Proceedings of the European Control Conference 2007
Kos, Greece, July 2-5, 2007

Topical Problems of Adaptive Control

Brian D. O. Anderson

***Abstract*— Adaptive control is an adolescent technology, full of promise, given to over-enthusiasm, and sometimes short on experience. Many theoretical problems wait to be resolved before practitioners can use adaptive control with confidence. This paper summarizes several problems.**

***Index Terms*— Adaptive system, robust adaptive control**

I. The Problem

How many adaptive control textbooks contain theorems like the following? 'Under conditions X, Y, and Z, all signals in the closed loop are bounded, and as time goes to infinity, all becomes right in the universe.'

Such statements inflame tensions in the theory-practice gap, because they suppress recognition of very serious problems that continue to bedevil adaptive control today, as well as some problems that are now understood, like the bursting phenomenon when a plant does not have a persistently exciting input.

More generally, adaptive control has been oversold by theorists, through a failure to acknowledge potential shortcomings and continuing problems. Some are now mentioned.

II. The Problem Of Impractical Control Objectives

Suppose that a *known* plant with an associated linear-quadratic performance index delivers a controller and a closed loop with half a degree phase margin. Now suppose that the same plant is to be controlled, with the same performance index, but the designer does not know the plant, and so chooses to implement an adaptive control scheme. This task will end in tears. The reason is that the control objective being set is an impractical one. But the unpleasant truth is that the designer cannot know that fact when she/he is not given the plant. The problem is of course a generic one of adaptive control, not particular to adaptive LQ control.

How can one decide in the course of implementing an adaptive control algorithm on an unknown or partially known plant in order to achieve a prescribed objective whether or not the objective is practically feasible; and how can one gracefully abort the algorithm in the event of such a negative decision?

III. The Problem Of Transient Instability

If all is satisfactory when time goes to infinity, and all signals are bounded, what is the problem of transient instability? It is that one has one million amps of armature current at $t = 10$. Or more generally, that one connects at some point in the adaptive process a destabilizing controller, which although not remaining permanently in the loop, can do a lot of damage in the meantime. Safe adaptive control methods would generally make sure that a destabilizing controller is never connected, even temporarily.

We need a toolbox of techniques to predict instability/stability with a foreshadowed controller, under the circumstance that full knowledge of the plant is not available.

The toolbox is not completely empty; nor is it anywhere near full, and most people would regard the existing tools as on the primitive side.

IV. Stabilising A Suddenly Unstable Closed-Loop

The problem mentioned immediately above can also arise in the event that the plant has a stabilising controller connected, and then undergoes, perhaps without warning, some step change, due to a fault or major variation in a parameter such as a load. The previously stabilising controller becomes destabilising.

How can we deal as speedily as possible with a plant parameter step change inducing instability?

V. The Problem Of Changing Experimental Conditions In The Face Of Inexact Models

Adaptive control usually involves the generation of an explicit or implicit model of the plant. (An example of an implicit model is one where an estimate of the closed loop is known and the controller is also known). A good model is one where, *with the current controller*, the two closed loops defined using the real plant and using the model behave similarly—for a prescribed input class. The inputs might be white noise, steps, the exact inputs that have been used up to the present time etc.

Evidently, a model might be good for one set of inputs but not for others—especially if the first set of inputs is limited in its complexity. Suppose though that a model is a good model for a very wide class of inputs, certainly wide enough to include all of interest. Then it is important to recognise that if the current controller is changed to a new one, derived by calculations using the implicit or explicit model, the new controller-plant combination may behave nothing like what was expected. Put another way, the model ceases to be a good model for a different controller.

It is this sort of problem that gives rise to a need, or at least a preference, to change controllers by a small amount, or what is effectively the same thing, slowly. That much has

Research School of Information Sciences and Engineering, The Australian National University, Canberra ACT 0200, Australia. Work supported by an Australian Research Council Discovery Grant DP0342683.

National ICT Australia Ltd., Locked Bag 8001, Canberra, ACT 2601, Australia

ISBN: 978-960-89028-5-5

been recognised in recent years. Characterizing just what changes are small enough to be safe is quite another matter however. In the linear case, some results are available. But very few are available in the nonlinear case. 'Safety' here may refer to more than stability, but obviously includes stability. It can include requiring performance, however measured, not to deteriorate.

When and how can the algorithm assure that a proposed controller change is safe?

VI. THE ILLUSION OF MODEL-FREE ADAPTIVE CONTROL

Model free adaptive control sounds as if one knows nothing about the plant to begin with, and indeed never tries to construct a model. Actually, the seductive nomenclature is misleading. In every model-free approach, some structure for the allowed controllers is assumed, and if the set of allowed controllers does not include any that will be of practical utility, model-free adaptive control is pointless. Evidently, something must actually be known before one can define a set of allowed controllers. The deeper question is really, how much needs to be known (in order that an acceptable level of performance can be secured, during the learning phase and at the conclusion of learning)?

In model-free adaptive control, including those schemes based on multiple-model adaptive control, how much knowledge of the plant and inputs really is required to secure practical and effective adaptive control, avoiding for example the perils of bursting, transient instability or unsafe performance?

VII. A FINAL REMARK

The question asked at the end of the previous section is actually one that applies no matter what adaptive control algorithm is being used. And in particular contexts it has given rise to some debate. For example, there was a celebrated controversy built round the fact that a number of adaptive control algorithms relied for a convergence proof on a property of positive realness, which was unlikely ever to be achieved in practice. Notice that the question is about what will work in practice, and not what set of formal assumptions is the least set required for a certain theorem to be true.

If one is ever to successfully use adaptive control on a real plant, what a priori information is needed?