How Many Black Holes in the Milky Way?

Dr Michael McEllin

1 Why do we want to know?

Simple curiosity is a fairly good excuse in pure research, and I became curious after I was asked this question by a Year 12 at a local school, because I simply did not know the answer. However, a little thought made it seemed likely that there could be a fairly large number of undiscovered black holes, and an even larger number of neutron stars and a *very* much larger number of white dwarfs. (These are collectively known as ‘compact stellar remnants’ - the end points of life for most stars.) My subsequent investigation of the literature led me along interesting tracks through some of the most important current issues in modern astrophysics.

The theory of stellar evolution suggests that they have to be there in fairly large numbers, but are they common enough to be detectable (and how would we detect objects light-years away whose size may be measured in kilometres)? Are they important for the processes that go on within the galaxy and that affect its evolution? (It has, for example, recently been realised that collisions between neutron stars in binary systems may be very important for the creation of the heaviest elements—atoms like platinum and gold.) Any reliable observational information on the numbers of compact remnants may help us to distinguish between different theories of galaxy formation, but first we need to estimate whether it is even worth trying to do the experiment. Is it hopeless, or maybe just about possible?

One theory of galaxy formation proposes that the Milky Way (and all other galaxies) condensed from a large gas cloud shortly after the elements created in the Big Bang had had chance to cool and turn into mostly neutral hydrogen and helium. There would be relatively brief initial phase in which most of the gas was turned into stars. We have fairly good reasons for thinking that when stars form out of a large gas cloud contracting under its own gravity it breaks up in a particular way, determined by basic physics, such that there are many more low-mass stars than large stars. This is what theory suggests and also it is what we see *today* in regions where stars are forming. When we look around the galaxy at the older stars we also see that there are indeed many, many more low-mass stars than high-mass stars, so something similar has to have occurred in the past.

In today’s star forming regions the break-up of gas clouds contracting under their own gravity seems to occur according to universal rules giving in every case the same proportions of stars in each mass range. Theoretical analysis, however, suggests that the physics of gas cloud contraction depends to some extent on the small but important fraction of elements heavier than hydrogen and helium. (Elements like carbon, for example, condense into dust grains, and dust grains even in small amounts help to radiate away the energy released by gravitational collapse and help to accelerate the contraction. That is because as the contracting gas heats up it warms the dust grains and they radiate away heat in the infra-red, which is able to escape from the contracting cloud. Otherwise the contracting gas just gets hot and its pressure increases and that stops the contraction. You have to get rid of the heat somehow, and a small amount of dust, it turns out, make a very big difference.)

The primal gas from which the first stars formed, however, consisted of nothing but hydrogen and helium. which are very bad at radiating away heat, and current thinking suggest that star formation under those circumstances *should* favour the more massive stars—and that could lead to many more ancient black holes and neutron stars than we would expect if star formation works according to the present day empirical observations. We would like to know what is really going on, because it is crucial to theories of galaxy formation.

We know from the very well established theory of stellar evolution that high mass stars evolve through their life very much faster than low-mass stars. In particular, stars like the Sun and those more massive than our Sun would finish their lives by ejecting a large part of their mass into to the interstellar medium—for example, in supernova explosions or, less dramatically (but much more frequently) when they throw off their outer layers to form *planetary nebulae*. This gas, now enriched in heavier elements because of nuclear fusion with the stars, would then provide material for a second generation of stars. Subsequently, the heaviest stars from this generation would evolve quickly and recycle most of *their* mass to the interstellar medium, and so on, until the present day. So, the material in a star such as the Sun (and everything the Earth is made of) has perhaps been through the centre of another, now dead star perhaps two or three times.

Many details of this theory have to be substantially correct, but in the big picture of galactic evolution there are difficulties. In particular, stars just bit less massive than the Sun have lifetimes longer than the age of the Milky Way, so even those formed from the primal material are still around (and they have been observed, but they are not easy to find, even though they are very common, because they have very low luminosities). In fact any gas that ends up in such low-mass stars is effectively locked up and unavailable for recycling for perhaps hundreds of billions of years—effectively forever! Each generation of star formation, therefore, is taking gas out of the recycling loop, and we can ask whether the amount of gas we see around the galaxy today is consistent with what we might expect to see 12 or 13 billion years after the first stars formed. We shall see!

Alternative theories propose that star formation did *not* happen in one big burst, and that maybe, even, the galaxy is being continuously supplied with new material as gas falls in towards us from the intergalactic medium. (It turns out that there is a very large amount of gas between the galaxies, though it has an extremely low density.) In this picture, the centre of a galaxy like the Milky Way forms first and fairly quickly, but the outer regions (which have the spiral arms and now have most of the gas) grow later on, by mergers with smaller galaxies, or from large gas clouds falling into the galaxy from outside. (This is known as the 'inside-out' theory.)

There are several ways we might try to distinguish these between these theories. Ideally, we would like to look back across space to the earliest times after the Big Bank and watch galaxies forming through our telescopes. We cannot quite do that yet—though we are getting close to having the required technology!

Part of the problem is that these early galaxies are so far away that most of the light they emit is red-shifted to wavelengths that do not easily penetrate the Earth's atmosphere. Even the Hubble Space Telescope is not able to see sufficiently fat into the infra-rad. The *James Webb Space Telescope*—a replacement for Hubble—will attack this problem, because its sensors are design to operation in the far infra-red. Radio astronomers are also designing the *Square Kilometre Array* which is an enormous, World-wide telescope that will be able to map the primal gas from which the galaxies are forming.

The issue of galaxy formation is, in fact, probably the biggest research question in modern astronomy and is driving a large fraction of current research programmes.

In the mean time. we can think about the different models of galaxy formation and use them to predict the type of stars we should see today, the amount of gas we might see and also the numbers of stellar remnants that we should see around us today. We may be able to rule out some of the hypotheses just because they predict too many (or too few) of certain types of highly evolved object.

The whole question is complicated by observations of some galaxies that apparently have little remaining gas and therefore no current star formation, while others, such as our own Milky Way, have lots of gas and lots of current star formation. Why are they so different? We also have strong suspicions that some of the larger galaxies may have formed by 'eating-up' smaller galaxies, and that make un-picking history very complicated. Furthermore, many (probably all) larger galaxies seem to have super-massive black holes (SMBHs) at their centres. It is becoming clear that the giant super-energetic outbursts associated with SMBHs that appear to us as quasars and related phenomena can have a strong influence on star formation throughout the galaxy. One of the ways we can recognise when this might be happening is understanding the way the galaxy should evolve without the influence of the SMBH.

Hence, we should restate the important questions introduced earlier:

- Can we calculate the number of compact stellar remnants (neutron stars and black holes) we expect to see?
- Could we detect these remnants in sufficient numbers to rule out some models of

galaxy evolution?

It turns out that it is not that complicated to make some of the predictions—at least to a 'first approximation'[1].

Let us be honest: we are treading a well explored field, in which many researchers are using more powerful tools than we have at our disposal (including massive arrays of computers that can be used simulate in detail the evolution of whole populations of stars). What I hope we can achieve, however, in an understanding of the way a researcher takes first steps on the progression from an interesting but difficult question towards understanding whether it may be possible to make progress by following a certain line of investigation, using progressively more complicated tools. We should be able to see that these first steps can often be taken with no more than A-level knowledge.

1.1 Recommended Reading

There is a great deal of relevant material easily found on-line. Two books, however, provide a good introduction to the astrophysics of stars and galaxies at an introductory university standard: See:

- An Introduction to the Sun and the Stars (Green & Jones 2015).
- An Introduction to Galaxies and Cosmology (Jones, Lambourne & Serjeant 2015).

These are 'Level 2' Open University text books, but are well produced and excellently targeted at students who have not entered university education via the conventional progression from A levels. They assume intelligence and a willingness to work at understanding the contents, but little prior knowledge, use only a modest amount of fairly straightforward mathematics and ought to be entirely accessible in large part to a good A Level physics student. Nevertheless, they cover most of the more important ideas that are the basis of modern astrophysics in a way that, in my opinion, coveys good understanding. (It has recently been possible to pick up earlier editions, from about 2004, of these texts on-line secondhand at very modest prices often only £3 or £4. These earlier editions are still entirely suitable for the project work discussed here.)

---

[1] A 'first approximation' means a calculation that is imprecise but which is better than no calculation. It is also capable of being improved if the first result looks interesting. For example, we know that on 27 March 2011 (the time of the last census) the population of England and Wales was 56,075,912. We do not know the exact figure now, but by looking at the typical rates of births and deaths and migration we could get a value that good enough to use for government planning purposes (e.g. perhaps in order to work out how many school places to we might need in the next ten years). It is not precise, but it may be good enough, and it may tell us that we need to look more closely at the available evidence and do a more detailed calculation to get a 'second approximation'. Physicists have developed the process of successive approximations to a very fine art and used it as the basis of some of the most precise predictions ever performed in science. See (Mahajan 2008) and (Mahajan 2014) for good advice in getting answers by approximation.

I would also recommend the following two book about using mathematics, which are aimed at students about to enter university or in their initial undergraduate year. Copyright-legal PDF downloads are available free from the publisher and other websites. (See the PDF link on the publisher's web pages for these books, as given in the citation.)

- Street Fighting Mathematics (Mahajan 2008).
- The Art of Insight in Sciences and Engineering (Mahajan 2014)

These books promote a way of thinking about problems that is very familiar to professional physicists and mathematicians. (I certainly remember being taught this approach when I studied physics at Cambridge, and it is no surprise to find that the author of these books once taught on that course.) They are all about how to understand the real nature of a problem. When exact analytical solutions are impossible, one might do this by inventing similar problems that are easier to solve. You may not get the exact answer to the original problem, but your answer will probably have some of the characteristics of the real solution, and you will understand more about what is important in getting to a useful answer. Perhaps you can use several approximations to constrain an area of the solution space where it must nevertheless lie. They are all about getting *useful* answers rather than the formally 'right' answer.

2 Predicting the Number of Compact Stellar Remnants

I give below an outline of some of the astrophysics involved in this project, but you will certainly have to research each of the topics below in more detail, using your own Internet searches, and perhaps reading the recommended books.

2.1 End Points of Stellar Evolution

White dwarfs, neutron stars and stellar mass black holes are the end-points of stellar evolution. In broad terms,

- Stars with masses up to $8M_\odot$ (that is 8 times the mass of the Sun[2]) evolve into red giants, and eventually throw off their outer envelopes leaving the behind the core in which all the hydrogen has been burnt into helium—this is a white dwarf.
- Between $8M_\odot$ and about $20M_\odot$ stars reach a point where their cores will collapse into neutron stars.
- Above $20M_\odot$ the core is too big to be stable as a neutron star and black holes are possible.

---

[2] The traditional astrological symbol $\odot$ is used to label quantities that refer to our Sun, so $M_\odot$ is conventionally used to denote the mass of the Sun and $L_\odot$ denotes the luminosity of the Sun.

This is almost certainly an over simplification: we think that *some* stars of mass $> M_\odot$ might actually end up as neutron stars. There is also the possibility that some very massive stars blow themselves completely apart without forming black holes. We also know that about 10% of stars are in close binary systems and that can effect stellar evolution in complicated ways that may change the end-points. We just need to keep our assumptions in mind, and be ready to adjust them if we find contradictions with the evidence. Part of the project will working out how sensitive the answer is to these assumption by repeating the calculation using different assumptions about which stars turn into black holes. You will need to research stellar evolution and learn about some of the complications.

In conclusion, we think we know (in a broad sense) which stars are likely to turn into black holes. What we need to know now is the rate at which stars in each mass range have been and are being formed, and how quickly they evolve towards death, so we can predict the rate at which black holes are created. We will have to examine different scenarios for galaxy evolution, ranging from most of the galaxy's stars forming early in the Milky Way lifetime, to a linear rate of star creation, fuelled by new gas supplies, and everything in between. We should then be able to estimate the accumulated number of stellar remnants.

2.2 The Stellar Mass-Luminosity-Lifetime Relationship

The luminosity of a star for most of its lifetime is almost completely determined by its initial mass. Some fairly straightforward theory, for example (Longair 2006), suggests that the luminosity, L, must be approximately proportional to the cube of the star's mass. This theory involves some big simplifications, and though we can see that it has to be more-or-less right, we do not expect exact agreement. In fact observations and sophisticated computer calculations suggest that the relationship between mass and luminosity should be $L \propto M^\alpha$ with $\alpha \approx 3.5$ over the mass range 0.1..50 $M_\odot$ but with a slightly different power-law index at higher and lower end of the mass range. On the whole, our simple theory is surprisingly good.

The luminosity relationship has implications, because massive stars with high luminosity are burning their fuel supply at a rate roughly proportional to $M^{3.5}$ but their total fuel supply is only proportional to their mass. Hence, we should suspect that the lifetime of a star depends on its initial mass as something like: $t_\star \propto M_\star / M_\star^{3.5} = M_\star^{-2.5}$. So, massive stars live fast and die young. While stars like the Sun may live for 10 billion years, stars that are 100 times as massive might only last a million years. (Once again really detailed computer calculations produce result that are a little different from this approximation, but we are not that far wrong.)

Think about this: it means that if we take a bunch of stars that all formed at about the same time (say, for example, a globular cluster—see Figure 1 on page 7 which shows and image of the cluster Messier 80) after a billion years of so, all the more massive stars

Figure 1: Globular Cluster Messier 80. (Credit: NASA Hubble Archive).

will have reached the end of their life and turned into neutron stars or black holes, while the less massive stars will still be around. Figure 2 on page 8 shows a colour-magnitude diagram for Messier 80, which I plotted from Hubble Space Telescope observations, where the 'colour' variable on the bottom of the graph more-or-less corresponds to stellar mass with low mass, low luminosity, red stars on the right and high mass, luminous and blue stars on the left. (The Sun has a colour of 1 on this scale.) The vertical scale corresponds to luminosity.

Most of the stars lie on a 'main sequence' in a diagonal band from top left to bottom right, and this is what we would expect for stars that have not yet consumed their fuel supplies. In the upper half of the plot, however, we can see a bend in the curve, and in fact there are no stars at all in the top left. The brightest stars have already evolved through their full life time and there are none left behind. Those stars above the bend are 'red giants'—they have already consumed much of their fuel and are in a late stage of stellar evolution, still very luminous but with large fairly cool outer envelopes. (You can see the sprinkling of reddish bright stars in Figure 1.) Eventually they will mostly throw off their extended outer layers, turning into 'planetary nebulae' and their cores will drop down to the bottom left of the diagram as white dwarfs. (These are so dim that they are often difficult to see in globular clusters.)

When we look at stars in the main part of the Milky Way, all around us, we *do* see the

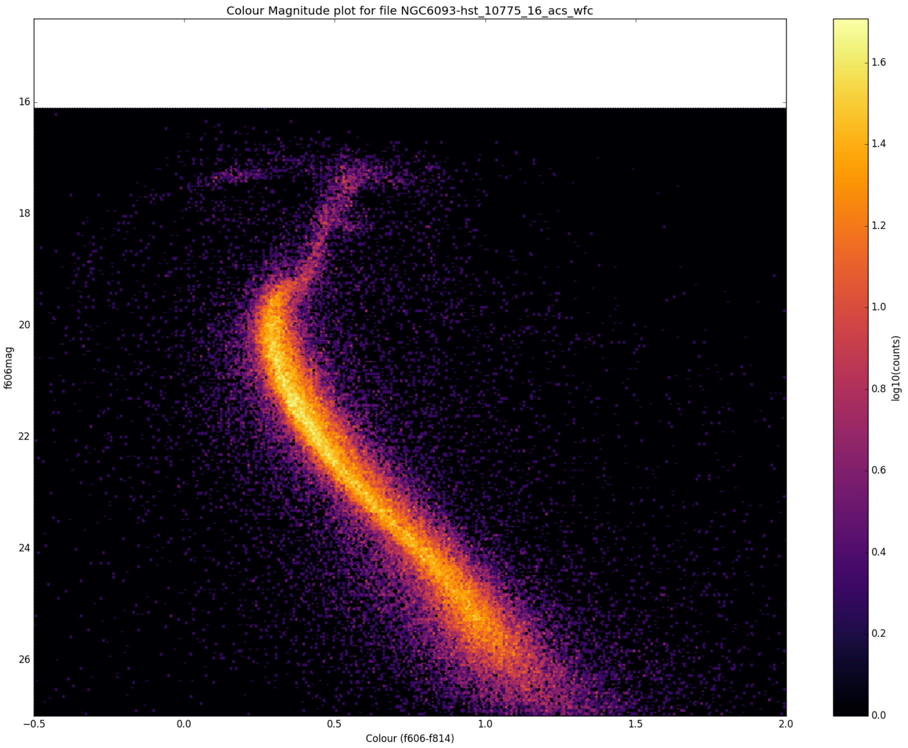

Figure 2: Colour Magnitude Diagram for Messier 80.

main sequence extending up to the top left of the diagram. (See Figure 3 on page 9.) This means that stars must still be forming in the plane of the Milky Way, because the hot blue luminous stars in the top left have very short lifetimes.

The Mass-Luminosity-Lifetime relationship also means that stars a little less massive than the Sun (with its 10 billion year predicted life) could easily have survived from the first wave of star formation after the Big Bang, and in fact will carry on much as they are now for many, many billions of years into the future.

We are almost there: just one more piece to the jigsaw and we have enough to predict the number of black holes in our galaxy. We need to know what proportion of stars that are born are big enough to become black holes.

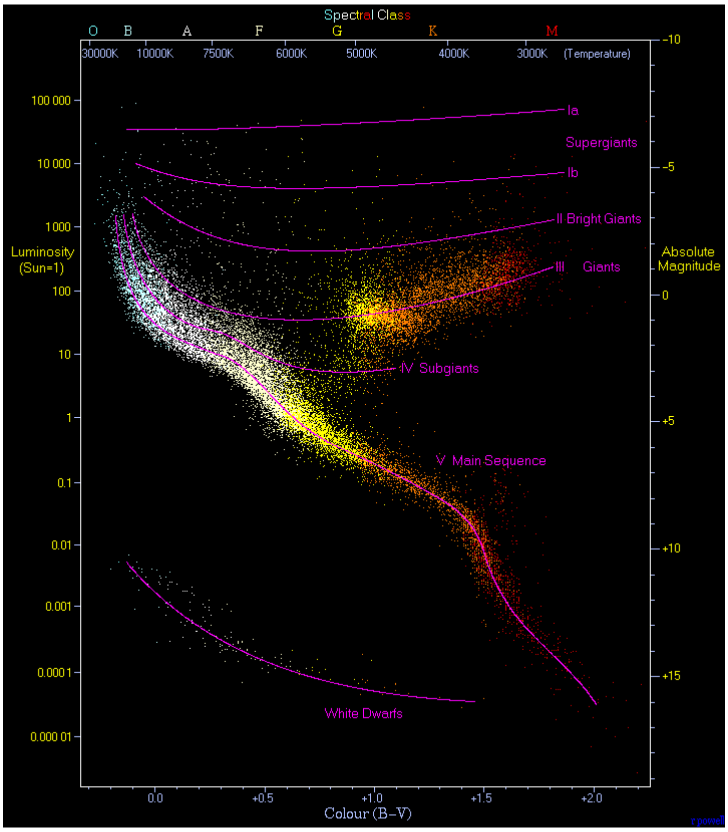

Figure 3: Colour Magnitude Diagram for Milky Way Stars

2.3 The Stellar Initial Mass Function

The final key astronomical concept that we need is the stellar *Initial Mass Function* (IMF) which describes the relative proportion of stars of various masses that form from a contracting cloud of gas. The simplest (and first published) approximation is due to Salpeter (Salpeter 1955) who proposed that the IMF, which is defined to be an expression giving the number of stars forming in a given mass range, should be described by the equation:

$$\xi(m)\Delta m = \xi_0 \left(\frac{m}{M_\odot}\right)^{-\beta} \left(\frac{\Delta m}{M_\odot}\right) \tag{1}$$

We need to pick apart the meaning of this equation. Firstly, we choose to measure mass on a scale that defines the mass of our Sun to be 1. ($M_\odot$ symbol means the mass of our Sun.) Hence, the ratio $m/M_\odot$ (or $\Delta m/M_\odot$) is merely saying that we are measuring all masses mass in terms of Solar masses. So,

- $\xi(m)\Delta m$ is the number of stars forming with a mass between m and $m + \Delta m$.

- ξ_0 is a constant that will depend on the things like the local density of the gas from which stars form. We do not know what it is in general, but it can often be worked out for particular stellar clusters where a whole bunch of stars has recently condensed from the interstellar medium.
- β also has to be derived by observing star forming clusters.

Salpeter analysis of observations of this type said showed that β seemed to have the same value of about 2.35 wherever he looked, though ξ_0 would depend on things like the local density of gas. So, β is telling us something about the physics of the way gas clouds break up when they are contracting under gravity—something that does not depend on things like the local density of the gas[3].

The IMF formula therefore says that there are many more low-mass stars than high-mass stars.

Using this equation we can easily compare the predicted formation rate of with the mass of the Sun to that for stars ten times more massive because in this case, since we are looking at a ratio, that actual value of ξ_0 cancels out. We would then estimate that the more massive stars are $(10/1)^{-2.35} = 0.004$ times less common, while stars 1/10 of the Sun's mass would be 223 times *more* common.

Now, Salpeter's IMF was derived empirically and was fitted to observations over a limited range of stellar masses: he would not have seen very many massive stars (say, over 50 times the mass of the Sun, they are very rare) nor would he have been able to see many low-mass stars, because they are very dim and hard to find with the telescopes of the early 1950s (and we have to extrapolate over these 'observational selection effects' to account for the ones you miss). So, at best we can only claim that the equation works for a certain range of stellar masses, say from $0.5M_\odot$ to $50M_\odot$.

In fact, we know that this equation *has to be wrong* for an unrestricted range of stellar masses, because we should be able to integrate the formula to get the total number of stars within a defined range. The upper mass limit of the integral does not matter, because the number of stars goes rapidly to zero as we go to higher masses. However, the number of stars at the low-mass end tends to infinity as the mass goes down, and we cannot have an infinite number of stars. We can immediately deduce that the IMF has to have a different form for lower mass stars.

Later researchers have looked very carefully at the low-mass end (generally thought of

---

[3] The form of the IMF is a *power law.* These turn up all over physics (for example, in the energy distribution of cosmic rays) and at first one might wonder why so many very different phenomena seem to be describable by very similar mathematical formulae. In fact, this is one of those situations where intuition lets us down: we have to do the maths properly. Power laws, it turns out, tend to occur when the probability of something happening depends on a multiplication of individual randomly distributed effects. Some university-level statistical mathematics shows that power-laws tend to fall out of these situations providing you multiply together enough random factors, even when the individual probability distributions do not, in fact, look anything like power laws. Mathematical magic—non-intuitive but true.

as $m < 0.5M_\odot$ based on observations with much more sensitive modern telescopes that count many more low luminosity stars. There are, indeed, a whole range of proposals some using fairly complex mathematical forms. One which seems to work not too badly (Kroupa 2001) assumes that we can divide the mass range into three parts where we have three different values of β, where the value $\beta = 2.3$ can be applied for stellar masses $m > 0.5M_\odot$, $\beta = 1.3$ applies for $0.08M_\odot < m < 0.5M_\odot$ and $\beta = 0.3$ applies for $m < 0.0.08M_\odot$. This function converges at both ends and is relatively easy to manipulate mathematically. This will be good enough for our purposes, even though other, more complex and harder to manipulation representations may be a little more accurate.

Rather surprisingly, later work has confirmed what at least approximately seems to be a 'universal power law': star formation seems to follow the same rules wherever we look. Normally, when scientists come up with a simple theory, we find that nature introduces complications in different circumstances. Here, we have the opposite: all the theorists say that we should expect different rules to apply when stars form in different circumstances, but we see the *same* behaviour—and we do not really know why nature works like this! But remember the 'approximately'! The power-law index values are empirically fitted values and $\beta = 2.3$ really means we think that $2.2 < \beta < 2.4$. That apparently modest uncertainty in the power-law index may well make quite a big difference in the predicted number of the high mass stars that turn into black holes, because they are at the limit of extrapolation of this correlation.

Remember that we suggested earlier that the first generation of stars, forming from gas with no heavier elements, may well have favoured the production of somewhat more massive stars (that is, β may have had a small value). This first generation of stars could be responsible for the majority of black holes around today. You will therefore need to investigate the sensitivity to this parameter by repeating the calculations with a range of values. (Hence, you might like to think about how you would set up a spreadsheet to do the calculations in a way that allows you to tweak the controlling parameters up and down easily. This is a very common strategy in mathematical modelling.)

In fact, we will not need to work very much with the different power-law forms for $m < 0.5M_\odot$ because all the stars that form black holes are much heavier than this. We just need to know the proportion of the initial gas cloud mass that end up in these low-mass stars in each generation. We then work with the expression for the higher-mass stars normalised to the proportion with $m > 0.5M_\odot$.

Your first mathematical task, therefore, is to integrate the Kroupa form through the full range over which it applies (that is over all three segements) to obtain a normalisation (the integral corresponds to 100% of stars forming) which we can then use to derived the proportion of stars and the total stellar mass that end up in stars less massive than $0.5M_\odot$. You will, of course, have to make sure that the different power law forms in each of the three mass segments matches where they meet, by deriving an appropriate value of ξ_0 for each β.

3 Calculating the Number of Black Holes

3.1 First Approximation

3.1.1 Model 1

Assume that all the stars in the Milky Way formed at the same time.

- Work out the proportion of the stars that are greater than the $20M\odot$ mass limit for black-hole formation.
- Work out the proportion of the stars that are greater than $8M\odot$ and less than $20M\odot$ that have reach the end of their life and become neutron stars..

Please not that it is very useful at this stage to avoid introducing specific numerical values for the power indices on the luminosity relation and the initial mass function. Just handle the mathematics using symbols such as α and β standing for these values. When you have deduced your final results in fully symbolic form you will them be able to place the formulas into a spreadsheet and investigate when making slight variations in the values of α and β has much effect on the numbers of compact objects. If they have a relatively big effect we will need to know the values of α and β quite accurately.

Is this model likely to produce an over or under-estimate of the total number of black holes in the galaxy?

3.1.2 Model 2

Assume that the stars in the Milky Way have formed at a steady rate over time.

- Work out the proportion of the stars that are greater than the $20M\odot$ mass limit for black-hole formation that have reach the end of their lifetime at the present time.
- Estimate the number of neutron stars in the Milky way.
- Estimate the number of white dwarfs, that is stars with mass $m < 8M\odot$ but whose lifetime is shorter than the age of the galaxy. (The limit will be somewhere a bit less than the mass of the Sun.) Why do astronomers actually find it difficult to pick out white dwarf stars?

You could produce a fully analytic answer to this, using just A-level maths (I am not saying it is easy!) if you forgot about the recycling of gas ejected from exploding stars into new stars.

It is probably easier, however, to think about dividing the time since the formation of the galaxy into, say five or ten periods where all the stars that are going to be born in that period get born at the beginning. We then do our original calculation for each period (probably best with a spreadsheet. The answer will not be perfect, but it will probably be good enough.

You might even then consider how you would add the gas ejected by stars exploding in an earlier period to that available gas for star formation in a later period.

3.1.3 Model 3

Assume that the Milky Way's central bulge and galactic halo formed very early on, and that the disk has grown at a constant rate from then until now. Repeat the calculations above to estimate the number of black holes and neutron stars

3.1.4 Model 4

Now take account of the 10% of stars that form in binary systems, and which may have a different route to becoming black holes and neutron stars.

Why do binary systems have a particular importance in astrophysics—particularly for the fate of compact stellar remnants?

3.2 Second Approximation

Create a spread sheet with columns representing time, in intervals of, say, 1 billion years, and rows covering different mass ranges of stars. Implement the calculations performed previously but allowing the rate of star formation to change over time in an arbitrary way.

This is actually a substantial piece of work—but possible for the very committed students.

4 Detection of Neutron Stars and Black Holes

Research methods of detecting black holes and neutron stars.

In general, there are only three feasible methods of detecting black holes:

- A merger of two black holes can generate gravitational waves. These have recently been detected by the LIGO gravitational detector. (Almost certainly, the two black holes came into existence in from a binary star pair.)
- It is theoretically possible that a black hole moving across the line of sight to another star could distort the background light to an extent that would show up

in, say, observations of the GAIA satellite. What is the chance that this might happen? Have astronomers actually detected such ‘micro-lensing’ events?

- Energy can be released in an observable way when material falls into a black hole. (In general a swirling ‘accretion disk’ forms, which gets very hot and radiates in the X-ray.) Where might we see this phenomenon? How common is it? Have we seen it?

The methods of detecting neutron stars are almost the same—except that some neutron stars can also become pulsars.

5 Bibliography

An electronic version of this document is available via:

http://mcellin.me.uk/artfulcomputing/index.php/resources/crest-project-resources/astrophysics-resources.

You will note that in the main text and this list of citations, I have tried to follow the standard procedure that academic authors use for referencing other people’s work. You will be expected to follow something like this style for your work.

You are not expected to try to seek out all the material cited here. The professional research papers in particular are written in a very compact style that assumed a fairly high level of background astrophysics knowledge. Nevertheless, for Gold CREST projects, assessors like to see some attempt to find ‘primary’ research sources—that is, the original descriptions of research results, rather than the second or third hand digests you find in text-books.

The general rule, which you will certainly have to follow with a good deal of respect when preparing assignments at university, is that when you make a claim to support the argument in your written material that something is known, or some recognised expert has stated an opinion, you must point to the place where you got your information, so a reader can, if he or she wishes, check that it really is as you have said.

I do not, in general, follow up all the citations in the research papers I read. I do track them down from time to time, perhaps because I am intrigued by something I did not know, and I want to know more. Sometimes I do it because I suspect the author is making an invalid claim (they rarely do this deliberately: more frequently they have misunderstood some subtlety in the original research).

N.B. Had I been preparing this note as an academic document intended for peer-review and publication I would, in fact, have peppered it with about three or four times as many citations. In fact, I mainly provided the citations to give you an example of how it is done. See if you can spot the places where I make a claim of knowledge and where I have *not* provided a citation.

References

Green, S. F. & Jones, M. H. (2015), *An Introduction to the Sun and the Stars*, Cambridge University Press. (Open University Level 2 textbook. The 2004 edition also covers the necessary material.).

Jones, M. H., Lambourne, R. J. A. & Serjeant, S. (2015), *An Introduction to Galaxies and Cosmology*, Cambridge University Press. (Open University Level 2 textbook. The 2004 edition also covers the necessary material.).

Kroupa, P. (2001), 'On the variation of the initial mass function', *Monthly Notices of the Royal Astronomical Society* **322**(2), 231–246.
URL: *http://dx.doi.org/10.1046/j.1365-8711.2001.04022.x*

Longair, M. S. (2006), *The cosmic century: a history of astrophysics and cosmology*, Cambridge University Press. An excelent history of the field for those who want to know the details of how the many paradigm-changing astrophysical discoveries of the 20th Century really came about. This is a work of academic scholarship, and is not intended to be a science popularisation book. It is aimed at readers who already have a scientific background—preferably with some mathematics if they want to follow the explanations of the science in the technical appendices. Nevertheless, it tells a facinating and page-turning story with a good deal of insight—which is not all that surprising because the author was personally involved in some of the discoveries, and had worked with, or talked to, many of the other important players in the field. Some of the people with whom I studied—including Malcolm himself—figure largely, so I admit to a very biased view.

Mahajan, S. (2008), *Street-fighting mathematics*, MIT Opencourseware. Free PDF downloads are available from the publisher - MIT Press. Follow the URL and look for the PDF link on that page. I would strongly recomment this book to any student who intends to study physics or maths at university.
URL: *https://mitpress.mit.edu/books/street-fighting-mathematics*

Mahajan, S. (2014), *The art of insight in science and engineering: mastering complexity*, MIT Press. Free PDF download available from the publisher - MIT Press. Follow the URL and look for the PDF link on that page. This is a good insight into the way professional physicists and engineers actually start thinking about difficult problems. Any student intending to study physics or engineering would certainly benefit from reading this book.
URL: *https://mitpress.mit.edu/books/art-insight-science-and-engineering*

Salpeter, E. E. (1955), 'The Luminosity Function and Stellar Evolution.', *The Astrophysical Journal* p. 161.
URL: *http://adsabs.harvard.edu/abs/1955ApJ...121..161S*