New Directions in Adaptive Control:

Multiple Model Adaptive Control and Safety[1]

Author: Brian D.O. Anderson

Research School of Information Sciences and Engineering

The Australian National University, Canberra, ACT

This talk is an attempt to marry two recent concepts in adaptive control, multiple model adaptive control and safe adaptive control. We describe each of these concepts in turn.

Multiple model adaptive control postulates that the unknown true plant either belongs to a prescribed finite set of plants, or is in some way close to at least one of the members of that set. The set might for example comprise a set of linear time-invariant plants and closeness might be reflected in terms of an additive or multiplicative uncertainty model attached to each plant of the finite set, with this uncertainty model perhaps including nonlinearity. Call the finite set of plants the set of nominal plants. Each nominal plant has associated with it a controller, presumed to give satisfactory performance with that nominal plant, and any uncertainty that might lie around the plant. The **adaptive** controller includes a 'high level' element, the supervisor, as well

[1] Research supported by US Office of Naval Research Grant N 00014-97-1-0946

as the finite set of controllers. The supervisor switches between the members of the finite controller set, following some rule. The overall object is to end up with the best controller for the two unknown plant after some finite time, and then stick with it. The notion of best controller may be ambiguous when the true plant lies in an uncertainty ball around one or more nominal plants.

This very high level qualitative description deserves a number of qualifying remarks:

- The unknown plant may well be a member of a big if not infinite set. Determination of the set of nominal plants is itself then a non-trivial problem.
- The supervisor's task is essentially one of identification, or more accurately, hypothesis testing. Here hypothesis j is the hypothesis that the true plant lies in the uncertainty ball around the j-th nominal plant. The standard issues of hypothesis testing arise, ie effects of noise, decision errors, time to make a decision with adequately low error probability etc. The fact that a controller is connected which may be switched, thus changing the experimental conditions, is a complicating factor. The settings of the hypothesis testing algorithm will then need adjustment.
- It is possible to overlay a fine-tuning structure on top of the switching structure, whereby controller parameters are tuned continuously.

We shall indicate the structure of the supervisor. Broadly speaking, supervisors are constructed to try to select that particular plant from the set of nominal prescribed plants which is closest to the true plant, given the connection of a particular controller. Having selected the plant closest to the true plant, the index of that plant is selected as the index of the next controller to be switched in. The switching of the controller causes a change of experimental conditions, with the possibility that the index of the nominal plant which is closes to the true plant could change. Hence, a performance evaluation step which determines that nominal plant which would perform best with the existing controller, is not the same as a performance evaluation step which would determine that controller performing best with the true plant. The indices associated with the best plant and the best controller need not be the same.

We turn to safe adaptive control. In nearly all adaptive control algorithms, the plant is initially unknown, and in the course of executing the adaptive algorithm, an implicit or explicit identified model of the plant is used to design a controller, which is then connected. This means that (a) the true plant normally differs from the model used for controller design purposes and (b) that controller undergoes change. This change is dangerous in the following sense. Even if the closed-loop is not showing signs of instability, the change of the controller when the plant is not fully known has the potential to introduce instability. Of course, if this happens, the adaptive algorithm should be clever enough to discover the inappropriateness of the controller, and after improved identification of the plant would be likely to change the controller further.

This is a beguiling conclusion if at times the 'frozen' controller-plant combination is unstable. It is a conclusion that allows one MA of current in a one KW motor.

Safe adaptive control refers to adaptive control algorithms in which one guarantees a priori that any controller introduced will always yield with the only partially known plant frozen closed-loop stability. That is, if the controller introduced at any time were to remain unchanged from that time on, the resulting time-invariant system would be guaranteed stable. It is obvious that one could ask for a multi-model adaptive control algorithm to have the safe adaptive control property. In this talk, we shall explore how that may be achieved.