Oxford Physics: Part C Major Option Astrophysics

High-Energy Astrophysics

Garret Cotter

garret@astro.ox.ac.uk

Office 756 DWB

Michaelmas 2011 Lecture 6

Today's lecture

Synchrotron emission Part III

- Synchrotron self-absorption
- Synchrotron radiative timescales.

Radiosource lifetimes

- Spectral aging.
- FRI and FRII sources

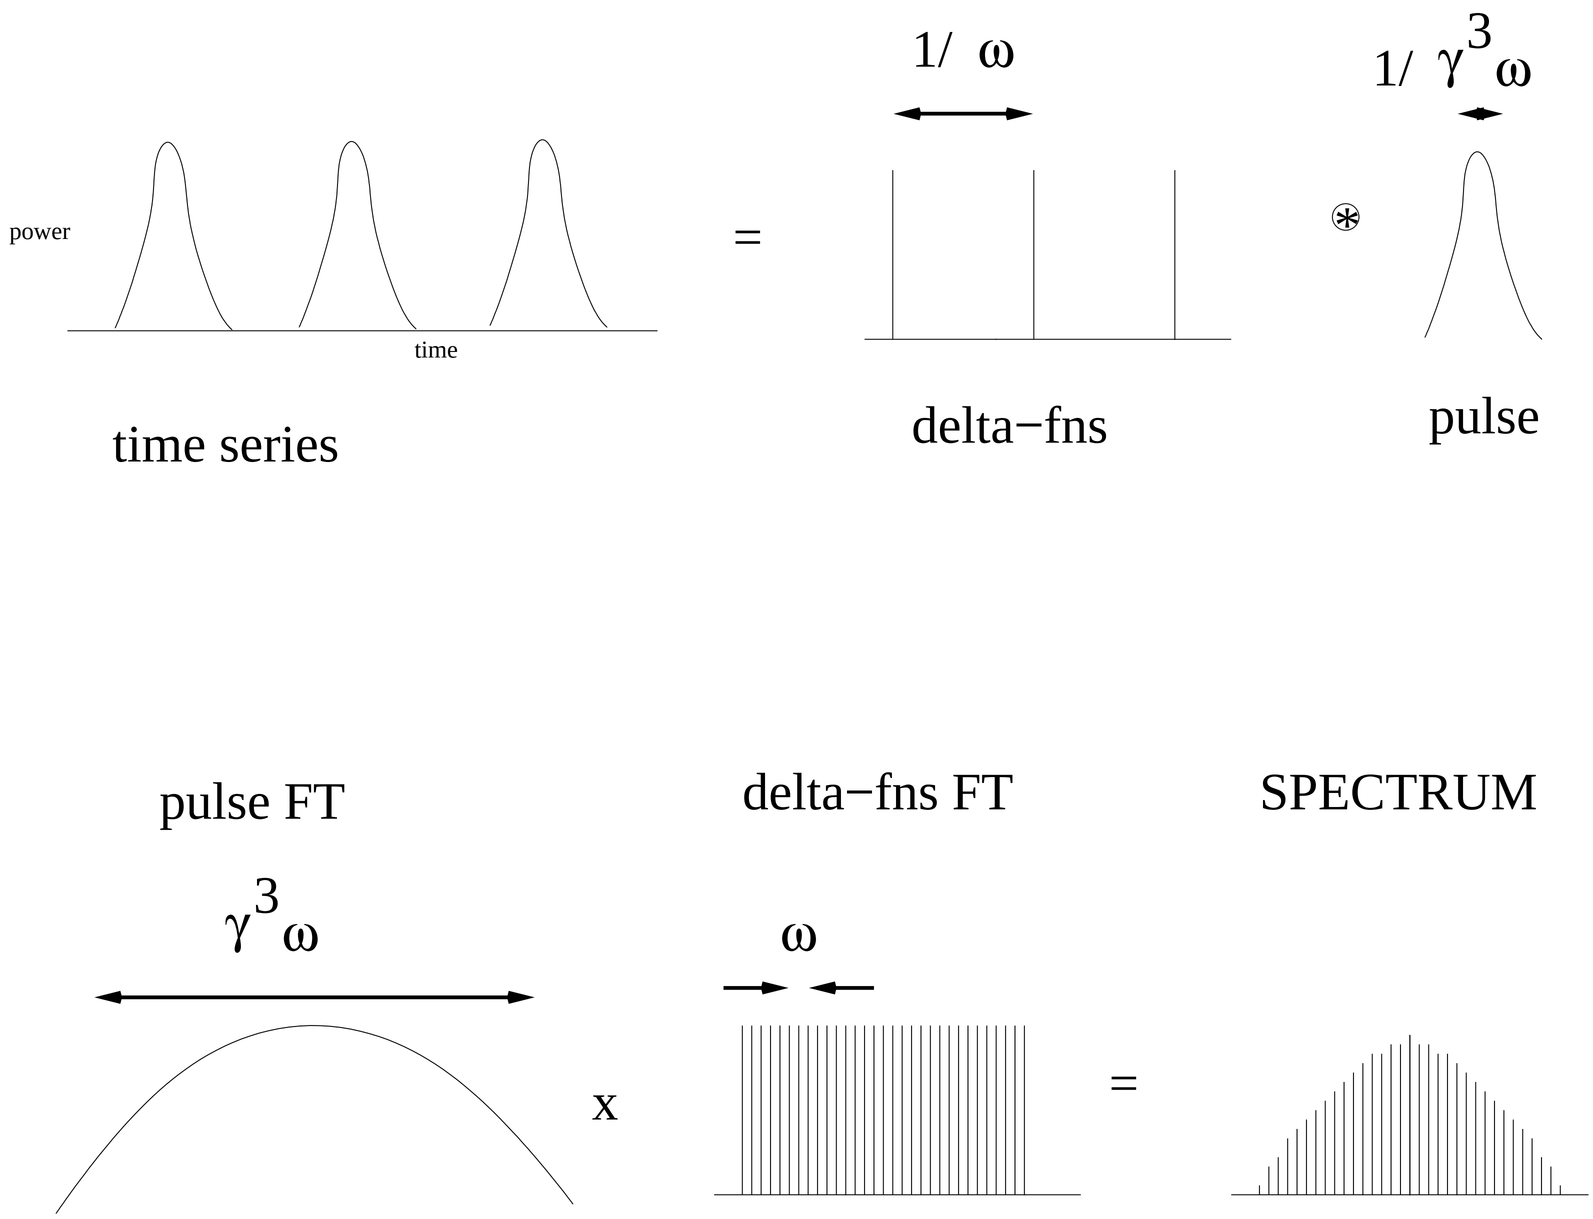
power
time
time series
=
1/ ω
delta−fns
1/ γ3ω
⊛
pulse
pulse FT
delta−fns FT
SPECTRUM
γ3ω
x
ω
=

Radio spectra

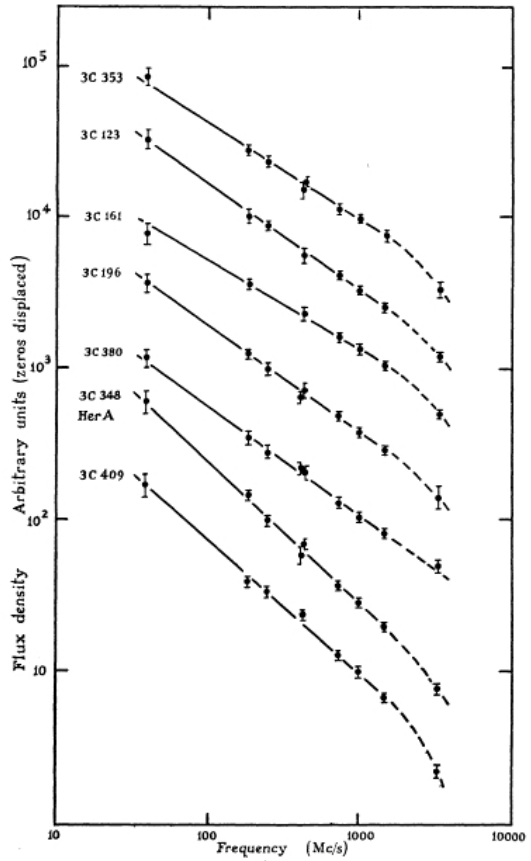

FIG. 1.—*Spectra of the seven calibration sources.*

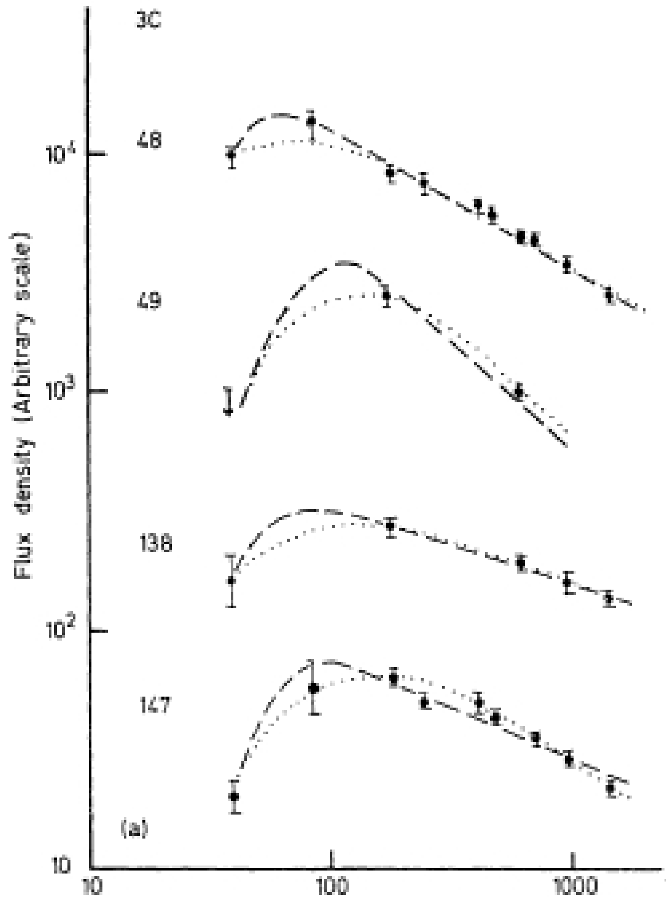

The power-law spectrum can be explained as synchrotron emission from a population of electrons whose energies have a power-law distribution $N(E)dE \propto E^{-k}dE$. Now we investigate the turnovers at high and low frqeuencies.

Scattering within synchroton plasma

In calculating the synchrotron spectrum so far, we have implicitly assumed that all of the synchrotron radiation emitted by each electron reaches the observer. However this is not necessarily the case: as a photon propagates through the plasma on its way out of the source, there is a chance that it will scatter off one of the synchrotron electrons. This is known as *synchrotron self-absorption*.

If such scattering occurs many times many times before the photon can get out of the source, the result is that an outside observer only “sees” emission from a thin layer near the surface of the source. Beneath this, the synchrotron electrons are simply exchanging photons, in a quasi-thermal-equilibrium fashion, and so the total flux the observer sees will be much smaller than if all the synchrotron photons escaped the source. This is analogous to looking at the surface of a dense object in thermal equilibrium—e.g. the surface of the Sun.

Scattering within synchroton plasma contd.

For a thermal distribution of particle energies, the calculation of the black-body spectrum observed is a fairly straightforward set piece. However we are dealing with a relativistic power-law distribution of particle energies, and each particle emits and receives radiation only at its characteristic wavelength. The detailed calculation of the self-absorption cross-section as a function of wavelength is difficult (Longair pp258–260) but the form of the result is intuitive: there is more absorption at longer wavelengths.

So for a source of a given size, we see very-long wavelength emission only from a very thin shell at the surface of the source. As our observing frequency increases, we see photons coming from regions of the source that are progressively deeper and deeper, and as we do so the toal flux density increases. However, of course, eventually we reach a point where we can “see all the way through” the plasma, and above this frequency we recover the underlying power-law distribution.

By analogy with the terms used in optical astronomy, we call the self-absorbed region of the spectrum, where the photon mean free path is much smaller than the source size, *optically thick*; and we call the remaining part of the spectrum, where the photon mean free path is greater than the source size, *optically thin*.

Synchrotron self-absorption: calculation

Without doing the detailed calculation we can however recover the spectral shape of the optically-thick spectrum, via a very simple (but in fact strictly correct) trick.

Let us assign an *effective temperature*, T_{Eff}, to the electrons in the synchrotron plasma. Although the particle energy distribution is not thermal, we can still define such an effective temperature according to

$$k_{\mathsf{B}} T_{\mathsf{Eff}} \sim \gamma m_{\mathsf{e}} c^2$$

We are now familiar with the result that the electron emits at a characteristic frequency, so we have

$$k_{\mathsf{B}} T_{\mathsf{Eff}} \propto E \propto \nu^{1/2}$$

Now here is the trick: even though the electrons are not in a thermal distribution, at any frequency they cannot emit radiation more effectively than a black-body. So we use the above functional form for $T_{\mathsf{Eff}}(\nu)$ in the Planck formula! This gives us the maximum brightness that a self-absorbed synchrotron plasma can have.

Synchrotron self-absorption: calculation contd.

Here is the full form for the brightness from a black-body (recall from lecture 1):

$$B_\nu = \frac{2h\nu^3}{c^2}\left\{\exp\left(\frac{h\nu}{k_{\mathrm{B}}T}\right) - 1\right\}^{-1}$$

We will make a simplification and use the Rayleigh-Jeans limiting form (as we are explicitly on the low-frequency tail here...)

$$B_\nu = \frac{2k_{\mathrm{B}}T\nu^2}{c^2}$$

Which, using $T_{\mathrm{Eff}} \propto \nu^{1/2}$, gives us

$$B_\nu \propto \nu^{5/2}$$

High-frequency cutoff

The upper limit to the particle energy that can be obtained by Fermi acceleration will depend on the properties of the shock in which the electrons are accelerated. In particular, a relativistic electron will have a gyro-radius of

$$r_g = \frac{\gamma m_{\mathrm{e}} v}{eB}$$

Once electrons reach energies sufficiently high that their gyro-radius is larger than the accelerating region, they won't get scattered up to still-higher energies. Nowadays we are fairly comfortable with first-order Fermi acceleration getting particles up to 10^{15} eV in supernova shocks, and all the way up to 10^{20} eV in AGN (we will later examine just how likely it is that such high-energy particles might survive their journey to us).

However, while such factors set a hard upper limit on the particle energies which can be achieved, the cut-off in the spectrum of the lobes of a radiosource is caused by energy loss in the particles *after* their initial acceleration: their synchrotron radiation drains the energy they acquired at the shock.

(Recall that this is why the intimal model of the atom with “orbiting” electrons was rejected...)

Synchrotron radiative timescale

We have seen that the power radiated by a synchrotron electron can be written as

$$P_{\mathrm{rad}} = \frac{4}{3}\sigma_{\mathrm{T}} c U_{\mathrm{mag}} \left(\frac{v}{c}\right)^2 \gamma^2$$

If the synchrotron power is the only source by which the electron loses energy, the implication, for highly relativistic electrons, is

$$-\frac{dE}{dt} \propto B^2 E^2$$

Consider then a population of electrons with some initial distribution of energies $N(E)$ in a uniform **B**: as time passes, the *higher energy electrons will radiate away their energy first*.

We can calculate a characteristic timescale for a synchrotron electron with energy E (and commensurate Lorentz factor γ) by simply taking the ratio of total energy to instantaneous power.

$$
\begin{aligned}
\tau &= \frac{E}{dE/dt} \\
&= \frac{E}{\frac{4}{3}\sigma_\mathrm{T} c U_\mathrm{mag} \left(\frac{v}{c}\right)^2 \gamma^2} \\
&= \frac{E}{\frac{4}{3}\sigma_\mathrm{T} c \frac{B^2}{2\mu_0}\gamma^2}
\end{aligned}
$$

Where in the last step we are approximating $\frac{v}{c} = 1$. To first order, this is the "lifetime" of a synchrotron electron: the time it would take to radiate away all of its initial energy via synchrotron radiation.

So whatever the highest energies of the electrons in the initial synchrotron population, if we return to a blob of synchrotron emitting plasma some time τ after it was accelerated, only electrons which have lifetimes longer than τ will remain.

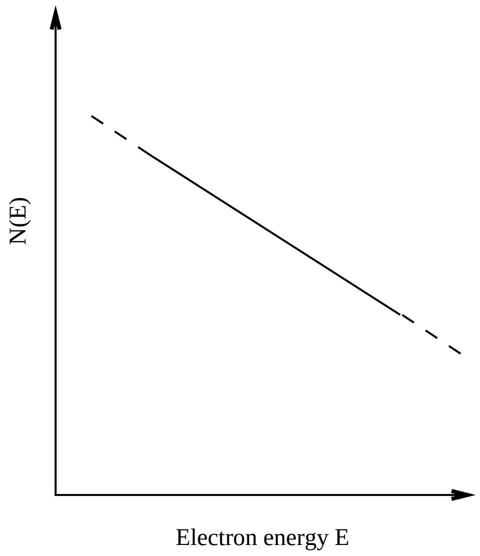

$t = 0$

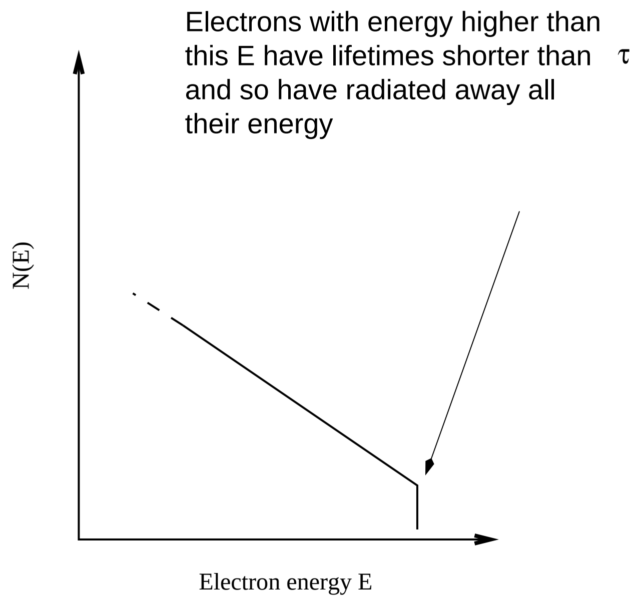

$t = \tau$

Now recall that we get the observed spectrum of synchrotron radiation by mapping convolving the electron energies with their characteristic emission frequencies.

$$\nu_{\mathrm{c}} = \frac{\gamma^2 e B}{2\pi m_{\mathrm{e}}}$$

So if our electron population is radiating in a uniform **B**, there will be *a characteristic frequency above which the emission falls sharply*. This is why we see the high-frequency cutoff in the lobes of a radiosource: the high-energy electrons are no longer present.

Synchrotron spectral "break"

The spectra we have examined so far are integrated over whole sources: the break becomes clearer if we examine the synchrotron spectrum from each part of a source individually.

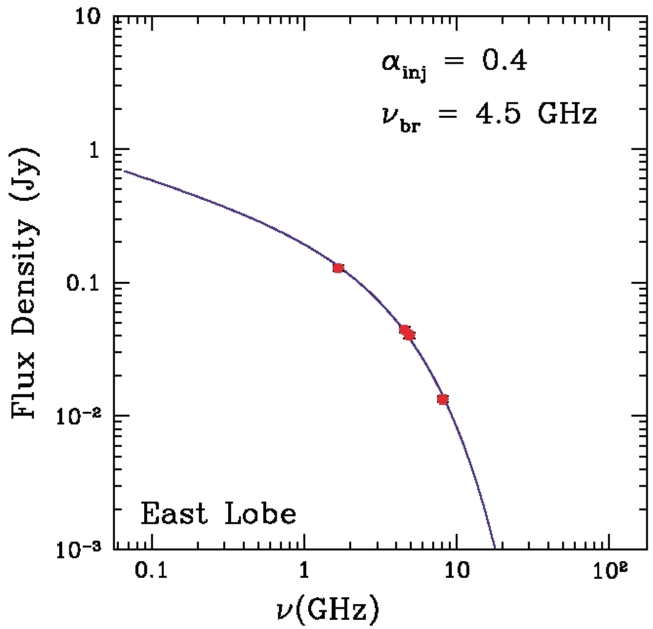

Spectrum of the lobe of a radiosource with a break at 4.5 GHz corresponding to the upper limit in the energy of the underlying electron population.

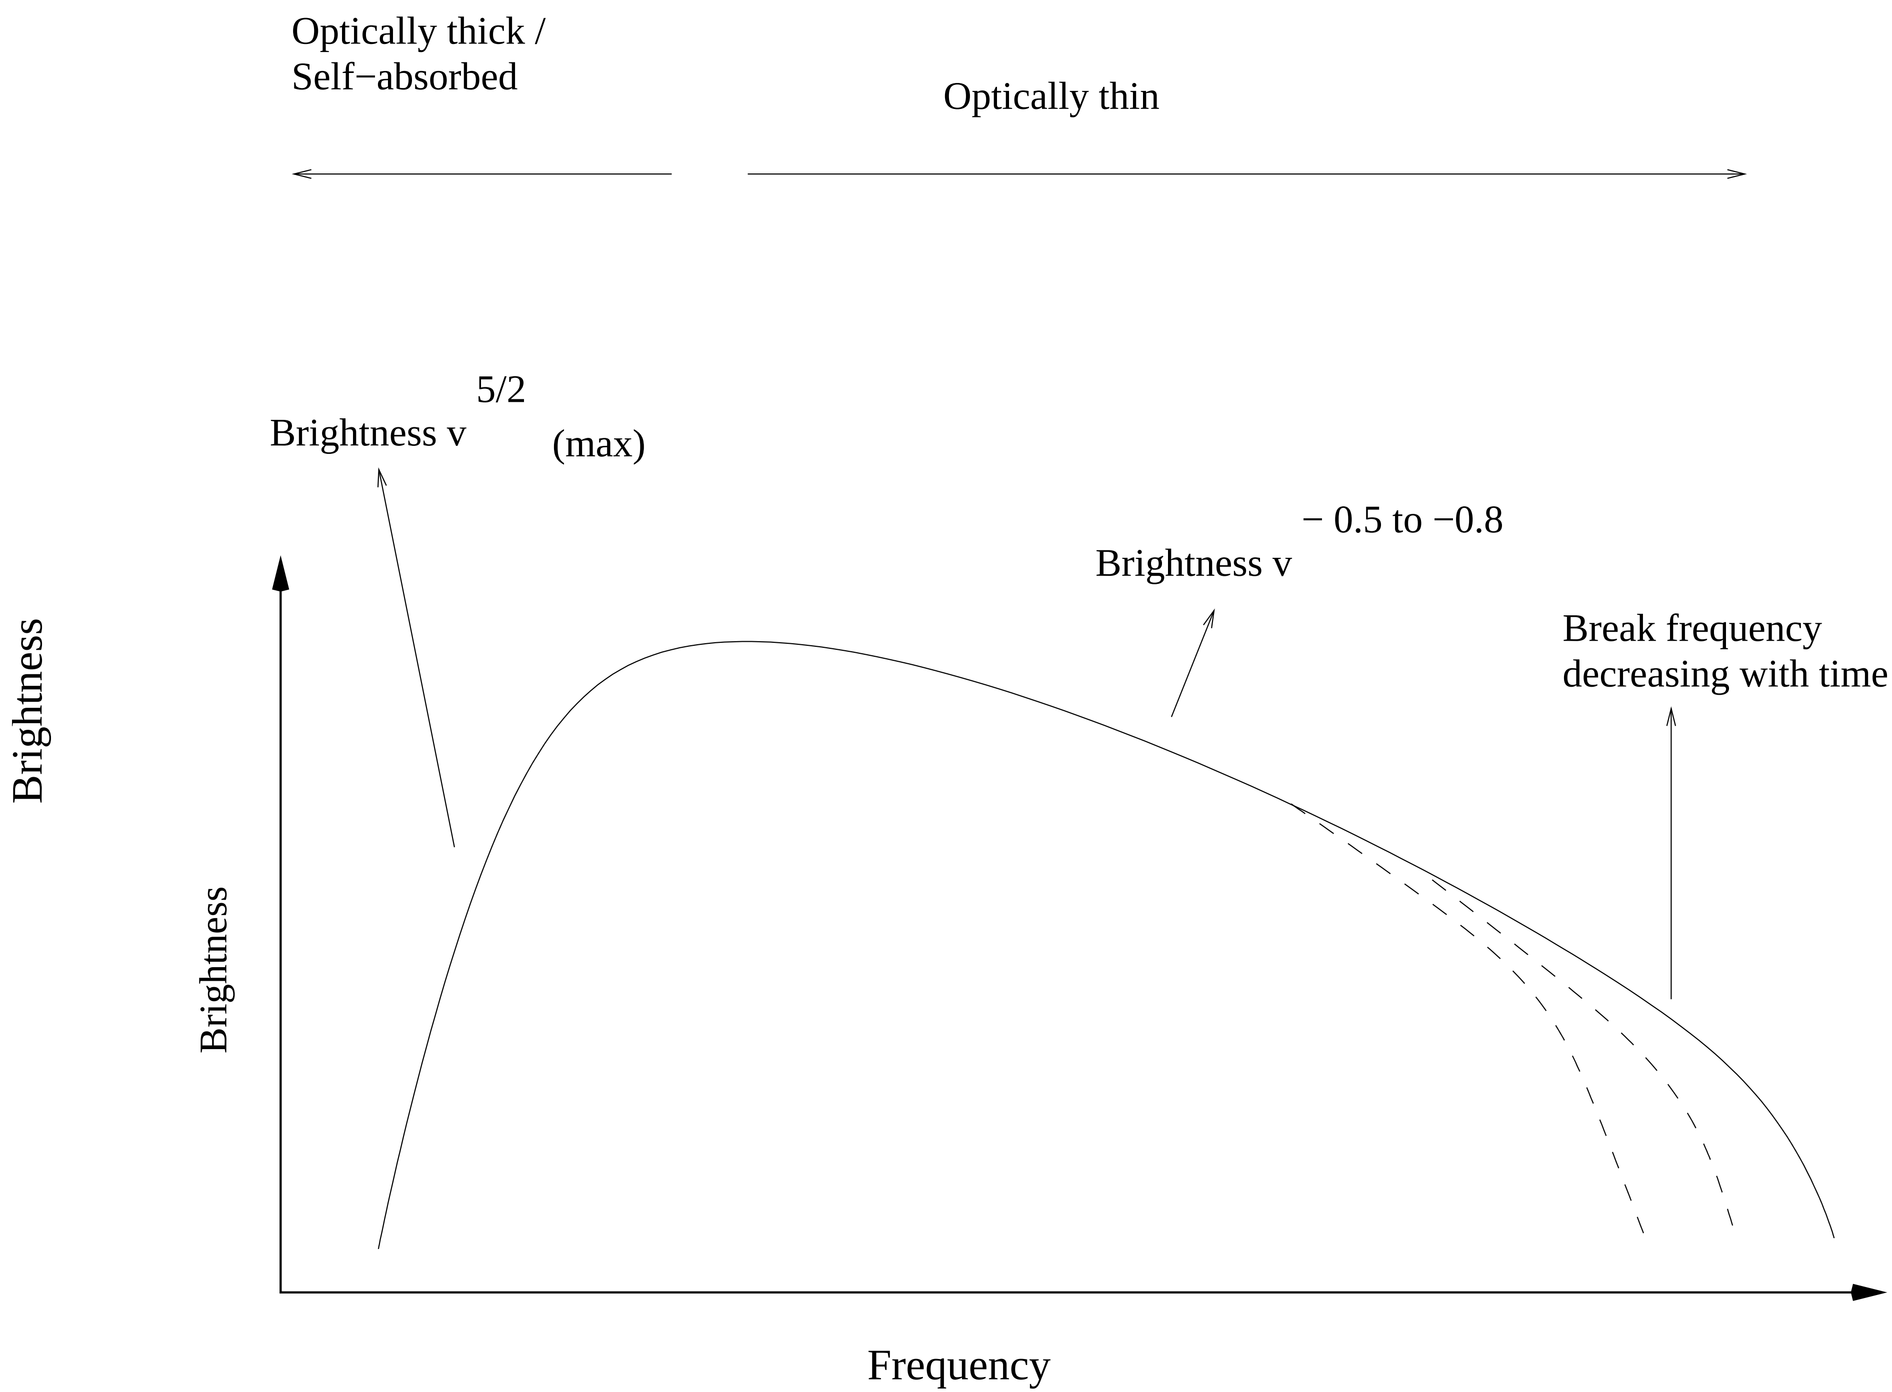

Optically thick /
Self−absorbed
Optically thin
Brightness v 5/2 (max)
Brightness v − 0.5 to −0.8
Break frequency
decreasing with time
Brightness
Brightness
Frequency

Synchrotron spectral ageing

We are now able to do something quite profound: for the sources with highest-power jets, we can estimate their ages. From high-resolution images, we have developed a picture where the jet ends in a strong shock against the intergalactic medium at the hotspot, and then the jet plasma flows out to fill the lobes of the source.

Let us assume that the last time the plasma is accelerated occurs at the hotspot, and that it is left behind as the source grows. We have already argued that **B** across the lobe should be reasonably uniform, and we can make a fair estimate of its value from minimum energy / equipartition arguments.

So we should see the spectral “break” shifting to lower and lower frequencies as we move back from the hotspot, to “older” plasma...

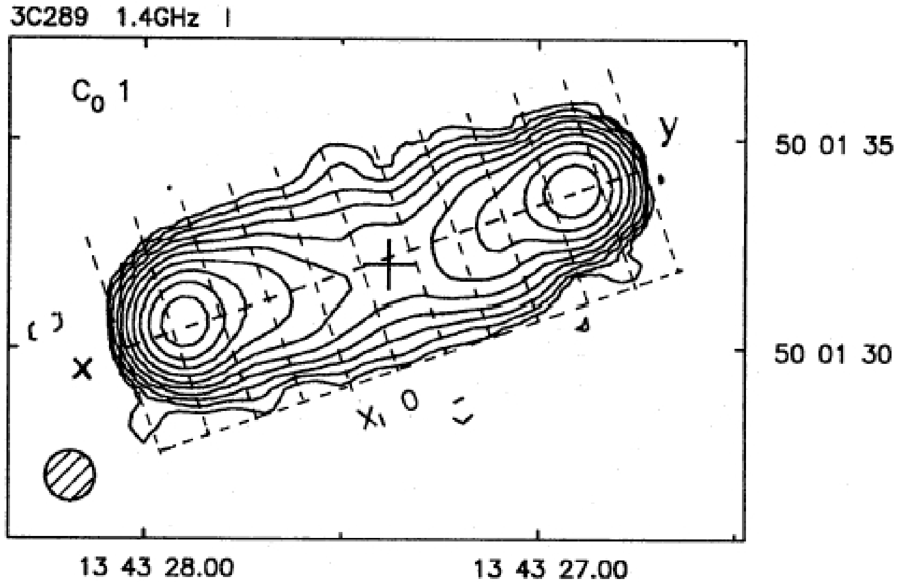

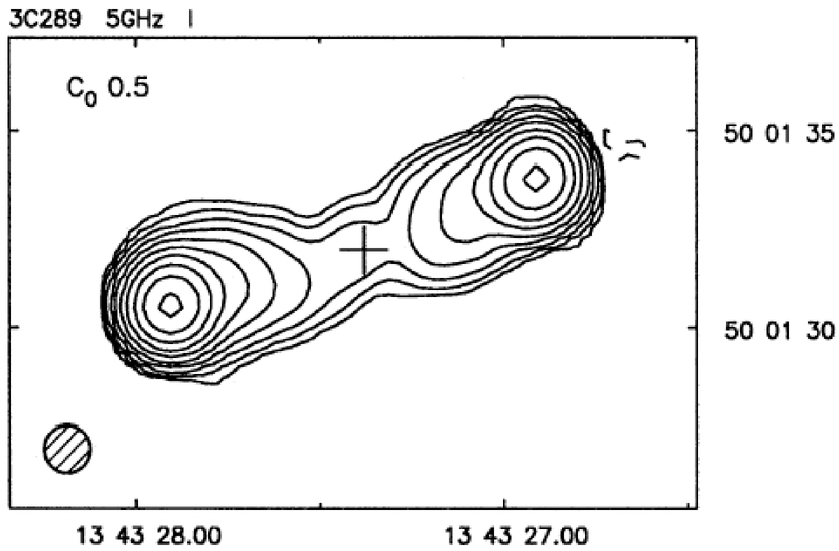

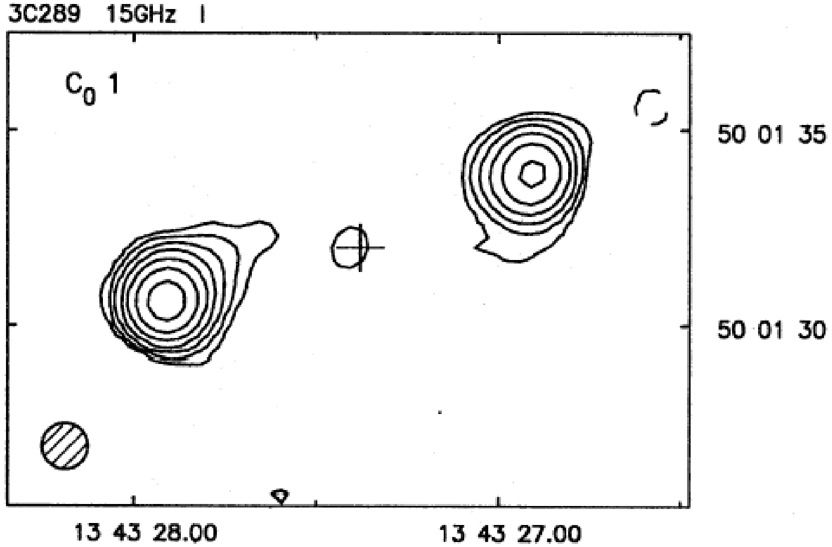

3C289 15GHz I
C0 1
50 01 35
50 01 30
13 43 28.00
13 43 27.00

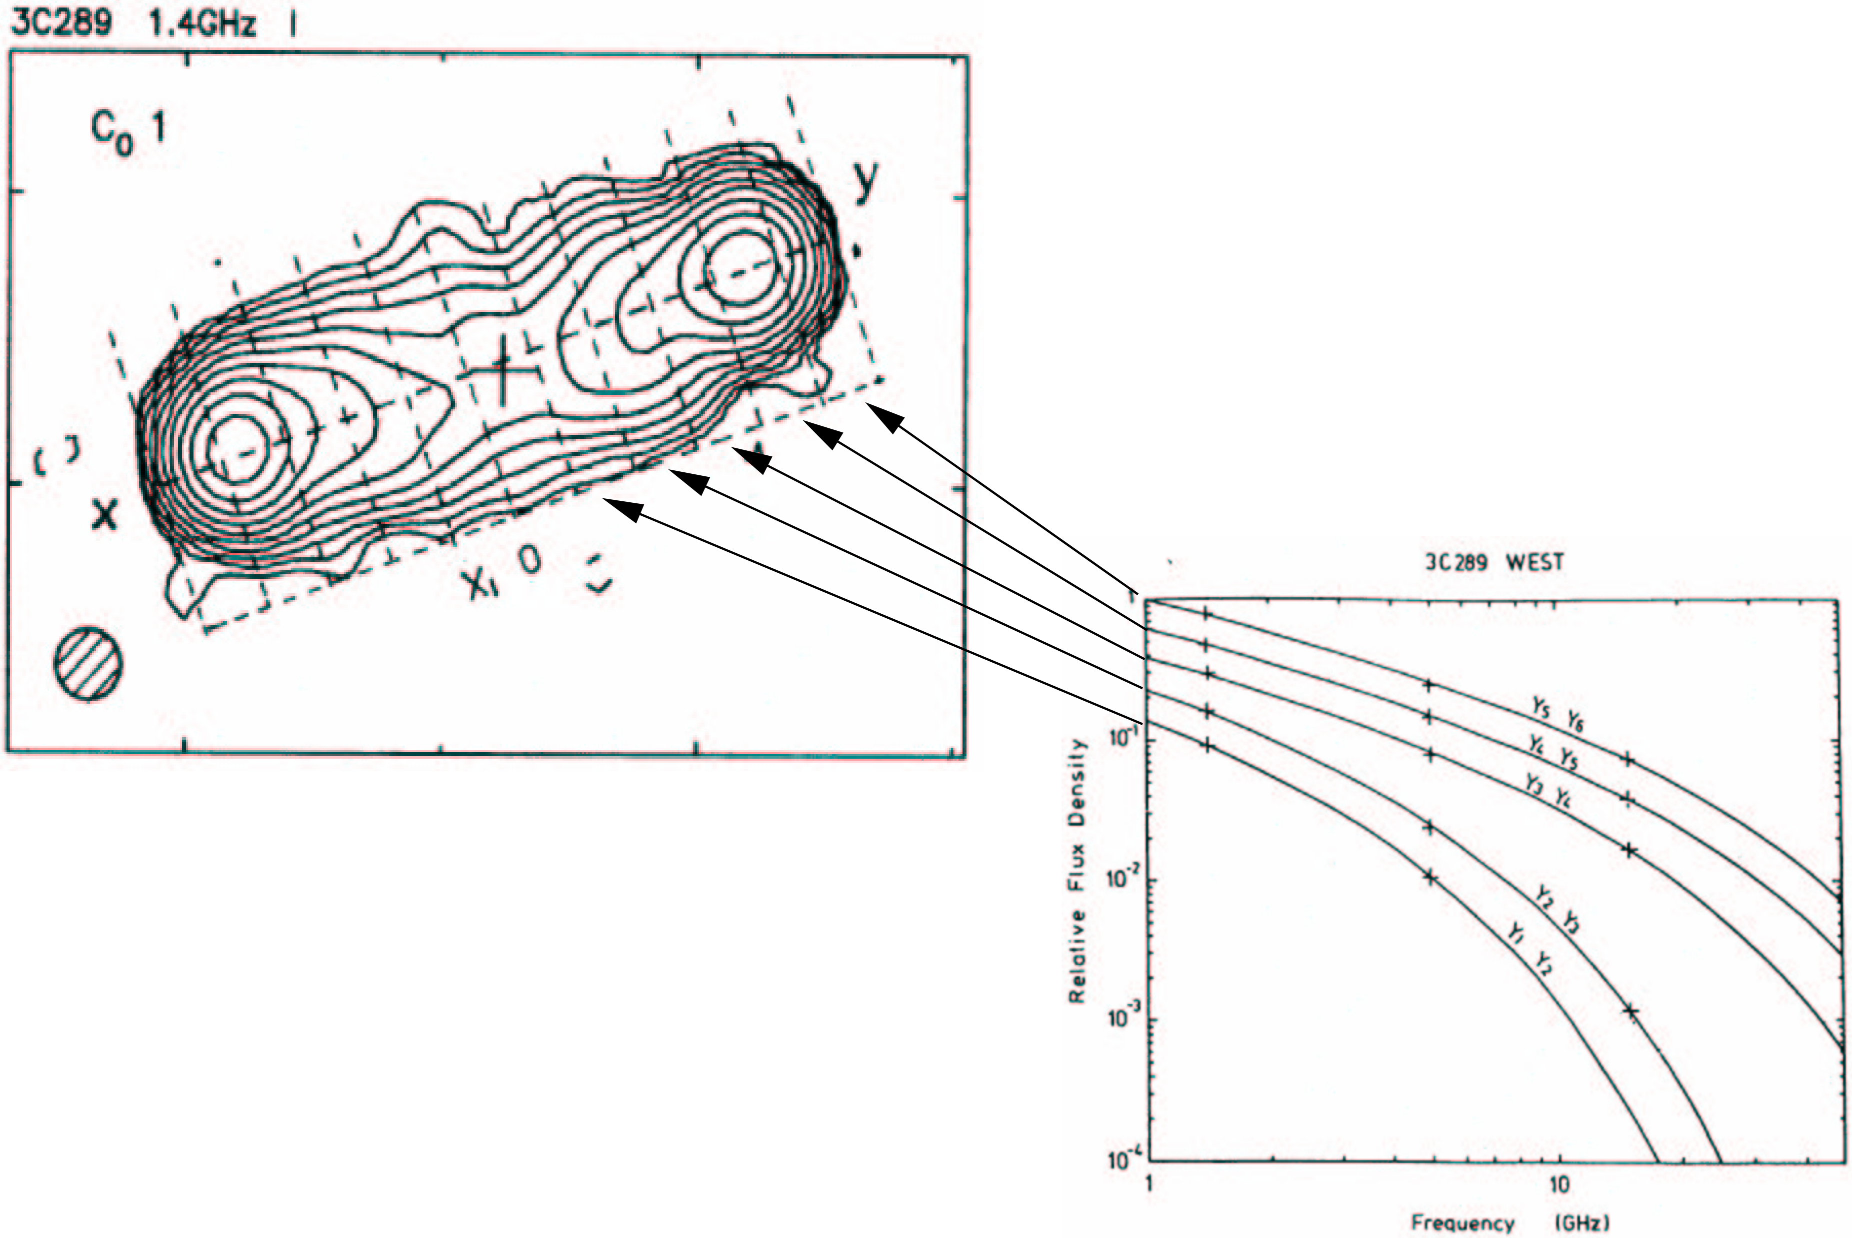
3C289 1.4GHz I
C_0 1
y
x
X_1
3C289 WEST
Y_5 Y_6
Y_4 Y_5
Y_3 Y_4
Y_2 Y_3
Y_1 Y_2
10^{-1}
10^{-2}
10^{-3}
10^{-4}
1
10
Relative Flux Density
Frequency (GHz)

Synchrotron spectral ageing: ingredients

- Measure the “break” frequency along the lobe from hotspot to nucleus.

- Use the total synchrotron luminosity to estimate the magnetic flux density in the lobes.

$$\nu_{\mathsf{c}} = \frac{\gamma^2 eB}{2\pi m_{\mathsf{e}}}$$

$$E = \gamma m_{\mathsf{e}} c^2$$

$$\frac{dE}{dt} = \frac{4}{3}\sigma_{\mathsf{T}} c U_{\mathsf{mag}} \left(\frac{v}{c}\right)^2 \gamma^2$$

$$\tau = \frac{E}{dE/dt}$$

Synchrotron spectral ageing: *caveat emptor*

Be aware that strictly we are not measuring the advance speed of the hotspot into the intergalactic medium, but the speed at which the hotspot and lobe plasma are separating.

Sometimes we see evidence that the lobe material has strong bulk flows back towards the nucleus, so the inferred speed would be an overestimate. Such bulk flows can, hower, sometimes be inferred directly from the map of the radiosource.

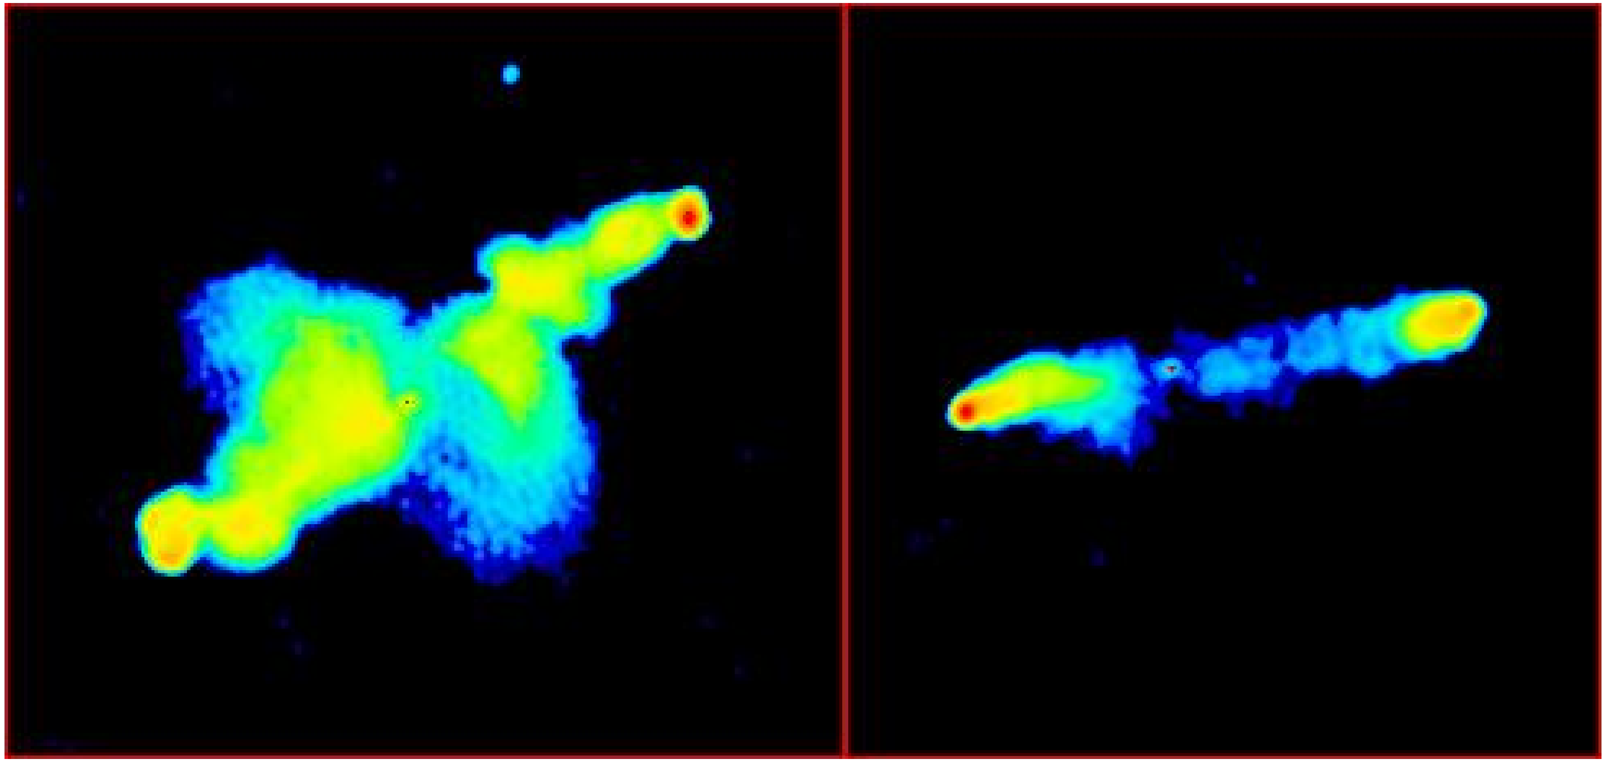

Synchrotron spectral ageing: results

- The powerful sources advance into the intergalactic medium at speeds of $\sim 0.01c$, with the most extreme examples approaching $0.1c$.

- The growth speed is roughly constant during the lifetime of an individual source.

- The oldest (and largest) sources observed have ages of a few 10^8 years.

- This age is a small fraction of the Hubble time.

The most powerful radiosources are short-lived beasts whose existence is probably limited by the fuelling of the AGN's central black hole. Given the overall number of radiosources that we see out to high redshift, it is likely that this is a phase which many of the most massive galaxies go through at some point in their evolution.

The biggest catch: lower-luminosity sources

Up to now we have been concentrating on radiosources whose jets are powerful enough to remain collimated as they extend into the intergalactic medium, terminating at a shock which defines the edge of the radio lobes.

In the 1970's, as high-resolution maps of radiosources became available, it was apparent that the lower-power sources had a different morphology: the jets do not end in a bright shock at the hotspot but are disturbed close to the host galaxy.

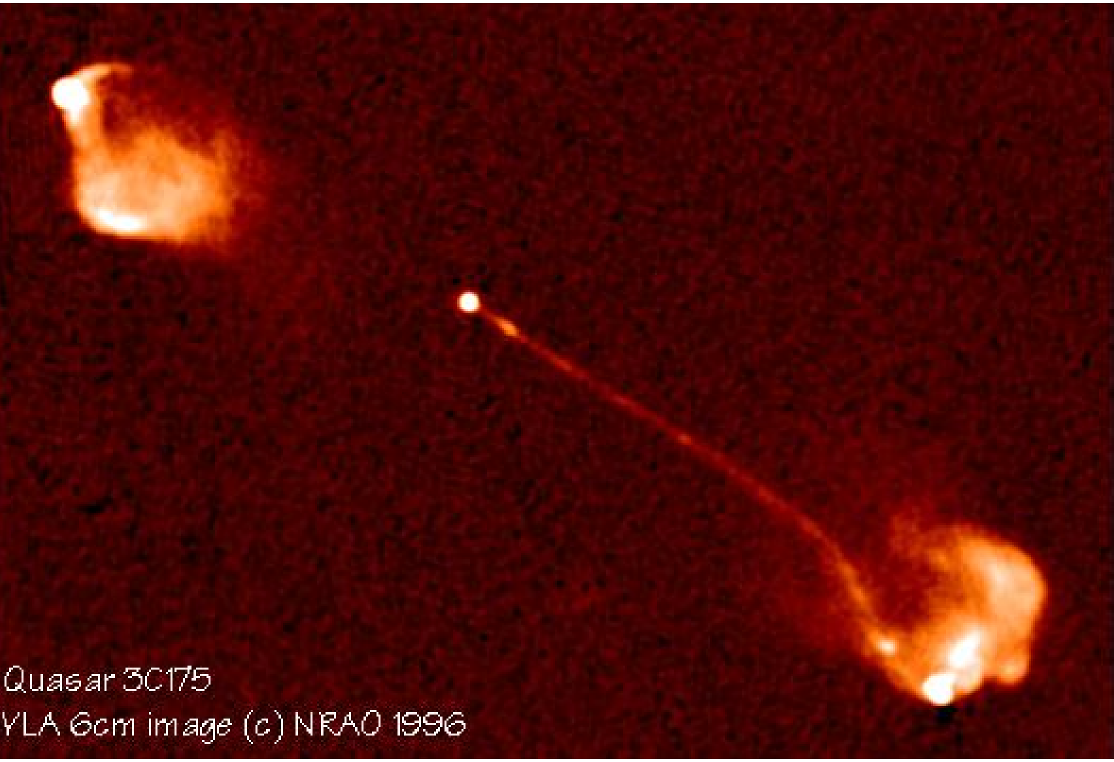
Quasar 3C175
VLA 6cm image (c) NRAO 1996

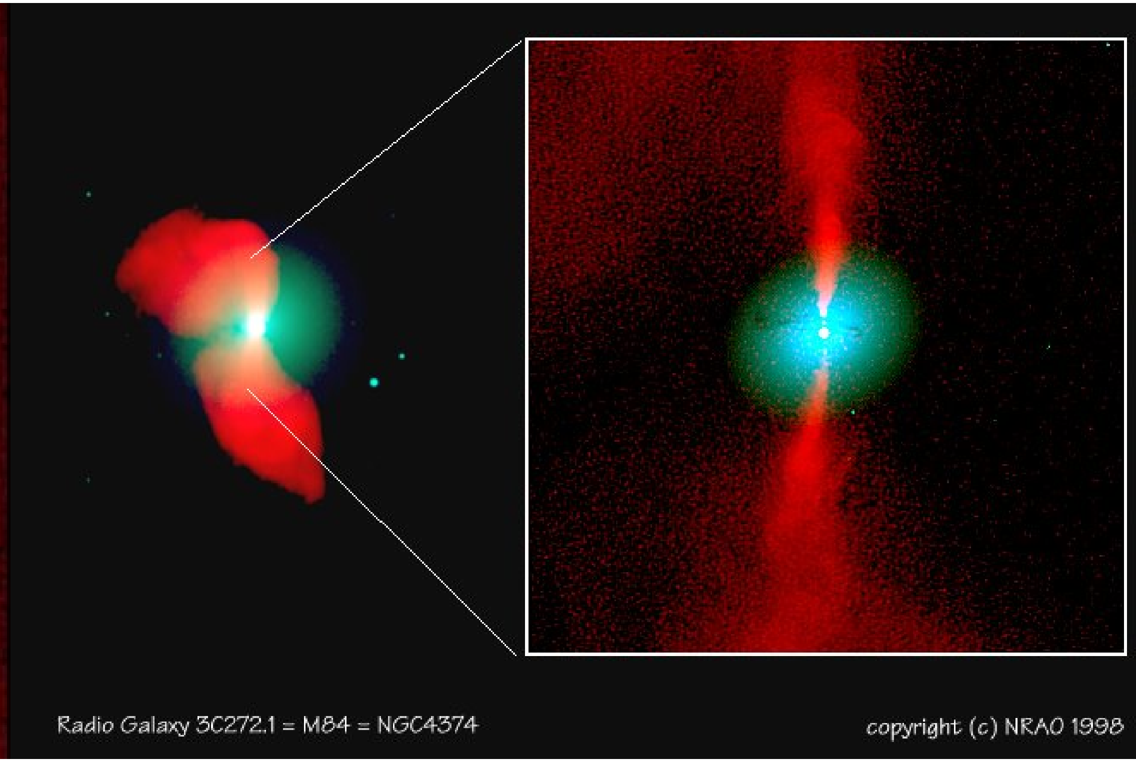
Radio Galaxy 3C272.1 = M84 = NGC4374
copyright (c) NRAO 1998

The nomenclature here was defined by Fanaroff and Riley who first made the connection between power and morphology. The lower-power sources with disturbed jets are type FRI, and the high-power sources are type FRII.

We will not be studying the (relativistic magnetohydrodynamics) physics of the decelerating, turbulent jets of the FRI sources in detail, but there are several important qualitative points.

 

Julia Riley & Bernie Fanaroff

- Near the host galaxy, the jet speeds are subsonic. The turbulence is a Rayleigh-Taylor instability—compare the radio images of FRI sources with pictures of smoke rising from chimneys and mushroom clouds.

- The plasma near the host galaxy is younger, and that further away from the source is older. Note that this is the inverse of FRII sources—the plasma is not being dumped at a shock front, but is gently “blowing away” from the host.

- These properties mean we have *no idea* just how old FRI sources are. Are they remnants of FRII sources, whose jets have powered down? If so, are they very long-lived? Or are they just short-lived sources with low-power jets?

Next: accretion and jet production