Harmonizing the use of optimization and feedback in process operations and control

James B. Rawlings [a,1] and Robert D. McAllister [a]

[a] Department of Chemical Engineering, University of California, Santa Barbara, CA 93106

Abstract

This paper addresses how to combine the two main engineering design tools, optimization and feedback, to obtain high performance process control or process operation systems. We discuss the historical development of process control, where feedback dominated the early practical designs, and optimization was then incorporated much later through technologies like model predictive control. In chemical production scheduling, on the other hand, optimization had a strong early influence on the problem formulation, and feedback has only recently made an appearance. The recent developments in both process control and scheduling are illustrated with specific examples emerging from this series of FOCAPO/CPC meetings. The paper next presents recent theoretical developments in nominal and stochastic model predictive control. The closed-loop properties that arise from these different open-loop optimal control problems are then compared. The paper closes with some discussion of when the improvements of the closed-loop properties are worth the added complexity of the stochastic optimal control problem.

Keywords

Feedback control, optimal control, on-line optimization, scheduling, stochastic optimal control, robust model predictive control.

1 Introduction

Creating systems to ensure reliable and high performance process control or process operations requires smart use of our two main engineering design[2] tools: optimization and feedback. The purpose of optimization is to ensure high performance, especially in the nominal case, and the purpose of feedback is to ensure reliability of this performance in the face of the inevitable unknown disturbances. The holy grail would be to perform an optimal feedback design; we call it the holy grail because that problem remains intractable for anything but the simplest situations: processes modeled by unconstrained linear dynamical systems and quadratic performance objectives.[3]

If the optimal feedback design is out of reach for most processes, then, unfortunately, we have to be smart. The problem of how to best combine optimization and feedback is now complex and many alternatives can be envisaged, each with particular strengths and weaknesses. And we cannot let every system design become its own special case. The development cost is too high, and the maintenance cost of many, one-off, complex designs is overwhelming. So we still require invention of *general* design approaches that handle large classes of processes. Creating that kind of general design requires *ingenuity* combined with careful analysis of performance and reliability tradeoffs. So one objective of this paper is to summarize our field's current level of ingenuity.

You can learn a lot about a civilization's mathematical development by observing how general is their concept of a function. Similarly you can learn a lot about a civilization's process systems engineering development by observing what combinations of optimization and feedback have they been able to reduce to practice. We shall attempt such an assessment of our field in this paper.

[1] Corresponding author. Email: jbraw@ucsb.edu. This work was supported by the National Science Foundation under Grant 2027091.

[2] Note that throughout this paper, *design* refers to the design of an *algorithm* or method used for control or scheduling. This meaning should not be confused with the field of *process design*, which we do not discuss here.

[3] See, for example, (Rawlings et al., 2020, pp. 90–93) for a brief statement of the difficulties of solving the optimal feedback problem using dynamic programming for nonlinear models of reasonable state dimension.

2 Historical context

History provides valuable context for any assessment of a current status and an outlook for the future. In the process control side of process systems engineering, the considerations of practical application mainly drove the field forward in its early history. Bequette (2019) provides an excellent summary of this period and describes the industrial developments for controlling temperature, pressure, and level during the period 1910–1930s. The opportunity to improve process operation through automation led to the development of feedback control analog hardware that implemented simple PID (proportional-integral-derivative) algorithms that responded to measurement signals, and made automatic adjustments to the available actuators to maintain process operations despite the unknown process disturbances taking place. Automation was essential because the hardware needed to respond in real-time in a continuous processing environment, and human operators were expensive and unreliable to handle large numbers of routine and low-level processes and units. In this early era of process control, there is clearly a large contribution from feedback and little or no contribution from optimization. There is also a human, cultural impact of this history: process and control engineers that lived through this period could not help but develop a high appreciation for the power of feedback. There was little regard for optimization as it was hardly even used in any practical, industrial process control application.

Optimization was not absent from the entire field of control, it was just not (yet) having much impact on practice. Control theory was under steady development, however, and optimization was mainstream to the development of control theory. The late 1950s and early 1960s were an explosive period of rapid development of optimal control, in particular. Tremendous progress was made on both the open-loop optimal control problem, summarized in results like Pontryagin's maximum principle (Pontryagin et al., 1961), as well as the feedback form of the optimal control problem, summarized under the names of dynamic programming and the Hamilton-Jacobi-Bellman equation (Bellman and Dreyfus, 1962; Bryson and Ho, 1975). But even the developers of optimal *open-loop* control during this period had a strong sense for the requirements to deploy a *feedback* version of the controller in applications. For example, the following passage from a leading optimal control textbook of the 1960s has an almost modern outlook (Lee and Markus, 1967, p. 24).

"In each optimal control problem our ultimate goal is to synthesize the optimal controller by an appropriately designed closed or feedback loop. The advantage of such a closed-loop control, as against an open-loop control, is that the process then becomes self-adjusting and self-correcting. A feedback control can often correct for unpredictable variations in the environment of the plant or for repeated perturbations or irregularities in the process."

A note on model predictive control (MPC). Many accounts have been written about the history of MPC, and there is little need to revisit those accounts here. A standard story line is that MPC emerged from industrial practice in the process industries in the 1970s with implementations such as IDCOM and DMC. While no one doubts that successful industrial implementation was a critical event—perhaps *the* critical event—leading to widespread *interest* in MPC, it does not address the development of the ideas and concepts underpinning MPC.

The critical intellectual idea to combine feedback and optimal control with a moving horizon certainly did not first emerge from 1970s industrial implementations in the process industries. It is difficult to pin down a single source for this idea. For example, consider the following passage, taken again from Lee and Markus's 1967 graduate level textbook, "Foundations of Optimal Control" (Lee and Markus, 1967, p. 423).

"One technique for obtaining a feedback controller synthesis from knowledge of open-loop controllers is to measure the current control process state and then compute very rapidly for the open-loop control function. The first portion of this function is then used during a short time interval, after which a new measurement of the process state is made and a new open-loop optimal control function is computed for this new measurement. The procedure is then repeated. In this way external disturbances and other unknowns are taken into account in much the same way as is done by a feedback controller."

So the authors first offer a succinct summary of model predictive control, and then go on to make an explicit connection to optimal feedback control.

"If no disturbance or other unknowns are encountered, the recomputed control function should agree with the appropriate portion of the previously computed controller. This is essentially the principle of optimality [Bellman] in the theory of dynamic programming, a feedback principle."

Here is a clearly articulated proposal to combine optimization and feedback, and the reader cannot help but be struck by the casual and offhand manner in which the idea is presented. The authors are not announcing a breakthrough research idea, but offering a simple reminder of some common lore of the control community, and stressing the importance of feedback in addition to open-loop optimal control. The least we can conclude is that the notion of combining open-loop optimal control with a moving horizon to obtain a feedback implementation was *mainstream* in the early 1960s control theory community. It seems pointless at best and misleading at worst to attempt to assign *that* critical idea to any individual or group of practitioners. It likely occurred to *all* of the 1950–1960 era control theory researchers who had been exposed in their education to both feedback controller design and optimal control theory.

What was missing in the 1960s was essentially a technology to make the phrase "compute very rapidly for the open-loop control function" deployable in applications. Fortunately for all of us working today in process systems engineering, the computing technology necessary for implementing open-loop optimal control as a feedback controller, entered a period of extremely rapid development shortly thereafter. The timing could not have been better. The combination of optimal control theory and inexpensive online computing technology enabled a revolution in what control theory could be applied in process control applications. The aftershocks of this rapid development in optimal control and fast and inexpensive online computing are still being felt today. The successful deployment of these ideas in process control applications led to successful deployments in many other industries such as flight control, robotics and mechatronics applications, HVAC systems, power systems, etc. It is not an exaggeration to say that the process control practitioners led the way to a revolution in advanced control technology that was deployed widely across many industrial sectors of the world economy.

Another illustrative example of the importance of judiciously combining feedback and optimal control arises when we consider accounting for uncertainty in the MPC design. One natural robust control formulation is the so-called min-max approach where we minimize over the manipulated variables, as in standard MPC, but maximize over the uncertainty in the inner problem. That approach was applied to the "Shell Standard Control Problem" by the Shell research team at the second Shell Workshop held in Houston in 1988 (Cuthrell et al., 1990). And what was the outcome? The bottom line was that there was *no improvement* in the closed-loop robustness compared to a *nominal* MPC design. All of us attending the workshop were scratching our heads over that outcome. David Mayne, an eminent control theorist from electrical engineering, pointed out the reason a few years later. One cannot achieve robustness by optimizing over an open-loop optimal control as a sequence of control actions. One must optimize over control *policies*. As David summarizes it (Rawlings et al., 2020, p. 199)

"The obvious and well-known conclusion is that feedback control is superior to open-loop control when uncertainty is present. Feedback control requires determination of a control *policy*, however, which is a difficult task if nonlinearity and/or constraints are features of the optimal control problem."

And David shows how to design a (robust) min-max MPC that does guarantee robustness (Rawlings et al., 2020, pp. 220-223). The trick is to *finitely* parameterize the policy to maintain a tractable online optimization problem. We illustrate this point also later in this paper when we define *stochastic* MPC, another form of MPC that addresses model uncertainty. David's contribution to MPC theory were especially noteworthy because he brought his vast expertise in optimal control and nonlinear control theory to bear on the MPC problem; he also brought the MPC problem to the attention of the broader control theory community, which sparked many different contributions to MPC from a much wider group of outstanding control theorists.

In summary, the interaction of these two communities, the practitioners in the process industries, and the control theory researchers, was indeed a special one. And that interaction took place largely at *this* series of meetings over the last thirty years. There was a lot of passion displayed by both groups, and there were many good ideas contributed by both groups as well. There was some heat, but there was also a lot of light. A common mythology discussed at these meetings centered around the existence of a supposed "gap" between academics doing control theory for theory's sake, and practitioners requiring reliable technologies to address their current control operational difficulties. That characterization does not do justice to the interactions that I (JBR) witnessed at these meetings. When you see one set of respected *practitioners* arguing to deploy large-scale model predictive control, and another set of respected *practitioners* resisting that notion and calling instead for better tuning guidelines to handle multi-loop PID controllers, that is hardly an academic-industrial gap between theory and practice. That is basically a conflict over what the future of industrial practice is going to become. And everyone had a stake in that outcome, and no one could predict whether either of these technologies or something else entirely was going to emerge as a clear winner. I (JBR) recall traveling home after several of the CPC meetings and thinking, "Wow, what just happened; where are we going next?" I could not wait to get back home to start working on the new ideas. For me, that is the legacy of the CPC side of these meetings.

A note on FOCAPO/CPC collaborations. The authors know much less about the historical development of process operations outside of process control. Consider chemical production scheduling, for example. The history of process scheduling is quite different than process control where the early practice was dominated by feedback solutions with little optimization. General scheduling (e.g., for discrete manufacturing) was studied in the 1950s. Chemical production scheduling is younger still—it appeared as a subdiscipline in the 1970s. Chemical production scheduling felt a stronger early influence from academic research focusing on optimization, i.e., solving for an optimal schedule. Feedback was not much considered until around the time of the merging of FOCAPO and CPC in 2012. The first joint FOCAPO/CPC meeting took place in Savannah, GA in January 2012. Christos Maravelias and I (JBR) were collaborating then on using feedback in optimal scheduling to add performance guarantees to the closed-loop schedule. We thought that collaboration would illustrate overcoming challenges in bringing these two communities into closer contact. The key for us turned out to be developing a common language for expressing the process models. That development took place over a few months in this way: Christos and I shared a PhD student, Kaushik Subramanian. Kaushik would meet with me, and I would say, "Those scheduling models do not make any sense; here's how we would express a dynamical system model." Kaushik would meet with Christos, and Christos would say, "Wait, the control people do what? That can't represent a scheduling model." Finally, after a few months of this treatment Kaushik said, "Here's how you translate a scheduling model into a dynamic state-space model, and here's how you then solve it."

What I (JBR) learned from this experience is that to bring two research communities together, you require... one really smart graduate student. The graduate student has conversations with both professors. Over a few months the graduate student becomes bilingual. Neither professor ever learns the other's language.[4] After a few months the graduate student can suddenly speak a new language to both professors, and you write down that new language. The beauty is that all three can now speak the new language. The language that Kaushik developed is summarized in (Subramanian et al., 2012). The translation provided enables the control community to design a feedback solution to the scheduling problem by formulating an MPC control problem for the dynamic model equivalent of the scheduling model.

Fruitful collaborations often impact everyone participating. Collaborators do not merely import ideas from other fields, they cross-pollinate. Control people expect that control theory might have something to contribute to chemical production scheduling because control exploits feedback to obtain robustness to disturbances and modeling errors, and feedback was largely missing in early optimal scheduling. But did scheduling have any impact on control theory? The answer is yes. By its nature, scheduling focuses attention on the *discrete* decision variables as an essential part of any application. Integers abound: which products should be made, in what sequence, in which pieces of equipment? Of course, discrete decisions appear in many process control problems, but traditionally they were handled with heuristics and supervisory logic, and not solved on-line as part of the optimal control problem. In MPC, for example, discrete decisions had been introduced in (Slupphaug et al., 1998; Bemporad and Morari, 1999), but the theory for the closed-loop properties for these systems was fragmentary and developed largely as special cases. Motivated by their importance in scheduling applications (and HVAC energy optimization), these integer decisions were fully integrated into mainstream MPC theory in (Rawlings and Risbeck, 2017). This development was the thesis topic of another smart graduate student, Michael Risbeck, who also did joint research with Christos. Michael took Kaushik's starting point to its logical conclusion and formulated the fully integrated scheduling *and* control problem, also expressed in state-space form. When tractable, that formulation can also be optimized directly in an economic MPC framework, and when a feasible reference trajectory is used as a terminal constraint, the closed-loop economic performance of the nominal system is shown to be at least as good as the reference trajectory (Risbeck et al., 2019). So we have a performance bound on the (nominal) closed-loop system that is better than the best available feasible, *open-loop* reference trajectory. From a theory perspective, that's a good starting point for integrating scheduling and control, but much work remains to reduce such an approach to practice.

The second author of this paper also participated as a PhD student in a (third) research collaboration with Christos. This one was a somewhat different but similar story to Kaushik's. In this collaboration we were trying to characterize and analyze the robustness of the feedback scheduling solution to unknown disturbances, such as equipment breakdowns, canceled orders, task delays, etc. But we already had a common modeling language, so conversations were easier at first. The ones with me then went like this, "Wait, what? Those disturbances are not small, like process and measurement noise, those disturbances are *large*; no controller can handle those. Koty, tell Christos that's hopeless." Koty would come back later and tell me, "Well, Christos says, 'Hopeless or not, those are the relevant disturbances. What can a feedback solution do about them?' " And that seemed like an impasse. But then Koty had a good idea. Robustness first had to be redefined from a stochastic perspective. The focus then changes from bounds on worst performance over all disturbances to bounds on *average* performance over all disturbances. And those large disturbances required a different probability distribution. And then we could say something about the stochastic robustness to those large (but rare) disturbances (McAllister et al., 2022; McAllister and Rawlings, 2021). What I learned from this collaboration was basically the same lesson. The graduate student talks to two different experts. The graduate student eventually internalizes what those experts know. The graduate student then has a new idea. Why couldn't Christos and I just generate the new idea directly without the third party? I have no explanation. Maybe psychologists have studied this issue and know the answer. Interestingly, we also never required a meeting with all three of us. Of course we had some of those meetings as well, but they were not the ones where the magic happened.

So in 2023, at the time of this *third* joint FOCAPO/CPC meeting, have the fields of optimal operation, production scheduling in this case, and feedback control been successfully integrated? A fair answer seems to be: No, not yet, but the situation remains fluid. Certainly such an integration has not been reduced to standard industrial practice as it has in process control. After a plenary lecture on process scheduling at the second joint meeting in Tucson (2017), I (JBR) asked the presenter if there was a reason one would not want to *reschedule* every time new information became available. The speaker's reaction was basically, "Mon dieux, we would never consider something crazy like that. Practitioners do not like changes to the schedule!" I remember thinking, "Well, at least the control practitioners were not telling to us to keep our hands off the valves in control applications."

But consider as well the recent and first textbook on chemical production scheduling(Maravelias, 2021, p. 365).

"In general, decisions obtained from a finite horizon planning model are implemented in a setting where the system operates indefinitely and under uncertainty, hence, what the model returns as *optimal predicted* solution may not be optimal for the long-term operation of the actual system.... As time passes by, more information becomes available, and this information should be accounted for as soon as possible to determine new decisions. Thus, real-time scheduling is a generalization of rescheduling, since it is based on a recomputation that is carried out not only upon the realization of trigger events but also periodically to consider new information."

So time will tell. Perhaps Maravelias (2021) will become

[4] Professors are notoriously stubborn and slow learners.

the Lee and Markus of the scheduling community and be quoted at some FOCAPO/CPC meeting 50 years from now as *proof* that the notion of combining feedback with optimal scheduling was a mainstream idea in the scheduling community as early as the 2020s.

3 Nominal and Stochastic Model Predictive Control

So after these admittedly subjective and selected historical comments, we would like to return to the main question posed in the title of the paper: how do we best combine or *harmonize* feedback and optimization to obtain high performance that is also robust to either unknown or imperfectly modeled disturbances affecting a system. And now we would like to be precise about what is known and what is not known about this issue. Grand generalizations may pass muster when summarizing some developments of the distant past, but such generalizations are dangerous and misleading when summarizing a complex, current state of the art. Also, given the limits on the authors' expertise, we have to restrict this discussion to the state of affairs in model predictive control and optimal stochastic control.

Wonham (1969) starts his optimal stochastic control paper with the sentence, "STOCHASTIC CONTROL is a convenient misnomer for the control of systems subject to stochastic disturbances." Similarly, stochastic MPC is applying MPC to systems where the *model* of the system includes stochastic disturbances, i.e., the process model is (1) rather than (2). But certainly we cannot expect the plant's true disturbances to be captured by the disturbance *model* chosen in stochastic MPC. That case reflects an unrealistic belief that sure, nature is random, but we somehow get to know the randomness exactly. If only life were that simple.

But *simulation* and case studies can tell us a lot about what might happen when we apply MPC to a plant with some unknown randomness. In simulation, the designer can test the performance of the control system in many different scenarios. For example, we can evaluate nominal MPC's performance using the disturbances as modeled in the stochastic MPC controller (is nominal MPC robust to disturbances?) We can evaluate stochastic MPC's performance when there are no disturbances (does stochastic MPC sacrifice nominal performance in search of robustness to disturbances?). Case studies of this sort are quite informative. But case studies can take us only so far. It is unlikely that we can draw general conclusions without theory and analysis of the control systems. We next present some recent theoretical developments on the question of whether stochastic MPC has obtained any closed-loop properties that are different than those achieved by nominal MPC.

Notation

Let $\mathbb{I}$ and $\mathbb{R}$ denote the integers and reals. Let superscripts on these sets denote dimensions and subscripts on these sets denote restrictions (e.g., $\mathbb{R}^n$ for real vectors of dimension n and $\mathbb{I}_{0:N}$ for integers from 0 up to and including $N \geq 0$.) Let $|\cdot|$ denote Euclidean norm. The function $\alpha : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is in class $\mathcal{K}$, denoted $\alpha(\cdot) \in \mathcal{K}$, if $\alpha(\cdot)$ is continuous, strictly increasing, and $\alpha(0) = 0$. The function $\alpha : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is in class $\mathcal{K}_\infty$ if $\alpha(\cdot) \in \mathcal{K}$ is unbounded, i.e., $\lim_{s\to\infty} \alpha(s) = \infty$. A function $\beta : \mathbb{R}_{\geq 0} \times \mathbb{I}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is in class $\mathcal{KL}$ if, for fixed $k \in \mathbb{I}_{\geq 0}$, the function $\beta(\cdot, k)$ is in class $\mathcal{K}$ and, for fixed $s \in \mathbb{R}_{\geq 0}$, the function $\beta(s, \cdot)$ is nonincreasing and $\lim_{k\to\infty} \beta(s, k) = 0$. Let $\mathbb{E}[\cdot]$ denote expected value of a random variable.

3.1 System Model

We consider discrete time, stochastic systems of the following form.

$$x^+ = f(x, u, w) \tag{1}$$

in which $x \in \mathbb{R}^n$ is the state, $u \in \mathbb{U} \subseteq \mathbb{R}^m$ is the input, $w \in \mathbb{W} \subseteq \mathbb{R}^q$ is the disturbance, and x^+ denotes the successor state. We treat the origin ($x = u = 0$) as the steady-state target (setpoint) of interest. Moreover, we assume that the disturbances $w \in \mathbb{W}$ are random variables that are independent and identically distributed in time (i.i.d.) with zero mean ($\mathbb{E}[w] = 0$). Let μ denote the probability distribution for w and let $\mu \in \mathcal{M}(\mathbb{W})$ denote the collection of all probability distributions on the support $\mathbb{W}$ that are zero mean. Let Σ denote the covariance matrix for w, i.e.,

$$\Sigma := \mathbb{E}[ww']$$

for all $\mu \in \mathcal{M}(\mathbb{W})$. To control this stochastic system, we consider two variations of MPC.

3.2 Nominal MPC

In nominal MPC, we use only a nominal model of the system, i.e.,

$$x^+ = f(x, u, 0) \tag{2}$$

to design the controller. For the horizon $N \in \mathbb{I}_{\geq 1}$, we use $\hat{\phi}(k; x, \mathbf{u})$ to denote the state trajectory of (2) at time $k \in \mathbb{I}_{0:N}$, given the initial state $x \in \mathbb{R}^n$ and input trajectory $\mathbf{u} \in \mathbb{U}^N$. We allow input constraints $u \in \mathbb{U}$, but do not enforce state constraint in the optimization problem. For a perturbed system, there is no guarantee that these state constraints can be satisfied. Instead, we convert these state constraints to penalty functions in the stage cost (Zheng and Morari, 1995; Scokaert and Rawlings, 1999). Thus, the optimizer avoids violating these constraints if possible, but does not produce an infeasible optimization problem otherwise. We do, however, require a terminal state constraint denoted by the set $\mathbb{X}_f \subseteq \mathbb{R}^n$. To characterize the performance objective for the controller, we define a stage cost $\ell(x, u)$. We also define a corresponding terminal cost $V_f(x)$ that is chosen to guarantee stability and robustness.

We define the set of admissible input trajectories, feasible initial states, and objective function, respectively, as

$$\mathcal{U}(x) := \{\mathbf{u} \in \mathbb{U}^N : x(N) \in \mathbb{X}_f\}$$
$$\mathcal{X} := \{x \in \mathbb{R}^n : \mathcal{U}(x) \neq \emptyset\}$$
$$V(x, \mathbf{u}) := \sum_{k=0}^{N-1} \ell(x(k), u(k)) + V_f(x(N))$$

in which $x(k) = \hat{\phi}(k;x,\mathbf{u})$. The nominal MPC optimization problem for any $x \in \mathcal{X}$ is defined as

$$\mathbb{P}(x): \; V^0(x) := \min_{\mathbf{u}\in\mathcal{U}(x)} V(x,\mathbf{u})$$

and the optimal solution is denoted $\mathbf{u}^0(x)$. We implement, however, only the first input in this optimal solution and the control law for nominal MPC is therefore

$$\kappa(x) := u^0(0;x)$$

The resulting closed-loop system is then

$$x^+ = f(x,\kappa(x),w) \tag{3}$$

and we use $\phi(k;x,\mathbf{w}_k)$ to denote the state of (3) at time $k \in \mathbb{I}_{\geq 0}$, given the initial state $x \in \mathcal{X}$ and disturbance trajectory

$$\mathbf{w}_k = (w(0),w(1),\dots,w(k-1)) \in \mathbb{W}^k$$

Thus, nominal MPC does not account for disturbances in the problem formulation directly, but does address disturbances through state feedback.

The following assumptions ensure that nominal MPC is inherently robust.

Assumption 3.1. *The system $f(\cdot)$, stage cost $\ell(\cdot)$, and terminal cost $V_f(\cdot)$ are continuous and satisfy $f(0,0,0) = 0$, $\ell(0,0) = 0$, $V_f(0) = 0$.*

Assumption 3.2. *The set $\mathbb{U}$ is compact and contains the origin. The set $\mathbb{X}_f$ is defined by $\mathbb{X}_f := \{x \in \mathbb{R}^n : V_f(x) \leq \tau\}$ for some $\tau > 0$.*

Assumption 3.3. *There exists a terminal control law $\kappa_f : \mathbb{X}_f \to \mathbb{U}$ such that*

$$f(x,\kappa_f(x),0) \in \mathbb{X}_f$$
$$V_f(f(x,\kappa_f(x),0)) \leq V_f(x) - \ell(x,\kappa_f(x))$$

for all $x \in \mathbb{X}_f$.

Assumption 3.4. *There exists $\alpha_\ell(\cdot) \in \mathcal{K}_\infty$ such that $\alpha_\ell(|x|) \leq \ell(x,u)$ for all $(x,u) \in \mathbb{R}^n \times \mathbb{U}$.*

3.3 Stochastic MPC

In stochastic MPC, we use a stochastic model of the system. For now, we assume this model is equivalent to the underlying plant in (1). Since we are considering all possible realizations of the disturbance in the optimization problem, we want to optimize over a trajectory of control *policies* for this system to account for all possible realizations of the state trajectory. To avoid the difficulties of dynamic programming and ensure that the optimization problem is tractable, however, we parameterize this control policy as $\pi(x,v)$ in which $v \in \mathbb{V}$ is the vector of parameters for the control policy, e.g., $\pi(x,v) = Kx + v$ in which K is a fixed feedback gain matrix. Thus, the system of interest is

$$x^+ = f(x,\pi(x,v),w) \tag{4}$$

We use $\hat{\phi}^s(k;x,\mathbf{v},\mathbf{w})$ to denote the state of (4) at time $k \in \mathbb{I}_{0:N}$, given the initial condition $x \in \mathbb{R}^n$, trajectory of control policies defined by $\mathbf{v} \in \mathbb{V}^N$, and disturbance trajectory $\mathbf{w} \in \mathbb{W}^N$.

Since we are considering the disturbance directly in the optimization problem, we can consider hard state as well as input constraints, i.e.,

$$(x,u) \in \mathbb{Z} \subseteq \mathbb{R}^n \times \mathbb{U}$$

We define the admissible control parameter trajectories, feasible initial states, and cost function, respectively, as

$$\begin{aligned}
\mathcal{V}(x) &:= \{\mathbf{v} \in \mathbb{V}^N : \\
&\qquad (x,\pi(x(k),v(k))) \in \mathbb{Z} \;\forall \mathbf{w} \in \mathbb{W}^N,\; k \in \mathbb{I}_{0:N-1} \\
&\qquad \text{and } x(N) \in \mathbb{X}_f \;\forall \mathbf{w} \in \mathbb{W}^N\} \\
\mathcal{X}^s &:= \{x \in \mathbb{R}^n : \mathcal{V}(x) \neq \emptyset\} \\
V^s_\mu(x,\mathbf{v}) &:= \mathbb{E}\left[\sum_{k=0}^{N-1} \ell(x(k),\pi(x(k),v(k))) + V_f(x(N))\right]
\end{aligned}$$

in which $x(k) = \hat{\phi}^s(k;x,\mathbf{v},\mathbf{w})$. Note that the cost function depends on the probability distribution for w, i.e., μ. The SMPC optimization problem for any $x \in \mathcal{X}^s$ is defined as

$$\mathbb{P}^s_\mu(x): \; V^{s0}_\mu(x) := \min_{\mathbf{v}\in\mathcal{V}(x)} V^s_\mu(x,\mathbf{v})$$

and the optimal solution is denoted $\mathbf{v}^{s0}_\mu(x)$. We again implement only the first control policy in this optimal solution and the control law for SMPC is therefore

$$\kappa^s_\mu(x) := \pi(x, v^{s0}_\mu(0;x))$$

The resulting closed-loop system is then

$$x^+ = f(x,\kappa^s_\mu(x),w) \tag{5}$$

and we use $\phi^s_\mu(k;x,\mathbf{w}_k)$ to denote the state of (5) at time $k \in \mathbb{I}_{\geq 0}$, given the initial condition $x \in \mathcal{X}^s$, disturbance sequence $\mathbf{w}_k \in \mathbb{W}^k$, and probability distribution $\mu \in \mathcal{M}(\mathbb{W})$. Note that the control law and therefore closed-loop system depend on the probability distribution μ used in the SMPC problem formulation.

For SMPC, we require modified versions of Assumption 3.2 and Assumption 3.3.

Assumption 3.5. *The set $\mathbb{Z}$ is closed and contains the origin. The sets $\mathbb{U}$ and $\mathbb{X}_f$ are compact and contain the origin. The set $\mathbb{X}_f$ contains the origin in its interior. The set $\mathcal{X}^s$ is bounded.*

Assumption 3.6. *There exists a continuous terminal control law $\kappa_f : \mathbb{X}_f \to \mathbb{U}$ such that*

$$f(x,\kappa_f(x),w) \in \mathbb{X}_f \qquad \forall w \in \mathbb{W}$$
$$V_f(f(x,\kappa_f(x),0)) \leq V_f(x) - \ell(x,\kappa_f(x))$$

for all $x \in \mathbb{X}_f$. Furthermore, $(x,\kappa_f(x)) \in \mathbb{Z}$ and $\pi(x,0) = \kappa_f(x)$ for all $x \in \mathbb{X}_f$.

Note that Assumption 3.6 implicitly restricts the size of $\mathbb{W}$ that may be considered for a specific system. Sufficiently large disturbances may render the construction of a suitable terminal control and a terminal set either difficult or impossible for nonlinear systems and open-loop unstable linear systems with input constraints. We also require the following assumption for the control law parameterization.

Assumption 3.7. *The set $\mathbb{V}$ is compact and contains the origin. The function $\pi(\cdot)$ is continuous.*

3.4 Properties

We introduce three potential definitions of robustness that include both deterministic and stochastic representations of the disturbance trajectory. Since the SMPC control law varies with μ, we use the generic control law $\kappa_\mu(\cdot)$ and corresponding closed-loop system $\phi_\mu(\cdot)$ in the following definitions to indicate this potential dependence. Since both MPC and SMPC rely on optimization to define the control law, we must first ensure that these optimization problems remain feasible for the closed-loop system. We characterize this property through robust positive invariance.

Definition 3.8 (RPI)**.** *A set $\mathcal{X}$ is robustly positive invariant (RPI) for the system $x^+ = f(x, \kappa_\mu(x), w)$, $w \in \mathbb{W}$ if $x^+ \in \mathcal{X}$ for all $x \in \mathcal{X}$, $w \in \mathbb{W}$, $\mu \in \mathcal{M}(\mathbb{W})$.*

If the feasible set of the MPC or SMPC optimization problem is RPI, then the optimization problem is robustly recursively feasible, i.e., the control law and closed-loop system are well defined. For this closed-loop system, we define deterministic robustness as follows in which $||\mathbf{w}_k|| := \max_{i \in \mathbb{I}_{0:k-1}} |w(i)|$.

Definition 3.9 (RAS)**.** *The origin is robustly asymptotically stable (RAS) for a system $x^+ = f(x, \kappa_\mu(x), w)$, $w \in \mathbb{W}$ in an RPI set $\mathcal{X}$ if there exist $\beta(\cdot) \in \mathcal{KL}$ and $\gamma(\cdot) \in \mathcal{K}$ such that*

$$|\phi_\mu(k; x, \mathbf{w}_k)| \leq \beta(|x|, k) + \gamma(||\mathbf{w}_k||) \tag{6}$$

for all $x \in \mathcal{X}$, $\mathbf{w}_k \in \mathbb{W}^k$, and $k \in \mathbb{I}_{\geq 0}$.

Thus, RAS ensures that the closed-loop state of the system converges to a neighborhood of the origin defined by a $\mathcal{K}$-function of the largest disturbance experienced up to time k, i.e., $\gamma(||\mathbf{w}_k||)$. Note that this bound must hold for any specific realization of the disturbance trajectory. As we intend to consider a stochastic representation of the disturbance, we are also interested in a similar bound based on the stochastic properties of the underlying system. We define stochastic robustness as follows.

Definition 3.10 (RASiE)**.** *The origin is robustly asymptotically stable in expectation (RASiE) for a system $x^+ = f(x, \kappa_\mu(x), w)$, $w \in \mathbb{W}$ in an RPI set $\mathcal{X}$ if there exist $\beta(\cdot) \in \mathcal{KL}$ and $\gamma(\cdot) \in \mathcal{K}$ such that*

$$\mathbb{E}[|\phi_\mu(k; x, \mathbf{w}_k)|] \leq \beta(|x|, k) + \gamma(\mathrm{tr}(\Sigma)) \tag{7}$$

for all $x \in \mathcal{X}$, $\mu \in \mathcal{M}(\mathbb{W})$, and $k \in \mathbb{I}_{\geq 0}$.

Similar to RAS, RASiE requires that the effect of the initial condition vanishes as $k \to \infty$ with a persistent term based on the disturbance w. In RASiE, however, (7) bounds a stochastic property of the closed-loop system ($\mathbb{E}[\cdot]$) based on a stochastic property of the disturbances ($\mathrm{tr}(\Sigma)$). Note that Σ depends on μ and (7) must hold for all $\mu \in \mathcal{M}(\mathbb{W})$.

These two definitions of robustness, one deterministic and one stochastic, consider the usual metric of performance in process control: distance to the origin (setpoint). In MPC, however, we define the control law by optimizing a performance metric for the system defined by the stage cost $\ell(\cdot)$. By requiring Assumption 3.4, we ensure that if $\ell(x, u) \to 0$, then $|x| \to 0$ as well. We nonetheless allow for significant flexibility in selecting $\ell(\cdot)$ that allows us to tune the stage cost to reflect the relative importance of the different elements of the state and input, or even include specific economic metrics of performance in the problem formulation. Given this flexibility and the benefits obtained from this more general definition of performance, we propose a definition of robustness with respect to the stage cost, i.e., the performance metric assigned to the MPC and SMPC problem formulation. We call this property robust asymptotic stability in expectation with respect to the stage cost, abbreviated as ℓ-RASiE.

Definition 3.11 (ℓ-RASiE)**.** *The origin is ℓ-RASiE with respect to the stage cost $\ell(\cdot)$ for a system $x^+ = f(x, \kappa_\mu(x), w)$, $w \in \mathbb{W}$ in the RPI set $\mathcal{X}$ if there exist $\beta(\cdot) \in \mathcal{KL}$ and $\gamma(\cdot) \in \mathcal{K}$ such that*

$$\mathbb{E}[\ell(x(k), \kappa_\mu(x(k)))] \leq \beta(|x|, k) + \gamma(\mathrm{tr}(\Sigma)) \tag{8}$$

in which $x(k) := \phi_\mu(k; x, \mathbf{w}_k)$ for all $x \in \mathcal{X}$, $\mu \in \mathcal{M}(\mathbb{W})$, and $k \in \mathbb{I}_{\geq 0}$.

Nominal MPC is known to be RAS for sufficiently small disturbances (Grimm et al., 2004; Pannocchia et al., 2011; Yu et al., 2014; Allan et al., 2017). Moreover, we can establish that nominal MPC also satisfies the two definitions of stochastic robustness introduced in this section, as detailed in the following theorem (McAllister and Rawlings, 2022d).

Theorem 3.12 (Nominal MPC)**.** *Let Assumptions 3.1 to 3.4 hold. For every $\rho > 0$ there exists $\delta > 0$ such that for all $\mathbb{W} \subseteq \{w \in \mathbb{R}^q : |w| \leq \delta\}$ and the system $x^+ = f(x, \kappa(x), w)$, $w \in \mathbb{W}$, and the set $\mathcal{S} := \{x \in \mathbb{R}^n : V^0(x) \leq \rho\} \cap \mathcal{X}$ we have that*

(i) The set $\mathcal{S}$ is RPI.

(ii) The origin is RAS in the set $\mathcal{S}$.

(iii) The origin is RASiE in the set $\mathcal{S}$.

(iv) The origin is ℓ-RASiE in the set $\mathcal{S}$.

For sufficiently small disturbances ($|w| \leq \delta$), nominal MPC satisfies all of the definitions of deterministic and stochastic robustness introduced in this section. This robustness is *inherent* to nominal MPC through feedback, despite the fact that nominal MPC does not consider the disturbance directly in the problem formulation. By contrast, SMPC does consider a stochastic description of the disturbance and optimizes a stochastic objective function, thereby affording stochastic robustness to the closed-loop system (Cannon et al., 2009; Kouvaritakis et al., 2010; Lorenzen et al., 2016; Chatterjee and Lygeros, 2014; Mayne and Falugi, 2019). In particular, we can establish the following result (McAllister and Rawlings, 2022a,d).

Theorem 3.13 (SMPC)**.** *Let Assumptions 3.1 and 3.4 to 3.7 hold with μ and $\mathbb{W}$ known exactly. Then for the system $x^+ = f(x, \kappa^s_\mu(x), w)$, $w \in \mathbb{W}$, we have that*

(i) The set X^s is RPI.

(ii) The origin is RASiE in the set X^s.

(iii) The origin is ℓ-RASiE in the set X^s.

Thus, nominal MPC and SMPC satisfy the same definitions of stochastic robustness, in terms of both distance to the origin and stage cost. In other words, SMPC is not providing some unique form of stochastic robustness not available to nominal MPC. On the contrary, SMPC is instead *sacrificing* deterministic robustness (RAS) for the possibility of improving the stochastic robustness of the closed-loop system.

Moreover, SMPC optimization problems are significantly more difficult to solve than nominal MPC optimization problems. These problems require evaluation of an expectation, which is often computed by sampling the random variable w and constructing all possible trajectories for these realizations of the disturbance, i.e., a scenario tree. Thus, the number of variables and therefore SMPC problem size is s^N times larger than the nominal MPC problem, in which s is the number of samples of w at each time step and N is the horizon length of the SMPC problem.

3.5 *A loss of nominal asymptotic stability*

We consider a small example of level control in two tanks, adapted from McAllister and Rawlings (2022d), as shown in Figure 1. The goal is to control the height of liquid in each tank (x_1, x_2) via the inlet flow rate to tank 1 and outlet flow rate from tank 2 (u_1, u_2). Tank 1 drains into tank 2 by gravity at a rate proportional to the height of tank 1. This proportionality constant is subject to uncertainty with $w \in \mathbb{W} := \{-0.3, 0, 0.3\}$. We consider the probability distribution $\Pr(w = -0.3) = \Pr(w = 0.3) = 0.35$ and $\Pr(w = 0) = 0.3$. The differential equations (in deviation variables) are

$$\frac{dx_1}{dt} = -(1+w)x_1 + u_1 - w$$
$$\frac{dx_2}{dt} = (1+w)x_1 - u_2 + w$$

Since the support for the disturbance is finite, we can discretize this differential equation (assuming a zero-order hold on the input and disturbance with a time step $\Delta = 1$) and solve the SMPC problem exactly by considering all possible disturbance trajectories.

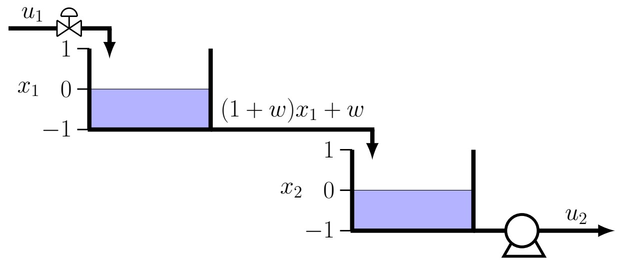

Figure 1: Taken from McAllister and Rawlings (2022d). Two tanks with gravity driven flow between tank 1 and tank 2.

We consider the input constraints $u_1, u_2 \in [-1, 1]$ and stage cost $\ell(x,u) = x'Qx + u'Ru$ with $Q = \text{diag}([0.1, 20])$ and $R = \text{diag}([0.1, 0.1])$. Note that these penalties strongly discourage any deviations in x_2, but nonetheless satisfy all the usual requirements for nominal MPC and SMPC. We use the LQR cost P and gain K for the nominal $(w = 0)$ unconstrained system to define the terminal cost $V_f(x) = x'Px$ and control law parameterization $\pi(x, v) := Kx + v$. We define the terminal constraint as $\mathbb{X}_f := \{x : |x_1| \leq 0.4,\ |x_2| \leq 0.4\}$ and verify that this formulation satisfies Assumptions 3.1 and 3.4 to 3.7 for SMPC with the terminal control law $\kappa_f(x) = Kx$.

For nominal MPC, we can establish that these same choices of stage cost, terminal cost, and terminal constraint render the origin asymptotically stable for the nominal closed-loop system and RAS for small disturbances, such as the disturbance $w \in \mathbb{W}$ considered in this example (Grimm et al., 2004). SMPC, however, does *not* render the origin asymptotically stable for the nominal closed-loop system. In Figure 2, we plot the closed-loop trajectory for SMPC subject to a nominal realization of the disturbance, i.e., $\mathbf{w}_k = \mathbf{0}$. Despite the fact that no disturbance occurs and the system is initialized at the setpoint $x(0) = 0$, the SMPC controller drives x_1 away from this setpoint. The value of x_1 converges to a different steady state at $x_1 \approx -0.6$ that is not even within the terminal set $\mathbb{X}_f$ defined for the SMPC controller. Thus, the SMPC controller does not render the origin (or the terminal set) RAS or nominally asymptotically stable.

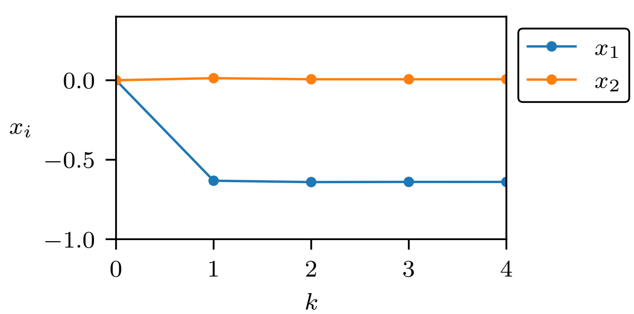

Figure 2: Taken from McAllister and Rawlings (2022d). The closed-loop trajectory for SMPC subject to a nominal realization of the disturbances, i.e., $\mathbf{w}_k = \mathbf{0}$.

The benefit of SMPC for this example is that the expected value of the stage cost evaluated for closed-loop state and input trajectory, i.e., $\mathbb{E}[\ell(x(k), u(k))]$, is lower for SMPC than nominal MPC. By lowering the value of x_1, SMPC reduces the effect of the disturbance on x_2. Since we have assigned a larger cost to deviations in x_2 than x_1, this approach reduces the expected stage cost for the system subject to this disturbance distribution. If this stage cost is closely related to the economic performance of the process and the disturbance distribution is well characterized, then SMPC may be preferred to nominal MPC despite the lack of nominal asymptotic stability.

3.6 Distributional robustness

We would like to briefly summarize some recent results on how robust stochastic MPC is to errors in the assumed probability distribution of the disturbances.

To investigate distributional robustness of the closed-loop system, we allow the distribution of the disturbance in the MPC model to be *different* from the distribution of the disturbance in the plant. To distinguish the model and the plant, we use $\hat{w} \in \hat{\mathbb{W}}$ with probability distribution $\hat{\mu}$ as the model disturbance, and $w \in \mathbb{W}$ with probability distribution μ as the plant disturbance. To measure the distance between these two disturbance models, we use the Hausdorff set distance $d_H(\mathbb{W}, \hat{\mathbb{W}})$ to measure the difference between the support sets $\mathbb{W}$ and $\hat{\mathbb{W}}$, and we use the Wasserstein metric to measure the distance between the probability distributions, denoted $W(\mu, \hat{\mu})$. McAllister and Rawlings (2022b,c) provide the precise definition of this metric and describe how it is computed given the two probability distributions. Note that the Wasserstein metric satisfies all the axioms of a distance on the set $\mathcal{M}(\mathbb{W})$, i.e., the metric $W(\mu, \hat{\mu})$ is symmetric, nonnegative, satisfies the triangle inequality, and $W(\mu, \hat{\mu}) = 0$ if and only if $\mu = \hat{\mu}$ for all $\mu, \hat{\mu} \in \mathcal{M}(\mathbb{W})$. Similarly, we have that the Hausdorff set distance satisfies all the axioms of a distance for nonempty, compact subsets of $\mathbb{R}^q$.

Definition 3.14 (DRASiE)**.** *The origin of the system $x^+ = f(x, \kappa_{\hat{\mu}}(x), w)$, $w \in \mathbb{W}$ is distributionally robustly asymptotically stable in expectation (DRASiE) in the RPI set $\mathcal{X}$ if there exist $\beta(\cdot) \in \mathcal{KL}$ and $\gamma_1(\cdot), \gamma_2(\cdot) \in \mathcal{K}$ such that*

$$\mathbb{E}\left[|\phi_{\hat{\mu}}(k; x, \mathbf{w}_k)|\right] \leq \beta(|x|, k) + \gamma_1(\hat{\mathbb{E}}[|\hat{w}|]) + \gamma_2(W(\mu, \hat{\mu})) \tag{9}$$

for all $x \in \mathcal{X}$, $\hat{\mu} \in \mathcal{M}(\hat{\mathbb{W}})$, $\mu \in \mathcal{M}(\mathbb{W})$, and $k \in \mathbb{I}_{\geq 0}$.

The first part of the upper bound in (9) is a $\mathcal{KL}$ function that ensures the effect of the initial condition $x \in \mathcal{X}$ (asymptotically) vanishes as $k \to \infty$. The second function $\gamma_1(\hat{\mathbb{E}}[|\hat{w}|])$ accounts for the persistent effect of the modeled disturbance ($\hat{w}$) in the control law design and the ideal system with $\mu = \hat{\mu}$. Note that if $\hat{\mathbb{E}}[\hat{w}] = 0$, we can replace $\hat{\mathbb{E}}[|\hat{w}|]$ with the upper bound $\text{tr}(\hat{\Sigma})^{1/2}$. The third function $\gamma_2(W(\mu, \hat{\mu}))$ accounts for the discrepancy between the disturbance distribution model $\hat{\mu}$, used in the SMPC optimization problem, and the true disturbance distribution μ. If $\mu = \hat{\mu}$, then $\gamma_2(W(\mu, \hat{\mu})) = 0$ and we recover the usual bound for idealized SMPC analysis. The most significant consequence of this result is that the effect of arbitrarily small errors between $\hat{\mu}$ and μ produce similarly small deviations from the closed-loop bound derived for idealized SMPC analysis.

Under suitable assumptions, SMPC is distribituionally robust (McAllister and Rawlings, 2022b,c). Specifically, we can establish that for sufficiently small errors in the disturbance support ($d_H(\mathbb{W}, \hat{\mathbb{W}}) \leq \delta$ for some $\delta > 0$), SMPC renders the closed-loop stochastic system DRASiE. Note that these errors include disturbances that are incorrectly modeled, unmodeled, or intentionally mismodeled via sampling-based approximations of the stochastic optimization problem. Thus, feedback is a crucial component of SMPC algorithms for essentially the same reason as for nominal MPC algorithms: To address the inevitable discrepancy between the stochastic or deterministic model used in the optimization problem and the plant.

4 Conclusion

We summarize the expectations and outcomes in comparing nominal MPC to stochastic MPC.

1. We expect nominal MPC to bring the system to setpoint from different initial conditions, because the optimal control problem is designed for that.

   This expectation is met.

2. We expect nominal MPC to be robust to small deterministic disturbances, although we did *not* design for that.

   This expectation is met.

3. We expect nominal MPC to be robust to small stochastic disturbances, although we did *not* design for that either.

   This expectation is met.

4. We expect stochastic MPC to handle small stochastic disturbances, because we *did* design for that.

   This expectation is met.

5. But did we expect stochastic MPC to lose nominal deterministic stability? Probably not, until we see the simulation example that demonstrates this loss (McAllister and Rawlings, 2022d).

   This is a surprise.

6. We expect stochastic MPC to be robust to small errors in the stochastic model that we are using although we did *not* design for that. This is the distributional robustness question.

   This expectation is met.

7. But since nominal MPC is a stochastic MPC with an unusual/trivial choice of stochastic model (zero), this distributional robustness also applies to nominal MPC.

   This is a surprise.

8. Open question. How robust are either of these MPC designs (nominal/stochastic) to deterministic model error, i.e., errors in $f(\cdot)$ in (1).

We summarize these conclusions in the following table.

	MPC	SMPC	SMPC always outperforms MPC?
RPI set	Yes	Yes	No
RAS	Yes	No	No
RASiE	Yes	Yes	No
ℓ-RASiE	Yes	Yes	Yes[5]
DRASiE	Yes	Yes	Unknown

Wonham (1969) concludes his paper on optimal stochastic control from 50 years ago with the following remarks.

"Since the mathematical model is usually greatly complicated by explicitly including stochastic features, it is always to be asked whether the extra effort is worthwhile, i.e., whether it leads to a control markedly superior in performance to one designed on the assumption that stochastic disturbances are absent. In the case of feedback controls the general conclusion is that only marginal improvements can be obtained unless the disturbance level is very high, in which case the fractional improvement from stochastic optimization may be large, but the system is useless anyway. That is, efforts to counter disturbances by simply warping the velocity field in state space are generally misplaced. For this reason I do not think there is much point in trying to develop low noise perturbation formalism for feedback controls."

So how well does this rather pessimistic conclusion hold up when applied to stochastic MPC? We think that today this conclusion warrants reconsideration. The control field is addressing more general classes of control problems today compared to fifty years ago, and the usefulness of stochastic control depends on the *type* of control problem that one is addressing. If one is interested primarily in *tracking* performance, Wonhanm's pessimism remains entirely justified. In fact, a recent surprise is that stochastic MPC does not provide robust asymptotic stability of the steady-state target in the *absence* of disturbances.

If one is interested primarily in *economic* performance, however, an *accurate* model of the randomness can pay significant dividends (Kumar et al., 2018). In these economic control problems, the stability of the steady state is not a goal of the control system design and therefore not a major concern.

References

Allan, D. A., C. N. Bates, M. J. Risbeck, and J. B. Rawlings (2017). On the inherent robustness of optimal and suboptimal nonlinear MPC. *Sys. Cont. Let. 106*, 68 – 78.

Bellman, R. E. and S. E. Dreyfus (1962). *Applied Dynamic Programming*. Princeton, New Jersey: Princeton University Press.

Bemporad, A. and M. Morari (1999). Control of systems integrating logic, dynamics, and constraints. *Automatica 35*, 407–427.

Bequette, B. W. (2019). Process control practice and education: Past, present and future. *Comput. Chem. Eng. 128*, 538–556.

Bryson, A. E. and Y. Ho (1975). *Applied Optimal Control*. New York: Hemisphere Publishing.

Cannon, M., B. Kouvaritakis, and X. Wu (2009). Probabilistic constrained mpc for multiplicative and additive stochastic uncertainty. *IEEE Trans. Auto. Cont. 54*(7), 1626–1632.

Chatterjee, D. and J. Lygeros (2014). On stability and performance of stochastic predictive control techniques. *IEEE Trans. Auto. Cont. 60*(2), 509–514.

Cuthrell, J. E., D. E. Rivera, W. J. Schmidt, and J. A. Vegeais (1990). Solution to the Shell standard control problem. In D. M. Prett, C. E. Garcia, and B. L. Ramaker (Eds.), *The Second Shell Process Control Workshop*, Stoneham, MA, pp. 27–58. Butterworths.

Grimm, G., M. J. Messina, S. E. Tuna, and A. R. Teel (2004). Examples when nonlinear model predictive control is nonrobust. *Automatica 40*, 1729–1738.

Kouvaritakis, B., M. Cannon, S. V. Raković, and Q. Cheng (2010). Explicit use of probabilistic distributions in linear predictive control. *Automatica 46*(10), 1719 – 1724.

Kumar, R., M. J. Wenzel, M. J. Ellis, M. N. ElBsat, K. H. Drees, and V. M. Zavala (2018, Jul). A stochastic model predictive control framework for stationary battery systems. *IEEE Trans. Power Systems 33*(4), 4397–4406.

Lee, E. B. and L. Markus (1967). *Foundations of Optimal Control Theory*. New York: John Wiley and Sons.

Lorenzen, M., F. Dabbene, R. Tempo, and F. Allgöwer (2016). Constraint-tightening and stability in stochastic model predictive control. *IEEE Trans. Auto. Cont. 62*(7), 3165–3177.

Maravelias, C. T. (2021). *Chemical Production Scheduling: Mixed-integer Programming Models and Methods*. New York: Cambridge University Press.

Mayne, D. Q. and P. Falugi (2019). Stabilizing conditions for model predictive control. *Int. J. Robust and Nonlinear Control 29*(4), 894–903.

McAllister, R. D. and J. B. Rawlings (2021). Inherent stochastic robustness of model predictive control to large and infrequent disturbances. *IEEE Trans. Auto. Cont.*.

McAllister, R. D. and J. B. Rawlings (2022a). Nonlinear stochastic model predictive control: Existence, measurability, and stochastic asymptotic stability. *IEEE Trans. Auto. Cont.*.

[5] No counter example yet.

McAllister, R. D. and J. B. Rawlings (2022b). On the inherent distributional robustness of stochastic and nominal model predictive control. IEEE Trans. Auto. Cont. Submitted 6/3/22, Conditionally accepted 10/3/22.

McAllister, R. D. and J. B. Rawlings (2022c). On using feedback control to contend with nature's randomness. Ind. Eng. Chem. Res. Submitted 8/19/22, Conditionally accepted 9/23/22, Revised 10/14/2022.

McAllister, R. D. and J. B. Rawlings (2022d). The stochastic robustness of nominal and stochastic model predictive control. IEEE Trans. Auto. Cont. Submitted 3/29/22, Revised 9/30/22.

McAllister, R. D., J. B. Rawlings, and C. T. Maravelias (2022). The inherent robustness of closed-loop scheduling. *Comput. Chem. Eng. 159*, 107678.

Pannocchia, G., J. B. Rawlings, and S. J. Wright (2011). Conditions under which suboptimal nonlinear MPC is inherently robust. *Sys. Cont. Let. 60*, 747–755.

Pontryagin, L. S., V. G. Boltyanskiy, R. V. Gamkrelidze, and Y. F. Mishchenko (1961). *Matematicheskaya Teoriya Optimalńykh Protsessov (Mathematical Theory of Optimal Processes)*. Moscow, Russia: State Publishing House of Mathematical Physics Literature.

Rawlings, J. B., D. Q. Mayne, and M. M. Diehl (2020). *Model Predictive Control: Theory, Design, and Computation* (2nd, Paperback ed.). Santa Barbara, CA: Nob Hill Publishing. 770 pages, ISBN 978-0-9759377-5-4.

Rawlings, J. B. and M. J. Risbeck (2017). Model predictive control with discrete actuators: Theory and application. *Automatica 78*, 258–265.

Risbeck, M. J., C. T. Maravelias, and J. B. Rawlings (2019). Unification of closed-loop scheduling and control: State-space formulations, terminal constraints, and nominal theoretical properties. *Comput. Chem. Eng. 129*, 106496.

Scokaert, P. O. M. and J. B. Rawlings (1999, Aug). Feasibility issues in linear model predictive control. *AIChE J. 45*(8), 1649–1659.

Slupphaug, O., J. Vada, and B. A. Foss (1998). MPC in systems with continuous and discrete control inputs. *Mod. Ident. Control 19*(4), 175–184.

Subramanian, K., C. T. Maravelias, and J. B. Rawlings (2012, Dec). A state-space model for chemical production scheduling. *Comput. Chem. Eng. 47*, 97–110.

Wonham, W. M. (1969). Optimal stochastic control. *Automatica 5*, 113–118.

Yu, S., M. Reble, H. Chen, and F. Allgöwer (2014). Inherent robustness properties of quasi-infinite horizon nonlinear model predictive control. *Automatica 50*(9), 2269 – 2280.

Zheng, A. and M. Morari (1995). Control of linear unstable systems with constraints. In *American Control Conference, Seattle, Washington*, pp. 3704–3708.