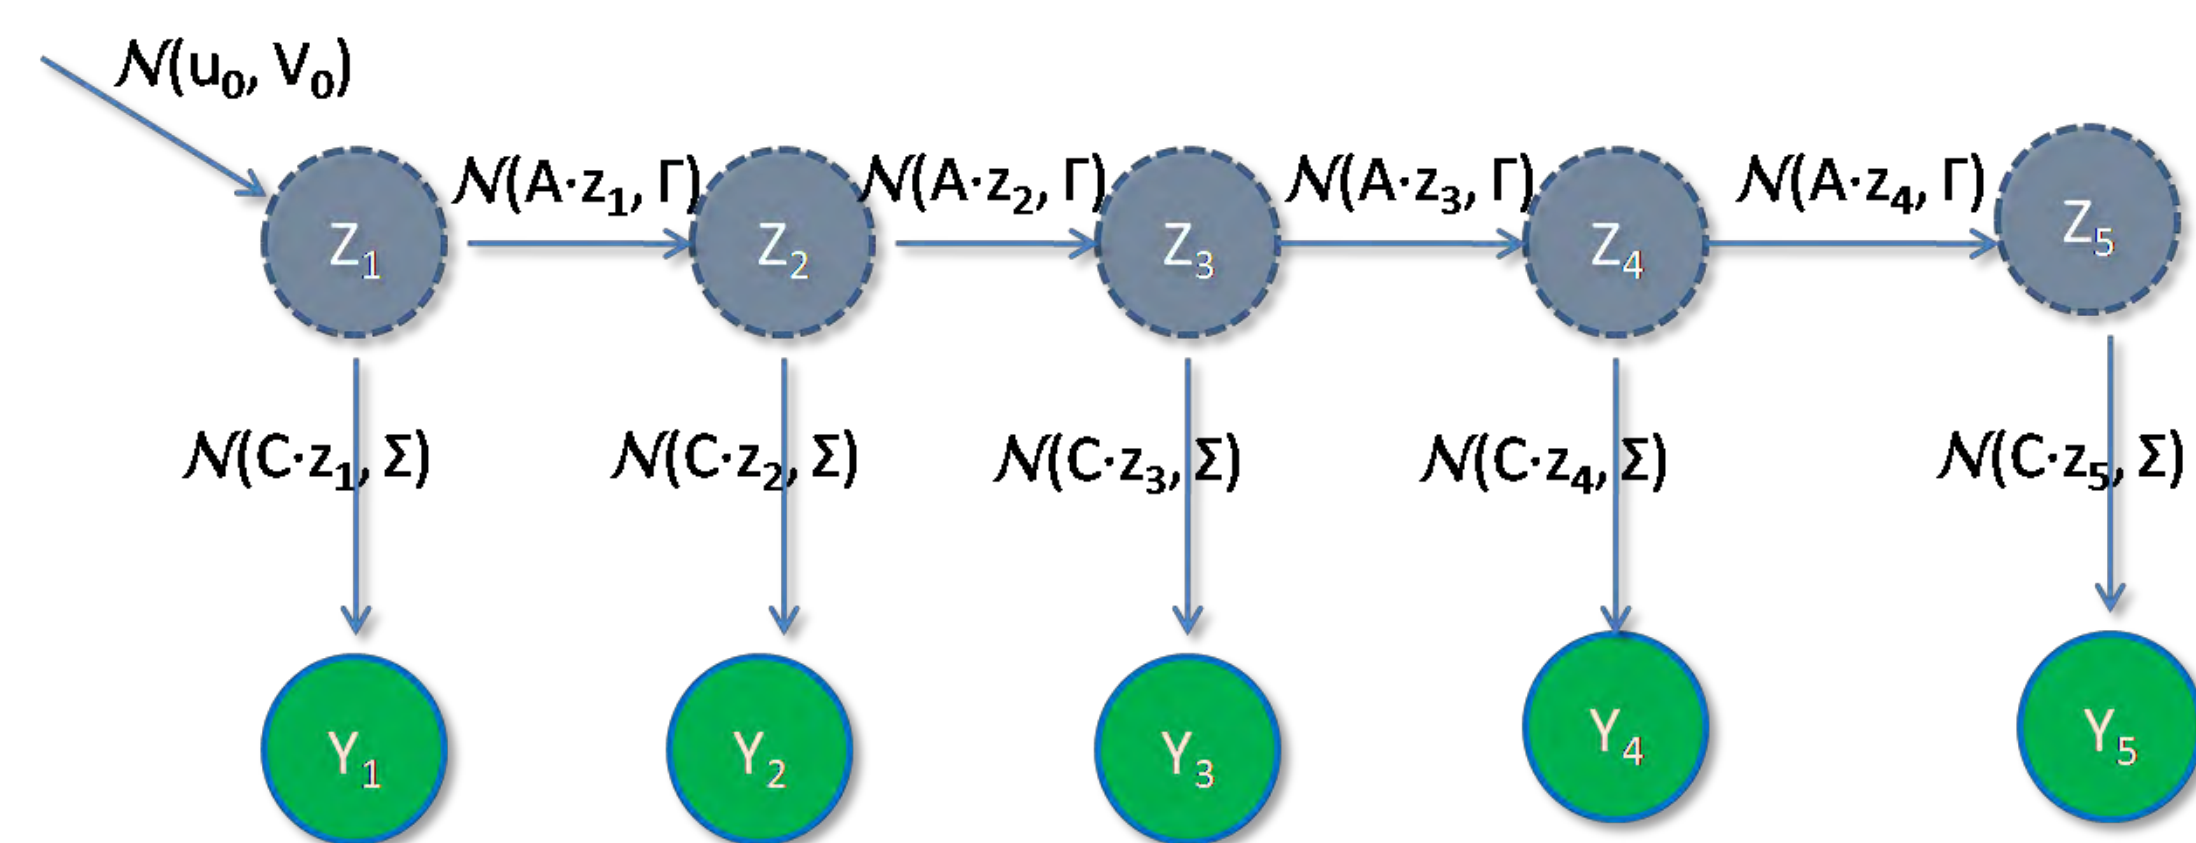




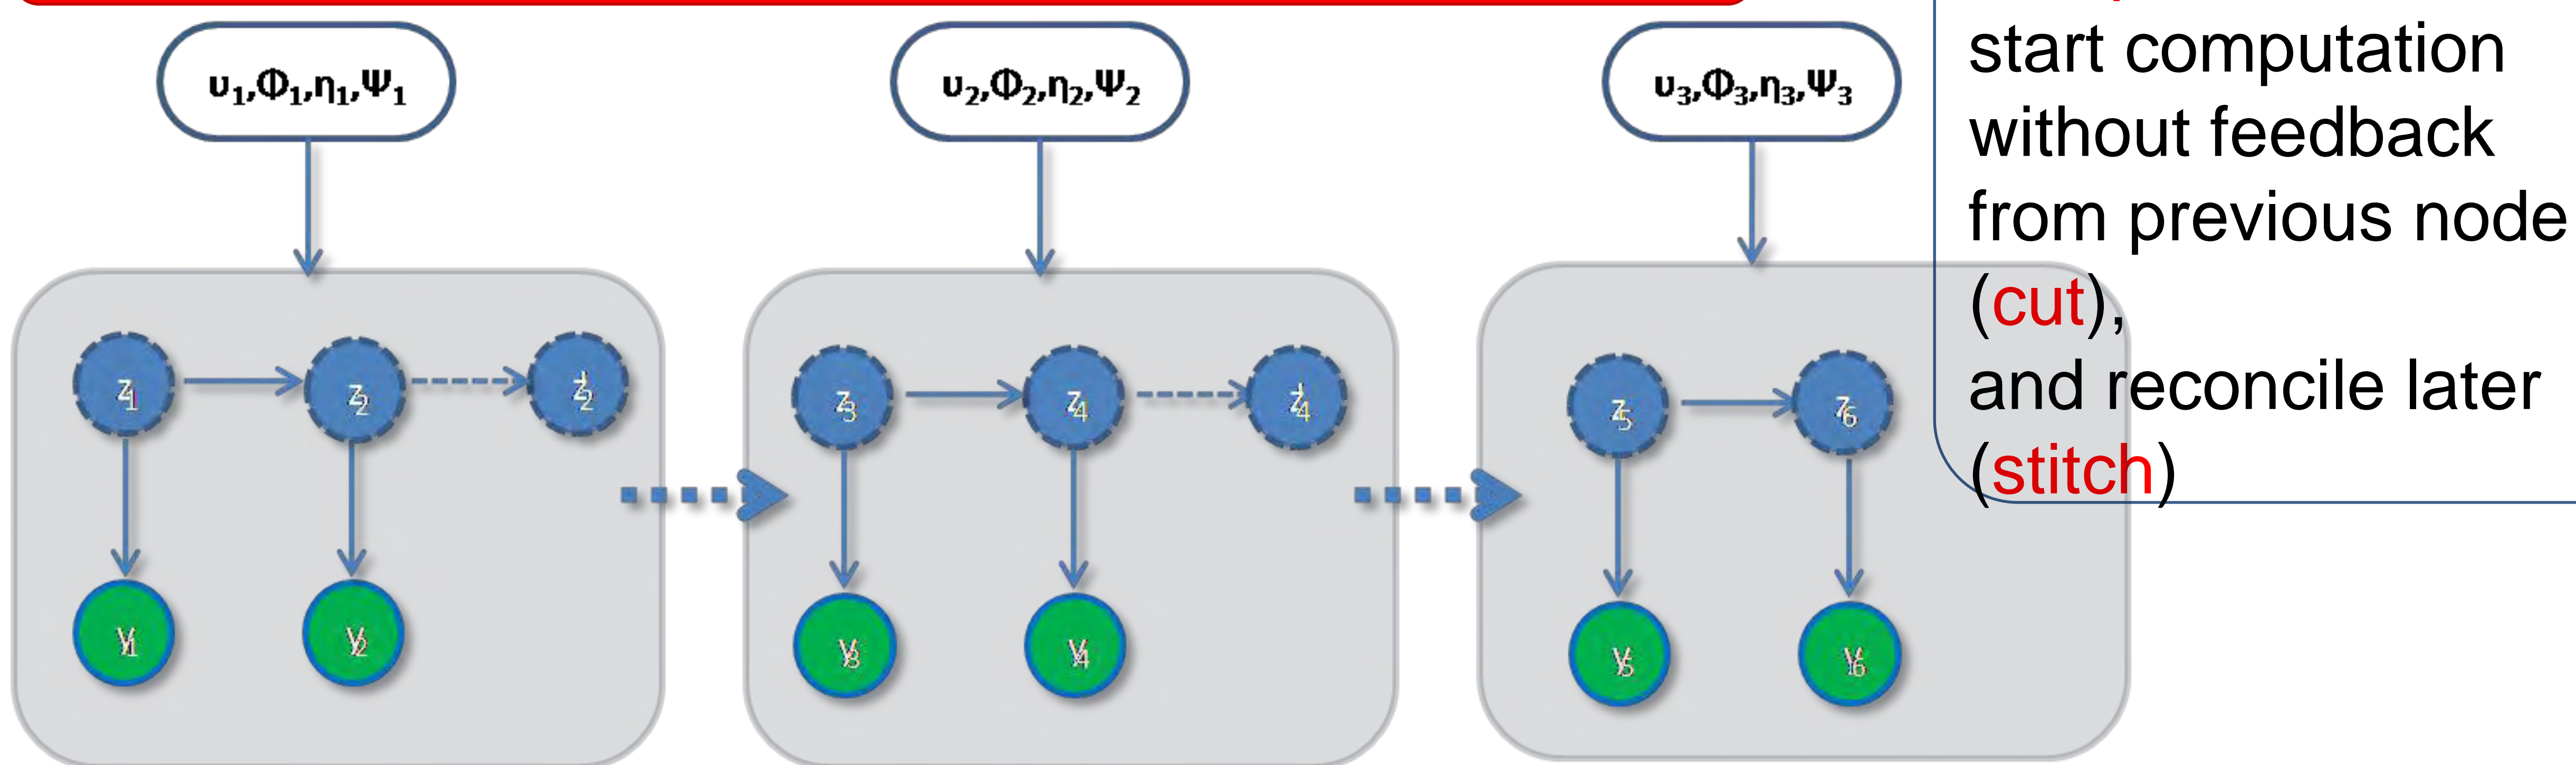
Problem Definition

- Problem:
Given a Linear Dynamical System, to find the model parameters.
- Traditional Method:
Maximum Likelihood Estimation via Expectation-Maximization(EM) algorithm
- Objective:
Parallelize the learning algorithm
- Tool:
OpenMP with shared memory parallelism

Parameters: $\theta=(u_0, v_0, A, \Gamma, C, \Sigma)$
Observation: $y_1 \dots y_n$
Hidden variables: $z_1 \dots z_n$




Proposed Method: Cut-And-Stitch



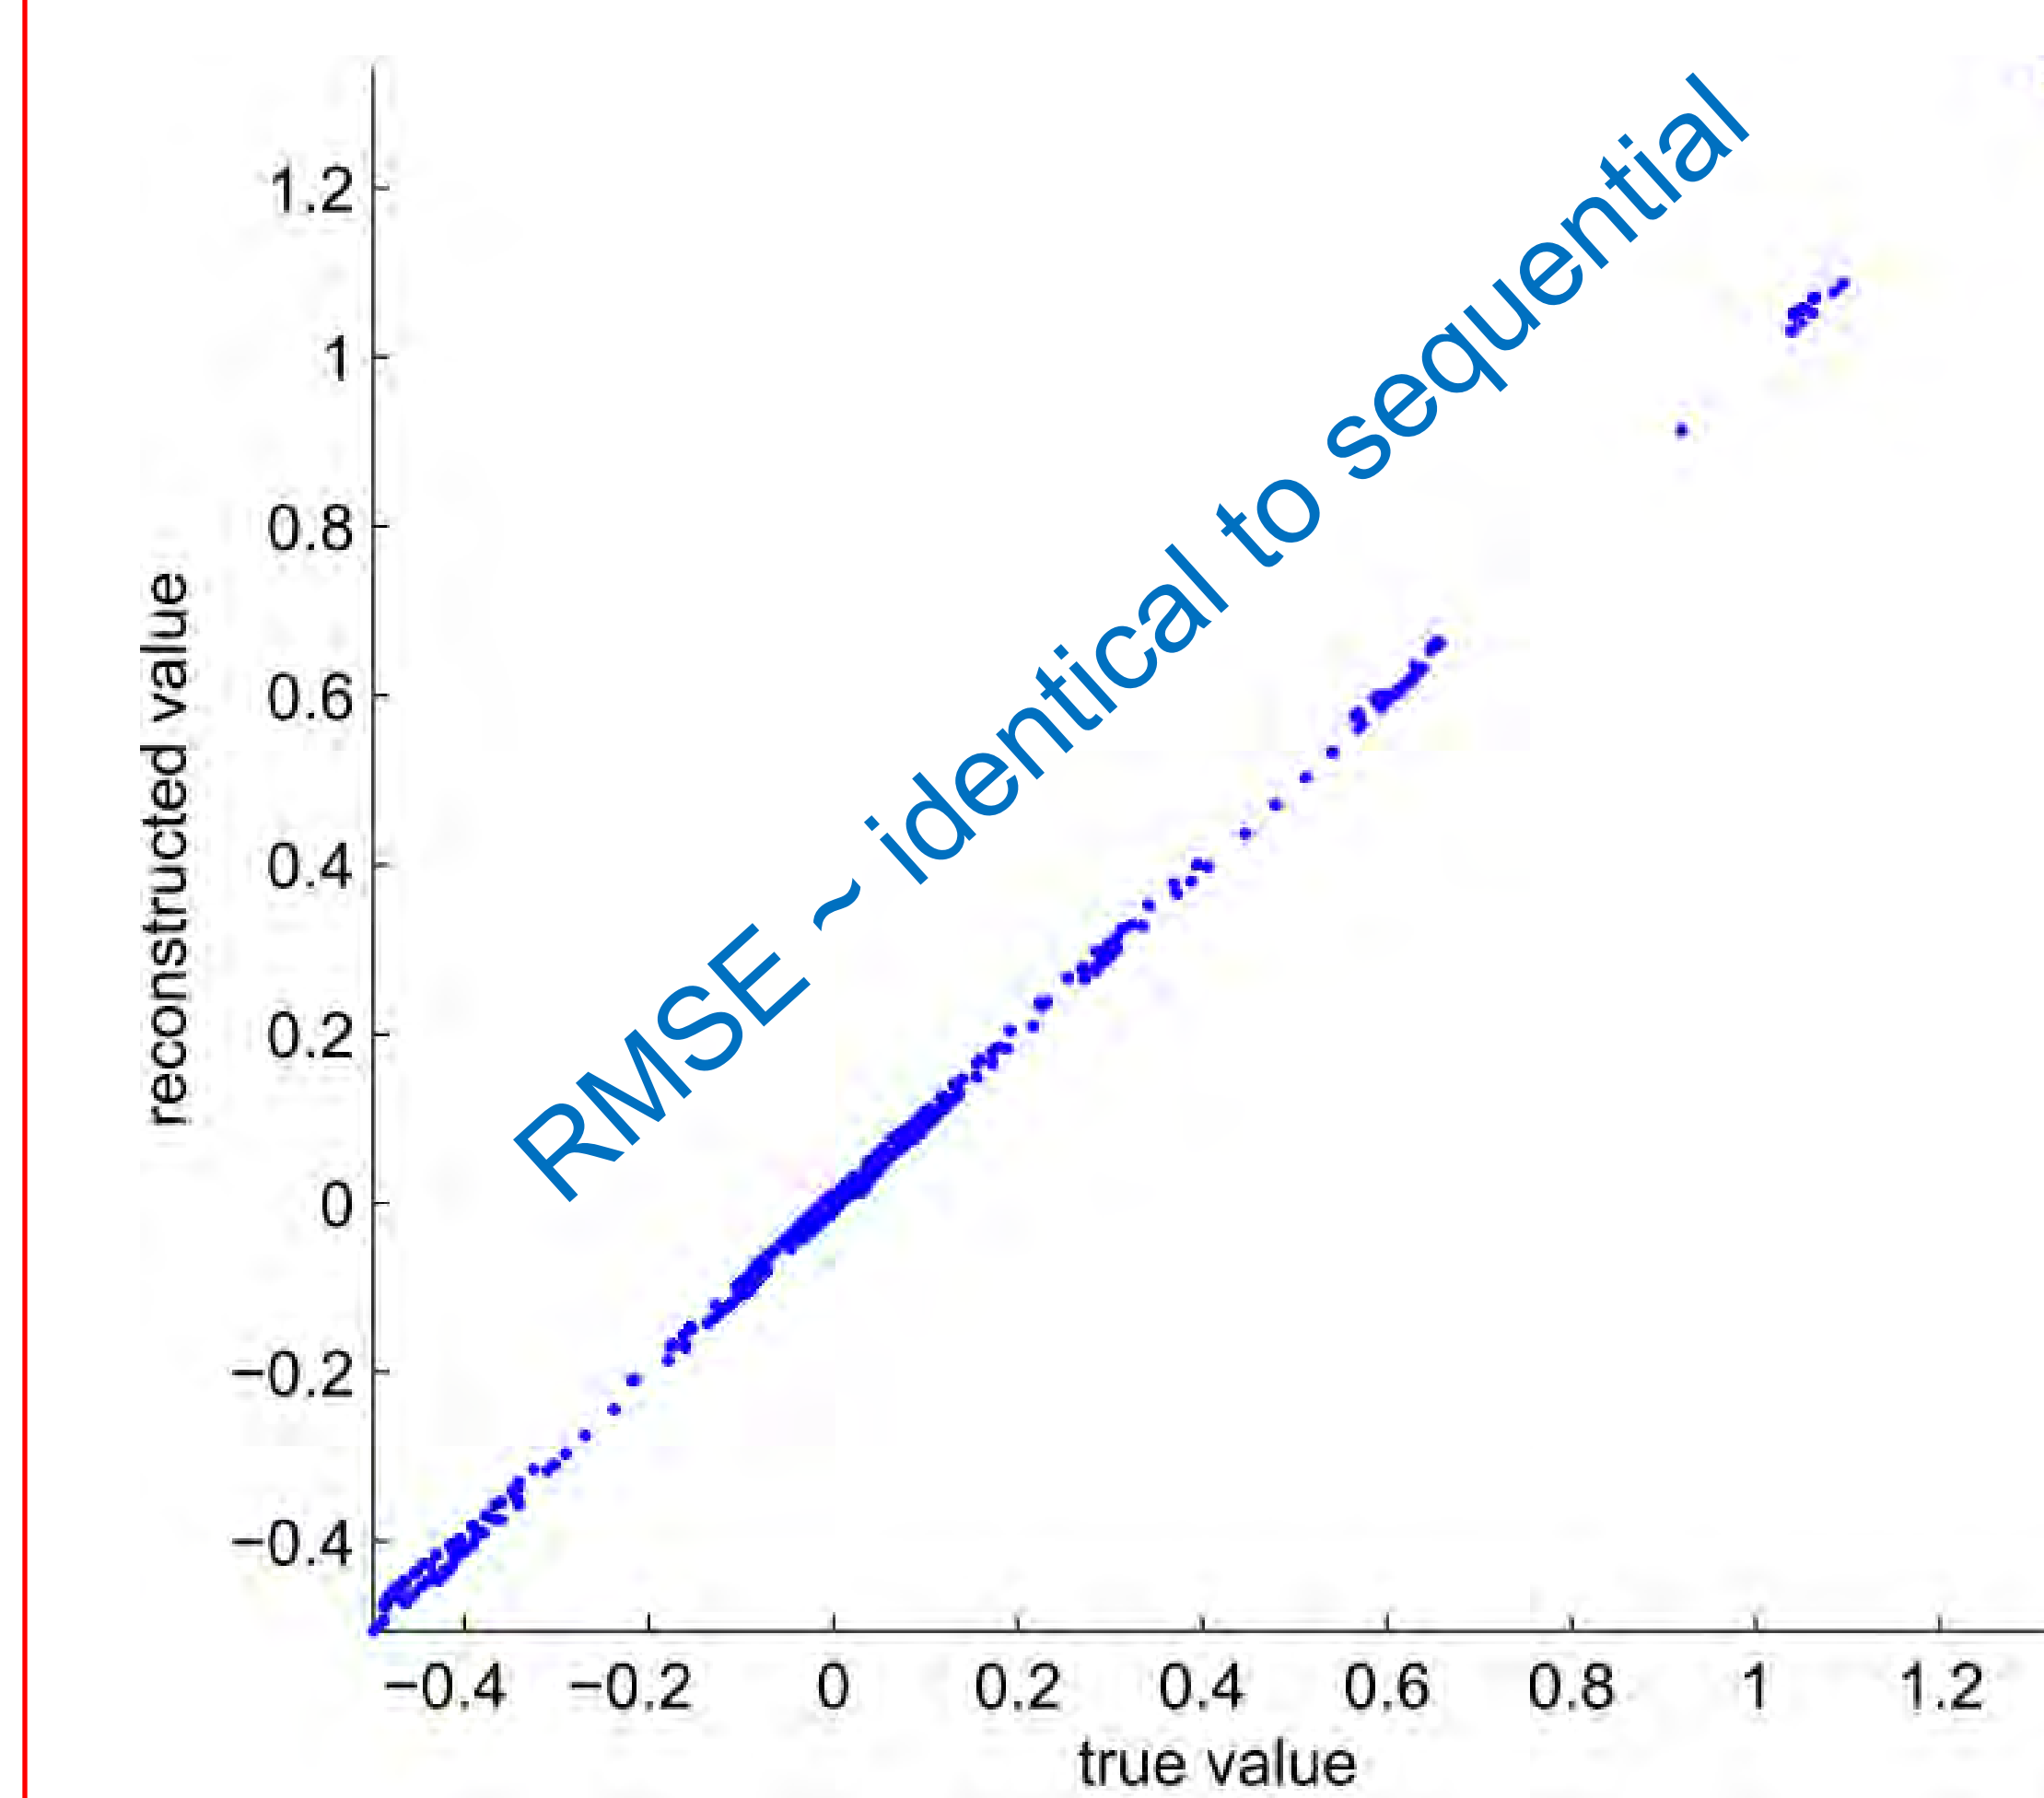
Experiments

- Dataset:
58 motion capture sequences.
Each with 93 dimensions (bone positions).
- Setup:
tested on both SMP (SGI Altix) and desktop multicore (4-core)

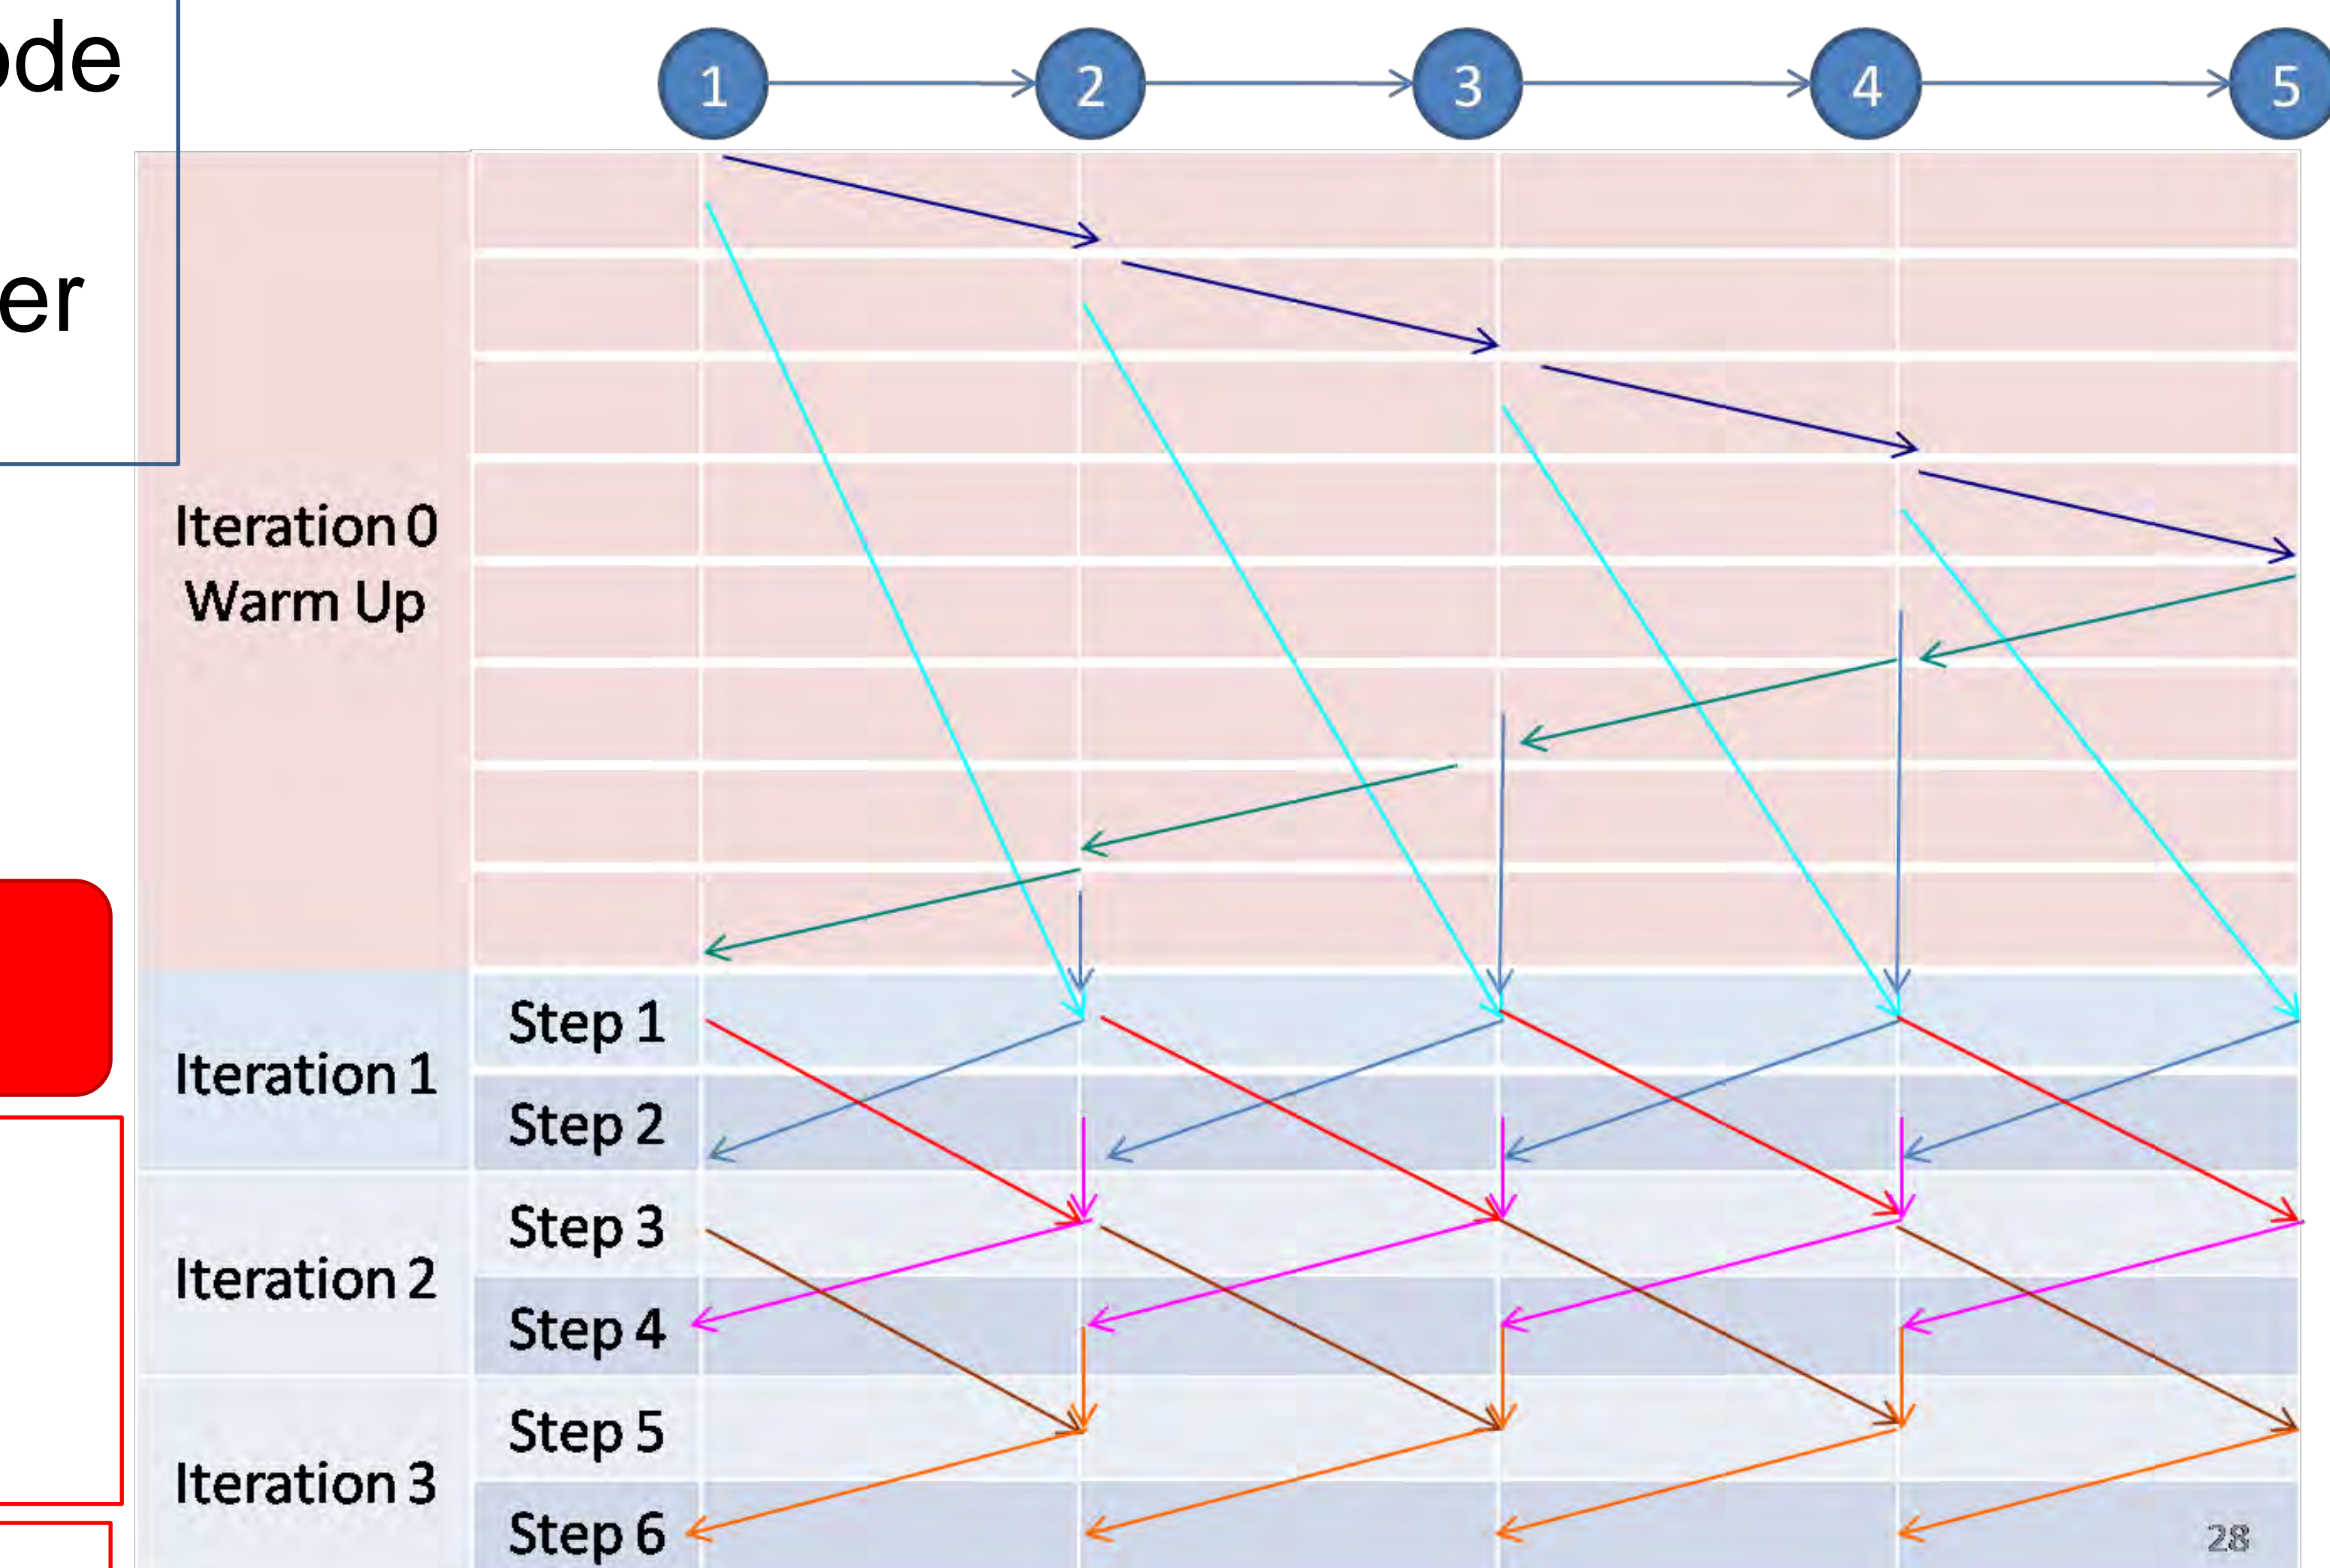
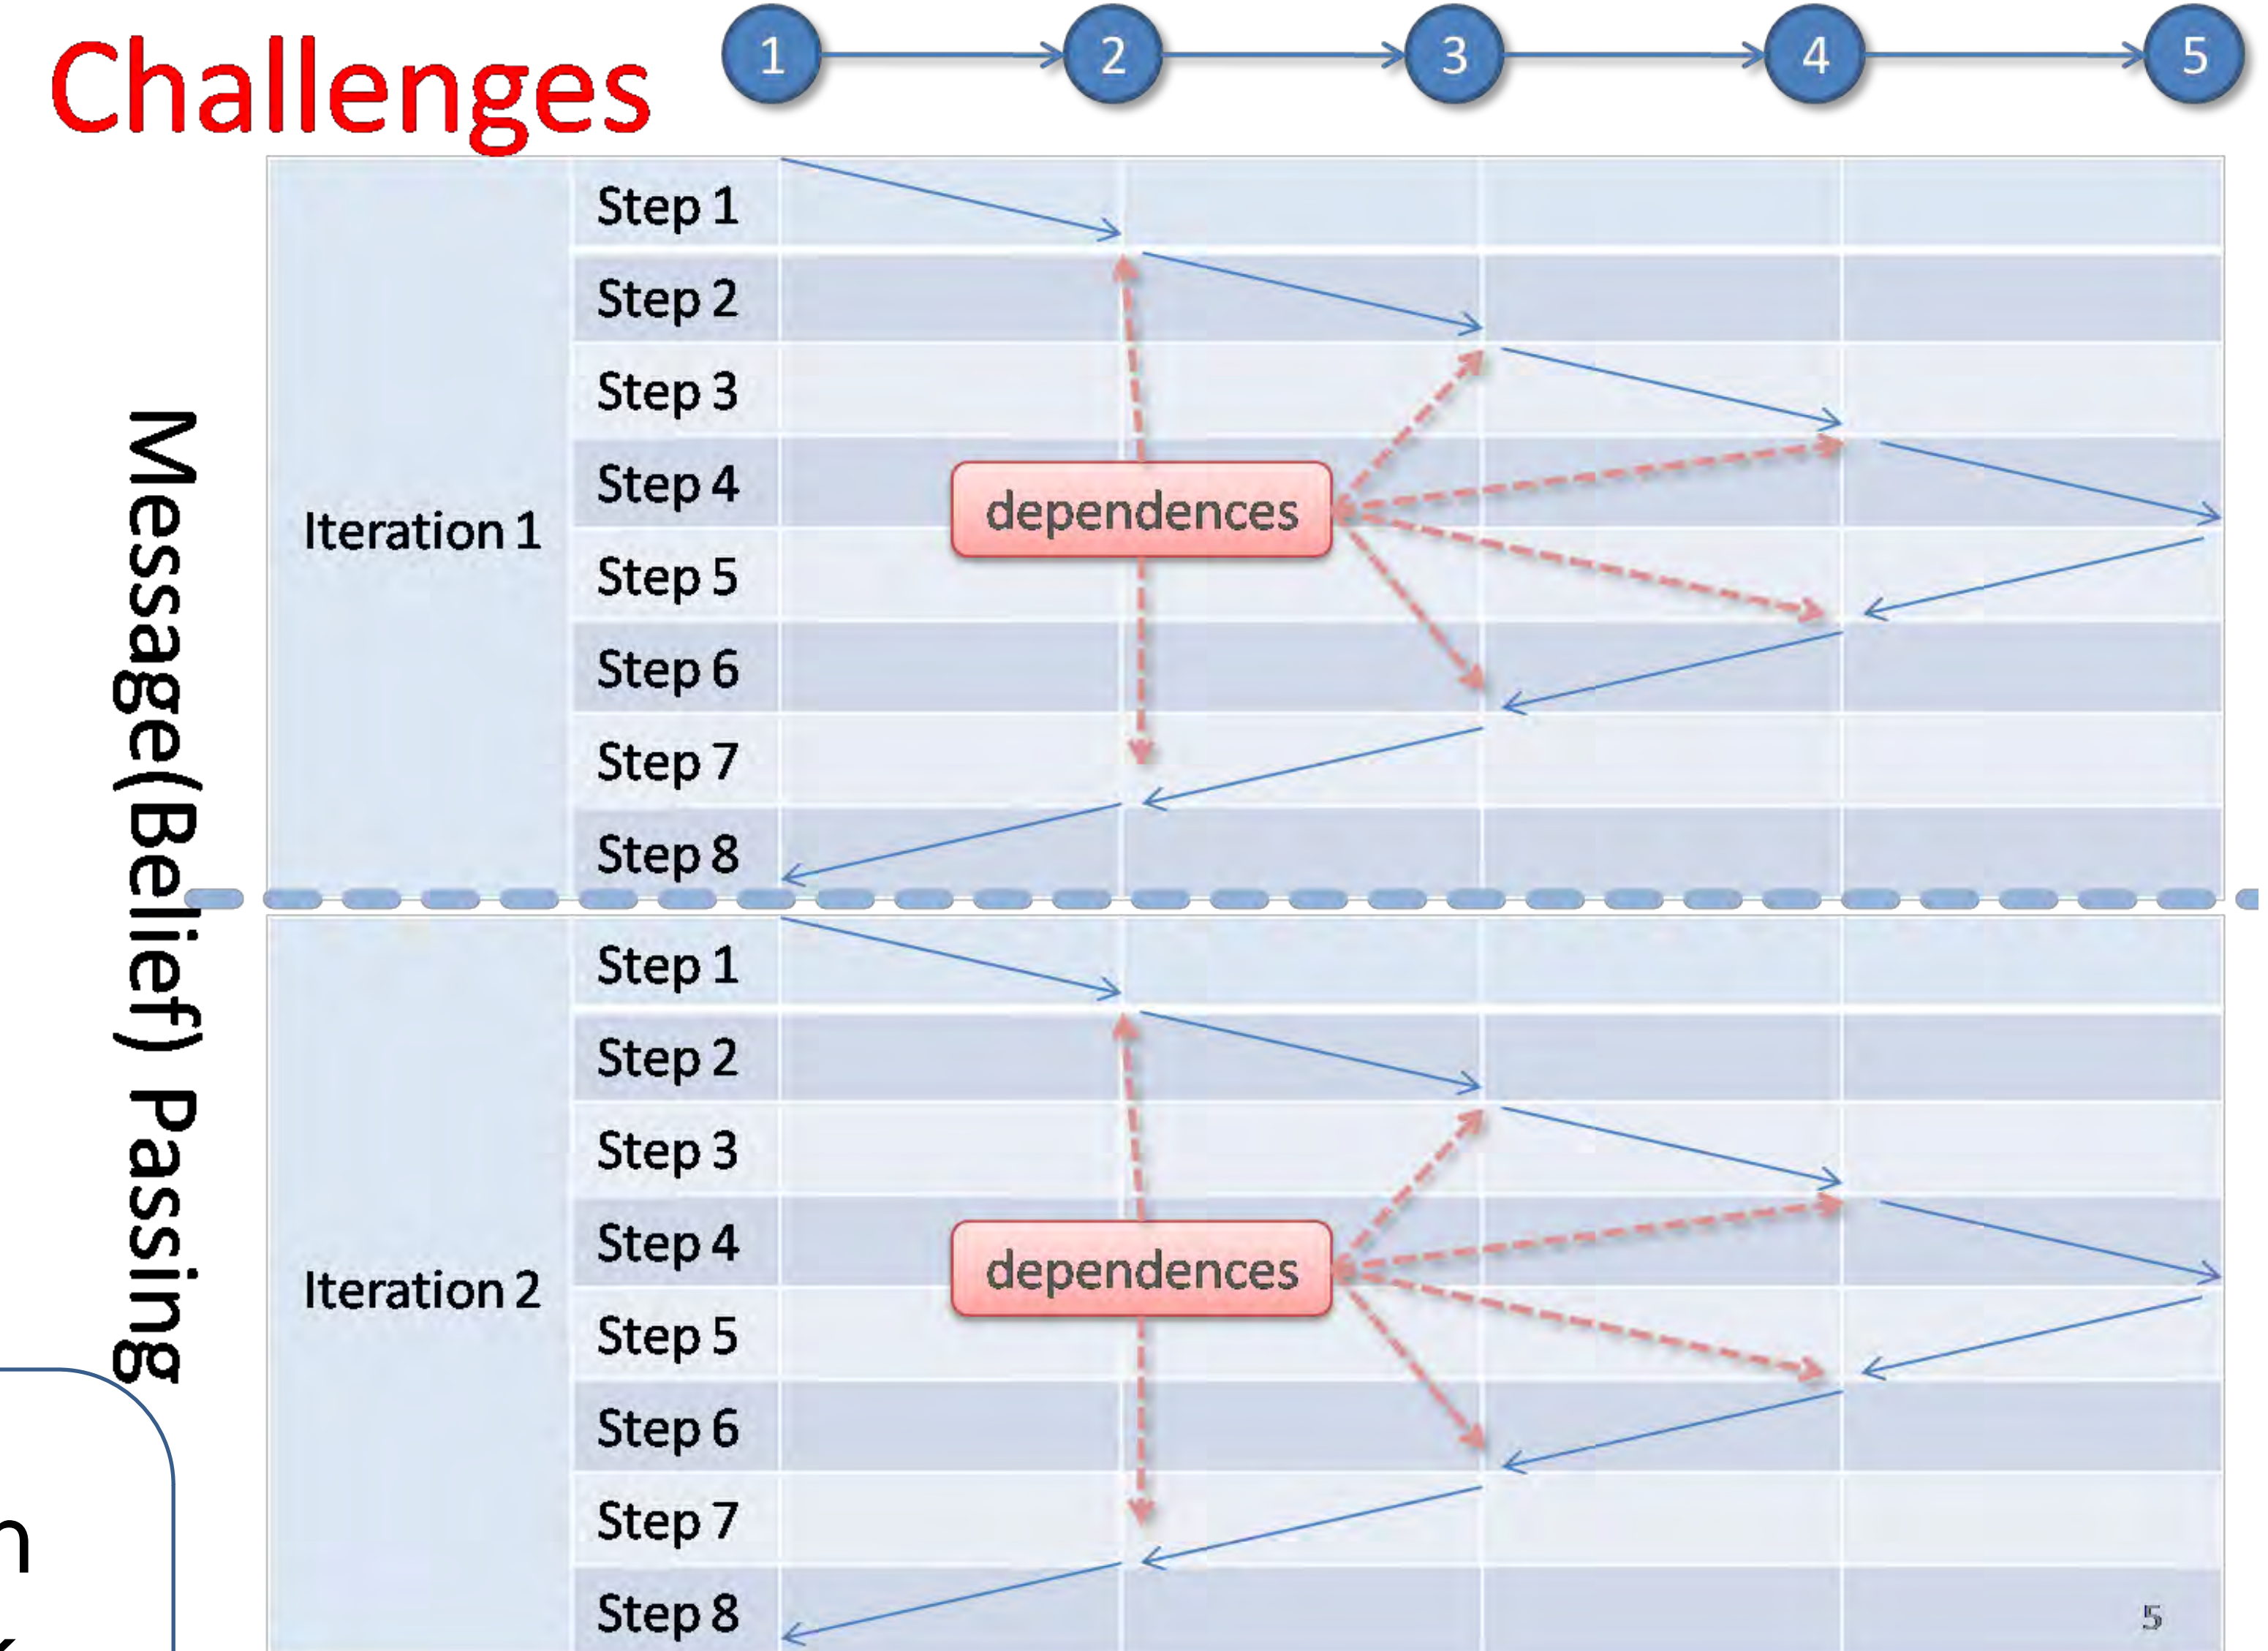
Speed up over the sequential algorithm



Quality: Reconstructed v.s. Real



Traditional Learning



Conclusion

- General approximate parallel learning algorithm for LDS
- Near linear speed up
- Quality (RMSE): ~ identical to sequential learning
- Promising Extensions:
 - Cut-And-Stitch for HMM
 - Cut-And-Stitch for general Markov models

Dataset and software available at
<http://www.cs.cmu.edu/~leili/paralearn/>