


Speedup of Kalman Filter fit and smoother
across architectures



Some conclusions

- We have achieved some good things
 - Cross architecture, fully vectorised Kalman Filter
 - Base 2x with respect to original implementation
 - Up to 6.7x on KNL
 - Our roofline model shows we are getting the most out of the machine
- There is more to come
 - Integration with present and future Gaudi
 - GPU implementation revamp
 - Vectorisation for Power8
 - Testing + vectorisation for ARM64