all of the available data and starts its computations, with no adaptation of the parameters
 488 being done until the solver's computations are finished. Afterwards, the solver waits until these new
 489 parameters are applied and the results for the corresponding run obtained, after which the new data is
 490 fed into the solver and the cycle restarts. The noted drawback of this approach is that we have to wait,
 491 on average, 2-3 runs (24-36 seconds) for an adaptation to take place, although the positive side of this is
 492 that the resulting data is generally less noisy due to the repeated experiments.

Figure 9. The generalization of “run-to-run” tuning to a system with a periodic setpoint trajectory. Only the setpoint is given here.



493 The controller employed is the discrete fixed-order controller given in (10), with the numerator and
 494 denominator coefficients being the (five) tuned parameters. The performance metric used is again a case
 495 of the general metric (5), but this time equal priority is given to tracking, controller aggressiveness, and
 496 the smoothness of the output trajectory, with $\lambda_1 := \lambda_3 := \lambda_4 := 1$ and $\lambda_2 := 0$.

497 As the poles of the controller are also being adapted (due to the adaptation of the denominator
 498 coefficients), controller stability constraints, as already derived in (11) and (12), are added to the ICT
 499 problem (with a tolerance of $\varepsilon := 0.01$):

$$\begin{aligned} |\rho_5| - 0.99 &\leq 0 \\ |\rho_4 - \rho_4\rho_5| - |1 - \rho_5^2| + 0.01 &\leq 0 \end{aligned} \quad ,$$

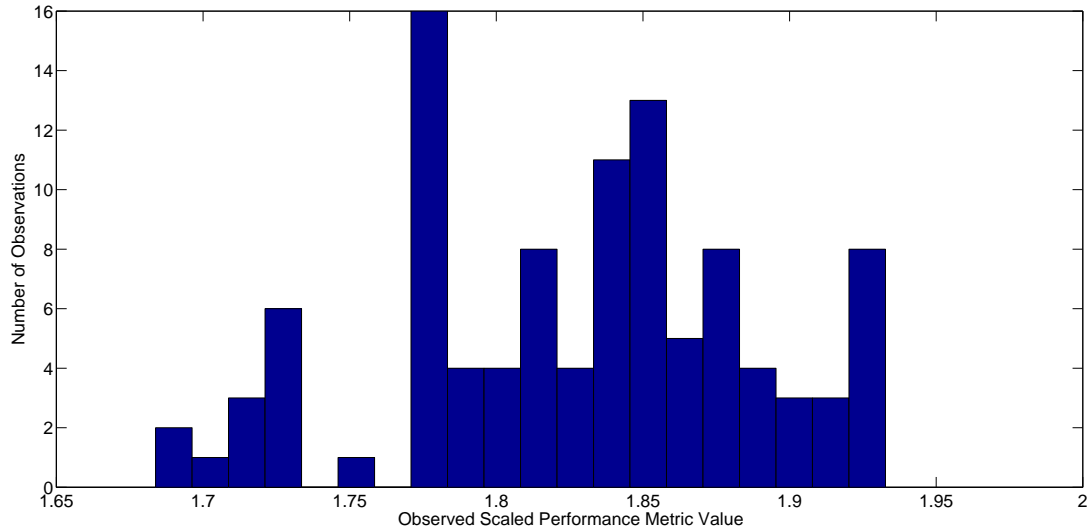
500 and are recast into differentiable form (as the solver requires \mathbf{G} to be differentiable):

$$\begin{aligned} \rho_5 - 0.99 &\leq 0 \\ -\rho_5 - 0.99 &\leq 0 \\ \rho_4 - \rho_4\rho_5 - (1 - \rho_5^2) + 0.01 &\leq 0 \quad , \\ -\rho_4 + \rho_4\rho_5 - (1 - \rho_5^2) + 0.01 &\leq 0 \end{aligned}$$

501 where we have used $|\rho_5| \leq 0.99 \Rightarrow |1 - \rho_5^2| = 1 - \rho_5^2$ in the reformulation of the second set.

502 The adaptive limits of (15) are used to constrain the individual parameters, thereby leading to the
 503 (adaptive) ICT-RTO problem:

Figure 10. A twenty-bin histogram representation of the observed scaled performance metric values for a hundred runs with the initial parameter set (Problem (19)).



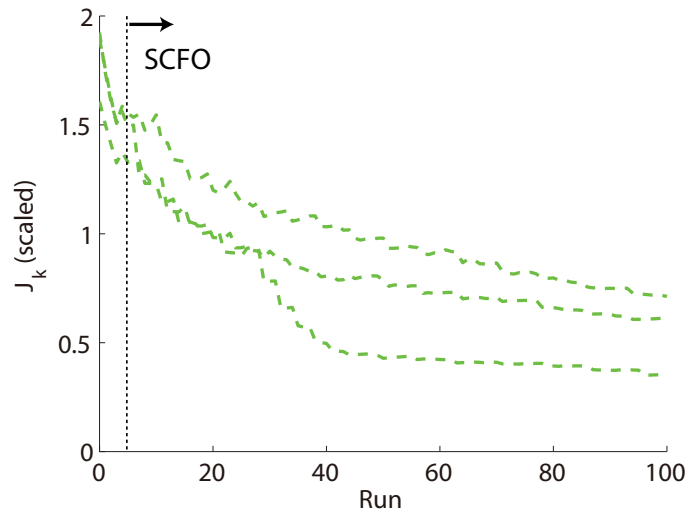
$$\begin{aligned}
 & \underset{\boldsymbol{\rho}}{\text{minimize}} && \frac{1}{100} \int_0^{12} ([y_{ref}(t) - y(t, \boldsymbol{\rho})]^2 + \dot{u}^2(t, \boldsymbol{\rho}) + \dot{y}^2(t, \boldsymbol{\rho})) dt && \left. \begin{array}{l} \\ \\ \\ \\ \\ \\ \\ \\ \\ \\ \end{array} \right\} \phi_p(\mathbf{v}) \\
 & \text{subject to} && \rho_5 - 0.99 \leq 0 && \left. \begin{array}{l} \\ \\ \\ \\ \\ \\ \\ \\ \\ \\ \end{array} \right\} \mathbf{G}(\mathbf{v}) \preceq \mathbf{0} \\
 & && -\rho_5 - 0.99 \leq 0 && \\
 & && \rho_4 - \rho_4 \rho_5 - (1 - \rho_5^2) + 0.01 \leq 0 && \\
 & && -\rho_4 + \rho_4 \rho_5 - (1 - \rho_5^2) + 0.01 \leq 0 && \\
 & && \rho_{k,1} - 0.5 \leq \rho_1 \leq \rho_{k,1} + 0.5 && \left. \begin{array}{l} \\ \\ \\ \\ \\ \\ \\ \\ \\ \\ \end{array} \right\} \mathbf{v}^L \preceq \mathbf{v} \preceq \mathbf{v}^U \\
 & && \rho_{k,2} - 0.5 \leq \rho_2 \leq \rho_{k,2} + 0.5 && \\
 & && \rho_{k,3} - 0.5 \leq \rho_3 \leq \rho_{k,3} + 0.5 && \\
 & && \rho_{k,4} - 0.5 \leq \rho_4 \leq \rho_{k,4} + 0.5 && \\
 & && \rho_{k,5} - 0.5 \leq \rho_5 \leq \rho_{k,5} + 0.5 &&
 \end{aligned} \quad , \quad (19)$$

504 where we scale the performance metric by 10^2 so as to make it vary on the magnitude of 10^0 .

505 An initial parameter set of $\boldsymbol{\rho}_0 := [1.00 \ 2.77 \ -2.60 \ 1.00 \ 0.50]^T$ was chosen and corresponds to an *ad*
 506 *hoc* initial design found by a mix of both simulation and hand tuning. To estimate the noise statistics of
 507 the non-repeatability noise term in the performance metric, the system was operated at $\boldsymbol{\rho}_0$ for 20 minutes,
 508 which produced a total of 100 performance metric measurements (see Figure 10). These were then offset
 509 by their mean to generate the estimated noise samples, with the latter being fed directly into the solver,
 510 which would then build an approximate distribution function for them.

511 Problem (19) was solved a total of three times for 20 minutes of operation (100 runs), with the
 512 performance improvements for the three trials given in Figure 11 and the visual improvement for the
 513 middle case (“middle” with regard to the final performance metric value) given in Figure 12. We note
 514 the variability in convergence behavior for the three cases (both in terms of speed and the performance
 515 achieved after 100 runs), which was largely caused by the solver converging to different minima, but

Figure 11. Performance improvement over 100 runs of operation for three different trials (dashed lines) of Problem (19).



516 note as well that all three follow the same “reliable” trend, in that performance is always improved with
 517 a fairly consistent decrease in the metric value over the course of operation.

518 4.3. PID Tuning for a Step Setpoint Change

519 We consider the problem previously examined in [13,22], where the parameters of a PID controller
 520 are to be tuned for the closed-loop system given by:

$$Y(s) = \frac{G_{ref}(s)G_p(s)}{1 + G_y(s)G_p(s)}Y_{ref}(s),$$

521 with the PID parameters K_p , τ_I , and τ_D being used to define $G_{ref}(s)$ and $G_y(s)$ as:


$$G_{ref}(s) = K_p \left(1 + \frac{1}{\tau_I s} \right)$$

$$G_y(s) = K_p \left(1 + \frac{1}{\tau_I s} + \tau_D s \right),$$

522 and $G_p(s)$ being the plant, whose definition will be varied for study purposes. The case of a setpoint
 523 step change ($Y_{ref}(s) = 1/s$) is considered, with only the tracking error to be minimized over a “masked”
 524 operating length, where a mask of t_m is applied so as not to penalize for errors on the interval $t \in [0, t_m]$,
 525 as proposed in [22].

526 Since the controller gain, K_p , is expected to vary on a magnitude of about 10^0 , it does not need
 527 scaling and so we define $\rho_1 \triangleq K_p$. For both τ_I and τ_D we assume the possibility of greater variations, on
 528 the magnitude of 10^1 (as has been suggested in both [22] and [13]), and thus define the scaled second
 529 and third parameters as $\rho_2 \triangleq \tau_I/10$ and $\rho_3 \triangleq \tau_D/10$. Since we do not know *a priori* what ρ^L and ρ^U for
 530 a PID controller should be, but do realize that both τ_I and τ_D should be positive, the adaptive definition
 531 of the lower and upper limits with the positivity constraints respected is chosen to yield the ICT problem
 532 in RTO form:

Figure 12. Difference in control input and output profiles between the first and final runs of Problem (19), with the dashed green line used to denote the input (motor voltage) values.



$$\begin{aligned}
 & \underset{\boldsymbol{\rho}}{\text{minimize}} && \frac{1}{J_0} \int_{t_m}^{t_b} [y_{ref}(t) - y(t, \boldsymbol{\rho})]^2 dt \\
 & \text{subject to} && \left. \begin{aligned}
 & \rho_{k,1} - 0.5 \leq \rho_1 \leq \rho_{k,1} + 0.5 \\
 & \max(\rho_{k,2} - 0.5, 0.01) \leq \rho_2 \leq \rho_{k,2} + 0.5 \\
 & \max(\rho_{k,3} - 0.5, 0.01) \leq \rho_3 \leq \rho_{k,3} + 0.5
 \end{aligned} \right\} \begin{aligned}
 & \phi_p(\mathbf{v}) \\
 & \mathbf{v}^L \preceq \mathbf{v} \preceq \mathbf{v}^U,
 \end{aligned} \quad (20)
 \end{aligned}$$

533 where we scale the cost function by dividing by the performance metric value for the original parameter
 534 set.

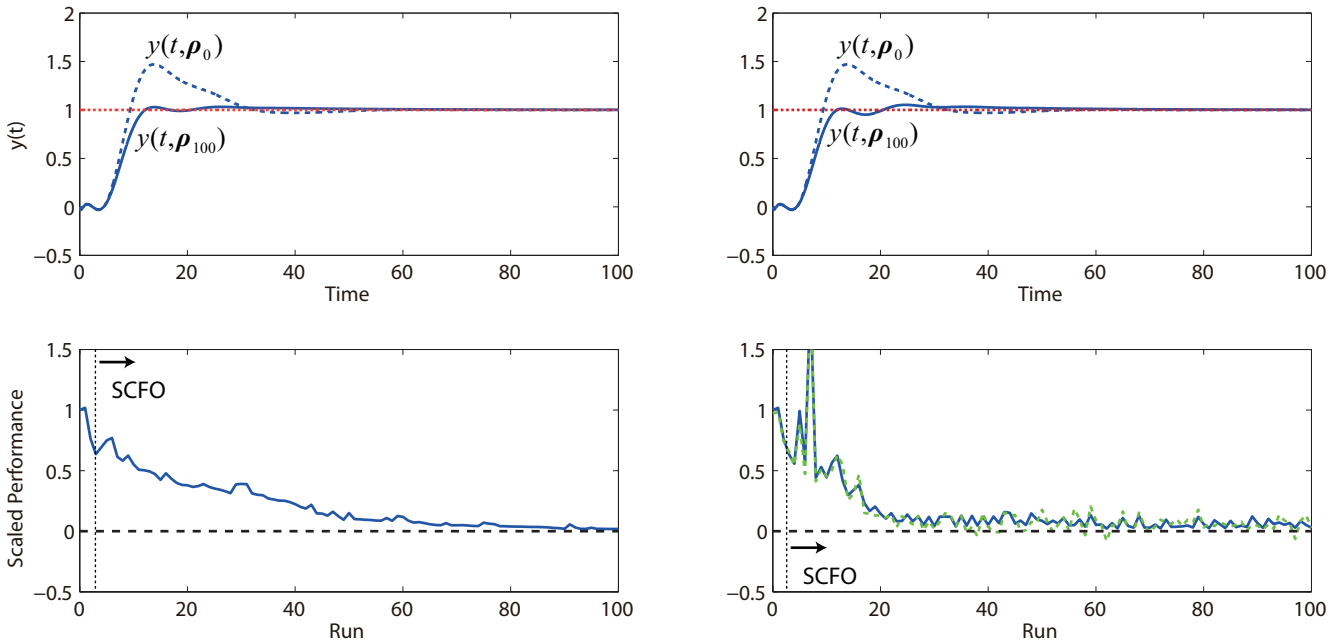
535 As done in [13], the original parameter set is chosen as the set found by Ziegler-Nichols tuning. The
 536 following three studies are considered here:

$$\begin{aligned}
 \text{Study 1 : } & G_p(s) = \frac{1}{1 + 20s} e^{-5s}, \quad t_m := 10, \quad t_b = 100, \quad \boldsymbol{\rho}_0 := [4.06 \ 0.93 \ 0.23]^T \\
 \text{Study 2 : } & G_p(s) = \frac{1}{1 + 20s} e^{-20s}, \quad t_m := 50, \quad t_b = 300, \quad \boldsymbol{\rho}_0 := [1.33 \ 3.10 \ 0.65]^T \\
 \text{Study 3 : } & G_p(s) = \frac{1}{(1 + 10s)^8}, \quad t_m := 140, \quad t_b = 500, \quad \boldsymbol{\rho}_0 := [1.10 \ 7.59 \ 1.90]^T
 \end{aligned}$$

537 So as to study the effect of non-repeatability noise, each observed performance metric value is
 538 corrupted with an additive error from $\mathcal{N}(0, (0.05J_0)^2)$, i.e., by an additive error with a standard deviation
 539 that is chosen as 5% of the original performance metric value (assumed known for solver configuration).
 540 Noiseless scenarios were simulated as well.

541 The results for the three studies are provided in Figures 13-15. On the whole, we see that the solver
 542 reliably optimizes control performance in both the noiseless and noisy scenarios, even though we note
 543 that the rate of improvement can vary from problem to problem. For the noisy cases, we generally
 544 see more “bumps” in the convergence trajectory, which should not be surprising given (a) the added
 545 difficulty for the solver in estimating local derivatives and (b) the reduced conservatism in the estimation
 546 of the quadratic bound constants M , for which the safety factor η in (16) is generally augmented less
 547 frequently when noise is present. However, for the latter point, we see that there is an upside with regard

Figure 13. Performance obtained by iterative tuning for both the noiseless (left) and noisy (right) cases of Study 1 of Problem (20), with the solid blue line used to denote the “true” performance of the closed-loop system and the green dashed line used to denote what is actually observed (and provided to the solver). In both cases, the SCFO solver brings the closed-loop performance metric value close to its global minimum of 0 (marked by the black dashed line in the lower plots).



548 to convergence speed. Because the values of M tend to be less conservative in the presence of noise, the
 549 algorithm tends to take larger steps and progresses quicker towards the optimum, as is witnessed in both
 550 Figures 13 and 15. We do note the occasional danger of performance worsening due to tuning, but this
 551 is almost always restricted to the earlier runs when the solver is relatively “data-starved”.

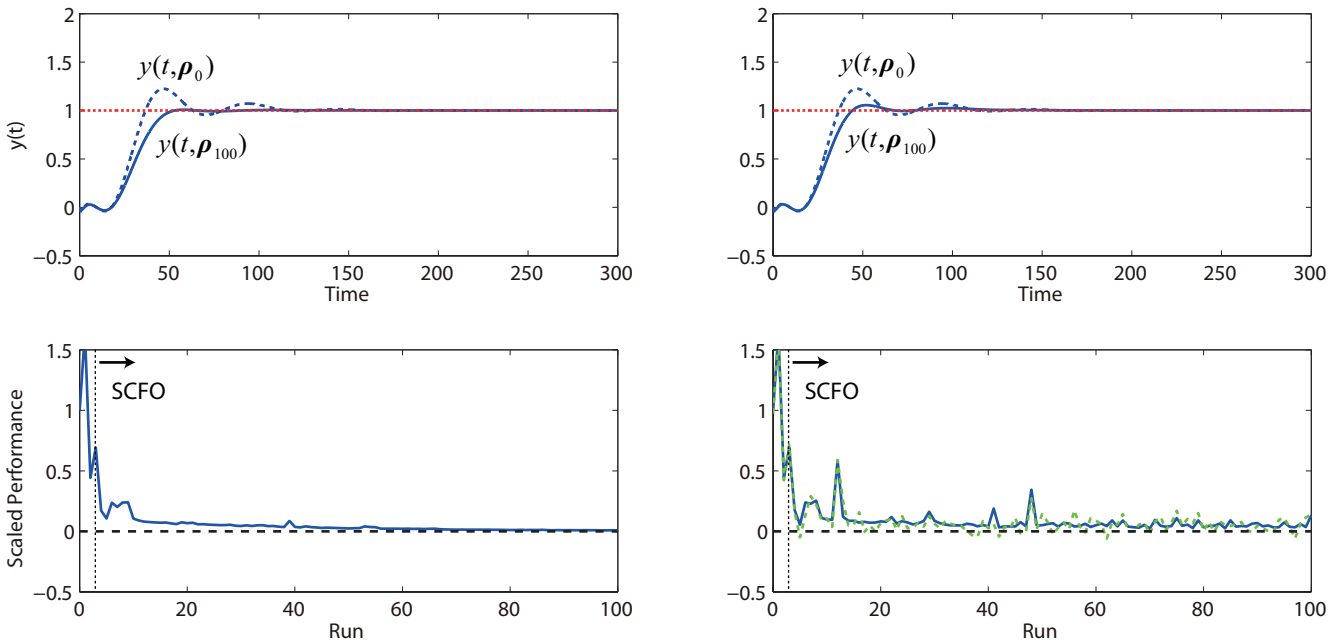
552 4.4. Tuning a System of PI Controllers for Setpoint Tracking and Disturbance Rejection

553 Here, we consider the following 5-input, 5-output dynamical system:

$$\begin{aligned}
 \ddot{y}_1(t) + \dot{y}_1(t) + y_1(t) &= u_1(t) - 0.033u_3(t) - 0.067u_4(t) - 0.1u_5(t) \\
 \ddot{y}_2(t) + 0.1\dot{y}_2(t) + y_2(t) &= 0.1u_1(t) + 2u_2(t) + 0.033u_3(t) - 0.033u_5(t) \\
 \ddot{y}_3(t) + 5\dot{y}_3(t) + y_3(t) &= 0.167u_1(t) + 0.133u_2(t) + 3u_3(t) + 0.067u_4(t) + 0.033u_5(t) \quad (21) \\
 \ddot{y}_4(t) + 2\dot{y}_4(t) + y_4(t) &= 0.233u_1(t) + 0.2u_2(t) + 0.167u_3(t) + 4u_4(t) + 0.1u_5(t) \\
 \ddot{y}_5(t) + 3\dot{y}_5(t) + y_5(t) &= 0.3u_1(t) + 0.267u_2(t) + 0.233u_3(t) + 0.2u_4(t) + 5u_5(t)
 \end{aligned}$$

554 While the user cannot be assumed to know the plant (21), we will assume that they have been able to
 555 properly decouple the system with the input-output pairings of $u_i \rightarrow y_i$, $i = 1, \dots, 5$ (as this is evidently
 556 the superior choice if one considers the relative gains). A system of five PI controllers is used for the
 557 pairings:

$$u_i(t) = K_{p,i} \left(e_i(t) + \frac{1}{\tau_{I,i}} \int_0^t e_i(t) dt \right), \quad i = 1, \dots, 5,$$

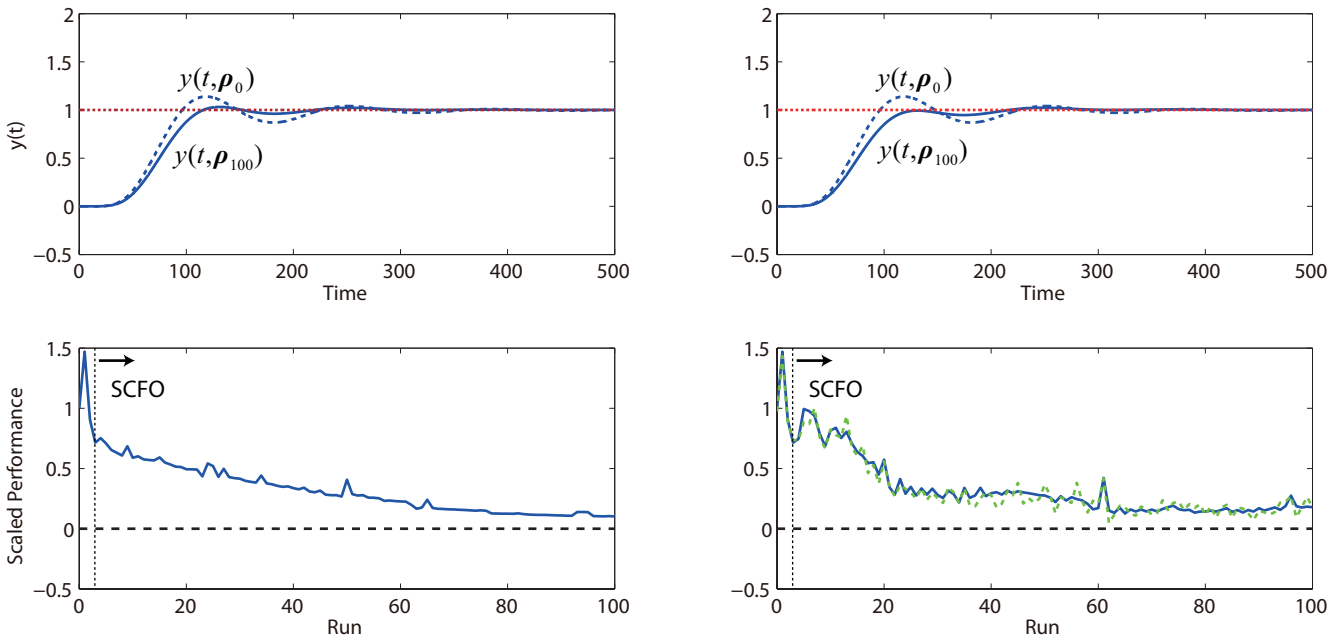
Figure 14. Performance obtained by iterative tuning for Study 2 of Problem (20).

558 which, of course, is not perfect since the decoupling is not either, and so what one controller does will
 559 inevitably affect the others.

560 The ICT problem that we define for this system consists of starting with all $y_i(0) = 0$ and defining
 561 the setpoints of y_1 , y_3 , and y_5 as 1 (which makes this a tracking problem with respect to these outputs)
 562 and the setpoints of y_2 and y_4 as 0 (which makes it a disturbance rejection problem with respect to these
 563 two outputs). The total sum of squared tracking errors for all of the outputs is used as the performance
 564 metric, with the interval of $t \in [2, 15]$ being considered in the metric computation (a “mask” of 2 time
 565 units being employed).

566 The first five tuning parameters are simply defined as the controller gains, with $\rho_i \triangleq K_{p,i}$, $i = 1, \dots, 5$.
 567 As in the previous example, we use a scaled version of the integral times to define the rest, with
 568 $\rho_{i+5} \triangleq \tau_{I,i}/10$, $i = 1, \dots, 5$. Once again, as we do not know *a priori* what lower and upper limits should
 569 be set on these parameters (save the positivity of the $\tau_{I,i}$), adaptive inputs with the positivity limitation
 570 (as shown in the previous case study) are used.

571 Furthermore, we suppose the existence of a safety limitation in the form of a maximal value that y_1 is
 572 allowed to take, with the constraint $y_1(t) \leq 1.2$ to be met at all times. Using the reformulation shown in
 573 Section 2.2, we may proceed to state this problem in RTO form as:

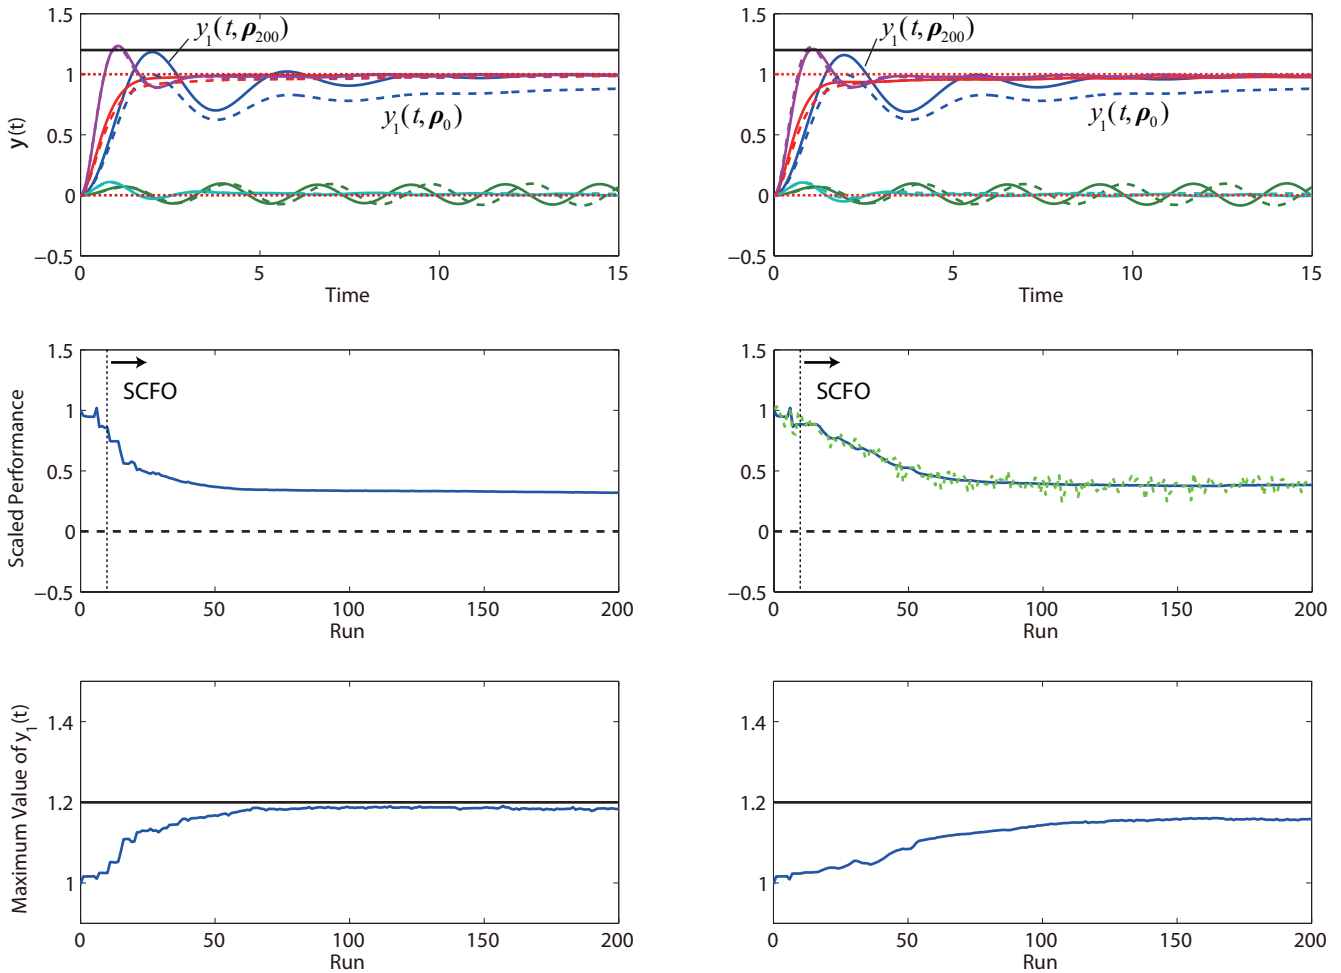
Figure 15. Performance obtained by iterative tuning for Study 3 of Problem (20).

$$\begin{aligned}
 & \underset{\boldsymbol{\rho}}{\text{minimize}} && \frac{1}{J_0} \sum_{i=1}^5 \int_2^{15} [y_{i,ref}(t) - y_i(t, \boldsymbol{\rho})]^2 dt \\
 & \text{subject to} && y_{1,max}(\boldsymbol{\rho}) - 1.2 \leq 0 \\
 & && \rho_{k,1} - 0.5 \leq \rho_1 \leq \rho_{k,1} + 0.5 \\
 & && \rho_{k,2} - 0.5 \leq \rho_2 \leq \rho_{k,2} + 0.5 \\
 & && \rho_{k,3} - 0.5 \leq \rho_3 \leq \rho_{k,3} + 0.5 \\
 & && \rho_{k,4} - 0.5 \leq \rho_4 \leq \rho_{k,4} + 0.5 \\
 & && \rho_{k,5} - 0.5 \leq \rho_5 \leq \rho_{k,5} + 0.5 \\
 & && \max(\rho_{k,6} - 0.5, 0.01) \leq \rho_6 \leq \rho_{k,6} + 0.5 \\
 & && \max(\rho_{k,7} - 0.5, 0.01) \leq \rho_7 \leq \rho_{k,7} + 0.5 \\
 & && \max(\rho_{k,8} - 0.5, 0.01) \leq \rho_8 \leq \rho_{k,8} + 0.5 \\
 & && \max(\rho_{k,9} - 0.5, 0.01) \leq \rho_9 \leq \rho_{k,9} + 0.5 \\
 & && \max(\rho_{k,10} - 0.5, 0.01) \leq \rho_{10} \leq \rho_{k,10} + 0.5
 \end{aligned}
 \left. \begin{array}{l} \phi_p(\mathbf{v}) \\ \mathbf{G}_p(\mathbf{v}) \leq \mathbf{0} \\ \mathbf{v}^L \preceq \mathbf{v} \preceq \mathbf{v}^U \end{array} \right\} \quad (22)$$

574 We note that this problem is a bit more challenging than the ones considered in the previous three
 575 studies due to the increased number of tuning parameters, and point out that, were the problem perfectly
 576 decoupled, we would be able to solve it as five 2-parameter RTO problems in parallel. However, seeing
 577 as all of the parameters are intertwined, we have no choice but to optimize over all ten simultaneously
 578 – the expected price to pay being a slower rate of performance improvement obtained by the solver.
 579 Alternate strategies that are based on any additional engineering knowledge, such as optimizing only the
 580 parameters of specific controllers or optimizing only the controller gains, could of course be proposed
 581 and are highly recommended.

582 As a somewhat arbitrary design, the initial set is chosen as $\boldsymbol{\rho}_0 := [2 \ 2 \ 2 \ 2 \ 2 \ 1 \ 1 \ 1 \ 1 \ 1]^T$. Like with the
 583 previous example, an additive measurement noise of $\mathcal{N}(0, (0.05J_0)^2)$ is added to corrupt the performance
 584 metric value that is observed for a given choice of tuning parameters. An additive measurement noise

Figure 16. Performance obtained by iterative tuning for the system of PI controllers in Problem (22) – the noiseless case is given on the left and the noisy case on the right. For the output profiles, we note that the initial profiles are given as dashed lines, with the final profiles given by solid lines of the same color.

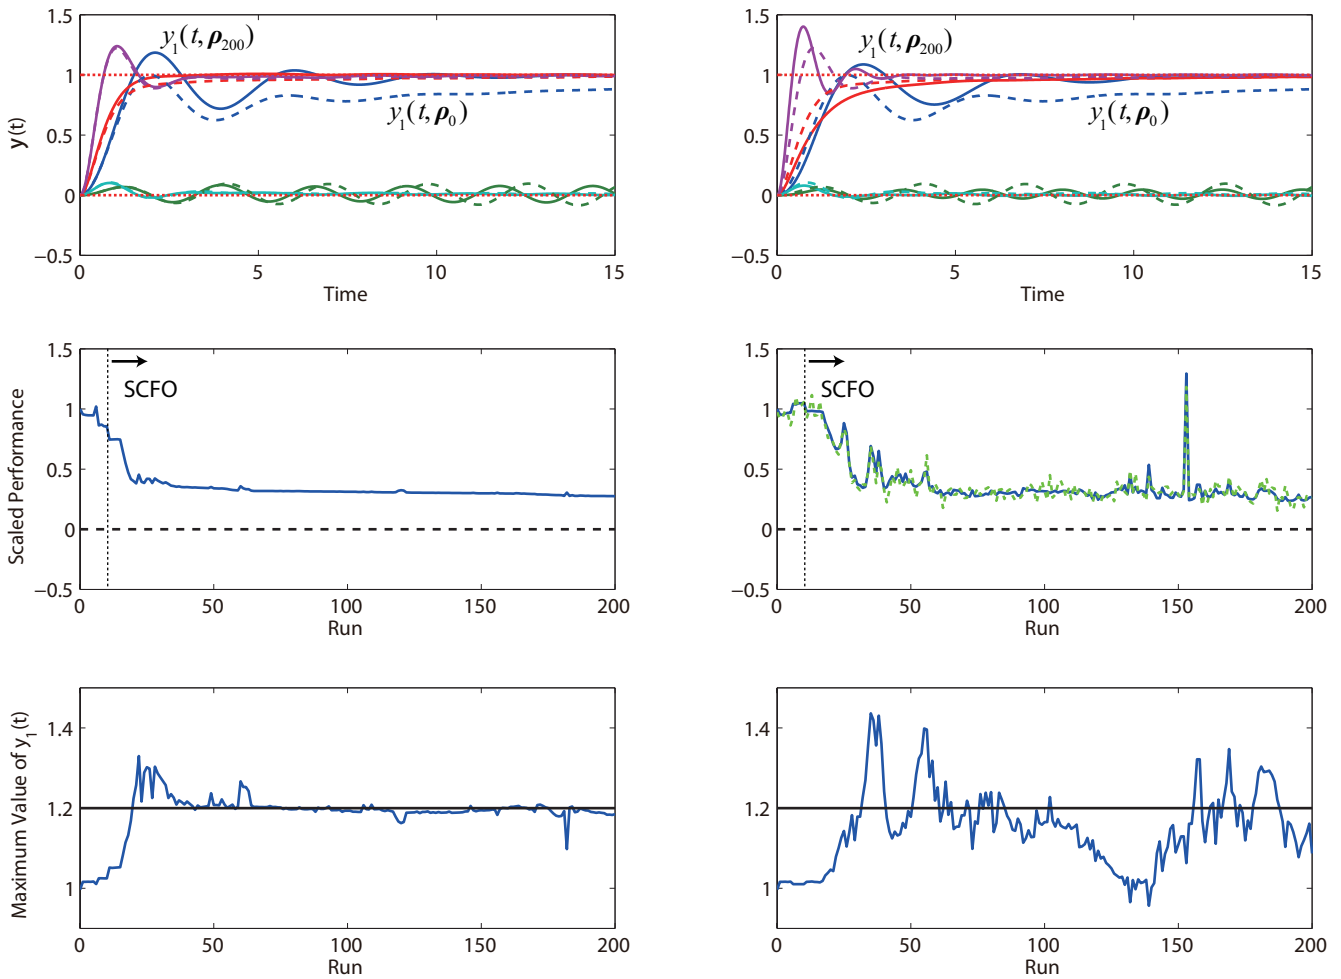


585 of $\mathcal{N}(0, 10^{-4})$ is added to corrupt the observed values of $y_{1,max}$. Both sets of statistics are assumed to
 586 be known for the purposes of SCFO solver configuration. As before, the noiseless scenarios are also
 587 considered.

588 We present the results in Figure 16, which show that the solver is able to obtain significant
 589 performance improvements within 50 iterations for both the noiseless and noisy cases without once
 590 violating the output constraint on y_1 . In this case, we see that the noise has the effect of slowing down
 591 convergence, which may be explained by the fact that the solver must take even more cautious steps so as
 592 not to violate the output constraint. Additionally, the performance that is observed after 200 iterations is
 593 a bit worse for the noisy case, which may be seen as being due to the back-off from the output constraint
 594 being larger (to account for the noise).

595 To test the effect of this constraint and to see if it is even necessary, we also run a simulation where the
 596 constraint is lifted from the problem statement. The results for this study are given in Figure 17, and show
 597 that not having the constraint in place certainly leads to runs where it is violated. This is not surprising,
 598 given that a lot of the performance improvement is obtained by tracking the setpoint of y_1 faster, which

Figure 17. Performance obtained by iterative tuning for the system of PI controllers in Problem (22) (without an output constraint).



599 is easier to do once there is no constraint on its overshoot. It is also seen that the performance obtained
 600 after 200 iterations is generally better than what would be obtained with the constraint – this is, again,
 601 not surprising, as removing a limiting constraint should allow for greater performance gains. We do
 602 note that the noisy case is more bumpy without the constraint, which is expected as there is less to limit
 603 the adaptation steps and more “daring” adaptations become possible. While some of the bumps may be
 604 quite undesired (particularly, the one noted just after the 150th run), the algorithm remains, on the whole,
 605 reliable as it keeps the performance metric at low values for the majority of the runs despite significant
 606 noise corruption.

607 5. Concluding Remarks

608 The goal of this paper has been to propose the idea of posing the iterative controller tuning (ICT)
 609 problem in the real-time optimization (RTO) framework, and it has been shown how one can easily
 610 formulate most ICT problems as RTO ones with the use of a repeatability assumption that, though it is
 611 only an approximation of reality in the presence of noise/degradation, appears to suffice for application
 612 purposes (at least, in the two experimental case studies considered here). A major advantage of this
 613 reformulation is that a number of previously unaddressed challenges in ICT, the majority of which

614 take the form of constraints in the performance metric minimization problem, may be addressed in a
615 fairly straightforward manner. To make the message more concrete, we have also shown how the ICT
616 problem may be solved by the SCFO real-time optimization solver, and have provided the reader with
617 the necessary solver settings to do so. Four case studies have shown the method to work very well for a
618 diverse range of problems.

619 Though we hope to have convinced the reader that the method proposed makes for a strong candidate
620 for solving general ICT problems in practice, its potential drawbacks should be clear:

- 621 • No solution has been proposed for how to treat the case where the repeatability assumption is not
622 a good approximation of reality. Instead of hoping that the approximation suffice in practice, it
623 would be beneficial to propose alternatives that would still allow one to use the RTO framework to
624 deal with the problem. In particular, one could attempt to make the repeatability assumption on the
625 input and output trajectories rather than making it directly on the performance metric. This could
626 allow one to establish a closer link between the lack of repeatability and the input/output noise in
627 the control system.
- 628 • Although the proposed configuration has been shown to be largely successful here, many of the
629 elements involved still remain heuristic in nature. Either improving on these heuristics or finding
630 ways to avoid them are desired.
- 631 • The method is currently limited to solving ICT problems where the control task remains the
632 same, which may significantly limit its domain of applicability. It would be interesting to
633 attempt to extend it to cases where the control tasks were similar, rather than identical, and then
634 somehow penalize the method based on the degree of similarity (e.g., one could attempt to lump
635 non-similarity into the noise element δ of the repeatability assumption).

636 We finish by noting that an abstract advantage of the RTO-ICT formulation is that we are now able
637 to attack the ICT problem from two directions – that of control and that of RTO. For the former, we
638 note that the proposed method applies very few control principles (unlike other direct tuning methods
639 [9,10], which make heavy use of control theory). While this is, in some sense, an advantage – as it
640 allows us to use the proposed method to tune almost any controller for almost any system – there
641 is undoubtedly something lost due to the “black-box veil” that the RTO formulation places on the
642 problem, and incorporating additional knowledge for specific controllers would very likely allow for
643 further improvements to the techniques discussed here. At the same time, the RTO methods themselves
644 are in a fairly nascent stage theoretically. Many improvements to both RTO theory and solution methods
645 are expected to appear in the coming years, which could only improve on the results presented here and
646 to make the solution of the ICT problem both faster and more reliable.

647 **Acknowledgements**

648 The authors would like to acknowledge Fernando Fraire Tirado, for having contributed to the birth
649 and development of the idea of using RTO approaches to solve the iterative controller tuning problem
650 via a number of student projects, the results of which ultimately led to this publication. We would also
651 like to thank Prof. Christos Georgakis, Prof. Kyongbum Lee, and Emily Edwards, as well as others at

652 the Chemical and Biological Engineering Department of Tufts University for allowing the first author to
653 spend two weeks there and to use the laboratory equipment to obtain the results for the first case study.
654 Finally, we want to thank Dr. Timm Faulwasser for several useful discussions.

655 References

- 656 [1] Ziegler, J.G.; Nichols, N.B. Optimum settings for automatic controllers. *Trans. ASME* **1942**,
657 *64*, 759–765.
- 658 [2] Ogunnaike, B.A.; Ray, W.H. *Process Dynamics, Modeling, and Control*; Oxford University
659 Press, 1994; pp. 648–665.
- 660 [3] Campi, M.C.; Lecchini, A.; Savaresi, S.M. Virtual reference feedback tuning: A direct method
661 for the design of feedback controllers. *Automatica* **2002**, *38*, 1337–1346.
- 662 [4] Garriga, J.; Soroush, M. Model predictive control tuning methods: A review. *Ind. Eng. Chem.*
663 *Res.* **2010**, *49*, 3505–3515.
- 664 [5] Burke, J.V.; Henrion, D.; Lewis, A.S.; Overton, M.L. HIFOO - a MATLAB package for fixed-
665 order controller design and H_∞ optimization. Fifth IFAC Symposium on Robust Control Design
666 (Toulouse), 2006.
- 667 [6] Khatibi, H.; Karimi, A.; Longchamp, R. Fixed-order controller design for systems with polytopic
668 uncertainty using LMIs. *IEEE Trans. Automatic Control* **2008**, *53*, 428–434.
- 669 [7] Landau, I.D.; Lozano, R.; M'Saad, M.; Karimi, A. *Adaptive Control*; Springer, 2011; pp. 11–18.
- 670 [8] Hjalmarsson, H.; Gunnarsson, S.; Gevers, M. Optimality and sub-optimality of iterative
671 identification and control design schemes. Proceedings of the American Control Conference
672 (Seattle), 1995, pp. 2559–2563.
- 673 [9] Hjalmarsson, H.; Gevers, M.; Gunnarsson, S.; Lequin, O. Iterative feedback tuning: theory and
674 applications. *Control Systems, IEEE* **1998**, *18*, 26–41.
- 675 [10] Kammer, L.C.; Bitmead, R.R.; Bartlett, P.L. Direct iterative tuning via spectral analysis.
676 *Automatica* **2000**, *36*, 1301–1307.
- 677 [11] Hjalmarsson, H. Iterative feedback tuning - An overview. *Int. J. Adapt. Control Signal Process.*
678 **2002**, *16*, 373–395.
- 679 [12] Karimi, A.; Mišković, L.; Bonvin, D. Iterative correlation-based controller tuning. *Int. J. Adapt.*
680 *Control Signal Process.* **2004**, *18*, 645–664.
- 681 [13] Killingsworth, N.J.; Krstić, M. PID tuning using extremum seeking: online, model-free
682 performance optimization. *Control Systems, IEEE* **2006**, *26*, 70–79.
- 683 [14] Wang, Y.; Gao, F.; Doyle III, F.J. Survey on iterative learning control, repetitive control, and
684 run-to-run control. *J. Process Control* **2009**, *19*, 1589–1600.
- 685 [15] Cutler, C.R.; Perry, R.T. Real time optimization with multivariable control is required to
686 maximize profits. *Comput. Chem. Eng.* **1983**, *7*, 663–667.
- 687 [16] Zhang, Y.; Monder, D.; Forbes, J.F. Real-time optimization under parametric uncertainty: A
688 probability constrained approach. *J. Process Control* **2002**, *12*, 373–389.
- 689 [17] Diehl, M.; Bock, H.G.; Schlöder, J.P.; Findeisen, R.; Nagy, Z.; Allgöwer, F. Real-time opti-
690 mization and nonlinear model predictive control of processes governed by differential-algebraic
691 equations. *J. Process Control* **2002**, *12*, 577–585.

- 692 [18] Mattingley, J.; Boyd, S. Real-time convex optimization in signal processing. *IEEE Signal Proc.*
693 *Mag.* **2010**, *27*, 50–61.
- 694 [19] Bunin, G.A.; François, G.; Bonvin, D. Sufficient conditions for feasibility and optimality of
695 real-time optimization schemes - I. Theoretical foundations. *arXiv:1308.2620v1 [math.OC]*
696 **2013**.
- 697 [20] Bunin, G.A.; François, G.; Bonvin, D. Sufficient conditions for feasibility and optimality of
698 real-time optimization schemes - II. Implementation issues. *arXiv:1308.2625v1 [math.OC]*
699 **2013**.
- 700 [21] Bunin, G.A.; François, G.; Bonvin, D. *The SCFO Real-Time Optimization Solver:*
701 *Users' Guide (version 0.9)*. Ecole Polytechnique Fédérale de Lausanne, 2013.
702 <http://infoscience.epfl.ch/record/186672>.
- 703 [22] Lequin, O.; Bosmans, E.; Triest, T. Iterative feedback tuning of PID parameters: Comparison
704 with classical tuning rules. *Contr. Eng. Pract.* **2003**, *11*, 1023–1033.
- 705 [23] Bunin, G.A.; Fraire, F.; François, G.; Bonvin, D. Run-to-run MPC tuning via gradient descent.
706 *Comput. Aided Chem. Eng.* **2012**, *30*, 927–931.
- 707 [24] Bunin, G.A.; François, G.; Bonvin, D. Iterative controller tuning by real-time optimization.
708 *Dynamics and Control of Process Systems (DYCOPS)* (Mumbai), 2013.
- 709 [25] Maciejowski, J.M. *Predictive Control : With Constraints*; Pearson, 2002; pp. 1–32.
- 710 [26] Gopal, M. *Digital Control Engineering*; New Age International, 1988; pp. 79–81.
- 711 [27] Åkerblad, M.; Hansson, A.; Wahlberg, B. Automatic tuning for classical step-response
712 specifications using iterative feedback tuning. *Proceedings of the 39th IEEE Conference on*
713 *Decision and Control (Sydney)*, 2000, Vol. 4, pp. 3347–3348.
- 714 [28] Kantas, N.; Maciejowski, J.M.; Lecchini-Visintini, A., *Nonlinear Model Predictive Control*;
715 Springer-Verlag Berlin Heidelberg, 2009; chapter Sequential Monte Carlo for Model Predictive
716 Control, pp. 263–273.
- 717 [29] Jang, S.; Joseph, B.; Mukai, H. On-line optimization of constrained multivariable chemical
718 processes. *AIChE J.* **1987**, *33*, 26–35.
- 719 [30] Brdys, M.; Tatjewski, P. *Iterative Algorithms for Multilayer Optimizing Control*; Imperial College
720 Press, 2005.
- 721 [31] Gao, W.; Engell, S. Iterative set-point optimization of batch chromatography. *Comput. Chem.*
722 *Eng.* **2005**, *29*, 1401–1409.
- 723 [32] Marchetti, A.; Chachuat, B.; Bonvin, D. Modifier-adaptation methodology for real-time
724 optimization. *Ind. Eng. Chem. Res.* **2009**, *48*, 6022–6033.
- 725 [33] Box, G.; Draper, N. *Evolutionary Operation: A Statistical Method for Process Improvement*;
726 John Wiley & Sons, 1969.
- 727 [34] Conn, A.; Scheinberg, K.; Vicente, L. *Introduction to Derivative-Free Optimization*; Cambridge
728 University Press, 2009.
- 729 [35] Alexandrov, N.; Dennis, J.; Lewis, R.; Torczon, V. A trust region framework for managing the
730 use of approximation models in optimization. Technical report, Langley Research Center, 1997.
- 731 [36] Myers, R.; Montgomery, D.; Anderson-Cook, C. *Response Surface Methodology*; John Wiley &
732 Sons, 2009.

- 733 [37] Marchetti, A.; Chachuat, B.; Bonvin, D. A dual modifier-adaptation approach for real-time
734 optimization. *J. Process Control* **2010**, *20*, 1027–1037.
- 735 [38] Bunin, G.A.; Lima, F.V.; Georgakis, C.; Hunt, C.M. Model predictive control and dynamic
736 operability studies in a stirred tank: Rapid temperature cycling for crystallization. *Chem. Eng.*
737 *Commun.* **2010**, *197*, 733–752.
- 738 [39] Educational Control Products. *Manual for Model 205/205a: Torsional Control System*, 2008.

739 Appendix

740 Description of the Initialization Scheme

741 The algorithm used to initialize the SCFO solver is as follows:

- 742 1. Initialize $\mathbf{P} \in \mathbb{R}^{n_v \times n_v}$ as a diagonal matrix with $P_{11} := 1$ and all other elements set to 0. Set
743 $k := 1$. Define by $\Delta \mathbf{v}_{pert} \in \mathbb{R}_{++}^{n_v}$ the perturbation vector and set $\Delta \mathbf{v}_{pert} := \Delta \mathbf{v}_{max}$.
- 744 2. Define $\mathbf{v}_k := \mathbf{v}_0 + \mathbf{P} \Delta \mathbf{v}_{pert}$ and compute the following matrix:

$$\Delta \mathbf{V} := \begin{bmatrix} (\mathbf{v}_0 - \mathbf{v}_1)^T \\ (\mathbf{v}_1 - \mathbf{v}_2)^T \\ \vdots \\ (\mathbf{v}_{k-1} - \mathbf{v}_k)^T \end{bmatrix}.$$

745 If the condition number of $\Delta \mathbf{V}$ is greater than 50, re-define \mathbf{v}_k as $\mathbf{v}_k := \mathbf{v}_{k-1} + \mathbf{R}_k \Delta \mathbf{v}_{pert}$, where
746 \mathbf{R}_k is a diagonal matrix of zeros with the sole k^{th} diagonal element equal to 1.

- 747 3. Obtain the corresponding $\hat{\phi}_p(\mathbf{v}_k) := J_k$ by running a closed-loop experiment with the controller
748 parameters $\boldsymbol{\rho}_k := \mathbf{v}_k$. Define:

$$\Delta \Phi := \begin{bmatrix} \hat{\phi}_p(\mathbf{v}_0) - \hat{\phi}_p(\mathbf{v}_1) \\ \hat{\phi}_p(\mathbf{v}_1) - \hat{\phi}_p(\mathbf{v}_2) \\ \vdots \\ \hat{\phi}_p(\mathbf{v}_{k-1}) - \hat{\phi}_p(\mathbf{v}_k) \end{bmatrix},$$

749 and compute:

$$\nabla \hat{\phi}_p := (\Delta \mathbf{V})^\dagger \Delta \Phi, \quad (23)$$

750 with \dagger denoting the Moore-Penrose pseudoinverse.

- 751 4. Re-define \mathbf{P} as a diagonal matrix with the diagonal elements set as:

$$P_{ii} := \begin{cases} 1, & \nabla \hat{\phi}_{p,i} \leq 0 \text{ and } i \leq k \\ -1, & \nabla \hat{\phi}_{p,i} > 0 \text{ and } i \leq k \\ 1, & i = k + 1 \\ 0, & i > k + 1 \end{cases},$$

752 where $\nabla \hat{\phi}_{p,i}$ denotes the i^{th} element of $\nabla \hat{\phi}_p$.

753 5. Set $k := k + 1$. If $k > n_v$, terminate. Otherwise, return to Step 2.

754 We make the following remarks:

- 755 • This scheme starts like the simple perturbation scheme where only one parameter is perturbed
756 at a time (only ρ_1 is perturbed for the first experiment), but adapts based on the results of the
757 perturbation. For example, if we see that setting $\rho_{1,1} := \rho_{0,1} + \Delta v_{pert,1}$ improves performance, then
758 we will maintain this perturbation while additionally perturbing ρ_2 in the following experiment. On
759 the other hand, if we see that this perturbation leads to worse control performance, then we simply
760 negate it for the following experiment, with this experiment being defined by the perturbations
761 $\rho_{2,1} := \rho_{0,1} - \Delta v_{pert,1}$ and $\rho_{2,2} := \rho_{0,2} + \Delta v_{pert,2}$. The (partial) linear estimate (23) of the gradient
762 acts as a guide for which directions to perturb in.
- 763 • Due to the pseudo-inversion of $\Delta \mathbf{V}$, it follows that we also require an additional safeguard to
764 ensure that the matrix remains well-conditioned, as not doing this could lead to a poor estimate
765 of the gradient (assuming the inputs \mathbf{v} to be well-scaled, which we do). Since the perturbation
766 scheme alone does not ensure this, an override is introduced where only a single input is perturbed
767 once the condition number goes over a certain threshold (chosen here as 50). This essentially
768 ensures that the conditioning does not get any worse as it forces $\Delta \mathbf{V}$ to be block diagonal.
- 769 • The choice of $\Delta \mathbf{v}_{pert} := \Delta \mathbf{v}_{max}$ is only a recommendation, as the recommended definition for
770 $\Delta \mathbf{v}_{max}$ as given in Table 1 (i.e., $0.1(\boldsymbol{\rho}^U - \boldsymbol{\rho}^L)$) tends to provide sufficient excitation without
771 perturbing “too far”. However, if there is a fear that applying perturbations of this size will
772 violate some of the problem constraints or destabilize the system, then $\Delta \mathbf{v}_{pert}$ should be reduced
773 accordingly.

774 *Data-Driven Estimations of the Performance Gradient and Hessian*

775 Estimates of the gradient and Hessian are obtained via response-surface modeling as follows:

- 776 • If $k < 2n_v + 1$, fit a linear model to all of the available data:

$$\phi_p(\mathbf{v}) \approx a_0 + \sum_{i=1}^{n_v} a_i v_i,$$

777 and define:

$$\left. \frac{\partial \hat{\phi}_p}{\partial v_i} \right|_{\mathbf{v}_k} := a_i, \quad H_{k,ij} := \begin{cases} 0.5\kappa_{\phi,i}, & i = j \\ 0, & i \neq j \end{cases},$$

778 i.e., the gradient is estimated as the coefficients of the linear model and the Hessian, in the absence
779 of more measurements, is defined as a diagonal matrix whose diagonals are equal to half of the
780 Lipschitz constants of the cost (we note that $\kappa_{\phi,i} = \bar{\kappa}_{\phi,i} = -\underline{\kappa}_{\phi,i}$ here – see Section 3.5 for how
781 these are chosen). The latter choice is justified as it (a) does not affect the relative scaling of the
782 different RTO input directions (the Lipschitz constants being equal for all inputs in this case –
783 see Section 3.5), and (b) yields a fairly small step size due to the expected conservatism of $\kappa_{\phi,i}$

(which may be desired since $\nabla \hat{\phi}_p(\mathbf{v}_k)$ is unlikely to be small for earlier runs). In the case where the data are not well-posed for linear regression and the coefficients of the linear model are poorly estimated, the following control step is applied to trim potentially bad estimates:

$$\begin{aligned} \frac{\partial \hat{\phi}_p}{\partial v_i} \Big|_{\mathbf{v}_k} > \kappa_{\phi,i} &\rightarrow \frac{\partial \hat{\phi}_p}{\partial v_i} \Big|_{\mathbf{v}_k} := \kappa_{\phi,i} \\ \frac{\partial \hat{\phi}_p}{\partial v_i} \Big|_{\mathbf{v}_k} < -\kappa_{\phi,i} &\rightarrow \frac{\partial \hat{\phi}_p}{\partial v_i} \Big|_{\mathbf{v}_k} := -\kappa_{\phi,i} \end{aligned} \quad (24)$$

- If $2n_v + 1 \leq k < 2n_v + 1 + \sum_{i=1}^{n_v-1} i$, fit a diagonal quadratic model to the data (quadratic without interaction terms):

$$\phi_p(\mathbf{v}) \approx a_0 + \sum_{i=1}^{n_v} a_i v_i + \sum_{i=1}^{n_v} a_{ii} v_i^2,$$

and define:

$$\frac{\partial \hat{\phi}_p}{\partial v_i} \Big|_{\mathbf{v}_k} := a_i + 2a_{ii} v_{k,i}, \quad H_{k,ij} := \begin{cases} 2a_{ii}, & i = j \\ 0, & i \neq j \end{cases},$$

where the trimming (24) is applied, as well as:

$$\begin{aligned} H_{k,ij} > 0.5\kappa_{\phi,i} &\rightarrow H_{k,ij} := 0.5\kappa_{\phi,i} \\ H_{k,ij} < -0.5\kappa_{\phi,i} &\rightarrow H_{k,ij} := -0.5\kappa_{\phi,i} \end{aligned} \quad (25)$$

where the latter supposes a certain degree of “flatness” in ϕ_p by supposing that no second derivative should ever be greater in magnitude than half of the maximal first derivative.

- If $k \geq 2n_v + 1 + \sum_{i=1}^{n_v-1} i$, fit a full quadratic model to the data:

$$\phi_p(\mathbf{v}) \approx a_0 + \sum_{i=1}^{n_v} a_i v_i + \sum_{i=1}^{n_v} \sum_{j=1}^{n_v} a_{ij} v_i v_j,$$

where $a_{ij} = a_{ji}$. Define:

$$\frac{\partial \hat{\phi}_p}{\partial v_i} \Big|_{\mathbf{v}_k} := a_i + \sum_{j=1}^{n_v} a_{ij} v_{k,j}, \quad H_{k,ij} := 2a_{ij},$$

and apply the trimmings (24) and (25) if necessary.

We note that while this scheme is not guaranteed to generate a positive-definite Hessian, the consequences of failing to do so are not expected to be very detrimental in our context, since the optimization target is, again, only a guide and does not affect the general reliability of the solver.

© August 23, 2013 by the authors; submitted to *Processes* for possible open access publication under the terms and conditions of the Creative Commons Attribution license <http://creativecommons.org/licenses/by/3.0/>.