New sampler for cosmology and New results from the Lyman-alpha forest

Anže Slosar, Brookhaven National Laboratory

München, 10/15/13

Introduction

- I have been giving talks on Lyman-α forest for the past year - very very bored by now
- Still, everybody expects me to give a Lyman-α talk...
- First part: A new sampler for cosmology - yet another alternative to MCMC
- Second part: Lyman-α spiel

The problem

- Question: How to get marginalised constraints for cosmological parameters when you have likelihood has $N > 10$ dimension and each evaluation is computationally expensive ($>$1s)
- Answer: use `CosmoMC` which runs Markov Chain Monte Carlo (MCMC)
- MCMC is an algorithm that "walks" around the likelihood and produces *samples*
- Integrals over likelihood can be converted to sums over samples

Problem with MCMC

Markov Chain Monte Carlo does not scale very well:

- Scales perfectly for small number of chains, but not on modern architectures with 1000s of cores
- To run a `CosmoMC` chain you still run on 64 cores and wait for two days, instead of running on 10000 cores and wait 20 mins.
- But can't you run 1000 chains?
- Yes, but burn-in is a constant time process: one always needs to throw away some $\sim$ 1000 steps, because either:
  - You start chains randomly – they need to burn in
  - You start chains at high-likelihood region - they need to decorrelate
- Both are inefficient: You take 1000 samples to burn-in, but then 100 samples on each chain is enough to get 100,00 samples – quite inefficient

Importance Sampling

Assuming that one can sample from a known distribution, then one can weight samples to recover the effective sampling from a target distribution (whose properties one would like to study)

$$w_i = A \frac{L_{\rm t}(\mathbf{x}_i)}{L_{\rm s}(\mathbf{x}_i)}, \tag{1}$$

- Used to add a dataset to existing chains
- People tried to use it to sample cosmological likelihood directly using a Gaussian, but it fails miserably with bananas:
  - Either your Gaussian does not encompass the banana: weights blow up at the edges
  - Your Gaussian covers the banana, but also empty volume around it: most weights zero.

Why naive importance sampling doesn't work

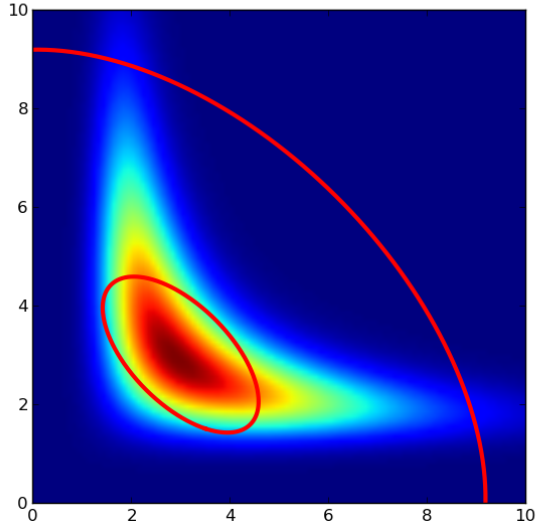

But if you could do something like this?

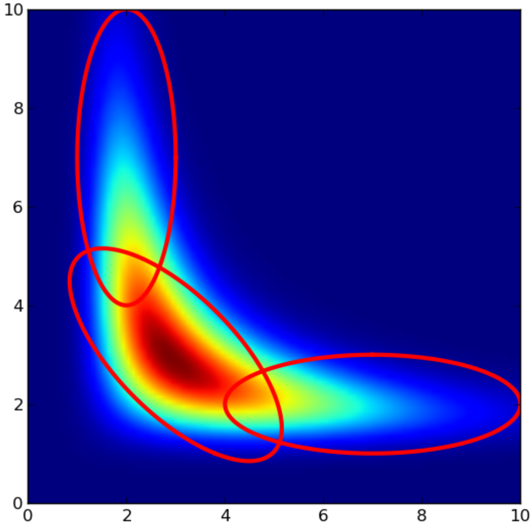

Guassian embedding sampler

1. Populate a list of Gaussians with a single Gaussian centered at a chosen starting point with suitable covariance.
2. Take N samples from the most recently added Gaussian in the list.
3. Calculate importance sample weights,

$$w_i = A \frac{L_{\mathrm{t}}(\mathbf{x}_i)}{\sum_{j=1 \ldots M} G_j(\mathbf{x}_i - \mu_j, \mathrm{C}_j)}, \tag{2}$$

4. Add new Gaussian at the position of the largest importance weight
5. GOTO step 2

Test 1: Gaussian

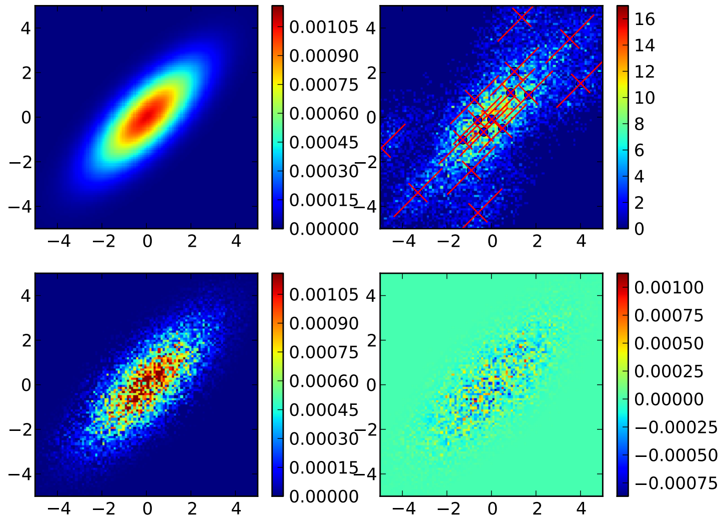

Test 2: Box

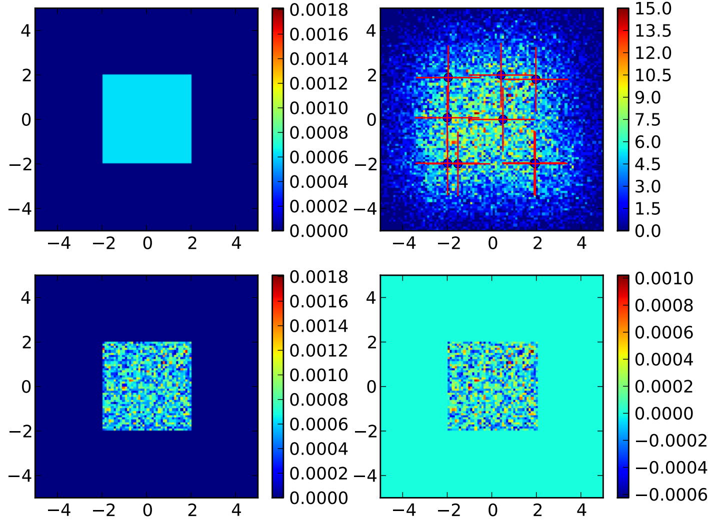

4
2
0
−2
−4
−4 −2 0 2 4
0.0018
0.0016
0.0014
0.0012
0.0010
0.0008
0.0006
0.0004
0.0002
0.0000
15.0
13.5
12.0
10.5
9.0
7.5
6.0
4.5
3.0
1.5
0.0
0.0010
0.0008
0.0006
0.0004
0.0002
0.0000
−0.0002
−0.0004
−0.0006

Test 3: Doughnut shaped banana

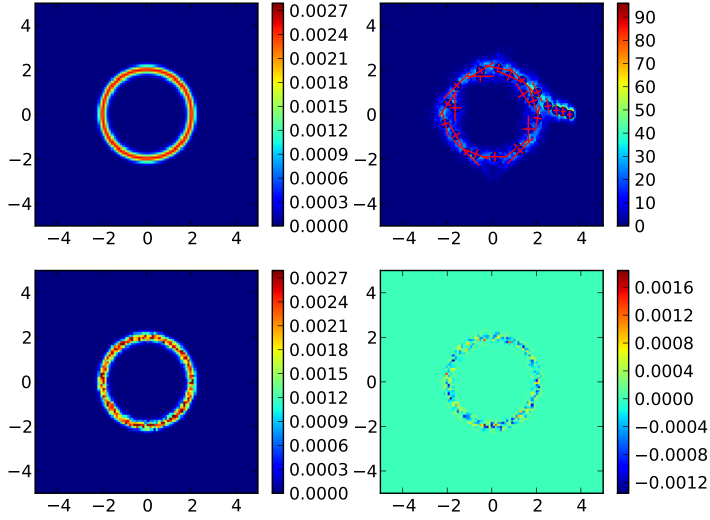

Convergence

Tried several convergence tests. Initially thought that demanding maximum weight to be less than unity would work: not clear what is the correct way to normalize weights.
One can define the effective number of samples

$$N_{\mathrm{eff}} = \frac{\sum w_i}{max(w_i)} \tag{3}$$

Demanding large N_{eff} proved to be very robust. If part of posterior not covered weights will blow up in that region, reducing the number of effective samples.

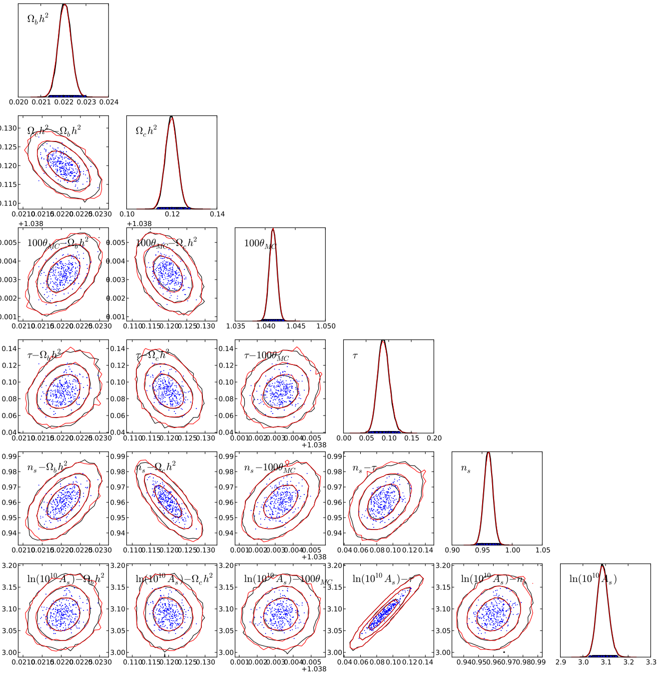
$\Omega_b h^2$
$\Omega_c h^2 - \Omega_b h^2$
$\Omega_c h^2$
$100\theta_{MC} - \Omega_b h^2$
$100\theta_{MC} - \Omega_c h^2$
$100\theta_{MC}$
$\tau - \Omega_b h^2$
$\tau - \Omega_c h^2$
$\tau - 100\theta_{MC}$
τ
$n_s - \Omega_b h^2$
$n_s - \Omega_c h^2$
$n_s - 100\theta_{MC}$
$n_s - \tau$
n_s
$\ln(10^{10} A_s) - \Omega_b h^2$
$\ln(10^{10} A_s) - \Omega_c h^2$
$\ln(10^{10} A_s) - 100\theta_{MC}$
$\ln(10^{10} A_s) - \tau$
$\ln(10^{10} A_s) - n_s$
$\ln(10^{10} A_s)$
+1.038

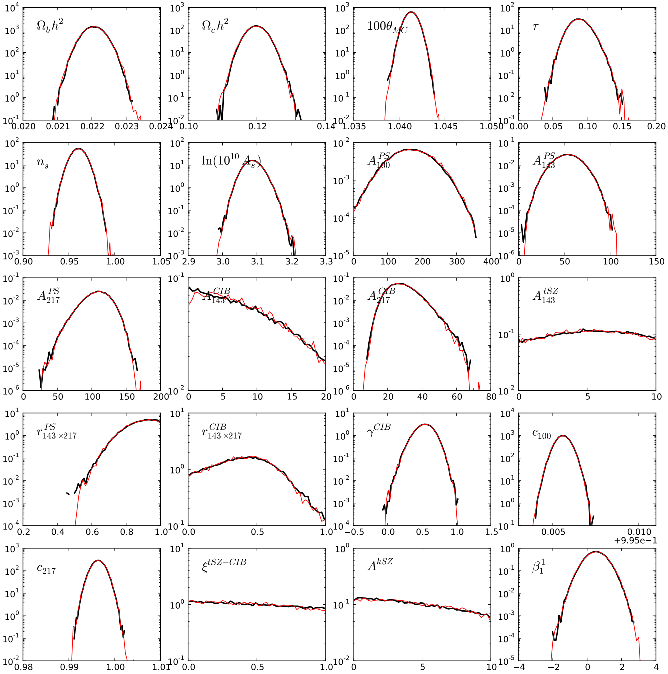

$\Omega_b h^2$
$\Omega_c h^2$
$100\theta_{MC}$
τ
n_s
$\ln(10^{10} A_s)$
A^{PS}_{100}
A^{PS}_{143}
A^{PS}_{217}
A^{CIB}_{143}
A^{CIB}_{217}
A^{tSZ}_{143}
$r^{PS}_{143\times217}$
$r^{CIB}_{143\times217}$
γ^{CIB}
c_{100}
+9.95e-1
c_{217}
$\xi^{tSZ-CIB}$
A^{kSZ}
β^1_1

Introduction

→

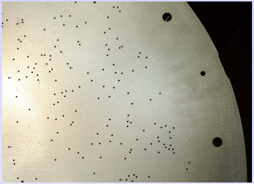

→

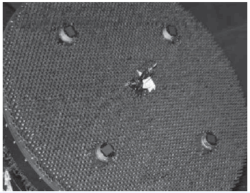

- Lyman-α forest is a unique field:
  - Got hammered with WMAP1 for running of spectral index
  - Very strong results in 2004 from 3000 SDSS QSOs by McDonald et al, but at a large emotional cost
  - Undergoing a revival now driven by BOSS
- Introduction to Lyman-α forest & the BOSS experiment
- Recent published BAO results
- Other interesting up-and-coming results

ALPINE 2.02(1266)   MESSAGE TEXT [A]   <RCF> archived   Msg 276 of 5,741 80% FWD +

Date: 9 Jan 2013 06:08:45 +0800
From: Am. J. Plant Sci. <ajps@elert-scirp.org>
Reply-To: ajps@scirp.org
To: anze@bnl.gov
Subject: Dr.An e Slosar, Special Issue on "The future of Forests" CFP

This message was sent to [anze@bnl.gov].Unsubscribe by clicking **here** [www.member.scirp.org].

Dear Dr.An e Slosar,

I am writing to let you know about our upcoming Special Issue on "The future of Forests", which will be published in American Journal of Plant Sciences (AJPS, ISSN:2158-2750), an open access journal. The deadline for submission is January 22nd, 2013, and the publication date is March 2013. You can find the Call for Papers for this Special Issue at the following website: **www.scirp.org/journal/ajps** [www.scirp.org]. The Special Issue is open to both original research articles as well as review articles.

I am contacting you about this Special Issue since I understand that you have published before in this area and I wanted to invite you to submit your manuscripts. American Journal of Plant Sciences (AJPS) is an open access journal, which means that all published articles are made freely available online without a subscription.

Please read over the journal's Authors' Guidelines [www.scirp.org] for more information on the journal's policies and the **submission process** [papersubmission.scirp.org]. Please include an indication of your intention to publish within the special issue to be entitled "Special Issue- The future of Forests". If you have any questions about this Special Issue, or about the journal, please do not hesitate to contact me.

Best regards,

On behalf of

? Help   < MsgIndex   P PrevMsg   - PrevPage   D Delete   R Reply
O OTHER CMDS   > ViewAttch   N NextMsg   Spc NextPage   U Undelete   F Forward

Measuring Density fields

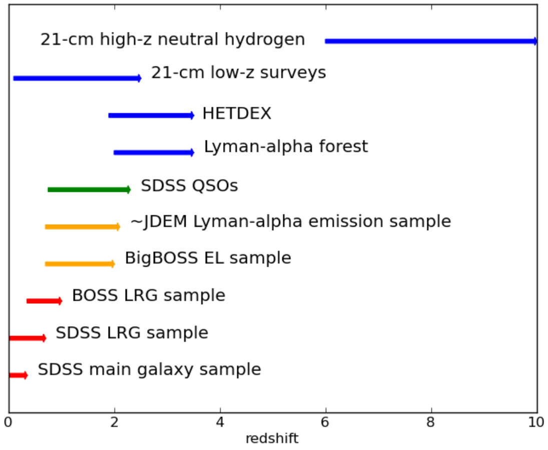

- To do cosmology, one needs a tracer of dark matter
- Lyman-α forest pretty unique in probing redhift 2-3 universe
- Volume probed is very, very large
- Systematics very different to galaxy surveys
- At $z < 2$ limited by forest moving into UV
- At $z > 3.5$ limited by faintness and number-density of quasars

What are quasars

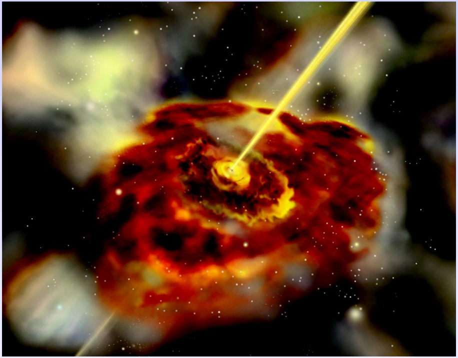

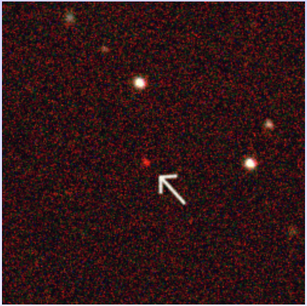

- Brightest things in the Universe
- Powered by energetic active galactic nuclei – can see them *very* far
- Featureless spectrum with a few broad emissions
- Understanding of underlying physics not important for our application.

Lyman-α forest

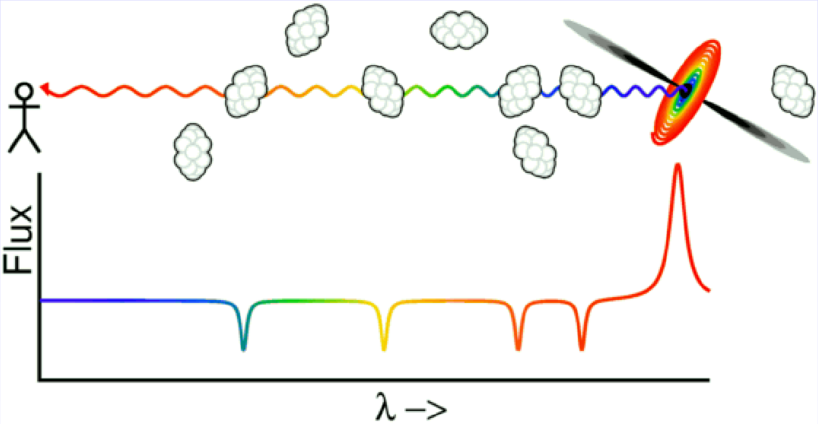

Neutral hydrogen absorbs light from distant quasars blue-ward of Lyman-α emission.

Lyman-α forest

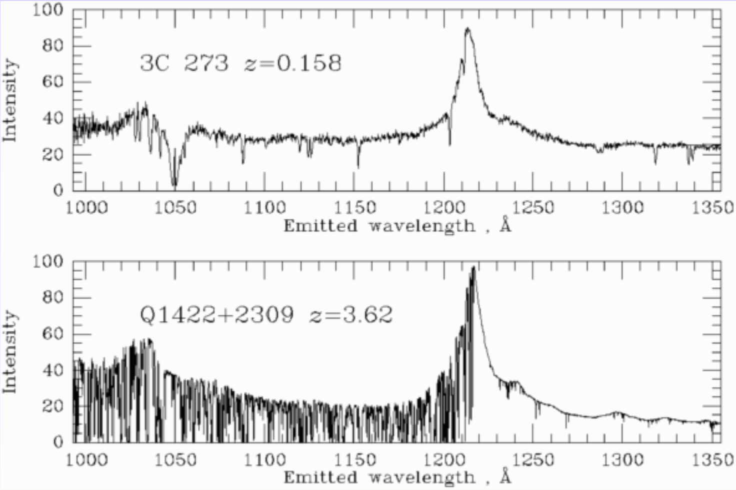

Neutral hydrogen absorbs light from distant quasars blue-ward of Lyman-α emission.

BOSS spectra

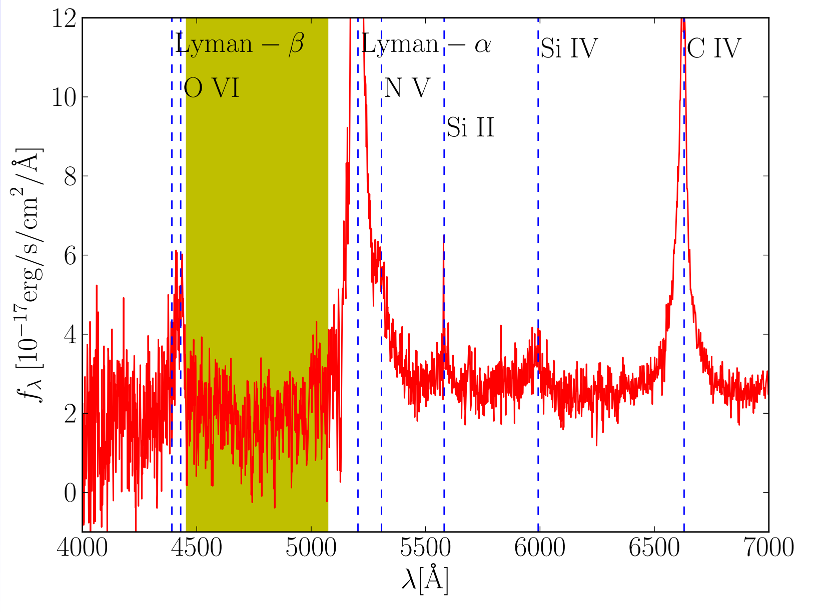

From baryons to flux

Absorption done by neutral hydrogen in photo-ionization equilibrium:

$$\Gamma n_{\rm HI} = \alpha(T) n_p n_e \tag{4}$$

$$n_{\rm HI} = \frac{\alpha(T)\rho_b^2}{\Gamma} \ll 1 \tag{5}$$

and so the absorbed flux fraction is given by

$$f = \exp\left(-\tau\right) \sim \exp\left(-A(1+\delta)^{1.7}\right) \tag{6}$$

What we are observing is a very non-linear transformation of the underlying density field.
On large scales this is simply a biased tracer. On small scales physics can be understood from first principles.

1D vs 3D

- Lyman-α forest is mapping the Universe through a very weird window function
- Historically: few and far apart high SNR measurements
- Quasars can be assumed independent in that limit: measure the 1D power spectrum of flux fluctuations
- With SDSS12: resolution down, noise up, quasar number up (from few tens to 15,000), but limited to 1D
- With SDSS3: resolution down, noise up, quasar number up (to 160,000): can finally measure correlations in three dimensions.

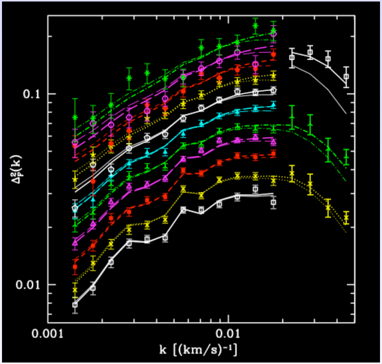

1D vs 3D

Power spectrum of Lyα measures:

- small scales (1D, $\sim 0.1\ \mathrm{Mpc/h}$): Effects of warm dark matter, sterile neutrinos, etc.
- medium scales (1D, $\sim 1\ \mathrm{Mpc/h}$): Inflation models, masses of light neutrinos, etc.
- large scales (3D, $> 10\ \mathrm{Mpc/h}$): Baryonic acoustic oscillations (dark energy, curvature of the universe), etc.

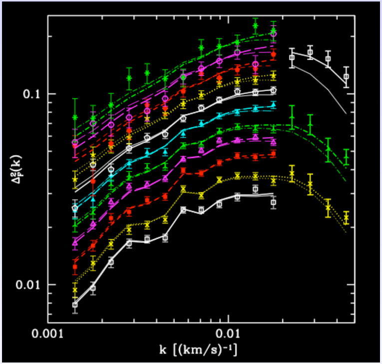

Baryon Oscillation Spectroscopic Survey (BOSS)

- BOSS is one of 4 experiments making up SDSS3.
- Uses 2.5m SDSS telescope
- Large etendue
- A 1000 fiber spectrograph
- Medium resolution: $R \sim 2000$
- Wavelength: 360nm (UV) 1000 nm (IR)

History of BOSS Lyman-α in 3D

- Nobody has done 3D Lyman-α to cosmological scales before BOSS
- We published first proof-of-concept paper in 2011
- Two papers with Lyman-α forest BAO appeared at the end 2012 with 60k quasars
- DR11 results (130k QSOs) already quite converged (but not yet public)
- Imagine what we could do if we had 20 million QSOs

3D sampling of the universe

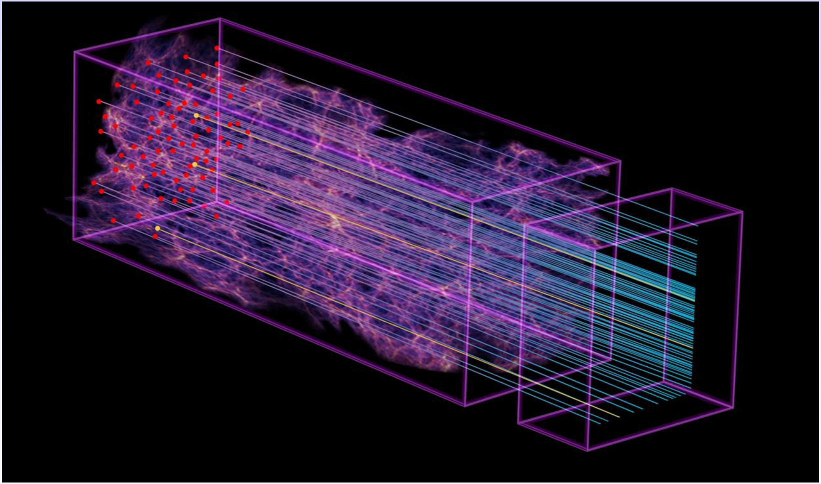

14k QSOs: ξ push

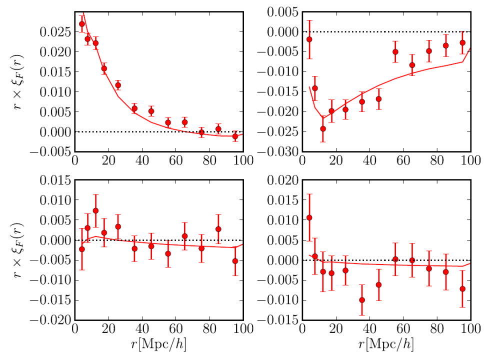

- Clear detection of correlations with no significant contamination
- The measured correlation function is distorted due to continuum fitting
- Analysis is harder than galaxy analysis:
  - Redshift-space distortions always matter
  - Redshift-evolution does matter

60k QSOs: BAO

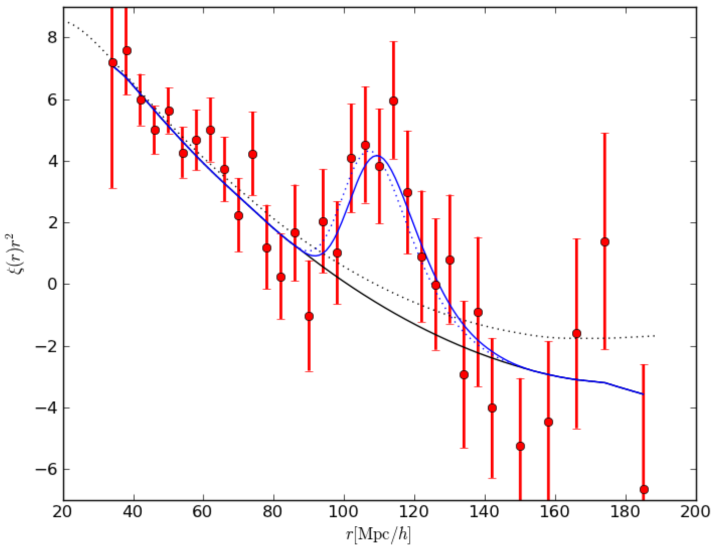

$\xi(r)r^2$
8
6
4
2
0
−2
−4
−6
20 40 60 80 100 120 140 160 180 200
$r[\mathrm{Mpc}/h]$

Cosmology fits

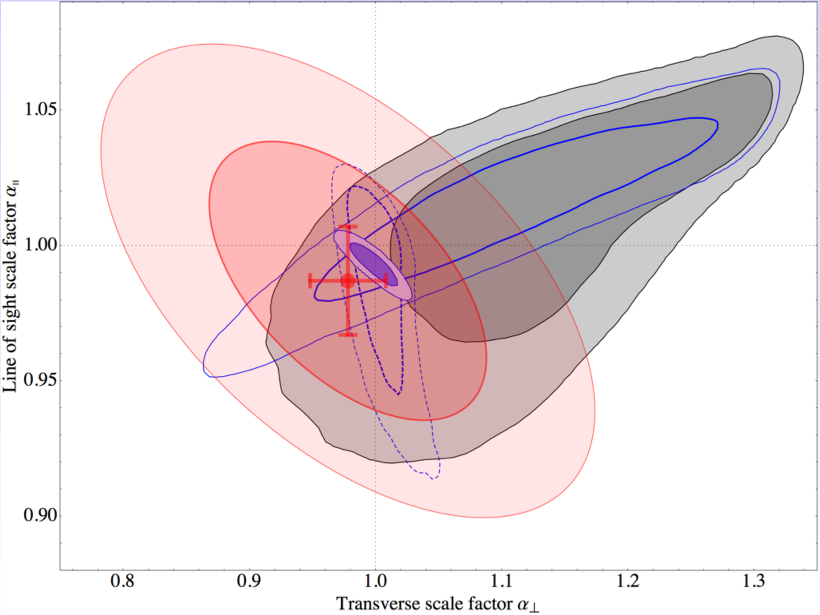

Distance plot

A cunning plot:

- Error-bars are distance errors
- bow-ties are Hubble-parameter measurements at central value: i.e. slopes
- Slanting of upper and lower errorbar is the correlations between parallel and perpendicular direction measurement.

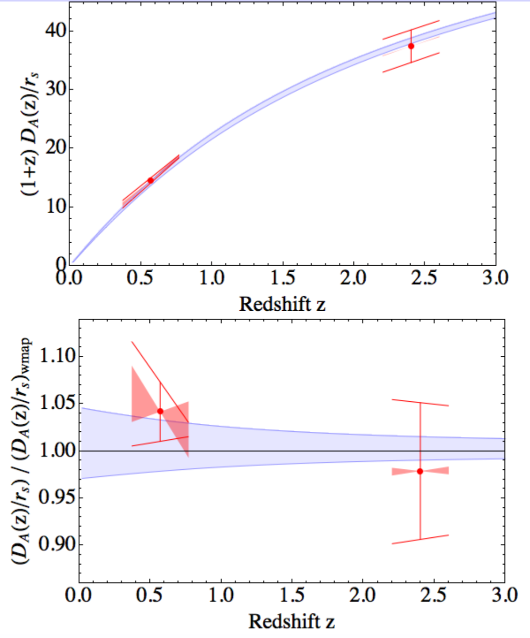

PRELIMINARY: quasar - forest cross-correlation

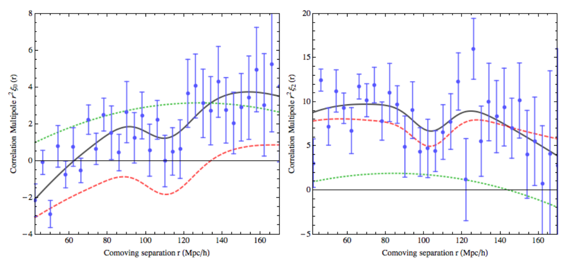

- Detection of the BAO in the cross-correlation between QSO and forest by Andreu Font & co.
- Ability for BOSS to do this has not been predicted, but constraining power nearly as powerful as with flux auto-correlation

1D power spectrum from BOSS

- Work done by Saclay group
- Palanque-Delabrouille et al, `arxiv:1306.5896`
- Selected $\sim 14,000$ quasars from $\sim 90,000$
- Using two methods: the FFT and likelihood maximization

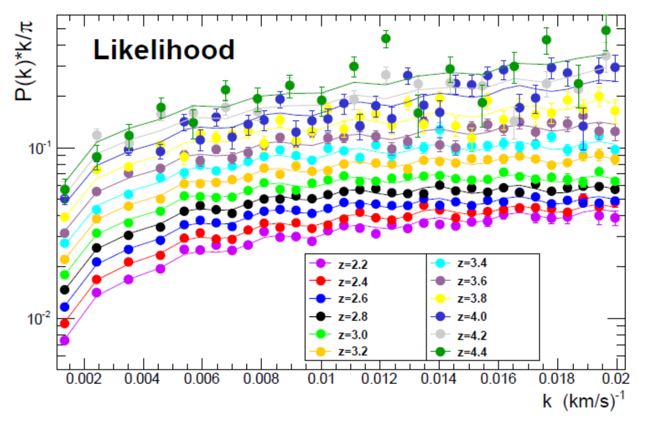

1D power spectrum from BOSS

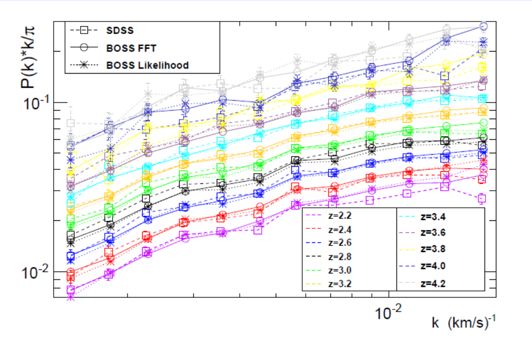

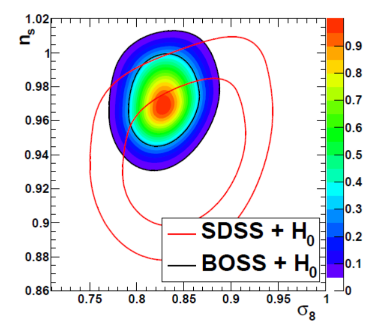

Latest WDM constraints

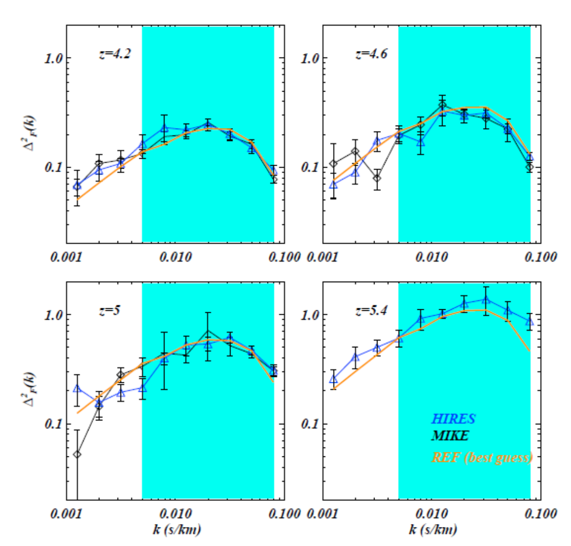

- Viel et al, `astro-ph/1306.2314`
- $m_{\rm WDM} > 3.3$keV at 2σ

Lyman-β forest

Why do we want to add Lyβ:

- The absorption cross section for Lyβ is smaller: $\sigma_\beta = \sigma_\alpha / 5.27$
- Lyβ forest measurements would be more sensitive to higher density and temperature where Lyα is saturated
- Increasing the effective path length in Lyα forest by nearly 20%

Transmitted flux:

$$F = \exp\left[-\tau\right] = \bar{F}\left(1 + \delta_F\right) = \bar{F}_\alpha \bar{F}_\beta \left(1 + \delta_\alpha\right)\left(1 + \delta_\beta\right)$$

In the linear approximation

$$\delta_F = \delta_\alpha + \delta_\beta + \delta_\alpha \delta_\beta$$

The cross correlation of two evolving fields yields both real and imaginary parts to the cross power spectrum $P^{tot}_{\alpha\beta} = P_{\alpha\beta} + iQ_{\alpha\beta}$

$$\xi_{\alpha\beta}(x,z) = \frac{1}{\pi}\int_0^\infty \left[P_{\alpha\beta}(k,z)\cos\left(kx\right) - Q_{\alpha\beta}(k,z)\sin\left(kx\right)\right] dk$$

Lyman-β forest

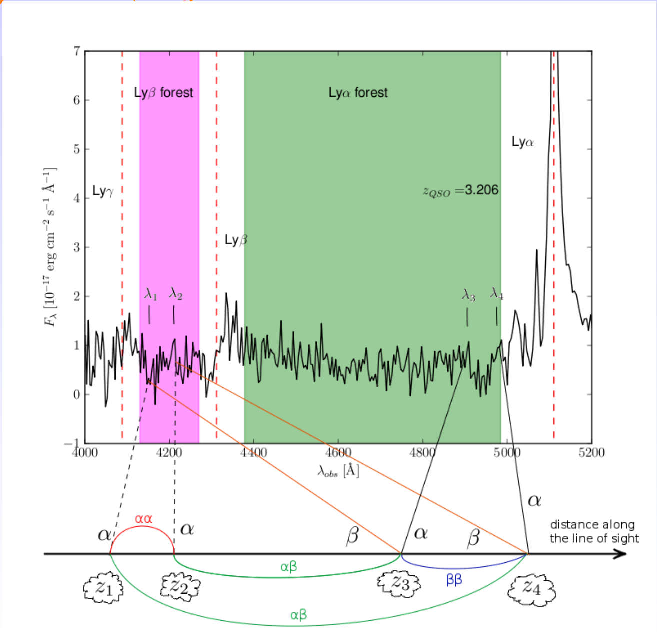

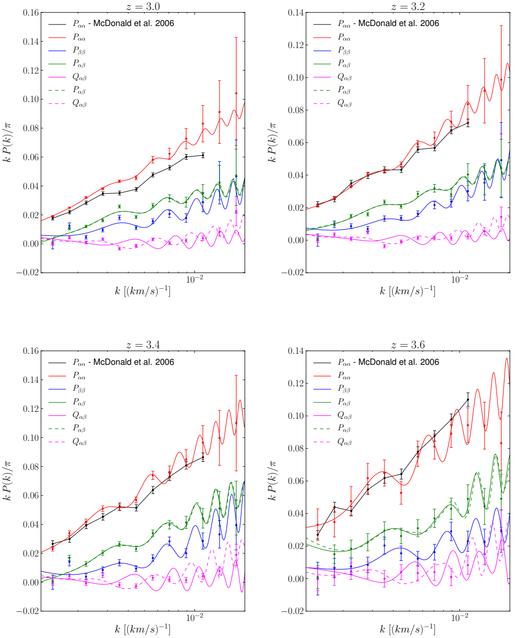

Similar to Si3 in Lyman-α , we see O5 in Lyman-β forest

Lyman-α simulation comparison

- Lyman-α suffers from a non-coherent simulation efforts - everyone has one or two Gadget runs...

- Running code comparison test with `Nyx` simulation code developed by LBL vs `Gadget3` run at BNL
- A subtle comparison analysis: smoothed particle hydro (SPH, `Nyx`) vs adaptive mesh refinement (AMR, `Gadget3`)

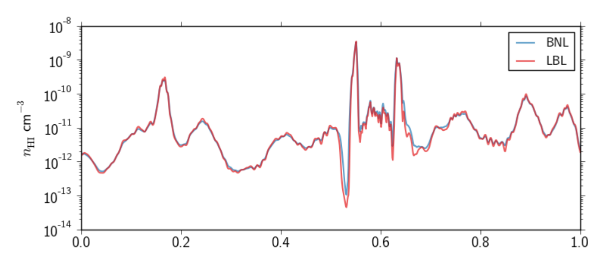

Comparison of neutral hydrogen density fluctutations along the line of sight for the same initial conditions

Large Scale Bias Parameters

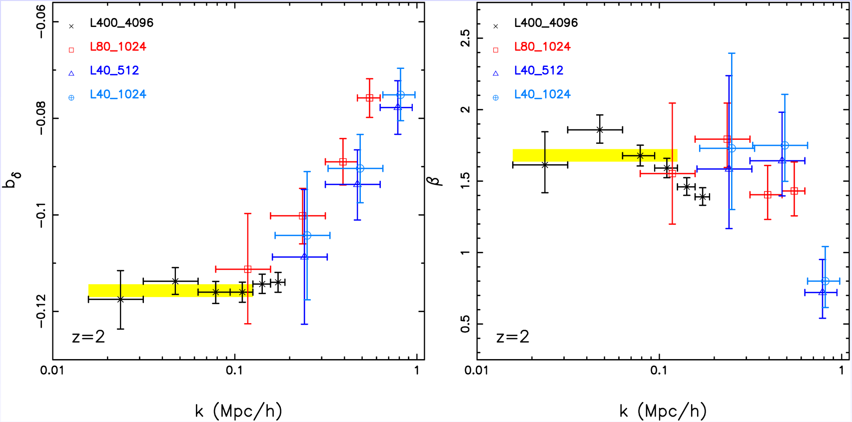

Large Scale Bias Parameters

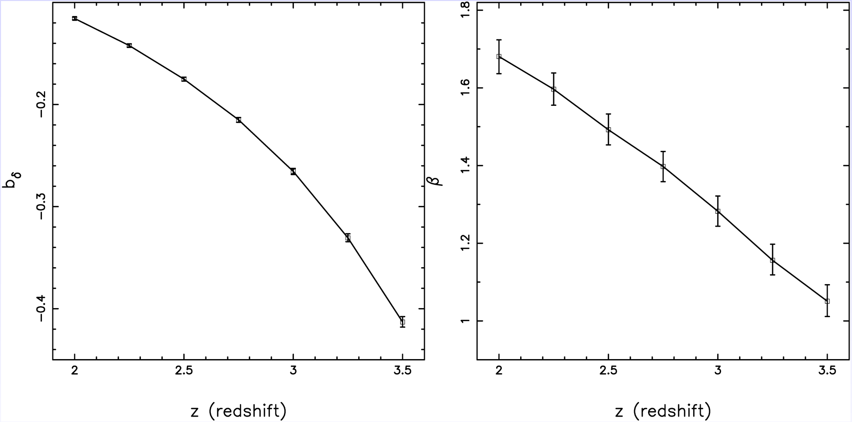

Large Scale Bias Parameters

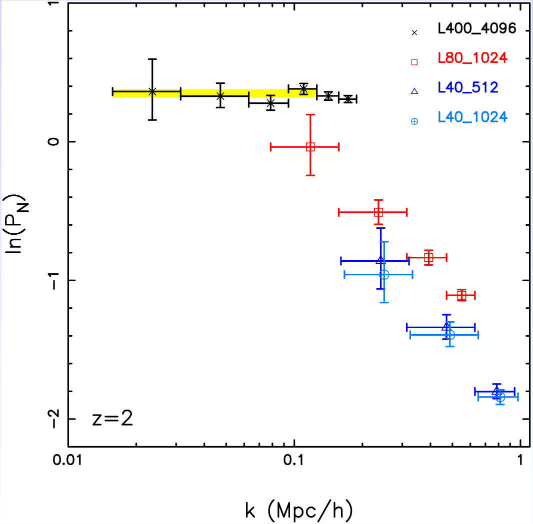

For a simple model

$$\delta_f = b(1 + \beta\mu^2)\delta + \epsilon$$

we get

$$\begin{aligned} P_{\delta f} &= b(1 + \beta\mu^2)P_{\delta\delta} \\ P_{ff} &= b^2(1 + \beta\mu^2)^2 P_{\delta\delta} + P_N \end{aligned}$$

Noise power converging very slowly. . .

Remember 2006?

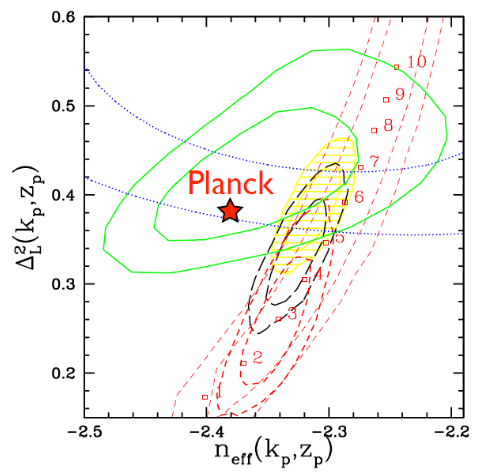

- Seljak et al, `astro-ph/0604335`
- With WMAP3, we head $\sum m_\nu < 0.17$eV at 2σ
- With Planck, constraint relaxes, but agreement much better.

Conclusions

- BOSS Lyman-α forest analysis progressing well
- BAO from forest-forest and forest-QSO correlations
- BOSS 1D power spectrum measured, we want to do something similar in 3D
- Much remains to be done regarding simulations
- Many other measurements possible: cross-correlations, 1D power spectra, Lyman-β forest, higher order correlations,. . .
- eBOSS / DESI will allow more of the same at unprecedented precision
- Neutrino masses and running of spectral index particularly interesting