A New Take on Probabilistic Robust Control

Anton Nyström

Department of Automatic Control, Faculty of Engineering, Lund University

May 28, 2024

Imagine you are on vacation and in the hotel shower. To your horror discover that the water is icy cold. You turn the hot water valve time and time again but nothing happens. Suddenly, the icy water turns scalding hot and you scream out in surprise and pain. Argh! How should you have turned the valves when you did not know what the result would be?

Controlling the temperature of the shower using the cold and hot valves can be a fickle thing. This is especially true when you do not know how the shower behaves. How much hotter will the water get when you turn the valve? How long is the delay before the water turns hot? This type of example is a common occurrence in control where the engineer wants to control a process output, such as the water temperature, without complete knowledge of the process behavior. Much like you having to turn the valves blindly, the engineer must find with a control law to regulate the process and achieve a desired performance; all without perfect knowledge of the process.

The disastrous result in the shower is the main thing to avoid; extreme temperatures are not at all what you want to experience. The same thing goes for the control engineer where the name of the game is to avoid instability in the process and achieve some performance. You realize that you acted stupidly in the shower, after all you have used many showers before and know that sometimes this is what happens. Interestingly enough you could have used your earlier experience with showers to influence and improve your decisions.

Gathering information by measuring a process allows you to create a model of its behaviour. But due to the information of the process being incomplete you are uncertain of its exact behaviour and the model will never perfectly describe it. Using the model to pick a control law and applying it to the true process can therefore yield unpredictable results. This led researchers to come up with ways of guaranteeing control performance as long as you are able to limit the uncertainty to a certain set of potential models. You can do this by gathering enough information to dismiss all those which would not fit the data you collected. This sounds great, but the approach lacks nuance and only provides a yes-no answer to whether the performance guarantee is satisfactory given the information you have gathered.

So the question is: can we add nuance by using the collected information about the uncertain system cunningly? The answer is yes! Specifically by also considering how rare or common each potential model is using a probability distribution. With this we can place sharper probabilistic guarantees on control performance. By picking an acceptable probability of failure you can now pick a control law that performs well most of the time and fails only with at most that probability. This makes the method of deciding on a control law very versatile to the situation as some situations demand very low probabilities of failure and some are more lenient and can afford a less cautious approach. For you taking a shower, it is all up to preference of how cautious you want to be. How lovely!

This new result provides extra nuance by replacing the yes-no answer with a trade-off between a guaranteed performance level and the rate of failure with respect to that level. As a consequence, this provides a fresh new perspective on probabilistic robust control and opens up for a wide range of future research to expand the work started here. Knowing how to take the perfect refreshing shower, you can now carry on with a care-free, relaxing and well earned vacation.

Read more in the Master thesis
Understanding Probabilistic Uncertainty using Nu-Gap
by Anton Nyström, TFRT-6230
Available for download from https://lup.lub.lu.se/student-papers/