XSPEC and MCMC

Keith Arnaud

Center for Research Excellence in Space Science and Technology

University of Maryland College Park
and
NASA's Goddard Space Flight Center

XSPEC status

XSPEC v12.9.1m and PyXspec v2.0 in HEAsoft v6.22 last week.
Next full release of XSPEC late '17/early '18.

PyXspec v2.0 added compatibility with Python 3.

ADS search shows almost 9,000 refereed papers including the word XSPEC. Currently about 500 are published per year.

XSPEC Facebook group has 614 members.

Challenges and Opportunities

XSPEC is used by a large number of people for a wide variety of spectra and models.

It has to work out of the box for people who might not have the knowledge or inclination to delve into details.

Because it is so widely used I have the chance to help push the field towards more statistical sophistication.

However people are intolerant towards changes which make their work harder so improvements must be as seamless as possible.

Markov Chain Monte Carlo

MCMC generates sets of parameter vectors which in the stationary limit are drawn from the posterior probability density.

All MCMC algorithms work by generating a proposed next link in the chain and then using an appropriate randomization process to decide whether to accept the new values or repeat the current ones.

It is important to run the chain(s) long enough for them to be stationary. It is also customary to throw away the start of the chain (burn-in) to remove transients based on the starting parameter vector(s).

Algorithms

Metropolis-Hastings: Possible new links are generated using a proposal probability distribution. Ideally this proposal distribution should be reasonably close to the target distribution.

Goodman-Weare (EMCEE): Effectively runs multiple simultaneous chains (called walkers) which interact. This requires only a decision about the number of walkers and their initial values.

XSPEC defaults to Goodman-Weare since this requires fewer decisions by the user.

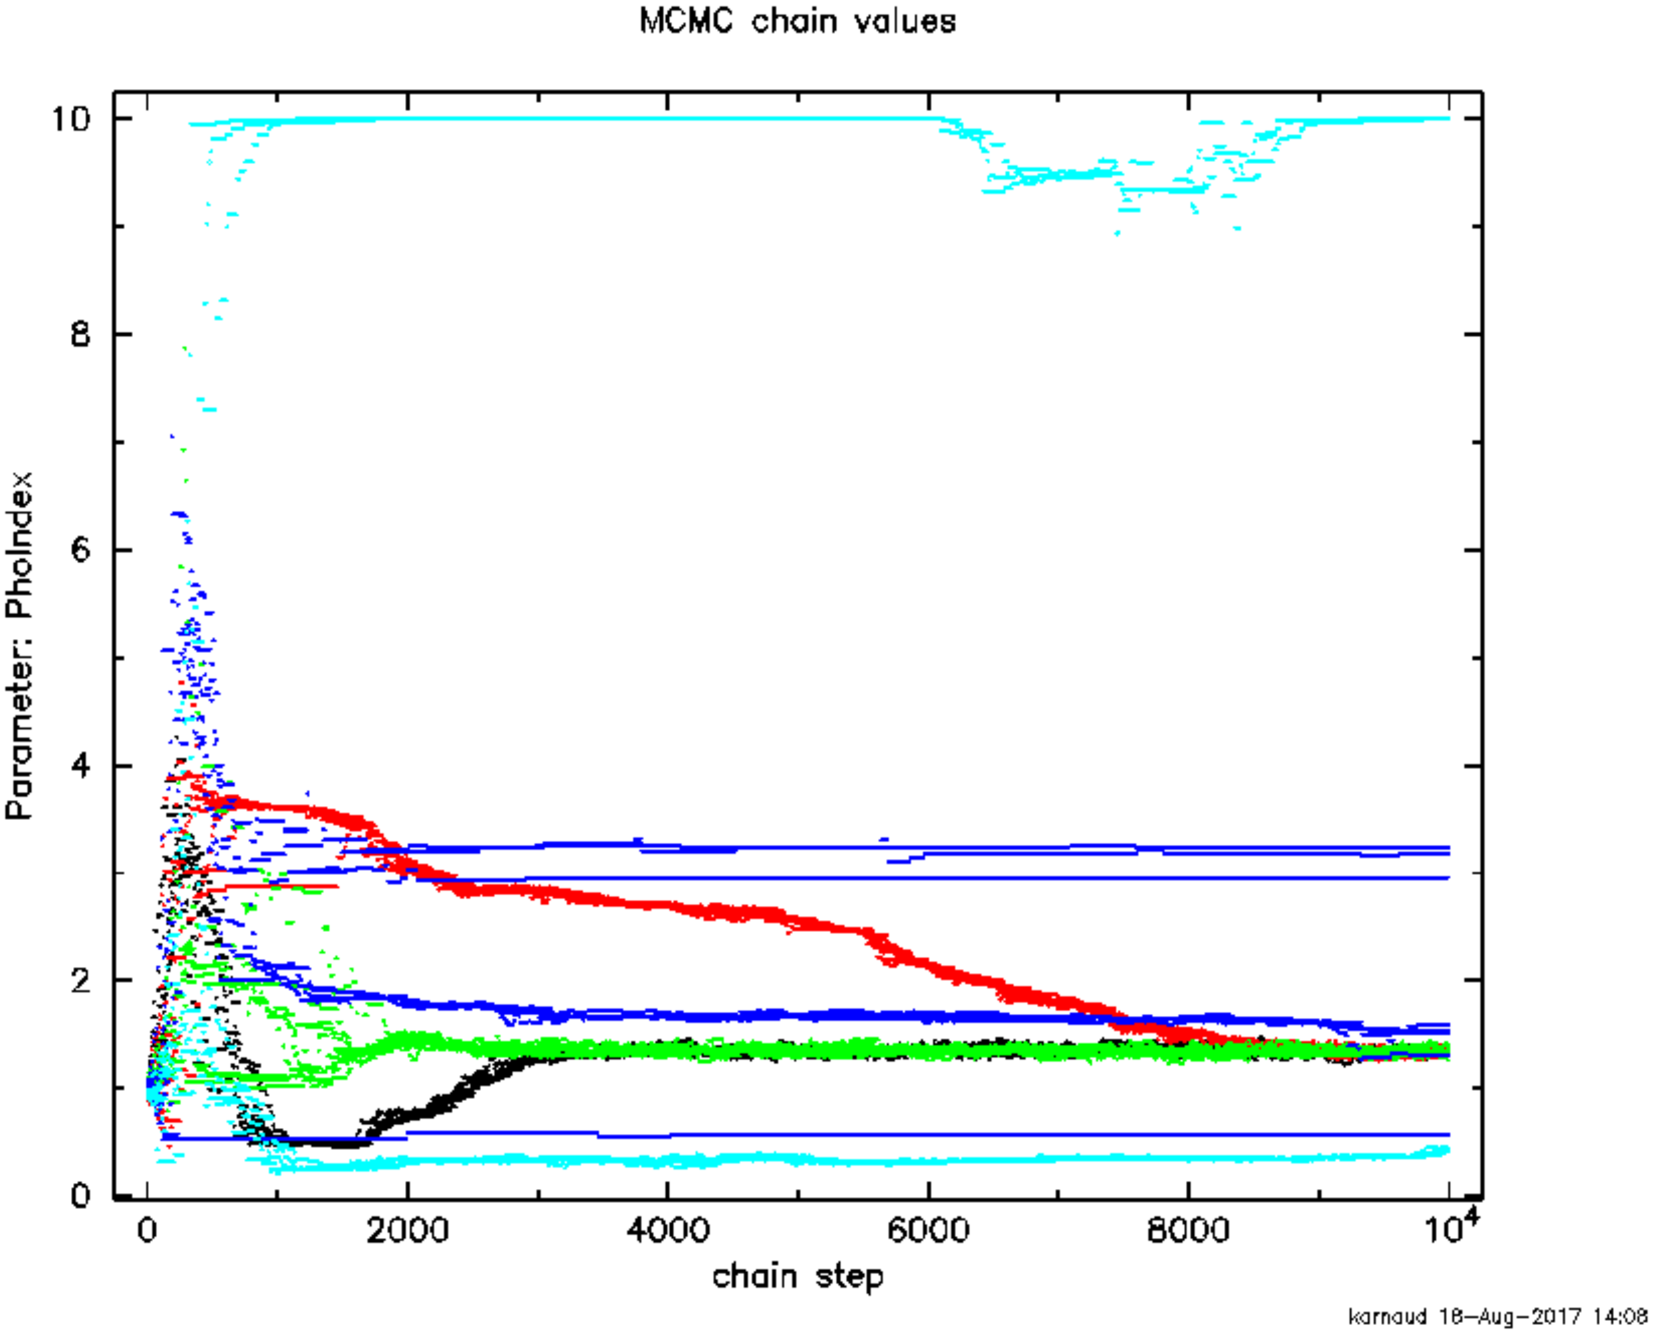
MCMC chain values
Parameter: PhoIndex
chain step
0
2
4
6
8
10
0
2000
4000
6000
8000
10^4
karnaud 18-Aug-2017 14:08

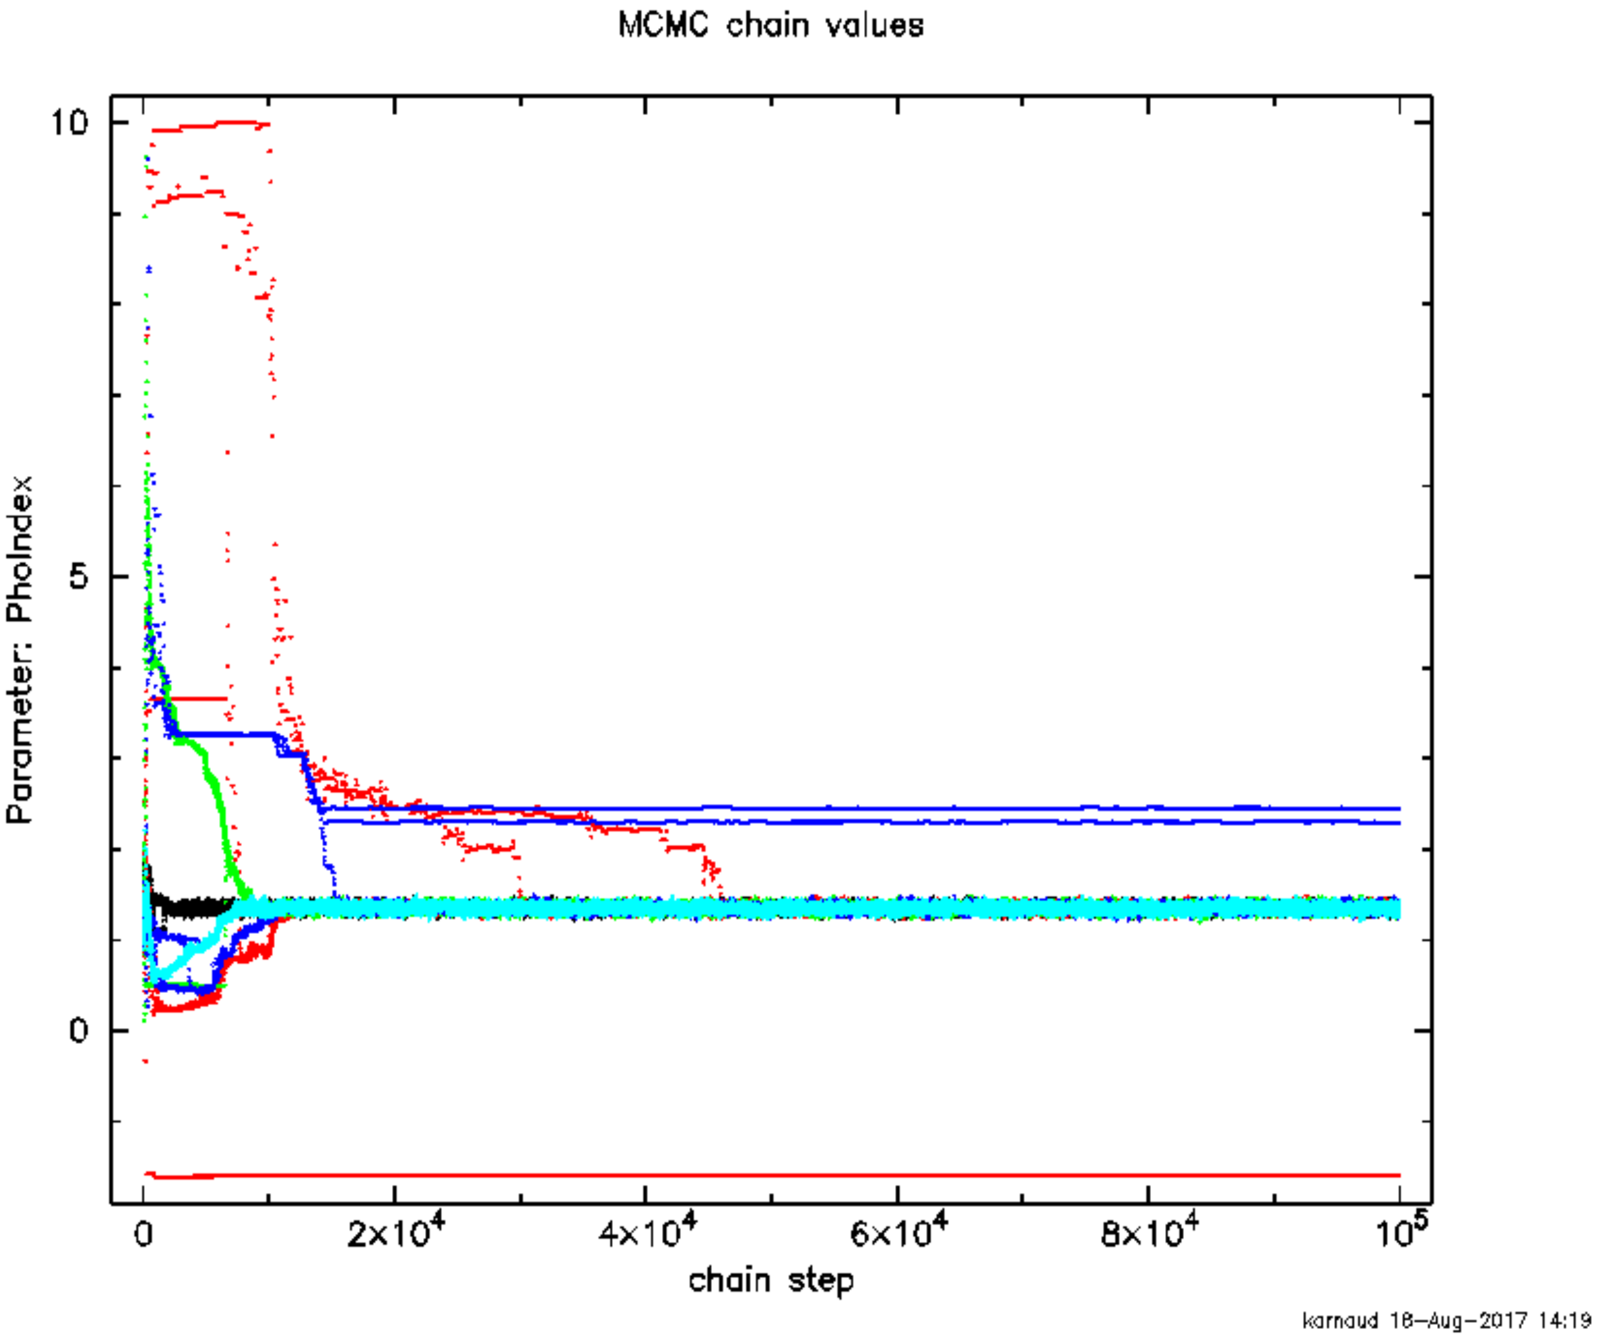
MCMC chain values
Parameter: PhoIndex
10
5
0
0
2×10^4
4×10^4
6×10^4
8×10^4
10^5
chain step
karnaud 18-Aug-2017 14:19

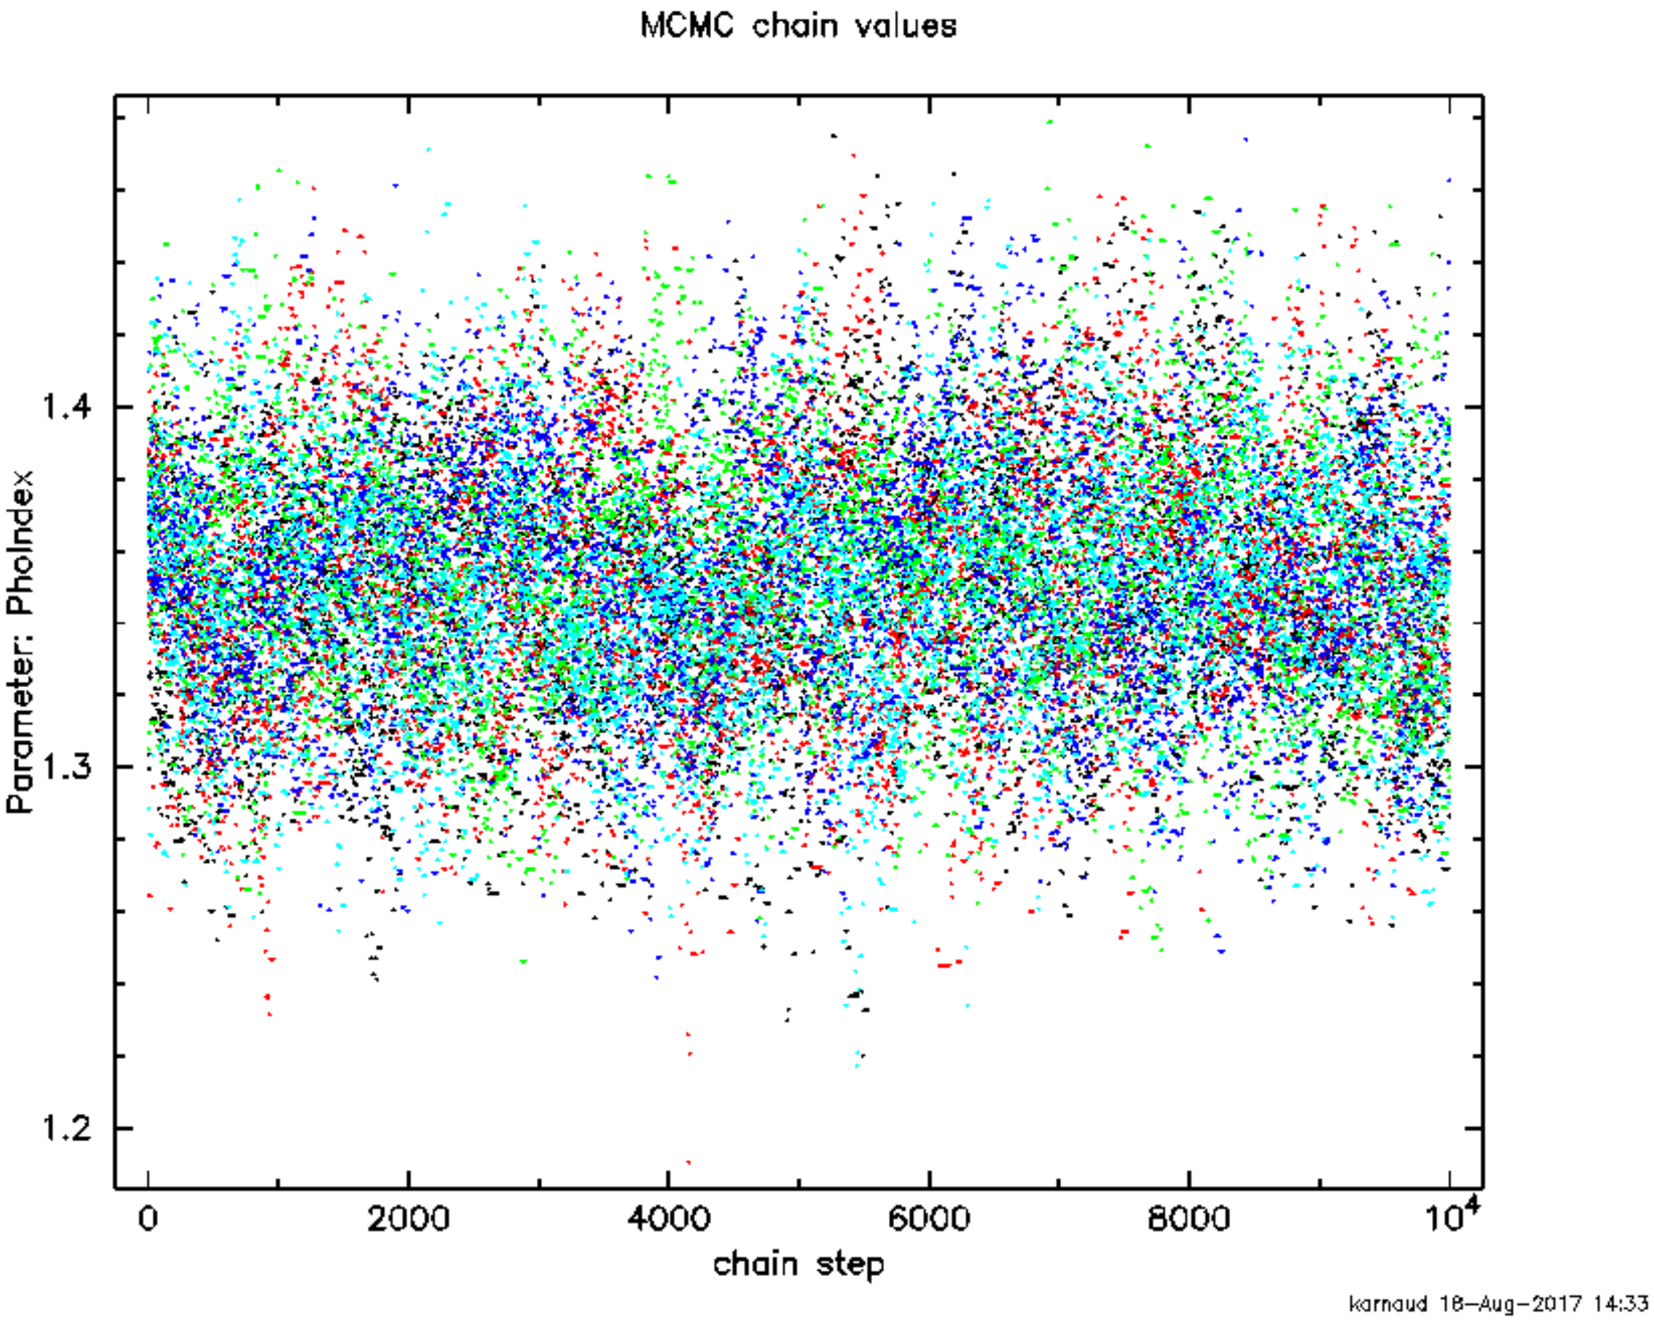
MCMC chain values
Parameter: PhoIndex
1.2
1.3
1.4
0
2000
4000
6000
8000
10^4
chain step
karnaud 18-Aug-2017 14:33

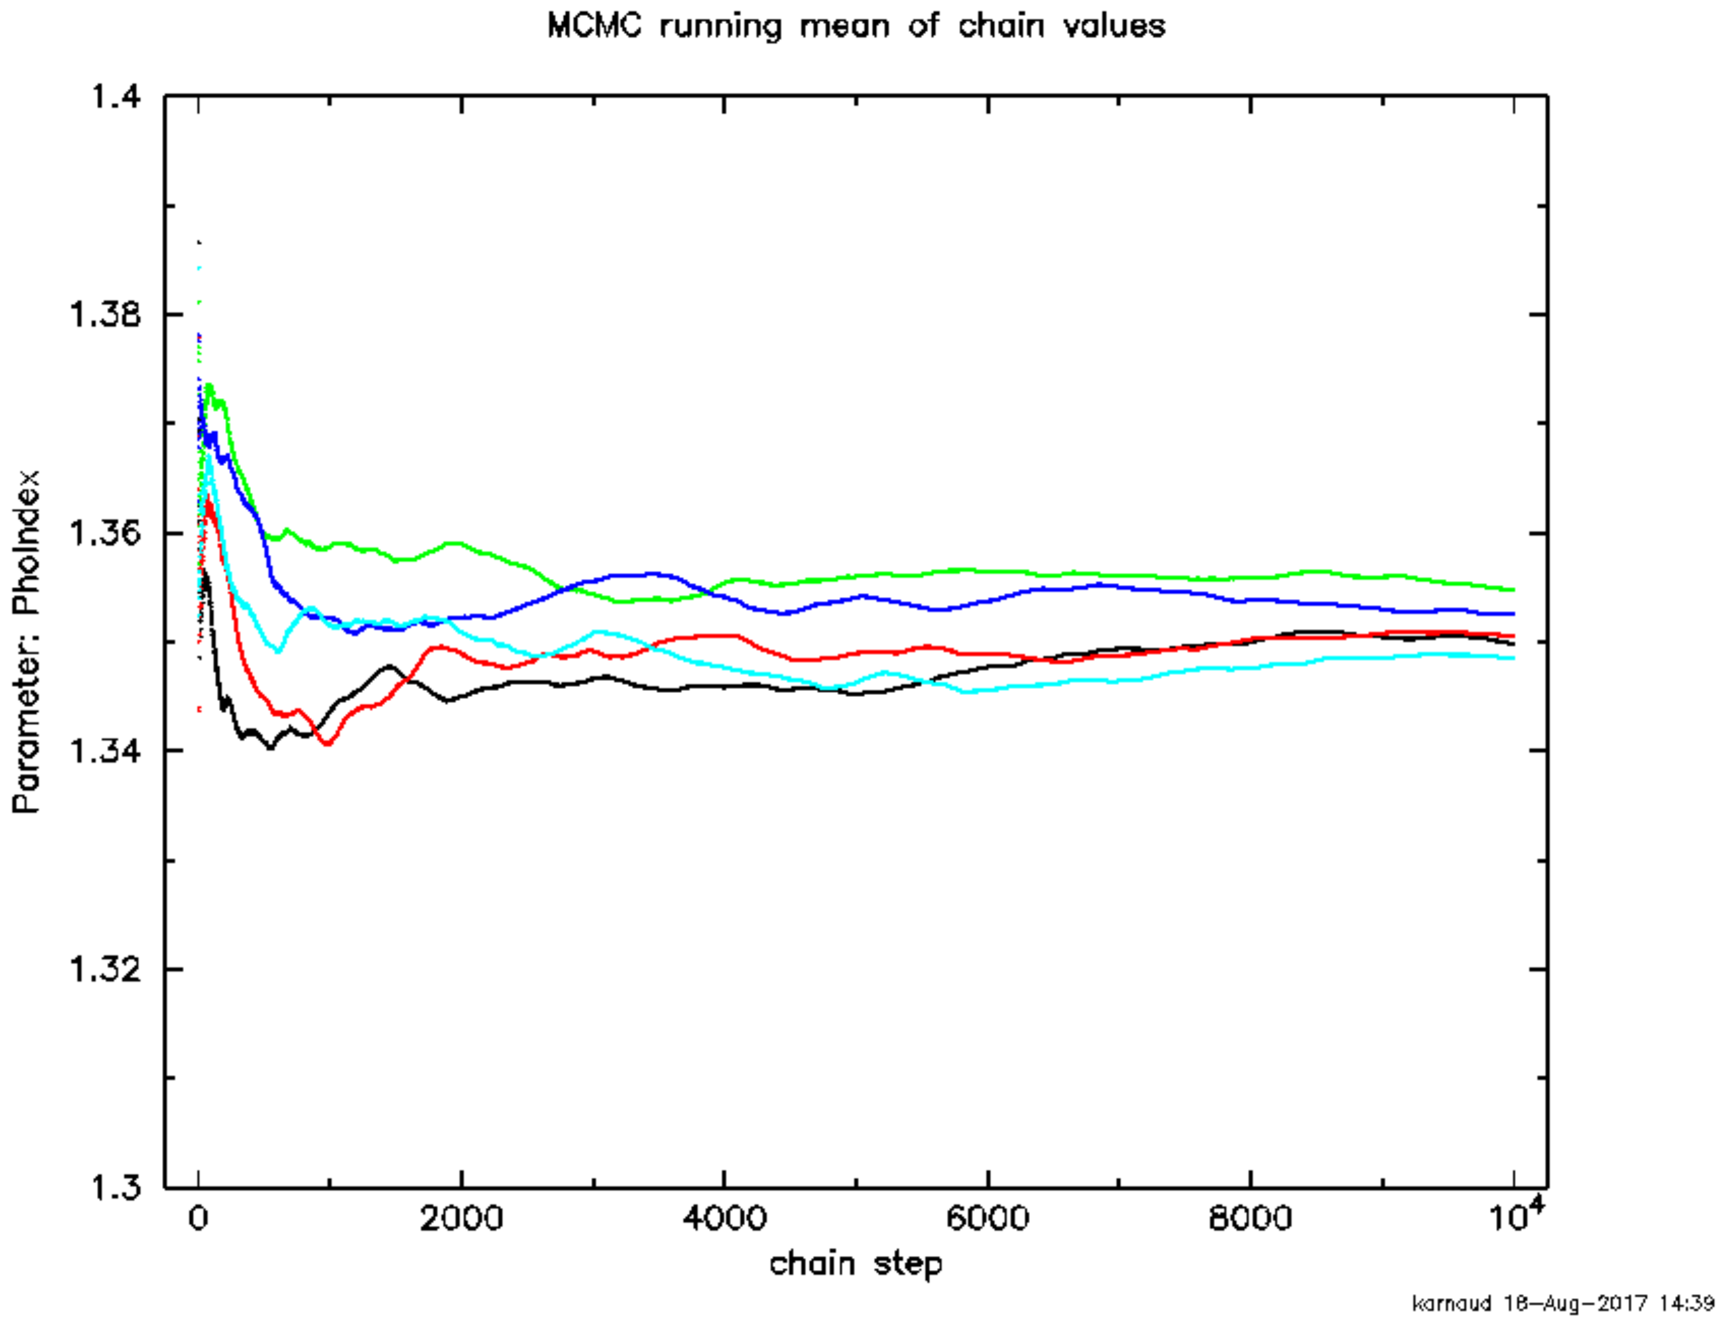
MCMC running mean of chain values
Parameter: PhoIndex
1.4
1.38
1.36
1.34
1.32
1.3
0
2000
4000
6000
8000
10^4
chain step
karnaud 18–Aug–2017 14:39

Using the chains

The chains are written out as they are calculated either as text or FITS files.

If chains are loaded XSPEC automatically uses them in the error, flux, and lumin commands.

tclout simpars and AllModels.simpars() can be used to extract a random set of parameters from the chain(s).

margin (with same arguments as steppar) gives probability density on a grid integrating over parameters you don't care about. Plot this using plot margin and plot integprob

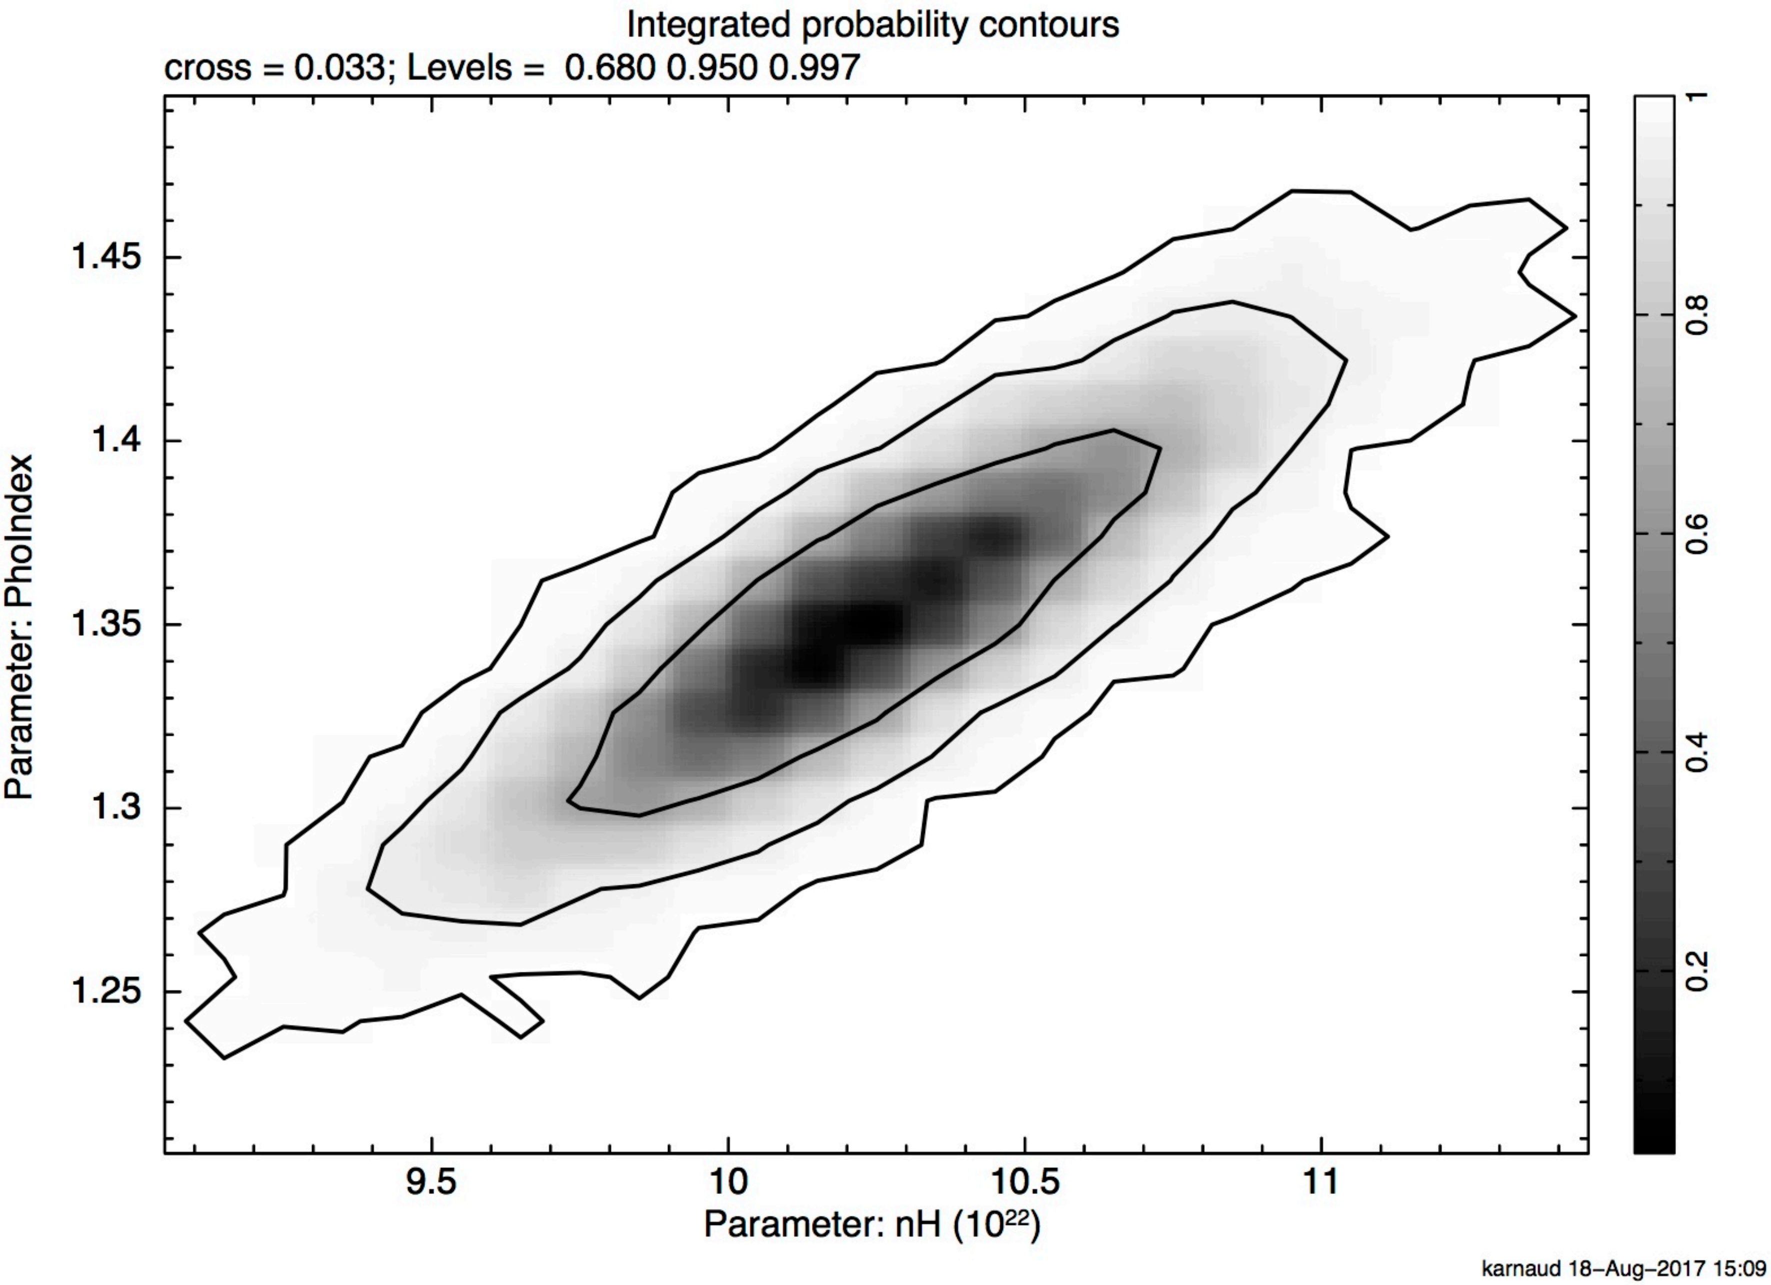
Integrated probability contours
cross = 0.033; Levels = 0.680 0.950 0.997
Parameter: PhoIndex
1.45
1.4
1.35
1.3
1.25
9.5
10
10.5
11
Parameter: nH (10²²)
1
0.8
0.6
0.4
0.2
karnaud 18–Aug–2017 15:09

An example – unabsorbed flux

The chain gives the probability distribution of the parameters hence can be used to derive the probability distribution of any function of the parameters.

Use simpars to generate a random set of parameter values

Set the model using these parameters. Set the Galactic absorption column to zero. Use the flux command and save the output.

Repeat about 10,000 times – you now have a sampling of the probability distribution of the unabsorbed flux.

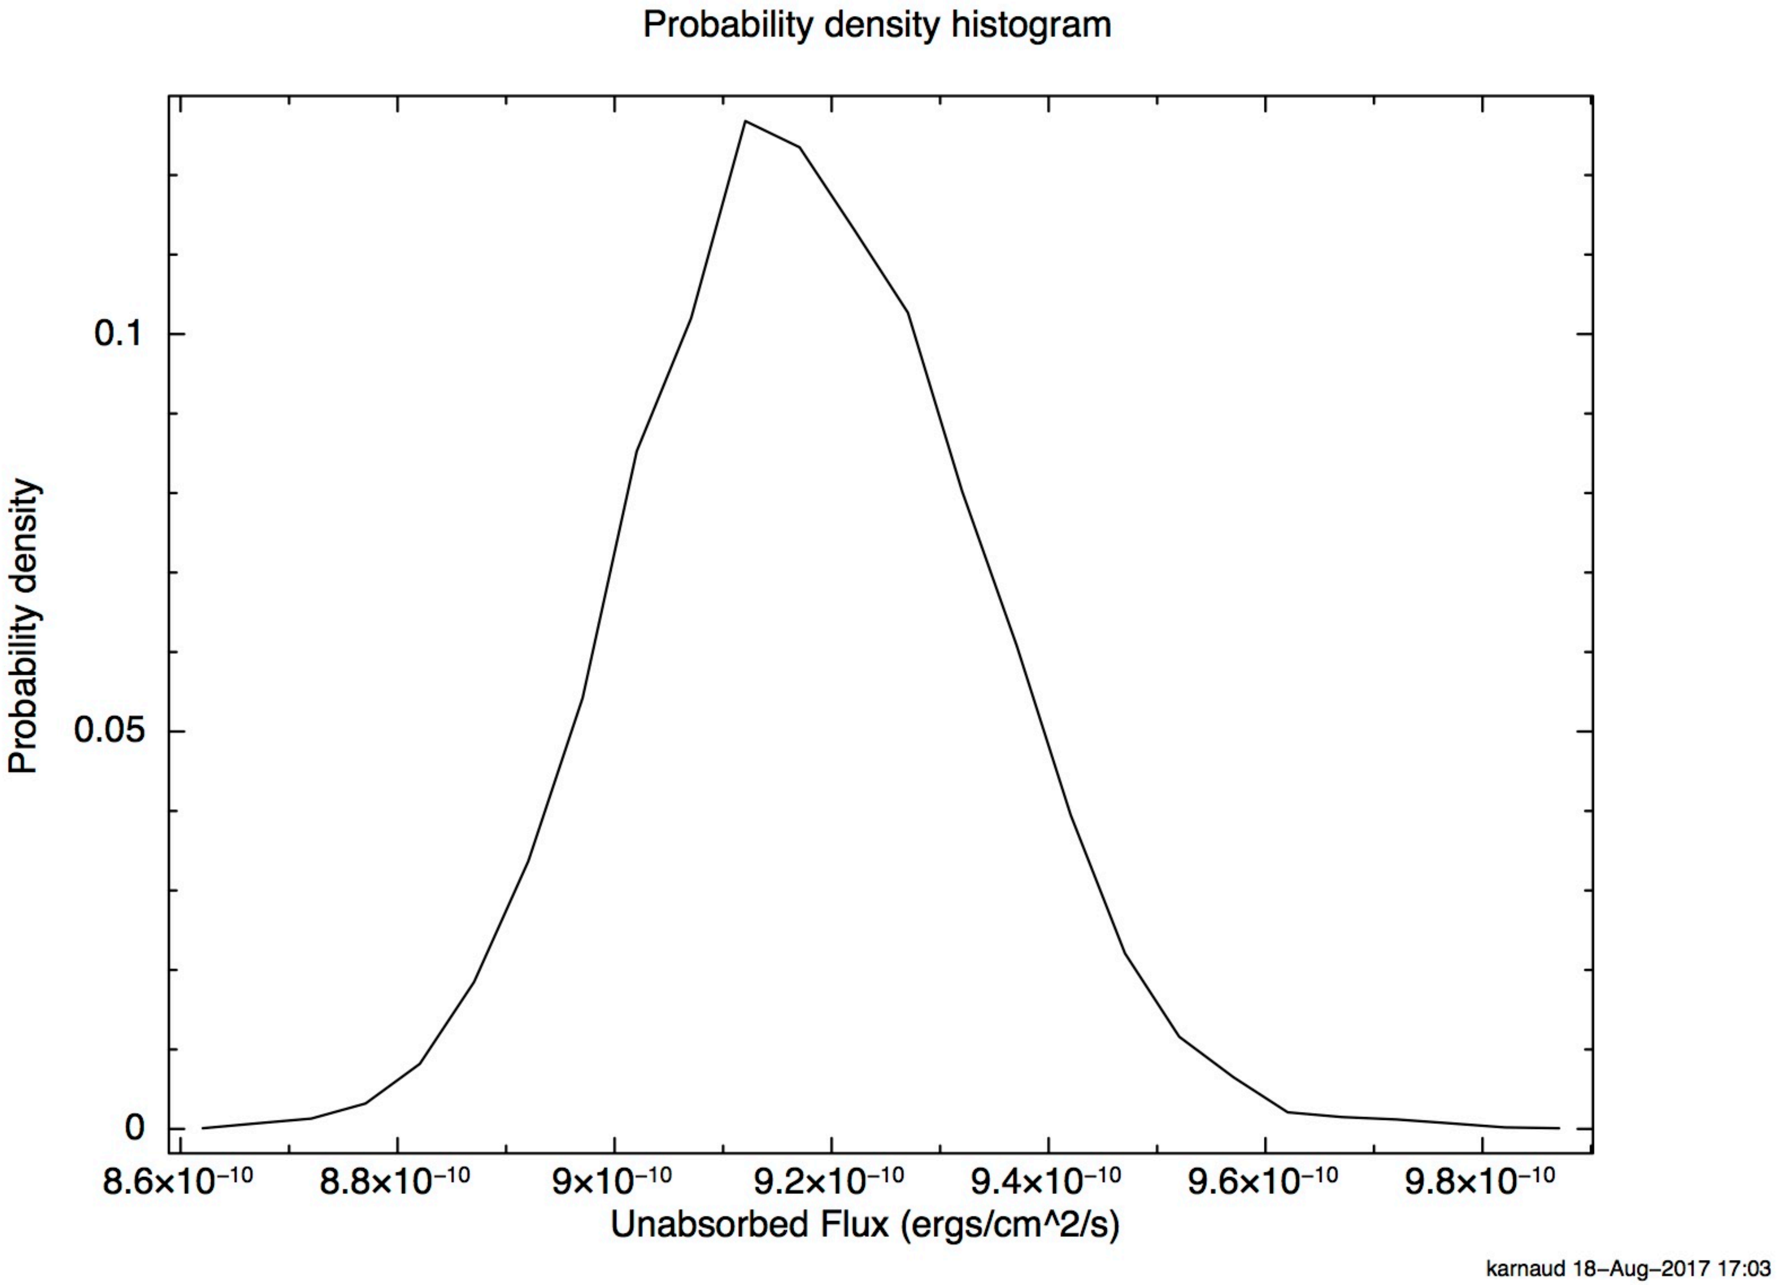
Probability density histogram
Probability density
0.1
0.05
0
8.6×10^−10
8.8×10^−10
9×10^−10
9.2×10^−10
9.4×10^−10
9.6×10^−10
9.8×10^−10
Unabsorbed Flux (ergs/cm^2/s)
karnaud 18−Aug−2017 17:03

Bayesian Priors

Individual parameters can be given priors – current options are constant, exponential, Jeffreys, and gaussian. e.g.

```
XSPEC12> bayes 5 gaussian 5 0.5
XSPEC12> bayes on
```

At present correlations between parameter priors cannot be included and hyperparameters cannot be fit.

Gelman's advice

If your paper depends on MCMC then start a new chain when you submit. By the time the referee's report comes back you will have a much longer chain available to answer any objections.

Wishlist

More general prior probability distributions. Define priors from an MCMC chain?

Integrate PyXspec with PyMC.

Simultaneous chains using parallel processing (Goodman-Weare option already has parallelization over walkers).

What else?

Thank you!

Questions?

Comments?