Can local particle filters beat the curse of dimensionality?

Ramon van Handel, Princeton University

The discovery of particle filtering methods almost two decades ago has enabled the use of nonlinear filtering in a wide array of applications. Unfortunately, the approximation error of particle filters typically grows exponentially in the dimension of the underlying model. This phenomenon, which has only been rigorously understood fairly recently, has rendered particle filters of limited use in complex data assimilation problems. In this talk, I will argue that it is often possible, at least in principle, to develop "local" particle filtering algorithms whose approximation error is in fact dimension-free. The key to such developments is the decay of correlations property, which is a spatial counterpart of the much better understood stability property of nonlinear filters. For the simplest possible algorithm of this type, our results provide under suitable assumptions an approximation error bound that is uniform both in time and in the model dimension. Methods from statistical mechanics play a central role in the proofs. (Joint work with P. Rebeschini)