

1

Black Holes and Galaxies

1.1 Basic Properties

A black hole is an object so dense that light cannot escape its gravity. The proper description of black holes requires general relativity (GR), and we discuss this in Chapter 2. But simple Newtonian ideas already give us some insight into their properties if we bear in mind the restrictions to velocities below the speed of light and energies smaller than the rest-mass value.

The escape velocity from the surface of a star of mass M and radius R is $v = (2GM/R)^{1/2}$, where G is the gravitational constant. This velocity reaches the speed of light, c , for a radius

$$R = \frac{2GM}{c^2}. \quad (1.1)$$

We see that for $M = 1M_{\odot}$ (the solar mass), the ‘star’ must have a tiny radius $R \lesssim 3$ km. The characteristic property of a black hole is that it is small for a given mass, making the gravitational field very strong in its immediate vicinity. But it is important to remember that at large distances from a black hole, the gravitational field strength is the same as for any gravitating object of the same mass. Black holes have their distinctive properties only because they are small enough to allow matter to get very close, so that orbital speeds approach that of light. The characteristic size of a black hole is its gravitational radius

$$R_g = \frac{GM}{c^2} \simeq 1.5 \times 10^{14} M_8 \text{ cm}, \quad (1.2)$$

where we have parametrized the black hole mass $M = 10^8 M_8 M_{\odot}$, as this is a typical SMBH mass.

Matter falling radially towards an object like this acquires very high speeds because of the large gravitational potential energy available near the black hole. If nothing intervenes to stop it getting very close to the hole the matter eventually

gains gravitational energy approaching $\sim 0.5c^2$ times its rest mass. If the matter is gaseous, it is very likely that the process is sufficiently untidy that much of this energy is dissipated as radiation. The luminosity this releases is far greater than would happen if the same mass was consumed in nuclear burning. Transmuting hydrogen to helium or heavier elements, which powers most stars, releases energy only $0.007c^2$ times the rest mass.

In reality two things complicate this comparison a little. First, in reality matter is almost certain to orbit the black hole with some angular momentum, and fall towards it more slowly, as it gradually loses this angular momentum to matter further out. This kind of configuration is called an *accretion disc*, and will appear everywhere in this book. Accretion through a disc gives up the gravitational binding energy of orbits close to the black hole, which is somewhat less than the radial infall kinetic energy (a factor of two for circular Newtonian orbits).

Second, GR changes these binding energies slightly from Newtonian values. A full GR calculation (see Section 4.1) refines the estimate of the infall energy $\sim 0.5c^2$ to a value $\sim 0.1\text{--}0.4c^2$, but this slightly reduced ‘accretion yield’ is still far higher than the nuclear yield $0.007c^2$. Matter–antimatter annihilation releases the full rest-mass energy, but is very unlikely on any scale larger than an atomic nucleus, so that accretion on to a black hole is the most efficient way of getting energy from normal matter. We conclude that

*accretion on to black holes must power the most luminous objects
in the Universe.*

The obvious candidates here are quasars and active galactic nuclei, collectively called AGN, which harbour the most massive black holes. Their luminosities can reach $10^{46}\text{--}10^{48}$ erg s^{-1} or even more.¹ At typical distances of Mpc, the angular sizes of their gravitational radii R_g are extremely small. But radio interferometry with extremely long baselines is now able to resolve some of the nearer AGN, giving spectacular images (e.g. Figure 1.1).

At a smaller mass scale, the same argument tells us that binary systems where a stellar-mass black hole accretes gas at a high rate from a companion star are good candidates for explaining some of the stellar-mass X-ray binaries, with luminosities up to $10^{38}\text{--}10^{39}$ erg s^{-1} . Of course there are no AGN analogues of neutron stars, which cannot have masses larger than about $3M_\odot$.²

¹ Gamma-ray bursts can for extremely short times exceed this luminosity, and reach $\sim 10^{53}$ erg s^{-1} . This approaches the total luminosity L_* of all the stars in the observable Universe, which contains $\sim 10^{11}$ galaxies, each with $\sim 10^{11}$ stars emitting roughly solar luminosities $\sim L_\odot \sim 10^{33}$ erg s^{-1} . This gives $L_* \sim 10^{55}$ erg s^{-1} . See Problem 1.1 at the end of the book which investigates if this is a coincidence.




Figure 1.1 The immediate surroundings of the supermassive black hole in the galaxy M87 as imaged in the radio by the Event Horizon Telescope (Event Horizon Telescope Collaboration, 2019). The scale of the solar system is shown for comparison. Credit: Randall Munro (2019).

Observations clearly distinguish between stellar-mass accretors and SMBH. X-ray binaries are in spatially extended populations in their host galaxies, while AGN are point sources close to the centres of their hosts, and are generally intrinsically far brighter than X-ray binaries. They cannot simply be unresolved collections of X-ray binaries because they are often observed to vary by factors of $\gtrsim 2$.

1.2 The Eddington Limit

The accretion luminosity L of a black hole must be related to its mass accretion rate \dot{M} by

$$L = \eta \dot{M} c^2, \quad (1.3)$$

² The minor complication here is that neutron stars are almost as compact as black holes – they have radii ~ 10 km for masses $\sim 1\text{--}3M_{\odot}$. Both black hole and neutron star X-ray binaries are fairly common, although there are probably far more neutron star systems in total.

where we see from the discussions of the last previous section that the *accretion efficiency* η is a dimensionless quantity of order at most a few times 0.1. But there is a limit to the luminosity that any gravitating object, accreting or otherwise, can emit, since radiation produces a pressure force which tends to disperse the matter producing the luminosity. This force acts on electrons because they scatter electromagnetic radiation, which carries momentum $(1/c)$ times its energy flux $L/4\pi r^2$. Protons have little effect on radiation, but make up most of the mass of the gas. For simplicity we consider a spherically symmetric situation, so that the radiation pressure force acts radially outwards. Its magnitude at radius r from the centre is

$$F_{\text{rad}} = \frac{L\sigma_T}{4\pi cr^2}, \quad (1.4)$$

where $\sigma_T \simeq 6.65 \times 10^{-25} \text{ cm}^2$ is the Thomson cross-section, the effective blocking area of an electron in a beam of radiation. The electron is not free to move in response to this outward force, since charge neutrality means that it is strongly bound by Coulomb attraction to a mass of gas carrying one proton charge. Most astrophysical gases are largely hydrogen, so this mass is of order the proton mass m_p . Then the gravity force resisting the radiation pressure is

$$F_{\text{grav}} \simeq \frac{G(m_p + m_e)}{r^2} \simeq \frac{GMm_p}{r^2}, \quad (1.5)$$

since the electron mass m_e is much smaller than m_p . Both F_{rad} and F_{grav} vary as r^{-2} , so we see that if one of them exceeds the other at any one radius, it does so at all radii. Then in spherical symmetry, accretion must be at least inhibited once L is large enough to make $F_{\text{rad}} = F_{\text{grav}}$. This defines the *Eddington luminosity* or *Eddington limit* as

$$L_{\text{Edd}} = \frac{4\pi GMc}{\kappa}, \quad (1.6)$$

where we have used the electron opacity $\kappa \simeq \sigma_T/m_p \simeq 0.34 \text{ cm}^2 \text{ g}^{-1}$ for an astrophysical gas of typical composition, giving

$$L_{\text{Edd}} \simeq 1.3 \times 10^{46} M_8 \text{ erg s}^{-1}. \quad (1.7)$$

(The opacity κ is roughly halved, so L_{Edd} doubled, if the accreting gas is hydrogen-poor.) This result, and the huge difference in the mass that is accreting gas, explains why AGN can be far more intrinsically luminous than X-ray binaries.

Although we have derived it here for spherically symmetric systems, the limit (1.6) holds to factors of order unity for almost any other geometry. In particular, this is true even for matter falling in through a sequence of orbits of decreasing angular momentum – that is, in an *accretion disc*. The appropriate form of the Eddington limit applies to any luminous object, whatever powers its luminosity.

The importance of the Eddington luminosity was first realized a century ago in the context of stellar structure, where the radiation comes from nuclear burning. Massive hot stars radiate luminosities close to the limit (1.7). A nuclear luminosity even slightly above L_{Edd} would make them expand a little, lowering the density ρ in the central nuclear-burning core. The nuclear luminosity L_{nuc} varies as ρ^2 and so drops below L_{Edd} . This self-limiting property means that hot stars can remain stably in equilibrium very close to L_{Edd} . The source of the luminosity reacts sensitively – and negatively – to the luminosity itself, rather like a thermostat.

But this kind of self-limiting behaviour does not apply to accretion-powered objects. The mass supply rate driving accretion is in general given by some process totally unaffected by changes in the accretion luminosity, and so is unlikely to adjust to respect the Eddington limit. For example, in close binary systems there is no reason why the evolution of the donor star, and so the resulting mass transfer rate \dot{M}_{supp} it supplies to a companion black hole, should know or care about the possibility that \dot{M}_{supp} might exceed the rate

$$\dot{M}_{\text{Edd}} = \frac{L_{\text{Edd}}}{\eta c^2} = \frac{4\pi GM}{\eta \kappa c} \quad (1.8)$$

that would make the accretion luminosity $L_{\text{acc}} = \eta c^2 \dot{M}_{\text{acc}} = \eta c^2 \dot{M}_{\text{trans}}$ greater than L_{Edd} . This possibility was already recognized in the very first papers discussing realistic accretion processes: it is entirely possible for a black hole (or any other accreting object) to be supplied with mass at rates $\dot{M} > \dot{M}_{\text{Edd}}$, and for this situation to persist over significant timescales. We will discuss in detail what happens in such cases in Section 4.6, but it already is clear that there are only two routes to dealing with the mismatch – either

- (a) preventing much of the matter getting too close to the hole, where it would gain and then radiate the full accretion energy, or
- (b) ensuring that the matter close to the hole has unusually low accretion efficiency and so does not radiate a strongly super-Eddington luminosity.

The mild complexity of these two possibilities has generated a cloud of confused and confusing language in the astrophysical literature. The phrase ‘super-Eddington accretion’ is deeply ambiguous if not carefully qualified, as it is often used to denote either one of the outcomes (a, b) (or in the worst cases, both simultaneously!). To avoid this ambiguity, this book uses the description super-Eddington *mass supply* (or *feeding*) to describe cases where $\dot{M}_{\text{supp}} > \dot{M}_{\text{Edd}}$. In treating these, it is vital to distinguish between the outcomes (a) and (b), which differ markedly.³

³ This muddle is maximal in discussions of the (stellar-mass) ultraluminous X-ray sources (ULXs). ULXs have very anisotropic radiation patterns (see Section 4.10), and when viewed from tightly defined directions appear

In case (a), the black hole mass cannot grow faster than the rate \dot{M}_{Edd} . Then the shortest e-folding timescale for mass growth is the Salpeter timescale

$$t_{\text{Sal}} = \frac{M}{\dot{M}_{\text{Edd}}} \simeq 5 \times 10^7 \eta_{0.1} \text{ yr} \quad (1.9)$$

(Salpeter, 1964), where $\eta_{0.1}$ is the efficiency of conversion of rest-mass energy to radiation in units of $0.1c^2$. We see that high radiative efficiency implies slower mass growth, as the limiting luminosity L_{Edd} is produced by a smaller accretion rate.

If instead (b) holds, the accretor mass grows faster than the Eddington rate, that is, on a timescale $< t_{\text{Sal}}$.

In both cases (a) and (b) it is difficult for the total luminosity (correctly evaluated over all directions in the case of anisotropy – see footnote 3) of an accreting object to exceed the Eddington luminosity L_{Edd} by large amounts, except in impulsive or explosive situations such as supernovae or gamma-ray bursts. If we are confident that a given object is not of this type, and its luminosity is not markedly anisotropic, its luminosity gives us a lower limit to its mass through (1.7).

1.3 SMBH Accretion

By the argument detailed previous section, observation places tight constraints on the total mass in SMBHs in the local (low-redshift) Universe. AGN spectra (cf. Figure 1.2) typically peak in the soft X-ray–far UV region, with almost always a significant component in the medium-energy X-rays. This latter component is relatively easy to observe, as it is often fairly immune to interstellar absorption or scattering. In addition, very few astronomical objects other than those that accrete produce substantial X-ray emission, so it is usually safe to assume that all the detected emission comes from the AGN itself.

Medium-energy cosmic X-ray detectors find non-zero fluxes even when not observing specific point sources – this is called the X-ray background. Since the emission from AGN in this spectral band is far more powerful than from anything else, such as X-ray binaries or supernovae, we can identify this background flux as the result of the collective emission from AGN – that is, growing SMBH – in the local Universe. From the typical X-ray spectrum (cf. Figure 1.2) this gives us the total SMBH growth in this region, and so a lower limit on the total SMBH mass there. This is in outline the *Soltan argument* (Soltan, 1982). It tells us that the

to have luminosities $\gg L_{\text{Edd}}$, if the observed flux is (wrongly) assumed to be isotropic. Confusingly, the indirect cause of the anisotropic emission, and so of the apparent super-Eddington luminosity, is super-Eddington *feeding*, since radiation pressure blows matter away from the accretion disc (see (a)) except near the rotational axes of the accretion flow, so collimating the escaping radiation tightly. The accretors in ULXs do not gain mass at significantly super-Eddington rates, even though the mass transfer rates from their companions are super-Eddington. Their luminosities are *apparently* super-Eddington, but in reality are not.




Figure 1.2 Average total spectrum (thick line) and main components (thin lines) in the X-ray spectrum of a type I AGN. The main primary continuum component is a power law with a high energy cut-off at $E \sim 100\text{--}300\text{ keV}$, absorbed at soft energies by warm gas with $N_H \sim 10^{21}\text{--}10^{23}\text{ cm}^{-2}$. A cold reflection component is also shown. The most relevant narrow feature is the iron $K\alpha$ emission line at 6.4 keV . Finally, a ‘soft excess’ is shown, resulting from thermal emission of a Compton thin plasma with temperature $kT \sim 0.1\text{--}1\text{ keV}$. Credit: Risaliti and Elvis (2004).

average ration of SMBH mass is $\gtrsim 10^8 M_\odot$ per medium-size galaxy. This mass scale agrees with the values we can deduce from the Eddington limit (1.7), and we shall see that it is similar to the masses found by various kinds of direct observation, so it cannot be concentrated in just a minority of galaxies. We conclude that

the centre of almost every medium to large-mass galaxy in the local Universe must host a supermassive black hole, whose mass grew via luminous accretion of gas.

Despite this, observations show that only a minority (less than 1%) of low-redshift galaxies have active nuclei, where the SMBH are currently growing their masses. But since we have just concluded that almost every galaxy has an SMBH, this must mean that accretion and SMBH growth occur only in short-lived phases – AGN must be strongly variable, not simply on the timescales we observe directly, but on longer ones also.

1.4 SMBH Locations

The reasoning of Section 1.3 implies that almost all low-redshift galaxies must host SMBH. Observations of the active minority where the SMBH is caught in the act of accreting always find the AGN close to the dynamical centre (hence the ‘nucleus’

part of ‘AGN’). This is no accident, but results because the gravity of an SMBH moving through a galaxy makes it pull large numbers of stars along behind it, in a kind of gravitational wake (see Figure 1.3). This process is called ‘dynamical friction’, and is directly analogous to the way the Coulomb attraction of a charged particle slows its motion through a plasma. The result (Chandrasekhar, 1943) is a drag force – the speed v of an SMBH of mass M moving through a galaxy with stellar mass density ρ obeys

$$\frac{dv}{dt} = -\frac{4\pi CG^2 M \rho}{v^2}. \quad (1.10)$$

Here $C \simeq 10$ is a constant (the Coulomb logarithm, measuring the cumulative effect of the weak drag forces of many distant stars (‘small-angle scattering’) compared with the individually stronger drag forces (‘large-angle scattering’) of a few nearby stars). If ρ is constant this equation integrates as

$$v^3 = v_0^3 \left(1 - \frac{t}{t_{\text{fric}}}\right) \quad (1.11)$$

where

$$t_{\text{fric}} = \frac{v_0^3}{12\pi CG^2 M \rho}, \quad (1.12)$$

with v_0 the initial velocity. In a time $\sim t_{\text{fric}}$ the SMBH is reduced to rest, which is only possible if it has spiralled in to the dynamical centre of the galaxy. Crudely modelling the central region of a galaxy as a uniform sphere of stars with total mass $M_* = 10^{11} M_\odot$ and radius $R_* =$ a few kpc, a $10^8 M_\odot$ SMBH with speed $v_0 \sim (2GM_*/R_*)^{1/2}$ has $t_{\text{fric}} \lesssim 10^8$ yr, much smaller than the age $\sim 10^{10}$ yr of a low-redshift galaxy.

So unless a galaxy has recently been disturbed, which is often obvious because it has an irregular shape, its SMBH is very likely to be at its dynamical centre. If it somehow has more than one SMBH, dynamical friction makes them collect rapidly in its centre, and then orbit under their mutual gravitational attraction. Here they lose energy and spiral inwards as they emit gravitational radiation, eventually merging. The merger may eject the lightest hole(s) if there are more than two, since these must be moving fastest if the holes have similar orbital energies, as is likely. So in most galaxies we expect to find just one SMBH, in its centre. (In very small galaxies the shallow gravitational potential may be unable to retain a merging pair of black holes as they recoil under anisotropic gravitational wave emission, so some may have no central black hole). By the same reasoning, a merger of two galaxies is likely to produce a single larger galaxy with a merged SMBH in its centre (see Figure 1.4).




Figure 1.3 dynamical friction: a massive object moving through a collection of stars raises a gravitational ‘wake’ and slows.




Figure 1.4 Galaxies in various stages of merging. Stages A and B are early mergers and stages C and D are late mergers. Credit: Ricci et al. (2017).

In summary, we expect almost every galaxy, except possibly a few dwarfs, to have a single SMBH at its centre. Observations appear to agree with this very simple one-to-one relation. Our own Galaxy’s central region is well studied. Infrared




Figure 1.5 Orbits of stars around the Galactic Centre supermassive black hole Sgr A* (position given by the white dot in the centre of the figure). All of these orbits, observed with great precision over decades, require a unique black hole mass $4 \times 10^6 M_{\odot}$. The closest orbits now show evidence for the Einstein precession (advance of the pericentre, as predicted by general relativity). Credit: S. Gillessen.

observations by groups in Germany and the United States have followed the proper motions of stars around it in exquisite detail (see Figure 1.5) for more than 25 years.⁴ Interpreting these motions as Kepler orbits shows that the moving stars orbit a central mass of order $4.5 \times 10^6 M_{\odot}$. Constraints on its size leave no room for doubt that this object (Sgr A*) must be the Galaxy's own SMBH.

This SMBH is remarkably inactive at the current epoch, but there is indirect evidence of activity in the past. In particular, two large gamma-ray-emitting lobes (the 'Fermi bubbles' – see Figures 1.6 and 1.7) are symmetrically placed each side of the Galactic plane with Sgr A* at the centre of symmetry. These are probably the result of an energetic outflow event from the SMBH about 6 Myr ago – the event was probably roughly isotropic, but the greater density of the Galactic plane means that propagation only occurs along the axis of the Galaxy.

⁴ Reinhard Genzel and Andrea Ghez shared half the 2020 Nobel Prize in Physics for leading this work.




Figure 1.6 Observations by the *Fermi*–LAT satellite of bubbles of gamma-ray emission symmetrically placed each side of the plane of the Milky Way. Credit: NASA/DOE/Fermi LAT/D. Finkbeiner et al.




Figure 1.7 Schematic picture of the Fermi bubbles superimposed on optical observations of the plane of the Milky Way. Credit: NASA Goddard.

1.5 The SMBH Environment

We have seen that dynamical friction ensures that most supermassive black holes are in the central bulges of spiral galaxies, or the centres of elliptical galaxies. Both of these environments are roughly spherical systems – often collectively called

spheroids – which are probably the outcomes of mergers with other galaxies. In a spiral the mergers were minor – that is, with smaller galaxies, whereas an elliptical probably resulted from a major merger of two galaxies of similar masses. The stellar motions in both types of spheroid are generally characterized by velocity distributions similar to the Maxwellian distribution of the particles of a hot gas. The velocity dispersions σ are spatially constant, as expected theoretically (Lynden-Bell, 1967). The resemblance to a constant-temperature gas means this kind of distribution is called isothermal. To remain in dynamical equilibrium under its mutual self-gravity, the stellar mass density obeys

$$\rho = \frac{\sigma^2}{2\pi Gr^2}. \quad (1.13)$$

The stellar mass inside radius R is then

$$M(R) = 4\pi \int_0^R \rho(r)r^2 dr = \frac{2\sigma^2 R}{G}. \quad (1.14)$$

The SMBH's own gravity controls stellar motions only in the region where $M(R) \lesssim M$, so inside a radius

$$R_{\text{inf}} \simeq \frac{GM}{\sigma^2}, \quad (1.15)$$

called its sphere of influence. In a typical galaxy with black hole mass $M = 10^8 M_8 M_\odot$, $\sigma = 200\sigma_{200} \text{ km s}^{-1}$, we find

$$R_{\text{inf}} \simeq \frac{8M_8}{\sigma_{200}^2} \text{ pc}. \quad (1.16)$$

The radius of influence R_{inf} is much smaller than the typical scale $R_b \gtrsim 5 \text{ kpc}$ of a spiral bulge, or the even larger scale of an elliptical galaxy. Other components of the galaxy, such as gas, presumably have velocity dispersions comparable with the stars, and so do not feel the black hole gravity outside the sphere of influence R_{inf} either.

Within R_{inf} stars move in elliptical orbits under the black hole's influence, but they also exchange energy with each other by repeated gravitational encounters, adjusting their orbits accordingly. After one relaxation time, stars of a similar mass should be distributed with number density $\rho \propto R^{-7/4}$, out to a radius $\sim 0.2R_{\text{inf}}$ (Bahcall & Wolf, 1976). This implies a sharply rising stellar 'cusp' near the SMBH. Similar results hold if the restriction to similar stellar masses is lifted. But so far no clear example of a cusp like this has been found observationally in any galaxy. For example, in our Galaxy the cusp should contain about 10^7 stars and extend out to about 0.4 pc (since $R_{\text{inf}} \simeq 2\text{--}3 \text{ pc}$). This would be easy to resolve from Earth, but the observed distribution is flat or even declining towards the SMBH. It may

be that the relaxation time here is in reality comparable to the age of the Galaxy, so that no cusp has yet had time to form, or that the cusp is present, but in a population of stars that is so far unobserved. We shall see in Chapter 6 that there may be other reasons why the cusp is apparently not present.

Other dynamical effects of the central SMBH are directly observed. The SMBH can disrupt stellar binary systems orbiting it. The more massive star of a pair is captured into a closer orbit around the SMBH, while the gravitational energy this liberates goes into propelling the lighter star away from the SMBH at very high speed (Hills, 1988). Examples of these hypervelocity stars have been observed in our Galaxy with speeds approaching $2,000 \text{ km s}^{-1}$, and velocities pointing directly away from the Galactic Centre.

Even more violently, SMBH can completely disrupt stars whose orbits get too close to it – for example, stars falling inwards on near-parabolic orbits. In these tidal disruption events (TDEs) the tides exerted by the SMBH overcome the stars' own self-gravity. This happens if a star's radius R_* is large enough to fill its tidal lobe, of size $R_{\text{tide}} \simeq 0.5(M_*/M)^{1/3}a$, where M_* is its mass and a its distance of closest approach to the SMBH. We shall see in the next chapter that there is a minimum value for a for stable orbits not falling directly into the hole, of the form $a_{\text{min}} \gtrsim zR_g \propto M$, with $z \sim \text{few}$. The parameter z here depends on the spin of the black hole. If the condition

$$R_* = R_{\text{tide}} \simeq 0.5(M_*/M)^{1/3}a_{\text{min}} \gtrsim 0.5(M_*/M)^{1/3}zR_g \quad (1.17)$$

fails, the star is never disrupted: it either orbits the hole ($a > a_{\text{min}}$) without being torn apart by tidal forces, or if $a = a_{\text{min}}$, simply falls into it 'silently', that is, is swallowed whole by it without disruption. Rearranging (1.17) gives the condition for disruption as

$$M \lesssim 5 \times 10^7 M_\odot \rho_*^{-1/2} z^{-3/2}, \quad (1.18)$$

where $\rho_* \simeq 3M_*/4\pi R_*^3 \sim 1/(M_*/M_\odot)^2 \text{ g cm}^{-3}$ for a low-mass main-sequence star. We see that tidal disruption is only possible for normal stars if the SMBH is fairly low, that is, $M \lesssim 10^7 M_\odot$. If disruption occurs, we will see (Section 5.8) that about half of the star's gas accretes on to the hole, while the other half is ejected. Observations have revealed a number of likely TDEs, and the associated SMBH masses are low, agreeing with (1.18). This, and the Soltan relation, tell us that TDEs cannot be the main growth channel for the most massive holes observed at low redshift. But it leaves open the possibility that TDEs may be important in growing SMBH with smaller masses, both in the local Universe, and possibly in the early growth of all SMBH at high redshift. Stars that do not quite fall close enough to be fully disrupted by the SMBH tides, or are the captured component of a binary disrupted by the Hills mechanism (Hills, 1988), are probably the cause

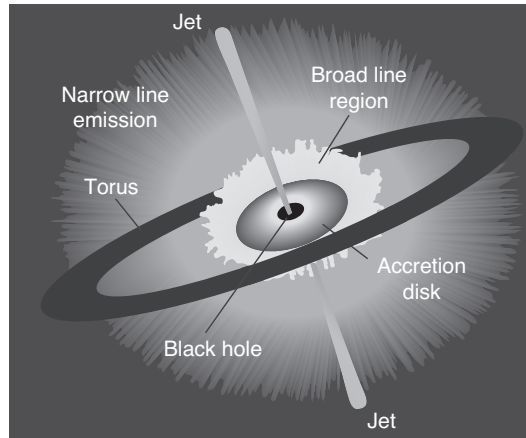


Figure 1.8 Cartoon of the surrounding of an active galactic nucleus (AGN). Note that the central accretion disc and the associated jet may be misaligned from the torus. The entire structure has in general no spatial alignment with the large-scale structure of the host galaxy. Credit: NASA.

of the quasi-periodic eruptions (QPEs) seen in X-ray emission from the nuclei of low-mass galaxies. Here a small star (white dwarfs are favoured by observational selection) is trapped in a highly eccentric orbit about the black hole. At pericentre the star fills its tidal lobe and gives a burst of mass transfer to the black hole, powering the QPE. The emission of gravitational radiation keeps the orbiting star close enough for this to continue (see Section 5.9).

Since our own SMBH is not currently active, it may not be a good guide to what happens when there is significant accretion, as in active galactic nuclei. The brightest of these objects are quasars, where the accreting nucleus outshines the host galaxy so strongly that the host is difficult to observe. Seyfert galaxies also have accreting nuclei, but their lower luminosities allow more observational insight into the galaxy. Seyferts are classified into Type I and Type II through the characteristics of their bright emission lines coming from highly ionized gas. Seyfert I's are luminous in the X-rays and ultraviolet, and show both broad (velocity widths $\Delta v \sim 10^4 \text{ km s}^{-1}$) and narrow ($\Delta v \sim 10^2 \text{ km s}^{-1}$) emission lines, while Seyfert II's show only narrow lines. Since higher velocities must come from gas orbiting close to the SMBH, an appealing picture (Antonucci & Miller, 1985; Antonucci, 1993) is that these two types of galaxy are essentially the same, but viewed from different inclinations to an obscuring torus of matter, probably containing significant amounts of dust, around the region producing the broad lines (see Figure 1.8).




Figure 1.9 The X-ray jet from the galaxy Cen A. Credit: NASA/SAO/R. Kraft et al. (2002).

Another striking property of accreting SMBH whose origin is still not fully settled is that many of them produce jets, usually observed in the radio. These are high-speed (often with velocities close to c), narrowly collimated outflows (see Figures 1.9 and 1.10), and can extend over scales of Mpc. Using the standard rule of thumb that the terminal velocity of an outflow is very close to the escape velocity at the point where it was launched, their relativistic velocities strongly suggest that they come from the close vicinity of the SMBH. Jets are not unique to AGN – every class of accreting objects is able to produce them, with terminal velocities again close to the escape value from the surface of the accretor. So stellar-mass black holes and neutron stars in accreting binaries also make relativistic jets, while protostars forming out of accreting discs of matter make much slower jets (see Figure 1.11). Despite the very different accretors, the similarities between the jets is obvious, suggesting that a very similar mechanism makes jets from all accretors. As we shall see in Chapter 7, there is a general expectation that jets from accreting objects arise from the combination of infalling matter, magnetic fields and rotation. Radio pulsars, which are rapidly spinning magnetic neutron stars, are not accreting, but also make jets, which are aligned with the pulsar spin axis. It may be that the plasma that we know must be present in their magnetospheres flows in ways that mimic accretion flows in all other jet sources.



Figure 1.10 The radio jets from the SMBH in the galaxy M87 and their effects on the surrounding gas, observed by the Very Large Array. The left-hand jet is partly obscured. Credit: NASA.




Figure 1.11 The jets from the protostar HH-47. Despite the very different spatial scales (pc/Mpc) and masses of the accreting objects ($1M_{\odot}/10^9M_{\odot}$), the similarity with the Cen A jet is striking. Credit: STScI & NASA.




Figure 1.12 The observed relations between black hole mass M and (left) host galaxy velocity dispersion σ , and (right) black hole mass and host bulge stellar mass M_b . Credit: Baldassare et al. (2015).

A marked characteristic of AGN is that although the jet appears to propagate close to the axis of the AGN torus, the orientation of this combined jet–torus structure is completely random when compared with the large-scale structure of the host galaxy (e.g. Nagar & Wilson, 1999, Kinney et al., 2000).

An important feature of AGN spectra is that they show that the gas near the AGN is strongly metal-enriched (Shields, 1976; Baldwin & Netzer, 1978; Hamann & Ferland, 1992; Ferland et al., 1996; Dietrich, Appenzeller & Wagner, 1999, 2003a, 2003b; Arav et al., 2007). This presumably means that this gas has been recycled through several generations of massive stars. We shall see how this may have come about in Chapter 7.

1.6 Supermassive Black Holes and Galaxies

The observations summarized in this chapter give a rich picture of the connections between SMBH and their hosts. In their brief active intervals we know that SMBH must accrete gas, and so are likely to show signatures of disc accretion. These include jets, emission lines from gas moving with a range of speeds, and obscuration by a dusty torus. Even non-active SMBH can show themselves in tidal disruption events, and there are many effects related to the central SMBH seen in our own Galaxy, such as hypervelocity stars, the Fermi bubbles, and a mass of detailed information from direct observation of the Galactic Centre.

But one class of observation above all suggests that SMBH are not simply interesting in their own right, but play a determining role in how all galaxies form and evolve. Since the beginning of this century we have known (Ferrarese & Merritt, 2000; Gebhardt et al., 2000) that the black hole mass M is tightly related to two properties of its host bulge (or spheroid) – its total stellar mass M_b , and its stellar velocity dispersion (Figure 1.12).

We can roughly summarize these relations as

$$M \simeq 10^{-3} M_b \quad (1.19)$$

and

$$M \simeq 3 \times 10^8 \sigma_{200}^4 M_\odot, \quad (1.20)$$

where $\sigma_{200} = \sigma / (200 \text{ km s}^{-1})$.

The connection between black hole and spheroid properties implied by both scaling relations appears very surprising at first sight, since we know that the direct gravitational influence of the SMBH is important only inside R_{sph} . With a size of a few parsecs, this is far smaller than the spheroid scale $\sim 1\text{--}10$ kpc, so for most purposes the host spheroid must be completely unaware of the SMBH. We shall try to answer these questions in Chapter 6.