Department of Physics

2022-2023 MSci research projects

Picture Winners

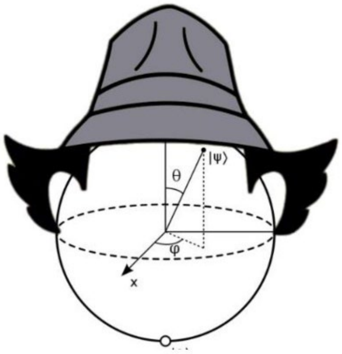

Inspector Q-Bit
by Joseph Keiren (p.22)

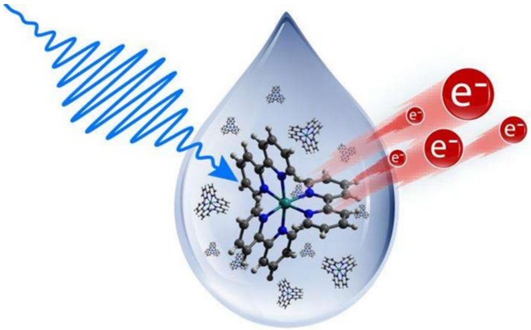

Forming thin liquid sheet jets
by Maria Filippova (p.17)

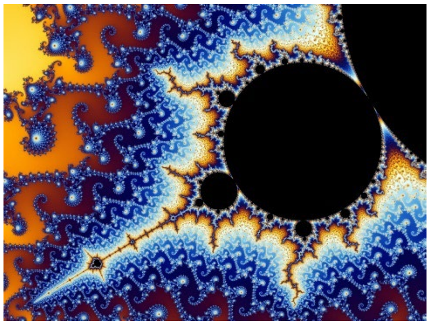

The Simplicity of Complexity
by Monte Ren (p.18)

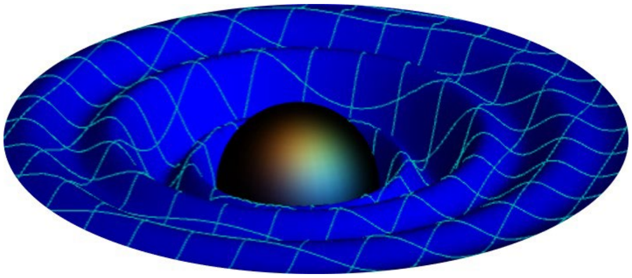

Gravitational Waves and Quantum Gravity
by Miguel Valera (p.26)

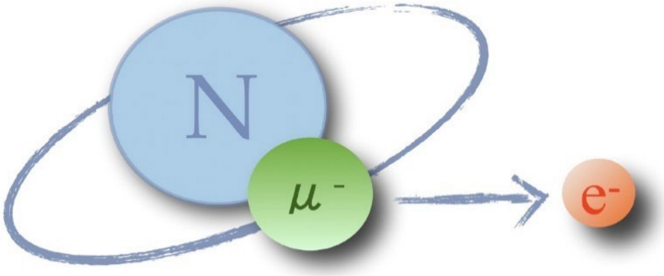

Muon to Electron Conversion
by Sam Zhou (p.12)

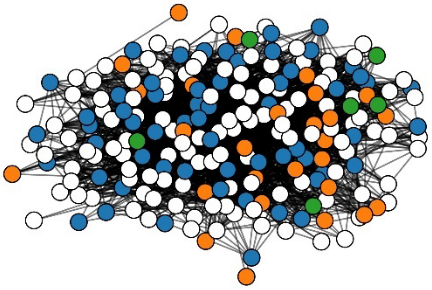

How to Spread a Rumor
by Thomas Cowperthwaite (p.21)

All students in the physics programme at Imperial College London take part in a research project. A selection of lay person summaries written by 2022-2023 fourth year MSci students is presented below and provide good examples of what to expect when doing a project in the physics department and how to engage with state-of-the art research.

The Physics Department Outreach Team offers a range of public lectures and school talks for all ages and tastes, as well as a work experience programme.

Please visit our web page on
https://www.imperial.ac.uk/physics/engage-with-us/outreach/

Selection of 2022-2023 Physics MSci
Lay Person summaries

Planets & Atmosphere

Fundamental Particles

Physics, Medicine, Surfaces and complex systems

Physics and Information

Quantum Gravity and Black Holes

Planets & Atmosphere

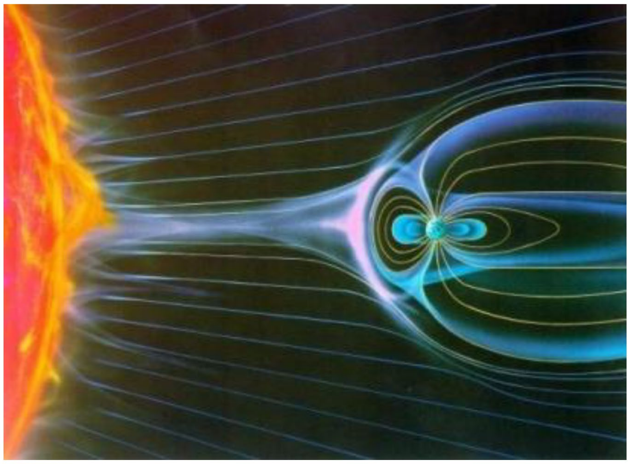

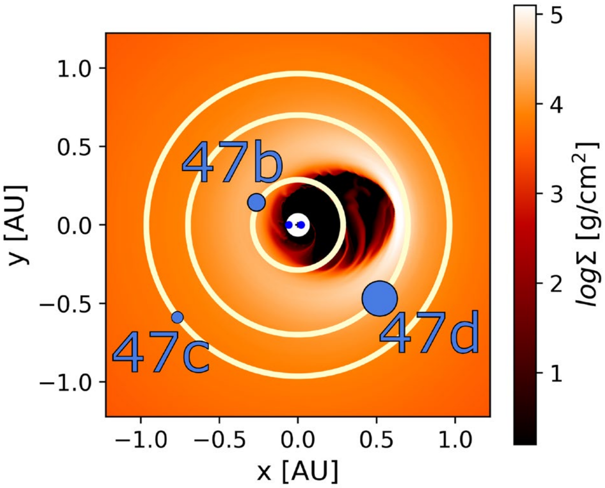

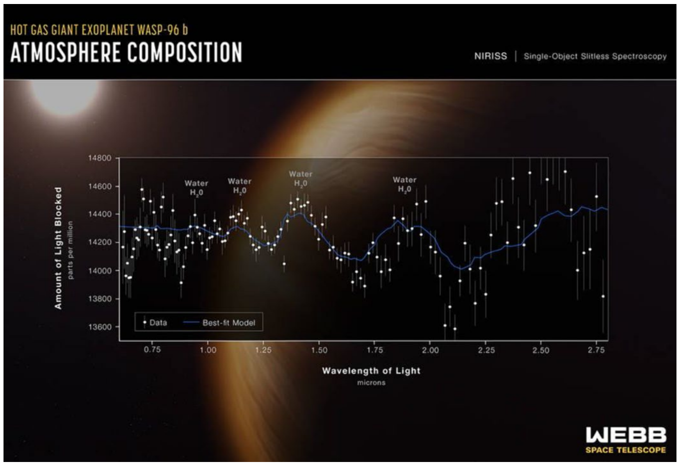

Nature's Greatest Glow-Up:

Modelling Electron Precipitation in the Upper Atmosphere of Jupiter

Jupiter's complex atmosphere has been a source of intrigue and inspiration ever since its infamous red spot was first glimpsed in the 17th century. Another one of the planet's great mysteries is the glowing *aurora* seen at its polar regions. The aurora on Jupiter are in fact much like those on Earth- a spectacular light display seen at locations close to the geomagnetic poles, especially during times of heightened solar activity. The southern lights are often only seen by either scientists, or penguins (who do not care very much for the underlying Physics). However, Jupiter's aurora remain a permanent fixture at its poles, with a power input three orders of magnitude greater than the terrestrial "polar lights". Jupiter's aurora has been imaged in various electromagnetic ranges, most notably by the Hubble Space Telescope (HST), and in stunning quality by the long-awaited James Webb Space Telescope (JWST).

Figure 1: Images of the aurora on both Earth and Jupiter. From left to right **1)** Emperor penguins (who may or may not enjoy Physics) looking at the southern lights, or "aurora australis". **2)** Image taken by the HST in the Ultraviolet range with the aurora in the northern hemisphere clearly visible. **3)** Near-infrared image as taken by the JWST in 2022 with the aurora clearly visible in both hemispheres.

What drives this mysterious aurora on Jupiter? Particles outgassed by volcanoes on the moon of Io become ionised are subsequently picked up by Jupiter's fast-rotating magnetic field. The electrons channelled along magnetic field lines and eventually impinge on the upper atmosphere are referred to as the "electron precipitation".

As these electrons permeate through the atmosphere, they lose energy via inelastic collisions with neutral atmospheric neutral molecules (mostly atomic hydrogen, molecular hydrogen, or methane). This is either by ionization: knocking an electron from the orbital of a neutral species, or by excitation: collisions promote the electron of a neutral to a higher energy level. When the electron relaxes to its ground state, it produces a photon, thus the light which the aurora is comprised of.

There are limitations in taking measurements of Jupiter's polar regions, therefore making understanding processes here challenging. We must instead resort to computational methods. The 90s Voyager and Pioneer missions to Jupiter provided a wealth of atmospheric data; in fact, the Voyager 2 spacecraft confirmed the existence of the electron precipitation. Although computational power has increased enormously since then, energy loss models which are 'multi-stream', simulating incoming electron beams over a range of angles over the planetary hemisphere, have remained sparse.

In this research, we present a multi-stream approach where we consider the extra-atmospheric electron beams in a one-dimensional column of atmosphere. The underpinning mathematics is based on treating an electron beam as a single bulk material, which loses energy continuously through the atmosphere. We also incorporate "*secondary electrons*": electrons produced from ionization reactions with the neutral atmosphere which subsequently cause collisions themselves. Other types of approach to this problem have been "kinetic" approaches. These can simulate individual collisions, but our model has the advantage of requiring less computational power. We compare our model to this approach, identifying cases in which our "bulk" model differs.

Successful modelling would ultimately solve the long-standing "energy crisis" (not to be confused with the one caused by Liz Truss). Theoretical predictions of the temperature of Jupiter's atmosphere by solar heating have not matched observations. The key to reconciling these are heating effects from the energy degradation of the incoming beams. Our research coincides with the launch of the ESA Jupiter Icy Moons Explorer (JUICE) which is set to be launched shortly (hopefully this April)! This will continue the legacy of the NASA Juno mission, which has provided substantial data from the polar regions, with its de-orbit set for 2025. We anticipate this research can be used as a pre-liminary tool for the large volumes of data available in the near future.

Assunta Sophia Felice
Supervisor: Prof Marina Galand

Space and Atmospheric Physics Group

Simulation says…

Earth's magnetic field is alive!

The Sun is an active ball of plasma (a magnetised gas). It constantly ejects plasma that reaches Earth in the form of solar wind, which drags field lines with it. How can it be safe to live on Earth then? The Earth has its own magnetic barrier: the magnetosphere (Image 1). Because magnetic field lines don't break or mix, the magnetosphere acts like a shield from the solar wind. But this shield is not perfect as magnetic reconnection can make field lines mix. Parallel field lines pointing in opposite directions can "break" and "re-connect", rearranging their shape as in Image 2. This is how solar wind field lines and magnetosphere field lines mix. The solar wind conditions that can impact Earth are called space weather, which include radiation exposure to satellites, and induced ground currents that could damage power networks, among others. Understanding the magnetosphere is key to minimise the effect of space weather on Earth.

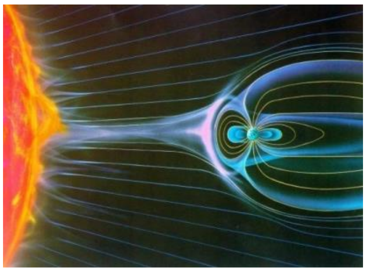

Image 1: Solar wind and Earth's magnetosphere interaction. Source: https://sci.esa.int/s/8r0oDLw

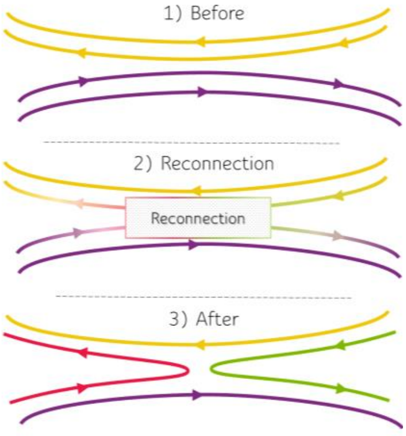

Image 2: Cartoon showing magnetic reconnection.

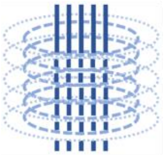

Image 3: Cartoon of a flux rope.

Image 4: Cartoon of two interlinked flux tubes.

Reconnection happens regularly in the dayside (the magnetosphere's closest part to the Sun) and in the magnetotail (the opposite side). Dayside reconnection allows solar wind to mix with the magnetosphere. It is then transported around Earth to the magnetotail, where it piles up. There, reconnection happens again, and solar wind plasma reaches Earth through magnetic field lines connected to the poles, which creates northern (and southern) lights!

Image 1 shows how reconnection occurs in 2D, but the magnetotail is in 3D. How does it look then? Is there a continuous reconnection line, or are there many individual events? It is not clear yet when, where and how often reconnection takes place in the magnetotail. What controls these factors? We try to answer these questions by simulating the magnetosphere using Gorgon, a magnetohydrodynamic (magnetic+fluids) code developed at Imperial.

We found reconnection to be highly scattered in the magnetotail, happening almost constantly and at different distances from Earth. In 3D, two field lines that are not perfectly parallel can reconnection in the form of component reconnection. This twists the fields after reconnection and allows different magnetic structures to form. The most typical ones are flux ropes. They are formed by a central group of field lines surrounded by a spiral of more field lines as in Image 3. Spacecrafts constantly observe flux ropes, and we have also found them in our simulation. Seeing flux ropes implies there is component reconnection in our simulation.

Most surprisingly, we observed a magnetic structure that has only been detected in the dayside and solar wind before, never in the magnetotail. These are interlinked flux tubes (IFTs). As seen in Image 4, they are formed by two hooked groups of field lines. Their formation is only possible when considering a complete 3D description of the magnetotail. These results can motivate further future missions. If spacecrafts detected IFTs in the magnetotail, we would have direct evidence that the magnetotail is governed by scattered, component reconnection.

During the simulation, we encountered an event where a huge amount of energy was suddenly released. This accelerated the flow of the magnetotail in the downtail direction (away from Earth). Even under this tremendous acceleration, IFTs stopped moving. How was this possible? When magnetic fields bend, they experience a tension force which straightens them. Both flux tubes experienced extremely large tension forces in opposite direction, but they couldn't mix without reconnection, so the IFTs remained stationary. Because of this, they behaved like obstacles, and their presence disturbed the plasma flow of the magnetotail. This means that the magnetotail itself could modify plasma flow internally!

But there is one last catch: our simulation only included constant solar wind. Because the solar wind didn't change, everything we observed –flux ropes, IFTs, plasma fluctuations…– couldn't be caused by an external source. The change had to come from the magnetosphere itself! These results suggest that the magnetosphere could be its own engine, which hasn't been considered before. Accounting for this could take us a step closer to achieve accurate space weather forecasting.

Adriana Bercebal Ruiz | Space and Atmospheric Physics Group | Supervisor: Dr. Jonathan Eastwood

How Did Tatooine Form? An Exploration of Planet Evolution in Binary System

Gabriel Swallow

The search for exoplanets, especially habitable ones, has long been the subject of science fiction. From the lush moon of Pandora, to the sentient ocean on Solaris, the amazing diversity of these fictional worlds beyond our own has excited audiences for centuries. One of the most exotic exoplanets from science fiction is Tatooine, the desert planet from Star Wars, which orbits around 2 stars in a so-called binary star system. Hence, Tatooine is called a circumbinary planet. Physics intuition would say such planets would be unstable, since there is not a constant gravitational force like in our solar system. The Kepler-47 system however has exceeded even the imagination of popular science fiction by having three circumbinary planets!

The birthplace of planets is in a protoplanetary disc, which orbits around the central star, or stars in the case of a circumbinary disc. The process of planet formation is already difficult, but in binary systems, it is immensely so. Disc matter is constantly falling on close orbits to the binary and being flung back out, colliding with the rest of the disc and generating a large central cavity. This extremely violent region of the disc is very unsuitable for planet formation, and yet, we find a small planet well within this region! A simulation of the surface density of this disc with no planets, along with an overlay of the observed orbits of the Kepler-47 planets, is shown in Figure 1. Understanding how such a system can form was the focus of this project.

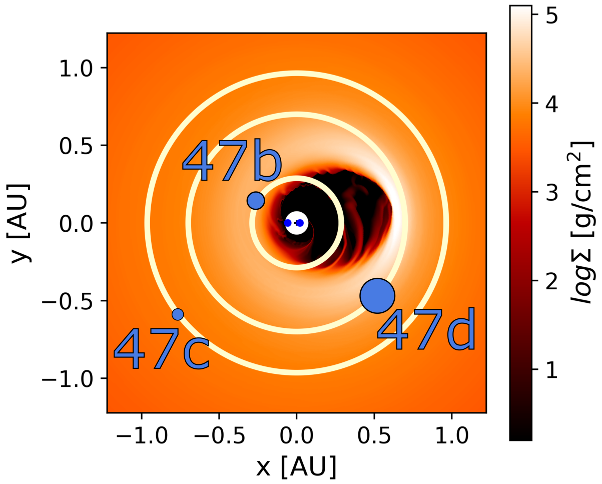

Figure 1: Log plot of the simulated disc surface density distribution Σ for the Kepler-47 binary with the three planets overlaid on their observed orbits.

We used fluid simulations to replicate the disc and attempted to understand the journey the three planets of Kepler-47 took to reach their final observed positions. The first was the inner planet, 47b. We showed it is too light to travel through the violent inner disc at < 26 Earth masses, $M_{\oplus}$. We then explored creative solutions for how it could. Jupiter mass planets ($318 M_{\oplus}$) can dominate the disc and make it suitable for smaller planets like 47b to migrate through. Their migration is unstable, as they end up on very close orbits leading to ejection. This means we have a way of getting 47b to its observed orbit, but we need a way for it to be able to stay there. Unfortunately, 47b is too light to remain on this close orbit on its own. Adding in the next planet, 47d, only added more problems. We find the two planets undergo a chaotic exchange of angular momentum, where 47b is usually ejected. So in fact, we have two problems: getting 47b on a close orbit, and keeping 47d on a far out orbit. Adding accretion to the model did not help, while changing the parameters of the disc to non-physical values had some success. We ultimately conclude there must be something missing from our model, or, observations of planet mass are wildly underestimated, and in fact 47b needs no inner planet or special disc to get to its final orbit.

With this project, we ruled out a number of possible solutions for the open problem of circumbinary planet formation. While we didn't find the solution, we have made the picture clearer for future work in the field. As observations of exoplanets better inform our models of planet formation and evolution, we also get one step closer to understanding the origin of life and whether we are alone, or just one of many observers in this universe.

What is in an exoplanet’s atmosphere and what can that tell us about their origin and formation? (Layperson’s Summary)

Exoplanets are planets that are outside of our solar system. The search for and study of these objects is one of the most prevalent fields in astrophysics today, with over 5,000 confirmed exoplanets having found by NASA to date. In this project, the study of exoplanet is furthered by investigating the atmospheric composition of the hot gas giant WASP-69b. This exoplanet is located at a staggering 163 light years away from Earth, has a radius of 1.11 times Jupiter’s radius, and has an orbital period of just 3 days. Its orbital radius is such that it is over 8 times closer to its host star than Mercury is to the Sun! But how can anyone measure the atmospheric composition of an object that is that far away? The answer: a technique called Transmission Spectroscopy.

As you may know, electrons within an atom orbit at different energy levels. Whenever an electron is hit by a photon with exactly enough energy for it to transition from its energy level to one of the higher ones, it will absorb the photon. Since a photon’s energy is related to its wavelength and different atomic species have different electron energy levels, each element or compound will absorb different wavelengths of light. We can utilise this by observing an exoplanet as it transits across its host star and note how much light is absorbed at each different wavelengths. If it absorbs more light at a wavelength where we know a compound also absorbs, we can conclude that that compound is present in the exoplanet’s atmosphere! This can be used to find an atmospheric absorption spectrum, like the one of WASP-96b taken by the James Webb Space Telescope, shown here.

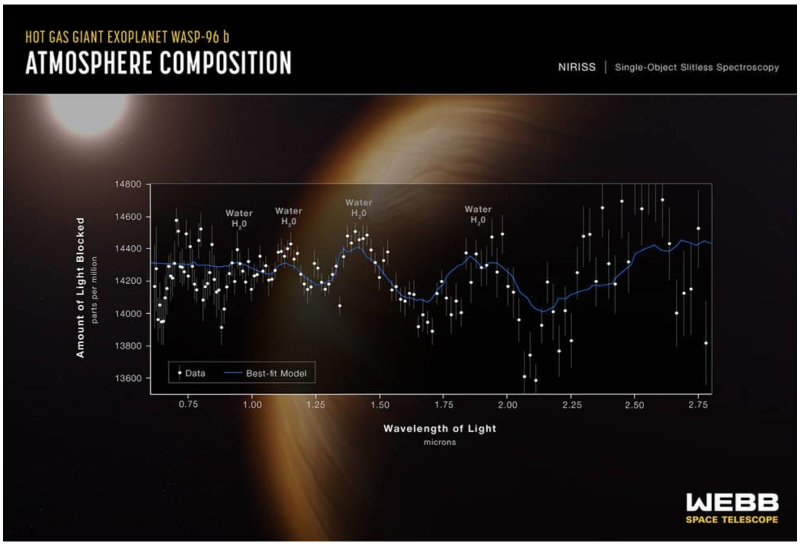

Atmospheric absorption spectrum of WASP-96b
Source: https://www.nasa.gov/webbfirstimages/

The result of our project showed that the exoplanet WASP-69b has sodium comparable to 50 times the solar abundance. This high metallicity indicates that the exoplanet most likely originated at the edge of a dense disk of gas and dust around the star and picked up sodium as it travelled inward to its final location in close orbit to the star – a process called Disk Migration. This research is important as it helps expand our understanding of exoplanets and their origin and formation, provides the perfect environment to study physics in extreme conditions, and could ultimately help in the search for habitable planets and extra-terrestrial life.

Lay Summary: Understanding the Unique Magnetic field at Ganymede

On the 14th April 2023, the European Space Agency's JUpiter ICy moons Explorer (JUICE) spacecraft began its journey around the solar system to uncover the secrets of Jupiter's largest moons. JUICE will take 8 years to reach Jupiter; once it reaches the system, it can begin its primary objective of taking a variety of magnetic, plasma, and spectroscopic measurements at three of the four Galilean moons, Callisto, Europa and Ganymede. Why are these moons of such great interest? Previous spacecraft missions to Jupiter (the Galileo and Juno missions) have provided data which suggests that oceans, 100 times deeper than the one here at Earth, could exist far beneath the icy surfaces of these moons. Understanding the nature of these subsurface oceans could provide insight into the discussion of whether life can exist outside of Earth.

How do we know there could be oceans beneath the surface of these moons? Magnetic induction is the answer. This is the process by which an induced magnetic response is generated in a conducting material due to the presence of a time-varying external magnetic field.

All three moons are well within the reach of Jupiter's large magnetic field; hence, they are exposed to a time varying magnetic field coming from the rotation of Jupiter. Furthermore, if the oceans beneath the surface have a sufficient salt content, this will result in a net conductance as electrolytes in the salts have a local charge. Therefore, when the conducting ocean is exposed to the time varying magnetic field of Jupiter, it will cause an induced response. However, this signal is incredibly weak compared to the background field, so we are sending JUICE with the most sensitive magnetic sensor ever put on a spacecraft (J-MAG) to try resolve the details of this signal to verify the existence of these oceans and perhaps even determine some of their characteristics.

In this project, we built a simulation of the magnetic environment of Ganymede. Ganymede is of particular interest as it is the only moon in the solar system to have its own intrinsic magnetic field. This is a magnetic field similar to the one here at Earth, generated in the core by a process called dynamo action. It is also the site of the final part of the JUICE mission, where the spacecraft will fall into a 500km altitude circular orbit around the moon to take measurements for the final few months of the mission lifespan.

In our model, we included magnetic field contributions from four of the main sources: the intrinsic field, the external field of Jupiter, the field of the Jupiter plasma sheet and the induced response of the ocean. By inputting the trajectories of JUICE around Ganymede, we can make predictions for what that spacecraft would measure as if it was at Ganymede!

We first compared our simulation to real Galileo spacecraft data and found that our model makes more accurate predictions at lower altitudes to Ganymede compared to higher altitudes. This indicated to us that we should focus on the aspects of the JUICE mission which have the lowest altitudes to Ganymede. JUICE's closest flybys of Ganymede's predicted field strengths ranging from 400-600 nT in strength with the induced response contributing ~2% to this. J-MAG can detect fields as small as 0.2 nT so this confirmed that we should be able to isolate its signal and confirm the oceans existence.

By performing Fourier analysis on the magnetic field time-series of the 500km circular orbit, we decomposed the signal into its composite frequencies which told us about the different dynamics of the system. We were able to identify all the key signals pertaining to the different magnetic field contributions at Ganymede, including that of the induced response, which is a promising sign for the JUICE mission as that is the signal they are looking to find.

Overall, our simulation can act as a foundational model to make predictions for magnetic measurements on the JUICE mission, while also being a powerful tool for analysis of the dynamics hidden within the Ganymedean system.

by Matthew Acevski

Fundamental Particles

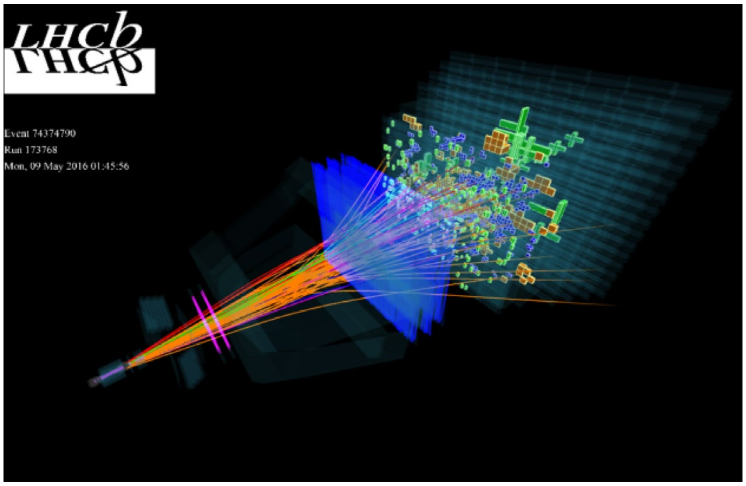

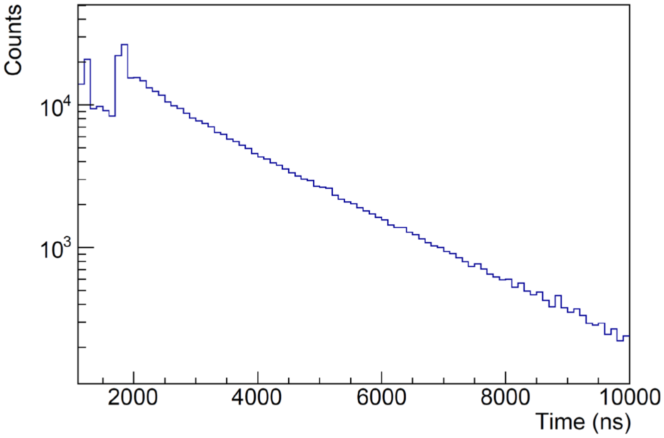

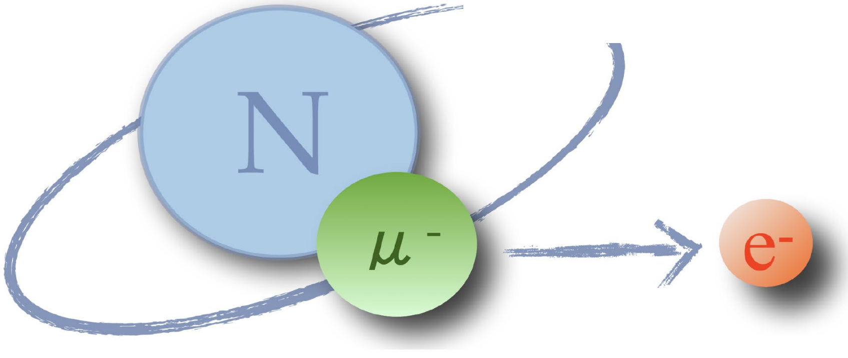

Searching for new physics with the LHCb Detector

For decades, physicists have sought after the holy grail of physics: an all-encompassing a theory capable of describing all physics at a fundamental level. This pursuit has eventually led us to the Standard Model (SM), culminating in the discovery of the Higgs Boson in 2012. Despite its success, many observations remain unexplained, necessitating the development of new theories. Axion-Like Particles (ALPs) are a class of hypothetical particles predicted by a wide range of beyond standard model theories. It is also speculated that ALPs are a constituent of dark matter, the unknown matter which constitutes 85% of the mass in our universe. Thus, by looking for ALPs, a wide range of physics beyond the standard model can be probed.

One approach to search for evidence of new physics is to collide protons near the speed of light in the hopes of producing new particles, and combing through the resulting debris to look for their decay products. This is the approach taken by many experiments, such as the LHCb detector at CERN. However, this is much easier said than done; up to 40 million proton-proton collisions occur every second, producing colossal amounts of data which we have to sift through. Moreover, other decays and interactions can produce the same combination of particles as a decaying ALP, further complicating the search. Because of this, searching for ALPs can be likened to searching for a needle in a haystack - with the added bonus of the needle looking like a straw of hay.

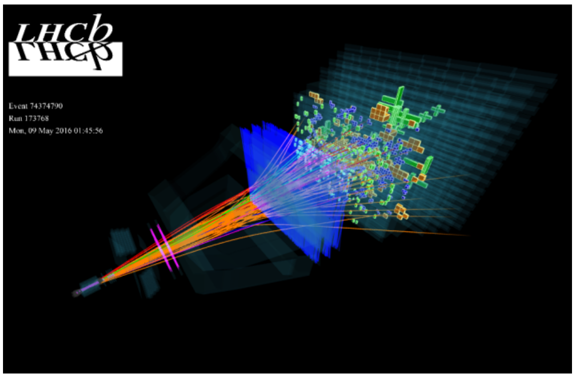

A typical event in the LHCb detector. Image: LHCb Collaboration

So how do we actually find them? For each decay event, we can determine the mass of the original particle by measuring the energy and momentum of its decay products. While the masses of background events are randomly distributed, events from a decaying particle have a distinctive mass. Each of these individual events can be likened to a dice roll; a single measurement does not tell us much, but by repeating the measurement and looking at the distribution of outcomes, we can tell if the dice is weighted. Thus, by searching for excesses in the mass spectrum, we can probe the existence of ALPs.

Even if no direct evidence for ALPs are found, we can still measure how many of these events are produced. This allows us to set limits on the properties of ALPs and narrow down the search space. Thus, by carrying out this search, we hope to constrain the existence of ALPs in an unexplored mass range. After all, the discovery of the Higgs boson was not made in an instant, but through decades of constraints provided by numerous null results.

Ho Sang Lee | High Energy Physics Group | Supervisor: Dr Michael McCann

Unlocking the Mystery of the Universe Through the Search for Muon-to-Electron Conversion

Imagine delving into the mysterious world of particle physics, where scientists tirelessly explore the tiniest building blocks of our universe: elementary particles. The Standard Model of particle physics is a testament to human ingenuity, but many questions still remain unanswered: What is dark matter? Why are there more particles than antiparticles in the universe? Why do neutrinos have mass? To solve these puzzles, physicists are searching for physics beyond the Standard Model.

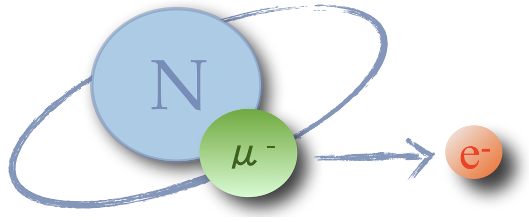

Figure 1: A diagram showing muon-to-electron cconversion.

The COherent Muon-to-Electron (COMET) experiment is one such search. It is an international collaboration in Japan on a quest to find a rare process called muon-to-electron conversion, as illustrated in the diagram. A muon is a fundamental particle in the Standard Model that has nearly identical properties to an electron, except for a much heavier mass. By special relativity's energy-mass relation, it is energetic enough to decay into an electron while also producing two extremely light particles called neutrinos. However, a muon-to-electron conversion that COMET is searching for is one where muon decays to electrons without producing neutrinos. If found, it would be a game-changing discovery and proof of new physics since this process is strictly forbidden in the Standard Model. However, it is like searching for a needle in a cosmic haystack, so COMET requires exceptional sensitivity to stand a chance.

As the saying goes, "Rome wasn't built in a day," and neither will COMET. This colossal task requires the combined effort of hundreds of brilliant minds. Our MSci project aimed to contribute to this endeavour by analysing data from COMET's preparatory stage, Phase-α. There, a detector called Range Counter produced electrical waveforms when struck by particles. It was designed to identify the particle and count the number of muon particles with specific momenta. The challenge was picking out useful information, such as peaks, from the noise in the waveforms.

We built a peak-finding algorithm that can process input waveforms, locate all peaks, which represent particle hits on the detector, and determine the time and height of each peak. To refine the information, we also fitted curves to the peaks. Unlike familiar fitting methods, like straight lines or Gaussian distributions, we crafted unique templates due to the complexity of the peak shapes. Specifically, we overlaid single peak waveforms from the data and took their average to create templates representative of all the peaks. We found this approach produced a better fit than the analytical functions found in the literature.

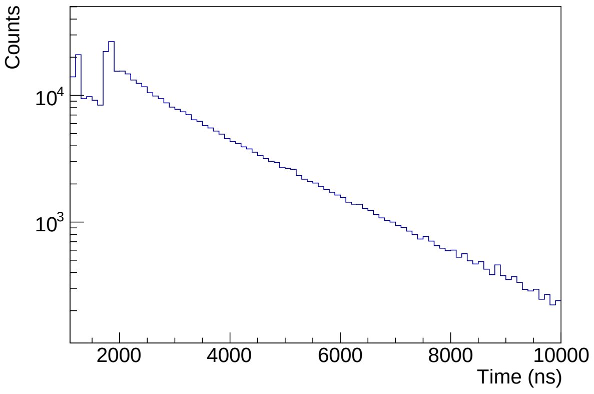

Figure 2: An almost perfectly straight line after the early time in the log-scaled plot indicated excellent agreement with the exponential decay law.

During testing, our algorithms distinguished two close-up peaks 98.5% of the time and achieved high accuracy in timing, height, and width measurements. Encouraged by these results, we applied our algorithm to experimental data, observing that electron and muon peaks clustered around different heights and widths, which means the information from our algorithms can be used to tell them apart. We also saw a smooth muon decay curve, almost perfectly following the exponential decay law, which not only validated our algorithm but also can help count muons.

These findings proved the validity and usefulness of our algorithms in data analysis and helping the Range Counter achieve its designed purpose. At this moment, our work is being integrated into the COMET data analysis code to serve its purpose. As the COMET experiment advances, refining and adapting these algorithms to similar detector signals will undoubtedly aid in its success and, ultimately, expand our knowledge of fundamental particle physics.

by Sam Zhou

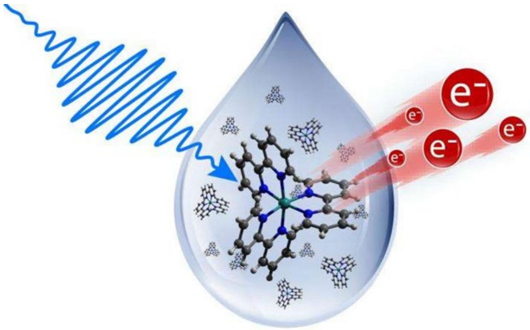

Physics, Medicine, Surfaces and complex systems

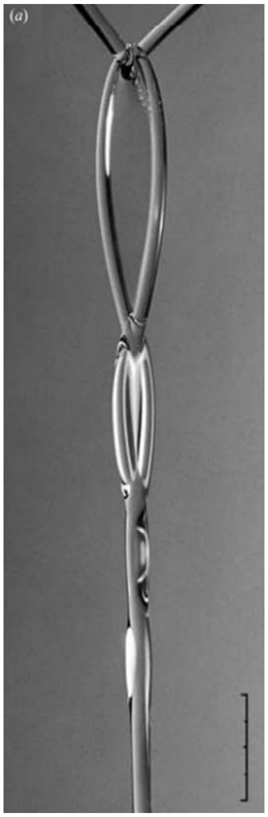

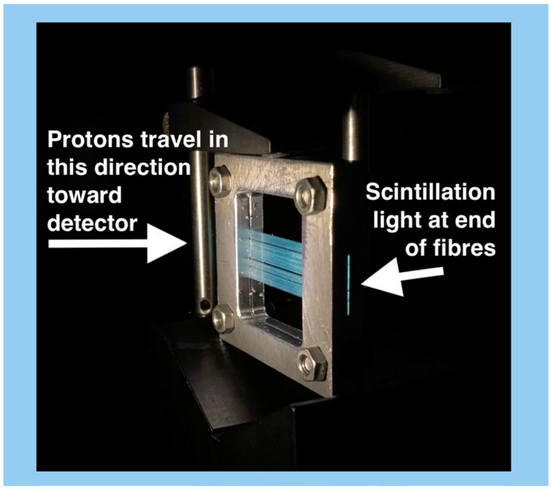

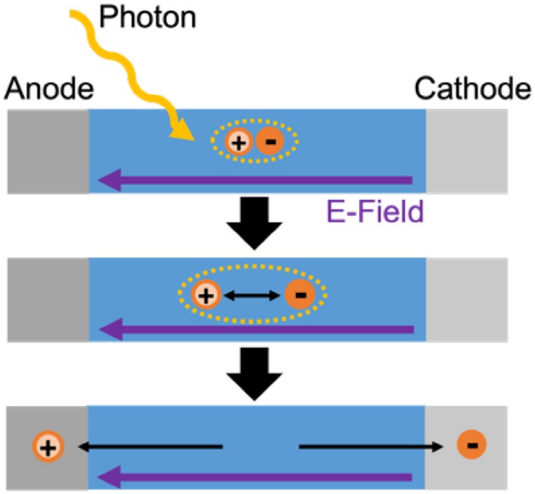

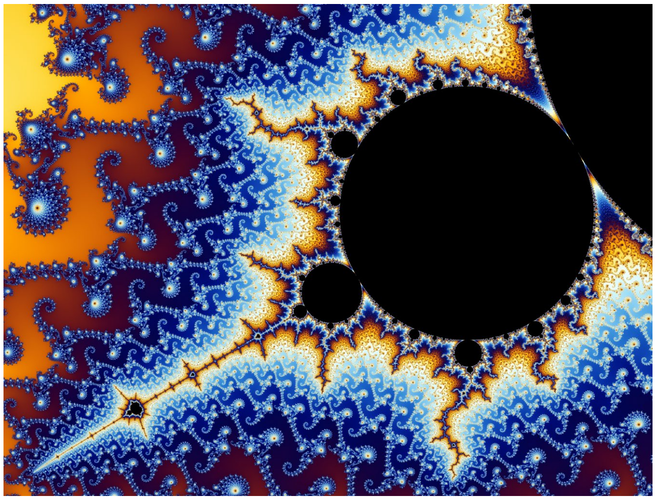

Fighting cancer with physics

Unfortunately, one in five of us will be affected by cancer in our lifetime. The most common treatments are radiotherapy, chemotherapy, and surgery. However, these treatments only guarantee recovery for about half of cancer patients. Change in cancer treatment is needed if we want to improve this statistic. The question is: what is the solution?

Radiotherapy: Treating tumours with beams of photons or particles (e.g. electrons).

Many experts believe that a breakthrough could be made if we are able to exploit a wider range of sub-atomic particles to their full therapeutic potential (now you start to see where the physics comes in!). Proton beam therapy (PBT) as a replacement for conventional photon or electron beams is of particular interest.

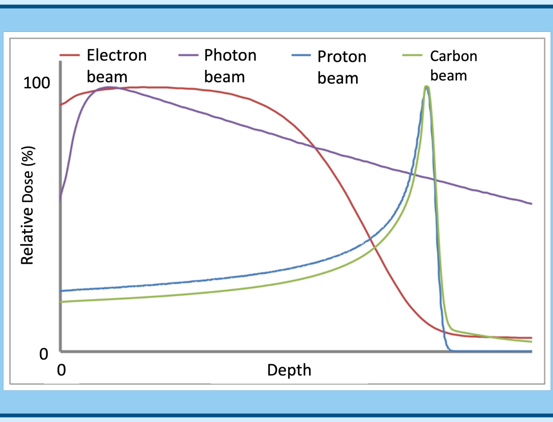

Figure 1: Depth-dose curves of different radiation beams, i.e., amount of dose (energy per unit mass) delivered by radiation as it travels into the patient. Image adapted from: Takada, 2020.

You might say, 'to-may-to, to-mah-to, radiation is radiation' – however, the key to understanding the promise of protons lies in the so-called depth-dose curves of these beams shown in Fig. 1. As you can see, most of a proton beam's dose is localised in a sharper peak than photon and electron beams. This ensures that the healthy tissue around the tumour doesn't get all of the nasty side effects of radiation. The peak of the dose can also be targeted deep into the patient, which is useful for deeply embedded tumours.

This advantage comes with a catch: we have to be *incredibly* careful that this intense peak of radiation aligns precisely with the target tumour – or else the patient's healthy tissue is at risk. We need some sort of detector which can reconstruct the dose deposition in all three-dimensions. This is the crux of my master's project: building a detector to map a proton beam's dose profile.

We designed, simulated and built a detector based on scintillation: the emission of light. The detector consists of 32 plastic fibres (like very thin spaghetti tubes) arranged in a strip and secured in a frame, as in Fig. 2. When we place the detector in the path of a proton beam, the protons excite a scintillating material in the fibres, releasing light (a photon). The photon propagates inside of the fibre by total internal reflection. As shown in Fig. 2, the light is visible at the ends of the fibres. The intensity of this light was recorded with a camera. The measured intensity is proportional to the radiation dose at that position. Using many detectors along the beam's axis allows us to build up a 3D dose profile.

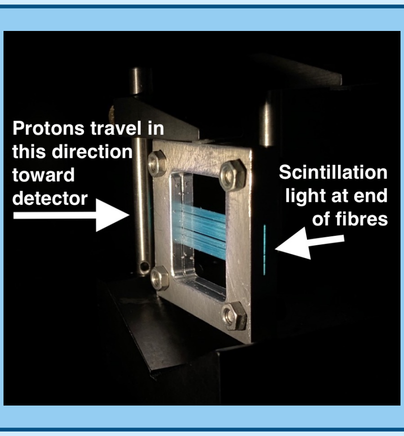

Figure 2: One of our scintillating fibre detectors.

With simulations, we showed that with just 4 detectors you can map the beam's dose profile very well. We tested our detectors with an ultraviolet source and found the fibres were not positioned as uniformly as we'd have liked (you can see uneven gaps between the fibres in Fig. 2).

The next challenge is to improve the alignment of the fibres and then to test it with a proton beam. Eventually these detectors will be used in a PBT research facility, the Laser-hybrid Accelerator for Radiobiological Applications (LhARA). We expect to see impactful cancer treatment discoveries coming out of LhARA in the next decade.

Anthea MacIntosh-LaRocque

Collapsing the Wavefunction of Cancer

Viraj Patel, Alexander Tiu

Over the last few decades, AI has helped advance research in various fields, ranging from physics to ecology. In particular, we've seen AI benefit the field of medical imaging with image recognition. We can now identify problems faster, which leads to earlier interventions by clinicians. It can also automate measurement tasks and monitor changes in these measurements, which is particularly useful in MRIs. Extensive training can also lead to AI being able to preliminarily diagnose patients more accurately. This is all great, but we require a deeper understanding of how AI works, more specifically the machine learning algorithms that AI are trained with.

Machine learning involves continuously passing data through an algorithm that learns what sections of data is important and how to process them. While generally regarded as a 'black box' type of process, the construction of such algorithms is deliberate and catered towards the specific problem. In this project, we constructed an image recognition algorithm with quantum computing techniques to identify tumours in brain scans.

There are two phases of image recognition: feature extraction and classification. Feature extraction involves the algorithm figuring out what the image contains, for example wheels and windows on a car. This is done by taking the convolution of the image. Classification is the process of using those features to determine if the image contains a car or not, and this is done via a neural network. By feeding the algorithm the 'correct' answer, it will adapt how it chooses to extract features and how it uses those for classification.

Quantum phenomena can be employed to potentially enhance an image recognition model. While regular bits can assume a 0 or 1 state, which are discrete and opposite to one another, quantum bits (qubits) are more versatile. They can possess both of these states at the same time, a concept known as superposition. It allows for more efficient information encoding, and can require fewer operations to be performed. For example, instead of performing a given operation on the 0 and 1 states individually, it only needs to be performed once on a state in superposition containing both 0 and 1 states. This can vastly reduce the computational load.

Entanglement is a feature of quantum mechanics that links qubits to each other. This would allow us to determine if different pieces of information are linked, which can be significant in image recognition as many pixels can all be representing a certain object. For example, if the image contains both windows and wheels, it is more likely to be a car, so the property of possessing both of these features is linked, or entangled.

Through building a quantum convolution algorithm, for feature extraction, we determined that the quantum model is more accurate than its classical counterpart for lower resolution images. This is due to the poor performance of the classical model at low image qualities, indicating that the quantum model is not greatly hindered by having limited amounts of data. Thus, our quantum algorithm is better than its classical counterpart at recognising smaller tumours, and tumours in low resolution scans (e.g. a zoomed in section of the brain).

We noticed that the quantum neural network performed better with more qubits. We also found that with just 8 quantum circuits each containing 4 qubits, we were obtaining a comparable performance to classical neural networks. Unfortunately, current quantum devices are limited in size, so larger neural networks designed for more complex problems aren't feasible yet. Researchers at Imperial College are currently building larger quantum computers, so it won't be too long before we can solve harder problems with quantum algorithms. Nevertheless, we ensured that our algorithms were scalable, so that they're well adapted for the near future.

Despite the technological constraints, the quantum convolution algorithm achieved close to 86% accuracy, compared to its classical counterpart that fell just short of 80%. The quantum neural network also managed to achieve a 1.6% accuracy advantage to its classical counterpart. With the limited number of qubits we had, we have shown that quantum machine learning can be applied to image recognition problems, and identified avenues where it has a tangible advantage over classical machine learning.

Optical Projection Tomography: CT scans but for mice?

Developmental biologists are people who are interested in studying how organisms grow. They may study embryos, larvae and even whole organisms. As you may well know, scientists cannot directly perform these experiments on humans due to certain ethical reasons (imagine the uproar!). This is why developmental biologists result to studying 'model organisms' such as mice and their embryos, that are structurally similar to humans.

There is a need to image these millimetre-sizeed specimens. The traditional way to do this has been cutting up the embryo into hundreds of slices, imaging them sequentially and computationally stitching them together. However, this is extremely time-consuming. Optical Projection Tomography (OPT) is a relatively new volumetric bioimaging technique that does this non-invasively and is the optical analogue of X-ray CT. A few hundred images of the sample are taken, as it undergoes a complete rotation (unlike medical CT scans when the source and detector rotate instead). The stack of images is then processed using a computer with what is called the filtered backprojection algorithm to do a slice-by-slice reconstruction of your sample— allowing you to look inside it without cutting it up. This is done using visible light because biologists know how to label specific proteins or genetic material with either precipitates that attenuate light or fluorescent markers that emit light when excited with photons that correspond to its energy gap. These markers do not exist with other similar techniques such as MRI or X-ray CT.

Biological tissue is usually milky in appearance due to the scattering of light. However, one requirement of OPT is that samples have to be see-through. Thankfully, a range of chemical techniques have been developed to perform optical clearing for us.

We constructed an OPT setup (think screws and Allen keys), wrote the image processing scripts (lots of debugging) and applied it to image an optically cleared mouse embryo (with gloves because the clearing substance is somewhat toxic).

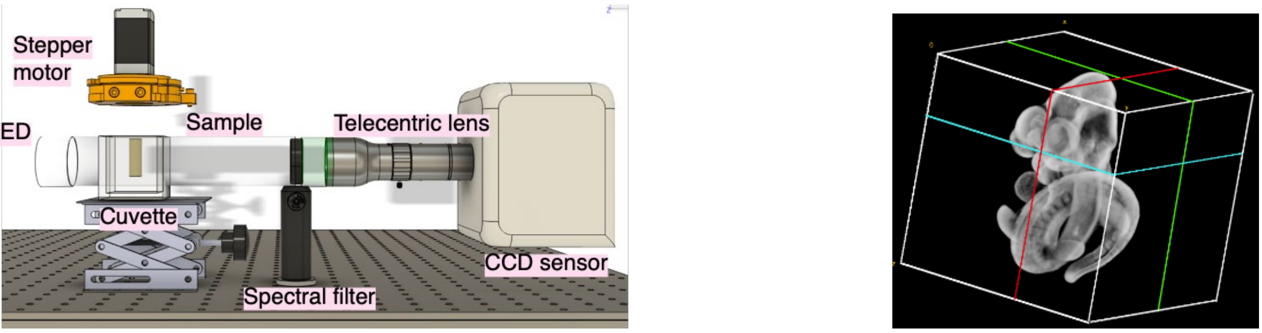

Left: CAD diagram OPT setup in transmission mode. Right: A 10-day-old mouse embryo that we imaged in fluorescent mode. The edge length of the cube is around 7 mm.

As physicists, we wanted to determine the limits of imaging performance. One such measure is image resolution, which is the minimum distance two signals can be separated for them to still be distinguished. The resolution is limited by diffraction from the aperture in the lens, which is best at large aperture sizes. As we needed the whole sample to be in focus when imaged, we also wanted to have a large depth of field, which is favoured with small aperture sizes. This unfortunately is the trade-off in OPT! We took images of known physical targets (a sharp knife edge and a fluorescent sample bead sample) to quantify this effect so OPT users can know the optimum aperture size for their sample.

In addition, we also prototyped a new system that uses an LED matrix as illumination. We discovered that by illuminating successive halves of the LED and taking images over a complete rotation, we can reconstruct the entire 3D refractive index of the sample! This is pretty exciting as existing methods that do this are complicated and expensive. It can be applied to label-free imaging of zebrafish, which developmental biologists know and love. We enjoyed this project as it's a manifestation of how physical principles, experimentation and computation all have a role in understanding complex biology— one embryo at a time.

Author: Kenton Kwok | Supervisor: Dr James McGinty

Going with the Flow - One Thin Liquid Sheet at a Time

Have you ever lay in bed and thought "Why?"…"Why am I here?"…"Why am I made the way I am and why do things work the way they do?" Life is still not fully understood after 3.7 billion years on Earth. But what if I told you there's a way we could get closer to the truth…

What is essential to life's survival? Water - the most abundant and common substance on Earth. It plays a crucial role in biochemical reactions, acting as an excellent solvent due to its polar nature. Unfortunately, scientists don't fully understand what is happening in these reactions. Yeah, molecules have been broken apart, electrons have changed states, but what is REALLY going on? To not only improve, but to understand, you must know how something happens. To improve the rate or even change the outcome of a biochemical reaction, scientists must investigate what happens between the final and initial states. To investigate the dynamics of these reactions, scientists use 'pump-probe' experiments which utilise super-short laser bursts, even smaller than the timescale of the reaction (~10^{-18}s). The first pulse can be used to 'pump' energy into the liquid to initialise a reaction and then a second pulse is used to 'probe' the liquid to see what's changed. Unfortunately, most of these experiments use laser pulses in the soft X-ray region which love being absorbed in liquids. To minimise absorption, we not only need a target which is thin and vacuum-stable but one that has self-healing properties to resist any damage from the pulses. Not asking for a lot, huh? However, I think I know just the thing…

Thin liquid sheet jets! These often look something like Figure 1, but it depends on how you want to form them. A popular way is to direct two cylindrical liquid jets at each other so that they form a sheet in the plane perpendicular to the collision plane. But how do these sheets behave? Where are they thinnest? Which part of the sheet should I do my pump-probe experiment on?

That's where this project came in. The behaviour and thickness of these sheets were characterised for both isopropanol and water for a variety of flow rates. Water tended to have shorter sheets than isopropanol, likely associated with its higher surface tension (sheet surface's resistance to stronger external forces). And as flow rate increased, water's jet length increased quadratically, but isopropanol's more linearly – this may be due to a difference in viscosities of the fluids, which is a measure of a fluid's ability to flow. At certain flow rates, the sheets began to misbehave and break up into droplets because they just couldn't handle the higher inertial forces or there wasn't enough of them. Very unstable.

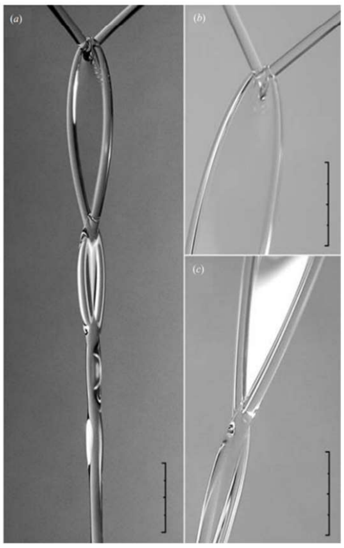

Figure 1: Thin liquid sheets formed using colliding jets. Scale: 1 cm. Image credit: J. Bush

Some time ago, two pals, Hasson & Peck, derived a relationship which indicated that the jet becomes thinner as you move further away from the collision point. Sometime quite recently, two pals, my partner and I, decided to take that theory to the test. But these liquid jets are micrometres thin so we couldn't gauge the relationship by eye. A ruler would disturb the sheet. So instead, we used thin-film interference! When light meets an air-liquid boundary it can be partially transmitted or reflected. Thin jets have two of those boundaries - the upper and lower surface. The light can take different paths and recombine at the surface to interfere constructively or destructively to produce fringes. Of course, the location of these fringes will be determined by the optical path difference between the two surfaces, which is determined by the sheet's thickness! This project used this phenomenon to find that as you get further away from the collision point, the relative thickness changes linearly by around 0.7 micrometres over 0.05 mm. However, some of the optical flatness was affected by standing waves creating ridges in the sheet from the shock of the colliding jets.

Generate and characterise a thin liquid sheet = answer life's biggest question: "Why do things work the way they do?"

By Maria Filippova

Layperson's Summary: The Simplicity of Complexity

“To see a World in a Grain of Sand
And a Heaven in a Wild Flower
Hold Infinity in the palm of your hand
And Eternity in an hour”
- William Blake, *Auguries of Innocence*

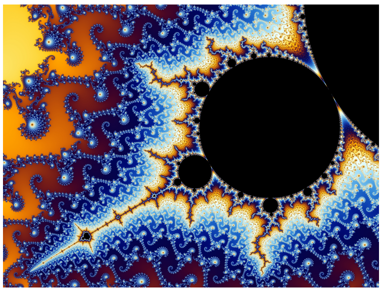

In the realm of physics, nature is a paradox - profoundly simple, yet infinitely complex. The concept of a circle, a square, a triangle is so intuitive to us but can you tell me where in nature I can find one of these objects? Show me one and I'll tell you to look a little closer. But this is obvious, you may say, clouds aren't *really* spheres, mountains aren't *really* cones and lightning doesn't *really* follow straight lines. So, what are they *really*?

In 1980, Benoit Mandelbrot stumbled on a mathematical object of profound natural beauty. An object that, from simple rules, consists of literally infinite complexity. The so-called Mandelbrot set (top right) exhibits finer and more elaborate recursive detail at ever increasing magnification. Search online for a zoom into the set and you will inevitably be struck by the likeness to a plethora of natural phenomena. Watch as spiral seahorses melt away into psychedelic islands and fjords, insect antennae morph into fragmented mountains and valleys cascade back to an ocean of spirals again. The Mandelbrot set is perhaps the most famous example of what is now known as a *fractal*.

Coined from the latin term “fractus” meaning “broken” or “fragmented”, a fractal is a special kind of geometrical object that contains self-similarity at whichever length scale you look at it. Unlike typical shapes which have integer dimensions (a line is one dimensional, squares and circles are two dimensional and so on), a fractal has a dimension that is somewhere in between - a *fractal dimension*. How bizarre! Despite this counter-intuitive notion, fractals are everywhere in nature. Look at a tree and see how the branches break off into smaller branches and so forth, see how rock features fragment into smaller rocks and granules. Fractals seem to elegantly capture this self-similarity in very simple rules. Examples of these objects have been around since the late 19th century, but it was only at the advent of computer graphics that the extent of the beauty of fractals could be seen. Nowadays, simple algorithms can construct highly detailed fractal trees and landscapes using iterative growth models (one of which, used in this thesis, made the background graphic on this page!).

However, fractals have gone far beyond niche examples of generative art. It turns out that fractals have direct applications, from modelling crystal growth to lightning strikes. By computer simulations, it is possible to model a wide range of phenomena that all fall under *universality classes*. Whether it be water percolating through a sponge, or a flame burning through paper, they can end up exhibiting the exact same fractal properties - despite a wildly different natural manifestation. This is where this thesis picks up. Cancer is an exceptionally deadly disease everyone is familiar with and there has been a flurry of recent research seeking to apply fractal techniques. What if cancer falls under one of these universality classes? Is it possible to find a growth model that encapsulates the essential characteristics of cancer growth? In this thesis, a new fractal growth model is proposed that accounts for the distinct features that make cancer so deadly. By analysing the fractal dimension of the growth model, a direct correspondence can be made with real cancerous tissue. With further work, patients' cancer growth could be simulated and predicted, possibly a deciding factor in surgical priority. The hope is that fractal simulations, like the one here, could be the standard in a doctor's extensive repertoire. (Mandelbrot set photo by Wolfgang Beyer, Wikipedia Commons).

by Monte Ren

The Quest for High-Performance Organic Solar Cells

Shi Wei Yuan

In the current global warming crisis, we face rising sea levels, more frequent natural disasters and the extinction of countless wildlife. To stop this, we need to quickly move away from burning fossil fuels to sustainable energy sources such as solar energy. Solar cells use photons to excite electrons, leaving a positively charged hole in its place, and transport these charges to its electrodes through an internal electric field. This process of "charge generation" is shown in Fig. A. When both electrodes are connected to an external device, an electrical current will flow and power it. Commercial solar cells use the elements silicon, but the increasing applications of renewable energy call for other technologies. Organic solar cells, made from molecular materials, are extremely exciting as they open the door for applications like power-generating windows, solar-powered vehicles and green-house compatible solar cells.

Organic solar cells have seen a rise in power conversion efficiency from 1% to almost 20% in the last two decades with the development of more and more complex molecules, see Fig. B. However, the specific properties that make a molecule good for solar cells are not well understood. Furthermore, in organic materials, the generated electrons-hole pairs are more strongly attracted to each other than in silicon, which is why state-of-the-art organic solar cells mix multiple organic materials together to split the electrons and holes. This makes the impact of individual materials on power conversion efficiency even more unclear.

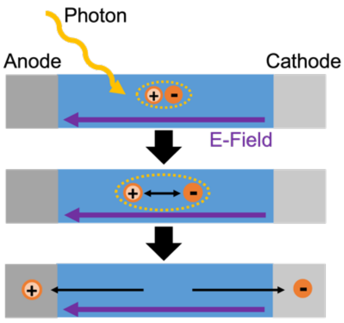

Figure A: Charge generation in an organic semiconductor

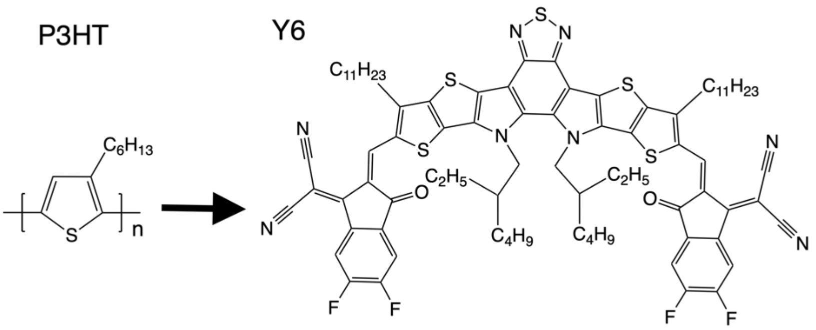

Figure B: Molecular structures of organic semiconductors P3HT and Y6

In my project, I investigated different organic materials by making single-material solar cells and measuring how many free charges they could generate with the help of an external voltage to split the electron-hole pairs. I discovered that materials like Y6 needed less voltage to effectively split electron-hole pairs than others, and that this caused these materials to generate higher currents in solar cells. Furthermore, I found that materials similar in molecular structure produced similar results. These findings help to explain why materials like Y6 perform better in state-of-the-art mixed solar cells than materials like P3HT. By measuring materials of similar molecular structure, we can now investigate which exact molecular properties improve charge generation in solar cells. By simulating these materials, my group and I found that the ease of electron-hole splitting is likely caused by how easily the electrons and holes travel through the material, or how strongly the electron is bound to the hole.

This newly developed method of investigating the charge generation in single materials will help to develop new, high-performing molecules that will pave the way for low-cost, high-efficiency organic solar cells in our fight against global warming.

Physics and Information

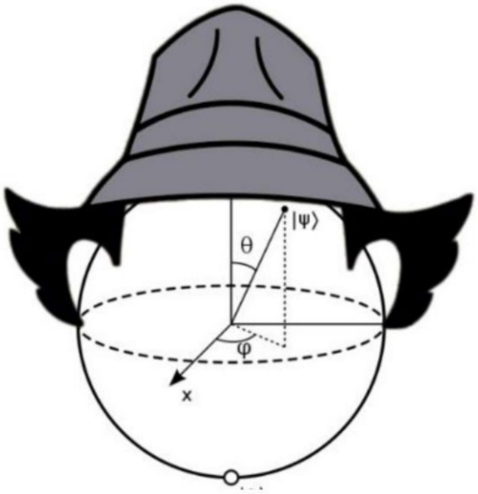

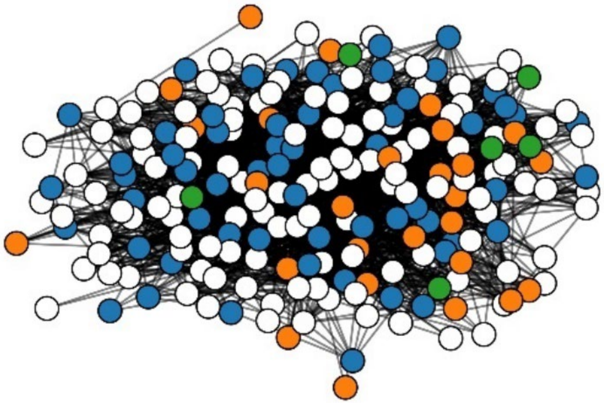

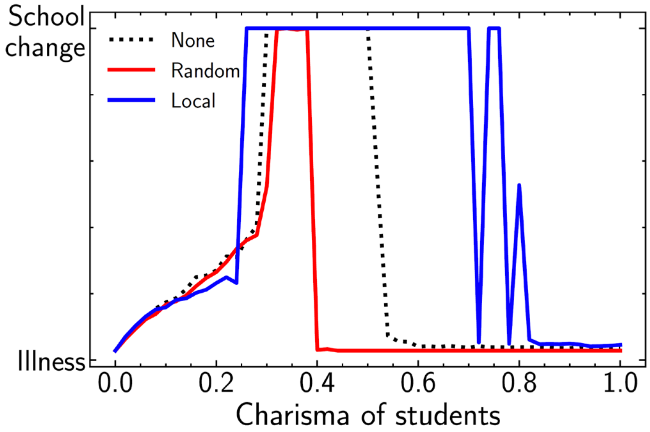

How to Spread a Rumour

By Tom Cowperthwaite

We've all been there – a tantalising piece of gossip spreads like wildfire through your school, regardless of whether or not it's true, but have you ever stopped to consider the mechanism through which it happens?

Research into this question falls into the field of sociophysics, which, among other things, is concerned with how ideas and opinions spread through groups of people. One common approach is to consider some population, such as a school, as a network of people connected according to whether or not they talk to one another.

Usually, scientists think about only the interactions between pairs of people, but clearly, that can't be the whole story, because we know that people sometimes interact in groups as well. The study of how opinions and ideas spread through groups is much less developed than its pairwise counterpart (frankly because the maths is much harder!), but researchers have found ways of modelling these group interactions through games.

Now these aren't the types of games you might be thinking of; these are a simple set of rules that people in the network follow each time they interact with their neighbours. One particularly interesting game is the naming game, where each person in the interaction is incentivised to agree with the others on the "name" they assign a particular object. This is quite a good assumption because, in real-world interactions, people may be put at a disadvantage if they disagree with their neighbours about the labels that should be assigned to things around them (imagine if you were dropped into a country where you didn't speak the language – you'd feel disadvantaged wouldn't you?).

In our recent work, we have not only taken group interactions into account but have also considered that the people you interact with may change over time. One way that people's social contacts change is when they have a disagreement with somebody; they may choose to move away from that person and stop speaking to them in future – if this happens, we say that the social network has rewired (i.e., changed its structure).

We call our model that includes this rewiring effect The Preacher Model, as the mechanism is like how a street-preacher attempts to convert people to the cause they advocate, and quickly move on if they are unsuccessful. Our work investigates the effect of different "preaching" strategies: the preacher may choose to move to a nearby group of people to attempt to convert OR they may simply select a random group to move to.

As an example, imagine you're sick and take a week off school without telling anyone. Most people at school assume you're sick, but a small group (say 3%) decide to spread the rumour that you've changed school and aren't coming back – how successful would they be at spreading this falsehood?

Clearly, success depends on many factors, and one we consider is how charismatic the people in your school are – if they are very persuasive, then information travels faster and vice versa. The image shows simulations of the status of the school's opinion by the time you get back, over the full range of charisma (0 to 1), for three different preaching strategies.

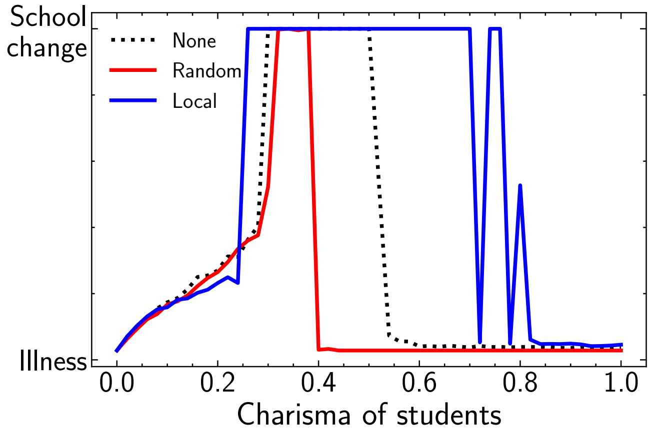

Different lines indicate the different strategies that the gossips could use, and we show that the most effective way to spread a rumour is by using the local preacher approach, and the least successful (by far) is the randomised preacher, as it is only successful over an extremely narrow range of charisma. This is because the rumour gets spread faster once it has a core of closely connected people who are convinced of the rumour and help to project it further. The opposite effect happens in the random preacher, where the minority who hold the rumour never manage to coordinate amongst themselves because they are always too far apart, sometimes the group even fragments completely!

Whether you are a victim or a spreader of rumours, further research into opinion dynamics on group interaction networks will be invaluable in coming years, as governments and companies rush to understand the effect of the massive group interactions that are made possible by social media□

Inspector Qubit and The Charged Black Hole

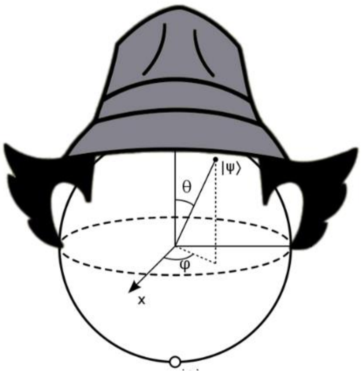

Inspector Qubit goes on many adventures through the universe. A quantum bit or qubit is a quantum analogue to the bit. Bits can be a 0 or 1 whereas a qubit, in addition to being either can be both at the same time. This unique wholly quantum property of the system can be taken advantage of to help with investigations into black holes! (Inspector Gadget hat found on stickpng.com and qubit found on Wikipedia).

Picture this: a tiny little qubit with a magnifying glass and an inspector's hat on going on adventures and investigating properties of black holes. You might think of a qubit in the context of quantum computers, but it turns out they have many applications throughout multiple areas of physics. Our little qubit is a detective and loves to investigate scalar fields - these are things which have values at every point in space and are not pointing in any particular direction. Think temperature, pressure, and humidity but with a quantum twist! Our qubit travels to many unusual places all throughout the universe to investigate how these scalar fields behave. Some areas that he investigates the scalar fields in are boring, but things get more interesting when mass is introduced - that's where black holes come in! Black holes are very heavy objects in space where gravity is so strong that not even light can escape once inside. The boundary where light can just about escape is called the event horizon and past this point is of interest to many physicists, including our qubit.

General relativity is a theory that says mass can bend space-time, and is one of the most accurate theories to grace physics. Think of it like putting a bowling ball on a trampoline - the mass of the ball causes bending of the trampoline's surface. However, instead of just up, down, left, right, back, and forth, there's also time to consider. It's like a fourth direction or surface. This theory is useful for things like GPS on your phone or in particle accelerators where its effects need to be accounted for.

Our qubit is a curious investigator, busy seeing how different terrains affect the scalar fields that occupy all of space. But the most exciting landscapes are the extreme ones, like charged black holes.

To measure the scalar field the qubit can be imagined as tying an unbreakable rope to himself and getting as close as possible to the black hole's event horizon, hovering just outside. If he was falling inwards, he would pass the event horizon (never to be seen again), so we want to prevent this. He then takes a measurement of how the gravitational forces affect him, as he is much simpler than our scalar field. This is because he is only at one point whereas the scalar field is everywhere.

The qubit's job is to measure the effects of the scalar field, so he can tell what is different about it compared to when there is no black hole. As it turns out, the black hole makes our scalar field act like a heat bath, which causes our qubit to go through something called "decoherence". The scalar field is essentially measuring our qubit just as much as our qubit is measuring him! This would cause someone else watching our qubit to be less and less sure of what the qubit's current state is. The qubit is usually a 0 or a 1 or some mixture of this, but the scalar field changes it, so the qubit is an exact 50:50 mixture of 0 and 1, so we are the least sure of what state he is in. This is exactly what heat baths do to qubits and it is caused by the black hole. How quickly this happens depends on the temperature of the black hole.

Overall, our little qubit detective has had a lot of fun exploring the properties of black holes and scalar fields. Who knows what other mysteries it might solve in the future!

By Kieran Joseph

Quantum Gravity and Black Holes

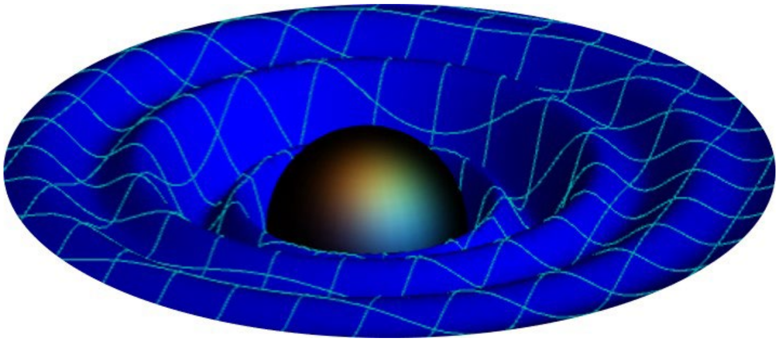

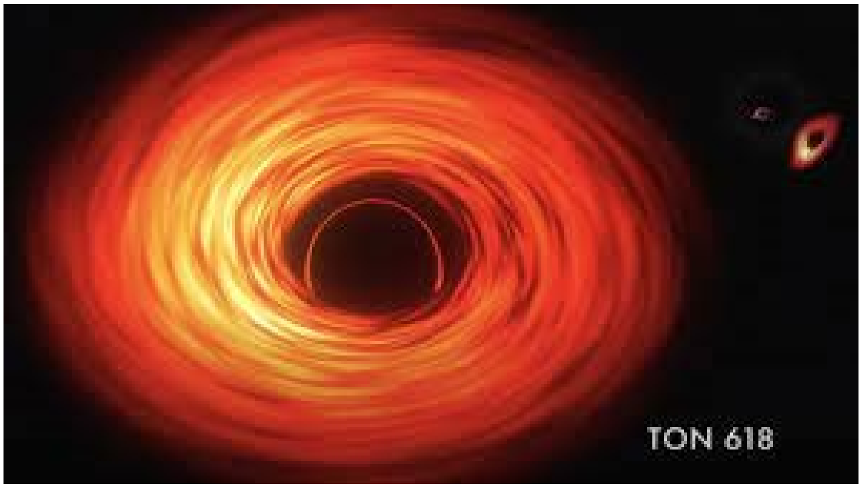

Constraints on a fully unified theory

Before I started my degree the two things in physics that captivated me were the promises of studying the very big and the very small. From the accelerating expansion of the universe to what makes up the fabric of its structure at the most fundamental level. However, you quickly learn that there is no theory in physics that can describe both. There is no grand unified theory. The Standard Model of particle physics comes close, unifying three of the four fundamental forces, but leaves out gravity. Gravity describes the very big. Although what we mean by very big is relative. Gravity describes both the evolution of galaxies and an apple falling to the ground. Obviously, these are not on the same scale. However, they are both relatively very big when compared to the very small. The very small is described by quantum mechanics, a world of weirdness, which according to classical physics doesn't make sense. Gravity is a theory of classical physics. The rules of quantum mechanics act on the scale of fundamental particles and gravity acts on much greater scales. However, what happens when there is a system which requires both gravity and quantum mechanics to describe it? A system that is incredibly heavy but also incredibly small. Such a system is called a black hole.

In 1916, Einstein completed his theory of gravitation, general relativity. General relativity posits that space and time exist on equal footing and, therefore, should be unified in the concept of spacetime. The theory goes on to state that this spacetime, which makes up the fabric of the cosmos, should bend and curve due to the presence of matter and energy. Furthermore, it is this curvature of spacetime that is responsible for gravitational phenomena. However, general relativity plagues us with infinities as soon as we apply it on a quantum level. In physics, when infinities crop up in calculations, if they can't be removed, it means your theory has gone wrong. This is where you learn that the very big and the very small cannot currently be described by one theory. Throughout history physicists have always tried to unify seemingly distinct phenomena in the quest for the *theory of everything*. So naturally, after Einstein published his paper on general relativity a flurry of research went into trying to unify general relativity with the rest of physics. However, the 1920s signified the dawn of quantum mechanics, general relativity's adversary in being a part of a fully unified theory.

Armed with the knowledge that general relativity and quantum mechanics are incompatible, physicists spent the twentieth century looking for a way to deal with this, looking for a theory of quantum gravity. It is a theme in physics that many theories tend to be suitable approximations of more general theories. Newtonian gravity is an approximation of general relativity, and it was discovered that general relativity can be thought of as an approximation to an underlying complete theory of quantum gravity. However, even though general relativity can be thought of as an approximation to some unknown unified theory, that does not mean it can be completed into this unknown *theory of everything*. There are several constraints on whether a theory can be completed into a full theory of quantum gravity. One, is the weak gravity conjecture (WGC), proposed by Arkani-Hamed et al in 2007. It promotes the long-standing observation that gravity is the weakest force, into a principle and if this conjecture is not satisfied then a theory cannot be completed.

In our project we investigate this conjecture. We did this by studying gravitational waves, ripples in the fabric of Einstein's spacetime, around black holes and calculating their speed of propagation. By studying this for black holes with a specific charge to mass ratio, one is able to see how causality effects the WGC. Therefore, one is able to see how causality plays a role in a small step to finding a *theory of everything*. We found that causality has no effect on the WGC in the case we investigated! However, in order to draw the conclusion that causality has no impact on this conjecture, other cases have to be considered.

Kathryn Jones

Lay Summary: Discrete Spacetime and Black Holes

The mysterious black holes. You have heard of them before, but what are they? Imagine you are flying through the universe in your rocket and have decided to go black hole sightseeing. You can see it from afar, but you want a better look. As you travel towards it, you suddenly start feeling weird and discover that you are being torn apart! While you turn around, full throttle trying to escape, unfortunately, your trip ends on a sour note. What happened?

Essentially, as you get too close to a black hole, you cross the event horizon, the “point of no return”, which defines the surface of a black hole. Once inside, nothing can escape the black hole’s gravitational pull, not even light. You are doomed to fall into the singularity at its centre. Since light cannot escape into the outside world, and nothing can travel faster than the speed of light, this implies that all communication is cut-off. That is why black holes are black. Hence, the information about the things that disappear into the black hole is lost. We can describe any swallowed object on a microscopic level by the distribution of the energy configurations of the molecules making up the object. These have some disorder associated with them and hence entropy, therefore we conclude that also black holes must have entropy.

As described before, entropy is usually related to the distribution of configurations of the molecules inside the object of interest. However, nothing inside a black hole can communicate with the outside, so it would not make sense that entropy resides inside the event horizon. On the other hand, if it is outside, we are not dealing with a black hole. Therefore, the entropy of a black hole can only possibly live on its surface! However, the black hole’s surface is nothing but empty spacetime, as you tragically discovered during your trip. Hence, what does black hole entropy correspond to? In our MSci thesis, we propose that it is encoded in the distribution of “horizon molecules” making up the spacetime on the surface of a black hole.

We assume that the universe is fundamentally discrete, like pixels on a TV screen, and conjecture that the spacetime continuum arises merely as an approximation to an underlying fundamental discrete structure, a causal set. One can imagine a causal set as a network permeating spacetime. The nodes of this network represent events, coordinates in time and space, and the edges dictate which events in the future can be influenced by the events in the past. That gives spacetime its causal structure. Hence, only specific causal set “networks” would correspond to, e.g. the flat spacetime we experience daily or to a curved spacetime surrounding a black hole we study in our thesis. Since this network-like structure permeates the entirety of spacetime, some edges must connect events across the spherical event horizon of a black hole, forming either simple links between two events or possibly more complicated connected structures when e.g. one event inside connects with more events outside. That is what we call “horizon molecules”. We can study the distribution of different types of molecules and count them to obtain entropy, which is proportional to their number.

Our thesis is the *first-ever* study of the black hole entropy in the Causal Set Theory. Building a highly-parallelised `C++` simulation framework, we used Imperial’s computing cluster to simulate causal sets $\sim$ 8300x larger than the previous largest recorded causal set in curved spacetime. As with any theory, we would need to be able to make predictions that agree with the general scientific consensus. Our results confirm that the entropy of a black hole scales proportionally to its area as predicted by the Bekenstein-Hawking formula, which is a well-established result. Moreover, we are able to show that the fundamental discreteness scale, the minimal distance between events in spacetime, is on the order of Planck length $l_p \sim 10^{-34}$m, where general physics is known to break down. Perhaps it is indeed because we breach the physical minimum when we try to probe the world beyond the Planck scale?

by Vid Homsak

What Gravitational Waves Tell Us About Quantum Gravity

Ever since an apple first fell on Isaac Newton's head, our understanding of the force that keeps us stuck to the ground has progressed astoundingly. This force is gravity, which Einstein completely redesigned from our friendly neighbourhood $g = 9.8\ \text{m/s}^2$ to a fundamental effect of the curvature of spacetime, i.e. the structure of the Universe itself. As counter-intuitive as that already was, it soon got worse. When Karl Schwarzschild found one of the many solutions to this theory, he discovered the mathematical description of what would become known as a black hole. These sound like something straight out of science fiction, with **nothing** being able to escape them past a point called the event horizon, being inevitably doomed to be crushed by the singularity at the centre. This point of infinite curvature makes physicists so uncomfortable that in 1969 Sir Roger Penrose hypothesised that no black hole should ever allow anyone to take a close look at its singularity and return home to tell the world.

Unfortunately for them, physics doesn't seem to care about what physicists prefer. When electric charge is added to Schwarzschild's black hole, it becomes a Reissner-Nördstrom (RN) black hole. Now a second event horizon is created close to the singularity, approaching the original horizon as more and more charge is sucked in. At some point, the mass and charge balance out in such a way that both horizons combine into a single one, after which they disappear completely, leaving the central singularity exposed for the entire Universe to see! We call these **extremal black holes**, and even though we are yet to observe any of these, they are a thorn in the side of modern physics, which researchers tend to love and hate equally.

Einstein's gravity presents yet another new phenomenon - **gravitational waves**. These are ripples in spacetime, analogous to those we see spread through a pond when if we throw a rock (or ourselves) into it. In Einstein's theory, they propagate at the speed of light, which is the universal speed limit for anything which wants to move from point A to point B. Quite fast... but never faster.

Sadly, this gravity still has its handful of issues. Worst of all, it does not get along with quantum mechanics, which describes all tiny things. So is that it for physics? Maybe not so fast. It turns out that there are quite a few ways we can solve this. However, not only are we unable to prove any of them experimentally, but we can't even agree on which one to put our money on (at least for now). Of course, these new theories further redesign our old understanding in unbelievably messy ways. Thankfully, we can look at the smallest effects these new-age gravities have on our not-so-classic gravity and still draw exciting conclusions about what path to follow, or at least shed some light on the branching paths ahead of us.

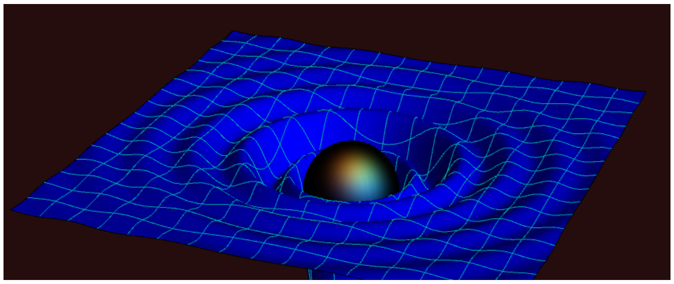

Figure 1: Representation of the generation of gravitational waves around a spherical body.

Our project aims to determine what effects these modifications bring to the RN black hole, as they change the description of spacetime around it, similarly to what other researchers have done in the past. These effects include changing the mass-charge relation which defines the point past which we expose the central singularity. But this proves nothing, apart from the fact that we may still have the same old dreaded extremal behaviour, now with even uglier equations to scare students away. This is where our research comes in - poking these new black holes and seeing what kind of ripples they throw at us. If these ripples go wild or move so fast that they break the universal speed limit independently of how we tweak the parameters of our new theory, then this could mean that quantum gravity provides further protection from those nasty exposed singularities.

Although our results haven't uncovered any pathological rule-breaking, forcing them to follow the speed limit can be used to propose additional constraints on the effects of the new theory we considered. This, together with other past published work, helps further narrow down our search for the coveted Theory of Everything.

Miguel Barroso Varela Supervisor: Prof. Claudia de Rham Theoretical Physics Group

Layperson’s Summary - Gravity as the Weakest Force

We are often interested in probing the extremes of nature. The theory of gravity describes the universe at a large ‘macroscopic’ scale. Matter and energy distort space, and time - the fourth dimension of one continuum known as spacetime. The bigger the object, the bigger the distortion. Picture the universe as an infinite trampoline. When we sit on this trampoline, we cause it to bend. However, if a tiny twig is placed on the trampoline, it causes hardly any distortion at all. A ball on the trampoline would spiral into us and fall into the dip we have created. This is analogous to how the bending of spacetime causes motion in the bodies of the universe. Locally, we can consider ourselves to be like the twig, and gravity on a small energy scale reduces to the familiar Newtonian, flat-space, gravity.

At the largest energy scale, objects can cause space-time to curve so much it collapses in on itself and the curvature becomes infinite at one point - a black hole. The gravitational attraction of a black hole is so huge that nothing can escape, not even light! The point of no return is called the event horizon. The fact that light cannot escape is significant. It means that no signal will ever be able to reach us from the singularity. In other words, it can never be observed. Our notion of time and space break down approaching the singularity, and we can no longer use our theory to make predictions. Therefore, it was postulated that a ‘naked’ singularity can never exist - there must always be an event horizon protecting it from view.

For the black holes we see in nature, this is always true. However, if we look at charged black holes, we find this is no longer the case. These have two event horizons; an internal Cauchy horizon, at which determinism breaks down, and the regular event horizon. As the charge of the black hole increases, the event horizons get closer together, and when they reach each other they actually disappear! Since this is forbidden, we get a limit on the amount of charge a black hole can store. A black hole at this limit is an extremal black hole.

Although nothing can escape a black hole, Stephen Hawking worked out that they can still decay via a process known as Hawking radiation. We need to make sure that an extremal black hole can still decay, without being pushed over the extremal limit as it loses mass. This means that there exists at least one particle where its charge is greater than its mass to allow the black hole to decay. There is at least one particle where gravity acts on it as the weakest force. This is known as the weak gravity conjecture.

Consider the other end of the extreme, the quantum. In this theory, energy takes only discrete values. This is in contrast to the smooth continuum in general relativity. It is well known that these theories are incompatible. When we try and promote general relativity to a quantum field theory, there are an infinite number of terms we must add to make these theories compatible. However, when we work in the low energy regime, only the first few are significant. This is similar to how we can ignore the interactions of atoms when we are looking at the behaviour of waves in the ocean. This is our low energy effective field theory of quantum gravity. Our project looks at how the behaviour of a charged black hole changes in this landscape, and looks at how the weak gravity conjecture is affected.

We find that the weak gravity conjecture can be automatically satisfied, but this is dependent on a coefficient in our low energy approximation. Ripples can travel through spacetime, like the surface of a parachute, these are gravitational waves. We found the speed of these waves around the extremal event horizons. Since nothing can travel faster than light, we wanted to see if the speed depended on the coefficient, and if we could therefore determine its sign. However, we found that the speed was always equal to the speed of light.

By Nadia Cooper

Lay Summary: Probing Black Holes in an Atomic Spacetime

At the beginning of XXth century, Einstein revolutionised theoretical physics by developing the theory of General Relativity. First, he unified space and time into a single substance, spacetime. Then, he explained gravity as the curvature of said spacetime, in turn following from the presence of mass: **mass tells spacetime how to curve, spacetime tells mass how to move**. Relativity predicted the existence of black holes, regions of spacetime where – you have probably already heard this – gravity is so strong that nothing, not even light, can escape. So black holes are heavy ... why do we care that much?

The first reason is that their centre is *infinitely heavy*! This is known as the singularity and has infinite density and curvature. It seemingly rips spacetime apart, breaking the trajectory of an object in space and time as if it existed no more. The other reason is that black holes are a unique combination of the main branches of theoretical physics. In fact, by combining Relativity and Quantum Mechanics, Hawking showed black holes are thermal objects with thermodynamic properties. He showed **they are indeed perfectly black, yet not perfectly holes**: some black body radiation can escape them. They can slowly deflate until disappear, making them cosmic erasers of information. In fact, imagine a black hole disappears: as far as we know, the left radiation only retains information about mass, charge and angular momentum. All other information about the absorbed objects vanishes. This has been labeled as the information paradox, since, loosely, this breaks some fundamental assumptions of modern theoretical physics ... quite a big deal.

Giving a coherent description of black holes falls within the problem of Quantum Gravity, i.e. of unifying Relativity and Quantum Mechanics and solving the yet unresolved mysteries of the universe. Our work was done in the context of the Causal Sets Theory (CST) approach. This assumes **spacetime is not continuous, but rather a network of events** randomly distributed and connected between them. It is not much different than the paper or screen you are reading this on: it *looks* continuous but is discrete, made of atoms. The distance between the "atoms of spacetime" would be of order of the Planck length: 10^{-35}m, i.e. as many times smaller than a proton, as a proton is smaller than the Earth!

We studied the surface of the black hole, known as the horizon. Anything crossing the horizon is doomed to fall into its centre. As black holes are thermal objects, and as temperature is usually associated with some molecules' distribution, following some previous proposals, we hypothesised **the horizon could be made of "horizon molecules"**. The horizon is just a region of spacetime, then, in our picture of spacetime as a network, these molecules are certain links between spacetime atoms crossing the horizon.

Our aim was to show that the horizon molecules could endow black holes with their known properties. We analysed the model both analytically and numerically. For the latter, we derived our algorithm and wrote our own code, available at `https://github.com/vidh2000/MSci_Schwarzschild_Causets`. This produced the very first black hole simulation in CST large enough for analysis to be made – single simulations used up to 750,000 points and 2 TB of RAM – and could represent an important tool for other investigations.

Our results showed that horizon molecules give the correct thermal properties and are indeed localised on the horizon. On the speculative side, horizon molecules could perhaps influence the outgoing radiation, with consequences on the information paradox. Alternatively, there are some hints that they could even prevent the disappearance of black holes. Further studies are required.

by Stefano Veroni