Bayesian calibration of the Community Land Model using a multi-chain Markov chain Monte Carlo method

Laura Swiler[1], J. Ray[2], M. Huang[3], Z. Hou[3]

[1]Sandia National Laboratories, Albuquerque, NM, [2]Sandia National Laboratories, Livermore, CA, and [3]Pacific Northwest National Laboratory, Richland, WA

OBJECTIVE

Perform Bayesian calibration of three hydrological parameters in the Community Land Model (CLM) using observed latent heat fluxes (LH) from the US-ARM site

- Parameters of interest Θ = {subsurface runoff decaying factor, maximum subsurface drainage, Clapp-Hornberger exponent}.
- Quantify the uncertainty in the parameter estimation .
- Use Markov chain Monte Carlo (MCMC) to solve the calibration problem and develop probability density function for parameters.

BACKGROUND

- Community Land Model (CLM) is the land component of the Community Earth System Model. It simulates bio-geophysical processes such as energy and water fluxes from canopy and soil; heat transfer in soil and snow; hydrology of soil, canopy, and snow; and stomatal physiology and photosynthesis.
- Sensitivity analysis of LH simulated by CLM to 10 parameters was performed: the top three parameters contributing to the variance of LH are F_{drai} (subsurface runoff decaying factor), Q_{dm} (maximum subsurface drainage), and b (Clapp-Hornberger exponent in the soil water retention curve).
- MCMC requires $O(10^4)$ evaluations of CLM version 4. The expense of running CLM becomes prohibitive for these calculations, especially when one needs to run CLM for multiple years to compare latent heat fluxes.
- Surrogates have been explored, but the accuracy of the surrogates could be questionable depending on the variables and/or sites of interest.
- Our solution is to examine parallel MCMC chains. This will allow exploration of the parameter space using multiple communicating chains so that surrogates could be avoided.

Bayesian Inverse Problems

- Let $y^{obs}(t)$ be the observed monthly latent heat at month t; let $M(t;\Theta)$ be the CLM model prediction for the same month.
- Observation model: $y^{obs}(t) = M(t;\Theta) + \varepsilon$, $\varepsilon \sim N(0,\sigma^2)$ is a model-data mismatch modeled as i.i.d. Gaussian.
- Calibration problems involves developing an expression for the posterior probability distribution $P(\Theta \mid y^{obs}(t))$ using Bayes' rule.
  - $\Pi(\Theta)$ is our prior belief in Θ i.e., uniform distributions over the ranges of F_{drai}, $\log_{10}(Q_{dm})$ and b.

$$P(\Theta,\sigma^2 \mid \mathbf{y}^{obs}) \propto P(\mathbf{y}^{obs} \mid \Theta)\Pi(\Theta)$$
$$\propto \frac{1}{\sigma}\exp\left[-\frac{\left\|\mathbf{y}^{obs}-M(t;\Theta)\right\|_2}{\sigma^2}\right]\Pi(\Theta)$$

- This is a 3-parameter estimation.
- Will use an adaptive MCMC method to compute estimates in the form of a multidimensional posterior distribution.

SaChES

Scalable Adaptive Chain-Ensemble Sampling

SaChES is a hybrid method that incorporates:

- DREAM to utilize multiple chains to obtain high-quality proposal densities; useful in the early part of the sampling epoch.
- DRAM to obtain posterior distributions efficiently; useful only after a few samples have been collected by DREAM.
- Parallel chains to accelerate computations.

DREAM: DiffeRential Evolution Adaptive Metropolis

- DREAM [1] uses information from multiple chains to construct a new chain position that is a weighted value of previous chain positions.
- DREAM chains communicate in pairs every MCMC step but these communication are unconnected to the communication of other pairs. There is no need for synchronization.

DRAM: Delayed Rejection Adaptive Metropolis

- DRAM [2, 3] is an MCMC algorithm which uses accepted samples from a MCMC chain to construct a very efficient proposal covariance. The covariance is updated periodically and requires all-to-all communication
- DRAM performs best if there are a few informative samples to begin with. This is difficult to do in practice. DREAM can perform the task of collecting the initial informative samples.

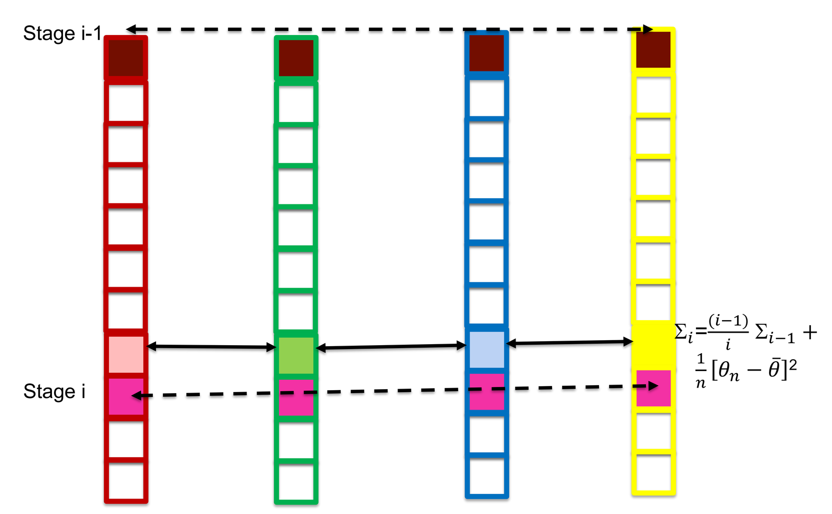

Schematic of 4 SAChES chains and their communication patterns for updating proposal covariance.

MPI Implementation

- SAChES is implemented using one-sided MPI-2 communication.
  - Done to ensure scalability, given that chains are loosely coupled.
  - Also, for resilience – MPI-2 allows for re-spawning of dead processes.
    - Useful if chains have to be restarted.
- DREAM requires chain-to-chain communications between arbitrary, mutually disjoint chain pairs during differential evolution.
  - Implemented using lock synchronization (MPI_Win_[lock(),unlock()].
- DRAM requires all-to-all communications between processors to construct proposal covariances.
  - MPI_allreduce() is the preferred implementation.
  - Alternative implementation with active target synchronization also exists.
  - Uses MPI_Win_[start(), complete(), post(), wait()].
- Actual data transfers done using MPI_Get().

Test Problem Results

Test #1: Sampling a 2D Rosenbrock distribution.

- Analytical expression of a very twisted probability distribution; sampling is a challenge.
- Figure (top left) shows that our chains are exploring the correct distribution.
- Figure (bottom right) shows that our 4 chains (different colors) are sampling uniformly i.e., not stuck in one part of the distribution.
- This is also shown in the marginal PDFs.

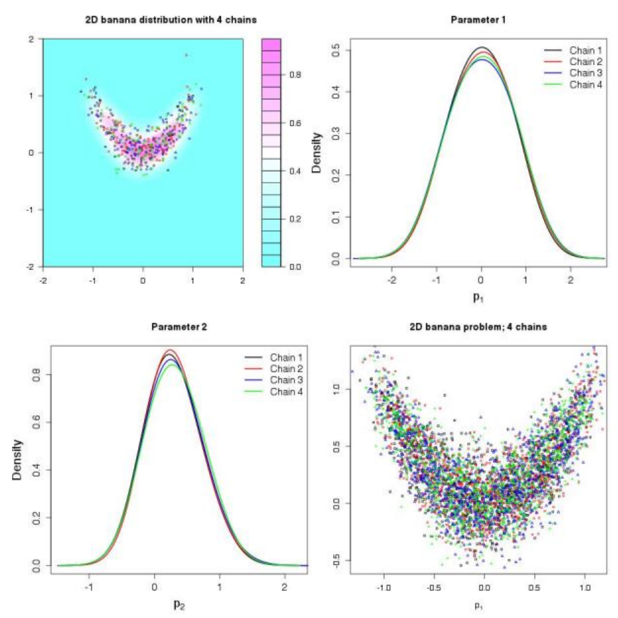

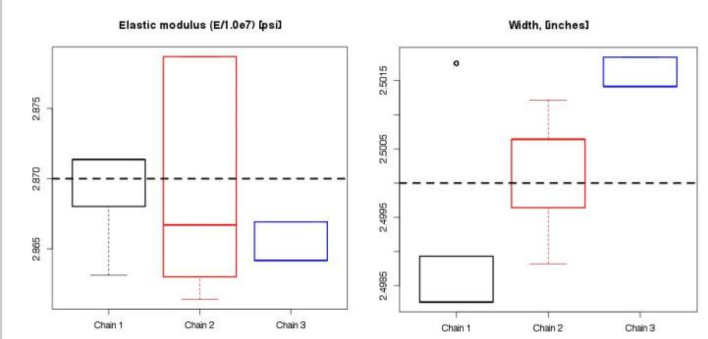

Test #2: Solving a statistical inverse problem to estimate elastic modulus (E) and width (w) of a cantilever beam.

- measurements are of stress and deflection, collected over 10 repeated experiments.
- Dashed line = true value.

CLM Inversion Results

- We modeled four years of monthly latent heat (LH) data, from the US-ARM flux tower site. The four years of data were climatologically averaged, to obtain 12 months of average LH flux. These 12 monthly averaged observations were then compared to the same four years of CLM predictions which were also climatologically averaged.
- Range of the parameters
  - 0.1 < F_{drai} < 5.0; default: 2.5
  - -6 < $\log_{10}(Q_{dm})$ < -2;; default : $\log_{10}(5.5 \times 10^{-3})$
  - 1< b < 15 ; default : 9.76
- The parallel 8-chain MCMC is still in progress, but early chain results show that the parameter b appears to be converging, F_{drai} and Q_{dm} are not yet.

CLM Posterior Traces from 8 Chains

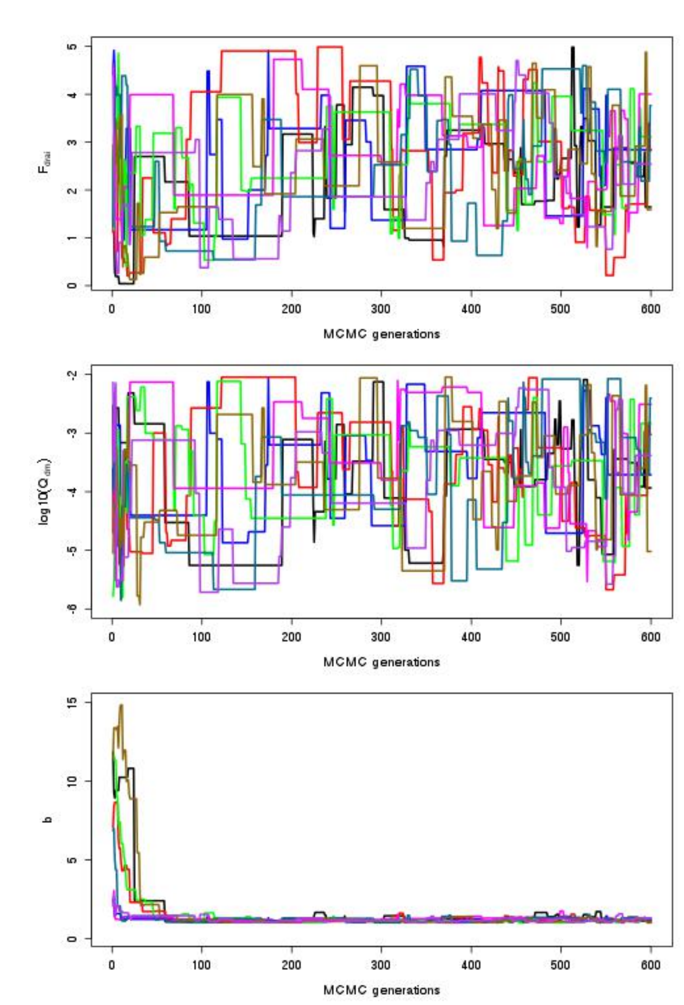

Comparing Model Predictions Against Observations

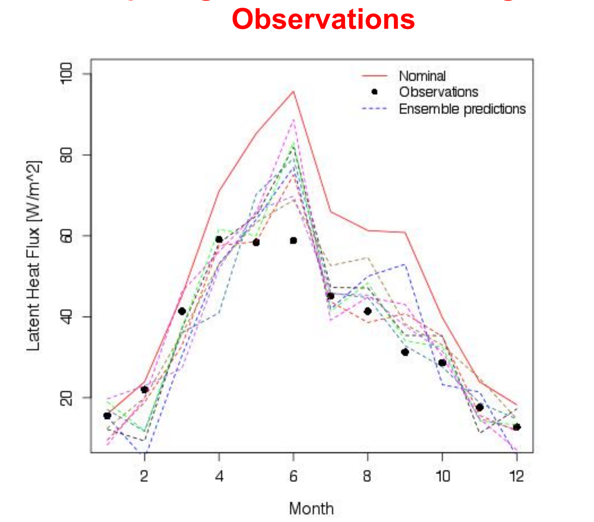

- The LH predictions based on posteriors (not yet final) are substantially better than those based on default CLM parameters.
  - Figure (right) shows an ensemble plot from each of the chains, plotted against observations.
  - Predictions with the default values of the parameters are in red.

Posterior Densities

Posterior Histograms (across chains)

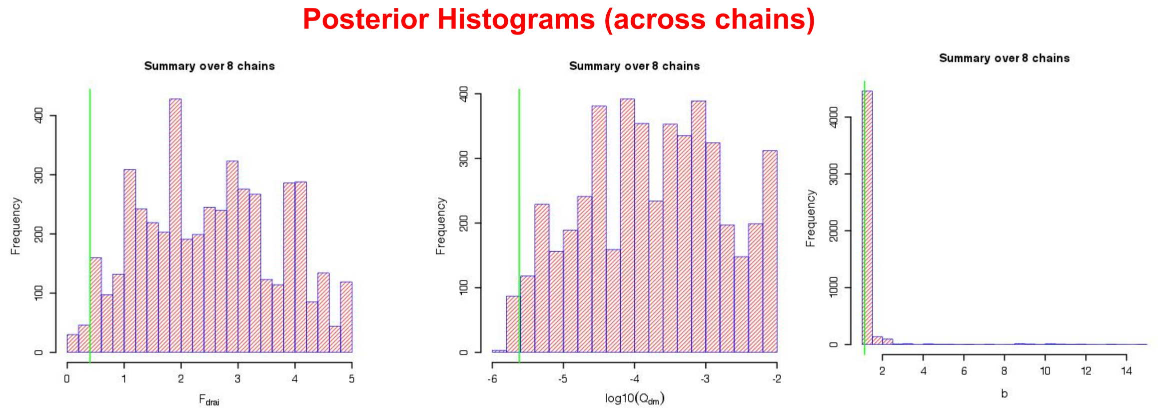

	Fdrai	Qdm	b
Nominal Value	2.500	-5.203	9.760
Median Posterior with surrogate	0.400	-5.620	1.120
Median from SaChES Parallel runs thus far	2.589	-3.769	1.171
1st quartile from SaChES Parallel runs thus far	1.528	-4.452	1.096
3rd quartile from SaChES Parallel runs thus far	3.471	-2.979	1.281

- Chains for F_{drai} and Q_{dm} have not converged; the median values computed using a surrogate are quite different.
  - Inference using surrogate models are plotted in green in the figures.
- The chain for b has converged and recovered the value estimated using surrogate models.

SUMMARY

- We have a version of SaChES working to estimate CLM parameters.
- The estimated parameters are more predictive than the default values.
- The b parameter appears to have converged for our CLM problem, still need more runs to determine convergence for F_{drai} and Q_{dm}.
  - Parameter inferences using CLM have been compared to those obtained using surrogates. They agree for b.
- Next step: Investigate efficiency of parallel methods with coordinating chains vs. multiple independent chains.

Acknowledgements

The project was funded by the Department of Energy Office of Science, via the Office of Advanced Scientific Computing Research (OASCR).

References

1. J. A. Vrugt, C.J.F. ter Braak, C.G.H. Diks, B. A. Robinson, J. M. Hyman, and D. Higdon. "Accelerating Markov Chain Monte Carlo Simulation by Differential Evolution with Self-Adaptive Randomized Subspace Sampling." International Journal of Nonlinear Sciences and Numerical Simulation, 10(3, March 2009, pages 271-288.
2. H. Haario, M. Laine and A. Mira, "DRAM: Efficient adaptive MCMC", Statistical Computing, 16:339-354, 2006.
3. A. Solonen, P. Ollinaho, M. Laine, H. Haario, J. Tamminen, and H. Jarvinen, "Efficient MCMC for climate model parameter estimation: Parallel adaptive chains and early rejection", Bayesian Anal., 7 (2012), pp. 715–736.

For additional information, please contact:

J. Ray, Sandia National Laboratories, jairay@sandia.gov

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

PNNL is operated by Battelle Memorial Institute for the U.S. Department of Energy under contract DE-AC05-76RLO1830.

SAND2015-10732C