

Vinnicombe Metric As A Closed-loop Nonlinearity Measure

Guan Tien Tan*,

Mihai Huzmezan†,

K. Ezra Kwok*

*Department of Chemical and Biological Engineering

†Department of Electrical and Computer Engineering

University of British Columbia, Vancouver, BC, Canada

Abstract

The focal point of this paper is to develop a measure for closed-loop nonlinearity. In this work, the Vinnicombe metric and the quasi-linear parameter varying representation of nonlinear systems are exploited for this purpose. It is expected that the proposed measure can serve as a decision making tool for control engineers when deciding whether a linear or a nonlinear control strategy should be employed to solve their control problems. A continuous stirred tank reactor (CSTR) simulation example is used to illustrate the proposed measure.

Keywords: Nonlinearity measure, Vinnicombe metric, quasi-LPV systems, \mathcal{H}_∞ loop-shaping, robust control.

1 Introduction

In model-based control design, it is often desirable to check the adequacy of a linear model before any attempts to use a nonlinear controller are made. To achieve this in a systematic manner, linearity tests are required to ascertain whether a nonlinear controller is really needed. Over the past few decades, various linearity tests were proposed. Among them, the statistical approach (Ramsey, 1969; Subba Rao and Gabr, 1980; Hinich, 1982; Brock et al., 1987; Hjellvik and Tjøstheim, 1995; Billings and Voon, 1983), the norm-bounded error approach (Nikolaou, 1993; Ogunnaike et al., 1993) and the geometrical approach (Guay et al., 1995) are often encountered in the literature.

In general, the statistical approach is normally based on the test of a hypothesis using statistical inference. For example, the regression error specification test (RESET) (Ramsey, 1969) checks the hypothesis of a signal-to-noise ratio admitting an asymptotic χ^2 distribution. On another hand, the Brock-Dechert-Sheinkman test (BDS) (Brock et al., 1987) exploits the concept of correlation dimension arising from chaos theory and statistical properties of an independent and identically distributed (i.i.d.) noise sequence. If the nonlinearity exists, the null hypothesis of the correlation dimension being an i.i.d. sequence is rejected.

The second class of nonlinearity tests involves measures of norm-bounded deviation of a nonlinear plant from its ideal linear counterpart in the vicinity of a particular operating point, see for example (Nikolaou, 1993). A slightly different approach, which also falls into this class, is proposed by (Ogunnaike et al., 1993). In this approach, the changes in the local gain are obtained over a specific operating region. However, this method is known to suffer from sensitivity to the choice of perturbation directions.

The third class consists of geometrical approaches proposed by (Guay, 1995; Guay et al., 1995, 1997a,b). Unlike the methods mentioned in the first two classes, which measure the open-loop nonlinearity, methods in this class are establishing measures for closed-loop nonlinearity, particularly

for systems with a unity feedback. It is noted that since feedback control is known to modify system's open-loop nonlinearity, the closed-loop nonlinearity measure is found to be more appropriate from the control design perspective. In (Guay et al., 1997b), differential geometry interpretation of the relative gain array was used to assess the degree of closed-loop nonlinearity for a given plant. Unfortunately, the accuracy of this method deteriorates due to process noise.

Since its debut, the Vinnicombe (or ν -gap) metric notion developed by (Vinnicombe, 1993) has attracted much attention particularly in robust control and system identification. In principle, the ν -gap metric measures the aperture of two closed Hilbert sub-spaces representing the bounded input-output pairs (i.e. the plant model graph) of two linear (possibly unbounded) operators. Together with a homotopy condition, the ν -gap metric gives the least conservative robust stability results (Vinnicombe, 1993). This means that if the ν -gap between two plants is big, then a controller that gives satisfactory robust stability for one plant will show poor robust stability or even destabilize the other plant. Like wise, if the ν -gap between two plants is small, then a controller which guarantees robust stability of one plant implies that it robustly stabilizes the other.

Recently, several attempts have been made to extend the idea of the ν -gap metric to nonlinear systems. For instance, in (Vinnicombe, 1998, 1999) a set of integral quadratic constraints (IQCs) (Megretski and Rantzer, 1997) was used to describe the system's nonlinearity. Then, conditions to which the linear controller fails to stabilize the nonlinear plant were established in terms of IQCs and the ν -gap metric. Other nonlinear extensions of the ν -gap metric include (Georgiou and Smith, 1997; Anderson et al., 1998; Vinnicombe, 1999; James et al., 2000; Anderson et al., 2002). In general, the distance between a nonlinear plant and its linear model can be obtained using the above measures. To do so, the gap between the graphs of a nonlinear plant and a defined linear model is computed. The computed gap is then compared against what a linear controller can handle. If the gap is small in the sense of ν -gap metric, then any linear controller designed for the identified linear plant is also claimed to maintain the closed-loop stability of the nonlinear plant. This also implies indirectly that the degree of closed-loop nonlinearity is manageable by the aforementioned linear controller. Of course any statements about performance would be still conservative. Unequivocally, ν -gap metric can be seen as a potential powerful tool to assess the degree of closed-loop nonlinearity indirectly. At the best knowledge of the authors, none of the metrics mentioned above are exploited to provide a reliable closed-loop nonlinearity measure. Hence, the ν -gap metric is used in a completely different context in this work.

A nonlinear model (or operator), which captures system's nonlinearity, is crucial to the success of the proposed algorithm. In this development, a Quasi-Linear Parameter Varying (quasi-LPV) (Shamma and Cloutier, 1993) is an appealing candidate owing to the following reasons: (i) plant's nonlinearity can be captured by selecting appropriate scheduling parameters; (ii) it is not a linearized version of the nonlinear plant, instead it is derived through a state transformation; (iii) a family of local linear models can be easily obtained by merely freezing the scheduling parameters. In this light, the focal point of this paper is to develop an indirect closed-loop nonlinearity measure by exploiting the ν -gap metric notion and the special structure of systems which admit a quasi-LPV transformation.

This paper is organized as follows. Section 2 gives a brief review on the quasi-LPV representation, the \mathcal{H}_∞ loop-shaping controller design procedure and the linear ν -gap metric notion. Next, a computational algorithm is presented in Section 3. In Section 4, an example involving the control of a continuous stirred tank reactor (CSTR) is used to illustrate the proposed measure in Section 4. Finally, some concluding remarks are drawn in Section 5.

2 Briefing on Quasi-LPV, Vinnicombe Metric and \mathcal{H}_∞ Loop-Shaping

2.1 Quasi-LPV Transformation

Any plant exhibiting output nonlinearity such as the one in Eq. (1)

$$\frac{d}{dt} \begin{bmatrix} y \\ x_r \end{bmatrix} = \phi(y) + \begin{bmatrix} \tilde{A}_{11} & \tilde{A}_{12} \\ \tilde{A}_{21} & \tilde{A}_{22} \end{bmatrix} \begin{bmatrix} y \\ x_r \end{bmatrix} + \begin{bmatrix} \tilde{B}_1(y) \\ \tilde{B}_2(y) \end{bmatrix} u, \quad (1)$$

can be easily recast into a quasi-LPV representation as shown in Eq. (2) provided that $x_{r,eq}$ is differentiable with respect to the scheduling parameter y :

$$\begin{aligned} \frac{d}{dt} \begin{bmatrix} y \\ x_r - x_{r,eq}(y) \end{bmatrix} &= \begin{bmatrix} 0 & \tilde{A}_{12}(y) \\ 0 & \tilde{A}_{22}(y) - \frac{dx_{r,eq}(y)}{dy} \tilde{A}_{12}(y) \end{bmatrix} \begin{bmatrix} y \\ x_r - x_{r,eq}(y) \end{bmatrix} \\ &+ \begin{bmatrix} \tilde{B}_1(y) \\ \tilde{B}_2(y) - \frac{dx_{r,eq}(y)}{dy} \tilde{B}_1(y) \end{bmatrix} (u - u_{eq}(y)) \end{aligned} \quad (2)$$

In the above, $x_{r,eq}(y)$ denotes a family of the equilibrium points obtained by setting the derivatives in Eq. (1) to zero. Note that for plants that do not exhibit output nonlinearity, the quasi-LPV representation approximates the actual plant up to the first order approximation of all other states except the scheduling state.

In order to use Eq. (2) for control purposes, the state dependent $u_{eq}(y)$ needs to be known. Any incorrect estimation of $u_{eq}(y)$ may jeopardize the robust property of the closed-loop system. To avoid this problem, an integrator at the plant input, which stores the trim input value $u_{eq}(y)$, can be added as suggested in Shamma and Cloutier (1993). As a consequence Eq. (2) can be rewritten as follows:

$$\begin{aligned} \frac{d}{dt} \begin{bmatrix} y \\ x_r - x_{r,eq}(y) \\ u - u_{eq}(y) \end{bmatrix} &= \begin{bmatrix} 0 & \tilde{A}_{12}(y) & \tilde{B}_1(y) \\ 0 & \tilde{A}_{22}(y) - \frac{dx_{r,eq}(y)}{dy} \tilde{A}_{12}(y) & \tilde{B}_2(y) - \frac{dx_{eq}(y)}{dy} \tilde{B}_1(y) \\ 0 & -\frac{du_{eq}(y)}{dy} \tilde{A}_{12}(y) & -\frac{du_{eq}(y)}{dy} \tilde{B}_1(y) \end{bmatrix} \begin{bmatrix} y \\ x_r - x_{r,eq}(y) \\ u - u_{eq}(y) \end{bmatrix} \\ &+ \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} v \end{aligned} \quad (3)$$

2.2 \mathcal{H}_∞ Loop-Shaping

Proposed by McFarlane and Glover (1992), the \mathcal{H}_∞ loop-shaping controller design method is based on the \mathcal{H}_∞ robust stabilization and classical loop-shaping technique. The \mathcal{H}_∞ loop-shaping consists of two major steps:

1. The open-loop plant is shaped using pre- and post-compensators to give a desired open-loop shape. Normally, it is desirable to shape the plant such that the maximum singular value frequency plot has a -20dB/decade slope at the crossover frequency.
2. Denoted by $P_s = W_1 P W_2$, the shaped plant is then robustly stabilized with respect to coprime factor uncertainty using a controller synthesis method based on an \mathcal{H}_∞ optimization.

It is noted that the \mathcal{H}_∞ norm of the closed-loop transfer function is minimized in the above \mathcal{H}_∞ robust stabilization synthesis. Denoted by $b_{P,C}$, the reciprocal of Eq. (4) is often called the generalized stability margin which has a close relationship with the ν -gap metric. Mathematically, the

generalized stability margin is defined as:

$$b_{P,C} \triangleq \begin{cases} \left\| \begin{bmatrix} P \\ I \end{bmatrix} (I - CP)^{-1} \begin{bmatrix} -C & I \end{bmatrix} \right\|_{\infty}^{-1} & \text{if } \begin{bmatrix} I & P \\ C & I \end{bmatrix}^{-1} \in \mathcal{H}_{\infty} \\ 0, & \text{otherwise.} \end{cases} \quad (4)$$

For a more detail treatment of the \mathcal{H}_{∞} loop-shaping, see (McFarlane and Glover, 1992).

2.3 Vinnicombe Metric

Given two (possibly unbounded) linear operators, the ν -gap metric is defined in terms of P_1 and P_2 as:

$$\delta_{\nu}(P_1, P_2) \triangleq \begin{cases} \|(I + P_2 P_2^*)^{-\frac{1}{2}} (P_1 - P_2) (I + P_1 P_1^*)^{-\frac{1}{2}}\|_{\infty}, & \text{if } \text{Index}(P_1, P_2) = 0 \\ 1, & \text{otherwise} \end{cases} \quad (5)$$

where, $P^* = P^T(-s)$ and $\text{Index}(P_1, P_2) \triangleq \eta(P_1, P_2^*) - \text{deg}(P_2)$. In the above, η and deg denote the number of open RHP poles and McMillan degree, respectively.

Note that the index can be determined from the “ A ” state matrix of a state-space realization of $[P_1, -P_2^*]$ which is given by:

$$A_{P_1, P_2^*} = \begin{bmatrix} -(A_2 - B_2 W D_1^T C_2)^T & C_2^T Y C_1 \\ B_1 W^T B_2^T & A_1 - B_1 D_2^T Y C_1 \end{bmatrix}, \quad (6)$$

where $W := (I + D_1^T D_2)^{-1}$, $Y := (I + D_1 D_2^T)^{-1}$ and (A_1, B_1, C_1, D_1) and (A_2, B_2, C_2, D_2) are state-space realizations of P_1 and P_2 , respectively, see (Vinnicombe, 1999, 2001). It follows that

$$\text{Index}(P_1, P_2) = 0 \iff A_{P_1, P_2^*} \text{ has precisely } \text{deg}(P_2) \text{ eigenvalues with a positive real part.}$$

Together with the $b_{P,C}$, the following theorem is one of the main results arising from the ν -gap metric theory.

Theorem 1. *Given a nominal plant $P_1 \in \mathcal{P}^{p \times q}$, a compensator $C \in \mathcal{P}^{q \times p}$ and a constant β , then: $[P_2, C]$ is stable for all plants $P_2 \in \mathcal{P}^{p \times q}$ satisfying $\delta_{\nu}(P_1, P_2) \leq \beta$ iff $b_{P_1, C} > \beta$.*

For a quasi-LPV system $\mathcal{P}^{p \times q}(y)$ with a gridding space Ω of a scheduling parameter y , the above theorem can be restated as follows:

Corollary 1. *Given a nominal plant $P_i(y) \in \mathcal{P}^{p \times q}(y)$ obtained by freezing the scheduling parameter $y \in \Omega$ at operating point y_i , a compensator $C \in \mathcal{P}^{q \times p}$ and a constant γ , then: $[P_j(y), C]$ is stable for all plants $P_j(y) \in \mathcal{P}^{p \times q}(y)$, $\forall y \in \Omega$ satisfying $\delta_{\nu}(P_i(y), P_j(y)) \leq \gamma$ iff $b_{P_i(y), C} > \gamma$.*

3 Vinnicombe Metric As A Nonlinearity Measure

In this section, a nonlinear measure algorithm based on a ν -gap metric and a quasi-LPV representation is proposed. The basic idea behind the proposed approach is that if the uncertainty ball induced by plant’s nonlinearity is larger than what the best available linear controller can cope with, then the degree of closed-loop nonlinearity is not manageable by the designed linear controller. This means that the designed controller stabilizes the plant at some operating points where the nonlinearity is mild and it destabilizes the plant at other points where the effect of the nonlinearity becomes prominent. Under such a condition, the control engineer might want to consider a nonlinear controller design approach to address the control problems. Conversely, if the aforementioned uncertainty ball

is small and within the tolerance of the designed linear controller, then this controller should be sufficient to tackle the stabilization problem.

However, to implement this idea, one needs to define how large the uncertainty ball really is and how best the nominal model is. In this approach, the ν -gap metric is used to quantify the size of an uncertainty ball and also to justify what is meant by the best nominal model. A computational algorithm summarizing the proposed measure can be defined as follows.

3.1 The Computational Algorithm

1. Recast the nonlinear system into a quasi-LPV representation.
2. Grid the scheduling parameter space. A set of linear models is then obtained by simply freezing the scheduling parameter.
3. For each model at operating point $y_i \in \Omega$, the ν -gaps to all other models are obtained (i.e. $\delta_i = \{\delta_\nu(P_i(y), P_j(y)), \forall P_j(y) \in \mathcal{P}^{p \times q}(y), \forall y \in \Omega\}$).
4. Denoted by $P_0(y)$, the best nominal model for closed-loop control is the one that has the smallest ∞ -norm δ^* in δ_i , $\forall i$.
5. Apply pre- and post-compensators to the best nominal model $P_0(y)$ (i.e. $P_s(y) = W_1 P_0(y) W_2$) such that the maximum singular value of the shaped plant has a desired loop-shape.
6. Since the ν -gap metric is sensitive to scaling, the same compensators obtained from step 5 are applied to all plants in $\mathcal{P}^{p \times q}(y)$. Repeat steps 3 to 5 until the weights W_1 and W_2 and the choice of the best nominal model converge.
7. Design a robust linear controller using the \mathcal{H}_∞ loop-shaping for $P_s(y)$ and compute $b_{PC, \max}$, the maximum uncertainty ball that the resulting linear controller can tolerate.
8. If the $b_{PC, \max}$ is small ($b_{PC, \max} < 0.25$), go to step 5. (This often indicates that the chosen loop shape is incompatible with robust stability requirements – see McFarlane and Glover (1990)).
9. Find the farthest point $P'(y) \in \{\mathcal{P}^{p \times q}(y) \cap y \in \Omega\}$ (in the ν gap metric sense) in the polytope centered at $P_0(y)$. The ν -gap between $P_0(y)$ and $P'(y)$ is denoted by δ' .
10. By Corollary 1, it is obvious that if the $b_{PC, \max}$ is greater than δ' , the nonlinearity is manageable by the designed linear controller.
11. If the $b_{PC, \max} < \delta'$, the nonlinearity is larger than what the linear controller can cope with and hence a nonlinear controller becomes mandatory for a stable closed-loop.

4 A Continuous Stirred Tank Reactor Example

A schematic diagram of a CSTR is depicted in Figure 1. Consider an irreversible, first-order, exothermic reaction from component A to component B that occurs in the reactor. Assuming constant liquid volume, the following nonlinear ordinary differential equations describe the CSTR process dynamics (Henson and Seborg, 1997).

$$\begin{aligned} \frac{dC_A}{dt} &= \frac{q}{V}(C_{Af} - C_A) - k_o \exp\left(-\frac{E}{RT}\right)C_A \\ \frac{dT}{dt} &= \frac{q}{V}(T_f - T) + k_2 C_A + k_1(T_c - T) \end{aligned} \quad (7)$$

with $k_1 = \frac{UA}{V\rho C_p}$ and $k_2 = \frac{(-\Delta H)}{\rho C_p} k_o \exp\left(-\frac{E}{RT}\right)$. Where C_A , T , and T_c represent reactor effluent concentration of component A , reactor temperature, and coolant temperature, respectively. The control objective is to provide a good control of C_A in the temperature range from $T = 300$ K to 373 K by manipulating T_c .

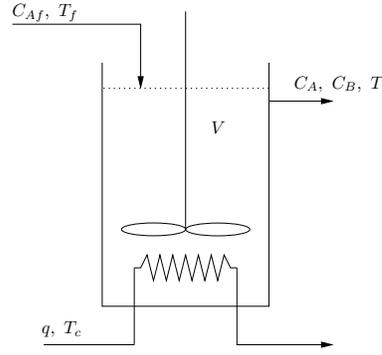


Figure 1: Continuous stirred tank reactor

To begin, Eq. (7) is first recast into a quasi-LPV representation via a state transformation. In this example, the reactor temperature is chosen as scheduling parameter. The resulting quasi-LPV model is

$$\frac{d}{dt} \begin{bmatrix} T \\ C_A - C_{A,eq} \\ T_c - T_{c,eq} \end{bmatrix} = \begin{bmatrix} 0 & k_2 & -\frac{k_1}{dT} \\ 0 & -\left(\frac{q}{V} + k_0 \exp\left(-\frac{E}{RT}\right)\right) - \frac{dC_{A,eq}}{dT} k_2 & -\frac{dC_{A,eq}}{dT} k_1 \\ 0 & -\frac{dT_{c,eq}}{dT} k_2 & -\frac{dT_{c,eq}}{dT} k_1 \end{bmatrix} \begin{bmatrix} T \\ C_A - C_{A,eq} \\ T_c - T_{c,eq} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} v \quad (8)$$

By using a 50 scheduling parameter grid and employing the computational algorithm proposed in Section 3, the best model $P_0(y)$ is the one corresponding to $T = 341$ K. Based on this nominal model, a pre-compensator $W_1 = 10 \frac{30s+400}{s+50}$ and a unity post-compensator are applied to $P_0(y)$ such that the closed-loop bandwidth is 20 rad/s. As required by the \mathcal{H}_∞ loop-shaping procedure, the slope of the loop-shape at the crossover frequency is -20dB/decade as in Figure 2. Since it is known that

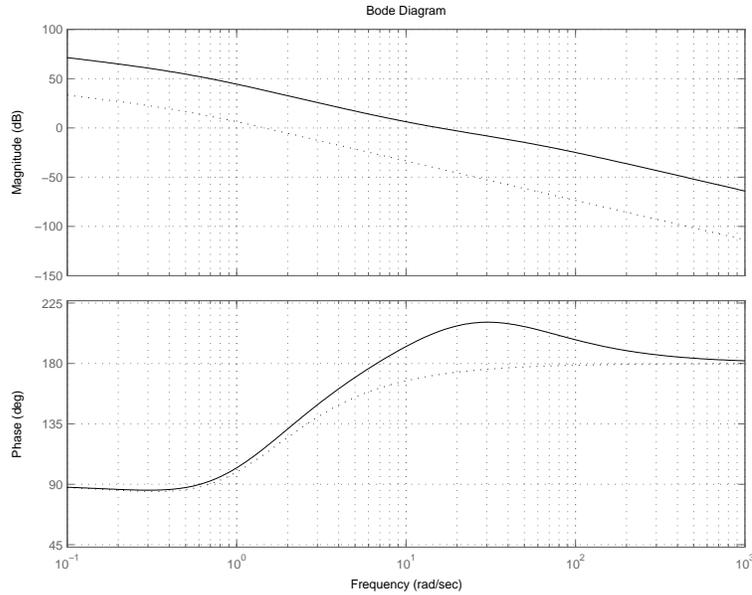


Figure 2: Unshaped (dotted) and shaped (solid) loop gains

the ν -gap metric is sensitive to scaling, it is desirable to check how does these compensators affect the choice of the nominal design model. Figures 3 and 4 show the ν -gap between a chosen model at temperature T_i and a model at temperature $T_j \forall T_i, T_i \in [300 \ 373]\text{K}$. A significant reduction in the ν -gap is observed when the compensators are applied to the quasi-LPV plant at all gridding points. However, it is also interesting to note that the best model remains the same. In addition, the model $P'(y)$ that was located at the farthest point (in the ν -gap sense) along the scheduling parameter trajectory also remains unchanged. In this case, that model is the one at $T = 300$ K.

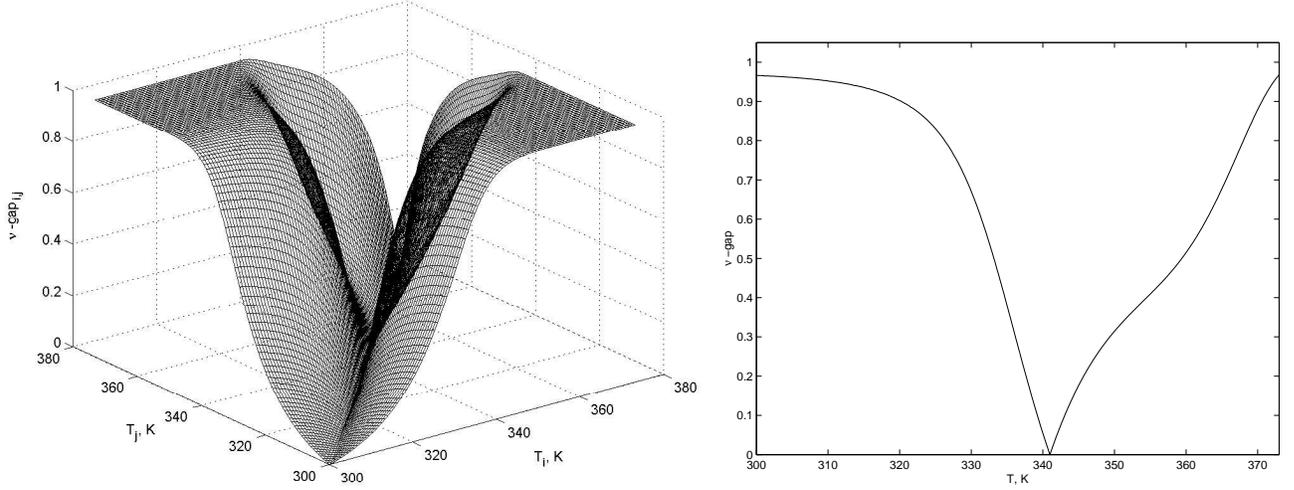


Figure 3: Left: Unshaped ν -gap over all gridding space. Right: Unshaped ν -gap for nominal model at $T=341$ K

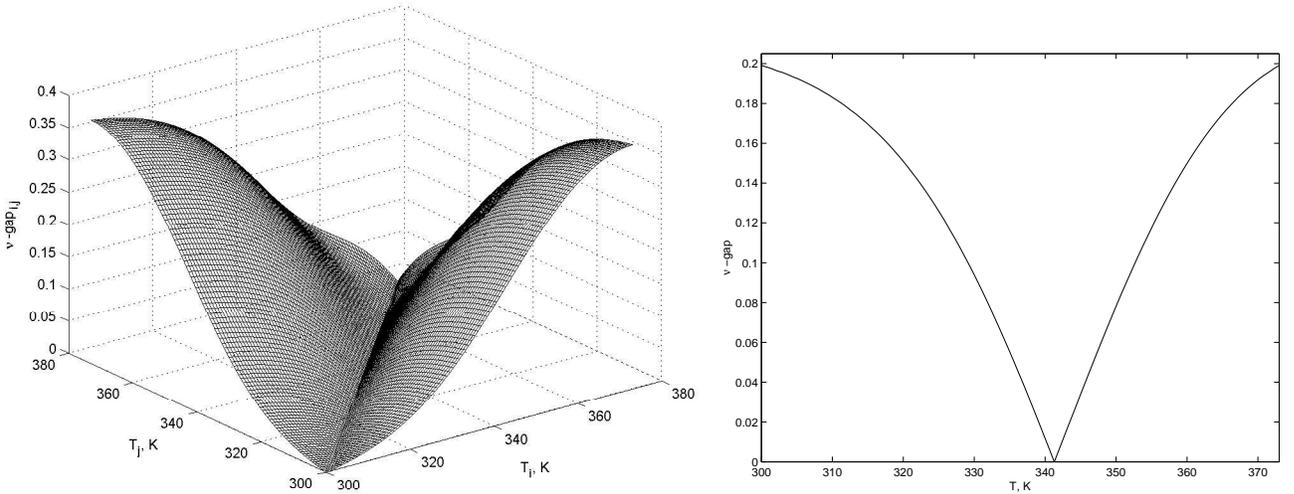


Figure 4: Left: Shaped ν -gap over all gridding space. Right: Shaped ν -gap for nominal model at $T=341$ K

At first glance, Figure 3 suggests that the ν -gap between $P_0(y)$ and $P'(y)$ is 0.9666. This implies that the uncertainty induced by the closed-loop nonlinearity is very large and may not be manageable by a single linear controller. As shown in Figure 5, the servo responses of the CSTR under a unity feedback are unacceptable (i.e. the process is at the brink of instability). Note that this observation is consistent with the result obtained from the ν -gap metric calculations.

However, as one can see in Figure 4, the aforementioned ν -gap was reduced to 0.1964 after applying the appropriate compensators. Based on the best model and the loop-shape given in

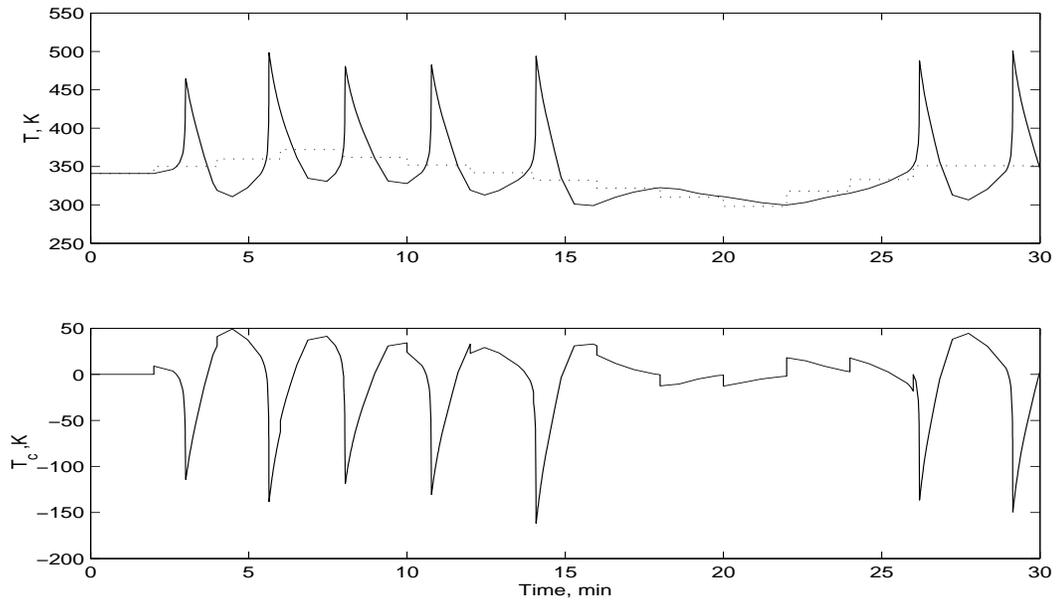


Figure 5: Upper: Servo responses of the CSTR under unity feedback. T (solid), setpoint (dotted). Lower: The corresponding system command, T_c (solid)

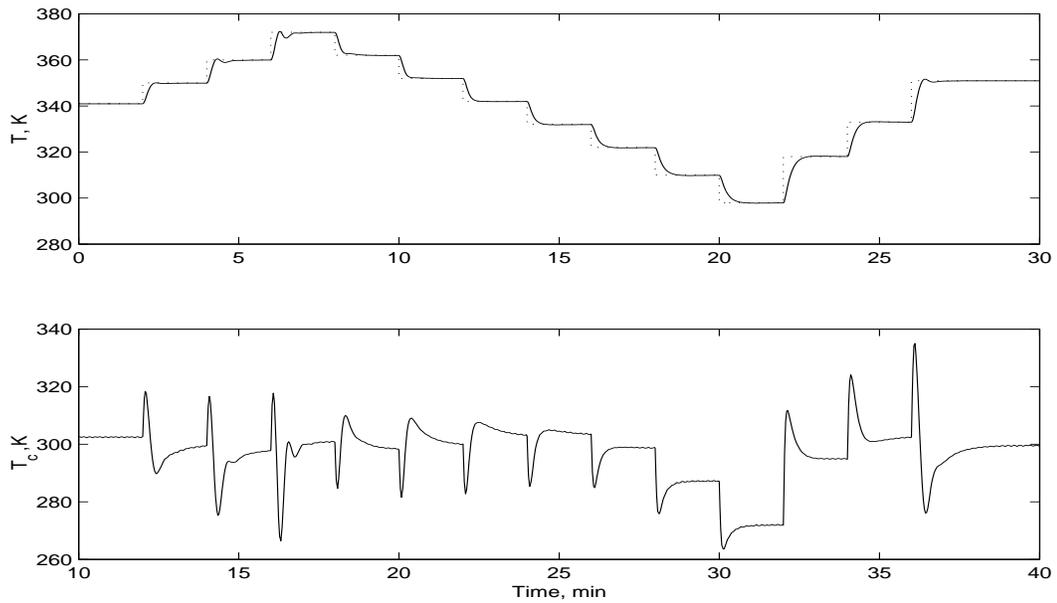


Figure 6: Upper: Servo responses of the CSTR under a linear \mathcal{H}_∞ controller. T (solid), setpoint (dotted). Lower: The corresponding system command, T_c (solid)

Figure 2, a linear \mathcal{H}_∞ controller is obtained via the \mathcal{H}_∞ loop-shaping design method discussed in Section 2. For this design, the $b_{PC,\max}$ is found to be 0.4440 which is equivalent to 44% coprime factor uncertainty. This suggests that the designed controller is sufficient to cope with the closed-loop nonlinearity when the plant is pre- and post-compensated. Simulation results, as shown in Figure 6, confirm this claim.

5 Conclusion

An indirect closed-loop nonlinearity measure using the ν -gap metric and the quasi-LPV representation is proposed. The contribution of this work is two-fold. Firstly, it acts as an effective decision making tool for the control engineers when they are faced with the problem of deciding whether to stick to the linear control strategy or use a nonlinear control approach in solving their problems. Secondly, for a certain class of nonlinear systems, the proposed measure can be used as a way to design compensators which reduce the closed-loop nonlinearity. However, a systematic approach to closed-loop nonlinearity reduction needs a more indepth study. As an alternative to the gridding of the scheduling parameter space, a coprime factorization of the quasi-LPV system can produce promising results. This is the path on which our research is moving onto. Dealing with non-differentiable nonlinearity such as hysteresis can be difficult. The proposed method can be extended to handle this type of nonlinearity by embedding such a nonlinearity using an IQC followed by a quasi-LPV transformation of the remaining model. This will be an extension of the work by (Vinnicombe, 1998).

References

- Anderson, B., Brinsmead, T., and Bruyne, F. D. (2002). The Vinnicombe metric for nonlinear operators. *IEEE Trans. Automat. Contr.*, 47:1450–1465.
- Anderson, B., James, M., and Limebeer, D. (1998). Robust stabilization of nonlinear systems via normalized coprime factor representations. *Automatica*, 34:1593–1599.
- Billings, S. and Voon, W. (1983). Structure detection and model validity tests in the identification of nonlinear systems. *IEEE Proceedings, Part D*, 130:193–199.
- Brock, W., Dechert, W., and Scheinkman, J. (1987). A test for independence based on the correlation dimension. Technical report, University of Wisconsin, Social Systems Research Institute, Madison.
- Georgiou, T. T. and Smith, M. C. (1997). Robustness analysis of nonlinear feedback systems: an input-output approach. *IEEE Trans. Automat. Contr.*, 42:1200–1221.
- Guay, M. (1995). Curvature measures for multiresponse regression models. *Biometrika*, 82:411–417.
- Guay, M., McLellan, P., and Bacon, D. (1995). Measurement of nonlinearity in chemical process control systems: The steady-state map. *Can. J. Chem. Engng.*, 73:868–882.
- Guay, M., McLellan, P., and Bacon, D. (1997a). Measurement of dynamic nonlinearity in chemical process control. In *Proceedings of ADCHEM 97*, Banff, Alberta, Canada.
- Guay, M., McLellan, P., and Bacon, D. (1997b). On a measure of interaction for nonlinear chemical processes. *AIChE. J.*, 43:2261–2278.
- Henson, M. and Seborg, D., editors (1997). *Nonlinear Process Control*. Prentice Hall, Inc., New Jersey.

- Hinich, M. (1982). Testing for gaussianity and linearity of a stationary time series. *Journal of Time Series Analysis*, 3:169–176.
- Hjellvik, V. and Tjøstheim, D. (1995). Nonparametric tests of linearity for time series. *Biometrika*, 82(2):351–368.
- James, M., Smith, M., and Vinnicombe, G. (2000). Gap metrics, representations, and nonlinear robust stability. In *Proc. of the 39th IEEE Conf. on Decision and Control*, pages 2936–2941.
- McFarlane, D. and Glover, K. (1990). *Robust Controller Design Using Normalized Coprime Factor Plant Descriptions*. Vol. 138 of Lecture Notes in Control and Information Sciences. Springer-Verlag, Berlin.
- McFarlane, D. and Glover, K. (1992). A loop shaping design procedure using \mathcal{H}_∞ synthesis. *IEEE Trans. on Automatic Control*, 37:759–769.
- Megretski, A. and Rantzer, A. (1997). System analysis via integral quadratic constraints. *IEEE Trans. Automat. Contr.*, 42:819–830.
- Nikolaou, M. (1993). When is nonlinear dynamic modelling necessary? In *Proc. of the American Control Conf.*, San Francisco, CA.
- Ogunnaike, B., Pearson, R., and Doyle, III, F. (1993). Chemical process characterization: With applications in the rational selection of control strategies. In *Proc. Europ. Cont. Conf.*, pages 1067–1071, Groningen, The Netherlands.
- Ramsey, J. (1969). Test for specification errors in classical linear least squares regression analysis. *Journal of the Royal Statistical Society B*, 31(350-371).
- Shamma, J. and Cloutier, J. (1993). Gain-scheduled missile autopilot design using linear parameter varying transformation. *Journal of Guidance, Control and Dynamics*, 16:256–263.
- Subba Rao, T. and Gabr, M. (1980). A test for linearity of stationary time series. *Journal of Time Series Analysis*, 1:145–158.
- Vinnicombe, G. (1993). Frequency domain uncertainty and the graph topology. *IEEE Trans. Automat. Contr.*, 38:1371–1383.
- Vinnicombe, G. (1998). On IQCs and the ν -gap metric. In *Proc. of the 37th IEEE Conf. on Decision and Control*, pages 1199–1200.
- Vinnicombe, G. (1999). A ν -gap distance for uncertain and nonlinear systems. In *Proc. of the 38th IEEE Conf. on Decision and Control*, pages 2557–2562.
- Vinnicombe, G. (2001). *Uncertainty and Feedback – \mathcal{H}_∞ loop-shaping and the ν -gap metric*. Imperial College Press, London, UK.