Subject: Resubmission of 3C31 paper
From: Robert Laing <rlaing@eso.org>
Date: Sun, 17 Aug 2008 00:57:40 +0200 (CEST)
To: Alan Bridle <abridle@nrao.edu>, Paola Parma <parma@ira.inaf.it>, Matteo Murgia <murgia@ira.inaf.it>

Dear Alan, Paola and Matteo

Here is what I hope is very close to the final version of the 3C31 paper. MNRAS has a six-month deadline for resubmission, which is getting very close (approx Aug 24th). I'm not sure either of the exact date or whether there is scope for a short extension. It would help a lot if you could give me any final comments (or an agreement to send off this version) by the end of the week (22nd).

Regards

Robert

We thank the referee for a very thoughtful reading. As will become apparent, this has caused us to think carefully about our error analysis and to make significant changes in some areas. We think that the results are significantly more robust as a consequence.

This paper is a tour de force. The rotation measure and depolarization images in figures 2 and 3 are an impressive achievement. They require matched resolutions over a wide range of wavelengths and excellent image fidelity of a complex, mostly low surface brightness, source (to say nothing about accurate polarization calibration). And there are enough resolution elements across the source to make a concerted investigation of the statistical properties of the Faraday screen feasible and very worthwhile.

Because both the observational material and the analysis and simulations are new and extensive, the paper ends up being very long and perhaps tries to do too much for a single paper. The second half of the paper has a more didactic flavour and is, in part, a primer on the statistical analysis of random screens and how to simulate them. This is necessary because it is new, and such excellent observational material has not previously been available. The paper ends by applying these ideas to Hydra A. By now it has wandered somewhat from the title and original focus of the paper and the reader is exhausted. This suggests a natural division into two papers: the first on the magneto-ionic medium around 3C31, and the second on the statistical analysis, simulations and application to other sources, with room to expand as much as the authors wish (the writing is becomes quite compacted by the end of the present paper). This is just a suggestion. The authors may well not agree, and I leave it between them and the editor. My interest is simply in getting all parts of this work the attention they deserve.

We have considered this suggestion carefully. It has a number of attractions for us (even excluding the base motivation of another publication): we agree that there might be a broader audience for the basic observational results than for full details of the modelling; each paper would be shorter and more likely to be read by its target audience. However, we concluded that split papers would require a lot of repetition and cross-referencing, to the extent that the total would be much longer. We also felt that it would be very hard to understand some of the conclusions without at least speed-reading about the methods. It is instructive to compare with the two papers by Ensslin & Vogt (2003; A&A 412, 373 and 401, 835), which adopted precisely this split approach: there is indeed a lot of repetition. We therefore feel that the paper will be shorter in total and easier to read if left in its current form.

On reflection, we feel that we should offer a better route map in the abstract, which originally focused on the results, downplaying the methods.

We have therefore rewritten it. We have also changed the title in the hope that this will also give a more accurate guide to the content of the paper.

Detailed questions and comments, starting from the beginning:

Note that there is a new figure (4) and that the order of other figures has changed. We refer to the old numbering in what follows.

p2 paragraph 2, item (ii) last sentence: it would be helpful to name the sources.

We have now done this.

Some questions regarding Figures 2 and 3:

Although they donâ€™t say so, the agreement between 2b and 2c shows that the RM determined from the 4 Lband wavelengths, with their rather short lever arm, agrees very well with the 5 wavelength measurements. This is a testimony to the internal consistency and quality of the images. It also says that (with care) one can make RM maps â€œin a single bandâ€ if properly scaled array observations in a different band are not available. Not necessarily worth commenting on, but gratifying to note, nonetheless.

Indeed. We already said that "the two images are consistent with each other where they overlap" and we now also refer to this issue briefly in the section on further work, where we point out that the next generation of correlators will make in-band RM determinations much easier.

In Figure 3, the RM plots use different algorithms. PACERMAN clearly works better at lower SNR. A couple of sentences saying what it does differently, and why it does better at low SNR would be helpful. Should one always use it? Was it used in Figure 2?

We have added some text to give a short description of the Pacerman algorithm. We are not prepared to be dogmatic about the circumstances in which it should be used, as we have only tested it for a limited range of problems. It was not used in Fig 2., where we have higher s/n and very few problems with npi ambiguities (we say this specifically). It did not work very well for the S lobe of Hydra A, as Vogt et al. (2005) point out.

Figure 4 shows selected depolarization plots at 1.5 arcsec resolution. As they state, only 5e shows significant depolarization. Three other plots have good SNR but 5 are consistent with a range of slopes including zero. Does this mean that in Figure 3a we are mainly looking at the errors in the slope? It is clear from Figure 8f that there is still significant depolarization in the southern jet at 1.5 arcsecond resolution, but that is after some averaging. The difficulty with plots like Figures 2 and 3 is that every point has a different error (unlike an I map), and I donâ€™t know if there is any way to show that. It is compounded in 3a by dividing by p(0). One might infer from Figure 4 that perhaps half the pixels in 3a are essentially noise (depending strongly on location). I assume the authors have made sensible choices about clipping levels etc, but some discussion of the problem would be helpful.

The s/n is indeed low for any depolarization estimator at 1.5 arcsec resolution: we ae trying to measure a rather subtle effect. We have emphasised the spatial coherence evident in the south of the source, which indicates that the variations of depolarization are real. The referee's comments caused us to re-evaluate the use of p'(0)/p(0) as an estimator of depolarization. Modelling of the error distribution showed that it was rather asymmetric except in the regions of very high s/n, potentially biasing our averages. We therefore decided to fit a Burn law instead. This has two advantages: it avoids the need to divide the gradient by the zero-wavelength polarization to get a quantity physically related to depolarization and we expect it to be a better description of (most of) the data, as we already explained later in the paper.

Although this approach reduces the bias, modelling of the fitting process shows that there are still some residual problems. We have, conservatively, used only data with s/n > 4 in the degree of polarization for quantitative comparison with models and averaged profiles. This primarily affects the 1.5-arcsec data, as the lower-resolution images had s/n > 4 at essentially all of the points. We still show wider areas in Figs 2(a) and 3(a).

We now discuss these issues in the text.

As a consequence, we have:
- Replaced Figs 2(a) and 3(a) with the equivalent images of Burn k.
- Replaced the corresponding profiles in Fig 8 (now 9) e and f,
- Revised all of the relevant text
- Referred to the discussion of rotation and depolarization by an almost resolved foreground screen in Section 2.1, to motivate use of the Burn law.

The choice and order of Figures 4,5,6 is curious. Fig 6 might go better next to Fig 4, since they refer to the same resolution images in Fig 3. There are no plots of depolarization at 5.5 arcsec resolution to go with the present Fig5. They would show clearly the depolarization in the southern jet, and how it largely goes away at high resolution.

We think that the figures were in a sensible order, but we have responded to the referee's comment by adding a plot of p against \lambda^4 at low resolution. All four plots (as functions of \lambda^2 or \lambda^4 as appropriate) are now on the same page, allowing comparison either between p and position angle at the same resolution or between the same quantities at different resolutions. We have reduced the number of panels in the low-resolution plots to save space, as six examples are enough to make our points.

In section 2 they correctly emphasize the linearity of the RM plots as strong evidence for an external screen. Any internal rotation would cause deviations from linearity. Is it possible to put limits on that which are small enough to be interesting?

The limits are still a factor of >100 or so higher than the internal densities we infer from a conservation-law analysis (Laing & Bridle 2002 MNRAS 336, 1161), even in the optimistic case of a completely ordered magnetic field. The limits depend on the number of reversals in the field and the details of the geometry, of course. The limits from the lack of depolarization in the North jet are actually a bit more stringent, so we have given (approximate) results for both methods in a new subsection at the end of Section 2. We now give a rough estimate for a distance of about 30 arcsec from the nucleus, where we can compare directly with our earlier number. This is probably the most interesting location: further out, the increase in path length through the jet is essentially counterbalanced by the decrease in magnetic field (assumed to be close to the equipartition value) and the s/n on the position-angle measurements is worse.

The display of observational results continue with Figures 7 and 8, which are quite striking (especially 8 e and f), but are buried in the analysis section, far from the images they summarize. Can they be moved earlier? Perhaps they could be shrunk somewhat to facilitate that. (By the way, in the caption to Figure 8, is the â€œmainâ€ jet the north jet?)

This was a latex problem: the figure references were at the correct places in the text. We have moved figures around and changed shapes and sizes so that Figs 7 and 8 (now 8 and 9) appear closer to where they are referenced. We now say "the North jet" rather than "the main jet" in the caption to Fig 8 (now 9).

Section 3. The overview of the analysis is helpful.

In 3.1 (iii), they say the field is isotropic because there is no evidence for anisotropy in the RM distribution. But Figure 8a seems to show that the large RMâ€™s are in the south. Are they saying that any gradient is not significant compared to the fluctuations? In the middle of page 10 they say the RM in the north tail is primarily Galactic, which would seem to imply that there is indeed a local (to 3C31) gradient in RM. Or is a gradient different from an anisotropy?

We have clarified what we mean by "anisotropy" in this context, i.e. that the field has no preferred direction when averaged over a sufficiently large volume. We see no preferred direction in the RM variations on scales up to 100 arcsec or so. Given that we infer a power spectrum with significant amplitude on larger scales, we expect gradients: these will appear anisotropic due to imperfect sampling. In addition, of course, there are global variations in density.

We have rewritten the paragraph to clarify this point.

In 3.1 (iv) they assume that the amplitude of the magnetic field power spectrum varies with thermal gas density, but its shape is everywhere the same. The shape is determined by the slope and cutoffs in the spectrum. The characteristic length scales are interesting and there is not much discussion of them. A high frequency cutoff is expected if there is dissipation to dampen the fluctuations, and a low frequency cutoff or change of slope may be associated with the paddle that stirs the turbulence. The scales found here in section 4.5 are 4 kpc and 17 kpc. The closely related â€œmagnetic autocorrelation lengthâ€ depends on the poorly determined outer scale and is listed in Table 2. How do these numbers compare to what is expected or what is known about the cluster in which 3C31 lies? In other words, what physical significance should the reader attach to them?

We have added a short subsection at the end of Section 3 describing the various scales in the problem and their possible physical meaning, cautioning that there is no generally accepted theory. Energy input could occur over a range of scales from the size of the radio source down to the jet bending radius. Dissipation is expected to occur on the resistive scale, which is tiny compared with our beam. There is a possibility that a change of slope in the power spectrum might occur on the folding scale in some fluctuation dynamo models, and we give references for this idea.

The caption to Figure 11 is very confusing. It says these are simulated images, but the rest of the caption reads as if they are observed. If they are simulated, what should the reader compare them to? In the text it seems to imply that 11b is a simulation of 11a. But this canâ€™t be so because a stochastic process cannot reproduce any specific pattern in such detail.

We agree that this was confusing and have split the figure. Panels (a) - (c) all result from the same simulated dataset. (a) and (b) are a test of the short-wavelength approximation as described in the text and (c) is the associated depolarization. Panel (d) - now a separate figure - shows a portion of the observed RM distribution, included purely to show that it displays the same sort of artefacts as are seen in the simulation.

On reflection, we decided that the material on deviations from lambda^2 rotation and the wavelength dependence of polarization in the South of the source at 5.5 arcsec, while useful as a consistency check and not documented in the literature, was getting in the way of our main argument. We have therefore moved it to a new Appendix C.

On page 15, they point out that the Kolmogorov spectrum predicts too much depolarization on small scales because it doesnâ€™t have a high frequency cutoff. They do not want to add a cutoff because it would narrow the range of scales with a Kolmogorov slope, and dilute the reason for choosing it in the first place. I am not sure this is quite fair. If the depolarization on small scales requires a high frequency cut-off, even with a spectrum as steep as 11/3, then this is simply a result demanded by the observations. It is not evidence against the processes that might lead to a Kolmogorov slope at lower frequencies, and it is no more inelegant than cutting off the power law spectrum at high frequencies. In either case the cutoff tells us something about the physics of the Faraday screen (see above).

We no longer try to make this point: a more careful error analysis shows that both models for the RM power spectrum predict depolarizations consistent with those observed (to within 1 sigma) and we make this point clear.

On page 14 they quote best fit values for fbreak and qlow for the Kolmogorov spectrum. I would expect them to be quite strongly correlated. Is that the case? Or in other words, what is the range of values that fit the data adequately? The same question applies to the values of q and fmax for the broken power law.

This comment prompted us to formalize our fitting procedure, which had previously been done "by eye". We now minimise chi-squared (summed over the four non-overlapping source regions with good data and excluding SP2). The parameters we derive are slightly different from those we quoted originally and the improvement in fit was sufficient to cause us to redo subsequent steps in the analysis with the new values.

The values of q and fmax for the cut-off power law (we think that the referee means this, rather than broken power law) and fbreak and qlow for the broken power law are correlated in the expected sense: a flatter power spectrum requires a cut-off at a higher frequency. We now point this out and tabulate appropriate limits.

Figure 14 shows the simulated and observed RM measurements in selected regions. It would help the reader to be guided as to what features indicate this is a successful model, since there is obviously no one-to-one match. What would constitute an unacceptable simulation? Is there some sort of measure of â€œgoodness of fitâ€?

The structure function fit (using the error bars we derive from Monte Carlo simulations) gives a quantitative estimate of how well we describe the spatial statistics of a given region. However, this assumes that the RM is actually a Gaussian random variable. Our data are not extensive enough for us to look at higher-order correlations or other methods for the detection of non-Gaussianity, but a visual comparison is a good sanity check - the eye is very good at detecting correlations with structure and preferred directions in the data. This is why we show selected realisations.

We have made these points in the text.

In section 5 they are looking for a 3-dimensional model for the Faraday screen. They sometimes use a single scale model for ease of computing and sometimes the power law spectrum of fluctuations. It is not easy for the reader to figure out which model they are using in which simulations (Figures 17 and 18), and under what circumstances the single scale model is inadequate.

We have rewritten 5.2 to emphasise that a single-scale model does not always give us something which can be compared directly with observations. The reason is that it assumes averaging over many cells. With a realistic power spectrum, which has power on large scales, this assumption may be incorrect. For our data at 5.5 arcsec resolution, the assumption breaks down close to the nucleus (where the source is narrow) - this is why the simulations predict lower rms than the single-scale approximation close to the nucleus. If one averages over a large enough region, then the magnetic field strength can be derived from the single-scale approximation, setting the scale equal to the magnetic autocorrelation length (as we say). But we cannot average over large enough areas, because we are trying to determine changes in the normalization of the RM power spectrum across the source. A second, critical, problem is that we cannot derive realistic sampling errors, which are vital to our chi-squared analysis. Our main comparison is between simulations and observations, which automatically takes into account the irregular sampling and we use the single-scale model merely as a means of exploring parameter space simply and quickly.

We have moved the subsection after the description of the 3-d simulations so any residual confusion over fitting methods should be removed.

For Figure 17, I make the same comments as for Figure 14 (two paragraphs back).

Our criterion for an acceptable 3D model is that we fit both the shape and the normalization of the RM structure function to within errors set primariy by sampling variance. Given our 2D analysis, we know that we can fit the shape of the structure function with a power spectrum whose normalization varies across the source. To fit the normalization, we need to average over an area large enough to get a reliable number but small compared with the variations. Our best compromise is to fit the rms RM at 5.5 arcsec resolution. We now do this formally, again evaluating chi-squared between simulated and observed data using the rms of multiple realizations as the error bars. Given a geometrical model (e.g. a spherical distribution of hot plasma with a cavity excised) and a dependence of magnetic field on density, the only free parameter is the overall normalization, which then gives the field strength.

We now show the predicted and observed rms RM profiles superposed. We have used red and blue colours for clarity, but (following advice from our current referee on a previous occasion) we have also used different symbols for the predicted and observed points, to minimise problems for colour-blind readers.

For the spherical model, is there not automatically a cavity where the source is (even though the x-ray data do not show one)? If there isnâ€™t, then you have Faraday rotating material mixed in with the synchrotron emitting particles. Section 2.2 argues convincingly for the absence of internal Faraday rotation. If they have filled source with thermal plasma, how does that affect those arguments?

We agree that there should be a cavity, but emphasise that ours is the first analysis to treat the effects of a cavity on Faraday rotation quantitatively. We therefore want to retain the spherically-symmetric model for comparison with earlier work, and because the true extent of the cavity is unclear.

Figure 19 is confusing. If the tick mark at the bottom represents 100 arcsec, then the cone extends for 250 arcsec or so. Not the 140 arcsec stated in the text. What are the three concentric circles?

The original figure was plotted correctly, but it was perhaps not obvious that the two panels were on different scales (our attempt to indicate this was probably not clear enough). We now use ellipsoidal cavities, and have had to redraw the diagram. We now have the two panels on the same scale, and have noted in the caption that the arcs of circles represent isodensity contours.

The cavity model seems to generate the required asymmetry by virtue of the cone angle matching the angle to the line of sight. This seems contrived.

Yes, we agree. But this is necessary in order to generate the sharp change in RM fluctuation amplitude across the nucleus. Unless we have been misled by small-number statistics (which remains possible), a cavity with a wide opening angle is necessary to explain the observed profile of rms RM. We have switched to using an ellipsoidal cavity geometry, mainly for consistency with Hydra A (where we have X-ray images), but the qualitative point remains valid.

Note that there is significant diffuse emission at low surface brightness which may fill in the required volume. We have replaced Fig 1 with a different representation of the same image (also repeated from Laing et al. 2008) in order to emphasise this point.

The disk distribution seems equally ad hoc. A diagram similar to Figure 19 would be very
helpful here. Is there any precedence for a Faraday screen in the shape of an equatorial
disk? Does any galaxy have a large disk of material similar to what they visualize?

We do not favour the disk hypothesis, and in retrospect probably gave it too much prominence by putting it in a separate subsection and showing simulation results. We now refer to it more briefly in the "alternative explanations" subsection. Given that, we think that a sketch is not essential.

We are aware of one apparently elongated distribution of hot gas in a radio galaxy: 3C 403 (Kraft et al. 2005, ApJ 622, 149). Although this has only been detected out to ~5 kpc from the nucleus, it may well extend much further - the optical isophotes are very elliptical on larger scales. The measured ellipticity is ~0.6. It is possible that the gas is in a disk. We now mention this reference.

Otherwise, the closest observed analogues are probably the disks in Cen A (HI, molecular and ionized gas) and NGC 612 (HI, ionized gas and stars). We now refer to these briefly. We do not think it likely that 3C31 has a structure similar to either of these. The only circumstance in which such a disk would be missed in 3C 31 would be if it were close to face-on, in which case the geometry is unsuitable for the purpose of creating a Faraday rotation asymmetry, even if the magnetoionic properties of such a disk prove to be consistent with our observations.

Faraday-rotating "superdisks" in more powerful radio galaxies have been suggested previously [Gopal Krishna & Nath 1997, A&A 326, 45; Gopal Krishna & Wiita 2000, ApJ 529, 189]. It is not clear to us whether such structures really exist. Certainly, the measured spatial extent of the disk and molecular gas disk in 3C31 is orders of magnitude smaller than the radio lobes. But we now refer briefly to this possibility.

Section 5.7 contains reasonable suggestions for other explanations. At this point the 3dimensional modeling, realistically, has been inconclusive, and perhaps the abstract should reflect that. Given that, comparison to Hydra A seems premature and makes the paper too long.

Here we disagree: although the 3D modelling is not definitive, we think that the cavity idea is much the most plausible hypothesis - cavities are observed in X-rays, and this must mean that the density of thermal particles within them is low. In Hydra A (unlike 3C31), the geometry of the innermost cavities is well-determined from Chandra observations, allowing us to test the ideas developed for 3C31. Our models fit remarkably well, and we think that this section complements the work on 3C31. We therefore wish to retain Section 6, but we have clarified its relation to the rest of the paper.

We have thoroughly revised the paper in the light of the improved error analysis: this led us to make a fairly large number of minor changes to the text and diagrams without seriously modifying the structure of the paper or our conclusions.

3c31ls_2.pdf	**Content-Description:**	Paper
	Content-Type:	APPLICATION/PDF
	Content-Encoding:	BASE64

---

response.txt	**Content-Description:**	Response to referee
	Content-Type:	TEXT/PLAIN
	Content-Encoding:	BASE64

Subject: Re: 3C31 paper 2 - work so far
From: Robert Laing <rlaing@eso.org>
Date: Wed, 5 Mar 2008 18:38:54 +0100
To: Alan Bridle <abridle@nrao.edu>

Dear Alan

I've attached my working version of the paper and response. Still quite a few things to do, but please comment on alterations so far.

Regards

Robert

This paper is a tour de force. The rotation measure and depolarization images in figures 2
and 3 are an impressive achievement. They require matched resolutions over a wide range

of wavelengths and excellent image fidelity of a complex, mostly low surface brightness,
source (to say nothing about accurate polarization calibration). And there are enough resolution elements across the source to make a concerted investigation of the statistical
properties of the Faraday screen feasible and very worthwhile.

Because both the observational material and the analysis and simulations are new and extensive, the paper ends up being very long and perhaps tries to do too much for a single
paper. The second half of the paper has a more didactic flavour and is, in part, a primer on
the statistical analysis of random screens and how to simulate them. This is necessary because it is new, and such excellent observational material has not previously been available. The paper ends by applying these ideas to Hydra A. By now it has wandered somewhat from the title and original focus of the paper and the reader is exhausted. This
suggests a natural division into two papers: the first on the magneto-ionic medium around
3C31, and the second on the statistical analysis, simulations and application to other sources, with room to expand as much as the authors wish (the writing is becomes quite compacted by the end of the present paper). This is just a suggestion. The authors may well
not agree, and I leave it between them and the editor. My interest is simply in getting all
parts of this work the attention they deserve.

We have considered this suggestion carefully. It has a number of attractions for us (even excluding the base motivation of another publication): we agree that there might be a broader audience for the basic observational results than for full details of the modelling; each paper would be shorter and more likely to be read by its target audience. However, we concluded that split papers would require a lot of repetition and cross-referencing, to the extent that the total would be much longer. We also felt that it would be very hard to understand some of the conclusions without at least speed-reading about the methods. It is instructive to compare with the two papers by Ensslin & Vogt (2003; A&A 412, 373 and 401, 835), which adopted precisely this split approach: there is indeed a lot of repetition. We therefore feel that the paper will be shorter in total and easier to read if left in its current form.

On reflection, we feel that we should offer a better route map in the abstract, which currently focuses on the results, downplaying the methods.
....

We have also changed the title in the hope that this will give a more accurate guide to the content of the paper.

Detailed questions and comments, starting from the beginning:

p2 paragraph 2, item (ii) last sentence: it would be helpful to name the sources.

We have now done this.

Some questions regarding Figures 2 and 3:

Although they donâ€™t say so, the agreement between 2b and 2c shows that the RM determined from the 4 Lband wavelengths, with their rather short lever arm, agrees very

well with the 5 wavelength measurements. This is a testimony to the internal consistency
and quality of the images. It also says that (with care) one can make RM maps â€œin a single
bandâ€ if properly scaled array observations in a different band are not available.
Not necessarily worth commenting on, but gratifying to note, nonetheless.

Indeed. We now refer to this point briefly in the section on further work, where we point out that the next generation of correlators will make this much easier.

In Figure 3, the RM plots use different algorithms. PACERMAN clearly works better at lower SNR. A couple of sentences saying what it does differently, and why it does
better at low SNR would be helpful. Should one always use it? Was it used in Figure 2?

We have added some text to give a terse description of the Pacerman algorithm. We are not prepared to be dogmatic about the circumstances in which it should be used, as we have only really tested it for a limited range of problems. It was not used in Fig 2., where we have higher s/n and very few problems with npi ambiguities (we say this specifically). It did not work very well for the S lobe of Hydra A, as we note.

Figure 4 shows selected depolarization plots at 1.5 arcsec resolution. As they state, only 5e shows significant depolarization. Three other plots have good SNR but 5 are consistent with a range of slopes including zero. Does this mean that in Figure 3a we are
mainly looking at the errors in the slope? It is clear from Figure 8f that there is still
significant depolarization in the southern jet at 1.5 arcsecond resolution, but that is after
some averaging. The difficulty with plots like Figures 2 and 3 is that every point has a
different error (unlike an I map), and I donâ€™t know if there is any way to show that. It is
compounded in 3a by dividing by p(0). One might infer from Figure 4 that perhaps half the
pixels in 3a are essentially noise (depending strongly on location). I assume the authors
have made sensible choices about clipping levels etc, but some discussion of the problem
would be helpful.

The data in Fig 3a are inevitably noisy - individual points are often consistent with no variation of polarization with wavelength. The error distributions are indeed complicated. What we want to draw attention to is the coherence of the the negative patches in the South of the source. This is a visual demonstration that depolarization is still significant there even at the higher resolution - a fact that is confirmed quantitatively by averaging.

We have added a note to clarify this point.

The choice and order of Figures 4,5,6 is curious. Fig 6 might go better next to Fig 4,

since they refer to the same resolution images in Fig 3. There are no plots of depolarization
at 5.5 arcsec resolution to go with the present Fig5. They would show clearly the depolarization in the southern jet, and how it largely goes away at high resolution.

We think that the figures were in a sensible order. What was perhaps illogical was the absence of a plot of p against lambda^2 at low resolution. We have now added this. All four plots against lambda^2 are now on the same page, allowing comparison either between p and position angle at the same resolution or between the same quantities at different resolutions. We have reduced the number of panels in the low-resolution plots to save space, as six examples are enough to make our points.

In section 2 they correctly emphasize the linearity of the RM plots as strong evidence for an
external screen. Any internal rotation would cause deviations from linearity. Is it possible
to put limits on that which are small enough to be interesting?

The display of observational results continue with Figures 7 and 8, which are quite striking
(especially 8 e and f), but are buried in the analysis section, far from the images they
summarize. Can they be moved earlier? Perhaps they could be shrunk somewhat to facilitate that. (By the way, in the caption to Figure 8, is the â€œmainâ€ jet the north jet?)

This was a latex problem: the figure references were at the correct places in the text. We have moved figures around and changed shapes and sizes so that Figs 7 and 8 appear closer to where they are referenced. We now say "the North (main) jet" in the caption to Fig 8.

Section 3. The overview of the analysis is helpful.

In 3.1 (iii), they say the field is isotropic because there is no evidence for anisotropy in the
RM distribution. But Figure 8a seems to show that the large RMâ€™s are in the south. Are
they saying that any gradient is not significant compared to the fluctuations? In the middle
of page 10 they say the RM in the north tail is primarily Galactic, which would seem to

imply that there is indeed a local (to 3C31) gradient in RM. Or is a gradient different from
an anisotropy?

The use of the word "anisotropy" without qualification seems to have led to confusion. What we mean at this point is that the field has no preferred direction when averaged over a sufficiently large volume. We see no preferred direction in the RM variations on scales up to 100 arcsec or so. Given that we infer a power spectrum with significant amplitude on larger scales, we expect gradients: these will appear anisotropic due to imperfect sampling. In addition, of course, there are global variations in density.

We have rewritten the paragraph to clarify this point.

In 3.1 (iv) they assume that the amplitude of the magnetic field power spectrum varies with
thermal gas density, but its shape is everywhere the same. The shape is determined by the
slope and cutoffs in the spectrum. The characteristic length scales are interesting and there
is not much discussion of them. A high frequency cutoff is expected if there is dissipation
to dampen the fluctuations, and a low frequency cutoff or change of slope may be associated with the paddle that stirs the turbulence. The scales found here in section 4.5 are

4 kpc and 17 kpc. The closely related â€œmagnetic autocorrelation lengthâ€ depends on the
poorly determined outer scale and is listed in Table 2. How do these numbers compare to
what is expected or what is known about the cluster in which 3C31 lies? In other words,
what physical significance should the reader attach to them?

The caption to Figure 11 is very confusing. It says these are simulated images, but the rest
of the caption reads as if they are observed. If they are simulated, what should the reader
compare them to? In the text it seems to imply that 11b is a simulation of 11a. But this canâ€™t
be so because a stochastic process cannot reproduce any specific pattern in such detail.

On page 15, they point out that the Kolmogorov spectrum predicts too much depolarization
on small scales because it doesnâ€™t have a high frequency cutoff. They do not want to add a
cutoff because it would narrow the range of scales with a Kolmogorov slope, and dilute the
reason for choosing it in the first place. I am not sure this is quite fair. If the depolarization
on small scales requires a high frequency cut-off, even with a spectrum as steep as 11/3,
then this is simply a result demanded by the observations. It is not evidence against the
processes that might lead to a Kolmogorov slope at lower frequencies, and it is no more

inelegant than cutting off the power law spectrum at high frequencies. In either case the
cutoff tells us something about the physics of the Faraday screen (see above).

On page 14 they quote best fit values for fbreak and qlow for the Kolmogorov spectrum. I
would expect them to be quite strongly correlated. Is that the case? Or in other words, what
is the range of values that fit the data adequately? The same question applies to the values
of q and fmax for the broken power law.

Figure 14 shows the simulated and observed RM measurements in selected regions. It would help the reader to be guided as to what features indicate this is a successful model,
since there is obviously no one-to-one match. What would constitute an unacceptable simulation? Is there some sort of measure of â€œgoodness of fitâ€?

In section 5 they are looking for a 3-dimensional model for the Faraday screen. They sometimes use a single scale model for ease of computing and sometimes the power law spectrum of fluctuations. It is not easy for the reader to figure out which model they are
using in which simulations (Figures 17 and 18), and under what circumstances the single

scale model is inadequate.

For Figure 17, I make the same comments as for Figure 14 (two paragraphs back).

For the spherical model, is there not automatically a cavity where the source is (even though the x-ray data do not show one)? If there isnâ€™t, then you have Faraday rotating
material mixed in with the synchrotron emitting particles. Section 2.2 argues convincingly
for the absence of internal Faraday rotation. If they have filled source with thermal plasma,
how does that affect those arguments?

Figure 19 is confusing. If the tick mark at the bottom represents 100 arcsec, then the cone
extends for 250 arcsec or so. Not the 140 arcsec stated in the text. What are the three

concentric circles?

The cavity model seems to generate the required asymmetry by virtue of the cone angle matching the angle to the line of sight. This seems contrived.

The disk distribution seems equally ad hoc. A diagram similar to Figure 19 would be very
helpful here. Is there any precedence for a Faraday screen in the shape of an equatorial
disk? Does any galaxy have a large disk of material similar to what they visualize?

Section 5.7 contains reasonable suggestions for other explanations. At this point the 3dimensional
modeling, realistically, has been inconclusive, and perhaps the abstract should reflect that. Given that, comparison to Hydra A seems premature and makes the paper too

long.

3c31ls_2.pdf	**Content-Type:**	application/pdf
	Content-Encoding:	base64

---

response.txt	**Content-Type:**	text/plain
	Content-Encoding:	base64

Subject: Revised version of 3C31 paper
From: Robert Laing <rlaing@eso.org>
Date: Mon, 4 Feb 2008 00:17:25 +0100 (CET)
To: Alan Bridle <abridle@nrao.edu>, Paola Parma <parma@ira.inaf.it>, Luigina Feretti <lferetti@ira.inaf.it>, Gabriele Giovannini <ggiovann@ira.inaf.it>, Matteo Murgia <murgia@ira.inaf.it>, Rick Perley <rperley@aoc.nrao.edu>

Dear All

Here is a revised version. I need to give it a final read-through, but I have tried to address all of John's points.

Main issues:
- Better version of Fig 11. Looks good on my laptop, but I don't currently have access to a printer, so cannot tell what it might look like on paper.
- Bandwidth smearing comments. A slightly subtle point.
- Revised arc descriptions.
- HST overlay. I have chosen to make this a little more elaborate, but actually would not mind if it went.

Comments please. It would be nice to resubmit this tomorrow from Santiago before I get on the next plane.

Regards

Robert

We thank the referee for his kind words and helpful comments (see below for detailed replies).

Reviewer's Comments:

Reviewer: John Wardle

This paper presents extraordinarily high quality images of the total intensity and polarization of 3C31. The careful work of Laing and Bridle and their collaborators are turning this source into a Rosetta Stone for extragalactic radio sources.

I have only minor requests for clarification and suggestions that might improve the presentation.

I have one overall request. Most of the images are labeled in RA and Dec. But they cover a wide range of resolutions and fields of view, and the reader has to pay attention to whether the numbers on the declination axis are arcminutes or arcseconds (and similarly in right ascension). Since almost all the discussion is in terms of distances from the core, it would be a service to the reader if the images were labeled as they are figure 12. There the core is at (0. 0) and the axes are labeled in arcseconds, whether tens or hundreds. This makes it much easier for the reader to follow the changing scales and to know which part of the source is being shown. If this would be unreasonably burdensome on the authors, I will relent, because all the information is of course present in the labels and captions.

This suggestion would require remaking almost all of the diagrams, and would involve a significant amount of work. It would also make comparison with other published images (e.g. optical, X-ray) a little more difficult. We prefer to use labelling in RA, Dec for unrotated images. When we rotate to have the jets along an an axis (as we do in Fig. 12 and other papers on this object) then labelling in arcsec from the core is much more appropriate. We would prefer, therefore, to leave the diagrams as they are.

Yes. We haved checked carefully and fixed all occurrences.

Page 15, 12 lines down the second column. Don't start a sentence with a lower case letter. It looks like a typo. Try "The value of s .. " or something similar.

Agreed. We now say "The asymptotic value of s is only approached ..."

Two lines lower down, Gamma > 10. The discussion needs a little clarifying as to whether these are shock velocities or flow velocities.

Actually both, as pointed out by Kirk (2005). We have clarified this point and added the new reference.

Page 16: table 3 refers to the spectra shown on the previous page. Should it not be on the same page as Figure 13?

Yes, it should. With some dificulty, we have persuaded latex to allow this, and will make sure that the two remain on the same page in the published version.

The third decimal place is surely inconsistent with their claimed errors on the spectral index (~ 0.01).

We have reduced the number of significant figures.

Same page, final paragraph: they refer to Monte Carlo calculations, but I don't seem to be able to find any mention of them in the main part of the paper.

This was an overly cryptic reference to the work by Lemoine & Pelletier (2003; quoted earlier, in 5.2.3). The detail is not relevant in a summary, and we have deleted it.

But these are all small points. This is a very fine paper.

Thank you

We have also corrected one small error ("3C 66B" was typeset incorrectly).

3c31ls_1.pdf	**Content-Description:**	Revised version
	Content-Type:	APPLICATION/pdf
	Content-Encoding:	BASE64

MNRAS__MN_07_2001_MJ	**Content-Description:**	Response
	Content-Type:	TEXT/PLAIN
	Content-Encoding:	BASE64

On page 4, they discuss the effect of limited short spacings in the u-v coverage. For the largest scale maps which show the diffuse tails, do delay and bandwidth smearing affect the images or are they still close enough to the pointing center?

This is an important (and subtle) point, which we now discuss. The effects of delay smearing are completely negligible. Bandwidth smearing is potentially important, however, and indeed affects images of background sources in Fig 1b by a large amount. Fortunately, the structure of 3C31 is such that the effects of bandwidth smearing on our images are very small. The reason is that flux density is conserved and 3C31 has almost uniform surface brightness on the scales where bandwidth smearing is at all significant. For the spectral fits, the same argument applies (although the worst reduction in peak flux flux density in the region we use for quantitative work is in any case only ~5%.

We now discuss this point in Section 2.2 and reiterate it briefly in Section 5.1, to emphasise that it has no significant effect on our spectral indices. We note the effect on background sources in the caption for Fig 1(b), which is badly affected in its outer regions.

On page 8, it would be helpful to define again what is meant by type (i) and type (ii) arcs. It is unclear to me what the difference is between a type (ii) arc and just a sharp edge to the jet. I was unsure what they meant by "lozenge" and went to Laing et al 2006b to find out, but the word is not used there.

We used "lozenge" in the dictionary and heraldic sense (essentially a synonym for "rhombus"). As this seems to be confusing even to some of the co-authors, we now instead restate the definitions from our earlier paper. The sharpest brightness gradients of type (ii) arcs occur well within the outer envelope of the jet emission. We now say this explicitly, and point to Fig 5(b), where this is most obvious.

I was unable to glean very much from figure 7. Perhaps the complex dust lane and the molecular disk could be indicated. But since they seem to have little to do with the radio source, I wonder if figure 7 adds anything to the paper.

We do want to make the point that the inner X-ray gas distribution and the cold (dust/molecular) disk have very similar scales. We think that the former is directly associated with the flaring and recollimation of the jets, but the latter merely appears projected on the jets and has no direct effect on them. It has been suggested that the flaring of the jet of 3C 264 results from interaction with the dust disk (Baum et al. 1997), so the point may not be universally accepted.

We have improved the figure by labelling the dust disk and showing the core radius of the X-ray component and have also clarified this point in the text.

In figures 9, 10 and 11 the underlying grayscale images are very hard to see, and that makes it very hard to see the connections between morphological features in total intensity and the apparent magnetic field structure. I am not sure what to suggest. Perhaps the contrast on the Sobel filtered images can be cranked up, or perhaps there are too many B vectors.

We have attempted to improve these figures. In Fig. 9, we found nothing better than the existing grey-scale representation of the Sobel-filtered image, but we have increased the line thickness of the vectors slightly. Altering the transfer function in Fig. 10 has increased the visibility of the total intensity image. We have also used fewer vectors. Finally, we could only make a better version of Fig. 11 by using colour for the Sobel-filtered I image. The areas overed in the two panels are smaller than in the previous version and the vector spacing is larger. Ww think that the result is a significant improvement.

I hesitate to include this, but in my dictionary "further" indicates se paration in time and "farther" indicates separation in distance. Twice on page 14 further was used when farther was meant.

This paper is a tour de force. The rotation measure and depolarization images in figures 2 and 3 are an impressive achievement. They require matched resolutions over a wide range of wavelengths and excellent image fidelity of a complex, mostly low surface brightness, source (to say nothing about accurate polarization calibration). And there are enough resolution elements across the source to make a concerted investigation of the statistical properties of the Faraday screen feasible and very worthwhile.

Because both the observational material and the analysis and simulations are new and extensive, the paper ends up being very long and perhaps tries to do too much for a single paper. The second half of the paper has a more didactic flavour and is, in part, a primer on the statistical analysis of random screens and how to simulate them. This is necessary because it is new, and such excellent observational material has not previously been available. The paper ends by applying these ideas to Hydra A. By now it has wandered somewhat from the title and original focus of the paper and the reader is exhausted. This suggests a natural division into two papers: the first on the magneto-ionic medium around 3C31, and the second on the statistical analysis, simulations and application to other sources, with room to expand as much as the authors wish (the writing is becomes quite compacted by the end of the present paper). This is just a suggestion. The authors may well not agree, and I leave it between them and the editor. My interest is simply in getting all parts of this work the attention they deserve.

Detailed questions and comments, starting from the beginning:

p2 paragraph 2, item (ii) last sentence: it would be helpful to name the sources.

Some questions regarding Figures 2 and 3:

Although they don't say so, the agreement between 2b and 2c shows that the RM determined from the 4 Lband wavelengths, with their rather short lever arm, agrees very well with the 5 wavelength measurements. This is a testimony to the internal consistency and quality of the images. It also says that (with care) one can make RM maps "in a single band" if properly scaled array observations in a different band are not available.
Not necessarily worth commenting on, but gratifying to note, nonetheless.

In Figure 3, the RM plots use different algorithms. PACERMAN clearly works better at lower SNR. A couple of sentences saying what it does differently, and why it does better at low SNR would be helpful. Should one always use it? Was it used in Figure 2?

Figure 4 shows selected depolarization plots at 1.5 arcsec resolution. As they state, only 5e shows significant depolarization. Three other plots have good SNR but 5 are consistent with a range of slopes including zero. Does this mean that in Figure 3a we are mainly looking at the errors in the slope? It is clear from Figure 8f that there is still significant depolarization in the southern jet at 1.5 arcsecond resolution, but that is after some averaging. The difficulty with plots like Figures 2 and 3 is that every point has a different error (unlike an I map), and I don't know if there is any way to show that. It is compounded in 3a by dividing by p(0). One might infer from Figure 4 that perhaps half the pixels in 3a are essentially noise (depending strongly on location). I assume the authors have made sensible choices about clipping levels etc, but some discussion of the problem would be helpful.

4e

excelnt in green & blue regions fig 3 statistical differences NOT noise.

The choice and order of Figures 4,5,6 is curious. Fig 6 might go better next to Fig 4, since they refer to the same resolution images in Fig 3. There are no plots of depolarization at 5.5 arcsec resolution to go with the present Fig5. They would show clearly the depolarization in the southern jet, and how it largely goes away at high resolution.

— These are done in 2(c) or 3(a) logic is clear Fig. 4 shows depol. derived in Fig. 3.

In section 2 they correctly emphasize the linearity of the RM plots as strong evidence for an external screen. Any internal rotation would cause deviations from linearity. Is it possible to put limits on that which are small enough to be interesting?

The display of observational results continue with Figures 7 and 8, which are quite striking (especially 8 e and f), but are buried in the analysis section, far from the images they summarize. Can they be moved earlier? Perhaps they could be shrunk somewhat to facilitate that. (By the way, in the caption to Figure 8, is the "main" jet the north jet?)

— yes, should be moved earlier

Section 3. The overview of the analysis is helpful.

In 3.1 (iii), they say the field is isotropic because there is no evidence for anisotropy in the RM distribution. But Figure 8a seems to show that the large RM's are in the south. Are they saying that any gradient is not significant compared to the fluctuations? In the middle of page 10 they say the RM in the north tail is primarily Galactic, which would seem to imply that there is indeed a local (to 3C31) gradient in RM. Or is a gradient different from an anisotropy?

In 3.1 (iv) they assume that the amplitude of the magnetic field power spectrum varies with thermal gas density, but its shape is everywhere the same. The shape is determined by the slope and cutoffs in the spectrum. The characteristic length scales are interesting and there is not much discussion of them. A high frequency cutoff is expected if there is dissipation to dampen the fluctuations, and a low frequency cutoff or change of slope may be associated with the paddle that stirs the turbulence. The scales found here in section 4.5 are 4 kpc and 17 kpc. The closely related "magnetic autocorrelation length" depends on the poorly determined outer scale and is listed in Table 2. How do these numbers compare to what is expected or what is known about the cluster in which 3C31 lies? In other words, what physical significance should the reader attach to them?

The caption to Figure 11 is very confusing. It says these are simulated images, but the rest of the caption reads as if they are observed. If they are simulated, what should the reader compare them to? In the text it seems to imply that 11b is a simulation of 11a. But this can't be so because a stochastic process cannot reproduce any specific pattern in such detail.

On page 15, they point out that the Kolmogorov spectrum predicts too much depolarization on small scales because it doesn't have a high frequency cutoff. They do not want to add a cutoff because it would narrow the range of scales with a Kolmogorov slope, and dilute the reason for choosing it in the first place. I am not sure this is quite fair. If the depolarization on small scales requires a high frequency cut-off, even with a spectrum as steep as 11/3, then this is simply a result demanded by the observations. It is not evidence against the processes that might lead to a Kolmogorov slope at lower frequencies, and it is no more inelegant than cutting off the power law spectrum at high frequencies. In either case the cutoff tells us something about the physics of the Faraday screen (see above).

On page 14 they quote best fit values for f_{break} and q_{low} for the Kolmogorov spectrum. I would expect them to be quite strongly correlated. Is that the case? Or in other words, what is the range of values that fit the data adequately? The same question applies to the values of q and f_{max} for the broken power law.

Figure 14 shows the simulated and observed RM measurements in selected regions. It would help the reader to be guided as to what features indicate this is a successful model, since there is obviously no one-to-one match. What would constitute an unacceptable simulation? Is there some sort of measure of "goodness of fit"?

yes, good point

In section 5 they are looking for a 3-dimensional model for the Faraday screen. They sometimes use a single scale model for ease of computing and sometimes the power law spectrum of fluctuations. It is not easy for the reader to figure out which model they are using in which simulations (Figures 17 and 18), and under what circumstances the single scale model is inadequate.

For Figure 17, I make the same comments as for Figure 14 (two paragraphs back).

For the spherical model, is there not automatically a cavity where the source is (even though the x-ray data do not show one)? If there isn't, then you have Faraday rotating material mixed in with the synchrotron emitting particles. Section 2.2 argues convincingly for the absence of internal Faraday rotation. If they have filled source with thermal plasma, how does that affect those arguments?

if we believe the linearity, yes. The sphere is illustrative

Figure 19 is confusing. If the tick mark at the bottom represents 100 arcsec, then the cone extends for 250 arcsec or so. Not the 140 arcsec stated in the text. What are the three concentric circles?

The cavity model seems to generate the required asymmetry by virtue of the cone angle matching the angle to the line of sight. This seems contrived.

Yes, brings reasons for it to work! That's the dilemma

The disk distribution seems equally *ad hoc*. A diagram similar to Figure 19 would be very helpful here. Is there any precedence for a Faraday screen in the shape of an equatorial disk? Does any galaxy have a large disk of material similar to what they visualize?

Section 5.7 contains reasonable suggestions for other explanations. At this point the 3-dimensional modeling, realistically, has been inconclusive, and perhaps the abstract should reflect that. Given that, comparison to Hydra A seems premature and makes the paper too long.

but also contrived! So depends on what one preconceives as reasonable

don't see HyA as a stand-alone though.

Subject: Response
From: Robert Laing <rlaing@eso.org>
Date: Mon, 7 Apr 2008 23:41:39 +0200
To: abridle@nrao.edu

R

This paper is a tour de force. The rotation measure and depolarization images in figures 2
and 3 are an impressive achievement. They require matched resolutions over a wide range

of wavelengths and excellent image fidelity of a complex, mostly low surface brightness,
source (to say nothing about accurate polarization calibration). And there are enough
resolution elements across the source to make a concerted investigation of the statistical
properties of the Faraday screen feasible and very worthwhile.

Because both the observational material and the analysis and simulations are new and
extensive, the paper ends up being very long and perhaps tries to do too much for a single
paper. The second half of the paper has a more didactic flavour and is, in part, a primer on
the statistical analysis of random screens and how to simulate them. This is necessary
because it is new, and such excellent observational material has not previously been
available. The paper ends by applying these ideas to Hydra A. By now it has wandered
somewhat from the title and original focus of the paper and the reader is exhausted. This
suggests a natural division into two papers: the first on the magneto-ionic medium around
3C31, and the second on the statistical analysis, simulations and application to other
sources, with room to expand as much as the authors wish (the writing is becomes quite
compacted by the end of the present paper). This is just a suggestion. The authors may well
not agree, and I leave it between them and the editor. My interest is simply in getting all
parts of this work the attention they deserve.

We have considered this suggestion carefully. It has a number of attractions for us (even excluding the base motivation of another publication): we agree that there might be a broader audience for the basic observational results than for full details of the modelling; each paper would be shorter and more likely to be read by its target audience. However, we concluded that split papers would require a lot of repetition and cross-referencing, to the extent that the total would be much longer. We also felt that it would be very hard to understand some of the conclusions without at least speed-reading about the methods. It is instructive to compare with the two papers by Ensslin & Vogt (2003; A&A 412, 373 and 401, 835), which adopted precisely this split approach: there is indeed a lot of repetition. We therefore feel that the paper will be shorter in total and easier to read if left in its current form.

On reflection, we feel that we should offer a better route map in the abstract, which currently focuses on the results, downplaying the methods.
....

We have also changed the title in the hope that this will give a more accurate guide to the content of the paper.

Detailed questions and comments, starting from the beginning:

p2 paragraph 2, item (ii) last sentence: it would be helpful to name the sources.

We have now done this.

Some questions regarding Figures 2 and 3:

Although they donâ€™t say so, the agreement between 2b and 2c shows that the RM determined from the 4 Lband wavelengths, with their rather short lever arm, agrees very

well with the 5 wavelength measurements. This is a testimony to the internal consistency and quality of the images. It also says that (with care) one can make RM maps â€œin a single bandâ€ if properly scaled array observations in a different band are not available. Not necessarily worth commenting on, but gratifying to note, nonetheless.

Indeed. We now refer to this point briefly in the section on further work, where we point out that the next generation of correlators will make this much easier.

In Figure 3, the RM plots use different algorithms. PACERMAN clearly works better at lower SNR. A couple of sentences saying what it does differently, and why it does better at low SNR would be helpful. Should one always use it? Was it used in Figure 2?

We have added some text to give a short description of the Pacerman algorithm. We are not prepared to be dogmatic about the circumstances in which it should be used, as we have only really tested it for a limited range of problems. It was not used in Fig 2., where we have higher s/n and very few problems with npi ambiguities (we say this specifically). It did not work very well for the S lobe of Hydra A, as Vogt et al. (2005) point out.

Figure 4 shows selected depolarization plots at 1.5 arcsec resolution. As they state, only 5e shows significant depolarization. Three other plots have good SNR but 5 are consistent with a range of slopes including zero. Does this mean that in Figure 3a we are mainly looking at the errors in the slope? It is clear from Figure 8f that there is still significant depolarization in the southern jet at 1.5 arcsecond resolution, but that is after some averaging. The difficulty with plots like Figures 2 and 3 is that every point has a different error (unlike an I map), and I donâ€™t know if there is any way to show that. It is compounded in 3a by dividing by p(0). One might infer from Figure 4 that perhaps half the pixels in 3a are essentially noise (depending strongly on location). I assume the authors have made sensible choices about clipping levels etc, but some discussion of the problem would be helpful.

The choice and order of Figures 4,5,6 is curious. Fig 6 might go better next to Fig 4, since they refer to the same resolution images in Fig 3. There are no plots of depolarization at 5.5 arcsec resolution to go with the present Fig5. They would show clearly the depolarization in the southern jet, and how it largely goes away at high resolution.

We think that the figures were in a sensible order. What was perhaps illogical was the absence of a plot of p against lambda^2 at low resolution. We have now added this. All four plots against lambda^2 are now on the same page, allowing comparison either between p and position angle at the same resolution or between the same quantities at different resolutions. We have reduced the number of panels in the low-resolution plots to save space, as six examples are enough to make our points.

In section 2 they correctly emphasize the linearity of the RM plots as strong evidence

for an
external screen. Any internal rotation would cause deviations from linearity. Is it possible
to put limits on that which are small enough to be interesting?

The limits are still a factor of 200 or so higher than the internal densities we infer from a conservation-law analysis (Laing & Bridle 2002 MNRAS 336, 1161), even in the optimistic case of a completely ordered magnetic field. The limits depend on the number of reversals in the field and the details of the geometry, of course. We now give a rough estimate for a distance of about 30 arcsec from the nucleus, where we can compare directly with our earlier number. This is probably the most interesting location: further out, the increase in path length through the jet is essentially counterbalanced by the decrease in magnetic field (assumed to be close to the equipartition value) and the s/n on the position-angle measurements is worse.

The display of observational results continue with Figures 7 and 8, which are quite striking
(especially 8 e and f), but are buried in the analysis section, far from the images they
summarize. Can they be moved earlier? Perhaps they could be shrunk somewhat to facilitate that. (By the way, in the caption to Figure 8, is the â€œmainâ€ jet the north jet?)

This was a latex problem: the figure references were at the correct places in the text. We have moved figures around and changed shapes and sizes so that Figs 7 and 8 appear closer to where they are referenced. We now say "the North (main) jet" in the caption to Fig 8.

Section 3. The overview of the analysis is helpful.

In 3.1 (iii), they say the field is isotropic because there is no evidence for anisotropy in the
RM distribution. But Figure 8a seems to show that the large RMâ€™s are in the south. Are
they saying that any gradient is not significant compared to the fluctuations? In the middle
of page 10 they say the RM in the north tail is primarily Galactic, which would seem to

imply that there is indeed a local (to 3C31) gradient in RM. Or is a gradient different from
an anisotropy?

The use of the word "anisotropy" without qualification seems to have led to confusion. What we mean at this point is that the field has no preferred direction when averaged over a sufficiently large volume. We see no preferred direction in the RM variations on scales up to 100 arcsec or so. Given that we infer a power spectrum with significant amplitude on larger scales, we expect gradients: these will appear anisotropic due to imperfect sampling. In addition, of course, there are global variations in density.

We have rewritten the paragraph to clarify this point.

In 3.1 (iv) they assume that the amplitude of the magnetic field power spectrum varies with
thermal gas density, but its shape is everywhere the same. The shape is determined by the
slope and cutoffs in the spectrum. The characteristic length scales are interesting and there
is not much discussion of them. A high frequency cutoff is expected if there is dissipation
to dampen the fluctuations, and a low frequency cutoff or change of slope may be associated with the paddle that stirs the turbulence. The scales found here in section 4.5 are
4 kpc and 17 kpc. The closely related â€œmagnetic autocorrelation lengthâ€ depends on the
poorly determined outer scale and is listed in Table 2. How do these numbers compare to

what is expected or what is known about the cluster in which 3C31 lies? In other words, what physical significance should the reader attach to them?

The caption to Figure 11 is very confusing. It says these are simulated images, but the rest
of the caption reads as if they are observed. If they are simulated, what should the reader
compare them to? In the text it seems to imply that 11b is a simulation of 11a. But this canâ€™t
be so because a stochastic process cannot reproduce any specific pattern in such detail.

We agree that this was confusing and have rewritten the caption. Panels (a) - (c) all result from the same simulated dataset. (a) and (b) are a test of the short-wavelength approximation as described in the text and (c) is the associated depolarization. Panel (d) is a portion of the observed RM distribution, included purely to show that it displays the same sort of artefacts as are seen in the simulation.

On page 15, they point out that the Kolmogorov spectrum predicts too much depolarization
on small scales because it doesnâ€™t have a high frequency cutoff. They do not want to add a
cutoff because it would narrow the range of scales with a Kolmogorov slope, and dilute the
reason for choosing it in the first place. I am not sure this is quite fair. If the depolarization
on small scales requires a high frequency cut-off, even with a spectrum as steep as 11/3,
then this is simply a result demanded by the observations. It is not evidence against the
processes that might lead to a Kolmogorov slope at lower frequencies, and it is no more

inelegant than cutting off the power law spectrum at high frequencies. In either case the
cutoff tells us something about the physics of the Faraday screen (see above).

On page 14 they quote best fit values for fbreak and qlow for the Kolmogorov spectrum. I
would expect them to be quite strongly correlated. Is that the case? Or in other words, what
is the range of values that fit the data adequately? The same question applies to the values
of q and fmax for the broken power law.

Figure 14 shows the simulated and observed RM measurements in selected regions. It would help the reader to be guided as to what features indicate this is a successful model,
since there is obviously no one-to-one match. What would constitute an unacceptable simulation? Is there some sort of measure of â€œgoodness of fitâ€?

In section 5 they are looking for a 3-dimensional model for the Faraday screen. They sometimes use a single scale model for ease of computing and sometimes the power law spectrum of fluctuations. It is not easy for the reader to figure out which model they are
using in which simulations (Figures 17 and 18), and under what circumstances the single

scale model is inadequate.

For Figure 17, I make the same comments as for Figure 14 (two paragraphs back).

For the spherical model, is there not automatically a cavity where the source is (even though the x-ray data do not show one)? If there isnâ€™t, then you have Faraday rotating
material mixed in with the synchrotron emitting particles. Section 2.2 argues convincingly
for the absence of internal Faraday rotation. If they have filled source with thermal

plasma,
how does that affect those arguments?

Figure 19 is confusing. If the tick mark at the bottom represents 100 arcsec, then the cone
extends for 250 arcsec or so. Not the 140 arcsec stated in the text. What are the three

concentric circles?

The cavity model seems to generate the required asymmetry by virtue of the cone angle matching the angle to the line of sight. This seems contrived.

The disk distribution seems equally ad hoc. A diagram similar to Figure 19 would be very
helpful here. Is there any precedence for a Faraday screen in the shape of an equatorial
disk? Does any galaxy have a large disk of material similar to what they visualize?

Section 5.7 contains reasonable suggestions for other explanations. At this point the 3dimensional
modeling, realistically, has been inconclusive, and perhaps the abstract should reflect that. Given that, comparison to Hydra A seems premature and makes the paper too

long.

response.txt	**Content-Type:** text/plain **Content-Encoding:** base64

Subject: 3C31 paper 2
From: Robert Laing <rlaing@eso.org>
Date: Tue, 24 Jun 2008 17:28:33 +0200
To: Alan Bridle <abridle@nrao.edu>

Dear Alan

Not quite finished yet, but close enough that I'd like to get your comments on some of the major changes. Please ignore the abstract, Sections 6 and 7 and Appendix C for now. Otherwise, do your worst - particularly if you can find ways of tightening up arguments, removing dangling sub-clauses and making the rough places plain. I don't anticipate many changes in the abstract or Section 7, but the analysis for Hydra needs to match what is done for 3C31 and App C is a dumping ground for various bits of depolarization analysis, not yet in any logical order.

Various things for you to consider:

1. I found that moving the the subsection on the effects of partial resolution to section 2 actually created some problems that we failed to notice at the time. The trouble is that it justifies the convolution relation which is fundamental to our 2D analysis as well. I've restored it to Section 3 with an additional motivating sentence and put a forward reference in Sec 2.

2. There's a new 2.4 on limits to internal rotation. Maybe too long and an encouragement for others to revert to old, bad arguments?

3. 2D section is now more rigorous, I think. Unfortunately, the depolarization argument is inconclusive. Fig 12 is pretty convincing though. I have deliberately swept under the carpet the problem of the number of degrees of freedom for chi-squared in these fits. I think the quoted error ranges are reasonable, but would hate to have to defend them in detail. I think this is acceptable - OK?

4. 3D section: extensively rewritten. This is what I have just finished, so expect most rough edges here. Main changes:
   - ellipsoidal cavities - see new sketch, which also uses deep I image to emphasise emission from N spur.
   - Fit of rms RM profiles. Chi-squareds are more rigorous here.
   - Disks relegated to "other explanations"; superdisks mentioned and sneered at.
   - Stuff on deviations from axisymmetry, intrinsic asymmetry etc. moved to end of cavities section.
   - I think the cavity fit is actually pretty reasonable except for the -300 arcsec bin when you think about the sampling errors, so have talked this up a bit.
   - Single-scale model stuff moved later and problems with it emphasised.

In general, there is much scope for incorrect cross-referencing, especially with panel labels in the more complicated figures, for confusion between broken and cut-off power laws and for leftover sentences referring to out-of-date analyses or opinions. Please watch out for this - I think my brain no longer notices it.

I'll deal with Hydra next. Shouldn't take too long, as the machinery is now exactly the same as for 3C31.

I've attached pdf and what should be a complete source tarball.

Have fun.

Regards

Robert

```
The original MIME headers for this attachment are:
Content-Type: text/ai; name="3c31ls_2.pdf"
Content-Transfer-Encoding: base64
Content-Disposition: attachment; filename="3c31ls_2.pdf"
```

3c31ls_2.pdf	**Content-Type:** text/ai

source_20080624.tgz	**Content-Type:** application/x-gzip **Content-Encoding:** base64

We thank the referee for a very thoughtful reading. As will become apparent, this has caused us to think carefully about our error analysis and to make significant changes in some areas. We think that the results are significantly more robust as a consequence.

This paper is a tour de force. The rotation measure and depolarization images in figures 2 and 3 are an impressive achievement. They require matched resolutions over a wide range of wavelengths and excellent image fidelity of a complex, mostly low surface brightness, source (to say nothing about accurate polarization calibration). And there are enough resolution elements across the source to make a concerted investigation of the statistical properties of the Faraday screen feasible and very worthwhile.

Because both the observational material and the analysis and simulations are new and extensive, the paper ends up being very long and perhaps tries to do too much for a single paper. The second half of the paper has a more didactic flavour and is, in part, a primer on the statistical analysis of random screens and how to simulate them. This is necessary because it is new, and such excellent observational material has not previously been available. The paper ends by applying these ideas to Hydra A. By now it has wandered somewhat from the title and original focus of the paper and the reader is exhausted. This suggests a natural division into two papers: the first on the magneto-ionic medium around 3C31, and the second on the statistical analysis, simulations and application to other sources, with room to expand as much as the authors wish (the writing is becomes quite compacted by the end of the present paper). This is just a suggestion. The authors may well not agree, and I leave it between them and the editor. My interest is simply in getting all parts of this work the attention they deserve.

We have considered this suggestion carefully. It has a number of attractions for us (even excluding the base motivation of another publication): we agree that there might be a broader audience for the basic observational results than for full details of the modelling; each paper would be shorter and more likely to be read by its target audience. However, we concluded that split papers would require a lot of repetition and cross-referencing, to the extent that the total would be much longer. We also felt that it would be very hard to understand some of the conclusions without at least speed-reading about the methods. It is instructive to compare with the two papers by Ensslin & Vogt (2003; A&A 412, 373 and 401, 835), which adopted precisely this split approach: there is indeed a lot of repetition. We therefore feel that the paper will be shorter in total and easier to read if left in its current form.

On reflection, we feel that we should offer a better route map in the abstract, which originally focused on the results, downplaying the methods. We have therefore rewritten it. We have also changed the title in the hope that this will also give a more accurate guide to the content of the paper.

Detailed questions and comments, starting from the beginning:

Note that there is a new figure (4) and that the order of other figures has changed. We refer to the old numbering in what follows.

p2 paragraph 2, item (ii) last sentence: it would be helpful to name the sources.

We have now done this.

Some questions regarding Figures 2 and 3:

Although they don't say so, the agreement between 2b and 2c shows that the RM determined from the 4 Lband wavelengths, with their rather short lever arm, agrees very well with the 5 wavelength measurements. This is a testimony to the internal consistency and quality of the images. It also says that (with care) one can make RM maps "in a single band" if properly scaled array observations in a different band are not available. Not necessarily worth commenting on, but gratifying to note, nonetheless.

Indeed. We already said that "the two images are consistent with each other where they overlap" and we now also refer to this issue briefly in the section on further work, where we point out that the next generation of correlators will make in-band RM determinations much easier.

In Figure 3, the RM plots use different algorithms. PACERMAN clearly works better at lower SNR. A couple of sentences saying what it does differently, and why it does better at low SNR would be helpful. Should one always use it? Was it used in Figure 2?

We have added some text to give a short description of the Pacerman algorithm. We are not prepared to be dogmatic about the circumstances in which it should be used, as we have only tested it for a limited range of problems. It was not used in Fig 2., where we have higher s/n and very few problems with npi ambiguities (we say this specifically). It did not work very well for the S lobe of Hydra A, as Vogt et al. (2005) point out.

Figure 4 shows selected depolarization plots at 1.5 arcsec resolution. As they state, only 5e shows significant depolarization. Three other plots have good SNR but 5 are consistent with a range of slopes including zero. Does this mean that in Figure 3a we are mainly looking at the errors in the slope? It is clear from Figure 8f that there is still significant depolarization in the southern jet at 1.5 arcsecond resolution, but that is after some averaging. The difficulty with plots like Figures 2 and 3 is that every point has a different error (unlike an I map), and I don't know if there is any way to show that. It is compounded in 3a by dividing by p(0). One might infer from Figure 4 that perhaps half the pixels in 3a are essentially noise (depending strongly on location). I assume the authors have made sensible choices about clipping levels etc, but some discussion of the problem would be helpful.

are

The s/n is indeed low for any depolarization estimator at 1.5 arcsec resolution: we ae trying to measure a rather subtle effect. We have emphasised the spatial coherence evident in the south of the source, which indicates that the variations of depolarization are real. The referee's comments caused us to re-evaluate the use of p'(0)/p(0) as an estimator of depolarization. Modelling of the error distribution showed that it was rather asymmetric except in the regions of very high s/n, potentially biasing our averages. We therefore decided to fit a Burn law instead. This has two advantages: it avoids the need to divide the gradient by the zero-wavelength polarization to get a quantity physically related to depolarization and we expect it to be a better description of (most of) the data, as we already explained later in the paper.

Although this approach reduces the bias, modelling of the fitting process shows that there are still some residual problems. We have, conservatively, used only data with s/n > 4 in the degree of polarization for quantitative comparison with models and averaged profiles at 1.5 arcsec resolution, although we still show wider areas in Figs 2(a) and 3(a).

We now discuss these issues in the text.

As a consequence, we have:
- Replaced Figs 2(a) and 3(a) with the equivalent images of Burn k.
- Replaced the corresponding profiles in Fig 8 (now 9) e and f,
- Revised all of the relevant text
- Referred to the discussion of rotation and depolarization by an almost resolved foreground screen in Section 2.1, to motivate use of the Burn law.

The choice and order of Figures 4,5,6 is curious. Fig 6 might go better next to Fig 4, since they refer to the same resolution images in Fig 3. There are no plots of depolarization at 5.5 arcsec resolution to go with the present Fig5. They would show clearly the depolarization in the southern jet, and how it largely goes away at high resolution.

We think that the figures were in a sensible order, but we have responded to the referee's comment by adding a plot of p against \lambda^4 at low resolution. All four plots (as functions of \lambda^2 or \lambda^4 as appropriate) are now on the same page, allowing comparison either between p and position angle at the same resolution or between the same quantities at different resolutions. We have reduced the number of panels in the low-resolution plots to save space, as six examples are enough to make our points.

In section 2 they correctly emphasize the linearity of the RM plots as strong evidence for an external screen. Any internal rotation would cause deviations from linearity. Is it possible to put limits on that which are small enough to be interesting?

The limits are still a factor of >100 or so higher than the internal densities we infer from a conservation-law analysis (Laing & Bridle 2002 MNRAS 336, 1161), even in the optimistic case of a completely ordered magnetic field. The limits depend on the number of reversals in the field and the details of the geometry, of course. The limits from the lack of depolarization in the North jet are actually a bit more stringent, so we have given (approximate) results for both methods in a new subsection at the end of Section 2. We now give a rough estimate for a distance of about 30 arcsec from the nucleus, where we can compare directly with our earlier number. This is probably the most interesting location: further out, the increase in path length through the jet is essentially counterbalanced by the decrease in magnetic field (assumed to be close to the equipartition value) and the s/n on the position-angle measurements is worse.

The display of observational results continue with Figures 7 and 8, which are quite striking (especially 8 e and f), but are buried in the analysis section, far from the images they summarize. Can they be moved earlier? Perhaps they could be shrunk somewhat to facilitate that. (By the way, in the caption to Figure 8, is the "main" jet the north jet?)

This was a latex problem: the figure references were at the correct places in the text. We have moved figures around and changed shapes and sizes so that Figs 7 and 8 (now 8 and 9) appear closer to where they are referenced. We now say "the North jet" rather than "the main jet" in the caption to Fig 8 (now 9).

Section 3. The overview of the analysis is helpful.

In 3.1 (iii), they say the field is isotropic because there is no evidence for anisotropy in the RM distribution. But Figure 8a seems to show that the large RM's are in the south. Are they saying that any gradient is not significant compared to the fluctuations? In the middle of page 10 they say the RM in the north tail is primarily Galactic, which would seem to imply that there is indeed a local (to 3C31) gradient in RM. Or is a gradient different from an anisotropy?

We have clarified what we mean by "anisotropy" in this context, i.e. that the field has no preferred direction when averaged over a sufficiently large volume. We see no preferred direction in the RM variations on scales up to 100 arcsec or so. Given that we infer a power spectrum with significant amplitude on larger scales, we expect gradients: these will appear anisotropic due to imperfect sampling. In addition, of course, there are global variations in density.

We have rewritten the paragraph to clarify this point.

In 3.1 (iv) they assume that the amplitude of the magnetic field power spectrum varies with thermal gas density, but its shape is everywhere the same. The shape is determined by the slope and cutoffs in the spectrum. The characteristic length scales are interesting and there is not much discussion of them. A high frequency cutoff is expected if there is dissipation to dampen the fluctuations, and a low frequency cutoff or change of slope may be associated with the paddle that stirs the turbulence. The scales found here in section 4.5 are 4 kpc and 17 kpc. The closely related "magnetic autocorrelation length" depends on the poorly determined outer scale and is listed in Table 2. How do these numbers compare to what is expected or what is known about the cluster in which 3C31 lies? In other words, what physical significance should the reader attach to them?

We have added a short subsection at the end of Section 3 describing the various scales in the problem and their possible physical meaning, cautioning that there is no generally accepted theory. Energy input could occur over a range of scales from the size of the radio source down to the jet bending radius. Dissipation is expected to occur on the resistive scale, which is tiny compared with our beam. There is a possibility that a change of slope in the power spectrum might occur on the folding scale in some fluctuation dynamo models, and we give references for this idea.

The caption to Figure 11 is very confusing. It says these are simulated images, but the rest of the caption reads as if they are observed. If they are simulated, what should the reader compare them to? In the text it seems to imply that 11b is a simulation of 11a. But this can't be so because a stochastic process cannot reproduce any specific pattern in such detail.

We agree that this was confusing and have rewritten the caption. Panels (a) - (c) all result from the same simulated dataset. (a) and (b) are a test of the short-wavelength approximation as described in the text and (c) is the associated depolarization. Panel (d) is a portion of the observed RM distribution, included purely to show that it displays the same sort of artefacts as are seen in the simulation.

On reflection, we decided that the material on deviations from lambda^2 rotation and the wavelength dependence of polarization in the South of the source at 5.5 arcsec, while useful as a consistency check and not documented in the literature, was getting in the way of our main argument. We have therefore moved it to a new Appendix C.

On page 15, they point out that the Kolmogorov spectrum predicts too much depolarization on small scales because it doesn't have a high frequency cutoff. They do not want to add a cutoff because it would narrow the range of scales with a Kolmogorov slope, and dilute the reason for choosing it in the first place. I am not sure this is quite fair. If the depolarization on small scales requires a high frequency cut-off, even with a spectrum as steep as 11/3, then this is simply a result demanded by the observations. It is not evidence against the processes that might lead to a Kolmogorov slope at lower frequencies, and it is no more inelegant than cutting off the power law spectrum at high frequencies. In either case the cutoff tells us something about the physics of the Faraday screen (see above).

We no longer try to make this point: a more careful error analysis shows that both models for the RM power spectrum predict depolarizations consistent with those observed (to within 1 sigma) and we make this point clear.

On page 14 they quote best fit values for fbreak and qlow for the Kolmogorov spectrum. I would expect them to be quite strongly correlated. Is that the case? Or in other words, what is the range of values that fit the data adequately? The same question applies to the values of q and fmax for the broken power law.

This comment prompted us to formalize our fitting procedure, which had previously been done "by eye". We now minimise chi-squared (summed over the four non-overlapping source regions with good data and excluding SP2). The parameters we derive are slightly different from those we quoted originally and the improvement in fit was sufficient to cause us to redo subsequent steps in the analysis with the new values.

The values of q and fmax for the cut-off power law (we think that the referee means this, rather than broken power law) and fbreak and qlow for the broken power law are correlated in the expected sense: a flatter power spectrum requires a cut-off at a higher frequency. We now point this out and tabulate appropriate limits.

Figure 14 shows the simulated and observed RM measurements in selected regions. It would help the reader to be guided as to what features indicate this is a successful model, since there is obviously no one-to-one match. What would constitute an unacceptable simulation? Is there some sort of measure of “goodness of fit”?

The structure function fit (using the error bars we derive from Monte Carlo simulations) gives a quantitative estimate of how well we describe the spatial statistics of a given region. However, this assumes that the RM is actually a Gaussian random variable. Our data are not extensive enough for us to look at higher-order correlations or other methods for the detection of non-Gaussianity, but a visual comparison is a good sanity check - the eye is very good at detecting correlations with structure and preferred directions in the data. This is why we show selected realisations.

We have made these points in the text.

In section 5 they are looking for a 3-dimensional model for the Faraday screen. They sometimes use a single scale model for ease of computing and sometimes the power law spectrum of fluctuations. It is not easy for the reader to figure out which model they are using in which simulations (Figures 17 and 18), and under what circumstances the single scale model is inadequate.

We have rewritten 5.2 to emphasise that a single-scale model does not always give us something which can be compared directly with observations. The reason is that it assumes averaging over many cells. With a realistic power spectrum, which has power on large scales, this assumption may be incorrect. For our data at 5.5 arcsec resolution, the assumption breaks down close to the nucleus (where the source is narrow) - this is why the simulations predict lower rms than the single-scale approximation close to the nucleus. If one averages over a large enough region, then the magnetic field strength can be derived from the single-scale approximation, setting the scale equal to the magnetic autocorrelation length (as we say). But we cannot average over large enough areas, because we are trying to determine changes in the normalization of the RM power spectrum across the source. A second, critical, problem is that we cannot derive realistic sampling errors, which are vital to our chi-squared analysis. Our main comparison is between simulations and observations, which automatically takes into account the irregular sampling and we use the single-scale model merely as a means of exploring parameter space simply and quickly.

We have moved the subsection after the description of the 3-d simulations so any residual confusion over fitting methods should be removed.

For Figure 17, I make the same comments as for Figure 14 (two paragraphs back).

Our criterion for an acceptable 3D model is that we fit both the shape and the normalization of the RM structure function to within errors set primariy by sampling variance. Given our 2D analysis, we know that we can fit the shape of the structure function with a power spectrum whose normalization varies across the source. To fit the normalization, we need to average over an area large enough to get a reliable number but small compared with the variations. Our best compromise is to fit the rms RM at 5.5 arcsec resolution. We now do this formally, again evaluating chi-squared between simulated and observed data using the rms of multiple realizations as the error bars. Given a geometrical model (e.g. a spherical distribution of hot plasma with a cavity excised), the only free parameter is the overall normalization, which then gives the field strength.

[Actually this isn't quite true unless we fix the value of eta ....?]

We now show the results of these fits explicitly ....

For the spherical model, is there not automatically a cavity where the source is (even though the x-ray data do not show one)? If there isn't, then you have Faraday rotating material mixed in with the synchrotron emitting particles. Section 2.2 argues convincingly for the absence of internal Faraday rotation. If they have filled source with thermal plasma, how does that affect those arguments?

We agree that there should be a cavity, but emphasise that ours is the first analysis to treat the effects of a cavity on Faraday rotation quantitatively. We therefore want to retain the spherically-symmetric model for comparison with earlier work. The comment prompted us to consider a slightly different form for the cavity, in which we omitted the rapidly expanding part near the nucleus and kept the cylindrical section around the visible radio emission. This is the cavity we would expect in the absence of any evidence to the contrary from the RM distribution or X-ray imaging. We show results for this geometry, and compare them with our existing cavity model, which does a better job of explaining the rapid change in RM fluctuation amplitude across the nucleus.

Figure 19 is confusing. If the tick mark at the bottom represents 100 arcsec, then the cone extends for 250 arcsec or so. Not the 140 arcsec stated in the text. What are the three concentric circles?

The original figure was plotted correctly, but it was perhaps not obvious that the two panels were on different scales (our attempt to indicate this was probably not clear enough). We have replotted the figure so that the two panels are on the same scale, and have noted in the caption that the arcs of circles represent isodensity contours.

The cavity model seems to generate the required asymmetry by virtue of the cone angle matching the angle to the line of sight. This seems contrived.

Yes, we agree. But this is necessary in order to generate the sharp change in RM fluctuation amplitude across the nucleus. Unless we have been misled by small-number statistics (which remains possible), a cavity of this general shape is necessary to explain the observed profile of rms RM.

Note that there is significant diffuse emission at low surface brightness which may fill in the required volume. We have replaced Fig 1 with a different representation of the same image (also repeated from Laing et al. 2008) in order to emphasise this point.

..................?

The disk distribution seems equally ad hoc. A diagram similar to Figure 19 would be very helpful here. Is there any precedence for a Faraday screen in the shape of an equatorial disk? Does any galaxy have a large disk of material similar to what they visualize?

We do not favour the disk hypothesis, and in retrospect probably gave it too much prominence by putting it in a separate subsection and showing simulation results. We now refer to it more briefly in the "alternative explanations" subsection. Given that, we think that a sketch is not essential.

We are aware of one apparently elongated distribution of hot gas in a radio galaxy: 3C 403 (Kraft et al. 2005, ApJ 622, 149). Although this has only been detected out to ~5 kpc from the nucleus, it may well extend much further - the optical isophotes are very elliptical on larger scales. The measured ellipticity is ~0.6. It is possible that the gas is in a disk. We now mention this reference.

Otherwise, the closest observed analogues are probably the disks in Cen A (HI, molecular and ionized gas) and NGC 612 (HI, ionized gas and stars). We now refer to these briefly. We do not think it likely that 3C31 has a structure similar to either of these. The only circumstance in which such a disk would be missed in 3C 31 would be if it were close to face-on, in which case the geometry is unsuitable for the purpose of creating a Faraday rotation asymmetry, even if the magnetoionic properties of such a disk prove to be consistent with our observations.

Faraday-rotating "superdisks" in more powerful radio galaxies have been suggested previously [Gopal Krishna & Nath 1997, A&A 326, 45; Gopal Krishna & Wiita 2000, ApJ 529, 189]. It is not clear to us whether such structures really exist. Certainly, the measured spatial extent of the disk and molecular gas disk in 3C31 is orders of magnitude smaller than the radio lobes. But we now refer briefly to this possibility.

Section 5.7 contains reasonable suggestions for other explanations. At this point the 3dimensional modeling, realistically, has been inconclusive, and perhaps the abstract should reflect that. Given that, comparison to Hydra A seems premature and makes the paper too long.

Although the 3D modelling is not definitive, we think that the cavity idea is much more plausible ....

Additional changes

We have fixed a formatting error which caused a spurious equation A6 label and have updated a few of the references.