**PARALLEL FILTERING AND SMOOTHING METHODS FOR STATE-SPACE MODELS**

SIMO SÄRKKÄ

State space models (SSMs), including Gaussian state space models, non-linear/non-Gaussian state space models, and hidden Markov models (HMMs), are important tools in target tracking, time series analysis, machine learning, and various other fields. Bayesian filters and smoothers as well as their special cases such as Kalman filters and smoothers or approximations such as particle filters and smoothers are computationally optimal O(T) algorithms for state estimation in these models on classical CPU architectures. However, in parallel setting, such as on GPUs, they are no longer optimal, because they are inherently sequential algorithms. The aim of this talk is to discuss parallel versions of these algorithms. Most of the algorithms are based on so-called associative scans, which are computational primitives that can already be found, for example, in TensorFlow and JAX, and are easily implementable, for example, in Julia/CUDA.jl. These algorithms can be used to make state estimation optimally parallelizable leading to parallel O(log T) span complexity.