CARDIFF UNIVERSITY

On the search for intermediate duration gravitational waves using the spherical harmonic basis

by

Mark Edwards

A thesis submitted in partial fulfillment for the
degree of Doctor of Philosophy

in the
Gravitational Physics
School of Physics and Astronomy

December 2013

Declaration of Authorship

I, MARK EDWARDS, declare that this thesis titled, 'On the search for intermediate duration gravitational waves' and the work presented in it are my own. I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this University.
- Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work.
- I have acknowledged all main sources of help.
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself.

Signed: _____

Date: _____

Science - Its about cracking open the egg of opportunity, and scooping out the delights within.

Me

CARDIFF UNIVERSITY

Abstract

Gravitational Physics
School of Physics and Astronomy

Doctor of Philosophy

by Mark Edwards

The main focus of the work described here within, is the development of an analysis pipeline that can detect gravitational waves transients that have a duration from a large fraction of a second up to 100 seconds.

Firstly, we will present the basic theory of gravitational waves, their generation and their detection. We will then review the current state of the art of gravitational waves transient search algorithms and their limitations.

We also introduce the coherent search pipeline that has been developed from the ground up to find such intermediate length gravitational waves, and discuss in detail how it differs from other pipelines, and what makes it so well suited to this task. Further, we will describe the powerful glitch rejection algorithm and sky localisation solution that using the spherical harmonic basis made possible.

Finally, we demonstrate the use of this method on real data, and compare our results against a mature pipeline.

Acknowledgements

Everyone. There have been some hard times during the writing of this thesis, especially during the last 12 months, and I would especially like to thank Patrick Sutton for all of the understanding extended to me. James Clark, Valeriu Predoi and Ioannis Kamaretsos get special mention due to making the office a fun and fantastic place to work in, while Steve Fairhurst and Mark Hannam get bumped for being patient during the Amaldi months!

Everyone in the group has been helpful and knowledgable, and it has been a great pleasure to work with them all.

An extra shout goes out to my parents, who, even now, don't really understand why I gave up a well paid job to become a student again!

Latestly and most definitely not leasty, Helen my beautiful wife and Isaac my boisterous son. Both of which have given me support and cheered me along all the way through (but are secretly glad to see the back of it!)

Contents

List of Figures

List of Tables

Dedicated to my father, who passed away before we were ready to say goodbye.

Chapter 1

Gravitational Waves

"There is nothing new to be discovered in physics now. All that remains is more and more precise measurement."

– Unknown, *attributed to Lord Kelvin*

The use of the electromagnetic spectrum in Astronomy continues to reveal astonishing things about the universe (extrasolar planets, dark matter mapping), and will continue to do so for some time. However, it may be that a new way of viewing the universe is necessary to further our knowledge. There are other ways of imaging the universe such as neutrino counting, but the most likely candidate to present a new astrophysical window is the detection of gravitational waves. In this chapter we will provide a brief overview of linearised gravity and gravitational waves (for the full derivations see any good book on general relativity, for example [3], [4] or [5]).

1.1 The Einstein Field equations

The Einstein Field equations describe the relationship between matter and the geometry of spacetime

$$G_{\alpha\beta} = R_{\alpha\beta} - \frac{1}{2}Rg_{\alpha\beta} = \frac{8\pi G}{c^4}T_{\alpha\beta}. \tag{1.1}$$

In the Einstein Field equations, the stress-energy tensor, $\boldsymbol{T}$, is used to describe gravitational sources while the Einstein Tensor, $\boldsymbol{G}$, is used to describe the curvature of spacetime. Equation (1.1) comprises ten, non-linear partial differential equations, one for each of the ten independent metric coefficients, $g_{\alpha\beta}(x)$, for which there is no general solution. However, it is possible to form a complete analysis of the solutions for the Einstein equation for spacetimes that vary only slightly from flat spacetime. Details of which are given below.

Please note that from now on we will be using geometrized units where every time interval can be interpreted as the distance travelled by light during that time interval, so that G = c = 1.

1.1.1 The metric, $g_{\alpha\beta}$

The metric, $g_{\alpha\beta}$, produces the separation between infinitesimally separated events in spacetime

$$ds^2 = g_{\alpha\beta}dx^\alpha dx^\beta. \tag{1.2}$$

More specifically, it produces the scalar product of two vectors u^α, v^α, or the squared length of a single, u^α

$$(u, v) = g_{\alpha\beta}u^\alpha v^\beta, \tag{1.3}$$
$$(u, u) = g_{\alpha\beta}u^\alpha u^\beta. \tag{1.4}$$

There are 3 possible outcomes of the separation value:

"timelike" $\Rightarrow$ negative squared length (two events are separated in time)

"spacelike" $\Rightarrow$ positive squared length (two events are separated in space)

"null" $\Rightarrow$ zero squared length (two events lie on a null line or are connected by a light ray)

Note that the metric can be used to raise and lower indices

$$g^{\alpha\beta}g_{\beta\gamma} = \delta^\alpha_\gamma, \tag{1.5}$$

where $g^{\alpha\beta}$ are the components of the inverse metric.

We will be using the convention of $g_{\alpha\beta} = \text{diag}(-1, +1, +1, +1)$, so in flat spacetime the metric is as follows

$$g_{\alpha\beta} = \eta = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

1.1.2 The Riemann tensor, $R^\alpha_{\beta\gamma\delta}$

The Riemann tensor, $R^\alpha{}_{\beta\gamma\delta}$, can be thought of as a measure of curvature, and can be defined by

$$R^\alpha{}_{\beta\gamma\delta} = \frac{\partial\Gamma^\alpha_{\beta\delta}}{\partial x^\gamma} - \frac{\partial\Gamma^\alpha_{\beta\gamma}}{\partial x^\delta} + \Gamma^\alpha_{\mu\delta}\Gamma^\mu_{\beta\delta} - \Gamma^\alpha_{\mu\delta}\Gamma^\mu_{\beta\gamma}, \tag{1.6}$$

where $\Gamma^{\alpha}_{\beta\gamma}$ are the Christoffel symbols, which, in a coordinate basis are

$$\Gamma_{\mu\beta\gamma} = \frac{1}{2}(g_{\mu\beta,\gamma} + g_{\mu\gamma,\beta} - g_{\beta\gamma,\mu}). \tag{1.7}$$

There exist differential symmetries called the Bianchi identities that correspond to $R^{\alpha}_{\beta[\lambda\mu;\nu]} = 0$, where [] is the antisymmetric operation defined as $R^{\alpha}_{\beta\lambda\mu;\nu} + R^{\alpha}_{\beta\mu\nu;\lambda} + R^{\alpha}_{\beta\nu\lambda;\mu} = 0$. We can form the Ricci curvature tensor from the Riemann tensor via contraction, which (in a coordinate frame) is

$$R_{\mu\nu} \equiv R^{\alpha}{}_{\mu\alpha\nu} = \Gamma^{\alpha}_{\mu\nu,\alpha} - \Gamma^{\alpha}_{\mu\alpha,\nu} + \Gamma^{\alpha}_{\beta\alpha}\Gamma^{\beta}_{\mu\nu} - \Gamma^{\alpha}_{\beta\nu}\Gamma^{\beta}_{\mu\alpha}. \tag{1.8}$$

The scalar curvature is formed by contraction of the Ricci tensor

$$R \equiv R^{\mu}_{\mu} = g^{\mu\nu}R_{\mu\nu}. \tag{1.9}$$

1.1.3 Einstein curvature tensor, $G_{\alpha\beta}$

It can be seen from equation (1.1) that the Einstein tensor, $G_{\alpha\beta}$, is a piece of the Riemann tensor. Specifically,

$$G^{\mu}_{\nu} \equiv \frac{1}{2}\epsilon^{\mu\alpha\beta\gamma}R_{\beta\gamma}{}^{\rho\sigma}\frac{1}{2}\epsilon_{\nu\alpha\rho\sigma} = R^{\mu}_{\nu} - \frac{1}{2}\delta^{\mu}_{\nu}R = g^{\mu\alpha}(R_{\alpha\nu} - \frac{1}{2}g_{\alpha\nu}R) \tag{1.10}$$

where ϵ is the Levi-Civita tensor (see [3]), which is defined as

$$\begin{aligned}\epsilon^{\alpha\beta\gamma\delta} &= (-g)^{1/2}[\alpha\beta\gamma\delta],\\ \epsilon^{\alpha\beta\gamma\delta} &= g^{-1}\epsilon\alpha\beta\gamma\delta[\alpha\beta\gamma\delta]\\ &= -(-g)^{-1/2}[\alpha\beta\gamma\delta],\end{aligned} \tag{1.11}$$

where $[\alpha\beta\gamma\delta]$ is the complete antisymmetric symbol

$$[\alpha\beta\gamma\delta] \equiv \begin{cases} +1 \text{ if } \alpha\beta\gamma\delta \text{ is an even permutation of 0123,}\\ \text{-1 if } \alpha\beta\gamma\delta \text{ is an odd permutation of 0123,}\\ 0 \text{ if } \alpha\beta\gamma\delta \text{ are not all different,}\end{cases}$$

and g is the determinant of the matrix, $||g_{\alpha\beta}||$.

Also, the curvature, $G_{\alpha\beta}$, is generated by the stress-energy tensor, $T_{\alpha\beta}$. From (1.1):

$$G_{\alpha\beta} = 8\pi T_{\alpha\beta}. \tag{1.12}$$

This states that the curvature of spacetime, $G_{\alpha\beta}$, is defined by the stress-energy tensor of the matter, $T_{\alpha\beta}$. This includes the external spacetime curvature of a static source and the generation of gravitational waves by time-varying stress-energy distributions.

1.1.4 The stress-energy tensor, $T_{\alpha\beta}$

The stress-energy tensor is the source term for the field equations, representing the amount of energy, momentum, pressure and stress in the spacetime. Roughly, it corresponds to:

$$T = \begin{pmatrix} u & \rho_x & \rho_y & \rho_z \\ p_x & P_{xx} & \sigma_{xy} & \sigma_{xz} \\ p_y & \sigma_{yx} & P_{yy} & \sigma_{yz} \\ p_z & \sigma_{zx} & \sigma_{zy} & P_{zz} \end{pmatrix}.$$

The components of the stress-energy tensor relate to different aspects of the space-time they are describing. For instance:

T^{00} **(u):** This element represents the density of mass-energy as measured by an observer in the Lorentz frame (if we assume $g_{00} = -1$, then the following relations hold, $T_{00} = -T^0_0 = T^{00}$).

T^{j0} **(p_j):** These elements represent the density of energy in the three spatial coordinates- i.e. the momentum density.

T^{0j} **(ρ_j):** These elements represent the flow of energy in the three spatial coordinates i.e the energy flux.

T^{jk} **(P_{jj} & σ_{jk}):** Stress tensor from classical mechanics. These coefficients relate to the flow of spatial momentum for each of the spatial coordinates (i.e. the flux of the linear momentum), specifically:

T^{jj} **(P_{jj}):** these are the components of normal stress, i.e. that applied perpendicular to the region (pressure).

T^{jk} **(σ_{jk}, $j \neq k$):** these are the components of sheer stress, i.e. that applied tangentially to the region.

Note that T^{j0} = momentum density = T^{0j} = energy flux (due to the mass-energy relationship, $E = mc^2$), and more generally, $T^{\alpha\beta} = T^{\beta\alpha}$. The indices can be raised and lowered using the metric, as follows:

$$T_\alpha{}^\beta = g^{\beta\gamma} T_{\alpha\gamma} = g_{\alpha\gamma} T^{\gamma\beta}. \tag{1.13}$$

1.2 Linearised theory of gravity

If we apply the Einstein field equations (1.1) to a weak-field situation (i.e. a long distance from a gravitational source), we can define the spacetime as a Minkowski (flat) background spacetime plus a very small perturbation. The metric for flat spacetime is $g_{\mu\nu} = \eta_{\mu\nu}$ for Minkowski coordinates (t, x, y, z), where $\eta_{\mu\nu} = diag(-1, 1, 1, 1)$. Small perturbations to the geometry of this spacetime can be written as:

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu} \qquad\qquad |h_{\mu\nu}| << 1. \tag{1.14}$$

This approximation is valid, for example, in the solar system ($|h^{\mu\nu}| \sim |\Phi| < \frac{M_\odot}{R_\odot} \sim 10^{-6}$), or for a weak gravitational wave propagating through space. We can expand the field equations in powers of $h_{\mu\nu}$ in a coordinate system where (1.14) holds, and, without losing accuracy, only keep the linear terms (since $h_{\mu\nu}$ is very small). So, using (1.14) as the metric components the resulting connection coefficients, (1.7), when linearised, are:

$$\Gamma^{\mu}_{\alpha\beta} = \frac{1}{2}\eta^{\mu\nu}(h_{\alpha\mu,\beta} + h_{\beta\mu,\alpha} - h_{\alpha\beta,\mu}) \equiv \frac{1}{2}(h_{\alpha}{}^{\mu}{}_{,\beta} + h_{\beta}{}^{\mu}{}_{,\alpha} - h_{\alpha\beta}{}^{,\mu}). \tag{1.15}$$

Due to the approximation of a flat background spacetime, powers of $h_{\mu\nu}$ can have their indices raised and lowered using $\eta^{\mu\nu}$ and $\eta_{\mu\nu}$, rather than $g^{\mu\nu}$ and $g_{\mu\nu}$. The Riemann tensor becomes:

$$R_{\alpha\beta\mu\upsilon} = \frac{1}{2}(h_{\alpha\upsilon,\beta\mu} + h_{\beta\mu,\alpha\upsilon} - h_{\alpha\mu,\beta\upsilon} - h_{\beta\upsilon,\alpha\mu}) \tag{1.16}$$

which, to first order, is unchanged under small (gauge) transformations. The Ricci tensor, (1.8), now becomes

$$R_{\mu\nu} = \Gamma^{\alpha}_{\mu\nu,\alpha} - \Gamma^{\alpha}_{\mu\alpha,\nu} = \frac{1}{2}(h_{\mu}{}^{\alpha}{}_{,\nu\alpha} + h_{\nu}{}^{\alpha}{}_{,\mu\alpha} - h_{\mu\nu,\alpha}{}^{\alpha} - h_{,\mu\nu}), \tag{1.17}$$

where

$$h \equiv h^{\alpha}_{\alpha} = \eta^{\alpha\beta}h_{\alpha\beta}. \tag{1.18}$$

Contracting once again, we get the Ricci scalar, $R \equiv g^{\mu\nu}R_{\mu\nu} \approx \eta^{\mu\nu}R_{\mu\nu}$, which, when substituted into the Einstein field equation, (1.12) using the identity (1.10), we get

$$\begin{aligned} G_{\mu\nu} =& R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R && = 8\pi T_{\mu\nu} \\ =& \frac{1}{2}(h_{\mu\alpha,\nu}{}^{\alpha} + h^{\alpha}_{\nu\alpha,\mu} - h_{\mu\nu,\alpha}{}^{\alpha} - h_{,\mu\nu}) - \eta_{\mu\nu}(h_{\alpha\beta}{}^{,\alpha\beta} - h_{,\beta}{}^{\beta}) && = 8\pi T_{\mu\nu}. \end{aligned} \tag{1.19}$$

To reduce the number of terms, we define the trace-reversed metric perturbation as $\bar{h}_{\mu\nu} \equiv h_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu}h$, using a bar to imply a corresponding operation of any other symmetrical tensor. Thus, $G_{\mu\nu} = \bar{R}_{\mu\nu}$ to first order in $h_{\mu\nu}$, and $\bar{\bar{h}}_{\mu\nu} = h_{\mu\nu}$ ($h_{\mu\nu} = \bar{h}_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu}\bar{h}$). Therefore (see [3]

for more detail),

$$\underbrace{-\bar{h}_{\mu\nu,\alpha}{}^{\alpha}}_{\text{d'Alembertian}} \underbrace{-\eta_{\mu\nu}\bar{h}_{\alpha\beta,}{}^{\alpha\beta} + \bar{h}_{\mu\alpha,}{}^{\alpha}{}_{\nu} + \bar{h}_{\nu\alpha,}{}^{\alpha}{}_{\mu}}_{\text{gauge invariance}} = 16\pi T_{\mu\nu}, \tag{1.20}$$

which demonstrates that the components of $\bar{h}_{\mu\nu}$ obey wave equations. Nearly globally Lorentz systems can be connected to each other by two different types of coordinate transformations:

1. Global Lorentz transformations, $x^{\mu} = \Lambda^{\mu}_{\alpha'} x^{\alpha'}$, $\Lambda^{\mu}_{\alpha'}\Lambda^{\nu}_{\beta'}\eta_{\mu\nu} = \eta_{\alpha'\beta'}$:

   Which transform the metric via

   $$\begin{aligned}\eta_{\alpha'\beta'} + h_{\alpha'\beta'} = g_{\alpha'\beta'} &= \frac{\partial x^{\mu}}{\partial x^{\alpha'}}\frac{\partial x^{\nu}}{\partial x^{\beta'}} g_{\mu\nu} = \Lambda^{\mu}_{\alpha'}\Lambda^{\nu}_{\beta'}(\eta_{\mu\nu} + h_{\mu\nu}) \\ &= \eta_{\alpha'\beta'} + \Lambda^{\mu}_{\alpha'}\Lambda^{\nu}_{\beta'} h_{\mu\nu}.\end{aligned}$$

   Thus, $h_{\mu\nu}$ and $\bar{h}_{\mu\nu}$ transform like the components of a tensor in flat space-time

   $$h_{\alpha'\beta'} = \Lambda^{\mu}_{\alpha'}\Lambda^{\nu}_{\beta'} h_{\mu\nu}.$$

2. Infinitesimal coordinate transformations, e.g. 'ripples' in the coordinate system:

   $$x^{\mu'}(P) = x^{\mu}(P) + \epsilon^{\mu}(P), \tag{1.21}$$

   Where $\epsilon^{\mu(P)} \equiv 4$ arbitrary functions small enough to leave $|h_{\mu'\nu'}| \ll 1$. These tiny changes can be ignored in all quantities but the metric, where the (slowly varying) deviations from $\eta_{\mu\nu}$ contain all of the information about gravity. Using the usual transformation law for the metric

   $$g_{\rho'\sigma'} = g_{\mu\nu}\frac{\partial x^{\mu}}{\partial x^{\rho'}}\frac{\partial x^{\nu}}{\partial x^{\sigma'}},$$

   and combining with (1.21), and $g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}$, we can show that

   $$\begin{aligned}g_{\rho'\sigma'} = \eta_{\rho\sigma} + h_{\rho\sigma} - \epsilon_{\rho,\sigma} - \epsilon_{\sigma,\rho} \\ + (\text{negligible corrections } \sim h_{\rho\sigma,\alpha}\epsilon^{\alpha} \text{ and } h_{\rho\alpha}\epsilon^{\alpha}_{,\sigma'}).\end{aligned} \tag{1.22}$$

   Hence, the metric perturbation functions in the new ($x^{\mu'}$) and old (x^{μ}) coordinate systems are related by

   $$h^{new}_{\mu\nu} = h^{old}_{\mu\nu} - \epsilon_{\mu,\nu} - \epsilon_{\nu,\mu}. \tag{1.23}$$

   The functional forms of all other scalars, vectors and tensors are unaltered within the precision of linearised theory. This is a gauge transform, and the fact that it does not

affect the functional form of observables is called 'gauge invariance'. One can specify this gauge (i.e. the coordinates) such that $\bar{h}^{\mu\alpha}{}_{,\alpha} = 0$. This is called the Lorenz gauge, and is not uniquely fixed, but rather the gauge condition is left unaffected by any further gauge transformation for which $\epsilon^{\alpha,\beta}{}_{\beta} = 0$.

Once the gauge has been fixed, $h_{\mu\nu}$ and $\bar{h}_{\mu\nu}$ can be regarded as components of tensors in flat space-time and the field equations and the chosen gauge conditions as geometric, coordinate-independent equations in flat space-time. When using curvilinear coordinates $\eta_{\mu\nu}$ must be replaced by $g^{flat}_{\mu\nu}$, and ordinary derivatives by covariant derivatives whose connection coefficients come from $g^{flat}_{\mu\nu}$. By imposing the Lorenz gauge condition, $\bar{h}^{\mu\alpha}{}_{,\alpha} = 0$, to equation (1.20), they become

$$\bar{h}_{\mu\nu,\alpha}{}^{\alpha} = -16\pi T_{\mu\nu}, \tag{1.24}$$

where the other terms have vanished due to our choice of gauge.

The metric is defined as

$$g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu} = \eta_{\mu\nu} + \bar{h}_{\mu\nu} - \frac{1}{2}\eta_{\mu\nu}\bar{h}. \tag{1.25}$$

Equations (1.24) and (1.25) are the fundamental equations of linearised gravity in the Lorenz gauge. Note that equation (1.24) is simply the ordinary wave equation in flat spacetime. The general solution is in terms of a retarded integral over the stress tensor

$$\bar{h}_{\mu\nu}(t,x) = \int \frac{4T_{\mu\nu}(t - |\underline{x} - \underline{x}'|, \underline{x}')}{|\underline{x} - \underline{x}'|} d^3x'. \tag{1.26}$$

If we focus on nearly Newtonian sources (i.e. $T_{00} \gg |T_{0j}|, T_{00} \gg |T_{jk}|$) and velocities slow enough such that any retardation is negligible (i.e. a static solution), then

$$\bar{h}_{00} = -4\Phi, \tag{1.27}$$

$$\bar{h}_{0j} = \bar{h}_{jk} = 0, \tag{1.28}$$

$$\Phi(t,x) = -\int \frac{T_{00}(t,\underline{x}')}{|\underline{x} - \underline{x}'|} d^3x' \qquad \text{(Newtonian potential)}. \tag{1.29}$$

Then the metric far from the source is given by

$$ds^2 = -(1+2\Phi)dt^2 + (1-2\Phi)(dx^2 + dy^2 + dz^2) \tag{1.30}$$

$$\approx -(1-2M/r)dt^2 + (1-2M/r)(dx^2 + dy^2 + dz^2). \tag{1.31}$$

In the more general case of a source with weak internal gravitational fields and slow motion (i.e. $v \ll 1$), then the metric perturbation far from the source is

$$\bar{h}^{ij}(t,\ \underline{x}) = \frac{2}{r}\frac{d^2}{dt^2} I^{ij}(t-r), \tag{1.32}$$

where $r \equiv |\ \underline{x}|$ and I^{ij} is the second mass moment of the source:

$$I^{ij}(t) = \int d^3\ \underline{x'} T_{00}(t, \underline{x'}) x'^i x'^j. \tag{1.33}$$

Note that it is possible to make an order of magnitude estimate for a binary system as the source using this equation. The quadrupole moment $\propto MR^2$ and Kepler's law states that $\Omega^2 = M/R^3$, so, very approximately

$$|h| \equiv \varepsilon \frac{M}{r} \frac{M}{R}, \tag{1.34}$$

where ε measures the asymmetry, M is the total mass, R is the size of the source, and r is the distance to the source.

For example, an equal mass binary system, with components of total mass $2M_\odot$, at a distance of 1 Mpc , with the binary components separated by 100 km produces a gravitational wave with an amplitude of $|h| = \varepsilon \times 10^{-22}$. The gravitational waves carry energy away from the system which will cause the compact objects to loose angular momentum and eventually spiral into each other.

1.2.1 Plane wave solution

Outside of any source, $T_{\mu\nu} = 0$, and the linearised Einstein equations becomes $\bar{h}_{\mu\nu;\alpha}{}^{\alpha} = 0$, with the simplest solution being a monochromatic plane-wave

$$\bar{h}_{\mu\nu} = \Re\{A_{\mu\nu} \mathrm{e}^{iK_\alpha x^\alpha}\}. \tag{1.35}$$

Here, $A_{\mu\nu}$ is a tensor describing the wave amplitude and K_μ is the wave vector. These are constants satisfying

$$\begin{aligned} K_\alpha K^\alpha &= 0 && (K^\alpha \text{ is a null vector}) \\ A_{\mu\alpha} K^\alpha &= 0 && (A_{\mu\nu} \text{ is orthogonal to } K_\alpha). \end{aligned}$$

The wave has a frequency $w = K^0 = (K_x^2 + K_y^2 + K_z^2)^{1/2}$, and propagates at the speed of light. At first glance, the symmetrical tensor $A_{\mu\nu}$ appears to have 10 components, but the 4 orthogonality constraints ($A_{\mu\alpha}K^\alpha = 0$) reduce this to 6. There is also an arbitrariness due to the gauge freedom. A further gauge transformation (1.23) with $\epsilon^\mu = -iC^\mu \mathrm{e}^{iK_\alpha x^\alpha}$ where C^μ are 4 arbitrary constants, can be used to fix 4 of the 6 components of $A_{\mu\nu}$.

1.2.2 The transverse traceless (TT) gauge

Select a 4-velocity, u^α, for the whole of spacetime and using a specific gauge transform, impose $A_{\mu\alpha}u^\alpha = 0$. This gives 3 constraints, as $K^\mu A_{\mu\nu}u^\nu = 0$ is already satisfied. Then use $A^\alpha{}_\alpha = 0$. This gives 8 constraints, $A_{\mu\alpha}u^\alpha = 0, A_{\mu\alpha}K^\alpha = 0, A_\alpha{}^\alpha = 0$, so the coordinate system (i.e the gauge) is now rigidly fixed. The 2 remaining degrees of freedom represent the 2 polarizations. Restating the constraints in a Lorentz frame ($u^0 = 1, u^j = 0$) in terms of $h_{\mu\nu}$ directly gives

$$h_{\mu 0} = 0 \qquad \text{only spatial coordinates } h_{jk} \text{ are non-zero} \tag{1.36}$$

$$h_{kj,j} = 0 \qquad \text{spatial coordinates are divergent free} \tag{1.37}$$

$$h_{kk} = 0 \qquad \text{spatial coordinates are trace free} \tag{1.38}$$

Note that $h = h_\mu{}^\mu = h_{kk} = 0$, no there is no distinction between $h_{\mu\nu}$ and $\bar{h}_{\mu\nu}$. Any tensor satisfying these constraints is called a transverse-traceless (TT) tensor. It is transverse because it is purely spatial ($h_{0\mu} = 0$), and (if a wave) transverse to it own direction of propagation ($h_{ij,j} = h_{ij}K_j = 0$), and traceless because $h_{kk} = 0$. Note that any purely spatial tensor, S_{ij}, can be decomposed (see [3]) into a TT part S^{TT}_{ij}, a part $S^T_{ij} = \frac{1}{2}(\delta_{ij}f_{,kk} - f_{,ij})$ which is "transverse" ($S^T_{ij,j} = 0$) but determined entirely by one function, f, giving the trace of S ($S^T_{kk} = \nabla^2 f$), and a part $S^L_{ij} = S^L_{i,j} + S^L_{j,i}$ which is longitudinal and determined by a vector field, S^L_i.

If we now consider an arbitrary gravitational wave, it can be defined as a superposition of plane-waves, so, for each, we can introduce the TT gauge (which is linear) as follows: Pick a specific global Lorentz frame (i.e. choice of u^α); in this frame (where $u^\alpha = \delta^\alpha{}_0$) examine a gravitational wave of arbitrary form and find a gauge which satisfies the TT gauge. Since only $h_{jk} \neq 0$, impose $\Box h_{jk} = h_{jk,\alpha}{}^\alpha = 0$.

1.2.3 The transverse traceless gauge: General case

We note that the formula (1.26) for the general solution for the linearised Einstein equations (1.24) does not generally yield a result in the TT gauge. Given a gravitational wave $h_{\mu\nu}(t, x_j)$ in an arbitrary gauge, we can compute $h^{TT}_{\mu\nu}(t, x^i)$ as follows.

1) [Valid only for waves with $\Box\bar{h}_{\mu\nu} = 0$]

Calculate R_{j0k0} in the initial gauge, then integrate

$$h^{TT}_{jk,00} = -2R_{j0k0}, \tag{1.39}$$

to obtain h^{TT}_{jk}. If the wave is monochromatic, then $h_{\mu\nu} = h_{\mu\nu}(x^l)\exp^{-i\omega t}$ and the above has the simple form

$$h^{TT}_{jk} = 2\omega^{-2}R_{j0k0}. \tag{1.40}$$

2) [Valid only for plane waves]

Project the TT components in an algebraic manner

$$P_{jk} = \delta_{jk} - n_j n_k, \tag{1.41}$$

where $n_k = K_k/|\underline{K}|$ is the unit vector in the direction of propagation. We can verify that P_{jk} is the projection operator onto the transverse plane by checking that $P_{jl}P_{lk} = P_{jk}$, $P_{jk}n_k = 0$ and $P_{jj} = 2$. In that case, the transverse part of h_{jk} is $P_{jl}h_{lm}P_{mk}$ (PhP in matrix form), and the TT part is this quantity diminished by its trace

$$h^{TT}_{jk} = P_{jl}P_{mk}h_{lm} - \frac{1}{2}P_{jk}(P_{ml}h_{lm}) \qquad \Leftarrow \text{index notation} \tag{1.42}$$

$$h^{TT} = PhP - \frac{1}{2}P\,Tr(Ph) \qquad \Leftarrow \text{matrix notation.} \tag{1.43}$$

This actually removes 2 parts from h_{ij}. The first is

$$h^{T}_{jk} = \frac{1}{2}P_{jk}(P_{lm}h_{lm}), \tag{1.44}$$

which is transverse, but built from its own trace

$$h^T = Tr(PhP) = Tr(Ph) = P_{lm}h_{ml}, \tag{1.45}$$

while the second is the longitudinal part

$$\begin{aligned} h^L_{jk} &= h_{jk} - P_{jl}P_{mk}h_{lm} \\ &= n_l n_k h_{jl} + n_j n_l h_{lk} - n_j n_k (n_l n_m h_{lm}) \qquad \text{[index notation]} \end{aligned} \tag{1.46}$$

$$h^L = h - PhP. \qquad \text{[matrix notation]} \tag{1.47}$$

3) [General case]

Fourier analyse any symmetrical array,

$$h_{ij} = \int h_{ij}(K,t)\exp^{iK_m x^m}\, d^3K,$$

and apply equation (1.43) to each component separately. Written in direction independent form, the projection operator becomes

$$P_{jk} = \delta_{jk} - \frac{1}{\nabla^2}\partial_j\partial_k, \tag{1.48}$$

or

$$n_l n_m = \frac{1}{\nabla^2}\partial_l\partial_m. \tag{1.49}$$

Now, since $\partial_l = iK_l$ under the Fourier integral, $\frac{1}{\nabla^{-2}}$ can be evaluated by Green's functions as well as Fourier analysis ($\psi = \nabla^{-2}f$ is a solution of the Poisson equation $\nabla^2\psi = f$)

which has advantages in analytical calculations.

In linearised theory, h^L_{ij} is the purely gauge part of $h_{\mu\nu}$, whereas h^T_{ij} and h^{TT}_{ij} are gauge invariant. The special gauge which reduces $h_{\mu\nu}$ to its TT part is, as stated, called the TT gauge. The condition defining it is $h_{\mu\nu} = h^{TT}_{\mu\nu}$. Only "pure" waves can be reduced to this gauge (i.e. not more general solution of the linearised equations with $\Box\bar{h}_{\mu\nu} = -16\pi T_{\mu\nu}$. Under the TT gauge, the Riemann tensor space-time components ($R_{j0k0} = R_{0j0k} = -R_{j00k} = -R_{0jk0}$) have a simple form

$$R_{j0k0} = -\frac{1}{2}h^{TT}_{jk,00}. \tag{1.50}$$

The change in $h_{\mu\nu}$ due to the gauge transformation is

$$\delta h_{\mu\nu} = -(\partial_\nu \xi_\mu + \partial_\mu \xi_\nu). \tag{1.51}$$

The transverse part of this change is

$$P_{\mu l}P_{m\nu}(\delta h_{lm}) = -P_{\mu l}P_{m\nu}(\partial_l \xi_m + \partial_m \xi_l) = 0. \tag{1.52}$$

We note that for plane wave

$$\partial_l = i|\underline{K}|n_l, \tag{1.53}$$

$$P_{jl}n_l = 0, \tag{1.54}$$

and generally,

$$P_{jk} = \delta_{jk} - \frac{1}{\nabla^2}\partial_j\partial_k, \tag{1.55}$$

which gives $P_{jl}\partial_l = 0$. Therefore, h^{TT}_{ij} and h^T_{ij} of equations (1.42), (1.44) are gauge invariant: $\delta h^{TT}_{ij} = \delta h^T_{ij} = 0$. In empty space ($T_{\mu\nu} = 0$), h^T_{ij} vanishes by virtue of the field equations.

1.2.4 Polarisation of plane waves

Consider a plane gravitational wave propagating in the z-direction. In the TT gauge, the constraints are

$$\begin{aligned} h^{TT}_{0\mu} &= 0, \\ h^{TT}_{ij,j} &\equiv iK_j h^{TT}_{ij} = 0, \\ h^{TT}_{kk} &= 0, \end{aligned}$$

showing that the only non-vanishing components are

$$h_{xx}^{TT} = -h_{yy}^{TT} = \Re\{A_+ \exp^{-i\omega(t-z)}\},$$
$$h_{xy}^{TT} = h_{yx}^{TT} = \Re\{A_\times \exp^{-i\omega(t-z)}\},$$

giving two modes of polarisation. A given wave can therefore be resolved into two linearly polarised components. The basis tensors for these two components may be written as

$$\underline{e}_+ \equiv \underline{e}_x \otimes \underline{e}_x - \underline{e}_y \oplus \underline{e}_y; \tag{1.56}$$
$$\underline{e}_\times \equiv \underline{e}_x \otimes \underline{e}_y + \underline{e}_y \oplus \underline{e}_x, \tag{1.57}$$

where $\underline{e}_x$, $\underline{e}_y$ are orthonormal vectors transverse to the direction of propagation ($\underline{e}_z$). Alternatively, we may select non-orthogonal circularly polarised basis vectors

$$\underline{e}_{\text{right}} \equiv \frac{1}{\sqrt{2}}(\underline{e}_+ + i\underline{e}_\times); \tag{1.58}$$
$$\underline{e}_{\text{left}} \equiv \frac{1}{\sqrt{2}}(\underline{e}_+ - i\underline{e}_\times). \tag{1.59}$$

1.2.5 Stress-energy carried by a gravitational wave

As shown in [3], the stress-energy carried by a gravitational wave cannot be localised inside a wavelength, it can only be discussed within the terms of a larger region. In a, nearly, inertial frame, we can write

$$T_{\mu\nu}^{GW} = \frac{1}{32\pi} < h_{jk,\mu}^{TT} h_{jk,\nu}^{TT} >, \tag{1.60}$$

where the angle brackets denote an average over several wavelengths of the wave. For a general gauge (i.e. $\bar{h} \neq 0, \bar{h}_{\mu,\alpha}^{\alpha} \neq 0$ and $\bar{h}_{0\mu} \neq 0$), it is

$$T_{\mu\nu}^{GW} = \frac{1}{32\pi} < \bar{h}_{\alpha\beta,\mu}\bar{h}^{\alpha\beta}{}_{,\nu} - \frac{1}{2}\bar{h}_{,\mu}\bar{h}_{,\nu} - \bar{h}^{\alpha\beta}{}_{,\beta}\bar{h}_{\alpha\mu,\nu} - \bar{h}^{\alpha\beta}{}_{,\beta}\bar{h}_{\alpha\nu,\mu} > . \tag{1.61}$$

This is divergence free in vacuum, $T_{\mu\nu,\nu}^{GW} = 0$, but contributes to the large scale curvature of the background, $G_{\mu\nu} = 8\pi(T_{\mu\nu}^{GW} + T_{\mu\nu}^{\text{MATTER}} + T_{\mu\nu}^{\text{OTHER FIELDS}})$.

According to equation (1.61), the stress-energy tensor for a monochromatic plane wave propagating in the z direction is

$$h_{\mu\nu} = \Re\{(A_+ e_{+\mu\nu} + A_\times e_{\times\mu\nu}) \exp^{-\imath\omega(t-z)}\}, \tag{1.62}$$
$$T_{tt}^{GW} = T_{zz}^{GW} = -T_{tz}^{GW} = \frac{\omega^2}{32\pi}(|A_+|^2 + |A_\times|^2). \tag{1.63}$$

The stress-energy tensor, $T_{\mu\nu}^{GW}$, is only valid under certain constraints. For instance, the background radius of curvature, R, the mean reduced wavelength, $\widetilde{\lambda} \equiv \frac{\lambda}{2\pi}$ and amplitude, a,

of a gravitational wave satisfy:

$$
\begin{aligned}
R^{-2} &\sim \text{magnitude of components of } R^{B}_{\alpha\beta\gamma\delta} \\
&\sim T^{GW}_{\mu\nu} \sim \frac{a^2}{\widetilde{\lambda}^2} && \text{if T only contributes to background} \\
&\geqslant T^{GW}_{\mu\nu} \sim \frac{a^2}{\widetilde{\lambda}^2} && \text{if T not chief source} \\
\therefore a &\leqslant \frac{\widetilde{\lambda}}{R} && \text{dimensionless amplitude must be} \ll 1 \text{ and } \widetilde{\lambda} \ll R
\end{aligned}
$$

1.2.6 Generation of gravitational waves

While equations (1.60), (1.32) and (1.33) define the energy carried by a gravitational wave (and therefore the energy lost by the source), insight into the emission can be gained by comparing to electromagnetism (once again see [3]). In electromagnetism, electric dipole radiation dominates, thus, the luminosity is as follows:

$$L_{\text{electric dipole}} = (\frac{2}{3})e^2\underline{a}^2, \tag{1.64}$$

where single particle acceleration, $\underline{a}$, leads to the dipole moment changing as $\ddot{\underline{d}} = e\ddot{x} = e\underline{a}$. So, $L = (\frac{2}{3})\ddot{\underline{d}}^2$ for a general system with a dipole moment. The gravitational analog of the electric dipole moment is the mass dipole moment,

$$\underline{d} = \sum_{particles,A} m_A \underline{x}_A. \tag{1.65}$$

Now, the first time derivative of $\underline{d}$ is the total momentum of the system,

$$\dot{\underline{d}} = \sum_{particles,A} m_A \dot{\underline{x}}_A = \underline{p}, \tag{1.66}$$

while the second time derivative has to vanish due to the law of conservation of momentum, i.e. $\ddot{\underline{d}} = \dot{\underline{p}} = 0$. We can therefore conclude that there is no mass dipole radiation.

Next, consider magnetic dipole and electric quadrupole radiation. Magnetic dipole radiation is generated by the second time derivative of the magnetic moment, $\ddot{\mu}$. The gravitational analog of the magnetic moment is the angular momentum,

$$\underline{\mu} = \sum_A (\text{position of A}) \times (\text{current due to A}) = \sum_A \underline{r}_A \times (m\underline{v}_A) = \underline{J}. \tag{1.67}$$

Since angular momentum is a constant of the motion (but can decrease, see, for example, an inspiral), this cannot radiate, so there is no dipole radiation of any sort.

The electric quadrupole radiation has the form

$$L_{\text{eletric quadrupole}} = \frac{1}{20}\dddot{\underline{Q}}^2 \equiv \frac{1}{20}\dddot{\underline{Q}}_{jk}\dddot{\underline{Q}}_{jk}, \tag{1.68}$$

where $\underline{Q}_{jk} = \sum_A e_A(x_{Aj}x_{Ak} - \frac{1}{3}\delta_{jk}r_A^2)$. The gravitational counterpart is

$$L_{\text{mass quadrupole}} = \frac{1}{5} < \dddot{\underline{I}}^2 > \equiv \frac{1}{5} < \dddot{\underline{I}}_{jk}\dddot{\underline{I}}_{jk} >, \tag{1.69}$$

where

$$\underline{I}_{jk} \equiv \sum_A m_A(x_{Aj}x_{Ak} - \frac{1}{3}\delta_{jk}r_A^2) \tag{1.70}$$

$$= \int \rho(x_j x_k - \frac{1}{3}\delta_{jk}r^2)d^3x. \tag{1.71}$$

The 1/5 arises from the tensor calculations as can be verified by substituting Equations (1.32) and (1.33) into equation (1.60) and averaging over several wavelengths and emission directions. The averaging $< ... >$ occurs because we cannot localise the energy in a gravitational wave. We note that I is the trace-free portion of the second mass moment of the source, equation (1.33), also known as the reduced quadrupole moment.

1.2.7 Geodesic deviation

We now turn to understanding the physical effect of gravitational waves on a system of particles. This will lead us to an understanding of how the waves may be detected experimentally.

In the TT gauge an incident gravitational wave will induce an oscillating component to the light travel time between the freely falling masses, A and B. Let us examine the particles from the position of A, so that $x' = 0$ is attached to the world line of A (comoving coordinates). Note the use of hatted and non-hatted indices to relate to different coordinate systems. The coordinate time is equal to the proper time at A ($x^{\hat{0}} = \tau$ and $x^{\hat{j}} = 0$) and we orient orthonormal spatial axes along gyroscopes attached to A. This is a Lorentz frame all along the geodesic of the world line of A. The line element in these orthonormal coordinates is

$$ds^2 = -dx^{\hat{0}}dx^{\hat{0}} + \delta_{\hat{j}\hat{k}}dx^{\hat{j}}dx^{\hat{k}} + O(|x^{\hat{j}}|^2)dx^{\hat{\alpha}}dx^{\hat{\beta}}. \tag{1.72}$$

As a gravitational wave passes the particles, the oscillations alter the separation vector, $\underline{n}$, going from A to B. It can be shown that the geodesic deviation satisfies

$$\frac{D^2 n^{\hat{j}}}{d\tau^2} = -R^{\hat{j}}{}_{\hat{0}\hat{k}\hat{0}}n^{\hat{k}} = -R_{\hat{j}\hat{0}\hat{k}\hat{0}}n^{\hat{k}}. \tag{1.73}$$

Now, the coordinates of the separation vector are simply the coordinates of the particle B, since A is at the origin, so

$$n^{\hat{j}} = x_B^{\hat{j}} - x_A^{\hat{j}} = x_B^{\hat{j}}. \tag{1.74}$$

Moreover, for $x^{\hat{j}} = 0$ (where the calculations are being performed), the $\Gamma^{\hat{\mu}}_{\hat{\alpha}\hat{\beta}}$ vanish for all $x^{\hat{0}}$, so $d\Gamma^{\hat{\mu}}_{\hat{\alpha}\hat{\beta}}/d\tau$ also vanish, eliminating all of the Christoffel corrections. Therefore, equation (1.73) reduces to

$$\frac{d^2 x_B^{\hat{j}}}{d\tau^2} = -R_{\hat{j}\hat{0}\hat{k}\hat{0}} x_B^{\hat{k}}. \tag{1.75}$$

In the TT gauge, at least to first order in h_{jk}^{TT}, the TT coordinate time t is equivalent to proper time τ and $R_{j0k0}^{TT} = R_{\hat{j}\hat{0}\hat{k}\hat{0}}$. Therefore, equation (1.75) can be written as:

$$\frac{d^2 x_B^{\hat{j}}}{dt^2} = -R_{j0k0}^{TT} x_B^{\hat{k}} = \frac{1}{2}\left(\frac{\partial^2 h_{jk}^{TT}}{\partial t^2}\right) x_B^{\hat{k}}. \tag{1.76}$$

If the particles were stationary with respect to each other before the wave arrived we have $x_B^{\hat{j}} = x_{B(0)}^{j}$ where $h_{jk}^{TT} = 0$, so, integrating equation (1.76), we get

$$x_B^{\hat{j}}(\tau) = x_{B(0)}^{\hat{k}} \left[\delta_{jk} + \frac{1}{2} h_{jk}^{TT}(\tau, x_A^{\hat{m}})\right], \tag{1.77}$$

which is the oscillation of B relative to A. Note that for a plane wave, separation in the direction of propagation are unaffected

$$h_{jk}^{TT} x_{B(0)}^{\hat{k}} \propto h_{jk}^{TT} k_k = 0. \tag{1.78}$$

Only the transverse direction oscillates, so a gravitational wave is transverse physically as well as mathematically.

1.3 Gravitational Wave detection

Particles free of any influence other than the curvature of spacetime (e.g. electric or magnetic forces) only move in response to the geometry of that spacetime - they obey the geodesic equation

$$\frac{d^2 x^\alpha}{d\tau^2} = -\Gamma^\alpha_{\beta\gamma} \frac{dx^\beta}{d\tau} \frac{dx^\gamma}{d\tau}. \tag{1.79}$$

Due to the principle of equivalence, the motion of a single test body is not enough to detect the presence of a gravitational wave; this requires at least two. Consider two stationary free masses in unperturbed flat space, one at the origin $x_A^i(\tau)$= (0, 0, 0), and the other a distance D away along the x-axis at $x_B^i(\tau)$=(D, 0, 0). To calculate their new coordinates after the passing of a GW, we must solve the geodesic equation. Since the amplitude is small, we can just solve for the corrections, $\delta x_A^i(\tau)$ and $\delta x_B^i(\tau)$ to the motion, which are first order in the

amplitude of the wave. The geodesic equation for the spatial coordinates is:

$$\frac{d^2x^i}{d\tau^2} = -\Gamma^i_{\beta\gamma}\frac{dx^\beta}{d\tau}\frac{dx^\gamma}{d\tau}, \tag{1.80}$$

which only need to be evaluated to first order the wave amplitude. Now, as $\Gamma^i_{\alpha\beta}$ vanishes in the unperturbed flat spacetime, we can write:

$$\frac{d^2\delta x^i}{d\tau^2} = -\delta\Gamma^i_{\beta\gamma}u^\beta u^\gamma = -\delta\Gamma^i_{tt}, \tag{1.81}$$

where $\delta\Gamma^i_{\beta\gamma}$ are the first-order changes in the Γ's, and u^β and u^γ are the *unperturbed* four-velocities of the particles, e.g. $u^\beta = u^\gamma = (1, 0, 0, 0)$. For a linearized GW, Γ^i_{tt} vanishes, therefore, so does $\delta\Gamma^i_{tt}$, thus

$$\frac{d^2\delta x^i}{d\tau^2} = 0, \tag{1.82}$$

to first order in the amplitude of the wave.

Initially, $\delta x^i = 0$, and the test masses are at rest, so $d(\delta x^i)/d\tau = 0$, so, from equation (1.82), $\delta x^i(\tau) = 0$ for all τ , and all particles:

$$\delta x^i_A(\tau) = \delta x^i_B(\tau) = 0. \tag{1.83}$$

Therefore, the *coordinate* positions of the masses remain unchanged as the wave passes (to first order). However, the distance between them will change. In unperturbed flat space, they are separated by D, but in the presence of a gravitational wave traveling in the z-direction, they are separated by:

$$\text{proper distance} = \int |ds^2| = \int |g_{\alpha\beta}dx^\alpha dx^\beta|^{\frac{1}{2}} \tag{1.84}$$

$$= \int_0^D |g_{xx}|^{\frac{1}{2}}dx \tag{1.85}$$

$$= [1 + \frac{1}{2}h^{TT}_{xx}]D \tag{1.86}$$

assuming that the period of the wave is much greater than the light travel time between the A and B. This is consistent with equation (1.77).

This means that the flight time of a photon between the two test particles would vary, and it is this effect that is made use of in order to attempt to detect gravitational waves.

1.3.1 Laser Interferometers

Following the principle set out above, ground based gravitational wave detectors generally consist of a Michelson interferometer type design, with two arms at 90° to each other (see

Figure 1.1). The beam splitter, S, and the two mirrors, M, are vertically suspended so that they can freely move in the horizontal (i.e. in the plane of the detector).

The beam from a laser, L, is split into two beams at S, which travel in the perpendicular arms. At the far end of each of the arms hangs a mirror, M, which reflects the light back. This is recombined at S, before passing into the photodetector, D. Any difference in the relative length of the arms, such as would be caused by an incident gravitational wave, would be detected from the interference between the beams. In actuality, the real design also incorporates two partially

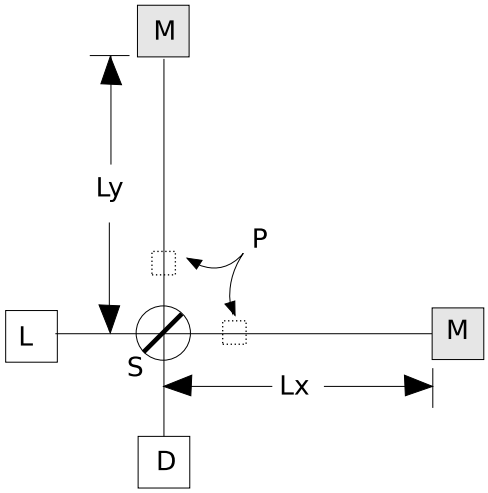

FIGURE 1.1: A simple Michelson interferometer design incorporating a Fabry-Pérot cavity within each arm.

reflecting mirrors, P, which creates a Fabry-Pérot cavity within each arm. This is tuned so that the beam is repeatedly reflected back and forth in the cavity, effectively increasing the length of the arms.

If the frequency of the laser is given by ω_{L}, then we can also define its wavenumber ($k_{\mathrm{L}} = \omega_{\mathrm{L}}/c$) and wavelength $\lambda_{\mathrm{L}} = 2\pi/k_{\mathrm{L}}$. Using these, we can write a spatial component of the electric field of the laser as

$$E_0 e^{-i\omega_{\mathrm{L}} t + i\mathbf{k}_L \cdot \mathbf{x}}. \tag{1.87}$$

As shown in Figure 1.1, the length of the arms of the interferometer in the x and y directions are given by L_x and L_y respectively.

We can determine the strain measurable by this setup (see Chapter 9 in [6]) by considering the effect on photons within the interferometer.

If we consider a photon arriving at the beam splitter at time t_0, then the part of the electric field that goes in the x arm is reflected from the mirror at distance L_x and returns to the beam splitter at a time $t = t_0 + 2L_x/c$. Similarly, the part that traverses the y arm arrives back at the beam splitter at time $t' = t_0 + 2L_y/c$. Thus, the recombined beam at a time t is made up a beam that entered the splitter at time $t_0^{(x)} = t - 2L_x/c$ (and then traversed the x

arm), and a beam that entered the splitter at a (potentially) different time, $t_0^{(y)} = t - 2L_y/c$ (and then traversed the y arm).

If we set the beam splitter at $\mathbf{x} = 0$, then the beam traversing the x arm has an initial phase of $\exp\{-i\omega_{\rm L} t_0^{(x)}\} = \exp\{-i\omega_{\rm L} t + 2ik_{\rm L} L_x\}$, while the initial phase of the beam in the y arm is $\exp\{-i\omega_{\rm L} t_0^{(y)}\} = \exp\{-i\omega_{\rm L} t + 2ik_{\rm L} L_y\}$ (using the equation for t and the fact that $k_{\rm L} = \omega_{\rm L}/c$). Using the fact that the phases are conserved during free propagation, but the fields acquire factors of -1/2 and +1/2 (due to reflection/transmission at the mirrors - see sidenote 2 on page 471 of [6]), we can write the electric fields at time t at the beam splitter as

$$E_1 = -\frac{1}{2} E_0 e^{-i\omega_{\rm L} t + 2ik_{\rm L} L_x}, \tag{1.88}$$

$$E_2 = +\frac{1}{2} E_0 e^{-i\omega_{\rm L} t + 2ik_{\rm L} L_y}, \tag{1.89}$$

$$E_{\rm out} = E_1 + E_2. \tag{1.90}$$

Inserting Equations (1.88) and (1.89) into (1.90), expanding, and using the following identities $2L_x = (L_x + L_y) + (L_x - L_y)$, $2L_y = (L_x + L_y) - (L_x - L_y)$ and $2i \sin x = e^{ix} - e^{-ix}$, we can write the total electric field as

$$E_{\rm out} = -iE_0 e^{-i\omega_{\rm L} t + ik_{\rm L}(L_x + L_y)} \sin[k_{\rm L}(L_y - L_x)], \tag{1.91}$$

and see that the power measured by the photodetector is proportional to

$$|E_{\rm out}|^2 = E_o^2 \sin^2[k_{\rm L}(L_y - L_x)]. \tag{1.92}$$

This demonstrates that any variation in the length of an arm directly results in a variation of the power detected, which is how this setup can be used as a gravitational wave detector.

In the TT gauge description, the mirrors and beam splitter can be assumed to be 'fixed' (i.e. their co-ordinates are not changed by the passing of a gravitational wave) as they are freely falling masses (or at least a good approximation of such). This allows us to define the origin of the unchanging coordinate system at the beam splitter, with the mirror terminating the x arm at (L_x,0) and the mirror terminating the y arm at (0, L_y). Under this description, the effect of the gravitational wave is seen in the propagation of light between these fixed points.

To simplify things, we will assume a gravitational wave with energy in the plus polarisation coming from the z direction. Thus, in the plane of the interferometer ($z = 0$) we have

$$h_+(t) = h_0 \cos \omega_{\rm gw} t, \tag{1.93}$$

and a space-time interval

$$ds^2 = -c^2 dt^2 + [1 + h_+(t)]dx^2 + [1 - h_x(t)]dy^2 + dz^2. \tag{1.94}$$

As we know, for photons, $ds^2 = 0$ (as they travel along null geodesics), so, for light in the x arm, we have (to first order in h_0)

$$dx = \pm cdt[1 - \frac{1}{2}h_+(t)], \tag{1.95}$$

where the positive sign corresponds to the trip from the beam splitter to the mirror, and the negative sign corresponds to the return journey. So, if we consider a photon that leaves the beam splitter at time t_0, which reaches the mirror (at coordinate $x = L_x$) at a time t_1, then, by integrating equation (1.95) with the positive sign,

$$\int_0^{L_x} dx = + \int_{t_0}^{t_1} cdt(1 - \frac{1}{2}h_+(t)), \tag{1.96}$$

we get

$$L_x = c(t_1 - t_0) - \frac{c}{2}\int_{t_0}^{t_1} dt' h_+(t'). \tag{1.97}$$

At this point, the photon is reflected by the mirror and returns to the beam splitter at a time t_2. We can follow the same procedure for this journey (i.e. integrating equation (1.95)), but with the negative sign and from $x = L_x$ to $x = 0$,

$$L_x = c(t_2 - t_1) - \frac{c}{2}\int_{t_1}^{t_2} dt' h_+(t'). \tag{1.98}$$

Summing the journeys, we get,

$$2L_x = c(t_1 - t_0) + c(t_2 - t_1) - \frac{c}{2}\int_{t_0}^{t^1} dt' h_+(t') - \frac{c}{2}\int_{t_1}^{t^2} dt' h_+(t'), \tag{1.99}$$

which, by dividing throughout by c and rearranging becomes,

$$t_2 - t_0 = \frac{2L_x}{c} + \frac{1}{2}\int_{t_0}^{t_2} dt' h_+(t'), \tag{1.100}$$

which tells us that the arrival time of a photon at the beam splitter after traveling down the x arm is $t_2 = t_0 + 2L_x/c$ plus a correction of order h_0. This can be substituted back into the upper limit of the integral in equation (1.100) (as the integrand is already $O(h_0)$ and we are neglecting terms of $O(h_0^2)$) to get

$$t_2 - t_0 = \frac{2L_x}{c} + \frac{1}{2}\int_{t_0}^{t_0+2L_x/c} dt' h_+(t') \tag{1.101}$$

$$= \frac{2L_x}{c} + \frac{h_0}{2\omega_{\rm gw}}\left[\sin(\omega_{\rm gw}(t_0 + 2\frac{L_x}{c})) - \sin\omega_{\rm gw}t_0\right]. \tag{1.102}$$

Using the identity $\sin(\alpha + 2\beta) - \sin\alpha = 2\sin\beta\cos(\alpha + \beta)$, this can be rewritten as

$$t_2 - t_0 = \frac{2L_x}{c} + \frac{h_0 L_x}{c}\frac{\sin(\omega_{\rm gw}L_x/c)}{\omega_{\rm gw}L_x/c}\cos(\omega_{\rm gw}(t_0 + \frac{2L_x}{c})), \tag{1.103}$$

and then,

$$t_2 - t_0 = \frac{2L_x}{c} + \frac{L_x}{c} h(t_0 + L_x/c) \frac{\sin(\omega_{\rm gw} L_x/c)}{\omega_{\rm gw} L_x/c}, \tag{1.104}$$

by using the fact that $h_+(t) = h_0 \cos \omega_{\rm gw} t$. Essentially the argument to $h(t)$, $t_0 + L_x/c$, is the time at which the photon touches the mirror on the x arm. Also,

$$\frac{\sin(\omega_{\rm gw} L_x/c)}{\omega_{\rm gw} L_x/c} \equiv \text{sinc}\left(\frac{\omega_{gw} L}{c}\right), \tag{1.105}$$

goes to 1 in the limit $\omega_{\rm gw} L_x/c \to 0$. Thus, if the period of the gravitational wave is large compared to L_x/c then the shift, Δt, in $t_2 - t_0$ with respect to its unperturbed value, $2L_x/c$ is simply $h(t_1)L_x/c$. However, if the opposite is true then Δt is suppressed - over the course of the photons travel time, the sign of $h(t)$ changes so much that it partially cancels out.

The analysis for the y arm is identical, except that the sign of $h(t)$ is reversed, so our result is

$$t_2 - t_0 = \frac{2L_y}{c} - \frac{L_y}{c} h(t_0 + L_y/c) \frac{\sin(\omega_{\rm gw} L_y/c)}{\omega_{\rm gw} L_y/c}. \tag{1.106}$$

We are generally more interested in the light that exits the beam splitter for a given value of the observation time, t, so, in order to return to the beam splitter at time t, the light that travelled thorough the x arm must have started at a time $t_0^{(x)}$. This means that, to first order in h_0, reversing equation (1.104) gives,

$$t_0^{(x)} = t - \frac{2L_x}{c} - \frac{L_x}{c} h(t - L_x/c) \text{sinc}(\omega_{\rm gw} L_x/c), \tag{1.107}$$

where t_2 has been replaced by t (the observation time, $h(t_0 + L_x/c)$ by $h(t - 2L_x/c + L_x/c) = h(t - L_x/c)$ and $\sin(x)/x$ by $\text{sinc}(x)$. Similarly, equation (1.106) becomes

$$t_0^{(y)} = t - \frac{2L_y}{c} + \frac{L_y}{c} h(t - L_y/c) \text{sinc}(\omega_{\rm gw} L_y/c). \tag{1.108}$$

Using the same coordinate setup as before, we can see that the phase of the field at the beam splitter at time $t_0^{(x)}$ is $\exp(-i\omega_{\rm L} t_0^{(x)})$, so the field is

$$E^x(t) = -\frac{1}{2} E_0 e^{-i\omega_{\rm L} t_0^{(x)}} \tag{1.109}$$

$$= -\frac{1}{2} E_0 e^{-i\omega_{\rm L}(t - 2L_x/c) + i\Delta\phi_x(t)}, \tag{1.110}$$

where

$$\phi_x(t) = \frac{\omega_{\rm L} L_x}{c} h(t - L_x/c) \text{sinc}(\omega_{\rm gw} L_x/c) \tag{1.111}$$

$$= h_0 \frac{\omega_{\rm L} L_x}{c} \text{sinc}(\omega_{\rm gw} L_x/c) \cos(\omega_{\rm gw}(t - L_x/c)). \tag{1.112}$$

Similarly, for the field that travelled through the y arm, we have

$$E^y(t) = +\frac{1}{2}E_0 e^{-i\omega_{\rm L} t_0^{(y)}} \tag{1.113}$$

$$= +\frac{1}{2}E_0 e^{-i\omega_{\rm L}(t-2L_y/c)+i\Delta\phi_y(t)}, \tag{1.114}$$

where

$$\phi_y(t) = h_0 \frac{\omega_{\rm L} L_y}{c} \mathrm{sinc}(\omega_{\rm gw} L_y/c)\cos(\omega_{\rm gw}(t - L_y/c)). \tag{1.115}$$

In practice, the length of the two arms are made as close as possible in order to cancel common noise. This means we can replace L_x and L_y in $\Delta\phi_x$ and $\Delta\phi_y$ by $L = (L_x + L_y)/2$ (as this is already of order h_0), but take account of the differences in length using the terms $t - 2L_x/c$ and $t - 2L_y/c$. If we write $2L_x = 2L + (L_x - L_y)$ and $2L_y = 2L - (L_x - L_y)$, the field equations become,

$$E^x(t) = -\frac{1}{2}E_0 e^{-i\omega_{\rm L}(t-2L_x/c)+i\phi_0+i\Delta\phi_x(t)}, \tag{1.116}$$

$$E^y(t) = +\frac{1}{2}E_0 e^{-i\omega_{\rm L}(t-2L_y/c)+i\phi_0+i\Delta\phi_y(t)}, \tag{1.117}$$

where

$$\begin{aligned}
\phi_0 &= k_L(L_x - L_y), \\
\Delta\phi_y &= -\Delta\phi_x, \\
\Delta\phi_x(t) &= h_0 k_L L \mathrm{sinc}(\omega_{\rm gw} L/c)\cos(\omega_{\rm gw}(t - L/c)), \qquad (1.118)\\
&\equiv |\Delta\phi_x|\cos(\omega_{\rm gw} t + \alpha), \qquad (1.119)
\end{aligned}$$

and $\alpha = -\omega_{\rm gw} L/c$ is a phase. The total phase difference that can be induced in the interferometer by incident gravitational waves is

$$\Delta\phi_{\rm TOT} \equiv \Delta\phi_x - \Delta\phi_y = 2\Delta\phi_x \tag{1.120}$$

In the limit $\omega_{\rm gw} L/c \ll 1$, equation (1.118) reduces to,

$$\begin{aligned}
\Delta\phi_x(t) &\simeq h(t - L/c)k_L L \\
&\simeq h_0 k_L \cos(\omega_{\rm gw} t) L,
\end{aligned} \tag{1.121}$$

but, $k_L = 2\pi/\lambda_L$, so this becomes,

$$\Delta\phi_x(t) \simeq \frac{2\pi}{\lambda_L} L h(t). \tag{1.122}$$

The interferometer is initially set up so that when the beams are recombined at D, they are 180° out of phase with each other (destructive interference). So, when the length of the arms change, the two beams will have some coherent phase. By measuring the amount of light detected at D, it is possible to measure the change in the arm length as a function of time (note that the sensitivity of the detector is limited by its ability to measure the phase difference). So, we can write

$$\begin{aligned}\Delta\phi_{\min} &\equiv 2\Delta\phi_x \\ &\equiv \frac{4\pi}{\lambda_L} L h_+(t)_{\min}.\end{aligned} \tag{1.123}$$

Rearranging we get

$$h_+(t)_{\min} \equiv \frac{\lambda_{laser}}{4\pi}\frac{\Delta\phi_{\min}}{L} \tag{1.124}$$

where $h_+(t)_{min}$ is the minimum sensitivity that can be measured, λ_{laser} is the wavelength of the laser, ϕ_{min} is the minimum phase difference that can be measured and L is the length of the arms. A typical value of $h_+(t)_{min}$ may be 10^{-21}, with arms of order $2-4$ km in length.

The LIGO (Laser Interferometer Gravitational-Wave Observatory) gravitational wave network consists of three interferometers. Two are co-located at the same site in Hanford, Washington, one with 4 km arms and one with 2 km arms[1]. The third interferometer is based 3000 km away at Livingston, Louisiana, and also has 4 km arms. There is also the Anglo-German detector GEO 600 based near Hannover, Germany, with arm lengths of 600 m. The LIGO Scientific Collaboration (LSC) also shares data with the Virgo detector in Cascina, Italy, which has 3 km arms. As we are discussing ground based detectors, it is important to remember that

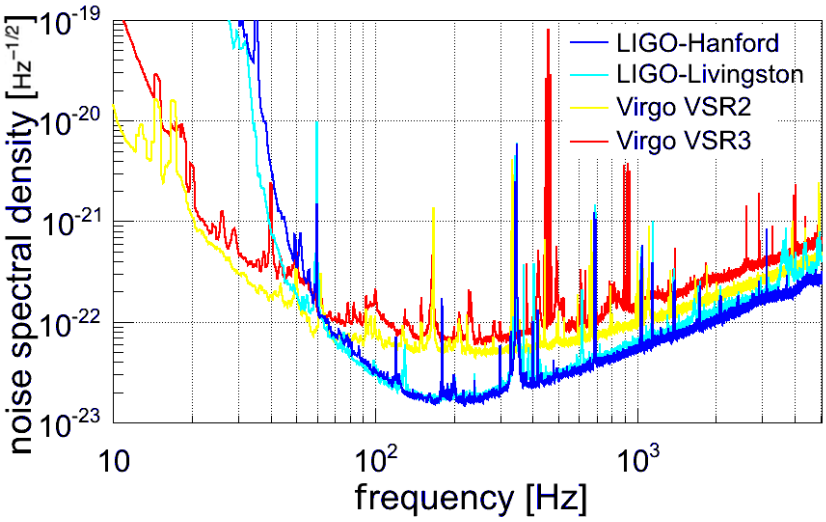

FIGURE 1.2: Detector noise curves for LIGO Science Run 6, Virgo Science Runs 2,3 (2009-2010). Taken from [1]

[1]Although the 2 km arm interferometer has been decommissioned, and did not take part in the 6^{th} Science run

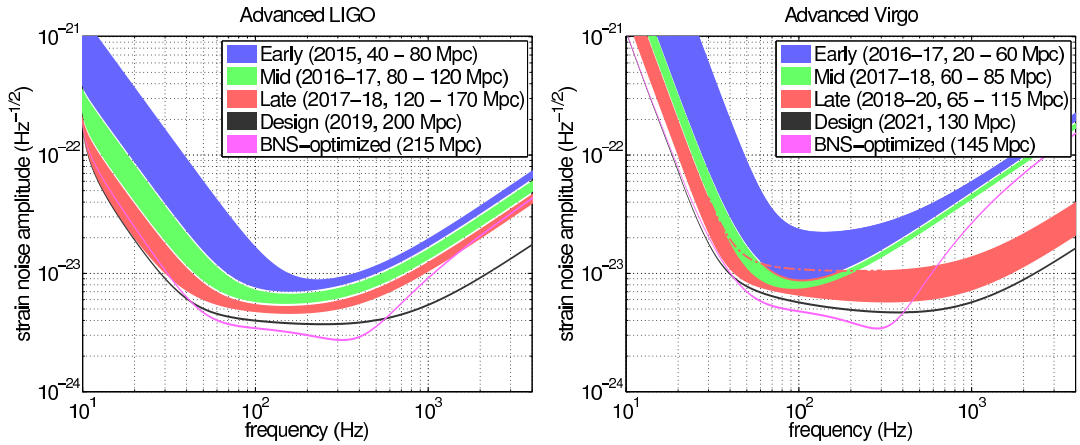

FIGURE 1.3: aLIGO (left) and AdV (right) target strain sensitivity as a function of frequency. The average distance to which binary neutron star (BNS) signals could be seen is given in Mpc. Current notions of the progression of sensitivity are given for early, middle, and late commissioning phases, as well as the final design sensitivity target and the BNS-optimized sensitivity (plots taken from [2]).

$h_+(t)_{min}$ will be affected by many things such as seismic noise, mirror "noises" and the stability of the laser. Figure 1.2 shows the sensitivity of these ground based detectors as a function of frequency for the latest science run of LIGO and Virgo. These detectors are currently being upgraded to their "advanced" states, aLIGO,[7], and AdV, [8, 9], respectively. They should be operational around 2015, and up to design sensitivity by around 2020, see [2]. aLIGO and AdV will have greeter sensitivities than the current generation by one order of magnitude; see Figure 1.3).

1.4 Astrophysical sources

By convention, searches for gravitational waves using ground-based detectors such as LIGO and Virgo are grouped into four categories.

1.4.1 Compact binary coalescence

This is the analysis concerned with the signal generated from the inspiral and merger of a compact binary system (black hole-black hole, (BBH), black hole-neutron star (NSBH), neutron star-neutron star (BNS)). It sweeps upwards in frequency and amplitude through the sensitive band of the gravitational wave detectors, see [10]. The detection of gravitational waves from this type of astrophysical source will provide information about many different areas of astrophysics, such as strong field gravity, dense matter, and the populations of such compact objects (neutron stars and black holes) in the Universe.

1.4.2 Burst signals

While certain sources (such as the CBCs above) have sophisticated models that are expected to produce physically accurate signals, there are a raft of other sources with uncertain waveforms, and therefore unknown amplitude and frequency evolution. This motivates performing a broad search for short duration gravitational wave transients, or bursts; such searches requires very robust detection methods that are insensitive to the waveform morphology.

There are many potential sources (see [11] for latest search results) that fall into this category, such as gamma-ray bursts, supernovas, neutron stars (e.g. pulsar glitches), and of course, the Unknown - things we don't know about yet.

Note that the focus has broadened away from unknown signals with durations shorter than $\sim$ 1 second, see [12], to longer signals and the incorporation of available signal waveform knowledge, to seek specific targets in addition to unknown phenomena and mechanisms of emission and propagation.

1.4.3 Continuous-wave signals

The previously summarised sources are transient in nature typically lasting for less than a second. On the other end of the scale are the sources that radiate continuous-wave signals; single band oscillations that can last months to years.

The most promising source of these continuous-wave (CW) gravitational signals are rapidly rotating neutron stars. The sources tend to fall into four broad categories: non-accreting known pulsars for which timing data is available, non-accreting known stars without timing data, unknown isolated stars, and accreting stars in known binary or stars in unknown binary systems. For each type of source, the properties of the source population is known or can be inferred. They are expected to emit gravitational radiation through a variety of mechanisms such as elastic deformations, [13], magnetic deformations, [14], unstable r-mode oscillations, [15] and [16], and free precession, [17] - all of which operate differently in accreting and non-accreting stars (see [18] for a review of these emission mechanisms).

1.4.4 Stochastic backgrounds

The last of the "standard" sources is very different from the others. A stochastic gravitational-wave signal is formed from the superposition of many different events or processes, individually too weak and too numerous to resolve, but which combine to produce a confusion-limited background. The stochastic gravitational-wave background (SGWB) can arise from cosmological sources (such as inflation, cosmic strings, and pre-Big-Bang models), or astrophysical sources

(such as CBCs, supernovae, neutron stars). One of the prime objectives of the Stochastic Working Group is to measure the SGWB. Recently, the group's focus has expanded, and long-lived gravitational wave transients (from e.g. rotational instabilities in neutron stars and accretion disks) are being investigated. The unifying feature of current stochastic analyses is that the signal they seek to measure is long-lived and difficult or impossible to model (as it is random) in the time domain, [19].

Chapter 2

Intermediate duration gravitational waves

As indicated by the summary of the "standard" searches in the previous chapter, to date, most gravitational wave searches by ground-based detectors have focused on either long duration signals (e.g. continuous wave and stochastic, see [18] & [19]) or transient signals with duration of order a second or less in length (e.g. see the CBC and Burst searches, [10] & [12]).

The shortest continuous wave signal searches feature a coherent integration with a duration of approximately 6 hours to 1 week [18], which means that a gap exists in the space of gravitational wave searches for durations of order 1 to 10^4 seconds.

There are several potentially interesting sources of gravitational waves in this range that we will discuss below, firstly however, we will outline the pitfalls of searching in this new regime and the issues that we will need to take into account.

2.1 Searching in a new regime

Tackling the unknown presents as many challenges as it does rewards, and there are factors that need to be taken into account before launching into unknown.

2.1.1 Waveform knowledge

The waveforms produced by such intermediate duration sources will either be unknown or badly understood, these are only issues for unmodeled burst searches in which the signal energy is spread out out over many pixels in a time-frequency map (as discussed in the next Chapter). The use of a matched-filter like those used in continuous wave and CBC searches

are unaffected by this provided that the waveform is fully know. Essentially, the matched filter template collects all of the signal energy no matter how it would be spread out in the TF plane (see Section 3.2.1, and [20]).

2.1.2 Signal parameters

The more mature and established searches, in both long and short duration regime, are looking for sources where there is a reasonable amount of knowledge about the physical mechanisms generating the gravitational wave - maybe not a huge amount, but usually enough to know how to limit the search (e.g. Burst search for GRBs has good hypotheses about their production mechanisms to compare to the electromagnetic output, CBC has numerical models etc.). The physical mechanisms driving the sources that emit gravitational radiation in the 1-10^4 second range may be relatively unknown, so we have a large signal space to search.

2.1.3 Signal length

If the total energy emitted by a short-duration signal was extended to a longer time span, then, unless it was particularly loud, it would more relatively more difficult to detect. This is the problem we face - for a longer duration signal to be detectable a more energetic engine would need be powering it. Therefore, to beat the noise limit, the signal may have to be integrated over a longer period, or the analysis would have to use larger Fourier transform lengths to capture as much of the signal as possible.

Before we can understand which of these are problems are relevant and if they can be overcome, it is necessary to determine the kind of astrophysical phenomena that may produce gravitational radiation in this time band.

2.2 Potential sources

We can theorise that any astrophysical source that can generate a *detectable* gravitational wave burst of intermediate duration (1-10^4 seconds), will probably have an exotic mechanism driving the emission. The reasoning behind this is simple - the source must radiate it's energy away slowly enough that it can emit for the required amount of time, *and* by energetic enough to be detected.

There are several models in the literature (see [21] for an overview) which may fit this hypothesis, but we need to determine which of them is most likely to generate a realistic opportunity of detection. We consider three in particular: microquasars, long gamma-ray bursts and neutron star instabilities. Note that although these models motivate the search, we are not tied

to them being correct. They exist to demonstrate that the parameter space we are interested in is not a barren wasteland.

2.2.1 Calculating Signal-To-Noise Ration

The detectability of a signal is mainly determined by its signal-to-noise ratio (SNR), ρ. To this end, we will use the parameters given in the literature to estimate ρ for each sources of interest. If we assume isotropic emission for a narrowband signal with central frequency, f_0, then, for a source at a distance r the energy emitted in gravitational waves is given by (taken from [20])

$$E_{\mathrm{GW}} = \frac{\pi^2 c^3}{G} r^2 f_0^2 h_{\mathrm{rss}}^2 \tag{2.1}$$

where h_{rss} is the root-sum-square amplitude,

$$h_{\mathrm{rss}} = \int_{-\infty}^{\infty} dt[h_+^2(t) + h_\times^2(t)]. \tag{2.2}$$

We can relate E_{GW} to ρ using (taken from [20])

$$\rho^2 \approx \frac{G}{2\pi^2 c^3} \frac{E_{\mathrm{GW}}}{S(f_0 f_0^2 r}, \tag{2.3}$$

where $S(f_0)$ is the one-sided noise power spectrum (i.e. Figures 1.2 and 1.3).

Note that the emission from the microquasar's "cannonballs" is a broadband signal rather than narrowband. The difficulty here is that once a signal has a wide bandwidth, its precise shape becomes important so we need to assume a waveform. One of the simplest short pulse waveforms is a linearly polarised Gaussian which can be written as

$$h(t) = A \exp(-t^2/(2\tau^2)), \tag{2.4}$$

where the shape is determined entirely by the decay time, τ.

Now, the equation for linear motion emission (taken from [20]) is

$$E_{\mathrm{GW}} = \frac{8}{15} \frac{\pi^2 c^3}{G} r^2 f_0^2 h_{\mathrm{rss}}^2 \tag{2.5}$$

which should work for Gaussians if we define the effective central frequency to be

$$f_0 = \frac{1}{(4\pi\tau)} \tag{2.6}$$

Although this is not an exact model, it should allow us to compute SNRs for each of the three models (microquasar, long GRB, r modes) in the same way so we may compare them.

2.2.2 Galactic microquasars

A galactic microquasar[1] is an X-ray binary system that mimics many of the properties of Active Galactic Nuclei (AGNs) - i.e. they have relativistic jets and are very luminous in the optical/X-ray frequency bands.

The engine powering this system is assumed to be mass accretion onto a compact object (either a neutron star or stellar mass black hole) from a companion star. Table 2.1 summarises the properties of this source, giving the rate/distances of microquasars in the galaxy, the parameters associated with the "blob" of ejected matter, and the expected values associated with gravitational wave emission for standard and advanced LIGO (aLIGO from now on). The double line separates the more likely parameters from those about which we have less certainty.

Kinematics of the X-ray binaries suggest that low mass black holes ($< 7M_{\odot}$) are formed by the delayed collapse of a neutron star with energetic supernovae, whereas stellar black holes with masses equal or greater than 10 solar masses are the result of the direct collapse of massive stars [22].

The shape and duration of the gravitational wave signal will depend on the matter accreted by the compact object, which, if comparable to flares in the optical [23], should produce GW bursts which can last $\gtrsim 100$s.

Upper limit on microquasars in galaxy	~ 100
Closest known microquasars, GRO J1655-40 1A 0620-00	 1 kPc 200 Pc
Time between blobs	~ 1s
Blob acceleration time, Δt	$\sim 10^{-4}$ s
Max frequency $\sim 1/(4\pi\tau)$	$\sim 800Hz$
LIGO Noise amplitude aLIGO Noise amplitude	$4 \times 10^{-23}\,\mathrm{Hz}^{-1/2}$ $5 \times 10^{-24}\,\mathrm{Hz}^{-1/2}$
Strain amplitude, h E_{GW}	$\sim 7 \times 10^{-23}$ $\sim 6.4 \times 10^{36} J$
LIGO SNR (distance $1kPc$) aLIGO SNR (distance $1kPc$)	~ 0.9 ~ 7.2

TABLE 2.1: Summary of microquasar properties associated with the cannonball model. The double line separates the more likely parameters from those about which we have less certainty.

[1]It is probable that this astrophysical phenomena exists in other galaxies, but to be detectable at those distances, it would need to be comparatively more energetic.

2.2.2.1 Production mechanism

One proposed mechanism to explain the existence of these microquasar jets is known as the "cannonball model". This was originally proposed by Dar and De Rujula [24] as a possible explanation for the production of long ($t_{burst} > 1s$) gamma-ray bursts (GRBs) associated with core collapse supernova and explains their main observational features (i.e. relativistic jets and superluminal motion of the afterglow).

Relativistic jets are generated when mass is accreted at a high rate onto a central compact object. The jets are not continuous streams, but instead consist of individual blobs of plasma that have been ejected via some mechanism (e.g. electromagnetic acceleration due to a dynamo effect, some form of magnetohydrodynamic (MHD) instability accelerating the blob of plasma, or radiation pressure). The firing of these blobs (or *cannonballs* as they are referred to) coincide with removal of matter from the surrounding environment.

Segalis and Ori [25] suggest that matter from the supernova explosion falls back onto the compact object forming a shell. Once a suitable reservoir of matter has been accreted, a plurality of cannonballs are ejected in each of the polar directions with a very large Lorenz factor ($\gamma \sim 10^3$). These then collide with the shell around the compact object to produce the GRB. They then continue, cooling and expanding to form the jet.

2.2.2.2 Gravitational wave properties

In the ultra-relativistic limit, the emitted gravitational waves have a wide angular distribution $(1 + \cos\theta)$ (θ being the angle between the blob's velocity and the direction vector from the source to the observer), compared to the small solid angle of γ^{-2} for the electromagnetic radiation. Note however, that the burst's frequency is Doppler shifted - the blue-shift factor being strongly beamed in the forward direction, meaning that the energy flux carried by the gravitational wave is also beamed .

Now, gravitational radiation is emitted from the system when the blob changes velocity, and, as several blobs are ejected in each polar direction and accelerated with a large Lorenz factor, there should be a distinctive emission of radiation for such a production mechanism.

To calculate the amplitude and frequency of the gravitational wave emission, we use the blob's energy, γm, and the fundamental time scale of the blob's acceleration, Δt.

From [25], we get the equation for the blob's gravitational wave amplitude (equation 12),

$$h_+ = (1/2)h_0 \sin^2\theta = \frac{2\gamma m\beta^2}{d}\frac{\sin^2\theta}{1-\beta\cos\theta} \tag{2.7}$$

where $\beta = v/c$ and $\gamma = (1-\beta^2)^{-1/2}$ (the Lorentz factor) represent the standard special-relativistic quantities, and d is the spatial distance to the observer.

If we consider the ultra-relativistic case ($\gamma \gg 1$ so $\beta \approx 1$) , Equation (2.7 reduces to the following two equations. For small θ

$$h_+ \approx \frac{4\gamma m}{d}\frac{\theta^2}{\gamma^{-2}+\theta^2} \qquad (\theta \ll 1), \tag{2.8}$$

and then large θ compared to $1/\gamma$

$$h_+ \approx \frac{2\gamma m}{d}(1+\cos\theta) \qquad (\theta \gg \gamma^{-1}). \tag{2.9}$$

Note that Segalis and Ori [25] use this result to show that while there is a small amount of "beaming", there is not a sharp enhancement in the narrow forward direction, rather there is a "hole" at $\theta < \gamma^{-1}$ (i.e. anti-beaming, so there is no emission in the direction of the jet). They also assume that these ejections are usually symmetrical (i.e. for every ejected blob, there is another in the opposite direction), so for a pair of ultra-relativistic blobs the contributions are direction independent, so,

$$\Delta h_+ \approx \frac{4\gamma m}{d}. \tag{2.10}$$

The characteristic frequency of the observed burst is given by

$$f_c \sim [(1-\cos\theta)\Delta t]^{-1}, \tag{2.11}$$

which leads to the following detection rate, [25],

$$N = \frac{128\pi}{3} n(\gamma m G/c^2 h_d)^3 w(\Delta t), \tag{2.12}$$

where w is either 1 (for $\Delta t < \Delta t_d$) or $1-(1-\Delta t_d/\Delta t)^4$) (for $\Delta t > \Delta t_d$), $\Delta t_d \equiv 1/(2f_d)$, n is the rate of micro-quasars per unit volume, and h_d is the strain value of the detector at the characteristic frequency, f_d, of the ejection (which is determined by Δt).

Using the event rate per unit volume, the sensitivity of the detector and the optimal frequency, [25] showed that is is possible to get a detection rate of approximately 1 a year, and blob energy, $E = \gamma mc^2 \sim 10^{52}$erg, which gives a maximal detection distance for LIGO of, $d_{max} \sim 15$ Mpc.

Considering the low frequency of events, and the low SNR for the LIGO detectors, it would appear that this source is not, in itself, sufficient motivation to carry out an intermediate duration search. This is a more "opportunistic" source that we may detect if the conditions were optimal.

Segalis and Ori [25] determined the amount of gravitational radiation that this model would

produce due to the acceleration of the cannonballs, disregarding the collision with the surrounding matter. They assumed that several (3-5) relatively energetic blobs are emitted in each polar direction; their contribution to the gravitational radiation do not cancel, but the two components will reach the detector at different times and with different Doppler factors. Those towards the observer will be blue-shifted and those away, red-shifted.

Note that the blob ejection mechanism is unimportant, so long as the system is at rest for $t < t_0$, the blob undergoes acceleration from $t = t_0$ to $t = t_1$ and for $t > t_1$ the blob has constant speed. Also, the amplitude of the gravitational wave only depends on the ejected particle's energy.

2.2.3 Long GRBs

Gamma-ray bursts (GRBs) are bright flashes of γ-rays that are observed approximately once per day. These bursts arrive from apparently random directions in the sky (i.e. an isotropic distribution) and last between tens of milliseconds and thousands of seconds.

Studies of GRBs observations have shown that there are two major sub-populations of GRBs dependent on their duration. This bimodal GRB duration distribution has a minimum at approximately 2 s. GRBs with durations shorter than 2 s are, on average, composed of higher energy (i.e. harder) photons than their longer brethren and make up approximately one quarter of the detected bursts. This demonstrates that the two populations represent distinctive physical phenomena. See the review by Meszaros [26] for more information.

One viable model for the origin of short hard GRBs is the coalescence of compact binary systems (e.g. double neutron stars or a neutron star and a black hole), which, if we can assume the gravitational wave emission is as energetic as the electromagnetic observations (by SWIFT and HETE-2 in 2005, see [27]) suggest strong gravitational-wave sources. Long GRBs are thought to originate from the collapse of massive stars (see for example, [28], [29] & [30] - the "collapsar" model), and have been associated with supernova (see, for example, [31] & [32]).

There are already established matched-filter searches targeting short GRBs, so we will focus here on long GRBs where the possible gravitational wave emission mechanisms are less well understood.

Long GRBs can last anywhere from 1s to $\sim 10^3$s and are often followed by an "afterglow" of X-rays that can last for hours or even days. This can allow the localisation of bursts to their host galaxy (and therefore give a measure of their distance) and the detection of associated optical and radio emission which can be observed for weeks to years.

We can therefore assume that the physical mechanisms generating long GRBs are catastrophic events involving a stellar-mass object or system (such as the formation of a stellar-mass black hole or magnetar) that releases a vast amount of energy ($\gtrsim 0.001\ M_{\odot}c^2$) in γ-rays in a compact region ($< 100km$) on time scales of seconds to minutes. This energy source, referred to as the central engine, accelerates an ultra-relativistic outow of matter to a Lorentz factor $\gtrsim 100$ and it is this outow that generates the observed γ-ray and X-ray prompt emission, and later the afterglow, at large distances from the source. This model is now generally accepted. In this context, microquasars that contain black holes could be fossils of gamma ray burst sources of long duration.

2.2.3.1 Production mechanism

A variation of the above production mechanism has been developed by van Putten, [33, 34], suggesting that the engine for the long bursts could actually be formed from a high angular momentum compact source with a thick torus of matter (formed from a hypernova or BH-NS coalescence). If the mass of the in-falling matter is similar to the mass of the newly created compact object, then a thick accretion torus will form around it, rather than a thin accretion disc.

The energy of the GRB is then roughly equivalent to the length of time between pulses, which suggests some form of discontinuous accretion, where a "reservoir" of matter must be replenished from the surrounding matter.

In this model, bursts of gravitational radiation can occur due to the large non-axisymmetry of the torus. The rapid spin of the black hole creates a state of suspended accretion which can support open magnetic field lines to infinity. This gives rise to jets of matter along the axial rotation axis. The gravitational wave emission from the lumpy torus is replenished via its magnetic connection to the central black hole. Most of its angular momentum is transferred to the torus, which carries on for the lifetime of the spin of the black hole, roughly 10-90 s. The torus may also precess, modulating the phase of the strain amplitude (this is one of the major differences with the standard model and may be observable).

Table 2.2 is a summary of the parameters associated with the van Putten model. We have the supernova rate, gravitational wave emission values and the expected SNRs for standard and enhanced LIGO. The double line separates the more likely parameters from those about which we have less certainty.

2.2.3.2 Gravitational wave properties

If we assume that the central engine of a long GRB is a rapidly spinning black hole with a torus of matter (may be hypernova/BH-NS coalescence), there is a very good possibility that

Event rate	$\sim$ multiple events out to 100MPc
Frequency	$\sim 1 - 2kHz$
Emission duration, Δt	$\sim 10 - 90$ s
LIGO Noise amplitude	$3 \times 10^{-23}\,\mathrm{Hz}^{-1/2}$
aLIGO Noise amplitude	$4.5 \times 10^{-24}\,\mathrm{Hz}^{-1/2}$
Strain amplitude, h	$\sim 3 \times 10^{-23}$
E_{GW}	$\sim 8.5 \times 10^{45} J$
LIGO SNR (distance $100MPc$)	0.7
aLIGO SNR (distance $100MPc$)	5.3

TABLE 2.2: Summary of long GRB properties associated with the van Putten model. The double line separates the more likely parameters from those about which we have less certainty.

the system will emit gravitational radiation. Whenever the torus becomes non-axisymmetric it emits gravitational radiation, and low-frequency modulated radio waves, which could be useful in correlated searches. The torus is strongly coupled to the spin energy of the black hole (through its magnetosphere), so asymmetries in the torus will produce gravitational radiation at twice the Keplerian angular frequency (1-2 kHz) with a duration of approximately 10-15s. The torus is likely to possess dynamical and radiative instabilities, and if massive enough, be unstable to self gravity and subject to the Chandrasekhar-Friedman-Shutz instability (similar to the emission mechanism of rapidly rotating neutron stars - see Section 2.2.4 below).

Emissions from the torus are replenished by spin up Maxwell stresses on the inner face of the torus through its magnetic connection to the black hole. Thus, an equilibrium state could emerge where radiative losses from the torus are balanced by this replenishing effect, leading to a suspended accretion state. The motion of the torus would still be subject to large shear effects between its inner and outer edge, with matter outflowing radially.

[35], [34] predict gravitational wave emission due to the precessing torus, with the total energy and characteristic frequency given by:

$$E_{GW} \approx 0.2 M_{\odot} \left(\frac{\eta}{0.1}\right) \left(\frac{M_H}{7M_{\odot}}\right), \tag{2.13}$$

where M_H is the black hole mass and η is the ratio of angular velocity of the torus and black hole, and

$$f_{GW} \approx 500Hz \left(\frac{\eta}{0.1}\right) \left(\frac{7M_{\odot}}{M_H}\right). \tag{2.14}$$

The duration of emission can be estimated as follows:

$$T_{GW} \sim 90s \left(\frac{\eta}{0.1}\right)^{8/3} \left(\frac{M_H}{7M_{\odot}}\right) \left(\frac{M_T}{0.03M_H}\right), \tag{2.15}$$

where M_T is the mass of the torus.

Here η denotes the ratio of the angular velocity of the torus to the black hole which represents the efficiency of the system converting black hole spin energy into emission by the torus, M_H is the black hole mass and $M_\odot = 2 \times 10^{54}$ erg. Note that the energy emitted by the torus, Equation 2.13, is very large, $\sim 1000\times$ more than emitted in γ-ray energy.

The energy budget of the system is determined by the assumption that 10 % of the black hole's spin energy could be radiated away in gravitational radiation, and that it is an approximately sinusoidal waveform with a frequency that wanders slowly in time as the black hole spins down.

This means that the gravitational wave emission from the torus (quadrupole and higher-order multipole mass-moments) changes over the lifetime of the black hole dependent on its spin. Emission is fundamentally at a single frequency, but Lense-Thirring precession of the torus will lead to this broadening. In [35], [34] the authors estimate a modulation of the strain amplitudes by $\sim 10\%$ of f_{GW}, so there would be sidebands at $f_{GW} \times (1 \pm 0.1)$. These might appear as just a broadening of the dominant emission at f_{GW}.

They also state that they expect the dominant frequency to decrease (due to spindown) by $\sim 10\%$ over the first 50% of the energy emission. This is of practical importance because fixed-frequency gravitational wave emission could easily be confused with fixed-frequency background noise sources. The change in the dominant frequency could help in distinguishing the gravitational wave from the background (see Section 5.2.7 and Figure 5.7 for a representation of this waveform).

The instantaneous strain amplitude is approximately

$$h = 3 \times 10^{-23} \left(\frac{\eta}{0.1}\right)^{4/3} \left(\frac{M^H}{0.7M_\odot}\right) \left(\frac{100Mpc}{d}\right) \left(\frac{M_T}{0.03M_H}^{1/4}\right), \tag{2.16}$$

from which we can determine the SNR for a waveform of this type - see Table 2.2 - once again, the low SNR and event rate restrict the usefulness of this source.

2.2.4 Neutron star instabilities

A neutron star (NS) can be formed when a star goes supernova, generally with a mass $\sim 1.3M_\odot$ concentrated into a radius of $\sim$ 10 km. Oscillating neutron stars could be interesting astrophysical sources of gravitational waves. Observations could give valuable information about the internal structure and equation of state for the NS, such as the mass, radius, magnetic field strength and rotation. However, the formation of a NS following a supernova or a starquake caused by an internal phase-transition may be violent enough to produce gravitational waves at a detectable level, so long as these events take place in our neighbourhood $\sim$ Local Group of Galaxies.

There is however, a second mechanism - instabilities. The neutron star may be inherently unstable to oscillation modes that will emit gravitational waves. These may be long-duration so we will focus on those here.

2.2.4.1 Production mechanism

These instabilities are associated with unstable modes of oscillation, where non-axisymmetric pulsations of the neutron star can lead to the production of gravitational waves. The most obvious scenario that could lead to this situation is neutron stars with large spins (e.g. generated from a supernova collapse, see [36], [16], [37] & [38] for more complete details).

Above a certain rotational frequency, Ω_K, there are no equilibrium states and so the NS must shed some of its angular momentum before it can settle down into a stationary state.

There are two main type of instability:

Type	Growth timescale	Effect on equilibrium
Dynamic (f-modes)	$\sim$ oscillations	Large (may die out quickly)
Secular (r-modes)	$\sim$ viscosity	Subtle

TABLE 2.3: Summary of the types of instability that may occur within Neutron Stars and effect they have on its equilibrium.

The f-mode (or bar-mode) instability essentially grows on the dynamical timescale of the NS, and leads to gravitational waves that may well prove to be observable.

However, the presence of dissipation mechanisms may lead to secular instabilities which grow on a relatively long timescale depending on the magnitude of the dissipation. The main secular instability is associated with viscosity which allows a transition to a lower energy state by violating the conservation of circulation (which holds in a perfect fluid star).

Similarly, rotating NS are generically unstable (r-modes[2]) due to the emission of gravitational waves (which leads to angular momentum of the star not being conserved). That such an instability would operate was first established for the Maclaurin spheroids - and it was shown that this instability is generic, all rotating perfect neutron stars are unstable. See [36] and [37] for a full description of the processes.

2.2.4.2 Mechanism for gravitational-wave instability

Consider a non-rotating star - the forwards and backwards propagating modes will have equal values. However, these two branches of modes are affected by rotation in different ways.

[2]analogous to the Rossby waves in the Earth's atmosphere and oceans, giving them their name

A backwards moving mode will be dragged forwards by the rotation of the NS, and if the star spins sufficiently fast the mode will move forwards with respect to the inertial frame. Meanwhile, the mode is still moving backwards in the rotating frame.

The gravitational waves from such a mode carry positive angular momentum away from the star, but since the perturbed fluid actually rotates slower than it would in absence of the perturbation the angular momentum of the retrograde mode is negative. The emission of gravitational waves consequently makes the angular momentum of the mode increasingly negative and leads to an instability. This class of frame-dragging instabilities is usually refered to as Chandrasekhar-Friedman-Schutz (CFS) instabilities.

The fact that the emission of gravitational radiation causes a growth in the mode energy in the rotating frame, despite the decrease in the inertial frame energy, may at first seem a bit strange. However, it can be understood (see [37] for more details) from the relation between the two energies

$$E_r = E_i - \Omega J \tag{2.17}$$

where E_r is the energy in the rotational mode, E_i is the energy in the inertial mode, J is the angular momentum and Ω is the frequency of rotational. From this we see that E_r may increase if both E_i and J decrease. In other words, when the mode radiates away angular momentum the star can find a rotational state of lower angular momentum *and* lower energy. Under these conditions the mode amplitude may grow.

A myriad of mechanisms have been suggested that could suppress the growth of these modes (e.g. the non-linear effects of mode-mode coupling (see [39]), dissipation at the crust-core boundary (see [40], [41], [42]), exotic bulk viscosity (see [43], [44], [45]) and the effects of magnetic fields ([46], [47], [48]), but there are a similar number that suggest ways that these modes could grow (see for example, [36], [37]). There are currently too many theoretical unknowns for a definitive answer either way, but this is no reason to ignore a potentially strong source of gravitational waves. Especially when observations could be used to distinguish between neutron stars and strange/quark stars [they have very different bulk viscosity which means the r-mode behaves very differently].

Table 2.4 is a summary of the parameters associated with instabilities in neutron stars. We have the event rate, gravitational wave emission values and the expected SNRs for standard and enhanced LIGO. The double line separates the more likely parameters from those about which we have less certainty.

The SNR values in Table 2.4 demonstrates that this would appear to be our candidate for the intermediate duration gravitational wave search. It has a reasonable SNR for LIGO, in an excellent frequency band and with a good length. The frequency of events may seem reasonable, but not every new born NS is expected to have such a high spin. However, this is a good source with plenty of interesting astrophysics waiting to be uncovered.

Event rate	$\sim$ several/year/20MPc
Frequency	$\sim 1kHz$
Emission duration, Δt	~ 150 s
LIGO Noise amplitude	$5 \times 10^{-23}\,\mathrm{Hz}^{-1/2}$
aLIGO Noise amplitude	$5.1 \times 10^{-24}\,\mathrm{Hz}^{-1/2}$
Strain amplitude, h	$\sim 5.6 \times 10^{-22}$
E_{GW}	$\sim 4.8 \times 10^{47} J$
LIGO SNR (distance $20MPc$)	7.9
aLIGO SNR (distance $20MPc$)	77.5

TABLE 2.4: Summary of the properties of Neutron Star instabilities at a distance of the Virgo cluster, $20MPc$. The double line separates the more likely parameters from those about which we have less certainty.

2.2.5 Summary

Table 2.5 summarises the SNRs for each of the sources described above. Although the strongest candidate for a detectable source would appear to be the NS instabilities, there are many theoretical unknowns in the production mechanism and some good luck is needed while the search is taking place(a supernova in our galaxy in the "right" part of the sky). However, just because we go into the search looking for one source, does not mean that other sources will be ignored, just that the search may not be optimally tuned to them. There is also the potential of discovering a whole host of new astrophysical phenomena that lie within this time range. As such, there appear to be enough evidence present to warrant looking for intermediate duration gravitational waves.

Source	LIGO SNR	aLIGO SNR
Microquasar ($1kPc$)	0.9	7.2
Long GRBs ($100MPc$)	0.7	5.3
NS instabilities ($20Mpc$)	7.6	77.5

TABLE 2.5: Summary of SNRs for all sources at their expected distances.

In summary, it would appear that the literature definitely indicates that there may some interesting astrophysics going on in our duration domain of choice. The search is now physically motivated, so we are not searching because we can, but because there is the potential for some exciting discoveries. Note that we are not relying on these models being correct, they are just to give us sufficient motivation to explore this parameter space.

All that is required now is a set of tools to enable the analysis to be carried out. A review of currently used techniques will be beneficial in shaping how the analysis will eventually work, even if things must be done in a new way due to the search space parameters.

Chapter 3

Gravitational Wave Burst Detection

If we take the hypothesis that a gravitational wave has passed through the Earth and affected the detectors, then we need to understand how they would behave, and what the data from them would look like. To this end, this chapter will focus on the data analysis methods that will allow us to extract unmodelled transient signals from the noise.

3.1 Coherent network analysis

Note that this section is adapted from the work by Sutton et al [49].

The output from a detector with an incident gravitational wave with polarization amplitudes, $h_+(t, \vec{x})$ and $h_\times(t, \vec{x})$ from a direction $\hat{\Omega}$, is a linear combination of the signal and noise, n

$$d(t + \Delta t(\hat{\Omega})) = F^+(\hat{\Omega})h_+(t) + F^\times(\hat{\Omega})h_\times(t) + n(t + \Delta t(\hat{\Omega})), \tag{3.1}$$

where $\Delta t(\hat{\Omega})$ is the time delay between the position of the detector ($\vec{r}$) and an arbitrary reference position ($\vec{r}_0$), which could, for example, be the centre of the Earth

$$\Delta t(\hat{\Omega}) = \frac{1}{c}(\vec{r}_0 - \vec{r}) \cdot \hat{\Omega}. \tag{3.2}$$

To carry out the following analysis, we will make a couple of simplifying assumptions and the following statement:

- We will assume that all of the detector data streams have been time-shifted prior to the analysis, so that we can suppress the explicit mention of the delay, $\Delta t(\hat{\Omega})$.
- We will write $h_{+,\times}(t, \vec{r}_0)$ as $h_{+,\times}(t)$.
- Also, note that the data output by the detector is sampled discretely, and, as such, we will be use discrete notation.

If the sampling rate is denoted as f_s, we can convert from continuous to discrete notation as follows:

- $x(t) \rightarrow x[j]$,
- $\tilde{x}(f) \rightarrow f_s^{-1}\tilde{x}[k]$,
- $\int dt \rightarrow f_s^{-1}\sum_j$,
- $\delta(t-t') \rightarrow f_s\delta_{jj'}$, and
- $\delta(f-f') \rightarrow Nf_s^{-1}\delta_{kk'}$.

The discrete Fourier transform, $\tilde{x}[k]$, of a time-series, $x[j]$, with N data points (and its inverse) is

$$\tilde{x}[k] = \sum_{j=0}^{N-1} x[j]e^{\frac{-i2\pi jk}{N}}, \tag{3.3}$$

$$x[j] = \frac{1}{N}\sum_{k=0}^{N-1} \tilde{x}[k]e^{\frac{i2\pi jk}{N}}. \tag{3.4}$$

We can write the one-sided power spectral density, defined as the distribution of average noise power as a function of the frequency, $S_\alpha[k]$, of the noise $\tilde{n}_\alpha$ as

$$\langle \tilde{n}_\alpha^*[k]\tilde{n}_\beta[k']\rangle = \frac{N}{2}\delta_{\alpha\beta}\delta_{kk'}S_\alpha[k]. \tag{3.5}$$

This means we can define the noise-spectrum weighted quantities as follows

$$\tilde{d}_{w\alpha}[k] = \frac{\tilde{d}_\alpha[k]}{\sqrt{\frac{N}{2}S_\alpha[k]}}, \tag{3.6}$$

$$\tilde{n}_{w\alpha}[k] = \frac{\tilde{n}_\alpha[k]}{\sqrt{\frac{N}{2}S_\alpha[k]}}, \tag{3.7}$$

$$F_{w\alpha}^{+,\times}(\hat{\Omega},k) = \frac{F_\alpha^{+,\times}(\hat{\Omega},k)}{\sqrt{\frac{N}{2}S_\alpha[k]}}. \tag{3.8}$$

The normalisation of the whitened noise is given by

$$\langle \tilde{n}_{w\alpha}^*[k]\tilde{n}_{w\beta}[k']\rangle = \delta_{\alpha\beta}\delta_{kk'}. \tag{3.9}$$

Using the above notation, we can now write equation (3.1) in matrix form as follows:

$$\tilde{\boldsymbol{d}} = \boldsymbol{F}\tilde{\boldsymbol{h}} + \tilde{\boldsymbol{n}}. \tag{3.10}$$

The explicit indices for frequency and sky position have been dropped, and the boldface symbols refer to the noise-weighted quantities that are vectors or matrices over the space of the detectors (note that $\tilde{\boldsymbol{h}}$ is not noise weighted), and are a function of frequency and sky position:

$$\tilde{\boldsymbol{d}} \equiv \begin{bmatrix} \tilde{d}_{w1} \\ \tilde{d}_{w2} \\ \vdots \\ \tilde{d}_{wD} \end{bmatrix}, \tilde{\boldsymbol{h}} \equiv \begin{bmatrix} \tilde{h}_{+} \\ \tilde{h}_{\times} \end{bmatrix}, \tilde{\boldsymbol{n}} \equiv \begin{bmatrix} \tilde{n}_{w1} \\ \tilde{n}_{w2} \\ \vdots \\ \tilde{n}_{wD} \end{bmatrix}, \tag{3.11}$$

$$\boldsymbol{F} \equiv \begin{bmatrix} \boldsymbol{F}^{+} & \boldsymbol{F}^{\times} \end{bmatrix} \equiv \begin{bmatrix} F^{+}_{w1} & F^{\times}_{w1} \\ F^{+}_{w2} & F^{\times}_{w2} \\ \vdots & \vdots \\ F^{+}_{wD} & F^{\times}_{wD} \end{bmatrix} = \begin{bmatrix} \frac{F^{+}_{1}}{\sqrt{\frac{N}{2}S_1}} & \frac{F^{\times}_{1}}{\sqrt{\frac{N}{2}S_1}} \\ \frac{F^{+}_{2}}{\sqrt{\frac{N}{2}S_2}} & \frac{F^{\times}_{2}}{\sqrt{\frac{N}{2}S_2}} \\ \vdots & \vdots \\ \frac{F^{+}_{D}}{\sqrt{\frac{N}{2}S_D}} & \frac{F^{\times}_{D}}{\sqrt{\frac{N}{2}S_D}} \end{bmatrix}, \tag{3.12}$$

where D is the number of detectors in the network. Note that the elements of F are functions of the sky position.

3.1.1 Time-frequency decomposition

The most common method of analysing the data from gravitational wave detectors for unmodelled burst searches is to generate a time-frequency map (see [50, 51]).

They are a useful tool for detecting unmodelled bursts for a number of reasons. In particular, many potential gravitational wave signals are band-limited or have significant signal power only at a subset of the frequencies, so converting the data to a time-frequency map allows us to easily separate out frequencies that contain only noise. It is also useful for visualising how different frequency bands evolve over time as the map contain both temporal and spatial information. In this thesis we use the discrete Fast Fourier Transform for time-frequency decompositions; some other analyses use wavelets [52],or sine-Gaussian decompositions [53].

To generate a time-frequency map from a discrete time series, $x[j]$, a Fourier transform length, N_s, must be chosen. There is a balance to be made here between the amount of spectra and temporal resolution required for the analysis. It is also possible to overlap sequential segments by a number of samples (usually $N_s/2$, so 50% overlap). Overlapping is required if the segments are windowed to negate the loss of information at the edges (see Figure 3.1). The procedure to generate a time-frequency map is as follows:

1. Select the first N_s of N samples.

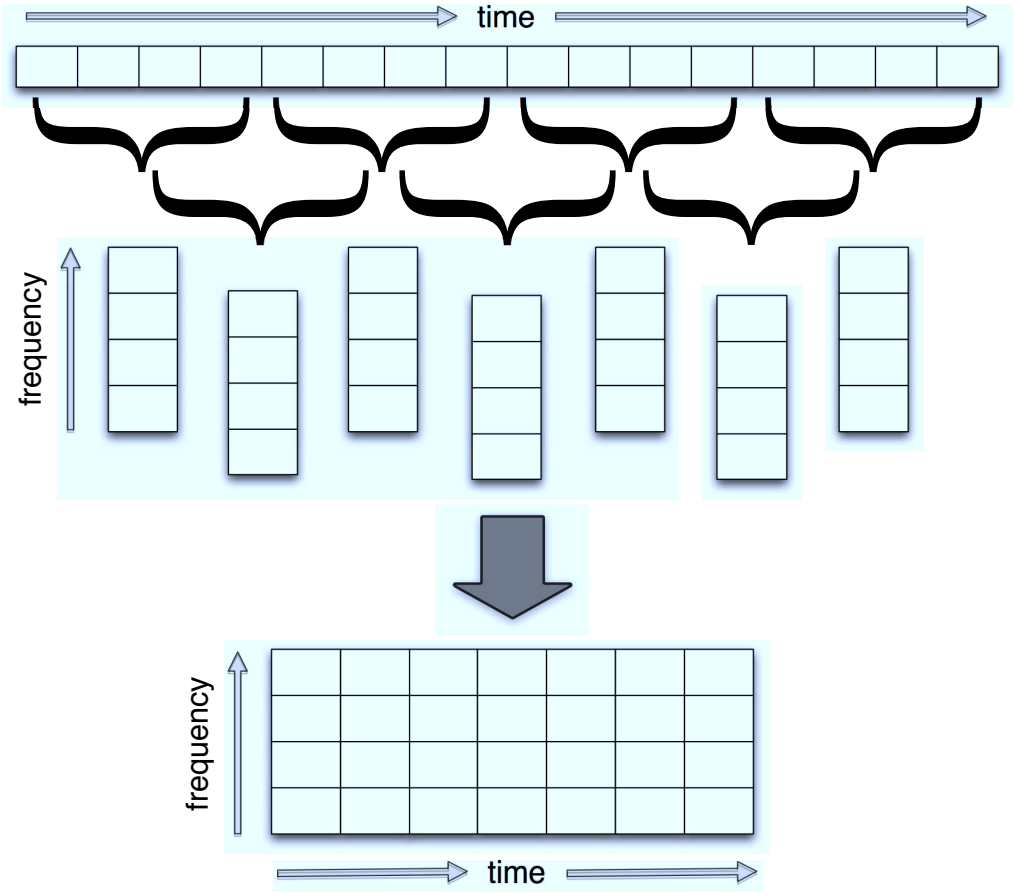

FIGURE 3.1: A visual representation of the generation of a time-frequency map from a time series. Intervals of length N_s samples are selected in turn. Each is FFTed to produce a vector of N_s frequency domain samples (columns). FFTed data from consecutive intervals are assembled as the columns of a matrix, representing the time-frequency map. Each box represents one time sample (top image) or time-frequency sample (lower images). In this example, $N_s = 4$ with a 50 % overlap.

2. Window the samples with an appropriate function to minimise the amount of frequency bin leakage.

3. Fourier transform (i.e. use the Fast Fourier Transform (FFT), equation (3.3)) these samples to produce N_s complex coefficients (of which only the first $N_s/2$ are useful, the rest being the complex conjugate of these first values). These values represent the frequencies from 0 to $f_s/2$ (the Nyquist frequency) separated by f_s/N_s Hz, where f_s is the sampling rate of the timeseries.

4. Select the next N_s samples, overlapping if necessary and repeat steps 2 - 4, until the N samples are exhausted.

5. FFTed data from consecutive intervals are assembled as the columns of a matrix, representing the time-frequency map. Individual elements of this matrix are referred to as "pixels"

Time-frequency maps can be made from individual or combined ("coherent") detector data streams so that a host of different statistics can be generated. In the next section we look at some of the simpler coherent likelihoods, those computed from projections of the data. The simplest of these is the *standard*, or *maximum likelihood*.

3.1.2 Standard Likelihood

If we denote $P(\tilde{\boldsymbol{d}}|\tilde{\boldsymbol{h}})$ as the probability of obtaining the whitened data, $\tilde{\boldsymbol{d}}$, in a time-frequency pixel in the presence of a known gravitational wave, $\tilde{\boldsymbol{h}}$, and from a known direction, then, assuming Gaussian noise, we can write

$$P(\tilde{\boldsymbol{d}}|\tilde{\boldsymbol{h}}) = \frac{1}{(2\pi)^{D/2}} \exp\left[-\frac{1}{2}\left|\tilde{\boldsymbol{d}} - \boldsymbol{F}\tilde{\boldsymbol{h}}\right|^2\right]. \tag{3.13}$$

If we assume that the signal occupies a set $\{\tilde{\boldsymbol{d}}\}$ of N_P time-frequency pixels, then

$$P(\{\tilde{\boldsymbol{d}}\}|\{\tilde{\boldsymbol{h}}\}) = \frac{1}{(2\pi)^{N_P D/2}} \exp\left[-\frac{1}{2}\sum_k \left|\tilde{\boldsymbol{d}}[k] - \boldsymbol{F}[k]\tilde{\boldsymbol{h}}[k]\right|^2\right], \tag{3.14}$$

where k indexes the pixels. The log-likelihood ratio, L, is then defined by the log-ratio of this probability to the corresponding probability under the null hypothesis (i.e. probability of observing the data $\{\tilde{\boldsymbol{d}}\}$, when there is no signal present, $\tilde{\boldsymbol{h}} = 0$, so $P(\{\tilde{\boldsymbol{d}}\}|\{0\})$)

$$L \equiv \log \frac{P(\{\tilde{\boldsymbol{d}}\}|\{\tilde{\boldsymbol{h}}\})}{P(\{\tilde{\boldsymbol{d}}\}|\{0\})} = \frac{1}{2}\sum_k \left[\left|\tilde{\boldsymbol{d}}[k]\right|^2 - \left|\tilde{\boldsymbol{d}}[k] - \boldsymbol{F}[k]\tilde{\boldsymbol{h}}[k]\right|^2\right]. \tag{3.15}$$

Generally, the signal waveform, $\tilde{\boldsymbol{h}}$, is not known, so it is not clear how to calculate the likelihood ratio in equation (3.15). One approach is to treat the waveform values $\tilde{\boldsymbol{h}} = (\tilde{h}_+, \tilde{h}_\times)$ in each t-f pixel as free parameters to be fit to the data. Then, the best fit values $\tilde{\boldsymbol{h}}_{\boldsymbol{max}}$ are those that maximises the likelihood ratio

$$0 = \left.\frac{\partial L}{\partial \tilde{\boldsymbol{h}}}\right|_{\tilde{\boldsymbol{h}}=\tilde{\boldsymbol{h}}_{max}}. \tag{3.16}$$

Naively, we would assume L to be maximum when $\left|\tilde{\boldsymbol{d}} - \boldsymbol{F}\tilde{\boldsymbol{h}}\right|^2$ is zero. This occurs when $\boldsymbol{F}\tilde{\boldsymbol{h}} = \tilde{\boldsymbol{d}}$, meaning $\tilde{\boldsymbol{h}} = \boldsymbol{F}^{-1}\tilde{\boldsymbol{d}}$. As $\boldsymbol{F}$ is not a square matrix, it does not have a well defined inverse, however, we can use the *pseudo-inverse* or *Moore-Penrose inverse*, $\boldsymbol{F}^{-1}_{MP}$.

$$\boldsymbol{F}^{-1}_{MP} \equiv (\boldsymbol{F}^\dagger \boldsymbol{F})^{-1}\boldsymbol{F}^\dagger \tag{3.17}$$

Here, † denotes the conjugate transpose. This was introduced by Rakhmanov [54]. For a standard linear polarisation basis such as set out in equations (1.56), (1.57), $\boldsymbol{F}$ and $\boldsymbol{F}^{-1}_{MP}$ are real. For other bases, such as the circularly polarisation bases of equations (1.58), (1.59), $\boldsymbol{F}$ and $\boldsymbol{F}^{-1}_{MP}$ may be complex.

Thus, the $\tilde{\boldsymbol{h}}_{max}$ that maximizes the likelihood ratio is

$$\tilde{\boldsymbol{h}}_{max} = \boldsymbol{F}_{MP}^{-1}\tilde{\boldsymbol{d}}. \tag{3.18}$$

Substituting this into equation (3.15) gives the *standard likelihood*:

$$E_{SL} \equiv 2L(\tilde{\boldsymbol{h}}_{max}) = \sum_k \tilde{\boldsymbol{d}}^* \boldsymbol{F}\boldsymbol{F}_{MP}^{-1}\tilde{\boldsymbol{d}}, \tag{3.19}$$

where we have used the fact that $\boldsymbol{F}\boldsymbol{F}_{MP}^{-1}$ is Hermitian.

3.1.3 Projection Operator and the Null Energy

$\boldsymbol{F}\boldsymbol{F}_{MP}^{-1}$ projects the data into the subspace spanned by $\boldsymbol{F}^+$ and $\boldsymbol{F}^\times$. The contribution to $\tilde{\boldsymbol{d}}$ by a gravitational wave from a specific sky position is restricted to this subspace (see, for example, equation (3.10)). Therefore, the standard likelihood is the maximum amount of energy in the whitened data that is consistent with the hypothesis of a gravitational wave from a given sky position - or more specifically, since it is defined in terms of noise-weighted data, the maximum possible squared signal-to-noise ratio, ρ^2.

Compare this with the *total energy* in the data

$$E_{tot} = \sum_k \left|\tilde{\boldsymbol{d}}\right|^2, \tag{3.20}$$

which is an incoherent statistic in that it only contains autocorrelation terms, and no-cross correlation terms. This is the quantity calculated for each time-frequency pixel in an excess power search (refer to Section 3.2.2).

The orthogonal projection operator ($\boldsymbol{I} - \boldsymbol{F}\boldsymbol{F}_{MP}^{-1}$) cancels the gravitational wave signal, yielding the *null stream* with an energy

$$E_{null} \equiv E_{tot} - E_{SL} = \sum_k \tilde{\boldsymbol{d}}^* (\boldsymbol{I} - \boldsymbol{F}\boldsymbol{F}_{MP}^{-1})\tilde{\boldsymbol{d}} \tag{3.21}$$

This null energy is the minimum amount of energy in the whitened data that is inconsistent with the hypothesis of a gravitational wave from a specific sky location.

An advantage of the coherent analysis is that projecting from the full data space, with energy E_{tot} down to the subspace spanned by $\boldsymbol{F}^+$, $\boldsymbol{F}^\times$, with energy E_{SL} removes noise, with energy E_{null}, without removing any of the signal.

3.1.4 Dominant Polarization Frame

For a time-frequency pixel, the data from the D detectors is a vector in a D-dimensional complex space. A basis for this space is formed by the set of single-detector strains (the basis used so far), however, this is not the best basis for writing detection statistics. Alternatively, we could use the 2-dimension subspace defined by $\boldsymbol{F}^+$ and $\boldsymbol{F}^\times$. In this space, there is a direction in which the detector network has the maximum antenna response, and an orthogonal direction in which it has minimum response. Thus, these directions can be used as basis vectors along with an orthonormal basis for the null space.

To simplify things further, it is possible to define the + and × polarizations so that $\boldsymbol{F}^+$ lies along the first basis vector, and the $\boldsymbol{F}^\times$ along the second. This choice is called the dominant polarization frame (DPF).

To construct the DPF, recall that the antenna response vectors of two frames separated by a polarization angle ψ are related by

$$\boldsymbol{F}^+(\psi) = \cos 2\psi \boldsymbol{F}^+(0) + \sin 2\psi \boldsymbol{F}^\times(0), \tag{3.22}$$

$$\boldsymbol{F}^\times(\psi) = -\sin 2\psi \boldsymbol{F}^+(0) + \cos 2\psi \boldsymbol{F}^\times(0). \tag{3.23}$$

For any direction on the sky, it is possible to choose a polarization frame such that $\boldsymbol{F}^+(\psi)$ and $\boldsymbol{F}^\times(\psi)$ are orthogonal and $|\boldsymbol{F}^+(\psi)| > |\boldsymbol{F}^\times(\psi)|$. Given $\boldsymbol{F}^+(0)$ and $\boldsymbol{F}^\times(0)$ in the original polarization frame, the rotation angle to the DPF is

$$\psi^{DPF}(\hat{\Omega}, k) = \frac{1}{4}\,\mathrm{atan2}\left(2\boldsymbol{F}^+(0)\cdot\boldsymbol{F}^\times(0), |\boldsymbol{F}^+(0)|^2 - |\boldsymbol{F}^\times(0)|^2\right), \tag{3.24}$$

where atan2 is the arctangent function with range $[-\pi, \pi]$, and is a function of sky position and frequency due to the noise weighting of $\boldsymbol{F}^+$ and $\boldsymbol{F}^\times$.

Denoting the antenna response vectors in the DPF as $\boldsymbol{f}^+$ and $\boldsymbol{f}^\times$, they have the following properties

$$|\boldsymbol{f}^+|^2 \geq |\boldsymbol{f}^\times|^2, \tag{3.25}$$

$$\boldsymbol{f}^+ \cdot \boldsymbol{f}^\times = 0. \tag{3.26}$$

The following unit vectors are also part of an orthonormal coordinate system,

$$\boldsymbol{e}^+ \equiv \frac{\boldsymbol{f}^+}{|\boldsymbol{f}^+|},\ \boldsymbol{e}^\times \equiv \frac{\boldsymbol{f}^\times}{|\boldsymbol{f}^\times|}. \tag{3.27}$$

However, the DPF does not specify any particular choice of basis for the null space, so any convenient choice can be made for that subspace. In the DPF, $\boldsymbol{F}\boldsymbol{F}_{MP}^{-1}$ takes on the simple

form of

$$\boldsymbol{F}\boldsymbol{F}_{MP}^{-1} = \boldsymbol{e}^{+}\boldsymbol{e}^{+\dagger} + \boldsymbol{e}^{\times}\boldsymbol{e}^{\times\dagger}, \tag{3.28}$$

and the standard likelihood (equation 3.19) becomes

$$E_{SL} = \sum_k \left[\left| \boldsymbol{e}^{+} \cdot \boldsymbol{d} \right|^2 + \left| \boldsymbol{e}^{\times} \cdot \boldsymbol{d} \right|^2 \right]. \tag{3.29}$$

As well as the standard likelihood, the convenience of writing detection statistics in the DPF allows us to define more. For instance, we can write the energy in the h_+ and $h_\times$ polarizations in the DPF as:

$$E_{+} \equiv \sum_k |\boldsymbol{e}^{+} \cdot \boldsymbol{d}|^2, \qquad \text{(plus energy)} \tag{3.30}$$

$$E_{\times} \equiv \sum_k |\boldsymbol{e}^{\times} \cdot \boldsymbol{d}|^2. \qquad \text{(cross energy)} \tag{3.31}$$

3.1.5 Statistics

The likelihoods $E_{\rm sl}, E_+, E_\times, , E_{\rm null}$ and $E_{\rm tot}$ are all defined in terms of projections of the data, and their statistical properties in Gaussian background noise are very simple - they follow a non-central χ^2 distribution. Specifically,

$$2E \sim \chi^2_{2N_p D_{proj}}(\lambda), \tag{3.32}$$

for a set of time-frequency pixels and a specified sky position. Here, N_p is the number of pixels, D_{proj} is the number of dimensions of the projection[1] and λ is the non-centrality parameter. The non-central χ^2 distribution has a mean of $(2N_pD_{proj} + \lambda)$ and a standard deviation of $\sqrt{2N_pD_{proj}}$ Thus, a signal should be detectable when

$$\frac{\lambda}{\sqrt{2N_pD_{proj}}} \gg 1. \tag{3.33}$$

λ, (the non-centrality parameter), is the expected squared signal-to-noise ratio of a matched filter for the waveform restricted to the time-frequency region in question and after projection by the appropriate likelihood projection operator, summed over the network (positive frequencies only - double the non-centrality parameter when including negative frequencies, $\lambda \to 2\rho$).

[1] 1 for $E_{+,\times}$, 2 for E_{SL}, D-2 for E_{null} and D for E_{tot}

Explicity

$$\lambda_+ = 2\sum_k |\boldsymbol{f}^+|^2|h_+|^2 = \frac{4}{N}\sum_\alpha\sum_k \frac{|F_\alpha^+(\psi_{DPF})\tilde{h}_+[k](\psi_{DPF})|^2}{S_\alpha[k]} =: \rho_+^2, \tag{3.34}$$

$$\lambda_\times = 2\sum_k |\boldsymbol{f}^\times|^2|h_\times|^2 = \frac{4}{N}\sum_\alpha\sum_k \frac{|F_\alpha^\times(\psi_{DPF})\tilde{h}_\times[k](\psi_{DPF})|^2}{S_\alpha[k]} =: \rho_\times^2, \tag{3.35}$$

$$\begin{aligned}\lambda_{SL} &= 2\sum_k \left[|\boldsymbol{f}^+|^2|h_+|^2 + |\boldsymbol{f}^\times|^2|h_\times|^2\right] \\ &= \frac{4}{N}\sum_\alpha\sum_k \frac{|F_\alpha^+\tilde{h}_+[k] + F_\alpha^\times\tilde{h}_\times[k]|^2}{S_\alpha[k]} =: \rho^2,\end{aligned} \tag{3.36}$$

$$\lambda_{tot} = \rho^2, \tag{3.37}$$

$$\lambda_{null} = 0. \tag{3.38}$$

Note that the antenna responses and waveforms for $\lambda_{+,\times}$ are defined in the DPF while equation (3.36), λ_{SL}, is independent of the polarization basis used.

3.2 Searching for transient signals

The generation and production of gravitational waves, and the physical requirements to detect them have been covered, but as real signals exist around the level of the noise, we need a way to extract them from the data. One method is to filter the data with the signal you are looking for (this is known as matched filtering), however many astrophysical phenomena are too complex, or the physics is too ill defined, to be precisely modelled (e.g. core collapse of high mass star in a supernova or accretion induced collapse of white dwarf.) Here, it is becomes necessary to use a technique that requires very little knowledge of the signal.

3.2.1 Matched filtering

The detection of transient signals comprising a known waveform in stationary Gaussian noise has an optimal strategy - the use of a matched filter (see [55]). Usually the waveform is not completely fixed, its form is specified by a set of parameters. Let $n(t)$ refer to a stationary Gaussian noise process, with one-sided power spectral density, $S_n(f)$, given by $\langle\tilde{n}(f)\tilde{n}^*(f')\rangle = \frac{1}{2}S_n(|f|)\delta(f-f')$, and $s(t)$ to be a data stream. The data stream may contain only detector noise, it which case $s(t) = n(t)$, or it may contain a signal in addition to the noise, so $s(t) = n(t) + h(t)$.

In either case, the matched filter output of $s(t)$ with a filter template $h_{template}(t)$ (i.e. the signal were are searching for) is

$$x(t) = 4\mathrm{Re}\int_0^\infty \frac{\tilde{s}(f)\tilde{h}^*_{template}(f)}{S_n(f)} e^{2\pi i f t} df. \tag{3.39}$$

This strategy is not as clean as it looks, as there are additional unknown parameters that can be used to describe the waveform (see for example [56]). To search for all of the possible waveforms, one constructs a 'bank' of matched filter templates to cover the entire parameter space. These are usually structured such that any real signal would have a waveform that would would match a template with a similar set of parameters, (see for example [57], [58], [59], [60], [61]).

This is a very powerful way to detect known signals, but in our search we will have very little knowledge about the signal waveform so this method cannot be used. Instead, unmodelled burst searches use the techniques of excess power detection which is sensitive to generic transient signals without requiring a template bank.

Also, by not restricted ourselves too much in the form of the waveform, the search has less chance of missing out on something interesting just because it doesn't match.

As such, we need a more generic search tool, something that can work with the bare minimum of signal knowledge.

3.2.2 Excess power detection

If it is possible to make assumptions about the approximate duration and frequency band of a astrophysical event, then this can be used to attempt to find the signal. Simply, the total power is computed for a "pixel" of size $\delta t \times \delta f$, and repeated for a range of start times. A detection can be claimed when the total power is significantly greater than that expected from detector noise alone. This "excess power" method is optimal in the absence of any other information about the signal, as shown by Anderson et al [51]. In fact, for signals with small time-frequency volumes (e.g. short duration occupying a narrow frequency band, $\delta t \times \delta f \sim 1$) it performs as well as matched filtering.

The technique works by exploiting the fact that the power distribution of stationary Gaussian noise follows a χ^2 distribution with the number of degree of freedom being twice the time-frequency volume as discussed in Section 3.1.5. When a signal is present, the power follows a non-central χ^2 distribution with the non-centrality being given by the signal power. A signal is therefore detectable when the excess power is much greater than the noise power fluctuations, which scale as the square root of the time-frequency volume. This limits the utility of the

excess power method for astrophysical phenomena that have large time-frequency volumes, but it is approximately as powerful as matched filtering in the limit of $\delta t \times \delta f \sim 1$.

3.2.2.1 Algorithm

The actual implementation of the excess power analysis is reasonably simple, and uses the following algorithm:

1. Choose a start time, t_s, a duration, δt (comprising N samples) and a frequency band $[f_s, f_s + \delta f]$.
2. Fast Fourier transform the time domain data from $[t_s \rightarrow t_s + \delta t]$.
3. Sum the power across the chosen frequency band.
4. Calculate the probability of obtaining this summed power from Gaussian noise alone.
5. Record a detection if the probability is significant.
6. Repeat the above procedure for all t_s, δt, f_s and δf.

The excess power algorithm can be thought of as a brute force method of determining the existence of a signal in the data. For example, in a short transient search δt might have values in the range 1/8 to 1/256 of a second, while $\delta f \geq 1/\delta t$ could be anywhere from 8 Hz up to the Nyquist frequency. The values of δt and δf must be chosen with care. If they are too large, then the signal to noise ratio will be bad due to the large number of pixels without a signal in them. Generally δt and δf should be chosen to be as close a match to the signal length and bandwidth as possible to give the best signal to noise ratio.

Note that while the excess power method can distinguish the presence of a potential signal in a time-frequency region, the algorithm proposed by Anderson et al. [51] does so only for a single detector. Thus, unless detectors are used in combination, the pipeline could be susceptible to loud glitches in the data stream. This "coherent analysis", introduced in Section 3.1, makes use of all of the detector data to produce statistics that comment on the viability of an event, as discussed further in Section 3.3 below.

To be clear, the original implementation proposed in [51] computed excess power in each detector data stream separately. Current applications (e.g. coherent WaveBurst [52], X-Pipeline [49]) differ in two important ways: 1) An excess power-like statistic is computed on the coherent combination of data from multiple detectors (e.g. X-Pipeline computes excess power on the h_+ and $h_\times$ data streams). However, they are still vulnerable to glitches, which typically cause excess power in these coherent data stream as well. 2) They do not use fixed

rectangular blocks, $\delta t \times \delta f$, on which to compute the excess power. Instead they apply a threshold to the time-frequency map and select clusters of loud pixels over which to sum the excess power. Clustering will be discussed in more detail in Section 4.4.1.

3.3 Glitch rejection

Up to now, the likelihood measures have been motivated under the assumption of stationary Gaussian noise. Obviously, this is not a valid assumption for real detector noise, as it contains *glitches*, short transients of excess strain that can appear to be gravitational-wave bursts. Therefore, a method of distinguishing noise glitches from true signals is necessary.

One can use the autocorrelation component of the coherent energies to construct tests that are effective in rejecting glitches (see [49], [62]). For instance, consider the null space, the subspace orthogonal to that used to define the standard likelihood. Projecting $\boldsymbol{d}$ onto this subspace contains only noise, thus the presence/absence of a GW should not affect the projection. Glitches do not couple to the data streams with any particular relationship to $\boldsymbol{F}^{+}$or $\boldsymbol{F}^{\times}$ so they are generally present in the null space projection. Thus, for a transient to be considered a GW signal, the energy in the null stream should be small.

To construct an effective test, first write equation 3.21 as:

$$E_{null} = \sum_{k} \tilde{\boldsymbol{d}}^{*} (\boldsymbol{I} - \boldsymbol{F}\boldsymbol{F}_{MP}^{-1}) \tilde{\boldsymbol{d}} \tag{3.40}$$

$$= \sum_{k} \sum_{\alpha,\beta} \tilde{d}_{\alpha}^{*} P_{\alpha\beta}^{null} \tilde{d}_{\beta} \tag{3.41}$$

where $\boldsymbol{P}^{null} = \boldsymbol{I} - \boldsymbol{F}\boldsymbol{F}_{MP}^{-1}$ is the null stream projection operator. The null energy is composed of cross-correlation terms $\tilde{d}_{\alpha}^{*}\tilde{d}_{\beta}$ and auto-correlation terms $\tilde{d}_{\alpha}^{*}\tilde{d}_{\alpha}$. Thus, if the transient is not correlated across detectors (as expected for a glitch), then the cross-correlation terms will be small compared to the auto-correlation terms. For glitches, the null energy will be dominated by the auto-correlation components, so:

$$E_{null} \simeq I_{null} \equiv \sum_{k} \sum_{\alpha} P_{\alpha\alpha}^{null} |\tilde{d}_{\alpha}|^{2}. \tag{3.42}$$

This auto-correlation part of the null energy is called the *incoherent energy*.

For a GW signal, the transient is correlated between the detectors, which cancel in the null stream, leaving only Gaussian noise. However, they cannot cancel in I_{null} as that is an incoherent statistic. So, for a strong GW signal,

$$E_{null} \ll I_{null}. \tag{3.43}$$

This suggests that we keep transients with $I_{null}/E_{null} >$ C, where $C > 1$. There are analogous tests for the other coherent energies, ($E_+, E_\times$ etc), which have the following incoherent energies

$$I_{+,\times} \equiv \sum_k \sum_\alpha |e_\alpha^{+,\times} \tilde{d}_\alpha|^2, \tag{3.44}$$

$$I_{SL} \equiv \sum_k \sum_\alpha \left[\left| e_\alpha^{+} \tilde{d}_\alpha \right|^2 + \left| e_\alpha^{\times} \tilde{d}_\alpha \right|^2 \right] = I_+ + I_\times. \tag{3.45}$$

For a glitch, the coherent energy, E, should be roughly equal to the incoherent energy, I, $E \simeq I$. Whereas for a strong GW, summing the signal over both polarizations should build coherently, so

$$E_{SL} > I_{SL}. \tag{3.46}$$

However, depending on the polarization of the GW signal, $E_+ > I_+$ or $E_+ < I_+$ (and similar for $E_\times$ and $I_\times$). For instance, if the signal is predominantly in the h_+ polarization, then

$$\begin{aligned} E_+ &> I_+, \\ E_\times &< I_\times, \end{aligned} \tag{3.47}$$

and if the signal is predominantly in the $h_\times$ polarization, then

$$\begin{aligned} E_+ &< I_+, \\ E_\times &> I_\times. \end{aligned} \tag{3.48}$$

As such, a GW will typically be characterised by at least one of the above - a build up of coherent energy in E_+, $E_\times$ or both.

The techniques summarised in this Chapter all have one thing in common - they are dependent on the sky position. For instance the antenna factors, $\boldsymbol{F}_+$ and $\boldsymbol{F}_\times$, have a dependence on sky position, Ω, (recall equation (3.1)). Similarly, the time shifts applied in forming a coherent combination of the detector data streams also depend on sky position.

It is this dependence on sky position that makes these techniques difficult to broaden to an all-sky coherent search effectively. Calculating a statistic over the sky is computationally expensive due to the fact that each sky point must be treated separately. It we could separate out this dependence on Ω then it would trivialise the calculations. We will describe one such technique for doing this in the next Chapter.

However, as a baseline, we will analyse the performance of running X-Pipeline all-sky (approximately 100 sky positions in real time on a single CPU). The sky has $\mathcal{O}(10^4)$ resolvable sky positions, so we would need $\mathcal{O}(100)$ CPUs to process data in real time. However, the dominant cost is the background estimation, which requires reprocessing the data with unphysically large time shifts (¿1seconds) between detectors. We require $\mathcal{O}(10^3)$ time shifts to

be able to assign false alarm rate probabilities at the 0.001 level. So nominally we would need $\mathcal{O}(10^5)$ CPUs dedicated to an all-sky search.

We could obviously make some efficiency savings by avoiding redundant calculations (e.g. computing the FFT once for each sky position and reusing it for different time shifts) but it is clear that very large improvements in speed is needed.

Chapter 4

The Spherical Radiometer

4.1 Current Limitations

Gravitational wave interferometers have an approximately omnidirectional response, so, for an untriggered search (i.e. one in which we do not already have some idea of the sky location), the standard procedure is to carry out the analysis for a multitude of grid points across the sky (see X-Pipline [49], cWB [63], STAMP [64] and Omega [65]). For a coherent analysis, in which the data must be time delayed by the appropriate amount prior to combining, this is a slow process.

Thus, the speed of the search is directly proportional to the granularity of the search pattern (i.e. the number of independent search positions). This can be circumvented somewhat by utilising the fact that each sky position calculation is independent and can therefore be easily parallelised.

However, these are merely work arounds, and what a coherent all-sky search actually requires is a method that can quickly process the data independently of sky position.

However, it is pertinent to note that coherent likelihoods **are** correlated across the sky - along directions in which the time-of-flight delay between two detectors is constant (see Figure 4.1). This correlation can be taken advantage of to minimise the computational burden.

4.2 Spherical Radiometry

We note from the discussions in Section 3 that coherent likelihoods can be constructed as a sum of cross-correlations and auto-correlations between the detector data streams. A computationally efficient method of correlating pairs of data streams over the whole sky has been developed by Cannon [66]. In this approach the correlation is performed in the spherical

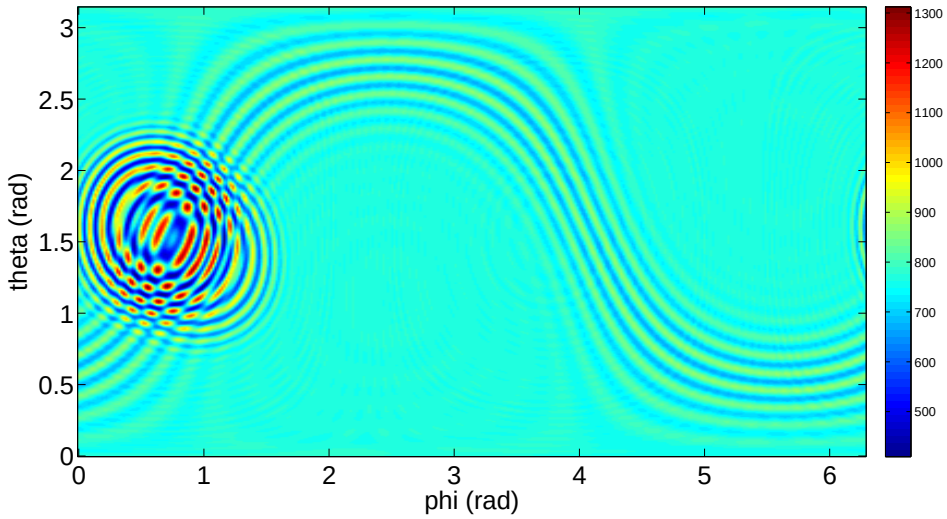

FIGURE 4.1: In this figure we see a skymap for a network of 3 detectors. For each pair of detectors it is possible to see the rings of constant delay - the crossing of these rings is the most probable sky direction of the signal.

harmonic domain (see Appendix C for an overview and the properties of spherical harmonic domain). This allows the correlation between detectors in a network to be performed extremely quickly - the output is a value corresponding to the amount of similarity between them (see [66] and [67]).

To generate the cross-correlation for each detector pair in our network, we need to compute the integral given in (4.1) (following the nonclementure giving in [66]):

$$\xi_{1,2}(\hat{s}) = \int_{t-T/2}^{t+T/2} g_1(t' - \vec{r}_1 \cdot \hat{s}) g_2(t' - \vec{r}_2 \cdot \hat{s}) dt'. \tag{4.1}$$

We can efficiently compute this integral by using the fact that the dependence on sky position,$\hat{s}$, can be separated out. Assuming that the data $g_{1,2}(t)$ has been discretely sampled (spacing Δt) and can be represented as a vector, $\vec{g}$:

$$\vec{g} = \begin{bmatrix} \vdots \\ g(t_j) \\ \vdots \end{bmatrix}, \tag{4.2}$$

where $t_j = j\Delta t$, the j^{th} element of the vector, we can write the linear transform (corresponding to the delay of the time series due to a arbitrary sky direction):

$$g(t) \to g(t - \vec{r} \cdot \hat{s}) \tag{4.3}$$

as

$$\vec{g}(\hat{s}) = \mathbb{T}(\vec{r} \cdot \hat{s}) \cdot \vec{g}, \tag{4.4}$$

where $\mathbb{T}(\vec{r}\cdot\hat{s})$ is the matrix that retards the time series according to the sky location. The integral (4.1) can now be written as an inner product

$$\begin{aligned}\xi_{1,2}(\hat{s}) &= \frac{1}{N}\vec{g}_1^T(\hat{s})\cdot\vec{g}_2(\hat{s}) \\ &= \frac{1}{N}\vec{g}_1^T\cdot\mathbb{T}^T(\vec{r}_1\cdot\hat{s})\cdot\mathbb{T}(\vec{r}_2\cdot\hat{s})\cdot\vec{g}_2.\end{aligned} \tag{4.5}$$

Note that the dependence on sky position, $\hat{s}$, has been decoupled from the data, $\vec{g}_i$, and folded into the two delay matrices, $\mathbb{T}$. By projecting these onto some choice of basis and computing expansion terms, their dependence, and that of $\xi_{1,2}$, can be encoded.

The choice of basis for the vectors $\vec{g}_i$ is determined by how close to diagonal and time independent $\mathbb{T}^T(\vec{r}_1\cdot\hat{s})\cdot\mathbb{T}(\vec{r}_2\cdot\hat{s})$ can be made, which means that the fewest coefficients can be computed.

For example, in an Earth-fixed coordinate system, $\mathbb{T}^T(\vec{r}_1\cdot\hat{s})\cdot\mathbb{T}(\vec{r}_2\cdot\hat{s})$ is independent of time, and taking the Fourier transform of $\vec{g}_i$ makes it exactly diagonal, and we could use the spherical harmonic basis. $\xi_{ij}(\hat{s})$ is a band-limited function on the sphere so can be accurately represented by an expansion to finite order. Note that this basis does not require the fewest number of coefficients to represent a general $\xi_{ij}(\hat{s})$, but the advantages (orthogonality and efficient computational algorithms) outweigh this.

Of course, the issue with an Earth-fixed coordinate system is that due to the rotation of the Earth a source at a static location on the celestial sphere will travel through a range of different directions, $\hat{s}$. To avoid this being an issue, the integration length should be small enough such that the angle moved through is smaller than the angular resolution of the network, or, to put it another way, this is not a problem when searching for signals with duration small enough that the rotation of the Earth over this time is smaller than the angular resolution of the network.

4.2.1 Geometric delay

An Earth-fixed coordinate system makes the position of the i^{th} instrument,$\vec{r}_i$, independent of time. In this framework, we can write the delay due to sky location as the functional, $\mathcal{T}$, so:

$$\mathcal{T}[g_i(t);\vec{r}_i\cdot\hat{s}] = g_i(t-\vec{r}_i\cdot\vec{s}) \tag{4.6}$$

Noting that the discretely sampled time series (where k is the k^{th} element of the vector), $g_i(t_k) = g_i(k\Delta t)$, $k\in\{-\infty,\ldots,0,\ldots,+\infty\}$ is obtained from a band limited function. So, by using the sampling theorem, we can obtain a continuous-time function, $g_i(t)$, by convolving with the sinc function, $\text{sinc}(x) = \sin(\pi x)/(\pi x)$:

$$g_i(t) = \sum_{k=-\infty}^{\infty} g_i(k\Delta t)\text{sinc}[(t - k\Delta t)/\Delta t] \tag{4.7}$$

Applying the geometric delay functional to (4.7), it becomes:

$$\mathcal{T}[g_i(t); \vec{r}_i \cdot \hat{s}] = \sum_{k=-\infty}^{\infty} g_i(k\Delta t)\text{sinc}[(t - \vec{r}_i \cdot \hat{s} - k\Delta t)/\Delta t] \tag{4.8}$$

The components of the matrix form of the geometric delay operator can be obtained by sampling (4.8) at $t_j = j\Delta t$ (where j is the j^{th} element of the vector), to construct a new discretely sampled time series:

$$\begin{aligned} g_i(j\Delta t; \hat{s}) &= \mathcal{T}[g_i(j\Delta t); \vec{r}_i \cdot \hat{s}] && (4.9)\\ &= \sum_{k=-\infty}^{\infty} g_i(k\Delta t)\text{sinc}[(j\Delta t - \vec{r}_i \cdot \hat{s} - k\Delta t)/\Delta t] && (4.10)\\ &= \sum_{k=-\infty}^{\infty} \mathbb{T}_{jk}(\vec{r}_i \cdot \hat{s}) g_i(k\Delta t) && (4.11) \end{aligned}$$

Therefore,

$$\mathbb{T}_{jk}(\vec{r}_i \cdot \hat{s}) = \text{sinc}(j - \vec{r}_i \cdot \hat{s}/\Delta t - k) \tag{4.12}$$

where $j, k \in \{-\infty, \ldots, 0, \ldots, +\infty\}$. In general, this matrix is nonzero everywhere, but the magnitude of the elements in each row decay proportionally to their distance from the diagonal. $\mathbb{T}$ can be approximated by choosing a number of elements, N_τ, centred on the diagonal to keep within each row, while setting the rest to zero. Thus, the j^{th} sample of the delayed time series for the direction $\hat{s}$ is:

$$g_i(j\Delta t; \hat{s}) \approx \sum_{k=j-(N_\tau-1)/2}^{j+(N_\tau-1)/2} \mathbb{T}_{jk}(\vec{r}_i \cdot \hat{s}) g_i(k\Delta t) \tag{4.13}$$

where N_τ is a positive odd integer.

Now, the geometrically delayed time series, (4.8), is obtained by convolution with a sinc interpolating kernel. This can be implemented using Fast Fourier Transform (FFT) techniques (see equation 3.3 in Chapter 3) to improve its speed (the cross correlation is superficially an operation of order $O(n^2)$, whereas the FFT is $O(n \log n)$)). This modification leads to the frequency-domain correlator, whose input is a frequency series rather than the time-series stated above. Consider a pair of discretely sampled time series with N samples and periodicity in N. From (4.12), it is clear that the elements of $\mathbb{T}_{jk}$ only depend on $j - k$, and can be writen

as $\mathbb{T}_{jk} = \mathbb{T}_{j-k}$; swapping out time for an integer index, $k\Delta t \rightarrow k$, (4.13) can be writen:

$$\begin{bmatrix} g_i(0;\vec{s}) \\ \vdots \\ g_i(j;\vec{s}) \\ \vdots \\ g_i(N-1;\vec{s}) \end{bmatrix} = \begin{bmatrix} \mathbb{T}_0(\vec{r}_1\cdot\hat{s}) & \mathbb{T}_{-1}(\vec{r}_1\cdot\hat{s}) & \dots & \dots & \mathbb{T}_{+1}(\vec{r}_1\cdot\hat{s}) \\ & & \vdots & & \\ & \dots & \mathbb{T}_{j-k}(\vec{r}_1\cdot\hat{s}) & \dots & \\ & & \vdots & & \\ \mathbb{T}_{-1}(\vec{r}_1\cdot\hat{s}) & \dots & \dots & \mathbb{T}_{+1}(\vec{r}_1\cdot\hat{s}) & \mathbb{T}_0(\vec{r}_1\cdot\hat{s}) \end{bmatrix} \begin{bmatrix} g_i[0] \\ \vdots \\ g_i[k] \\ \vdots \\ g_i[N-1] \end{bmatrix} \tag{4.14}$$

This means that a single index is used to access the matrix $\mathbb{T}$, $(j-k) \in \{-\frac{1}{2}(N_\tau - 1), \dots, 0, \dots, +\frac{1}{2}(N_\tau - 1)\}$, where N_τ is an odd integer and the periodicity of the input data has been used to wrap the rows of the matrix. Note that $\mathbb{T}_{j-k} = 0$ for $|j-k| > \frac{1}{2}(N_\tau - 1)$.

In the time-domain case, the rows of $\mathbb{T}$ were clipped to the N_τ elements around the diagonal to reduce the size of the computation. However, as the frequency domain data is assumed to be periodic, all of the elements of $\mathbb{T}$ can be computed, and the geometric delay can be performed exactly. However, in practice the data will not be periodic but will have been chopped into small segments prior to being Fourier transformed (see equation 3.3 in Chapter 3). The effect of this can be thought of as introducing noise at the boundaries of the input segments. Zeroing off-diagonal elements reduces this to just a few samples near the start and end of the output time series (see [66] for full details).

equation (4.14) can be rewritten as:

$$g_i[j;\hat{s}] = \sum_{k=0}^{N-1} \mathbb{T}_{j-k}(\vec{r}_i\cdot\hat{s}) g_i[k] \tag{4.15}$$

multiplying both sides by $\mathrm{e}^{-2\pi ijq/N}$, and summing over j, (4.15) becomes:

$$\sum_{j=0}^{N-1} g_i[j;\hat{s}]\mathrm{e}^{-2\pi ijq/N} = \sum_{j=0}^{N-1}\sum_{k=0}^{N-1} \mathbb{T}_{j-k}(\vec{r}_i\cdot\hat{s}) g_i[k]\mathrm{e}^{-2\pi ijq/N} \tag{4.16}$$

$$= \sum_{k=0}^{N-1} g_i[k] \sum_{j=0}^{N-1} \mathbb{T}_{j-k}(\vec{r}_i\cdot\hat{s})\mathrm{e}^{-2\pi ijq/N} \tag{4.17}$$

$$= \underbrace{\sum_{k=0}^{N-1} g_i[k]\mathrm{e}^{-2\pi ikq/N}}_{\tilde{g}_i[q]} \underbrace{\left(\sum_{j=0}^{N-1} \mathbb{T}_{j-k}(\vec{r}_i\cdot\hat{s})\mathrm{e}^{-2\pi i(j-k)q/N} \right)}_{\tilde{\mathbb{T}}_q(\vec{r}\cdot\hat{s})} \tag{4.18}$$

where the quantity in parentheses is the inverse Fast Fourier transform (see equation 3.4 in Chapter 3) of $\mathbb{T}_{j-k}(\vec{r}\cdot\hat{s})$, i.e:

$$\tilde{\mathbb{T}}_q(\vec{r}\cdot\hat{s}) = \sum_{j-k=0}^{N-1} \mathbb{T}_{j-k}(\vec{r}_i\cdot\hat{s})\mathrm{e}^{-2\pi i(j-k)q/N} \tag{4.19}$$

Substituting this into (4.16) and factoring out $\tilde{\mathbb{T}}_q(\vec{r}\cdot\hat{s})$ gives:

$$\sum_{j=0}^{N-1} g_i[j;\hat{s}]\mathrm{e}^{-2\pi ijq/N} = \tilde{\mathbb{T}}_q(\vec{r}\cdot\hat{s})\tilde{g}_i[q], \tag{4.20}$$

as expected from the convolution theorem. To make use of this equation for a pair of detectors, we first set i to 1 in (4.20), then $i = 2$. Taking the complex conjugate of the second equation, then the inner product of both sides, achieves:

$$\sum_{q=0}^{N-1}\left(\sum_{j_1=0}^{N-1} g_1[j_1;\hat{s}]\mathrm{e}^{-2\pi j_1 q/N}\sum_{j_2=0}^{N-1} g_2[j_2;\hat{s}]\mathrm{e}^{2\pi i j_2 q/N}\right) = \sum_{q=0}^{N-1}(\tilde{\mathbb{T}}_q(\vec{r}_1\cdot\hat{s})\tilde{g}_1[q])(\tilde{\mathbb{T}}_q(\vec{r}_2\cdot\hat{s})\tilde{g}_2[q])^* \tag{4.21}$$

Dividing by N and reordering gives:

$$\sum_{j_1=0}^{N-1}\sum_{j_2=0}^{N-1} g_1[j_1;\hat{s}]g_2[j_2;\hat{s}]\underbrace{\left(\frac{1}{N}\sum_{q=0}^{N-1}\mathrm{e}^{-2\pi(j_1-j_2)q/N}\right)}_{\delta_{j_1 j_2}} = \sum_{q=0}^{N-1}\tilde{g}_1[q]\left(\frac{1}{N}\tilde{\mathbb{T}}_q(\vec{r}_1\cdot\hat{s})\tilde{\mathbb{T}}_q^*(\vec{r}_2\cdot\hat{s})\right)\tilde{g}_2^*[q] \tag{4.22}$$

where the parenthesis on the left hand side contains the Kronecker δ, reducing the double sum over j_1 and j_2 to just a single sum over j:

$$\frac{1}{N}\sum_{j=0}^{N-1} g_1[j;\hat{s}]g_2[j;\hat{s}] = \frac{1}{N}\sum_{q=0}^{N-1}\tilde{g}_1[q]\left(\frac{1}{N}\tilde{\mathbb{T}}_q(\vec{r}_1\cdot\hat{s})\tilde{\mathbb{T}}_q^*(\vec{r}_2\cdot\hat{s})\right)\tilde{g}_2^*[q]. \tag{4.23}$$

An extra factor of N has been divided throughout, such that the left hand side is now the discrete form of the cross-correlator integral, (4.1), and the right side is the frequency-domain correlator.

Writing the left-hand side of (4.20) as $\tilde{g}[q;\hat{s}]$, the Fourier transform of the direction-dependent geometrically delayed time series, equation (4.23) becomes

$$\frac{1}{N}\sum_{q=0}^{N-1}\tilde{g}_1[q;\hat{s}]\tilde{g}_2^*[q;\hat{s}] = \frac{1}{N}\sum_{q=0}^{N-1}\tilde{g}_1[q]\left(\frac{1}{N}\tilde{\mathbb{T}}_q(\vec{r}_1\cdot\hat{s})\tilde{\mathbb{T}}_q^*(\vec{r}_2\cdot\hat{s})\right)\tilde{g}_2^*[q] \tag{4.24}$$

The left-hand sides of (4.23) and (4.24) are related by Parseval's theorem - that the square integral can be computed in the time, (4.23), or the frequency domain, (4.24). In this case, the quantity $\tilde{g}_1[q;\hat{s}]\tilde{g}_2^*[q;\hat{s}]$ is the cross-power in the baseline in the frequency bin q, for the time interval spanned by the time series that was Fourier transformed. As a vector of harmonic coefficients, it can be used to do a direction resolved time-frequency decomposition of the sky.

From equation 4.12 we see that the elements of the geometric delay matrix are oscillatory functions on the sphere will some finite frequency. These can be projected onto the spherical

harmonics basis with the expectation that we will be able to achieve a good approximation to the elements by retaining some finite number of coefficients.

Denoting the $(lm)$the coefficient of the expansion into spherical harmonics of the dependence of $\mathbb{T}_{jk}(\vec{r}_i \cdot \hat{s})$ on $\hat{s}$ as $\mathbb{T}_{jk}^{(lm)}(\vec{r}_i)$, we can write

$$\mathbb{T}_{jk}(\vec{r}_i \cdot \hat{s}) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \mathbb{T}_{jk}^{(lm)}(\vec{r}_i) Y_{lm}(\hat{s}), \tag{4.25}$$

where

$$\mathbb{T}_{jk}^{(lm)}(\vec{r}_i) = \int \mathbb{T}_{jk}(\vec{r}_i \cdot \hat{s}) Y_{lm}^*(\hat{s}) d\Omega. \tag{4.26}$$

As we are using an Earth-fixed coordinate system the instrument location, $\vec{r}_i$, is independent of time, so $\mathbb{T}_{jk}^{(lm)}(\vec{r}_i)$ is a constant, and,

$$\mathbb{T}_{jk}^{(l,-m)}(\vec{r}_i) = (-1)^m \left[\mathbb{T}_{jk}^{(lm)}(\vec{r}_i)\right]^* \tag{4.27}$$

since $\mathbb{T}_{jk}(\vec{r}_i \cdot \hat{s})$ is real valued.

This means we can write Equation(4.24) as

$$(\tilde{g}_1[q;\hat{s}]\tilde{g}_2^*[q;\hat{s}])^{(lm)} = \tilde{g}_1[q]\tilde{g}_2^*[q] \left(\frac{1}{N}\tilde{\mathbb{T}}_q(\vec{r}_1 \cdot \hat{s})\tilde{\mathbb{T}}_q^*(\vec{r}_2 \cdot \hat{s})\right)^{(lm)}, \tag{4.28}$$

where $()^{(lm)}$ refers to the $(lm)^{th}$ coefficient of the spherical harmonic expansion of the function on the sphere in parentheses. Which we can use to reframe the correlation equation, (4.1), in terms of the spherical harmonic expansion

$$\xi_{1,2}(\hat{s}) = \frac{1}{N} \sum_{l=0}^{l_{\mathrm{MAX}}} \sum_{m=-l}^{l} \sum_{q=0}^{N-1} \left[\tilde{g}_1[q] \left(\frac{1}{N}\tilde{\mathbb{T}}_q(\vec{r}_1 \cdot \hat{s})\tilde{\mathbb{T}}_q^*(\vec{r}_2 \cdot \hat{s})\right)^{(lm)} \tilde{g}_2^*[q]\right] \tag{4.29}$$

4.2.2 The algorithm

To evaluate (4.24), begin by computing the top row of $\mathbb{T}_q(\vec{r}_i \cdot \hat{s})$ as defined in (4.16), which has the element with the index $j - k = 0$ in the first position. Note that each element of this row is a function of the sky direction, $\hat{s}$. Perform a spherical harmonic expansion of each element of the row, creating a new matrix containing a vector of the coefficients of this expansion. The row is then DFT'ed to the frequency domain, where the DFT is carried out on the related coefficients in each vector -all of the monopole ($l = 0$, $m = 0$) coefficients for the row are taken together as a time series and DFT'ed, then the next ($l = 1$, $m = 0$) coefficients, and so on.

Recalling the equivalence shown in (4.19), we note that the DFT is frequency inverted to give the components of $\mathbb{T}_q(\vec{r}_i \cdot \hat{s})$, where $q \in \{0, \ldots, N-1\}$. This is carried out for both

instruments, then the product of the two arrays are computed frequency bin by frequency bin, giving $(\tilde{\mathbb{T}}_q(\vec{r}_1 \cdot \hat{s})\tilde{\mathbb{T}}^*_q(\vec{r}_2 \cdot \hat{s})^{(lm)}$. Each frequency bin, q, has been expanded as a vector of harmonic coefficients which can be stored for repeated use during correlation.

From this point, the correlation algorithm is trivial - for each segment of data having N samples, the two input time series (corresponding to one pair of detectors in the network of choice) are transformed to the frequency domain to give $\tilde{g}_1[q]$ and $\tilde{g}^*_2[q]$, then multiplied together, bin by bin. Each frequency bin of this vector multiples the harmonic coefficients in the corresponding frequency bin of $(\tilde{\mathbb{T}}_q(\vec{r}_1 \cdot \hat{s})\tilde{\mathbb{T}}^*_q(\vec{r}_2 \cdot \hat{s})^{(lm)}$ to yield the harmonic coefficients for the angular distribution of cross power for that frequency bin. Summing across bins generates the baseline's direction-dependent integrated cross-power in the form of a single vector of coefficients of its expansion in spherical harmonics (for more details see [66]).

One side effect of using the frequency domain correlator is that the length of the data to be processed, N, must be chosen before the analysis is carried out (we use 1 second FFTs in our search). This is so we can form the delay matrix, $\mathbb{T}_q$, at the frequency resolution that is needed (later in this chapter we will explore if it possible to circumvent this issue, see Section 4.4.3).

The raw unsummed output for a segment of data (one time bin of N samples) from the correlator is a 2-dimensional grid (see Figure 4.2). This is the spherical harmonic expansion for each frequency of interest. Along one axis we have the frequency span of our data, along the other we have the spherical harmonic coefficients, c_{lm}, of that expansion. The coefficients are ordered according to SpharmKit [68]; see Appendix B). Note that Figure 4.2 shows the

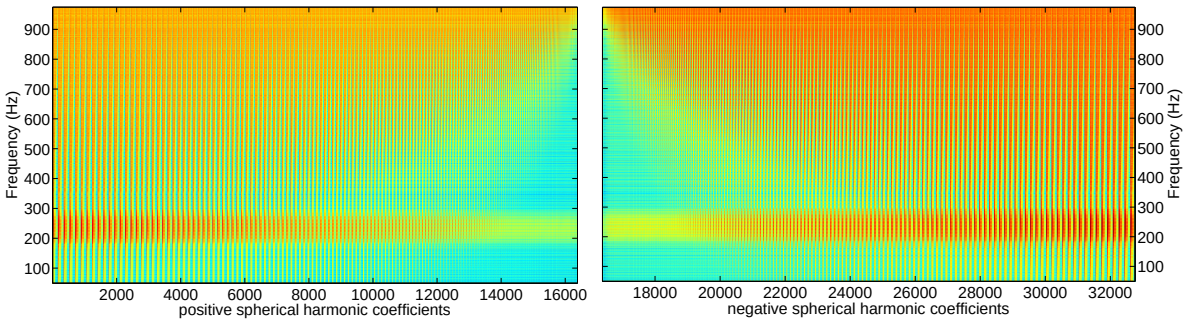

FIGURE 4.2: Example frequency-spherical harmonic map in log-scale colouring for a pair of detectors. The time bin analysed contained an injected sine-Gaussian signal. The injection is clearly visible as the horizontal line at $f = 235$ Hz, the central frequency of the injection. The background features are due to the plot ordering of the spherical harmonics and the log colouring.

cross-correlation between a single *pair* of detectors, and this would actually be generated for each detector pair in the network, and potentially auto-correlations, depending on the statistic we are trying to form.

The input data is processed in arbitrarily length blocks (generally 128-256 seconds long). These blocks are further subdivided into sections (which are processed separately by the correlator) corresponding to the FFT length,N, using a 50% overlap. Thus, the output from processing a single block (i.e. 128 seconds) of data for a single pair of detector is a time series of the matrix shown in Figure 4.2, as illustrated in Figure 4.3.

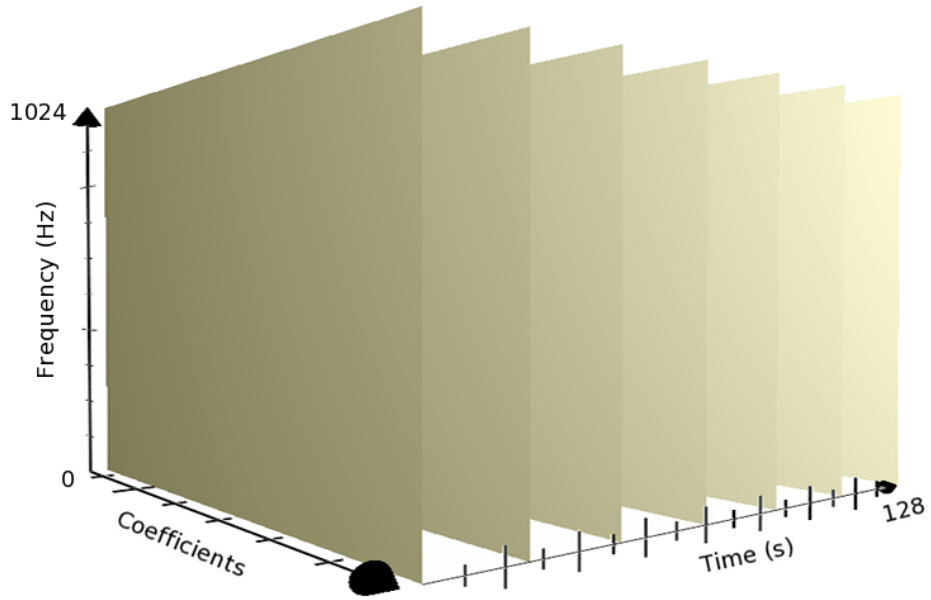

FIGURE 4.3: Illustration of a Time-Frequency-Spherical harmonic coefficients cube. Each plane represents the output of the correlator for a segment of data as in Figure 4.2

It is important to note that these demonstrate why the representation of the problem in spherical harmonics is so useful. For each time bin we have a 2-dimensional matrix (of frequency and spherical-harmonic coefficients) that can processed in a different ways to highlight various aspects of the correlation.

For instance, to generate the staple of the burst analysis, the time-frequency map, merely requires a summation along the spherical harmonic coefficients. For *each* frequency bin in a time slice, we calculate

$$P_{t,f} = \sum_{l=0}^{l_{\mathrm{MAX}}} \sum_{m=-l}^{l} |c_{lm}|^2, \tag{4.30}$$

where l_{MAX} is the spherical harmonic expansion order. One thing to remember here is that the value calculated from the spherical harmonic coefficients has been averaged over the whole sky (since we dealing with spherical harmonics), but used as a tool to find potential signals this should not cause issues. Generating $P_{t,f}$ for each of the time slices in our block, we get a coherent time-frequency map (see Figure 4.4):

Similarly, if we collapse in the frequency direction, i.e. sum over all frequencies for each spherical harmonic coefficient, we get the power in the sky for all frequencies. Generating this

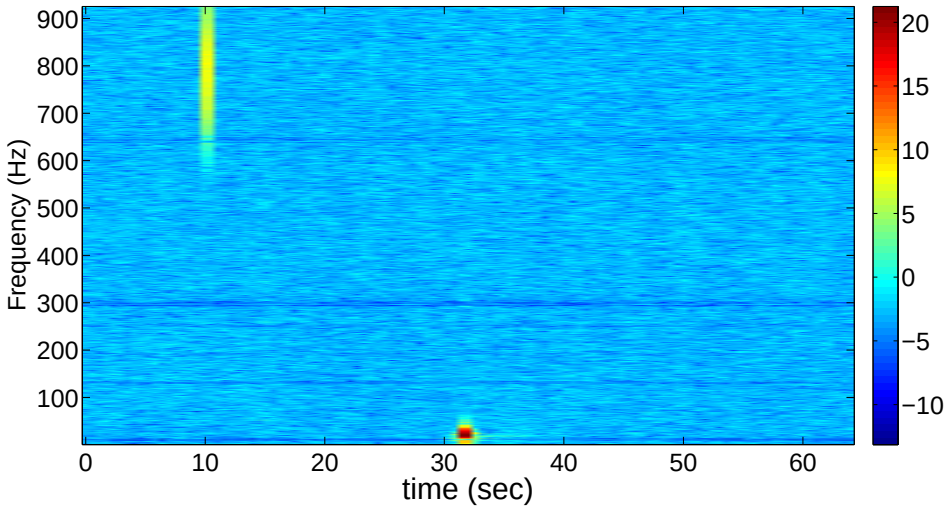

FIGURE 4.4: A time-frequency map of whitened data; to the upper left (10s, 800Hz) we see a short-duration white noise burst (WNB) signal, to the lower middle a glitch (32s, 25 Hz). The color scale is the logarithm of the coherent network power, as defined in equation (4.30)

.

for each time bin in our block we get a time-spherical harmonic coefficient map (see Figure 4.5). Each set of coefficients tell us something about the distribution of power over the sky for that time bin. Referring to Figure 4.5, it is pertinent to note that the power associated

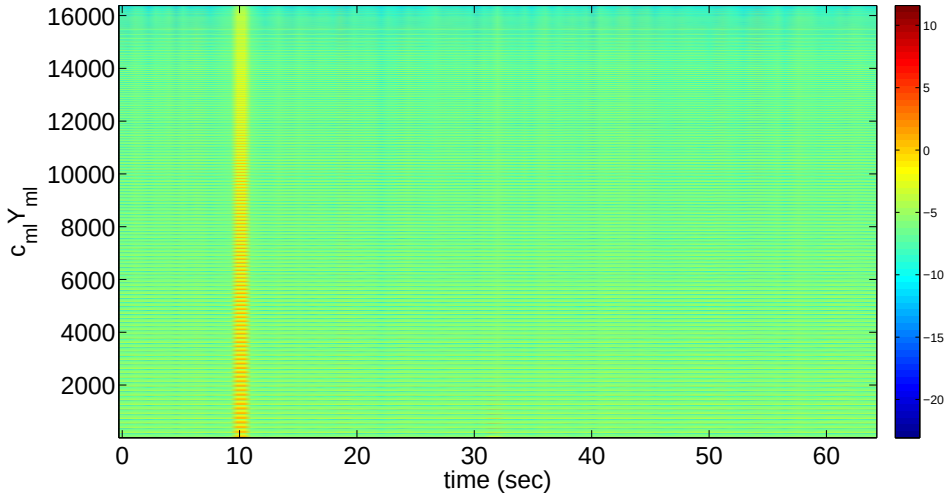

FIGURE 4.5: A time-harmonic map for the same segment shown in Figure 4.4. The injection is clearly visible as power is spread throughout the coefficents (the vertical line at 10s), while the glitch is just visible with energy in the $l = 0$ coefficient (at 32s) - compare to Figure 4.4.

with the glitch is concentrated in the $l = 0$ mode, just visible on this image. The way in which glitches behave in this representation will form the basis for a powerful glitch rejection statistic, as we shall see in Section 4.6.

Note that this 'dual view' of the data is very important, and is one of the major advantages of using spherical harmonics, the other is that we get all-sky for free.

4.3 Speed

One of the main motivations behind the development of yet another pipeline was the lack of an existing fast all-sky coherent analysis. If the speed of the methods outlined above did not translate into real world performance, then, while we would be still be left with a novel all-sky pipeline, there would be little else to differentiate it from the others. One thing that we have not really touched upon in this discussion of the pipeline is the startup lag from having to generate $\mathbb{T}_q(\vec{r}_i \cdot \hat{s})$. Separating the dependence on $\hat{s}$ (essentially the time delay between detectors) out of (4.1) by way of (4.14) allows us to pre-generate the delay matrix, $\mathbb{T}_q(\vec{r}_i \cdot \hat{s})$ for later use. However, it is a large, computationally expensive operation. To determine

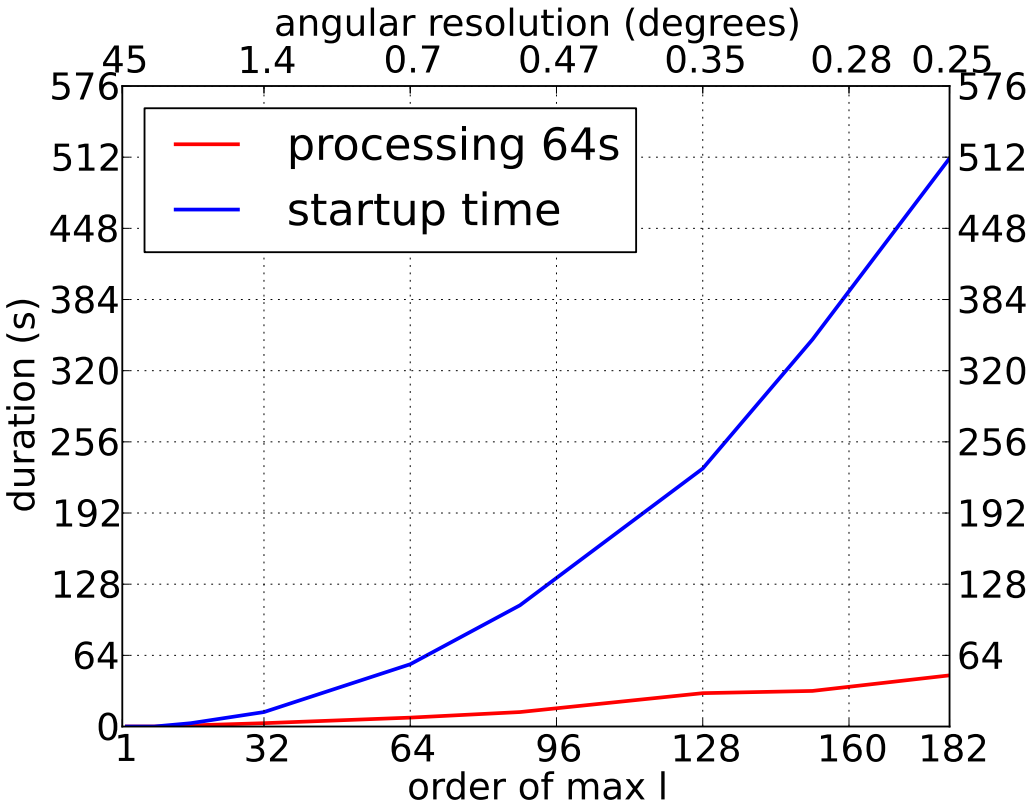

FIGURE 4.6: Processing time for startup and segment analysis against l order

the performance of the pipeline, it was run multiple times using different spherical harmonic expansion orders on the same 64 seconds of data from Science Run 4. This was on an iMac powered by an Intel 2.4 GHz Core 2 Duo. The results from this test are displayed in Figure 4.6. Note that the order of the spherical harmonic expansion needed to fully represent $\mathbb{T}_q(\vec{r}_i \cdot \hat{s})$ is given by, (see [66]),

$$l_\tau \sim \left[\frac{\pi |\vec{r}_i|}{\Delta t}\right]. \tag{4.31}$$

where $\vec{r}_i$ is the baseline distance between detectors, and Δt is the sample period.

Figure 4.6 shows the startup time (generation of $\mathbb{T}_q(\vec{r}_i \cdot \hat{s})$), and the subsequent segment processing time for a range of spherical harmonic expansion orders. As can be seen, the startup time can be substantial, approximately 512 seconds for the necessary order for sample frequency of 2048Hz ($l = 182$), note however, that this only has to be carried out at the beginning of an analysis run. Once created, it is cached and used for an entire segment (many thousands of seconds) of data, and is thus a small percentage of the entire run time.

The processing of a segment of data (64 seconds in this case) takes considerably less time, always less than real-time, and in many cases substantially faster. For example, for $l = 182$, the segment processing time is around 40 seconds (against 512 seconds for startup).

One important thing to note here is that this segment processing time includes correlating, clustering and event identification - the whole all-sky coherent analysis is carried out for multiple likelihoods. For comparison, Figure 4.7 shows the same graph

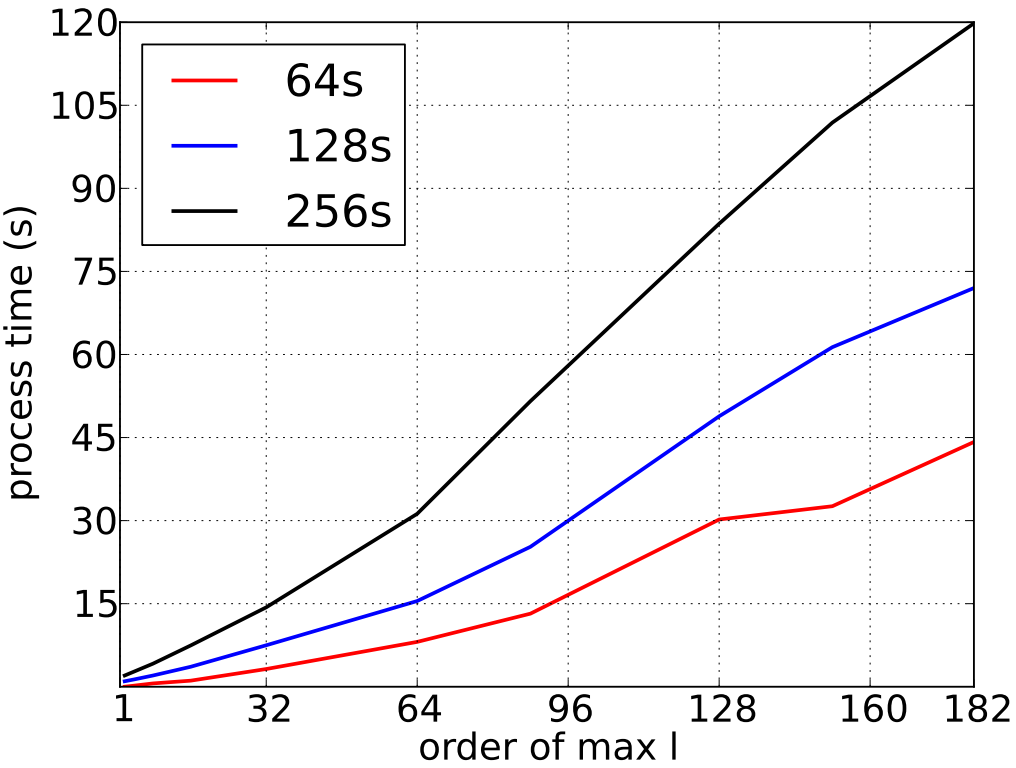

FIGURE 4.7: Processing time for segment analysis against l order for a 64, 128 and 256 second data segments

(only for segment processing) for 64, 128 and 256 seconds of S4 data. The processing time of a longer data segment is proportionally faster (per second) than the smaller data segments. This can be explained easily when you consider how the pipeline processes the data (see later), and is related to the amount of event followups carried out by the code. In each case here, 6 events are analysed in-depth, so the increased processing time can be directly attributed to the generation of a time-frequency and time-coefficient map at twice the time span.

In terms of memory usage, Figure 4.8 shows how it increases with the expansion order. This

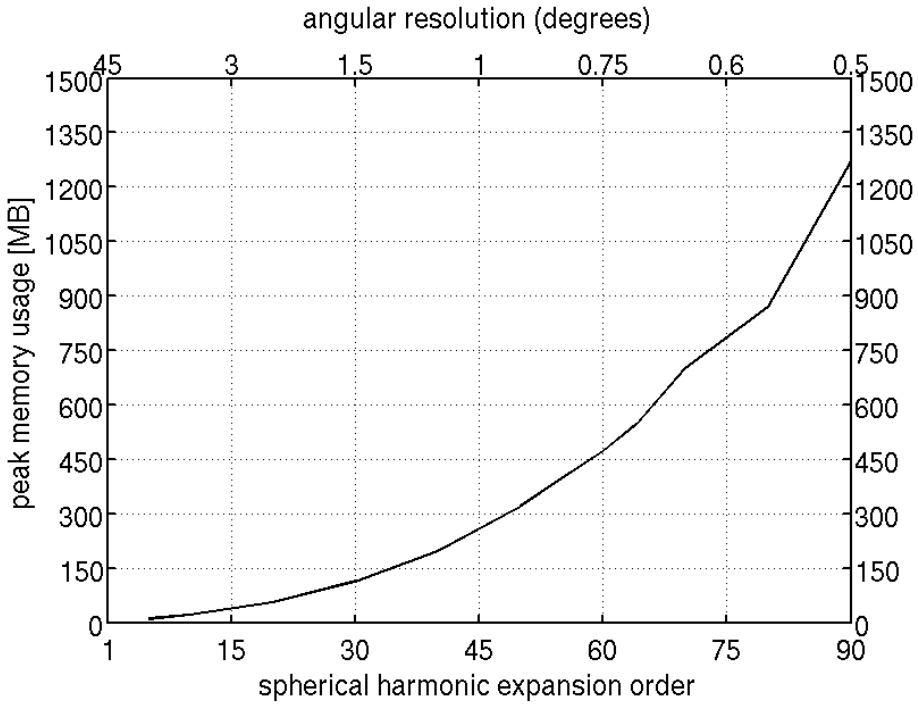

FIGURE 4.8: Maximum amount of memory used by the pipeline for each l order.

becomes a limiting factor when running on a cluster of machines - to run at the highest resolution requires nodes with a large amount of memory.

4.3.1 Background estimation

One of the last attributes of the pipeline to be explored was background generation. This is generally carried out by time-shifting detector data relative to each other (i.e. first stream no shift, second stream shifted by 2 seconds, third stream shifted by 4 seconds) then analysing these streams as normal. This shifting ensures that the data being processed is independent (as well as it can be) such that any features that do crop up cannot be a signal, and therefore gives a good indicator of the distribution of background noise.

In the Spherical Radiometer pipeline (referred to as SPHRAD from now on) there is the option to cary out these time-shifts internally - i.e. once the data segments have been read in, each segment is shifted by a certain amount of time and analysed as normal. The assumed benefit of carrying out the shifts internally is to bypass any initial setup time (populating Fourier transforms, loading the data etc). Figure 4.9 shows how the processing time scales for different orders of the spherical harmonic expansion and number of time lags carried out. A quick bit of division shows that, although quite variable, the time taken to process a single

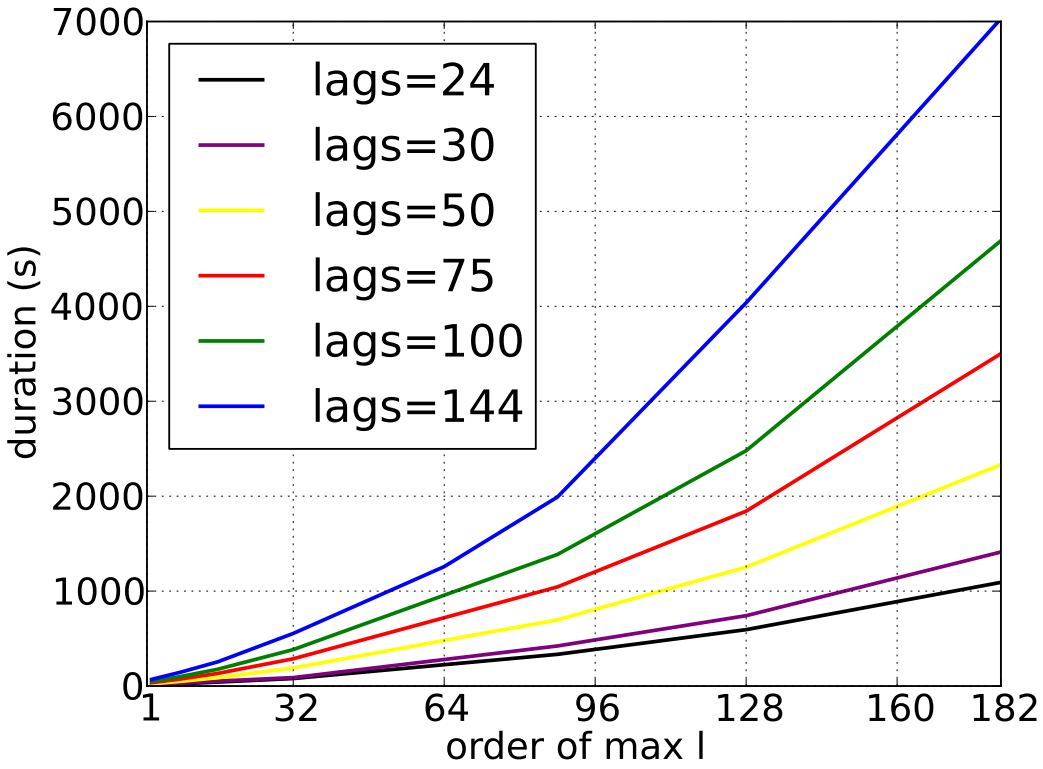

FIGURE 4.9: Processing time for a range of internal time-shifts, for a 64 second data segment, over a range of l orders.

time lag is a little smaller than processing the zero lag. The raw values can be seen in table 4.1.

l order	Number of time lags						
	0	24	30	50	75	100	144
2	0.9	11.7	14.7	24.5	36.7	48.9	70.4
4	1.3	15.3	19.2	32.3	48.4	64.5	94.9
8	2.0	23.1	29.6	49.4	73.7	98.6	143.5
16	3.6	42.3	52.7	88.2	133.8	177.4	255.2
32	7.5	89.2	89.9	190.2	287.3	384.7	554.4
88	25.3	334.4	422.5	697.4	1044.2	1386.5	1991.3
128	48.8	595.4	745.4	1252.4	1842.0	2479.2	3622.9
182	72.0	1093 .1	1403.4	2332.5	3501.3	4691.8	6823.5

TABLE 4.1: Table showing the time (in seconds) required to process internal time lags for a range of spherical harmonics orders. The numbers are for a 128 second block of data.

Decomposing the correlation integral into the spherical harmonic domain allows the dependence on sky position to be separated out from the processing of the data. This is the key feature of this pipeline - the dependency on sky position can be pre-computed (which takes a large startup time) and then the processing of the data segments can take place as fast as it is possible to send data to it.

4.4 Implementation

The Spherical Radiometer pipeline consists of two main parts: a Matlab framework based on the Omega pipeline and the spherical radiometer engine written in C, see Figure 4.10.

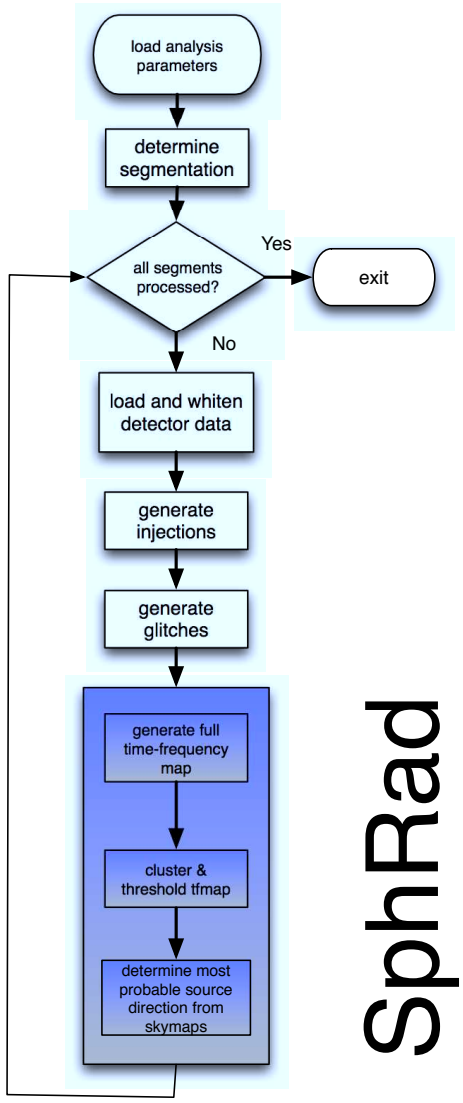

FIGURE 4.10: A flowchart showing where the Spherical Radiometer part of the pipeline fits in Omega

Omega is a burst pipeline that uses the principle of excess power to determine the existence of a gravitational wave in the detector data, see [65]. Unlike matched filtering (where the signal is known beforehand and the search can therefore be tweaked to be optimal, or templates can be used), burst searches attempt to find unmodelled gravitational-waves. It was an *incoherent* pipeline - so candidate events are constructed for each detector in isolation, then combined to find those with similar parameters in more than one detector - but it was modified to use the coherent statistics discussed in the previous chapter.

The spherical radiometry engine has its foundations in the radiometry library that Kipp Cannon wrote to demonstrate his ideas (see [66]). It replaces all of the functionality in Omega relating to the detection of gravitational wave signals.

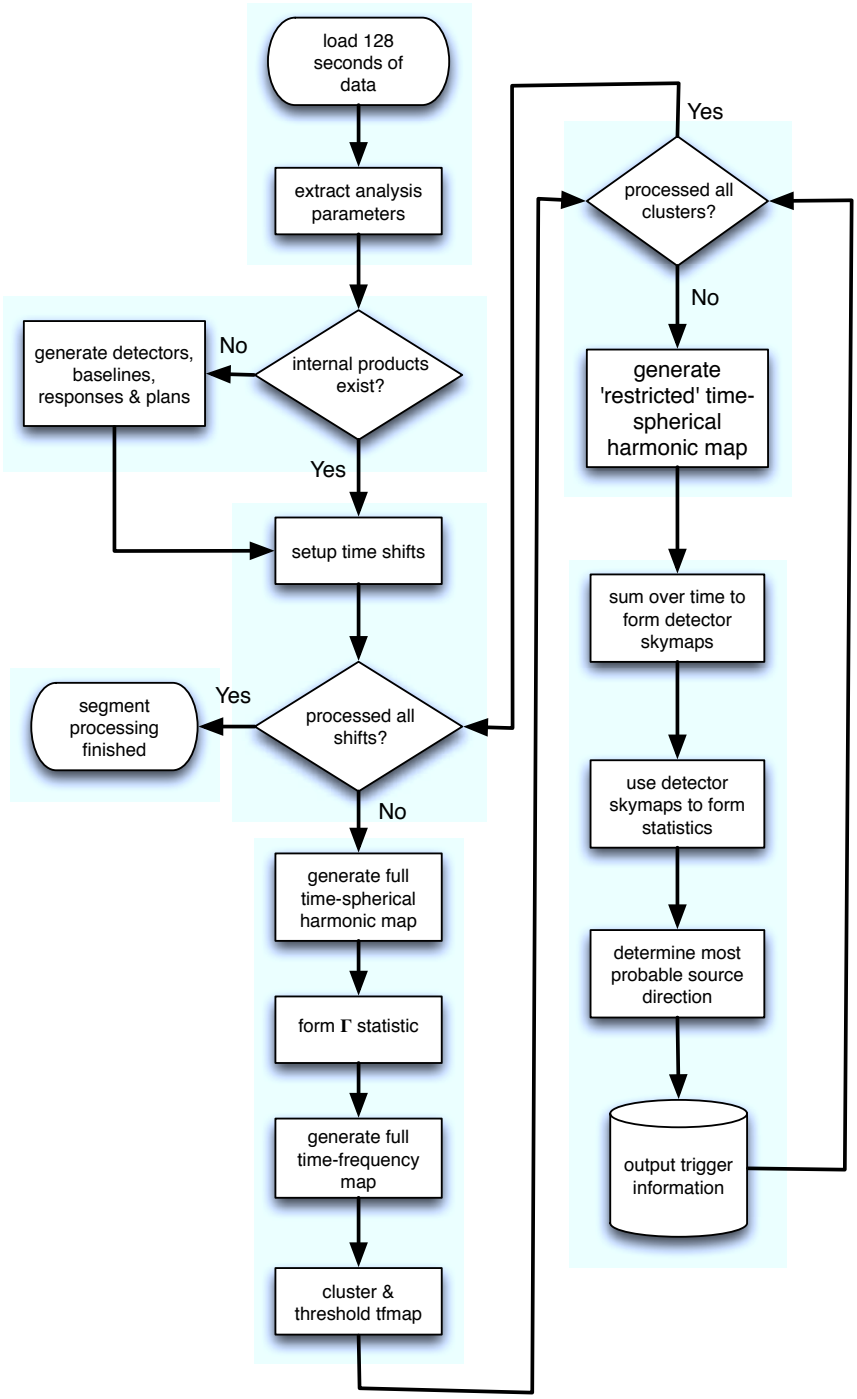

FIGURE 4.11: The flow chart for the Spherical Radiometer part of the pipeline

While the library provides the basic functions to calculate the correlation between two data streams, including handling the spherical harmonic coefficients, a large amount of code was written by the author. The majority of this extra code was to provide the necessary features

required to process the output from the correlator (handing multi-detector networks, formation of time-frequency and time-spherical harmonic coefficient maps, threshold and clustering, generating skymaps etc).

The flow of data through the SphRad pipeline is outlined in the flowchart shown in Figure 4.11. At its core is the frequency domain correlator, which takes two segments of whitened, Fourier transformed data of length N and produces the correlated output. This is processed and manipulated in different ways to produce the different views of the data (e.g. for each time bin, we can sum over a frequency range to produce a skymap (as spherical harmonic coefficients) or sum over the coefficients to produce a time-frequency map.

Specifically, for each time bin we compute the zonal (i.e. $m = 0$ for all l) spherical harmonic-frequency map for each pair of detectors (i.e. the power in the baseline joining the detectors). From this, we generate the Γ statistic (by summing over detectors, frequency and then using Equation (4.42)), and the time-frequency map (formed by summing over the coefficients and detectors). The time-frequency map is clustered to determine pixels of interest, with only the loudest $N_{\rm c}$ clusters continuing onto the next stage. As we are only interested in the amount of power present in each pixel on the time-frequency map, this approximation is both quick and accurate.

Once we have a cluster of pixels, its time-frequency span is used to restrict the full time-frequency-spherical harmonic map such that we reduce the amount of noise present in our selected voxels (3-d pixels) by disregarding those outside of the cluster. We then sum over the frequencies for each detector pair to generate the "restricted" spherical harmonic coefficients (so that we have a set of coefficients for each detector pair, representative of a sky map), generate the antenna responses for our frequency band of interest and form our sky localisation statistics. This multi stage process enables us to reduce the computational cost of the analysis (the first stage is very quick, the second stage slower, which is why we only fully process loud events), while attempting to maximise our efficiency (the first stage is an approximation, the second uses the full decomposition).

Note that the correlator's output is based on the data from two detectors, so even when we are combining the "raw" or zonal spherical harmonics coefficients there is still an element of a coherent analysis. Generally our analysis is concerned with 3 detector networks. This means that the steps above must be run on each pair of detectors (including single detectors with themselves). In this case, the total raw coherent network power would be given by the sum over all frequencies for all detectors in the network (decomposed into spherical harmonics)

$$N_{total} = \underbrace{\xi^{(lm)}_{1,2} + \xi^{(lm)}_{1,3} + \xi^{(lm)}_{2,3}}_{\text{cross-correlation}} + \underbrace{\xi^{(lm)}_{1,1} + \xi^{(lm)}_{2,2} + \xi^{(lm)}_{3,3}}_{\text{auto-correlation}} \tag{4.32}$$

In fact, most coherent statistics (see for example the standard likelihood [52]), make use of both of the cross- and auto-correlation terms for a network. For example, there are 6 terms that can be calculated for the HLV network, HH, LL, VV, HL, HV and LV. This means that any coherent statistic that makes use of any of these pairings can be formed.

For instance, to form the coherent energy in the plus polarisation, the following, equation 4.33, needs to calculated.

$$E_+ = \sum_{i=1}^{D}\sum_{j=1}^{D} |e_{ij}^+ \cdot d_{ij}|^2 \tag{4.33}$$

$$\begin{aligned} =& e_{11}^+ \cdot \tilde{d}_{11} \cdot e_{11}^+ \cdot \tilde{d}_{11}^* + e_{11}^+ \cdot \tilde{d}_{11} \cdot e_{12}^+ \cdot \tilde{d}_{12}^* + e_{11}^+ \cdot \tilde{d}_{11} \cdot e_{13}^+ \cdot \tilde{d}_{13}^* + \\ & e_{21}^+ \cdot \tilde{d}_{21} \cdot e_{21}^+ \cdot \tilde{d}_{21}^* + e_{21}^+ \cdot \tilde{d}_{21} \cdot e_{22}^+ \cdot \tilde{d}_{22}^* + e_{21}^+ \cdot \tilde{d}_{21} \cdot e_{23}^+ \cdot \tilde{d}_{23}^* + \\ & e_{31}^+ \cdot \tilde{d}_{31} \cdot e_{31}^+ \cdot \tilde{d}_{31}^* + e_{31}^+ \cdot \tilde{d}_{31} \cdot e_{32}^+ \cdot \tilde{d}_{32}^* + e_{31}^+ \cdot \tilde{d}_{31} \cdot e_{33}^+ \cdot \tilde{d}_{33}^* \end{aligned} \tag{4.34}$$

Here, the $\tilde{d}_{ij} \cdot \tilde{d}_{ij}^*$'s are the output from the correlator for the data from detectors i and j, while the e_{ij}'s are the orthonormal antenna factors for the same pair of detectors. The nature of equation 4.33 (a summation), means that each term can be calculated and added together to form the statistic.

4.4.1 Time-frequency maps & Event Identification

Having computed the correlations between the data from a network of detectors, the next task is to identify potential gravitational wave events in the output. One of the standard procedures is to form a time-frequency map and attempt to identify features that may be indicative of a signal (see [50]).

To generate a time-frequency map the spherical harmonic coefficients for each time and frequency bin must be summed. We can use Parseval's theorem for the spherical harmonic expansion,

$$\int_{-\infty}^{\infty} |x(t)|^2 dt = \int_{-\infty}^{\infty} |X(f)|^2 df, \tag{4.35}$$

which states that the total energy in the time-domain for the waveform $x(t)$ for all times is equal to the total energy of its Fourier Transformed waveform for all frequencies. We actual use the discrete version,

$$\sum_{n=0}^{N-1} |x[n]|^2 = \sum_{k=0}^{N-1} |X[k]|^2, \tag{4.36}$$

as we are dealing with finite time series.

Thus, we do the following:

$$\text{pixelPower} = \sum_{l=1}^{l_{\text{MAX}}} \sum_{m=-l}^{l} \left[\xi_{1,2})^{lm} \right]^2, \tag{4.37}$$

summing the coefficients to get a measure of the power in that time-frequency pixel.

An important thing to note here is that when we are generating the time-frequency map by summing over the spherical harmonics, we are doing so using a set of azimuthally symmetric coefficients (i.e. $m = 0$ for all l). The computational cost of rotating the spherical harmonic coefficients of the expansion for each frequency and time bin for each term of equation 4.32 is prohibitively high. The fast time-frequency map can be generated in less than a second, whereas the full time-frequency map takes approximately the same length of time as the amount of data we are processing. Figure 4.12 shows time-frequency maps for the same block

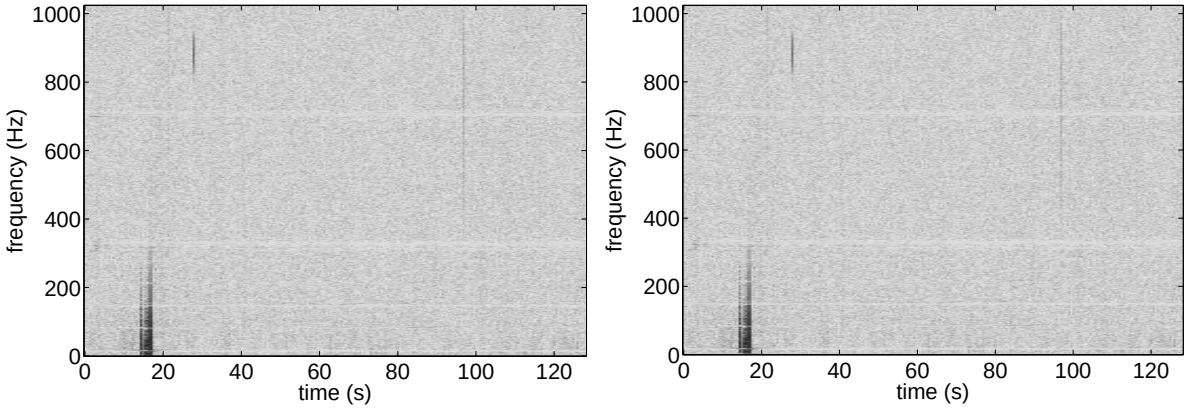

FIGURE 4.12: Plot showing two time-frequency maps generated in different ways: Left) Fast time-frequency map, Right) Full time-frequency map

of data. As can be seen, there is little difference between the two representations - all of the observable features are present in both.

The value in the pixel has been averaged over the whole sky, as the original spherical harmonic coefficients dealt with the cross-power over the sphere, and this summing discards the location information. However, as the main aim of forming the time-frequency map is to identify times and frequencies of interest, this is not an issue.

4.4.2 Clustering algorithm

While the time-frequency map is a good visual indicator as to how the power is distributed over time and frequency [for more information see [50]], an automated procedure is needed to identify features (essentially excess power, see [51]) in the map that could be due to a

gravitational wave. There are many different method of achieving this, but they tend to fall into two distinct categories (a good overview of the art can be found in [69]):

- **Shape searching -** If something is known about the signal *a priori*, such as its bandwidth, duration or time-frequency behaviour (for example, an inspiral or chirplet), then it is possible to construct a 'box' that would contain the pixels of the signal. By 'sliding' this shape over the time-frequency map and summing the pixels that it encapsulates, it is possible to determine which cluster of pixels give the greatest summed value. This method is useful *if* something is known about the signal, but even if the information is vague, many such shapes could be slid over the time-frequency map and compared.
- **Pixel clustering -** If nothing is known about the signal under investigation, or it can take many different disparate forms, then a method is required that heuristically generates clusters based on the pixel values. This is usually carried out in two distinct parts [see [69] for more information]: *i)* the thresholding, and *ii)* the clustering.

  Thresholding is the act of culling pixels by comparing their values to some defined threshold. This removes noisy pixels from the map and allows clustering to take place on a cleaner background.

  Clustering is the use of some characterising algorithm to group the pixels. This is usually a combination of location and pixel energy.

  The above processes produce a set of pixel clusters which match the search criteria best.

The clustering method implemented in SphRad is based on the pixel clustering method, but combines the thresholding and clustering steps together, along with a few other tweaks.

4.4.2.1 Hybrid clustering

Figure 4.13 is an illustrative time-frequency map with a handful of pixels:

- the blue pixels can be thought of as low-valued, but still above the background,
- the red pixels have very large values, while
- the green pixels are moderately valued.

Traditionally (see [49] and [65]), in the thresholding step, a lower bound value is generated based on all of the pixels in the time-frequency map. If, as we have in Figure 4.13, there is a patch of very loud pixels (such as a large-amplitude background feature), this can bias the threshold value such that the blue pixels would get completely culled, as well as some of the lower valued green pixels.

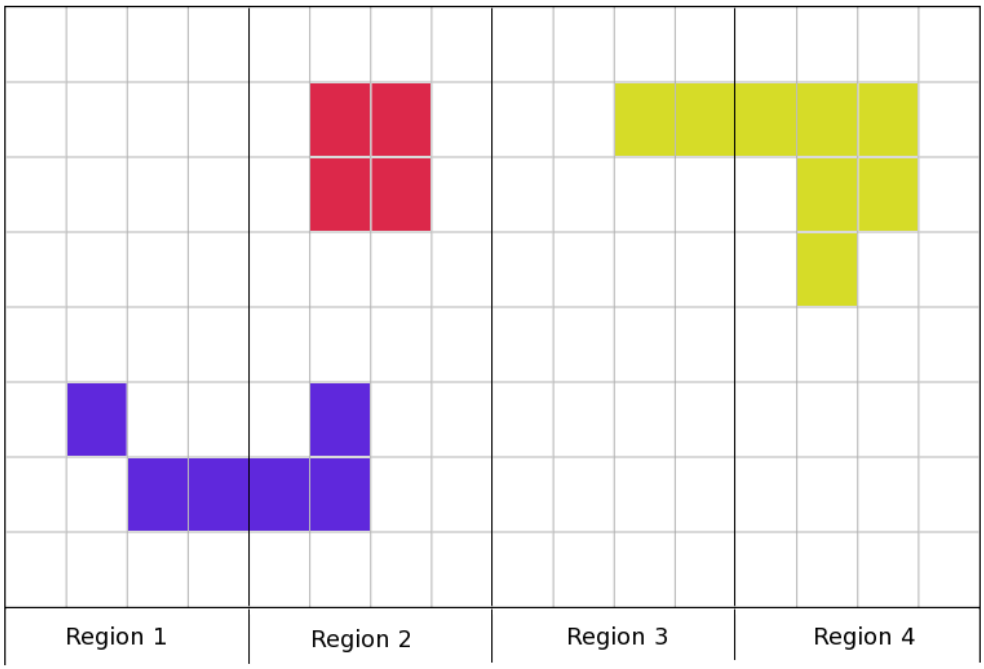

FIGURE 4.13: An example time-frequency map with a plurality of pixels - low-valued blue pixels straddling Regions 1-2, high-valued red pixels in Region 2 and moderately-valued green pixels in Regions 3-4.

The most obvious way of overcoming this issue is to subdivide the time-frequency map into smaller regions (as can be seen in Figure 4.13) and set individual thresholds for each region. A loud pixel cluster would then only effect its own region. However, the blue pixels in Region 2 would still be culled.

A more sophisticated approach would be to combine the threshold and clustering steps such that if a cluster starts in Region 1, it tests new pixels for cluster membership based on the threshold from Region 1 regardless of the location of the pixel under investigation. Thus, the three blue pixels in Region 2 would be assessed using the threshold from Region 1.

Finally, there may be pixels that have a value lower than the threshold but still belong to a cluster (for example, pixels in the decaying tail of a transient). Normally, these pixels would be culled, but to catch these, a secondary threshold is used (so called halo clustering, see [63]). Once the determination has been made that a pixel is the start of a cluster, the secondary threshold is used to check all adjoining pixels for cluster membership.

The clustering technique used in the SPHRAD pipeline contains all of the above features. Once a time-frequency map has been generated, it is then processed by the clustering algorithm as follows (see Figure 4.14):

1. Divide the time-frequency map up into R equal sized regions (can be in frequency, in time or both),

2. Calculate a first 'core' threshold, T_1, and a second 'halo' threshold,T_2, for each region, where the second threshold is less than or equal to the first,
3. Select the first pixel in the time-frequency map,
4. Compare the value of this pixel to the core threshold value,T_1, for that region,
5. If lower than T_1, move onto the next pixel and carry out step 4 again,
6. *START NEW CLUSTER* - Once a pixel value is greater than, or equal to, T_1, mark this pixel as the start of a cluster,
7. Compare the value of the neighbouring pixel against the halo threshold value, T_2 for the region in which the pixel initiated the cluster,
   (a) if higher, add pixel to cluster (and unfinished pixel stack), and check this pixel's neighbours against secondary threshold (i.e. recursively call step 7),
   (b) if lower, do not add into cluster, and check unfinished pixel stack:
      i. *END CLUSTER* - If empty, record cluster details, then move onto the next unanalysed pixel in the time-frequency map (back to step 4),
      ii. If pixels are present on the unfinished pixel stack, pop the top one (i.e. move back a pixel), and check the next unanalysed surrounding pixel (step 7 again).

The threshold values are calculated in a very simple, if laborious, manor - the pixels are sorted in order and then the value returned is that which is greater than a certain percent of the values (i.e. the median would be returned for 50%). When using the clustering algorithm set out above, the 1^{st} threshold is generally higher than it would be if we were using a single threshold clustering algorithm. This is because the combination of multiple regions and the 'mopping' up 2^{nd} threshold can otherwise select a large region of pixels. However, if set too overzealously, features may be missed. From testing, a value of 99% is used for the 1^{st} threshold and 96% for the 2^{nd} threshold appears to give the best balance between the two extremes.

This clustering scheme emerged from the requirements enforced by the analysis. When analysing data in blocks greater than a minute, it is common for at least one loud background event (e.g. a glitch of some kind) to be present in the analysed block of data. Also, because our time-frequency map is essentially presenting excess power, glitches do not get down weighted. Thus, a large glitch in the same block as a real signal would swamp the threshold calculation and the signal would be missed.

An easy to implement solution to this problem is to use region-based thresholding - if a glitch is present in one region, the thresholds may be very high, but as long as the signal originates in a different region (and even this can be fixed by tiling the time-frequency map in frequency as well as time) it will still be identified. Of course, this introduced its own problems - if a cluster

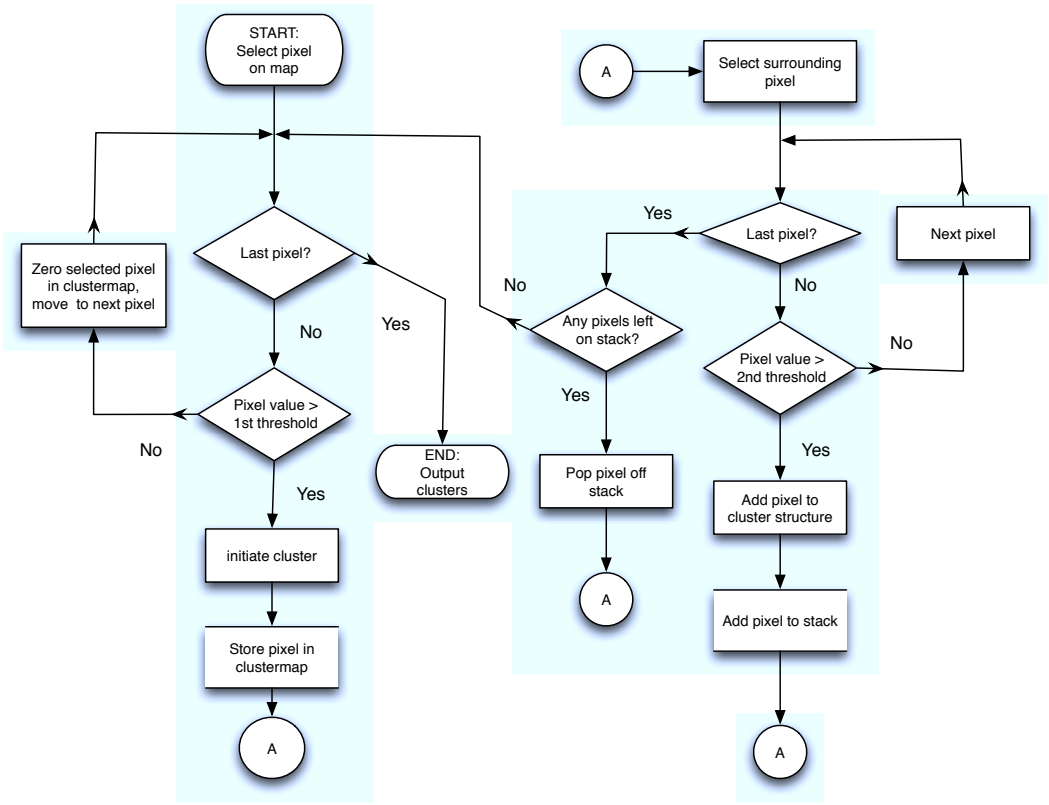

FIGURE 4.14: A flow chart detailing the algorithm used for the clustering algorithm.

passed across boundaries, there was the likelihood that the differing thresholds would stop the cluster at the boundary. Once again, this was solved easily - as we add pixels to a cluster, we continue to use the thresholds that were applied to the originating pixel, regardless of which region it is in. Testing has shown that this scheme works very well for the time-frequency maps produced by the correlator, and solves the issues that were found.

4.4.3 Multi-resolution time-frequency maps

Generally, time-frequency maps are generated at a range of Fourier transform lengths so that if a signal is present in the data one of the maps has the correct resolution for it to have maximal power (see [51], [50] & [49]). However, the architecture of SPHRAD does not naturally lend itself to generating and analysing multiple time-frequency maps of different resolutions. This is due to the long startup time required to generate the delay-products for each detector pair in the network, and its dependency on the Fourier transform length (i.e. each time-frequency map would require its own delay-product, see Table 4.2). There are, however, two methods that have been implemented in the code to facilitate the generation of multi resolution time-frequency maps, each of which has its own limitations and issues (demonstrated in Figure 4.15).

Constant l

In this method, the internal data products are generated at the maximum required l order, then the data is zero-padded (when required) to this length and FFTed so that we can form the products at the required frequency resolution. This method benefits from the fact that the internal data products are already generated at a high order, and therefore does not have to replicate any processing. Also, the time-frequency maps are generated using the required order to avoid aliasing. However, due to the size of the data products being processed (for $l = 182$ for a sample frequency of 2048Hz, so 33124 coefficients) these maps can take some time to generate ($\approx$ 30s for a 64s block of data with an expansion order of 182).

Low resolution l

This method replicates the actual procedure used to generate the multi-resolution time-frequency maps. Simply put, for each resolution all of the internal products are re-generated. However, to avoid the generating time becoming a major issue, very low order spherical harmonics are used to generate them, of the order of a few. The benefits for this method are obvious; the time-frequency maps are quickly generated (5 different Fourier resolutions from 1s down to 0.0625s) and the correct Fourier transform lengths are used for each map. The main disadvantage of this approach is that the use of such

FFT length (s)	order require
1	182
0.5	96
0.25	48
0.125	24

TABLE 4.2: This shows the minimum spherical harmonic expansion order required to completely represent an analysis of a certain Fourier length.

a low order spherical harmonic expansion will cause aliasing issues in the calculations - power from the higher spherical harmonic coefficients wrapping around into the lower. This will obviously cause problems when the coefficients are transformed into the image domain to form a skymap, as due to the aliasing, energy that would have been present in the higher order coefficients is forced down to the low order coefficients. This would alter the skymaps, but not the overall power in the map.

However, for transient gravitational wave signal identification in the time-frequency map, we are primarily interested in the total coherent power for each frequency band over the whole sky, and not the power in any specific direction. Once a cluster of interest has been identified then the skymap can be generated using the full spherical harmonic order, allowing sky pointing to be carried out correctly.

Thus, the use of this low-resolution l method may be a feasible method of a multi-resolution analysis.

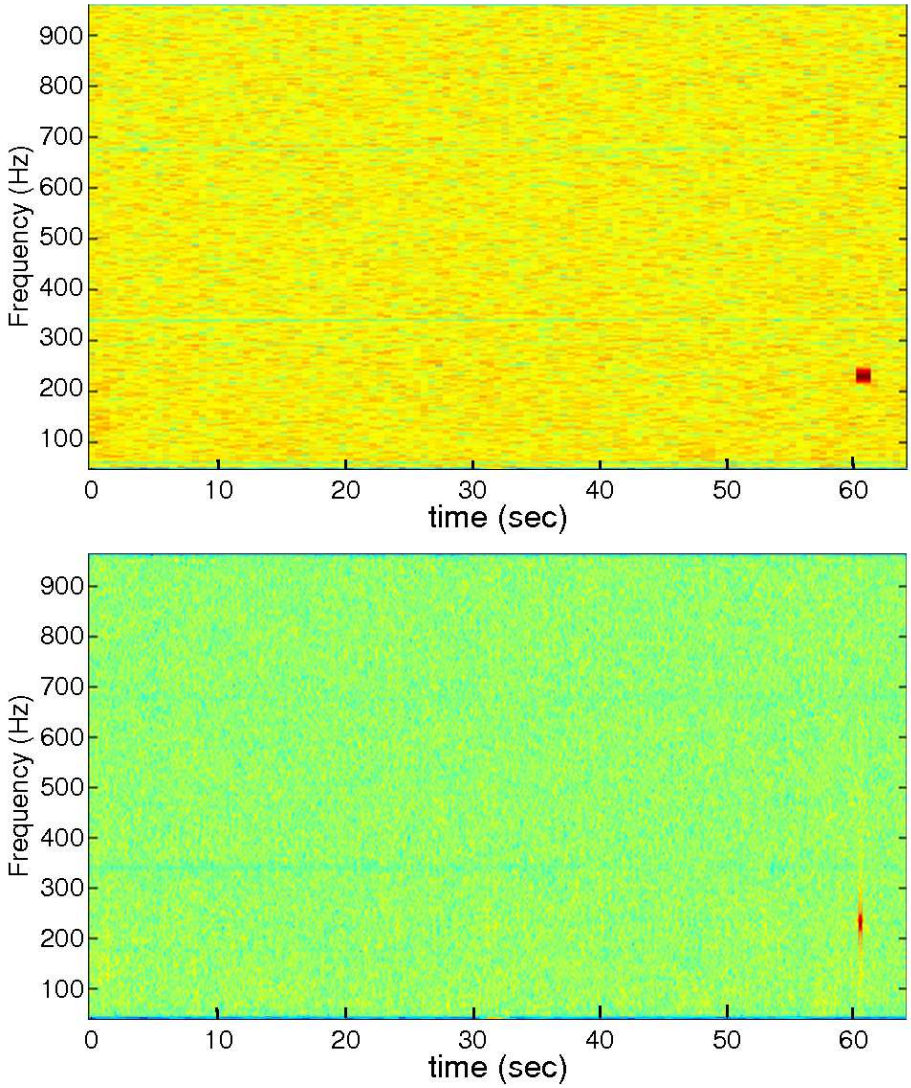

FIGURE 4.15: Two different resolution time-frequency maps in log-scale colouring containing an injected sine-Gaussian signal. The injection is clearly visible at 60 s and$f = 235$ Hz, the central frequency of the injection. The top image represents the use of a 1s FFT, while the bottom shows the a 1/8 s FFT, both for the constant l method.

Comparing the two different methods of generating multi-resolution time-frequency maps, we see that there is very little difference in the end result regardless of the method employed. Thus, there is no reason not to use the quickest version to generate the maps. The only point of concern is how to compare the pixels across different resolutions of the time-frequency maps. As the power in the spherical harmonic coefficients is summed, the normalisation of the differing numbers of coefficients means that it's a little more difficult to directly compare different size time-frequency maps.

As the present analysis is targeted at identifying long-duration signals, the multi-resolution analysis has a more limited use, so we will be using a single time-frequency resolution.

4.5 Sky localisation in an all sky search

The raw output of the radiometer is a series of spherical harmonic coefficients describing the correlation of each detector pair over the whole sky (both auto and cross).

The summation of these 'raw' spherical harmonics coefficients (the total network energy, equation 4.32) can be transformed using a Fast Spherical Harmonic Transform (see, for example, [70] &[71]) such as that implemented in the SpharmKit library [68].

As an example, consider the following results from an analysis of the S4 data. This is for a 128 second block of data (GPS time: 794611620 - 794611748) which contain the following waveforms (note that this is the same data as plotted in Figure 4.12):

Injection type	time (seconds)	sky position (phi, theta)	frequency (Hz)
Chirplet	28	(0.63, 1.18)	853
Sine-Gaussian	96	(5.66, 1.96)	853

TABLE 4.3: Parameters of the waveforms injected into the S4 segment 794611620 - 794611748

Plotting the spherical harmonic coefficients (see Figure 4.16), the two injected waveforms show up as vertical lines at approximately 28 & 96 seconds. Transforming the coefficients into the image domain generates the skymaps as shown above the time-spherical harmonic coefficient map in Figure 4.16.

As can been seen from Figure 4.16, the image domain data for the network power localises the injection to roughly the correct area of the sky. However, the auto- and cross-correlations that make up the network power (equation 4.32) can be coherently combined to form various statistics.

There are a plurality of different statistics that can be produced (and we can form anything that uses the auto- and cross- correlation terms), but not all of them will localise the event particularly well, if at all. To determine the statistic which performs most consistently, the pipeline can output a text file containing the sky location as determined by each included statistic. This allows a variety of plots to be made comparing the sky localisation of each. The statistics included in this analysis are:

standard likelihood

Used by X-pipeline and outlined in [see [49]) and [52]]. It is formed as the sum of the energy in the plus and cross polarisations.

$$E_{\rm sl} = E_{+} + E_{\times} \tag{4.38}$$

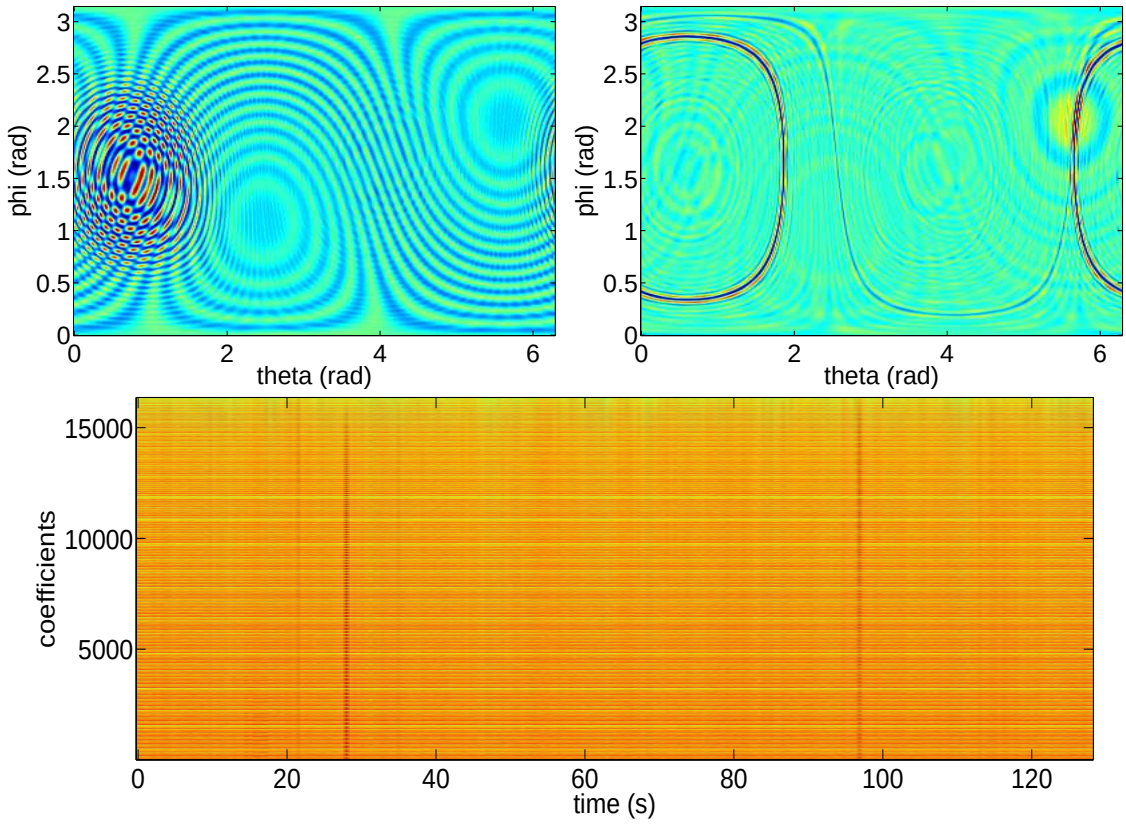

FIGURE 4.16: A time-coefficient map (bottom) for the 128 seconds of data shown in Figure 4.12, with corresponding skymaps for the injections (top). The "rippling" structure is caused by the constructive and destructive interference between the detector pairs in the presence of the signal.

modified standard likelihood

Used in the coherent waveBurst (cWB) pipeline [see [63]]. Modifies the standard likelihood using the null energy (total).

$$E_{\mathrm{cWB}} = \frac{E_{\mathrm{sl}} * E_{\times}}{E_{\mathrm{tot}} * (E_{\mathrm{tot}} - E_{\mathrm{sl}} + |E_{\times}|)} \tag{4.39}$$

ratio An adhoc statistic formulated by the author, which relies on the fact that a signal usually lies away form the diagonal on a coherent/incoherent energy plot. It is simply the sum of ratio of the coherent to incoherent energy for each of plus, cross and null (inverted and subtracted)

$$ratio = \frac{E_{+}}{I_{+}} + \frac{E_{x}}{I_{x}} - \frac{I_{\mathrm{null}}}{E_{\mathrm{null}}} \tag{4.40}$$

log ratio

Another adhoc statistic formulated by the author. This has its base in the ratio statistic, but it has been amended to make it one-sided (i.e. injections fall to one side of the background only). Essentially, we take the absolute log of each part of the sum (barring

the null energy as this is already one-sided) to build coherent energy. This should improve the sky localisation as we are more likely to get a bright spot on the skymap.

$$logratio = |\log\left(\frac{E_+}{I_+}\right)| + |\log\left(\frac{E_x}{I_x}\right)| + \frac{I_{null}}{E_{null}} \tag{4.41}$$

Tikhanov regularised statistic

Uses a Bayesian pre-factor to modify the standard likelihood to penalise regions with low antenna factor values [see [54]]. We also incorporate a range of amplitudes when calculating this statistic.

4.5.1 Testing the statistics

To determine which of the statistics was to be used to determine the sky-position for the S6 analysis, the pipeline was run over S4 data with a variety of injected signals with varying parameters. The recovered sky position for each injection was compared with its true sky position and the error plotted. In the figures that follow we have plotted the cumulative error in the sky position in two ways: 1) the top two plots show the curves for each of the statistics at a fixed injection scaling, and 2) the bottom two plots show the cumulative error for each injection scale for a single statistic (logratio and Tikhanov statistics). For each of the following figures, the curves demonstrate the fractional number of injections detected (i.e. number of injections detected/total number of injections) for certain error in the recovered sky position (in radians).

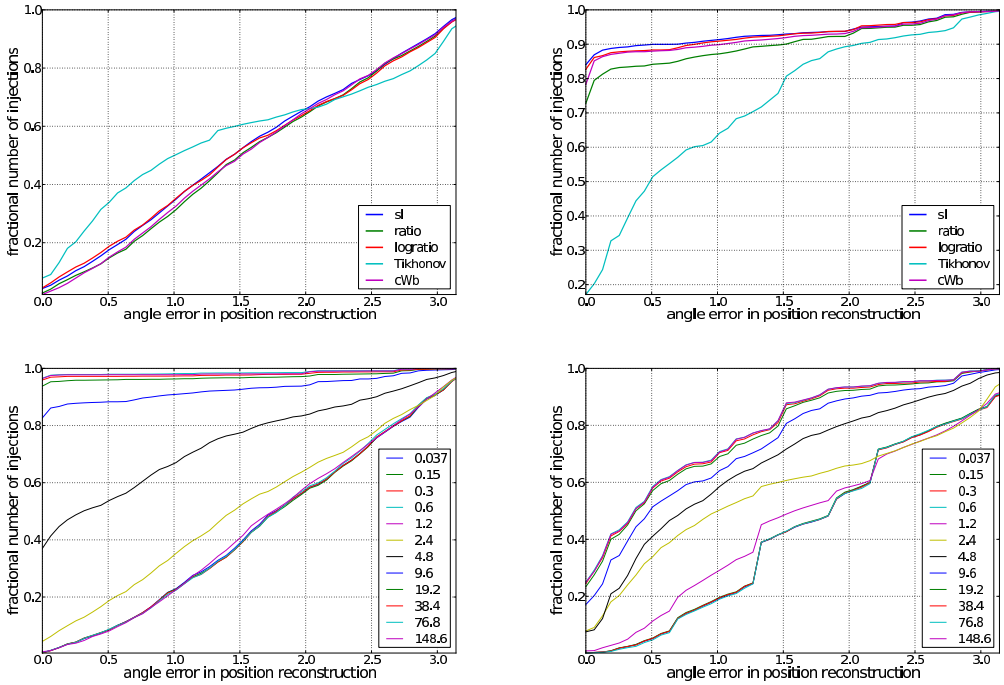

FIGURE 4.17: This figure show the cumulative error for a white noise burst, central frequency 300 Hz, and duration 0.1 s. The top two plots show the cumulative sky error for all of the statistics at an amplitude of $2.4 \times 10^{-21}\mathrm{Hz}^{-1}$ on the left, and $4.8 \times 10^{-21}\mathrm{Hz}^{-1}$ on the right. The bottom plots show the cumulative error for all of the injection scales (multiples of $1 \times 10^{-21}\mathrm{Hz}^{-1}$) for the logratio statistic on the left and the Tikhanov regularised statistic on the right

The plot in the top left of Figure 4.17 demonstrates how quickly the Tikhanov regularised statistic builds energy in the presence of a signal (in this case a 300 Hz white noise burst). It is tuned to detect quiet signals, which can be seen in the plot at the top right, where it has lost its advantage and is not doing as well as the other statistics.

The bottom two plots show the sky error for the logratio and Tikhanov statistics for the amplitude range $0.037 - 148.6 \times 10^{-21}\mathrm{Hz}^{-1}$. The bottom left plot shows the sky error for

the logratio statistic as the injection grows in amplitude - more of the injections are localised correctly. However, while Tikhanov statistic gets off to the better start (bottom right plot), it quickly saturates at a high amount of injections incorrectly localised.

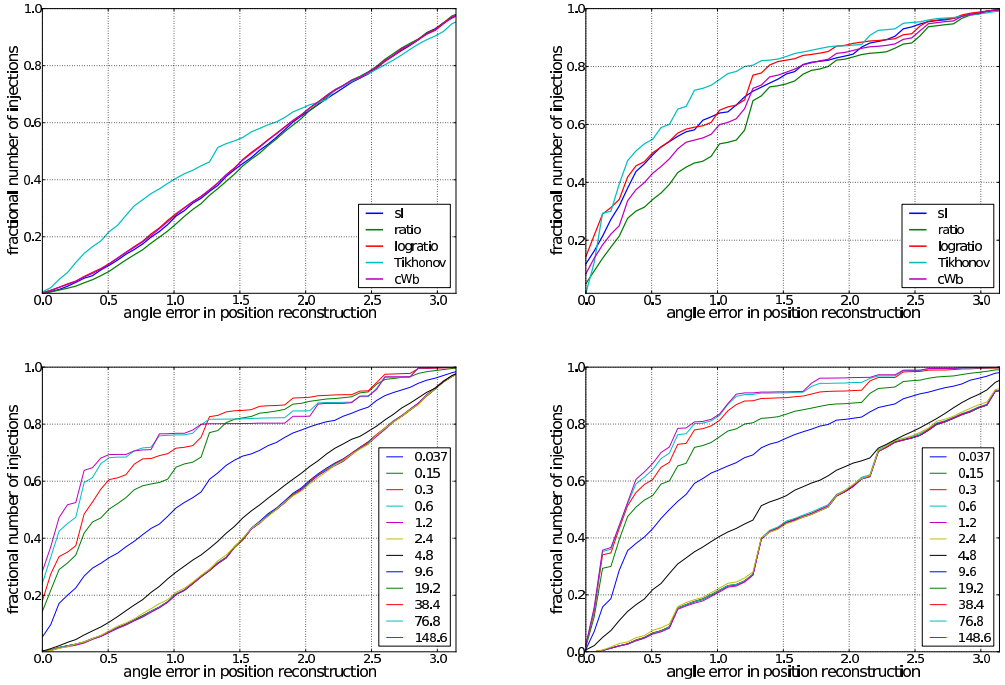

FIGURE 4.18: This figure show the cumulative error for a circularly polarised sine-Gaussian, central frequency 853 Hz, with Q=9. The top two plots show the cumulative sky error for all of the statistics at an amplitude of $4.8 \times 10^{-21}\mathrm{Hz}^{-1}$ on the left, and $19.2 \times 10^{-21}\mathrm{Hz}^{-1}$ on the right. The bottom plots show the cumulative error for all of the injection scales (multiples of $1 \times 10^{-21}\mathrm{Hz}^{-1}$) for the logratio statistic on the left and the Tikhanov regularised statistic on the right

Once again, the plot in the top left shows the Tikhanov regularised statistic reacting before the other statistics in the presence of a signal, Figure 4.18, in this case a circularly polarised sine-Gaussian with a central frequency of 843 Hz. The top right plot also demonstrates that this time, it keeps the lead too.

The bottom two plots show the sky error for the logratio (left) and Tikhanov (right) statistics for the amplitude range $0.037 - 148.6 \times 10^{-21}\mathrm{Hz}^{-1}$. Here, there is far more parity between the two statistics, although the logratio does localise more injection correctly.

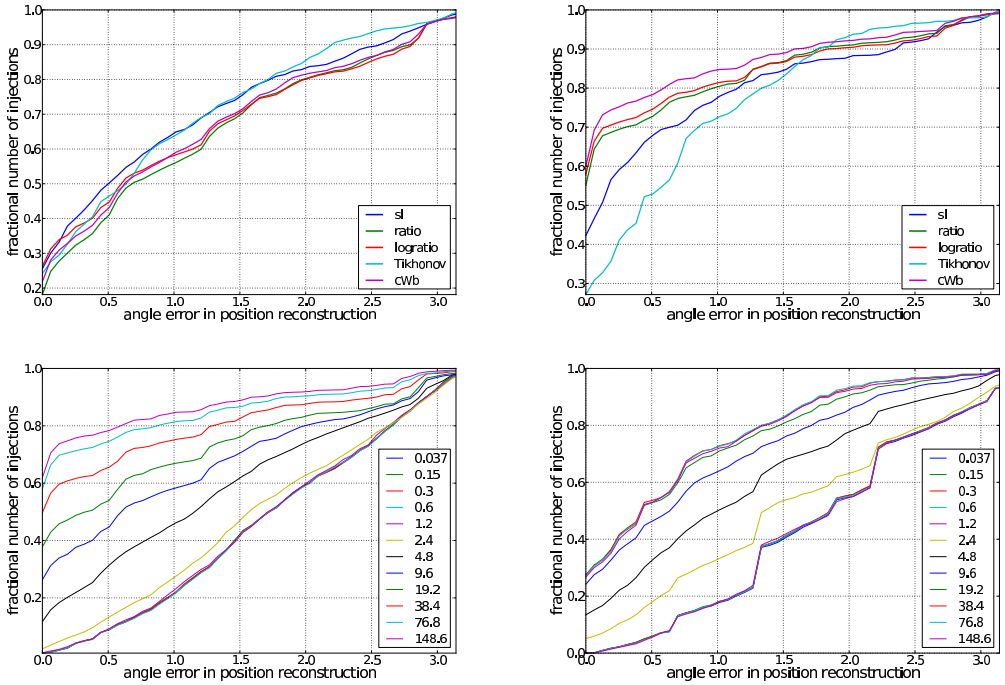

FIGURE 4.19: This figure show the cumulative error for a linearly polarised sine-Gaussian, with a central frequency of 235 Hz, and Q=3. The top two plots show the cumulative sky error for all of the statistics at an amplitude of $9.6 \times 10^{-21}\mathrm{Hz}^{-1}$ on the left, and $76.8 \times 10^{-21}\mathrm{Hz}^{-1}$ on the right. The bottom plots show the cumulative error for all of the injection scales (multiples of $1 \times 10^{-21}\mathrm{Hz}^{-1}$) for the logratio statistic on the left and the Tikhanov regularised statistic on the right

Figure 4.19 demonstrates that the Tikhanov statistic does not always get off such a good start (top-left), although, none of the statistics (top right) do a particularly good job for this 235Hz short-duration sine-Gaussian (due to the FFT duration mismatch). Once again, the bottom two plots show the sky error for the logratio and Tikhanov statistics for the amplitude range $0.037 - 148.6 \times 10^{-21}\mathrm{Hz}^{-1}$. Here, the bottom left plot shows the sky error for the logratio statistic with only 60% localised correctly. This is better than the Tikhanov statistic (bottom right plot) which can only achieve 30 % correct.

On initial inspection it would appear that the Tikhanov regularised statistic would make a good statistic for sky localisation but it does have some drawbacks, namely even when detecting signals well, it does not really localise their position accurately. It would make more sense to use it as a detection statistic rather than a sky localisation statistic. The performance of the cWB, standard likelihood and logratio statistics are quite similar in that they each take the lead in different situations and for different injection types and frequencies.

The logratio statistic was not the most accurate, but it appeared to have the most consistent performance. This is why it was chosen as the sky localisation statistic for the S6 analysis. Another method could have been to allow the pipeline to select the most appropriate statistic.

4.6 Glitch rejection in the Spherical Harmonic basis

The two main ways to view the output of the correlator are the time-frequency map (e.g. Figure 4.4)and the time-spherical harmonic coefficient map (e.g. Figure 4.5). Contrasting the two representations yields a rather interesting discovery (see Figure 4.20):

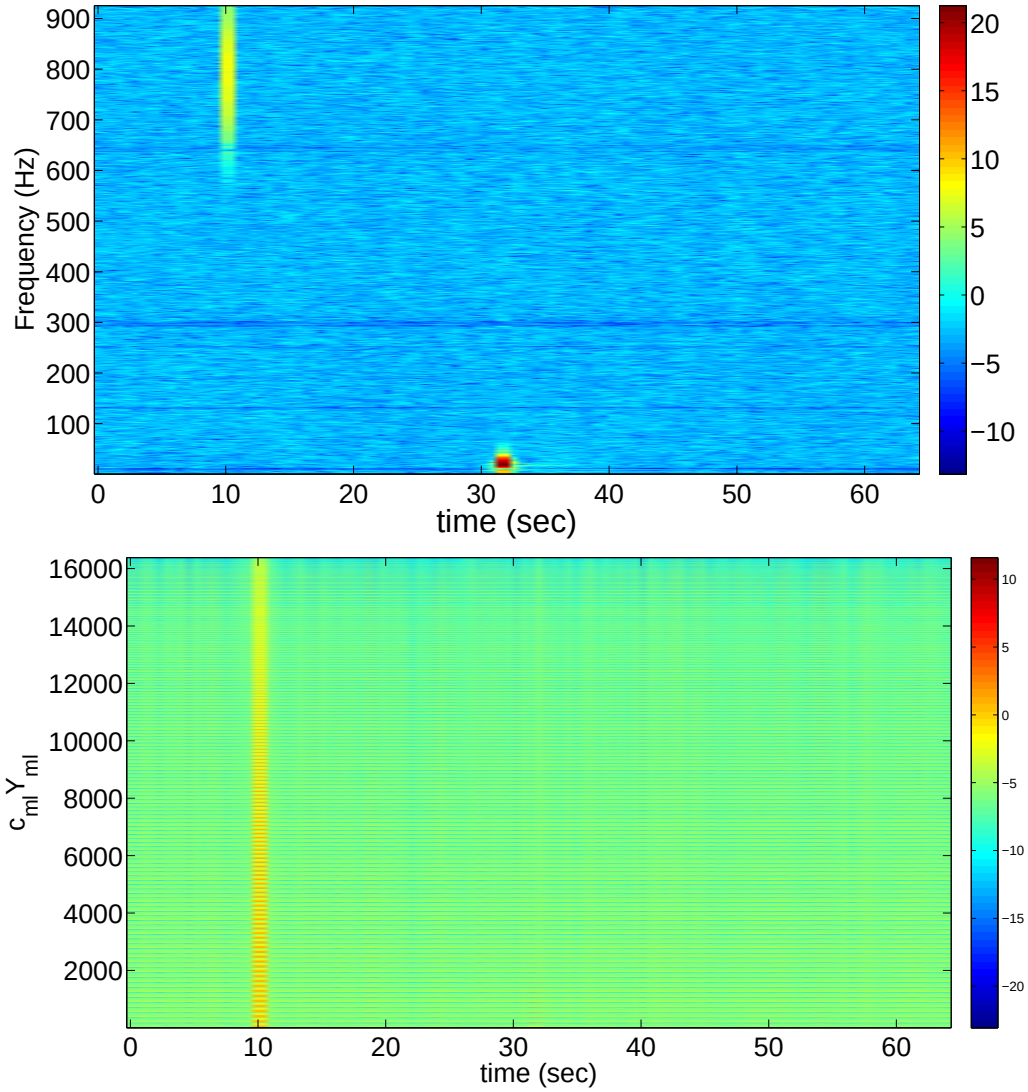

FIGURE 4.20: A dual-view of the same 64 seconds of data - time-frequency map top and time-spherical harmonic map bottom. An injection can be clearly seen at 10s (central frequency 835 Hz), while a low-frequency glitch is present at 32s.

The time-frequency map has two clear features, the injected signal at 10 seconds and the glitch at 32 seconds. In this representation there is nothing obvious that can be used to tell the two features apart, so when clustering the pixels the glitch may be chosen as the top trigger for this data block if it is loud enough. However, the time-spherical harmonic coefficient map only shows the injection at 10 seconds.

Now, both of the representations derive from the higher-dimensional structure shown in Figure 4.3, so they should contain the same total power. However, as can be seen from Figure 4.20, the distribution is very different, which implies something about the underlying structure. For simulated signals injected with a prescribed central frequency and peak time, the time-frequency map lights up for the duration and bandwidth expected. Similarly, for glitches, the time-frequency map also lights up the pixels around the expected time and frequency. In contrast, the time-spherical harmonic coefficient map shows structure along the coefficients at the time of an injection, but nothing at the time of a glitch.

This leads to the conjecture that the spherical harmonic coefficients could be used as some form of glitch discriminator, and a quick aside to think about how details on the sky map relate to the spherical harmonic coefficients should convince why this is so.

- **Behaviour with signal:** At the true sky location, a gravitational wave signal appears in phase in all of the detectors that can see it, giving strong correlations. As we move away from the true sky location, time shifts move the gravitational wave signal in and out of phase between different detectors, producing an 'interference pattern' in the sky image (see[72, 73]). The spacing of the fringes in the interference pattern is determined by the signal's frequency content and the detector pair baseline. The fine detail in the image directly corresponds to the power in the l^{th} mode of the spherical harmonic coefficients.

  In the time-frequency representation, we have essentially summed over the whole sky, so direction/time delays are not important. However, in the time-harmonic representation a signal is resolved to a patch of the sky meaning that power will be present in most of the coefficients - generally none will stand out more than any other.

- **Behaviour with glitch:** By contrast, independent glitches in different detectors are not correlated with each other, regardless of the time-shift applied to the data. Therefore the corresponding sky maps tend to be loud, but featureless as most of the power is concentrated into the $l = 0$ mode.

  As the time-frequency representation is independent of the time delay, we see similar behaviour to that with a signal. In spherical harmonic terms, the zeroth component, $c_{00}Y_{00}$, corresponds to the amount of energy present over the whole sphere, while the higher order terms relate to fine details. Thus, the zeroth component is a measure of amount of incoherent energy there is in the skymap, which directs relates to the glitch energy, as this is mostly present in the $c_{00}Y_{00}$ coefficient.

The new coherent statistic works on the principle that we can decouple the 'incoherent' and 'coherent' energy using the decomposition into spherical harmonics. Using the nomenclature c_{lm} to represent the spherical harmonic coefficients, the $|c_{00}|$ coefficient corresponds to the incoherent energy, while the sum of the rest of the coefficients corresponds to the coherent

energy. For a signal (which has a preferred sky location), most of the coefficients should have some measure of power, however, for a glitch, most of the energy is concentrated in the c_{00} coefficient.

This motivates the following statistic to be used for discriminating signals and glitches:

$$\mathbf{\Gamma} = \frac{\sum_{l=1}^{l_{\mathrm{MAX}}} \sum_{m=-l}^{l} |c_{lm}|}{|c_{00}|} \tag{4.42}$$

where l_{MAX} is defined using a combination of baseline length and the data sample frequency (see [66] for more details) and the statistic is computed for each time bin. When plotted, any time bin that rises above the background is a signal, and any that fall below is a glitch. Figure 4.21 shows how the statistic looks for the same segment of data that was used to produce the maps.

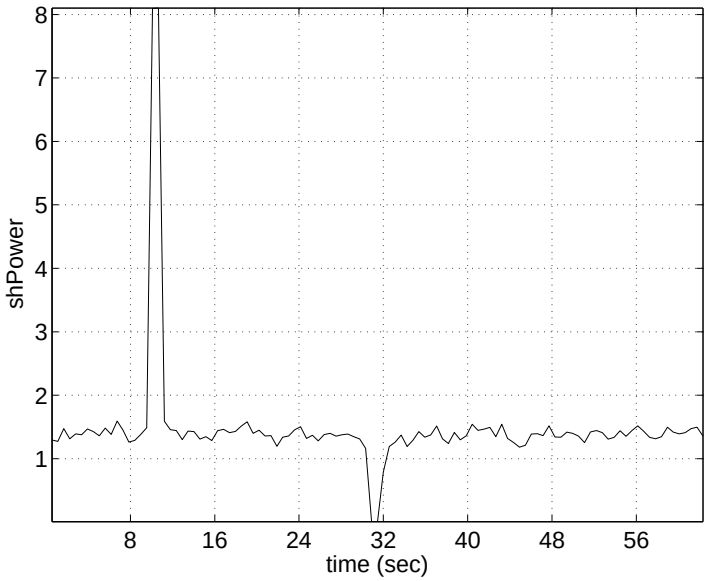

FIGURE 4.21: $\mathbf{\Gamma}$ for the same 64 seconds of data used to generate Figure 4.20. The injected gravitational-wave signal rises well above the background noise at t = 10 s, while the glitch drops below at t = 31s.

This statistic is useful because it is very fast to generate and we do not have to be concerned with sky localisation. To carry out the same analysis using E_+, $E_\times$ or some combination of the other statistics would require the energy associated with the most probable direction to be determined from a skymap for each time bin.

4.6.1 Performance of the statistic

To demonstrate how the new statistic, $\mathbf{\Gamma}$, performs, it was applied to both simulated and real data, sampled at 2048 Hz, and using the Hanford-Livingston-Virgo (HLV) network. White

Noise Bursts (WNB) containing 2 independent polarisations of approximately equal power were injected into the data with a central frequency of 300 Hz, a bandwidth of 50 Hz and a duration of 0.1 seconds. A range of root-sum-square injection amplitudes were used, from 3.7×10^{-23} $\mathrm{s}^{-1/2}$ to 1.5×10^{-19} $\mathrm{s}^{-1/2}$, defined as:

$$h_{RSS} = \sqrt{\int_{-\infty}^{\infty} (|h_+(t)|^2 + |h_\times(t)|^2) dt} \tag{4.43}$$

where $h_+(t)$ and $h_\times(t)$ are the plus and cross-polarisation strain functions, and h_{RSS} has units of $\mathrm{Hz}^{-1/2}$.

4.6.1.1 Simulated noise data

To test the SphRad pipeline under "controlled" conditions, it was used to processed one month of simulated noise data. Now, real detector data contains glitches - bursts of high power in a single detector caused by machine noise. Obviously, our simulated noise data is actually Gaussian and does not contain any real detector glitches, so a very simple model was used to generate a population of roughly one glitch per hour for the analysis. This consisted of inserting a family of sine-Gaussians with random parameters (central frequency, duration, peak time, amplitude) into the simulated data. To mimic the behaviour of real glitches, each was injected into one of the randomly chosen detector streams so that each had a population. In the course of the analysis, all of the standard coherent energies were generated (plus, cross, null and standard likelihood), as well as a timeline of the amount of power in the spherical harmonic coefficients.

In the following scatter plots, I_{SH} corresponds to the power in the zeroth spherical coefficient (i.e. $|c_{00}|$, the incoherent energy), while E_{SH} corresponds to the sum of the $\ell > 0$ coefficients. The energy in the null stream (i.e. the projection of the data in which a gravitational wave signal should be cancelled, E_{null} and I_{null}), is computed using the cross-product of the orthonormal antenna factors for the cross and plus-polarisation, see [49].

Figure 4.22 shows such a scatter plot. The bifurcation pattern is the defining feature of these type of plots, and shows good separation between the black crosses of the background (generated using the standard time slides method), and the simulated signals. Note that there are some coloured crosses present in the background. These are not injections, but are due to the post-processing selecting the "wrong" trigger to associate with an injected signal. There are many reasons this could happen, which is why a cut is usually applied to the trigger distribution. Here, we have just plotted the raw triggers to get a feel for their distribution. We will compare with known methods in the following section.

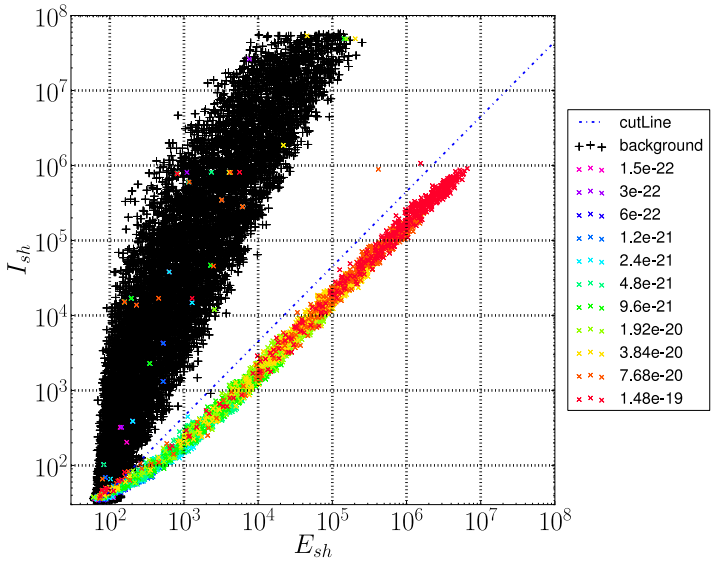

FIGURE 4.22: Scatter plot of the coherent spherical harmonic energy, E_{sh}, against the incoherent, I_{sh}, for simulated glitchy data, and injected WNB gravitational wave signal of various amplitudes. Note that there are some coloured crosses in the background. This is due to the post-processing selecting the "wrong" trigger to associate with an injected signal.

4.6.2 Science Run 4 data

In order to test the statistic on real data containing a real glitch population, the same analysis was run over a month of Science Run 4 (S4) data [74] when the three LIGO detectors were taking science data. In reality, the network used to take this data was the Hanford-Livingston network (H1L1H2), whereas the network we are interested in is Hanford-Livingston-Virgo (H1L1V1). To simulate this, we 'relabeled' the H2 data stream as V1 (i.e. once the H2 data was loaded, treat it as if it had been generated by the Virgo detector). As we were only interested in using the data as a source of noise and glitches, this was sufficient. We note, however, that in reality the H1 and H2 detectors were co-located; meaning noise between the data streams is correlated, while we are treating them as independent. This is, therefore, a "worse case" scenario.

Another thing of importance to note is that the values of the null energy, E_{null} & I_{null}, are calculated using the *best* recovered sky position as determined by the pointing statistic (in this case the standard likelihood, E_{sl}). The energy in the spherical harmonic coefficients is not limited in such a way (but could be), so it is less likely to be affected by sky position errors.

Once again, there is a clear separation between the background and the injected signals (Figure 4.23), but notice that the angle between them has reduced. This is a common feature of these plots when using real data and appears to be related to the amount of correlation in the noise between different detectors. This will be more of an issue here than in the analysis of the real V1 data, because there will be coupling between H1 and 'V1' (actually the relabelled H2).

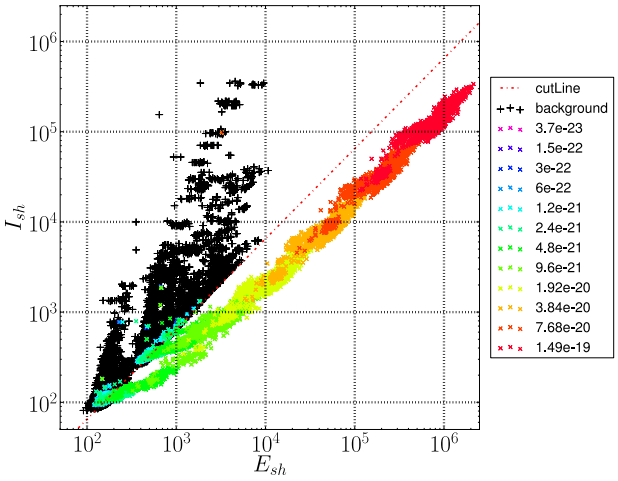

FIGURE 4.23: Scatter plot of the coherent part of the spherical harmonic coefficients versus the incoherent part for S4 data.

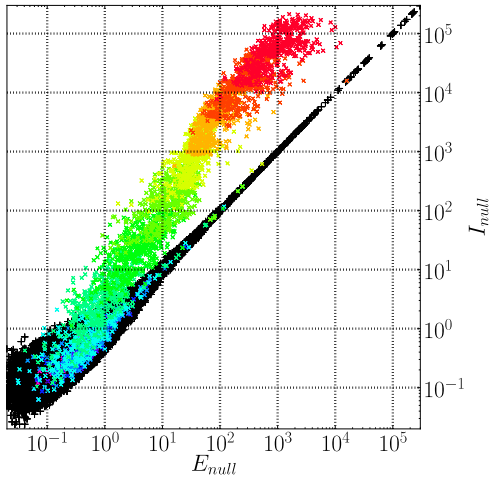

FIGURE 4.24: Scatter plot of the coherent and incoherent null energy for S4 data for best recovered sky position.

Generally, the values of the spherical harmonic energies in the background generate a reasonably clean edge with which to fit a straight line to. However, the orientation of the line can change. To cope with this, the following steps are used to generate the cut line

- From the background distribution generate lines following the median and minimum values of the incoherent energy,
- Determine the perpendicular distance of each point in the minimum line to the median line,
- If any of the points in the minimum line are "above" the median line, then cull those points, and recompute the minimum line from the points left,
- Calculate the average distance of the minimum line,
- Drop points from the minimum line that are closer to the median line than the average,
- Use least squares to generate a line of best fit for the minimum line.

This generates a close fitting line that culls most of the background triggers.

Comparing the distribution of the spherical harmonic energy, Figure 4.23, to the null energy, Figure 4.24, we see they have many elements in common - the well defined background, the spread of triggers associated with injections and the clear separation between loud injections and the background. But there is one main difference, the null energy must be calculated for the most probable sky direction, while the spherical harmonic power is all-sky, which we get as a by-product of the way in which the pipeline analyses that data. It can be generated

exceptionally quickly and be used as a first pass discriminator of signal/glitch. It is useful because it is unaffected by sky localisation errors, but this also makes it less sensitive as we are considering the whole sky.

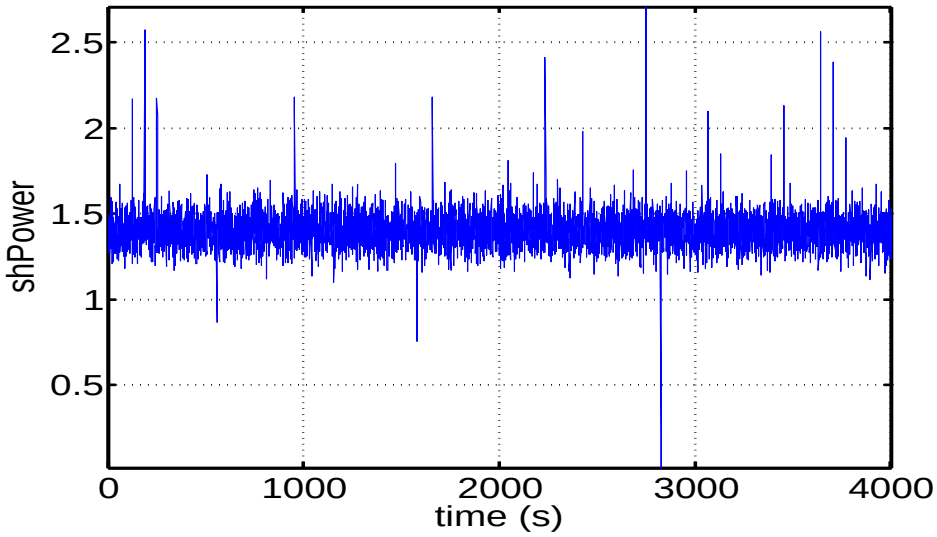

FIGURE 4.25: The Γ statistic for 4000 seconds of data starting at GPS time 793502657. Note that the peaks correspond to times at which a signal was injected, while the troughs are single detector glitches in the data stream.

As an additional demonstration, Figure 4.25 shows the plot of spherical harmonic power, $\mathbf{\Gamma}$, for a large stretch of data. There is a horizontal band at $\mathbf{\Gamma} \approx 1.4$ which is used to discriminate glitches and signals. Here, the injected signals have a hrss of 2.4×10^{-21} $s^{-1/2}$ and can seen rising above the band, while the three large, clear, glitches can be seen dropping below the band (at approximately 550s, 1600s and 2800s).

Compare this plot of Γ to any of those shown in Chapter 5 (e.g. panels d of Figure 5.1 or 5.2). Note that the "background" value (i.e. the average) of Γ appears to be higher in S6 data ~ 2, whereas in Fig 4.25 (i.e. S4 data) it is ~ 1.5. This appears to be measure of the noise content in the detector data and warrants further exploration (i.e. compare the average Γ values for white noise, S4, S5 and S6 data streams).

Chapter 5

Results

In this chapter we present results of our analysis of the LIGO-Virgo data from the most recent science run. We will demonstrate the performance of the SphRad pipeline on real gravitational wave data, and compare our results to those generated by existing analysis pipelines (see [11], [75] and [76]). Note that we are presenting the "closed box" results (efficiency of recovering injected signals and properties if the background) and not the zero-lag triggers. This is because the SphRad pipeline has not yet been reviewed for approval by LSC/Virgo. In Section 5.1 we will give a brief summary of the S6/VSR2,3 science run. In Section 5.2 we will introduce the different types of simulated gravitational waveforms used to test the SphRad pipeline. In Section 5.3 we will summarise the analysis procedure, while in Section 5.4, Appendix D & Appendix E we will present the results of the SphRad analysis of the three distinct epochs of S6/VSR2,3. Comparison of our results to the standard results from the coherent WaveBurst analysis pipeline will be presented in Section 5.5.

5.1 Science Run 6

In the time period between the end of the S5 run and the start of the S6 run, the Livingston and Hanford detectors were upgraded to their 'Enhanced' stage, while Virgo was similarly upgraded. This expanded the window of possible detection. Science run 6 was the last joint data taking run of the LIGO and Virgo detectors before they were decommissioned for the advanced detector upgrades. It was divided into 4 subsets, defined by commissioning breaks and detector/configuration changes. These were as follows:

Period A: Lasted from 21:00 UTC, July 7^{th} 2009, to September 1^{st}, 2009 at 00:00 UTC (GPS times: 931035615-935798415). All three detectors were operating (H1, L1 & V1).

Period B: Ran from 00:00 UTC, September 26^{th} 2009 to January 8^{th} 2010 at 22:00 UTC (GPS times: 937958415-947023216). Virgo had a commissioning break from 07:03:07

UTC on October 5^{th} 2009 to October 8^{th} 2009 at 17:37:00 UTC (GPS times: 938761402-939058635). Besides this commissioning time, all three networks were operating (H1, L1 & V1).

Period C: Started at 22:00 UTC, January 8^{th} 2010 to June 26^{th} 2010 at 00.00 UTC (GPS times: 947023216-961545615). The Virgo detector was undergoing upgrading during this period and did not participate in this portion of the run. Only H1 and L1 were operational.

Period D: Began at 00.00 UTC, June 26^{th} 2010 to October 20^{th} 2010 at 15.00 UTC (GPS times: 961545615-971622015). All three detectors were operational (H1, L1 & V1).

We test SphRad by analysing data from times when all three detectors were operational. It does limit the amount of usable data available for the analysis somewhat, but it is important to be able to test the sky localisation facilities of the pipeline. Triple coincident data would allow the pipeline to localise an event on the sky, with enough data available to perform a thorough analysis. The amount of live time for triple coincident data (network - H1L1V1) for each period is as follows:

Period A: Triple coincident data spans non-contiguously from GPS start time, 931081116, to GPS end time, 935147785, giving 1015068.0 seconds. After discarding segments less than 128 seconds i length (our analysis length), we end up with 949547.5 seconds of live time (just under 11 days). For this period, the duty cycle for triple coincident data was 32%.

Period B: Non-contiguous triple coincident segments spans from GPS start time, 938000353 to GPS end time, 946418415, giving 1416462.0 seconds in total. After removal of short segments, we were left with 1401561.0 seconds (just over 16 days). For this period, the duty cycle for triple coincident data was 31%.

Period C: no triple coincident data available.

Period D: The non-contiguous segments span from GPS start time, 965599233, to GPS end time, 971504457, giving 2328324.0 seconds of triple data. Removal of short segments left us 2311791.0 seconds to analysis (almost 27 days). For this period, the duty cycle for triple coincident data is slightly higher at 37%.

5.2 Simulated signals

To assess the performance of the pipeline, a series of simulated signals are injected into the data stream at different sky positions and injection scales. The data is analysed as normal, but

the triggers generated by the pipeline can be assessed in light of the injected signals and we can determine how many of them have been recovered. Aside from waveform specific parameters, each waveform is characterised by its peak time, τ_0 and its intrinsic signal strength, the root-sun-square signal amplitude, $h_{\rm rss}$:

$$h_{\rm rss} = \sqrt{\int |h_+(t)|^2 + |h_x(t)|^2 dt}. \tag{5.1}$$

For the S6 analysis, we inject white noise bursts, linearly and circularly polarised sine-Gaussians, inspirals, accretion disk instabilities, and poly-oscillating waveforms.

5.2.1 White noise bursts (WNB)

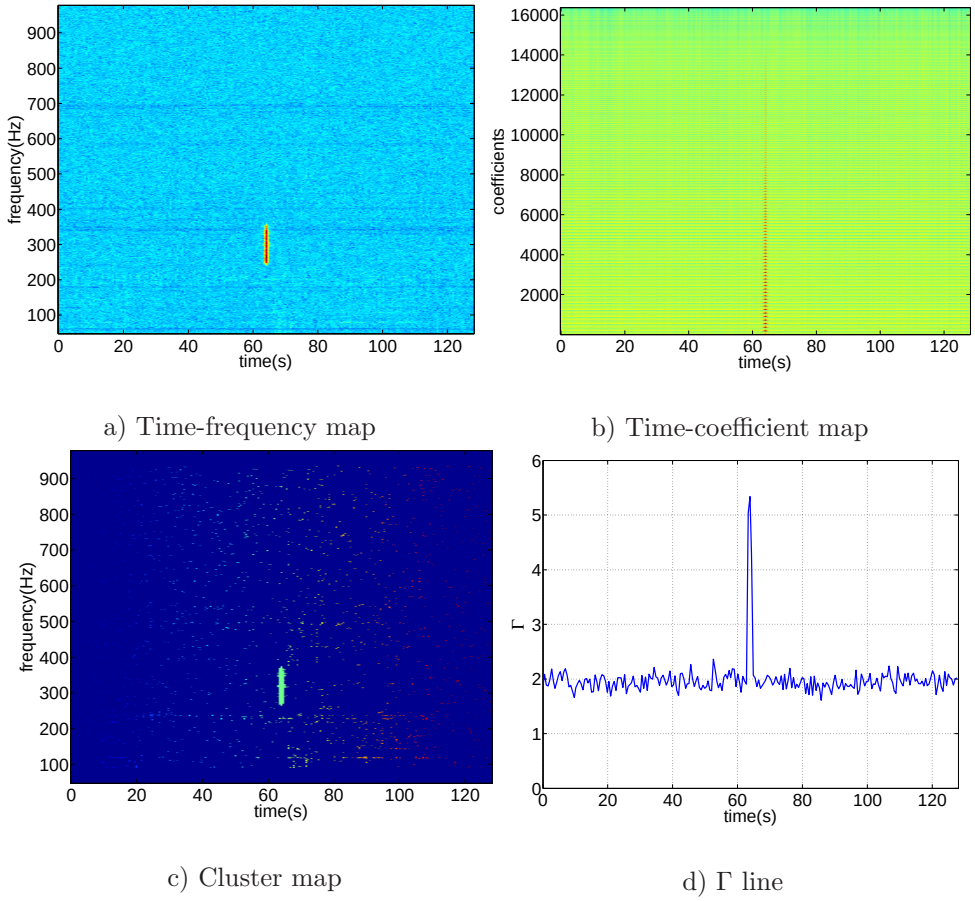

a) Time-frequency map

b) Time-coefficient map

c) Cluster map

d) Γ line

FIGURE 5.1: The plots in this figure demonstrate how a WNB with $h_{\rm rss}$= 7.68e-20 $Hz^{-1/2}$, central frequency of 300 Hz, bandwidth of 100 Hz, duration of 0.1 seconds and a peak time of 63 seconds, appears in a ; a) time-frequency map, as a small vertical rectangle centred on f_c, τ_0, b) time-spherical harmonic coefficient map, as a disjointed peak at τ_0, fading towards higher coefficients, c) clustermap, as a tightly defined rectangle, d) plot of Γ against time, which shows significant activity at the peak time.

The WNB family of waveforms consist of a stochastic white-noise time series that is band-limited over some frequency range and Gaussian windowed in time. These could be considered as a simple model of the stochastic emission in, for example, a core-collapsed supernova, see

([77]). The parameters that define the WNB waveform are its central frequency, f_c, bandwidth, Δf and width, Δt of the Gaussian time envelope (essentially its duration). Each polarization is an independent sample of noise that is white over the chosen band ($f_c - \Delta f$ to $f_c + \Delta f$ where the lowest and highest frequencies allowed are 0 and the Nyquist frequency respectively). Figure 5.1 shows some of the data products output by the SPHRAD pipeline in response to a WNB injected into S6A data with a $h_{\rm rss}$ of 7.68e-20 $Hz^{-1/2}$ with f_c=300 Hz, Δf=100 Hz (i.e. the frequency ranges from 250Hz-350Hz), $\Delta 4$=0.1 seconds and τ_0=63 seconds. The injection is clearly visible in each of the different representations, with the clustering algorithm (Figure 5.1(c)) only picking out the pixels associated with the WNB, and Γ showing a strong response (Figure 5.1 (d)) at the time of the injection.

5.2.2 Linearly polarised Sine-Gaussians (SGL)

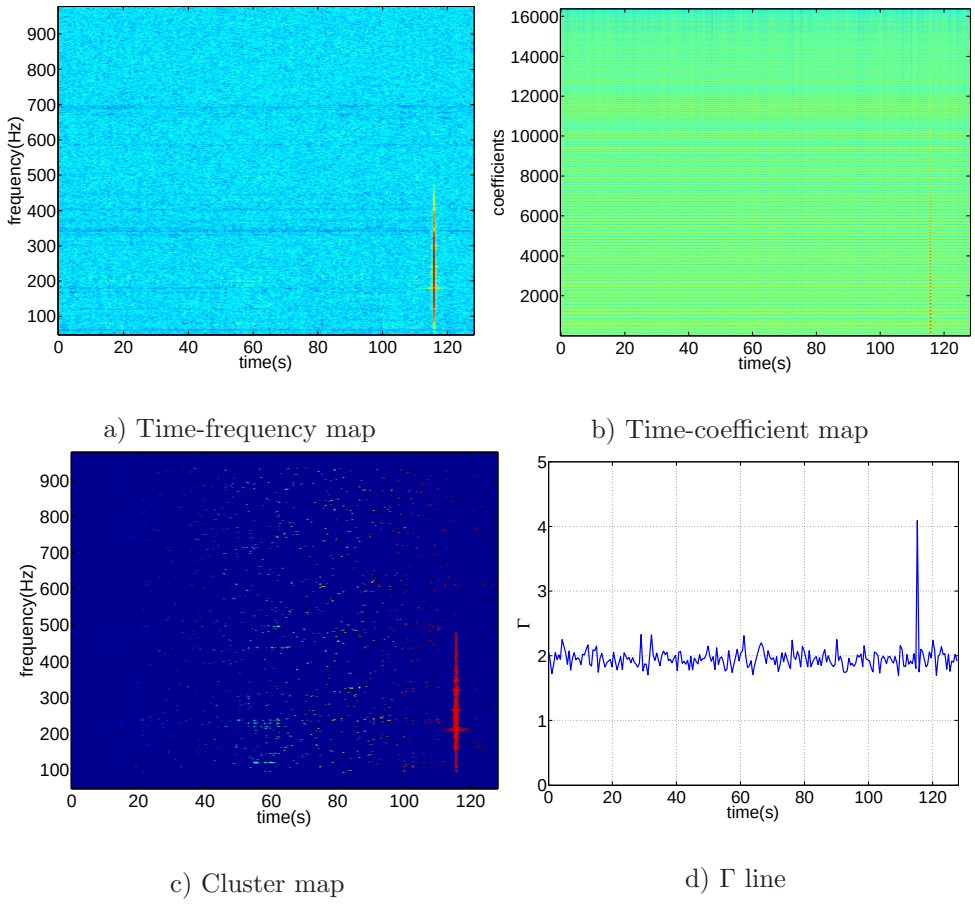

a) Time-frequency map

b) Time-coefficient map

c) Cluster map

d) Γ line

FIGURE 5.2: The plots in this figure demonstrate how a SGL with $h_{\rm rss}$ of 7.68e-20 $Hz^{-1/2}$, a central frequency of 235 Hz, a Q value of 3 (duration of 0.003 s), and peak time of 118 seconds, appears in a ; a) time-frequency map, as a fairly long vertical line centred at τ_0 with some width, b) time-spherical harmonic coefficient map, as a small disjointed peak at τ_0, fading towards higher coefficients, c) clustermap, as a narrow vertical line, d) plot of Γ against time, which is peaked at the peak time.

The SGL family of waveforms consist of linearly polarised Gaussian modulated sinusoids. They represent an ad-hoc (e.g. non-physical) model which is simple and narrowband (usually Q=9).

This is used as a standard waveform for testing burst methods (see [78]). The SGL waveform is defined by a set of parameters; its root-sum square signal amplitude, $h_{\rm rss}$, central frequency, f_c and the quality factor, Q (number of waveform cycles), as follows

$$h_+ = h_{\rm peak} \sin(2\pi t f_c)\, {\rm e}^{-t^2/\tau^2} \tag{5.2}$$

$$h_\times = 0, \tag{5.3}$$

where

$$h_{\rm peak} = h_{\rm rss} \sqrt{(4\sqrt{\pi} f_c)/(Q(1-\exp(-Q^2)))}, \tag{5.4}$$

$$\tau = Q/(\sqrt{2}\pi f_c). \tag{5.5}$$

The effective bandwidth is given by

$$bw = \sqrt{1}/(4\pi\tau). \tag{5.6}$$

The plots in Figure 5.2 demonstrate the results of the pipeline analysing a linearly polarised sine-Gaussian injected with a $h_{\rm rss}$ of 7.68e-20 $Hz^{-1/2}$ into S6A data. Once again, the injection is clearly visible in each representation (similar to the WNB, Figure 5.1), however, the peak Γ value (Figure 5.2(d)), isn't as high as for the WNB. This could possibly be explained due to the shorter duration of the SGL as compared to the WNB. We are using a fixed FFT length of 1 second, so do proportionally worse for shorter signals.

5.2.3 Circularly polarised Sine-Gaussians (SGC) - short duration

The SGC family of waveforms are similar to the SGLs, but are circularly polarised Gaussian modulated sinusoids as opposed to linearly polarised. These also represent an ad-hoc (e.g. non-physical but simple model), and are a standard waveform for testing burst methods (see [78]). The SGC waveform is defined by the following set of parameters; its root-sum square signal amplitude,$h_{\rm rss}$, central frequency, f_c and duration, τ, as follows

$$h_+ = h_{\rm peak} \cos(2\pi t f_c) {\rm e}^{-t^2/\tau^2} \tag{5.7}$$

$$h_\times = h_{\rm peak} \sin(2\pi t f_c) {\rm e}^{-t^2/\tau^2}, \tag{5.8}$$

$$\tag{5.9}$$

where

$$h_{\rm peak} = h_{\rm rss} \sqrt{(4\sqrt{\pi} f_c)/(Q(1-\exp(-Q^2)))} \tag{5.10}$$

$$\tau = Q/(\sqrt{2}\pi f_c). \tag{5.11}$$

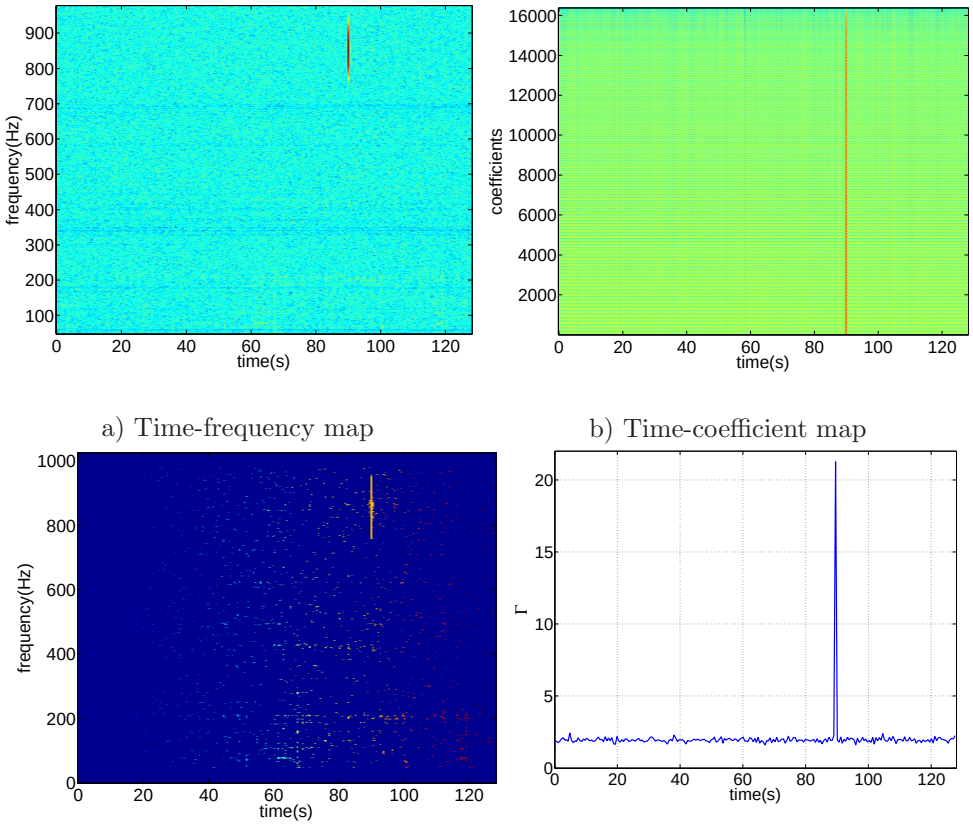

a) Time-frequency map

b) Time-coefficient map

c) Cluster map

d) Γ line

FIGURE 5.3: The plots in this figure demonstrate how a SGC with $h_{\rm rss}$ of 7.68e-20 $Hz^{-1/2}$, a central frequency of 853 Hz, a Q value of 9 (duration of 0.002 seconds), and peak time of 90 seconds, appears in a ; a) time-frequency map, as a bright compact rectangle centred at τ_0 and f_c, b) time-spherical harmonic coefficient map, as a tall vertical line at τ_0, with power in all of the coefficients, c) clustermap, as a small rectangle, d) plot of Γ against time, which is strongly peaked at the peak time.

Note that the bandwidth is given by

$$bw = \sqrt{1}/(4\pi\tau), \tag{5.12}$$

demonstrating that the longer the wavelength, the more it tends towards being monotonic in frequency. The plots in Figure 5.3 present the output from the pipeline when analysing S6A data containing a circularly polarised sine-Gaussian injected with a $h_{\rm rss}$ of 7.68e-20 $Hz^{-1/2}$, f_c of 853 Hz and a Q value of 9 (duration 0.002 seconds). The injection is clearly visible in each representation, but it is the significantly high peak Γ value (Figure 5.3(d)), which is most noticeable. This can be explained when you compare Equations (5.7) to (5.7) it is the addition of the non-zero $h_\times$ polarization that increases its detectability.

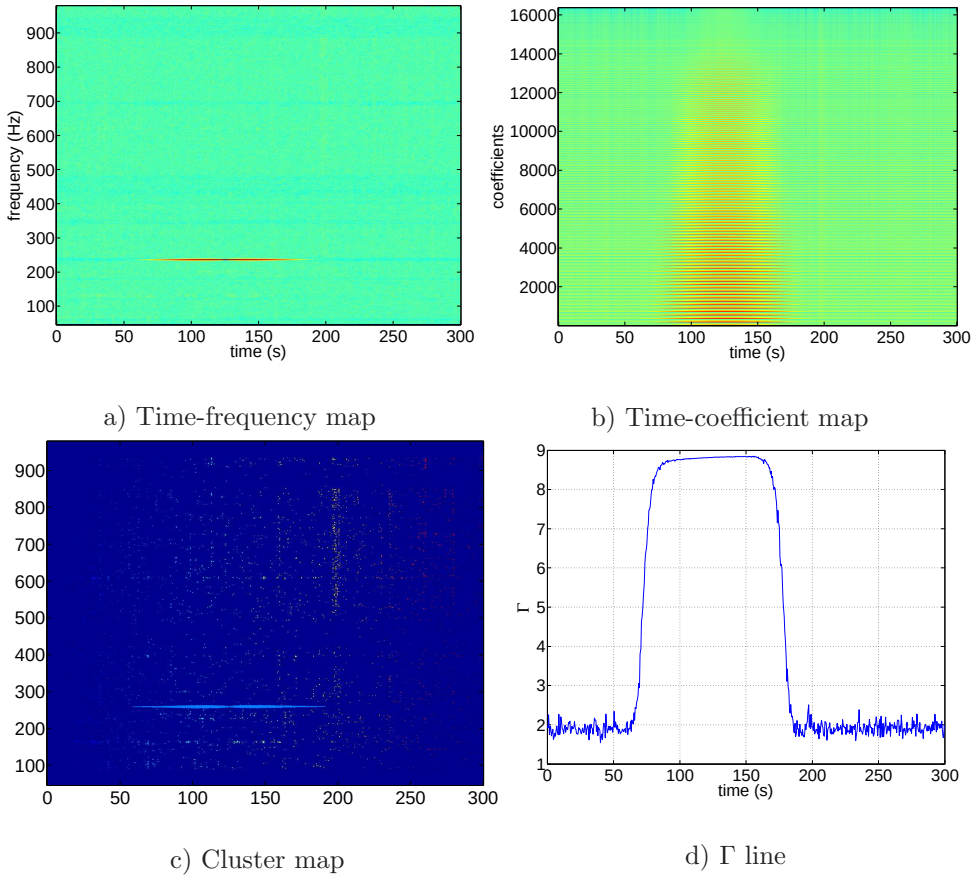

a) Time-frequency map

b) Time-coefficient map

c) Cluster map

d) Γ line

FIGURE 5.4: The plots in this figure demonstrate how a long duration SGC with $h_{\rm rss}$ of 1.48e-18 $Hz^{-1/2}$, a central frequency of 235 Hz, a duration of 96 seconds, and peak time of 126 seconds, appears in a ; a) time-frequency map, as a horizontal line centred at τ_0 and f_c, b) time-spherical harmonic coefficient map, as a broad segmented shape with power in the coefficients across the duration of the injection, c) clustermap, as a long thin horizontal line centred on τ_0, d) plot of Γ against time, as a steeped bell shape centred on τ_0.

5.2.4 Circularly polarised Sine-Gaussians (SGC) - long duration

The plots in Figure 5.4 present the output from the pipeline when analysing a long duration signal. S6A data containing a circularly polarised sine-Gaussian injected with a $h_{\rm rss}$ of 1.48e-18 $Hz^{-1/2}$, f_c of 853 Hz and duration of 96 seconds. The injection is clearly visible in each representation, but it is the large bell shape present in both the time-spherical harmonic coefficient and Γ plot (Figure 5.4(b) &(d)), which is most noticeable. Note that the reason this injection had such a high $h_{\rm rss}$ amplitude, and a large amount of data was analysed (300 seconds), was an attempt to overcome the whitening problem which "notches" out features that are monotonic in nature (see the discussion in Section E.1.4.3 on the long SGC waveforms).

5.2.5 Poly-oscillating waves (POW)

This is another ad-hoc (i.e. non-physical) waveform intended to test the long duration recovery performance of the pipeline. Designed as a complex long duration waveform, the POW waveform is based on a chirping signal, but rather than merely drifting up or down in frequency, it oscillates around the central frequency for the duration of the waveform. The waveform

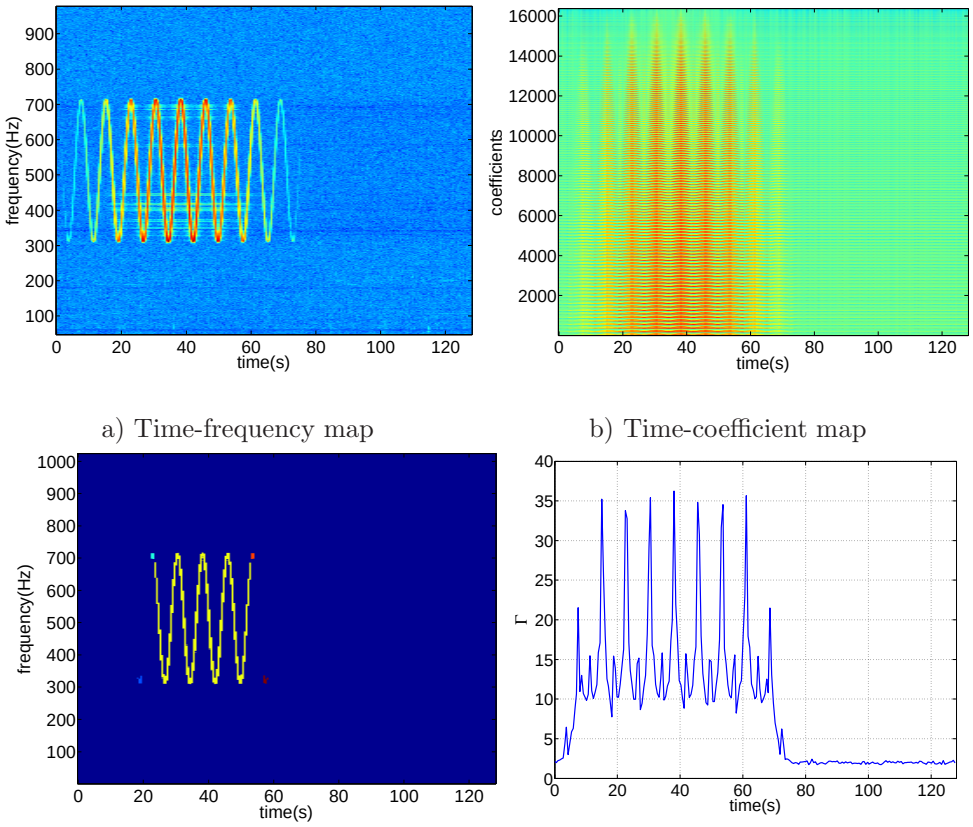

a) Time-frequency map

b) Time-coefficient map

c) Cluster map

d) Γ line

FIGURE 5.5: The plots in this figure demonstrate how a long duration poly-oscillating wave with an $h_{\rm rss}$ of 2e-19 $Hz^{-1/2}$, a central frequency of 512 Hz, a bandwidth of 400 Hz, a duration of 90 seconds, and peak time of 40 seconds, appears in a ; a) time-frequency map, as an oscillation around f_c for the duration of the signal, b) time-spherical harmonic coefficient map, as a crown shape, with power in most of the coefficients for the duration of the injection, c) clustermap, as a sine wave on the plane, picking out the strongest part of the signal d) plot of Γ against time, displaying the crown effect very strongly.

is defined by the following series of parameters: its root-sum square signal amplitude, $h_{\rm rss}$, central frequency, f_c, chirp number, Υ, oscillation parameters, Ξ, and the duration, τ. The waveform is generated using

$$h = h_{\rm rss} * \frac{1}{(2\pi\tau^2)^{1/4}} * \exp(2\imath\pi t f_c + \Upsilon + \Xi). \tag{5.13}$$

where

$$\Upsilon = -(1 - \imath\alpha) * t^2./(4\tau^2) + \imath\delta, \tag{5.14}$$

$$\Xi = \imath(\Delta f / f_Q) * \sin(2\pi t f_Q). \tag{5.15}$$

$$\tag{5.16}$$

Here, we have encapsulated all of the parameters associated the chirping behaviour into Υ. The chirp parameter, α, is used to cause the waveform to chirp up or down in frequency over its duration, while δ is used to alter the phase at the peak of the envelope. Similarly, the parameters for controlling the waveform's oscillations have been packaged as Ξ. The bandwidth of the signal in the time-frequency plane is given by Δf, while f_Q corresponds to the number of waveform cycles performed on the time-frequency map).

Figure 5.5 demonstrates the response of the pipeline to the poly-oscillating wave signal. Note that the injection appears shorter in duration due to the Gaussian envelope that decreases the amplitude towards the extremities of the wave.

5.2.6 Inspiral waveform (INS)

The inspiral waveform has its basis in the waveform generated by a compact object binary, widely considered to be the most likely source for detection in aLIGO and the advanced detectors. The inspiral waveforms are defined by a set of parameters, the mass of the first, m_1, and second, m_2, components of the binary system, the inclination angle of the system, ι, the distance to the system, d, and the coalescence time, t_c. The waveform used here was a 2-PN-order waveform. Figure 5.6 contains the graphical output of the pipeline for a 1.4-1.4

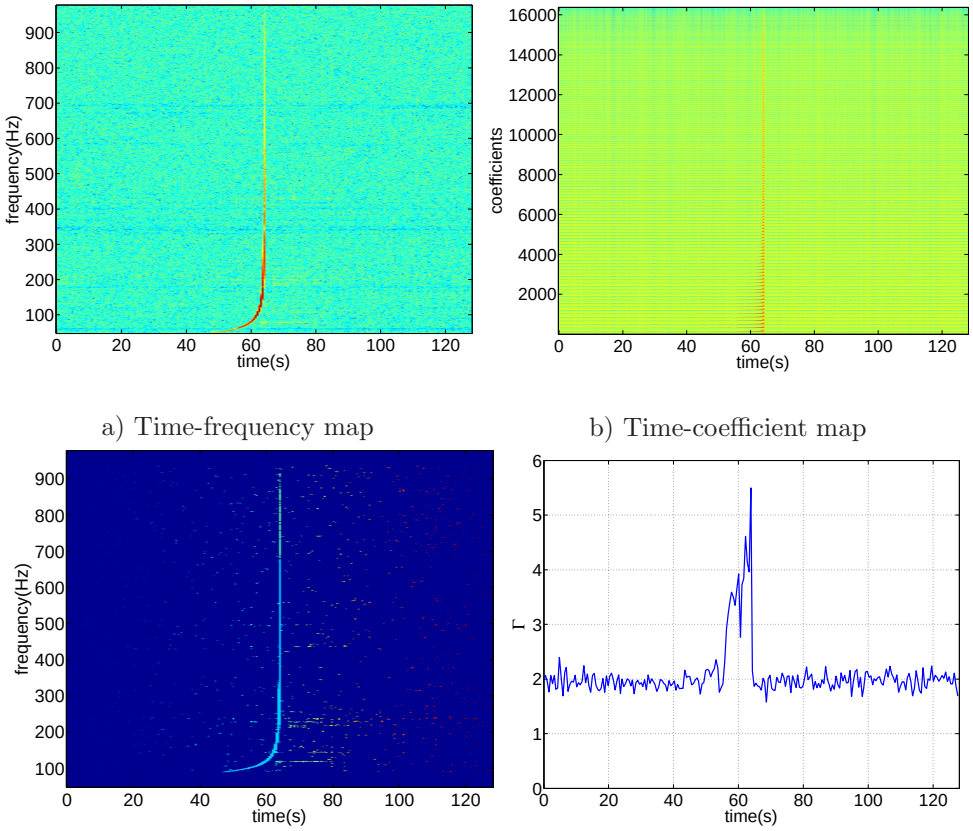

a) Time-frequency map

b) Time-coefficient map

c) Cluster map

d) Γ line

FIGURE 5.6: The plots in this figure demonstrate how a 1.4-1.4 $M_\odot$ inspiral waveform at 4 Mpc appears in a ; a) time-frequency map, with the familiar j shape, b) time-spherical harmonic coefficient map, with most of the power at the inspiral time, c) clustermap, clearly clustered together , d) plot of Γ against time, displaying the growth in energy as the frequency increases.

$M_\odot$ Inspiral at 4 Mpc.

5.2.7 Accretion disk instabilities, (ADI)

This waveform represents the emission from the precessing torus in the proposed model for long GRB emission by van Putten [79]. The gravitational emission from the torus changes with respect to the timescale of the black hole spin, in fact, the dominant frequency decreases by $\sim$

10% over the first 50% of the energy emission, so the frequency drops quickly then stabilises, see Section 2.2.3.2) for more details.

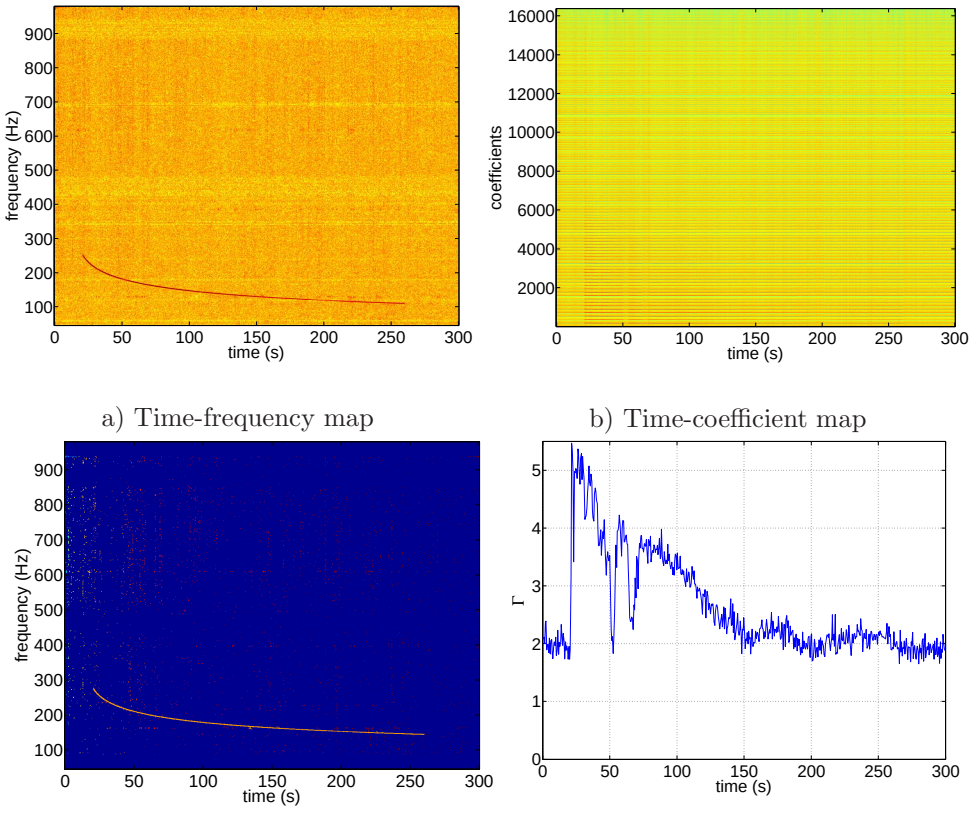

a) Time-frequency map

b) Time-coefficient map

c) Cluster map

d) Γ line

FIGURE 5.7: Plots demonstrating how the ADI signal is represented in the pipeline: a) time-frequency map, shows how this waveform is structured like a sidewards inspiral, starting at a known frequency (250 Hz) and slowly chirping down in frequency to 100 Hz over 220s, b) time-spherical harmonic coefficient map, not that useful due to the non-localisation of the signal, c) clustermap, all pixels of interest have been clustered, d) plot of Γ against time, shows the distribution of power over the duration of the signal.

5.3 Analysis procedure

Before we move onto the actual results, it is worth summarising the steps that were required to generate them. The analysis an be broken down into three main areas: background estimation, injections runs and zero-lag analysis.

5.3.1 Zero-lag analysis

The main output from the SPHRAD pipeline is a list of the loudest triggers in the detector data. For the "zero-lag" run (a single pass through the data without time-shifts or injections, but using the same analysis parameters) these triggers are our candidate list for a detection.

These are compared against the distribution of background events (from analysis of the time-shifted data) to determine how significant they are. We are typically looking for candidates with very low false alarm rates (FAR) of the order of once in a hundred years. The lower the FAR, the greater our confidence that we have a real detection.

All of the triple-coincident data (i.e. when H1, L1 and V1 were in science mode and taking data) for S6 was analysed with the pipeline. The only limitations on the data was the use of CAT2 vetoes to generate the data segments. These flagged periods when the data from a detector showed anomalous behaviour, correlated with known noise sources (e.g. earthquakes, planes etc, see ([78]) for more information). As stated previously, the live time was approximately 10 days for S6a, 16 days for S6b and 27 days for S6d, giving a total live time of approximately 53 days.

5.3.2 Background estimation

Before we are able to discuss the significance of any of the triggers produced by the zero-lag analysis, we need to be able compare the trigger statistics against triggers that we know have been produced from data that does not contain a signal. To ensure that any real gravitational wave signal in the data cannot produce a trigger that is correlated between detectors, we time-shift the data streams by amounts much larger than the time-of-flight between detectors before rerunning the analysis. Glitches and background anomalies will couple across the data streams generating large spikes of incoherent energy, which can be used to assess the significance of our events.

For each of the S6 periods, 135 time-shifts at multiples of 3 seconds were carried out. This gave us a total background livetime of 130979580 seconds for Period A, 123745850 seconds for Period B and 245383628 seconds for period D.

5.3.3 Injection runs

We can assess the performance of the pipeline by injecting simulated gravitational wave signals into the data and then analysing the results. The signals are injected with a range of parameters (strength, sky location, central frequency and bandwidth, duration and type) so that different aspects of the analysis can be assessed. The proportion of injections recovered for each strength is used to determine the strength that a real gravitational wave signal would need to have to be seen by the pipeline.

The injections are separated into two main categories, the short duration signals (usually less than a few second long: SGL, SGC, WNB, see Table 5.1 for the parameters used) and the long duration (greater than 10 seconds: SGC, ADI, POW, inspirals) Note that the ADI

Type	Centre freq (Hz)	Bandwidth (Hz)	Duration (s)	Q
WNB	300	50	0.1	-
WNB	750	250	0.1	-
SGL	235	27	0.003	3
SGL	235	9	0.009	9
SGL	853	100	0.0008	3
SGL	853	40	0.002	9
SGC	235	9	0.009	9
SGC	853	40	0.002	9

TABLE 5.1: Parameters for the short duration simulated signals (WNB, SGL and SGL), base amplitude $1e^{-21}$

Type	Centre freq (Hz)	Bandwidth (Hz)	Duration (s)	Cycles
POW	235	100	12	4
POW	512	50	24	1
POW	852	200	48	2

TABLE 5.2: Parameters used for the poly-oscillating wave, base amplitude $5e^{-20}$

Type	Mass 1 ($M_\odot$)	Mass 2 ($M_\odot$)	Inclination angle (radians)	Distance (Mpc)
Inspiral	1.4	1.4	0.7854	10
Inspiral	1.4	1.4	1.5708	10

TABLE 5.3: Parameters used for the inspiral waveforms, base amplitude $5e^{-20}$

Type	Centre freq (Hz)	Duration (s)
SGC	235	24
SGC	512	96
SGC	853	48

TABLE 5.4: Parameters used for the long duration SGC waveforms, base amplitude $5e^{-20}$

waveform is precomputed and has no free parameters. These signal types give a good mixture of frequencies, bandwidths, durations and time-frequency shapes to test the pipelines response at the lower end of its detectability and to also allows with comparison to other pipelines (such as coherent WaveBurst and X-Pipeline).

We performed runs for each of the previously mentioned simulated signal types being isotropic in sky direction, uniform over time (approximately 1 injection performed every 128 seconds), while the amplitude was varied in discrete steps from $(0.037 - 148.6) \times 1e^{-21} Hz^{-1/2}$ for the short duration signals, and $(0.037 - 148.6) \times 5e^{-20} Hz^{-1/2}$ for the long duration signals.

5.3.4 Tuning the analysis

Once the background estimation has been generated we can use it to determine how loud an event must be to stand out among this background. We can threshold on these values and use them to to 'tune' the analysis by setting the threshold high enough that only a few background triggers pass. We can fix some False Alarm Rate (FAR), i.e. one trigger of this strength in 100 years (if we have sufficient background), and set the thresholds so we achieve this. The recovered injections can be compared with this threshold and culled, demonstrating the sensitivity of the analysis.

From Sections 4.5.1 and 4.6.1 it is clear that the most useful statistics for distinguishing injections from background are the null energy and spherical harmonic coefficients. The use of the incoherent null energy against the coherent null energy for background rejection is well known (see X-Pipeline [49]), but the use of the spherical harmonic coefficients for background rejection is new to this work (see ([67]) by the author).

For each epoch of S6 (i.e. A, B and D) the background produced by time-shifting the triple coincidence data is randomly split into two sets. Using triggers from the first set, a 10^{th} order polynomial is fit to the distribution of background triggers in the $I_{\rm sh}$ vs $E_{\rm sh}$ and $I_{\rm null}$ vs $E_{\rm null}$ planes.

These polynomials are used to define cuts on their respectively planes, with triggers on the "wrong" side of the cut being rejected. These fixed cuts are applied to the background triggers in the second set, the zero-lag and injection runs to remove what we have defined as background. Whichever zero-lag triggers pass this cut have the potential to be a detection (and in the case of the injection runs, "recovered" simulated signals).

5.3.5 Defining upper limits

Event are ranked by the statistic Γ (derived from the spherical harmonic coefficients, see equation (4.42). By examining the distribution of the background triggers from the second set of triggers that pass these cuts, we can construct a rate vs Γ histogram. This can be used this to set a threshold on Γ, Γ_0, such that we have a False Alarm Rate (FAR) of $\approx 0.25{\rm yr}^{-1}$ (chosen to match the FAR used in the S6 coherent WaveBurst analysis (see [78]), to allow for comparison). We can then define the efficiency of the pipeline for each injection type as the fraction of injections at a given injection amplitude which survive the cut and have $\Gamma > \Gamma_0$. This procedure is carried out for both the null and spherical harmonic energies for comparison.

5.3.6 Detections

Before we can "open the box" on the zero-lag analysis, we must decide which cuts (null cut, spherical harmonic cut, or both) we wish to apply. If any triggers survive the cuts imposed on the zero-lag analysis, we can compute the FAR of the loudest using the background histogram. If any have a sufficiently low FAR we could consider it a candidate detection. For example, to have a "3σ" significance requires a false alarm rate probability of approximately 2.7×10^{-3}, which corresponds to a FAR of $0.019yr^{-1}$ in the 53 days of data. The following sections show the results for each epoch, S6A, B and D separately.

5.4 Science run S6A/VSR2

The "A" period began at GPS 931035615 (Tuesday July 7, 2009 at 21:00 UTC) and ended at GPS 935798415 (Tuesday Sept 1,2009 at 00:00 UTC). The total amount of background generated for S6A is 130979580 seconds, which works out at over 4 years of analysed data, and produced 8847615 triggers. As these are split into two sets for the purposes of tuning the analysis, it leaves us 4423802 background triggers - potentially meaning we can get down to a false alarm rate of $0.25yr^{-1}$. As stated previously, the analysis requires tuning, so we split the triggers into two, train on half, then apply the cut to the other half.

5.4.1 Background distribution of the null energy

Figure 5.8 shows how the null energy is distributed across the coherent/incoherent plane. It has the expected shape (see [49]), with triggers tending to the diagonal as they get louder. They pack together well and form a clean edge which the cuts track efficiently - there are only a handful of triggers that pass the cut, most of which lie away from the area that signals are expected to fall in. The cut is generated by slicing the E_{null} axis into small segments, calculating a median value of I_{null} for these points then fitting a 10^{th} order polynomial through them. One proviso is that a high valued"vanishing" point is added so that the the cut moves towards the diagonal (a feature of the E_{null}/I_{null} plot). Note that this cut is "tuned" on the 1^{st} trigger set and applied to the 2^{nd}. Figure 5.8 shows the result.

The blue line in Figure 5.9 is a curve demonstrating the number of triggers with a Γ value greater than that shown on the x-axis. For instance, all of the triggers will have a Γ value greater than zero. The red curve is the same background but after the cut from Figure 5.8 has been applied. The background has been suppressed significantly and from it we can read off the Γ value for the loudest surviving background trigger to use in calculating the upper limit.

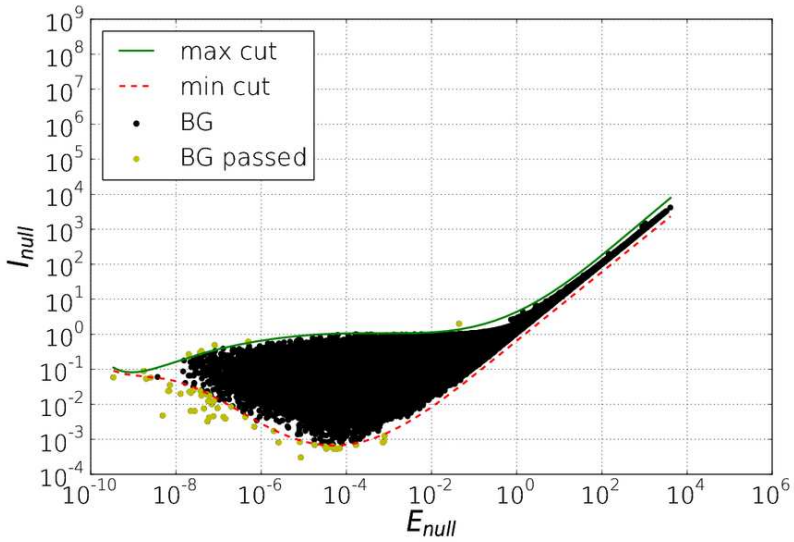

FIGURE 5.8: Scatter plot showing the distribution of $E_{\text{null}}/I_{\text{null}}$ for the 2^{nd} set of background triggers and the corresponding null energy cuts.

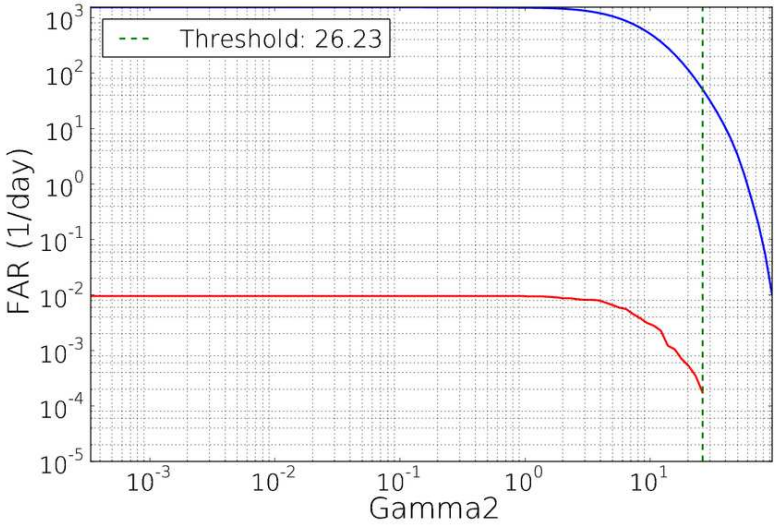

FIGURE 5.9: Plot of cumulative rate vs Γ for tuning and null cut background. Blue: before the null cut has been applied, Red: after null cut.

5.4.2 Background distribution of the energy in the spherical harmonic coefficients

Similarly, Figure 5.10 shows how the "coherent" and "incoherent" parts of the spherical harmonic coefficients are distribution on the plane. The triggers form a compact object which allows the tuning to produce a closely fitting track with very few surviving triggers. Figure 5.11 is the rate vs Γ plot similar to Figure 5.9 but the cut applied to the background is that generated by the tuning based on the spherical harmonic coefficients. As can be seen, a lower Γ threshold is produced.

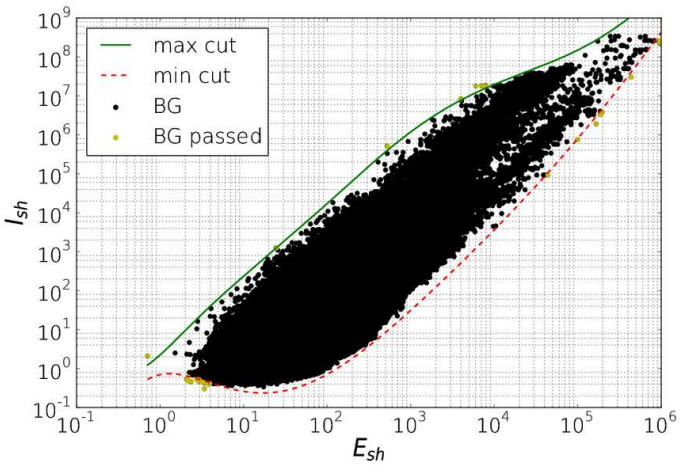

FIGURE 5.10: Scatter plot showing the E/I distribution for the 2^{nd} set background triggers and the corresponding cut for the spherical harmonic coefficients energy

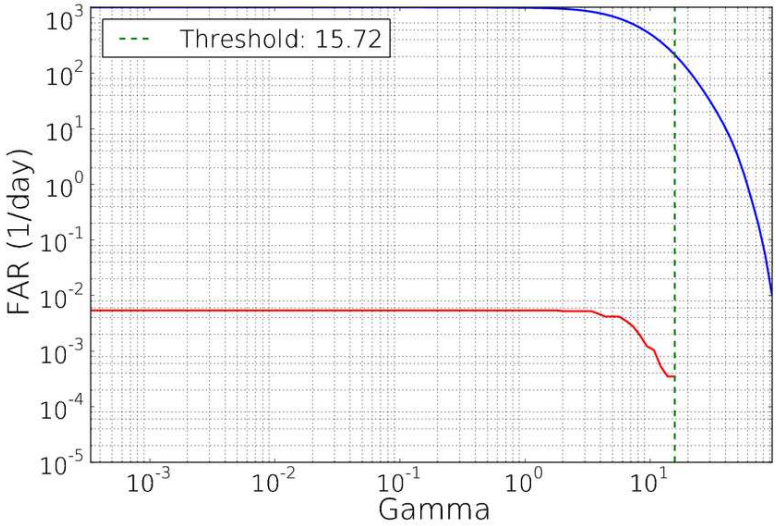

FIGURE 5.11: Plot of cumulative rate vs Γ for tuning and spherical harmonic cut background. Blue: before the spherical harmonic cut has been applied, Red: after spherical harmonic cut.

5.4.3 Short duration simulations

In the following sections, we will present figures demonstrating how triggers from the different injection types are distributed in the null and spherical harmonic plane (see Table 5.1 for the parameters used). Note that the paired scatter plots share a common legend showing the amplitude of the injected signals (e^{-21}), while the efficiency plots have their own legend giving the value of the 50% and 90% upper limits.

5.4.3.1 Behaviour of white noise bursts

The following figures demonstrate the distribution of triggers of white noise burst injections in both the null and spherical harmonic plane.

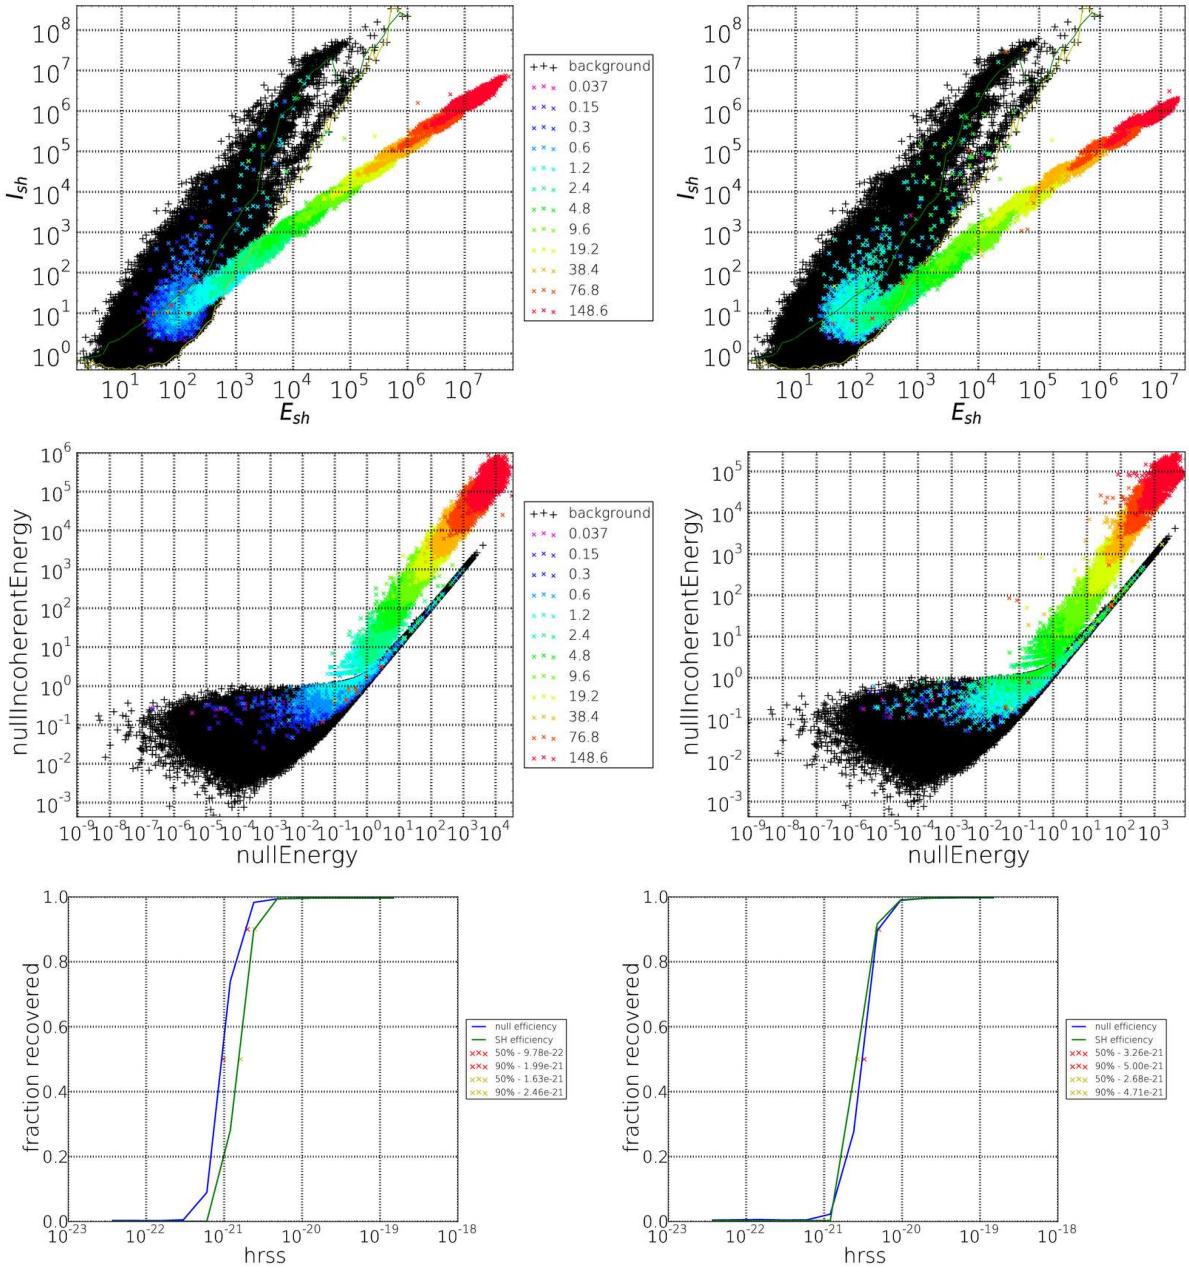

FIGURE 5.12: The plots in this figure show the distribution of triggers in the null and spherical harmonic plane for S6A, and the corresponding efficient curves generated by applying the appropriate cut to the triggers. The top pair of figures correspond to spherical harmonic energy, while the centre pair show the same for null energy. The left most figures relate to WNBs with f_c 300 Hz and Δf=50 Hz, while the right most figures show WNBs with a f_c of 750 Hz and Δf of 250 Hz. The bottom pair of plots show the efficiency curves for both the null (in blue) and spherical harmonic (in green) cuts.

The plots displayed in Figure 5.12 demonstrates the good separation between the loud injection triggers (plotted in red) and the background triggers (plotted in black). From the efficiency curves we can see how the null cut does better for lower frequency injections, while the spherical harmonic statistic recovers higher frequency injections better.

5.4.3.2 Behaviour of linearly polarised sine-Gaussians

The following figures demonstrate the distribution of triggers of SGL injections in both the null and spherical harmonic plane. Once again, Figure 5.13 demonstrates the good separation

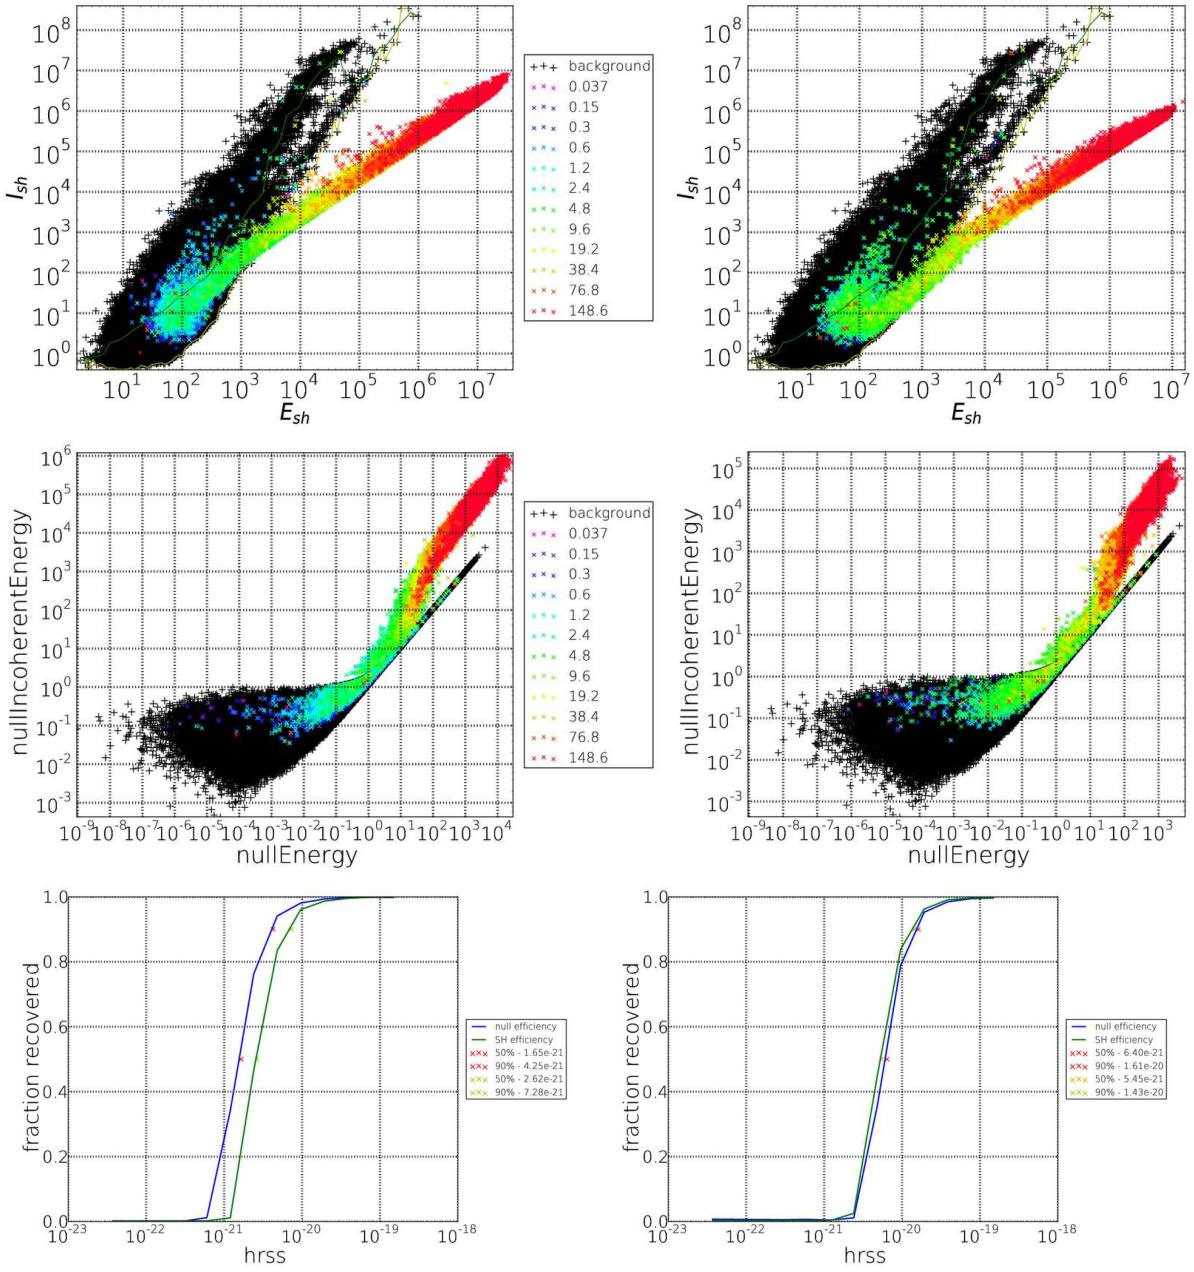

FIGURE 5.13: The plots in this figure show the distribution of triggers in the null and spherical harmonic plane for S6A, and the corresponding efficient curves generated by applying the appropriate cut to the triggers. The top pair of figures correspond to spherical harmonic energy, while the centre pair show the same for null energy. The left most figures relate to SGLs with f_c 235 Hz and Q=3, while the right most figures show SGL with f_c = 853 Hz and Q=3. The bottom pair of plots show the efficiency curves for both the null (in blue) and spherical harmonic (in green) cuts.

between loud injection triggers (plotted in red) and background triggers (plotted in black).

The efficiency curves show that the null cut does a better job at lower frequencies, while the spherical harmonic cut is slightly better at higher frequencies. By comparing Figures 5.13 and

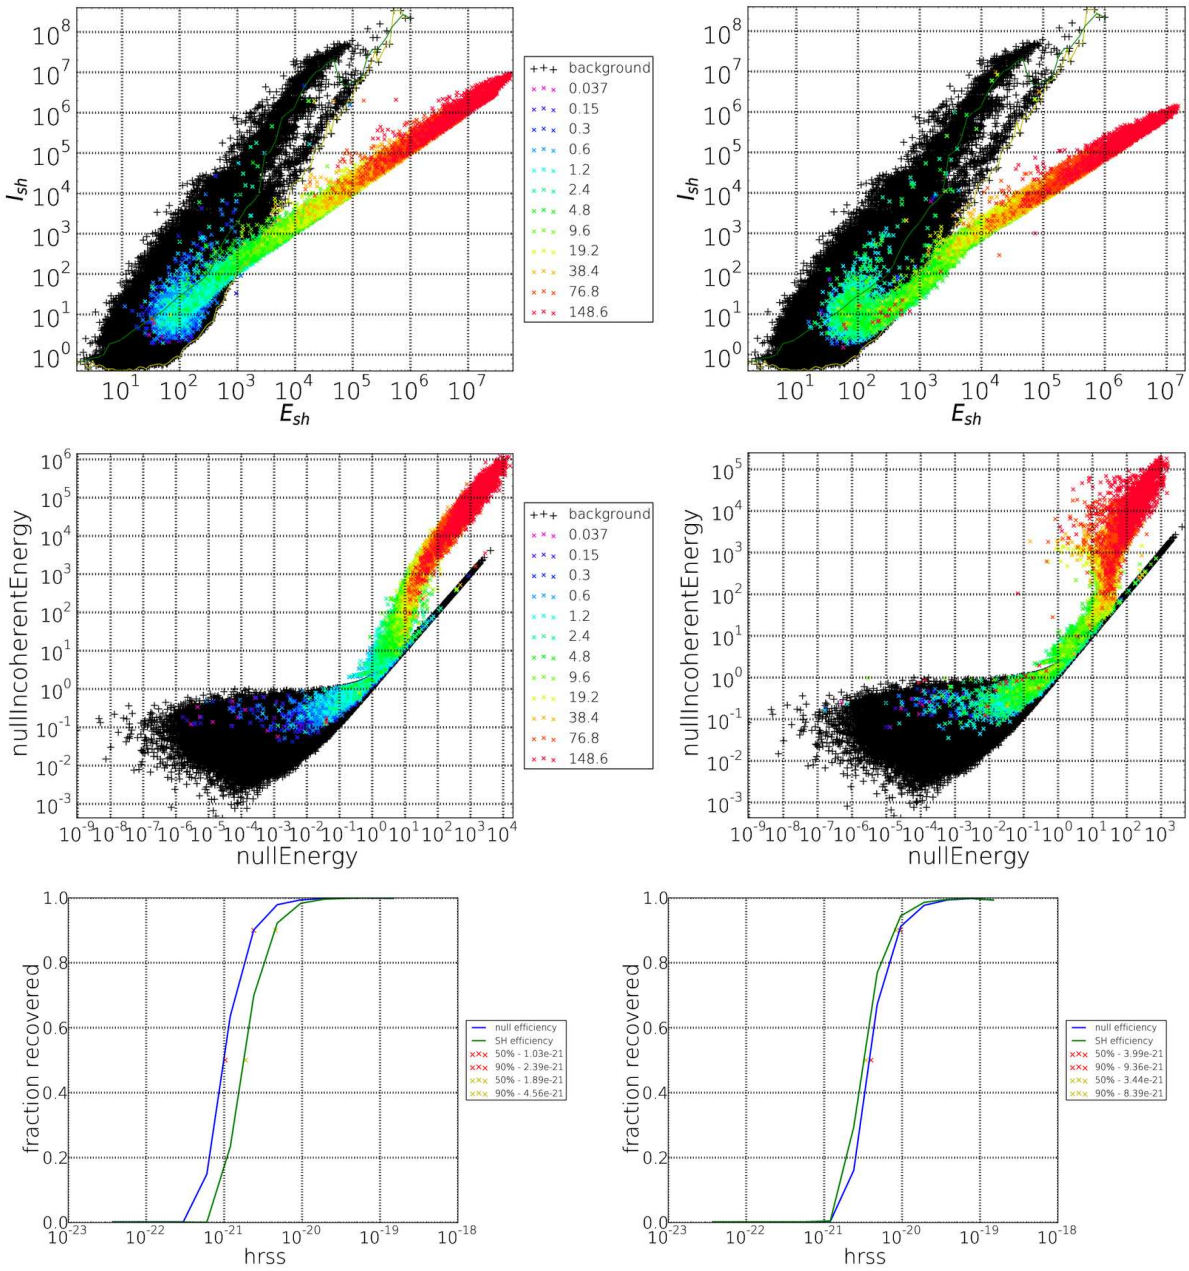

FIGURE 5.14: This figure shows the distribution of triggers in the null and spherical harmonic plane for S6A, and the corresponding efficient curves generated by applying the appropriate cut to the triggers. The top pair of figures correspond to spherical harmonic energy, while the centre pair show the same for null energy. The left most figures relate to SGLs with f_c 235 Hz and Q=9, while the right most figures show SGL with $f_c = 853$ Hz and Q=9. The bottom pair of plots show the efficiency curves for both the null (in blue) and spherical harmonic (in green) cuts.

5.14, we can see how the length of the injected waveform affects the efficiency of the SPHRAD pipeline. The injections are identical except for their Q value (which is related to duration; a

higher Q for the same frequency means a longer signal), and it is clear to see that the injections with the higher Q value have better upper limits, see Table 5.5.

Type	Centre freq (Hz)	Null cut ($Hz^{-1/2}$)		SH cut ($Hz^{-1/2}$)	
		50 % UL	90 % UL	50 % UL	90 % UL
SGLQ3	235	$1.65e^{-21}$	$4.25e^{-21}$	$2.62e^{-21}$	$7.28e^{-21}$
SGLQ9	235	$1.03e^{-21}$	$2.39e^{-21}$	$1.89e^{-21}$	$4.56e^{-21}$
SGLQ3	853	$6.40e^{-21}$	$1.61e^{-20}$	$5.45e^{-21}$	$1.43e^{-20}$
SGLQ9	853	$3.99e^{-21}$	$9.36e^{-21}$	$3.44e^{-21}$	$8.39e^{-21}$

TABLE 5.5: Table showing the consistently lower values for the peak amplitude UL for longer duration signals (SGL, Q=3 compared with SGL, Q=9).

5.4.3.3 Behaviour of circularly polarised sine-Gaussians

The following figures demonstrate the distribution of triggers for injected circularly polarised sine-Gaussians in both the null and spherical harmonic plane. As would be expected, Figure

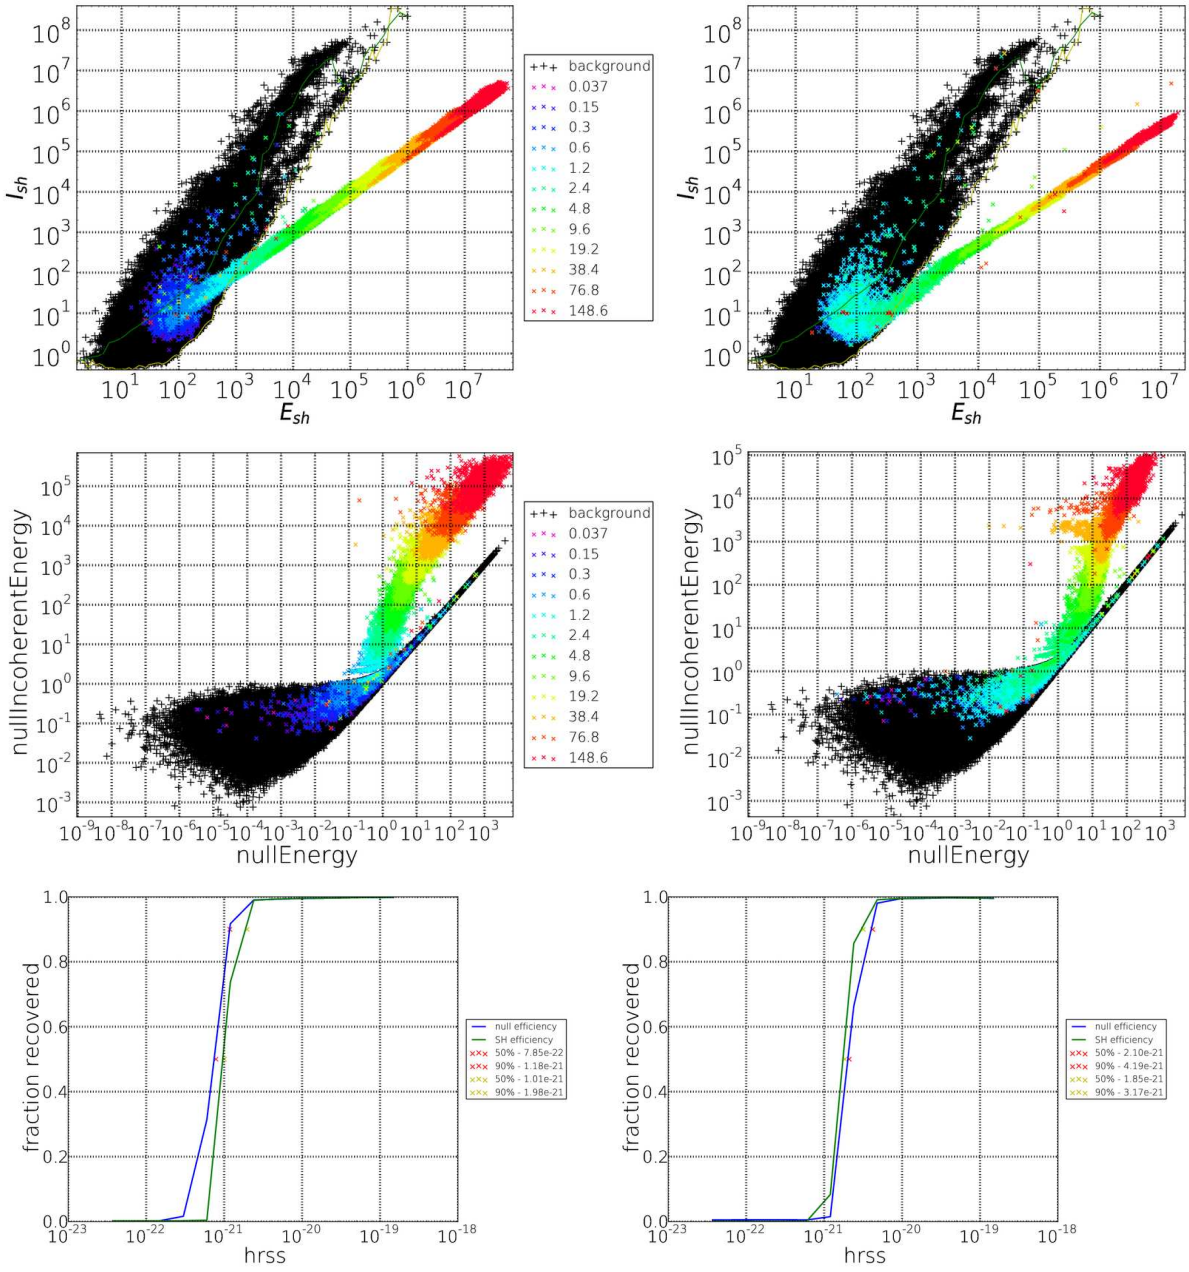

FIGURE 5.15: The plots of this figure show the distribution of triggers in the null and spherical harmonic plane for S6A, and the corresponding efficient curves generated by applying the appropriate cut to the triggers. The top pair of figures correspond to spherical harmonic energy, while the centre pair show the same for null energy. The left most figures relate to SGCs with f_c 235 Hz and Q=9, while the right most figures show SGCs with f_c = 853 Hz and Q=9. The bottom pair of plots show the efficiency curves for both the null (in blue) and spherical harmonic (in green) cuts.

5.15 demonstrates that the circularly polarised SGs are recovered best out the SGs. The

addition of power in the second polarization means that even if the signal is badly oriented with respect to the detectors, it can still be detectable.

5.4.4 Long duration simulations

In the following sections, we will present figures demonstrating how triggers from the different injection types are distributed in the null and spherical harmonic plane. See Tables 5.2, 5.3 and 5.4 for the parameters used for the POW, inspiral waveforms and the SGCs. Note that the paired scatter plots share a common legend while the efficiency plots have their own legend giving the value of the 50% and 90% efficiency.

5.4.4.1 Behaviour of poly-oscillating waveforms

The following figures demonstrate the distribution of triggers for injected POW in both the null and spherical harmonic plane. Regarding Figure 5.16, it demonstrates how well the SphRad pipeline can recover complicated long duration waveforms. There is a large separation between the loud injection triggers (plotted in red) and the background triggers (plotted in black). Once again the efficiency curves show how the null cut does better for lower frequency injections, while the spherical harmonic cut just edges it for higher frequency injections.

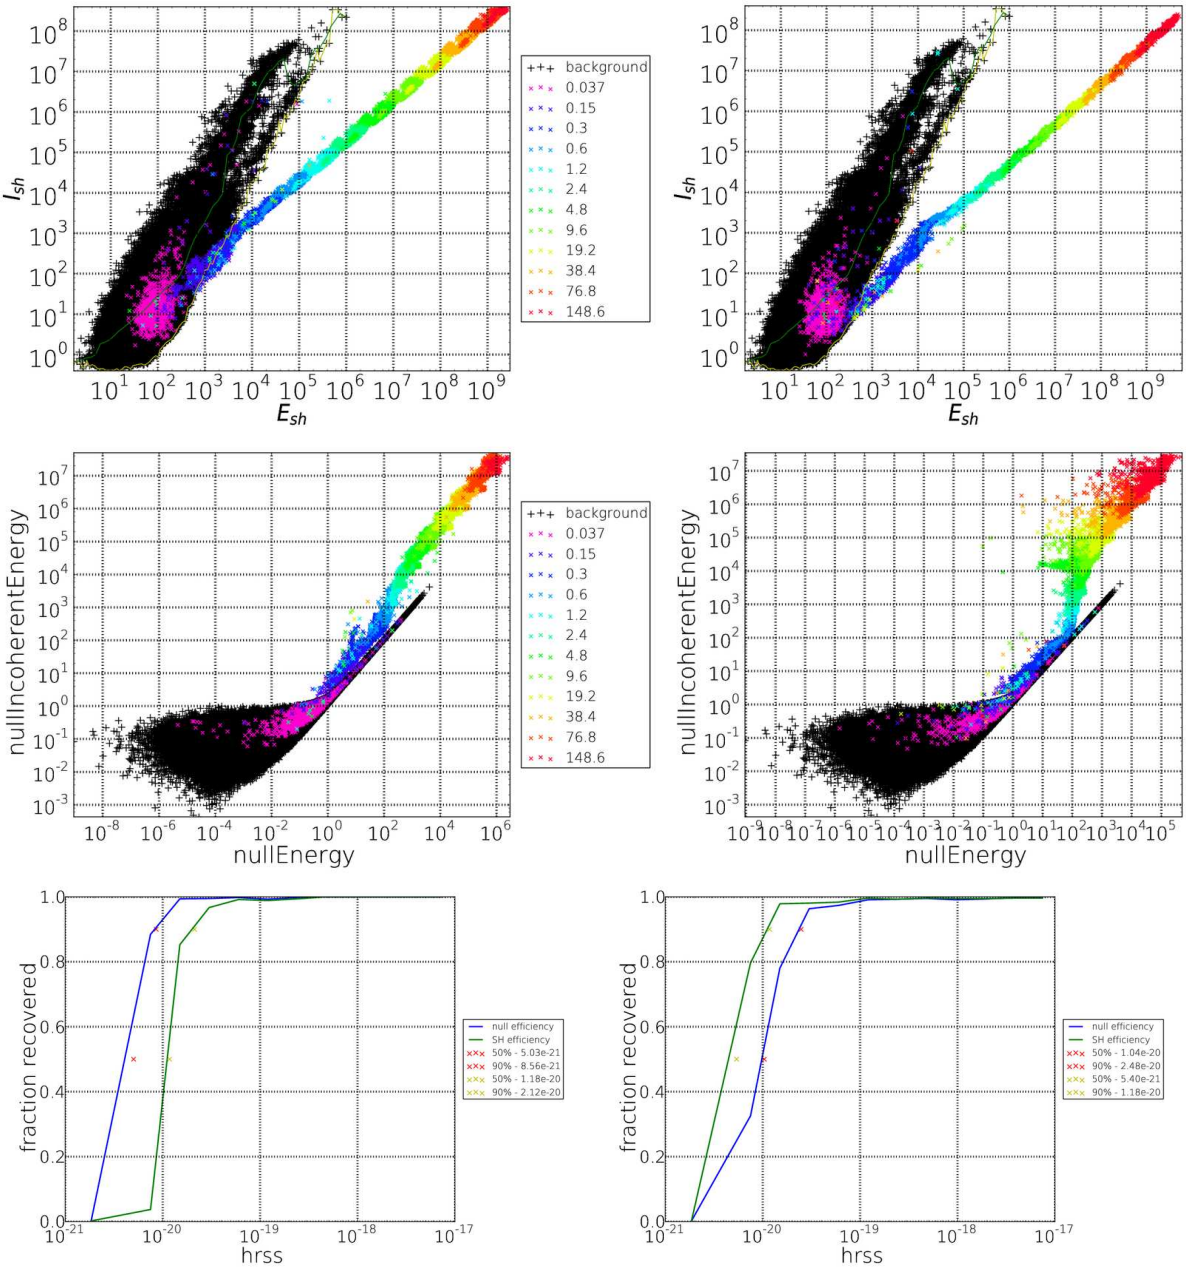

FIGURE 5.16: The plots of this figure show the distribution of triggers in the null and spherical harmonic plane for S6A, and the corresponding efficient curves generated by applying the appropriate cut to the triggers. The top figures correspond to spherical harmonic energy, while the centre show the same for null energy. The left most figures relate to POWs with $f_c = 235$ Hz, $\Delta f = 100$ Hz, and $\Delta t = 12$ seconds, while the right most figures show POWs with $f_c = 852$, $\Delta f = 50$ Hz, and $\Delta t = 24$ seconds. The bottom plots show the efficiency curves for both the null (in blue) and spherical harmonic (in green) cuts.

5.4.4.2 Behaviour of Inspiral simulations

The following figures demonstrate the distribution of triggers for injected inspiral waveforms in both the null and spherical harmonic plane. Regarding Figure 5.17, it shows that SPHRAD struggles to recover inspirals as well as some other waveforms. The separation between the

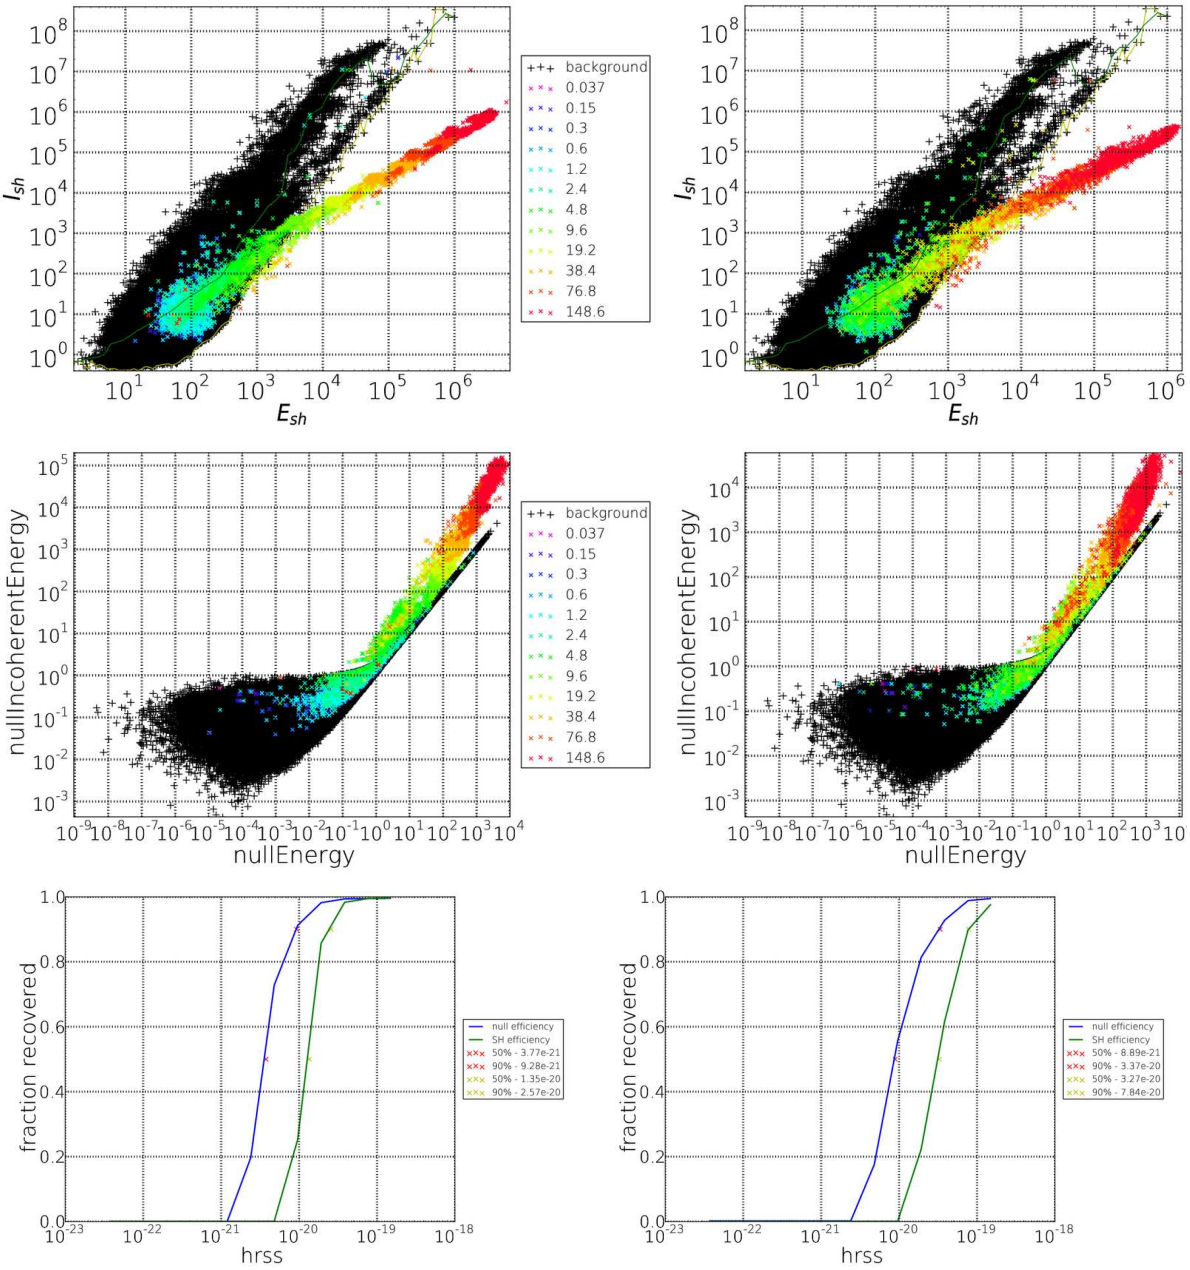

FIGURE 5.17: The plots of this figure show the distribution of triggers in the null and spherical harmonic plane for S6A, and the corresponding efficient curves generated by applying the appropriate cut to the triggers. The top pair of figures correspond to spherical harmonic energy, while the centre pair show the same for null energy. The left most figures relate to inspirals with an inclination angle of $\pi/4$, while the right most figures show inspirals with an inclination angle of $\pi/2$ (corresponds to edge on). The bottom pair of plots show the efficiency curves for both the null (in blue) and spherical harmonic (in green) cuts.

loud injection triggers (plotted in red) and the background triggers (plotted in black) is not as large as some of the other waveforms, and the null cut is consistently better. There does not appear to be an obvious reason for this, and so it warrants further investigation.

5.4.4.3 Behaviour of long duration circularly polarised sine-Gaussians

The following figures demonstrate the distribution of triggers for long duration circularly polarised sine-Gaussians in both the null and spherical harmonic plane. In most cases, the plots

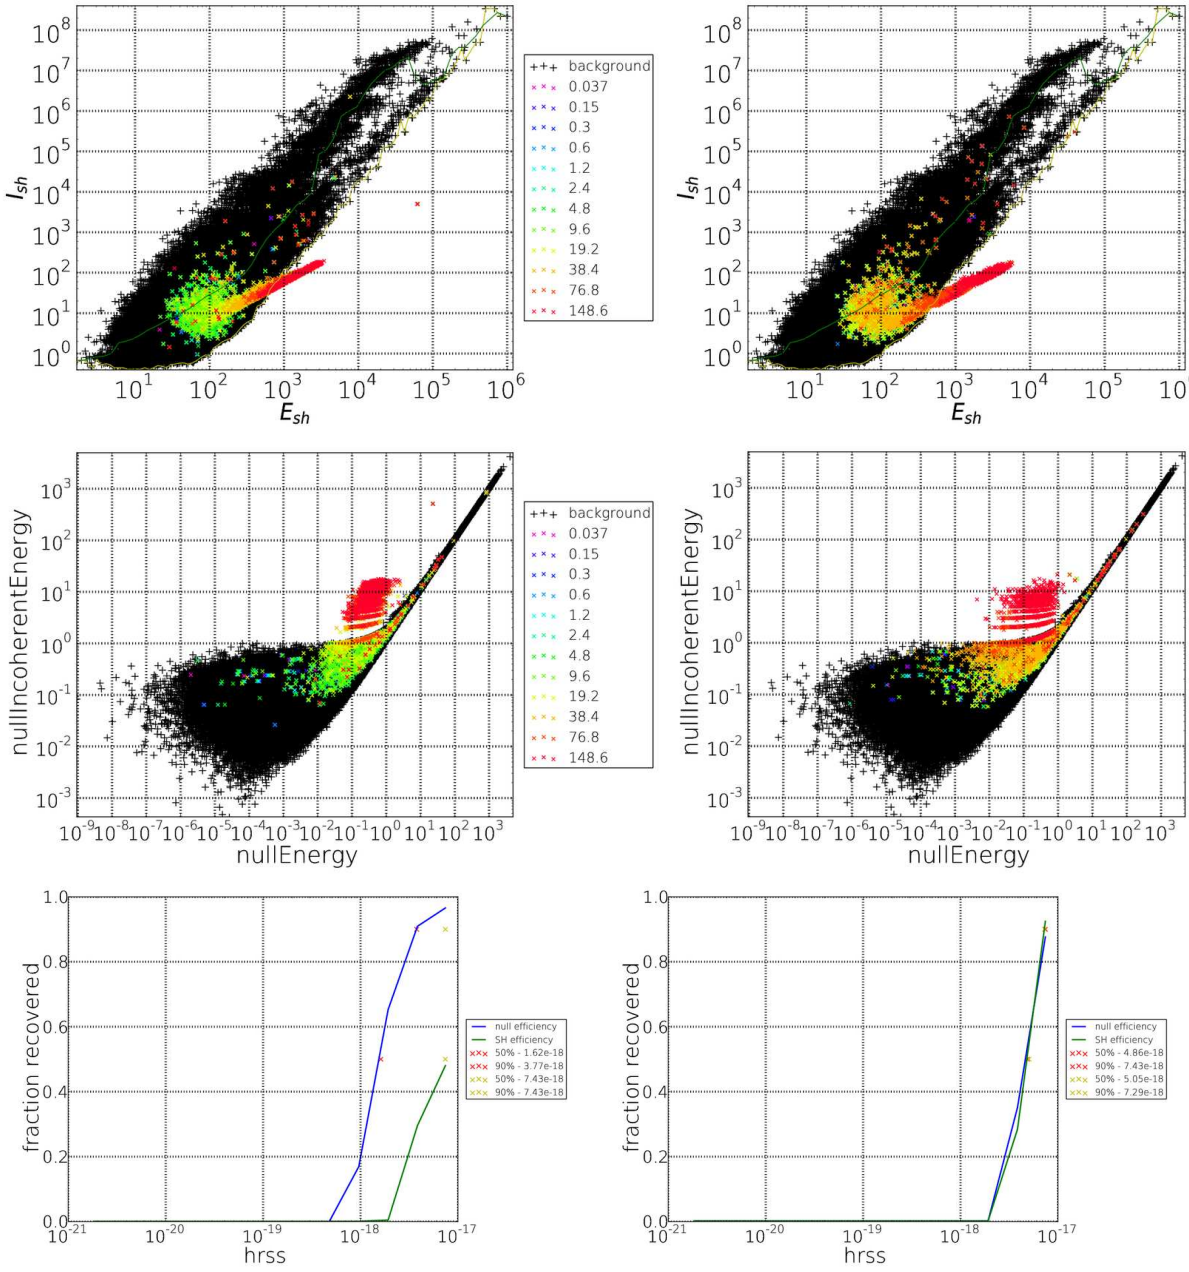

FIGURE 5.18: The plots of this figure show the distribution of triggers in the null and spherical harmonic plane for S6A, and the corresponding efficient curves generated by applying the appropriate cut to the triggers. The top figures correspond to spherical harmonic energy, while the centre show the same for null energy. The left most figures relate to SGCs with f_c 235 Hz and Δt=24 seconds, while the figures on the right hand side show SGCs with f_c 512 Hz and Δt=96 seconds. The bottom plots show the efficiency curves for both the null (in blue) and spherical harmonic (in green) cuts.

have shown that the pipeline is capable of recovering suitably loud injections with reasonable

ease, however, Figure 5.18 shows the distribution of triggers associated with the long duration SGCs. These range in duration from 24-96 seconds (see Table 5.4), and all but the most energetic of signals disappears into the background.

There is a simple reason for this, and it relates to the way that the data is whitened. For each analysis length (in this case 128s), the data is segmented, medianed, meaned and averaged (as defined in the FINDCHIRP paper, ([56])). Therefore, long signals that are monotonic in frequency would have power present in most of these segments and would cause a bias in the estimation of the noise floor, pushing it up for those frequencies in which the SGC has power. For short duration signals, or those that wander in frequency, this does not cause any significant problems, however, our SGCs contain most of their power in a single frequency.

This means that there is excess power at *all* times for this analysis length, so the injected signal can never rise above the background level. This can be clearly seen by comparing any of the plots in Figure 5.18 with those in Figure 5.15.

There are a few modifications we can make to the SphRad pipeline to overcome this issue:

1. The simplest solution is to use a longer stretch of data for the noise estimation. Once it is large compared to the length of the signal the noise estimation will no longer be biased.
   Pros: little modification to the pipeline is required, fast, noise estimation is still "local",
   Cons: signal could be extremely long and there is a limit to how much data we could process at once.

2. Estimate the noise floor using the analysis lengths immediately before and after the one we are processing.
   Pros: As above,
   Cons: As above, but the code changes will be more significant.

3. Maintain a "running estimate" of the noise floor,
   Pros: Should produce a robust, bias free estimation of the noise,
   Cons: Major pipeline changes, non-local noise estimation, at the start of the running estimate we are still faced with the same problem.

4. Use a completely new method of estimating the noise,
   Pros: Could produce a potentially better, more robust noise estimator which could be tailed to the exact problem,
   Cons: Need to find this new method of noise estimation, unknown issues.

There are clearly many ways to overcome this problem, but which solution is best depends on the effect that the changes to the whitening will have on the data and the analysis.

5.4.4.4 Behaviour of accretion-disk instabilities

The following figures demonstrate the distribution of triggers for very long duration ($\approx$ 220 seconds) accretion disk instability waveforms in both the null and spherical harmonic plane. Note that the ADI waveform is precomputed and so we only have control over the distance to the source. The waveform is still injected with a range of injection factors however, so we can generate the efficiency plots. We can compare and contrast the recovery of the long duration

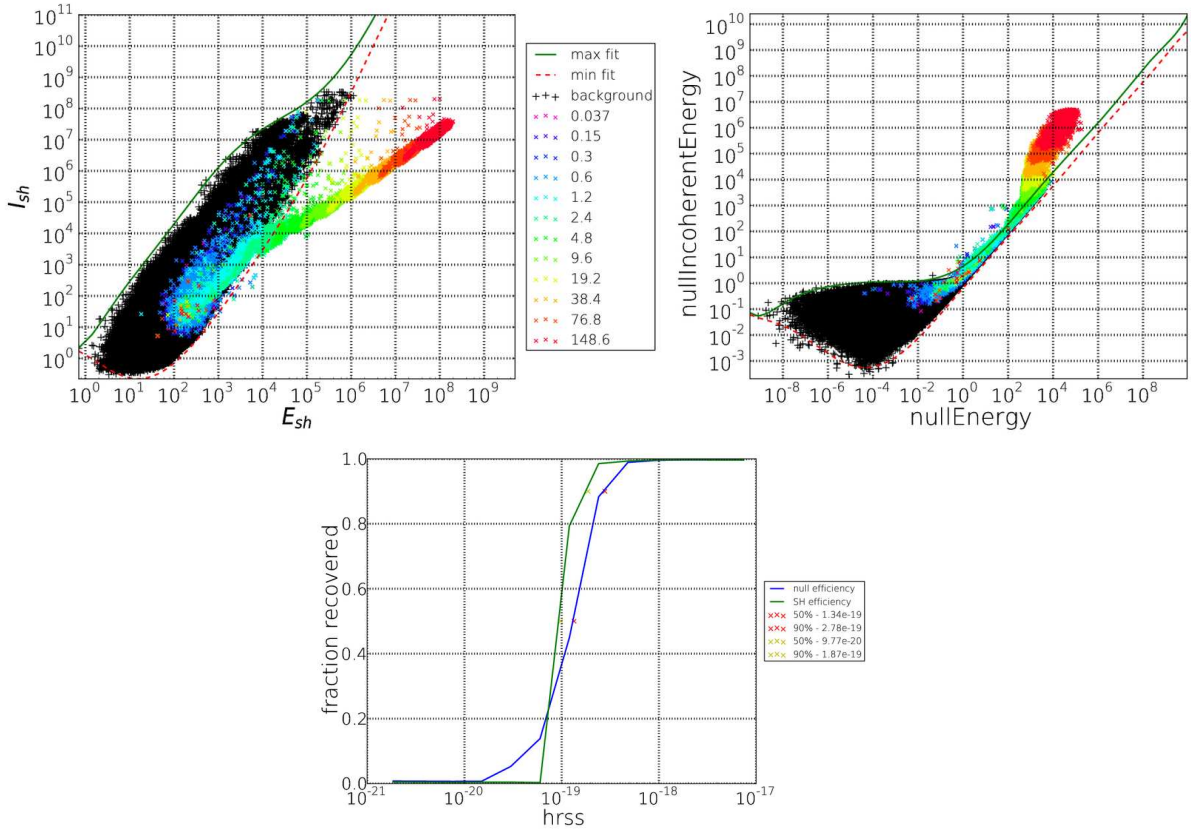

FIGURE 5.19: The plots of this figure show the distribution of triggers in the null and spherical harmonic plane for S6A, and the corresponding efficient curves generated by applying the appropriate cut to the triggers. The top left figure shows the distribution for spherical harmonic energy, while the top right shows the same for null energy. The bottom plot show the efficiency curves for both the null (in blue) and spherical harmonic (in green) cuts.

sine-gaussians directly with the ADI results. They are both very long duration signals, but the change in frequency of the ADI waveform circumvents the problem with the whitening procedure, leaving the signal at nearly the injected power. This means that we recover more of the injections. The spread between the injection and background triggers is not huge, but it is sufficient to distinguish the two.

Note that the pipeline was also used to process S6B (see Appendix D) and S6D (see Appendix E). The plots and results are very similar to those already shown for S6A so they have been placed in the Appendix. One of the primary differences between the analysis runs is that

some of the efficiency curves "turn" over for high $h_{\rm rss}$ (e.g. see Figures D.5 D.6 for S6B, and Figure E.5 for S6D). This could be due to many issues (pipeline changes between S6A and S6B/S6D, better or worse behaved noise, saturation due to the high power of the loud injections), but it appears to be a post-processing error. In some of the scatter plots, red crosses (i.e. loud injections) are present in the background bulk. This could be due to the post-processing assigning the wrong trigger to that injection, why it only falls at large signal power is something that needs to be investigated.

5.5 Remarks and comparison to cWB

In this section we will comment on some of the features seen in the results plots, and compare our efficiencies to those of coherent WaveBurst. Tables 5.6 and 5.7 collate all of the information on the injections, including how well the null and spherical harmonic cuts performed. We are using the 90% confidence level to determine the winner for each type, and how much it won by.

Waveform Type	Central freq (Hz)	Bandwidth (Hz)	Q	Duration (s)	Signal ($\times 10^{-21}$) Null	Spherical	Winner 90 %	Margin of victory
WNB	300	50	-	0.1	1.99	2.46	NULL	27 %
WNB	750	50	-	0.1	5.0	4.72	SHC	6 %
SGL	235	27	3	0.003	4.25	7.68	NULL	70 %
SGL	235	9	9	0.009	2.39	4.56	NULL	90 %
SGL	853	100	3	0.0008	16.1	14.3	SHC	13 %
SGL	853	40	9	0.002	9.36	8.39	SHC	11 %
SGC	235	9	9	0.009	1.18	1.98	NULL	68 %
SGC	853	40	9	0.002	4.19	3.17	SHC	32 %

TABLE 5.6: Table showing the comparison between the .

Waveform Type	Central freq (Hz)	Bandwidth (Hz)	Duration (s)	Signal ($\times 10^{-21}$) null	spherical	Winner	Margin of victory
POW	235	100	12	8.56	21.2	NULL	248 %
POW	852	50	24	24.8	11.8	SHC	111 %
INS	-	-	16	9.28	25.7	NULL	177 %
INS	-	-	16	33.7	78.4	NULL	233 %
SGC	235	6.40	24	3770	54300	NULL	1340 %
SGC	512	3.99	96	7430	7290	SHC	2 %
ADI	100	200	230	278	187	SHC	49 %

TABLE 5.7: Table showing the comparison between the 50% upper limit of coherent WaveBurst and SPHRAD. Once the "correction" factor due to the FFT length mismatch has been taken into account, it would appear that SPHRAD has the potential to perform as well as coherent WaveBurst.

It is quite clear from Table 5.7 that when the Null cut is better than the Spherical harmonic coefficient cut, it performs significantly better, whereas the Spherical harmonic cut just edges

ahead. There seems to be an even split of wins between the statistics, although it is clear that the null cut is generally the better cut for low frequency waveforms, whereas the spherical harmonic cut is marginally better for high frequency waveforms. There are two things to take into account here: 1) that the null statistic has been used for many years, and has been rigorously tuned and optimised. In comparison the spherical harmonic statistic is very new and has only been used on a very ad hoc basis, 2) the spherical harmonic statistic is all-sky, while the performance of the null statistic depends on the sky position being reconstructed correctly. Note that the fact that our new statistic works almost as well as the null statistic is very promising (especially when you consider that the sky position used to generate the null energy is calculated using the decomposition into spherical harmonics), but more work is required to turn it into a useful tool.

To fully understand how well the SphRad pipeline performs, we need to compare it to a mature and tested pipeline, such as coherent WaveBurst. This pipeline was run over the S6 data and has combined results for the H1L1V1 network, as shown in [78]. Note that coherent WaveBurst and [78] focussed on short duration (typically of the order of a few milliseconds in length) gravitational wave bursts, while the SphRad pipeline has been designed from the ground up to analyse signals with a duration $\geq$1 second.

The main reason that the SphRad pipeline is better suited to analyse longer duration signals is due to the application of time-shifts for sky position in the Fourier domain (as laid out in Section 4.2.2) which corrupts $\approx$ 20 ms at each end of the FFT. This means that we are restricted to FFT lengths $\gg$ than 20 ms so we can't use the FFT length that is optimal for the signals tested in [78].

However we can estimate the approximate loss of sensitivity from using non-optimal FFT lengths as we know that this spreads the signal power over more time-frequency pixels. From Anderson, ([51]), we expect the sensitivity loss to be $\propto N^{1/4}$ where N is the number of time-frequency pixels.

For a SG, we have

$$\tau = Q/(\sqrt{2}\pi f_c), \tag{5.17}$$

which, for a Q9 SG, becomes $\tau \sim 2/f_0$, and $\tau \sim 0.68/f_0$ for a Q3 SG. Our optimum FFT length should be double that, so

$$\text{Q3 :FFT}_{\text{len}} \sim \frac{1.35}{f_0}.\text{Q9 :FFT}_{\text{len}} \sim \frac{4}{f_0}. \tag{5.18}$$

which, for a 235 Hz SG gives us 17 ms. Our 1 second FFT is a factor of $1/0.017 \approx 59$ greater, which gives us a loss of about $59^{1/4} \approx 2.8$. It's a lot worse for a 853 Hz SG, which requires an FFT length of $\sim$ 5ms. A 1 second FFT is a factor of $1/0.005 \approx 213$ greater, which gives a loss of $213^{1/4} \approx 3.8$. A WNB with a $\tau \sim$ 100ms, requires an optimum FFT length of 200ms,

so a 1s FFT is a factor of ~ 5 greater. This gives us a loss in sensitivity of about 1.5. The same calculation produce a loss in sensitivity of about 3.63 for 235Hz, Q3 SG and 5 for SGL 853 Hz Q3.

Injection Type	Freq (Hz)	SphRad 50% $\times 10^{-21}$	Loss factor	Corrected SphRad 50% $\times 10^{-21}$	cWB 50% $\times 10^{-21}$
SGL Q3	235	1.65	3.6	0.46	0.57
SGL Q9	235	1.03	2.8	0.37	0.52
SGL Q3	849	6.40	5.0	1.28	1.51
SGL Q9	849	3.99	3.8	1.05	1.41
SGC Q9	235	0.79	2.8	0.28	0.49
SGC Q9	849	2.10	3.8	0.55	1.33
WNB	250	0.98	1.5	0.65	0.75

Table 5.8: Table showing the comparison between the 50% upper limit of coherent WaveBurst and SphRad. Once the "correction" factor due to the FFT length mismatch has been taken into account, it would appear that SphRad has the potential to perform as well as coherent WaveBurst.

Table 5.8 shows the 50% upper limits for both coherent WaveBurst and SphRad. Generally, SphRad is around a factor of two greater than the figures shown for cWB. However, once the "correction factor" due to the longer FFTs has been taken into account the numbers are startling. Obviously, this is only an estimated efficiency, but it shows that, even at this early stage, SphRad has great potential. It may be worth running both pipeline over the same waveform sets to get a direct comparison.

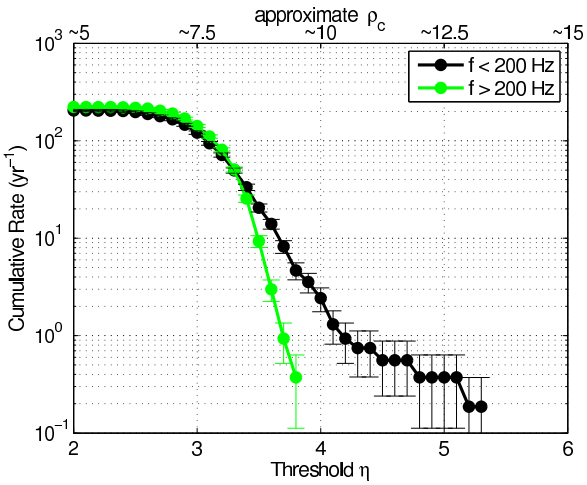

Figure 5.20: The coherent WaveBurst background for S6D and triple coincident data, H1L1V1. Taken from [2]

Figure 5.20 shows the cWB background for H1L1V1 in S6D, taken from [2]. Note that the low frequency ($f < 200$ Hz) has a tail at ~ 1 yr^{-1} level which prevented a high sensitivity/low FAR to be assigned to the "Big Dog" event, see S6 results paper [78]. So far there are no tails present in the background distributions generated by SphRad for S6A (see Figure 5.11),

S6B (see D.4) and S6D (see E.4). However, we only go to 1/year, so we cannot say with any certainty that there are no tails, but the initial results are very promising. More background needs to be generated so a more informed opinion can be made.

5.5.1 Conclusion

While these results are promising, note that they are only preliminary. Ideally, we would test both pipelines on common long-duration waveforms of the order of a few seconds. But, it does look like the performance of SphRad is "consistent" with cWB, which, due the fact that our pipeline has been tailor to find a different class of signal than have been considered before, and to do it all-sky and all-time, is fairly amazing.

This definitely motivates further study and continued development of the Spherical Radiometer pipeline.

Chapter 6

Conclusion

After recognising that there was a clear omission in the search parameter space, notably, signals with durations of between 1 to hundreds of seconds long, the next step was to determine if there were any potential sources to motivate this search. Microquasars, long gamma-ray burst and neutron star oscillations all have the potential to generate a signal detectable by the advanced LIGO and Virgo detectors.

There is a well-developed suite of burst analysis techniques that are able to detect gravitational waves without detailed prior knowledge of the waveform. Nothing appears to prevent these standard procedures from being applied in the analysis of intermediate duration signals. Further, the special case of intermediate-duration signals also offer opportunities for the application of new methods, in particular spherical radiometry. The advantage of this is that it is a computationally efficient way to compute likelihoods over the sky. Also, the decomposition of the data into spherical harmonics has given us a new detection statistic and background rejection test.

The detection statistic, Γ, (see Equation 4.42) is a product of the spherical harmonic coefficients used to decompose the detection problem. It is quick to generate and competitive with a mature statistic used by an existing pipeline that can't possibly run all-sky.

The raw coefficients themselves are also useful as a background rejection test, separating background, glitches and injections (a simple plot of $E_{\rm sh}$ against $I_{\rm sh}$), but in a Γ-time plot these elements are truly separated, above and below the average of the background noise. So, we have a new pipeline and a new analysis process; a new clustering algorithm too, chaining pixels together as it crawls across the time-frequency map. And, to top it off, we get an all-sky search for free.

This new pipeline was applied to analyse triple coincident LIGO-Virgo S6/VSR2,3 data. We were able to process all of the triple coincident data for S6A, S6B and S6D. We generated

background triggers using time-shifts and from their distribution determined cuts on the incoherent/coherent null energy and spherical harmonic coefficients that would suppress loud background triggers without falsely rejecting too many simulated gravitational wave signals. More importantly, we have demonstrated the potential of the pipeline to detect signals of amplitudes comparable (after considering the differences in FFT length) to current standard all-sky burst search pipelines, and our background appears to lack the tails often seen with other pipelines.

The Spherical Radiometer pipeline has already shown enormous potential. Enough to motivate continuing development of the pipeline. Now that the major development work has been carried out, efforts can be made into understanding how to make the most out of its uniqueness.

6.1 Future work

The analysis as presented in the this thesis is preliminary, and the SphRad pipeline could be enhanced in a number of ways now that time has been devoted into getting the foundation solid.

It would not take much effort to retool SphRad to process data in real-time. The Γ statistic can be very quickly generated (order of a couple of seconds for 128 seconds of data), so it could generate background (by time-shifting the incoming data) as it processed the zero-lag data. The sample frequency and Fourier Transform length of the analysis are user definable, so can be selected based on the events to be analysed. At the very least, it could be used to catalog and characterise large incoherent glitches.

While the pipeline has been optimised as much as possible, more performance could be gained in two simple ways:

1) Remove the dependencies on the LAL libraries

2) Remove the dependencies on MATLAB

While it was useful to make use of an existing code base, the improvements in memory usage and speed may be worth the effort in severing SphRad links to LAL and MATLAB. There are only a very small number of LAL functions used by SphRad, which could easily be written into the pipeline, similarly, any effort made to redo the data loading and conditioning (which is currently carried out in MATLAB) can only improve the speed. This also means that the pipeline could easily be ported.

While some rudimentary work has been carried out on configuring the pipeline to analyse long duration signals, more work could be done on tuning the pipeline to optimally detect them.

There are probably more discoveries to be uncovered by studying the representation of the detection problem in spherical harmonic coefficients and the Γ statistic could be harnessed in better ways.

As we demonstrated in the results chapters, specifically Figure 5.18, long single frequency waveforms are whitened out out the data unless they are unrealistically loud. Therefore a new whitening procedure is required that can whiten the data while leaving the injections untouched. One simple method could to carry out the whitening procedure on the blocks of data directly before and after the block under analysis. This gives a good estimation of the background, without being too influenced by injections in the analysis block. Another method would be to construct a "running" estimation of the background so that the segments containing the waveform represent a smaller proportion of the overall segments analysised for the background estimation.

The Γ statistic is such an excellent discriminator of signals and glitches, it could be used in combination with a fast incoherent search technique to instantly reject glitchy triggers.

For the current analysis SphRad has been run using a single length FFT, thus producing a fixed size time-frequency map. However, there is the potential to use the pipeline to carry out a fast multi-resolution analysis using very low order values of the spherical harmonic coefficients. More work is required to bring the code up to standard, but it would be able to cluster between time-frequency planes and select the most appropriate FFT length for the signal of interest.

There is a theoretical restriction on the minimum FFT length that the SphRad pipeline can safely analysis. This has not been fully explored, and could open up more of the lower duration signal space to us once we understand how low we can go.

We current have a selection of five statistics which are used to determine the sky position, it could be advantageous to look into other methods. Specifically if there is a way to manipulate the spherical harmonic coefficients to improve the pointing.

In conclusion, the Spherical Radiometer pipeline is full of possibilities. Further study and continued development of the Spherical Radiometer pipeline could produce the first next-gen analysis pipeline.

Appendix A

Null stream analysis in double whitened data

Consider the energy in the plus polarization in the DPF for single whitened data:

$$E_+ \equiv \sum_{pixels} |\boldsymbol{e}^+ \cdot \boldsymbol{d}^{SE}|^2 \qquad \text{(plus energy)} \tag{A.1}$$

There is an implicit sum over the detectors,

$$E_+ \equiv \sum_{pixels} \sum_{i=1}^{D} |e_i^+ \cdot d_i^{SW}|^2 \qquad \text{(plus energy)} \tag{A.2}$$

Now, making use of the following definitions:

$$\boldsymbol{e}^+ \equiv \frac{\boldsymbol{f}^+}{|\boldsymbol{f}^+|},\ \boldsymbol{d}^{SW} = \boldsymbol{d}^{DW}\sqrt{\boldsymbol{S}} \tag{A.3}$$

E_+ becomes:

$$E_+ = \sum_{pixels} \sum_{i=1}^{D} \left| \frac{f_i^+}{|f_i^+|} d_i^{DW} \sqrt{S_i} \right|^2 \tag{A.4}$$

Using:

$$f_i^+ = \frac{F_i^+(\psi_{PDF})}{\sqrt{S_i}} \tag{A.5}$$

We get:

$$E_+ = \sum_{pixels} \sum_{i=1}^{D} \frac{1}{|\boldsymbol{f}^+|^2} \left| \frac{F_i^+(\psi_{PDF}) d_i^{DW} \sqrt{S_i}}{\sqrt{S_i}} \right|^2 \tag{A.6}$$

$$= \sum_{pixels} \frac{1}{|\boldsymbol{f}^+|^2} \sum_{i=1}^{D} \left| F_i^+(\psi_{PDF}) d_i^{DW} \right|^2 \tag{A.7}$$

Where:

$$F_i^+(\psi) = cos2\psi F_i^+(0) + sin2\psi F_i^\times(0) \tag{A.8}$$

$$F_i^\times(\psi) = sin2\psi F_i^+(0) + cos2\psi F_i^\times(0) \tag{A.9}$$

The same can be carried out for the other energies, which reveal:

$$E_\times = \sum_{pixels} \frac{1}{|\boldsymbol{f}^\times|^2} \sum_{i=1}^{D} \left| F_i^+(\psi_{PDF}) d_i^{DW} \right|^2 \tag{A.10}$$

$$E_{SL} = \sum_{pixels} (E_+ + E_\times) \tag{A.11}$$

$$E_{null} = E_{tot} - E_{SL} \tag{A.12}$$

Appendix B

Spherical coefficient ordering

Spherical harmonics are an expansion of a function onto the sphere, expressed as the coefficients in a series of spherical harmonics up to and including order $l = l_{\max}$. The ordering of the coefficients in memory is as follows (expressed as (l, m) pairs, read left to right and top to bottom:

m value of row	coefficients stored as (l, m) pairs
m=0	(0, 0), (1, 0), (2, 0), (3, 0), $\cdots$, ($l_{\max}$, 0)
m=1	(1,+1), (2,+1), (3,+1), $\cdots$, ($l_{\max}$,+1)
m=2	(2,+2), (3,+2), $\cdots$, ($l_{\max}$,+2)
m=3	(3,+3), $\cdots$, ($l_{\max}$,+3)
$\vdots$	$\vdots$
m=$l_{\max}$	($l_{\max}$,+$l_{\max}$)
m=-$l_{\max}$	($l_{\max}$,-$l_{\max}$)
$\vdots$	$\vdots$
m=-3	(3,-3), $\cdots$, ($l_{\max}$,-3)
m=-2	(2,-2), (3,-2), $\cdots$, ($l_{\max}$,-2)
m=-1	(1,-1), (2,-1), (3, -1), $\cdots$, ($l_{\max}$,-1)

TABLE B.1: Table showing the order of spherical harmonic coefficients, expressed as (l, m) pairs, read left to right and top to bottom.

The coefficients are arranged in order of increasing l for fixed m. This ordering allows the series to be efficient summed, as there are recursive algorithms for evaluating spherical harmonics over a range of m at fixed l. Also, since real-valued functions can be described by specifying only the $m \geq 0$ coefficients, they can be represented using only the first half of the coefficient array. If the function is also azimuthally symmetric then only the $m = 0$ coefficients are present.

This ordering is the same as used in the S2Kit [68].

Appendix C

Spherical harmonics

C.1 Overview

These are a frequency-space basis used to represent functions defined over the sphere. Representing a direction vector with $\hat{s}$ -using spherical parameterisation, so $\hat{s} = (x, y, x) = (\sin\theta\cos\phi, \sin\theta\sin\phi, \cos\theta)$. The real-value spherical harmonic basis functions are defined as:

$$Y_l^m(\theta,\phi) = \left\{ \begin{array}{l} k_l^m P_l^{|m|}(\cos\theta)\cos(|m|\phi) \text{ if m} \geq 0, \\ k_l^m P_l^{|m|}(\cos\theta)\cos(|m|\phi) \text{ if m} < 0 \end{array} \right\} \tag{C.1}$$

In spherical coordinates, this is represented as:

$$Y_l^m(\theta,\phi) = k_l^m P_l^m(\cos\theta)e^{im\theta} \tag{C.2}$$

where, in both, P_l^m are the associated Legendre polynomials, and k_l^m are normalisation constants, given by:

$$k_l^m = \sqrt{\frac{2l+1}{4\pi}\frac{(l-m)!}{(l+m)!}} \tag{C.3}$$

The functions are indexed using the *order* (or *band*), l, and the *degree,* m (each band corresponds to a polynomial of that degree - so zero is constant, 1 is linear etc, and there are $2l+1$ functions in a given band). The functions with $l = 0$ are called zonal harmonics, and have rotation symmetry around the z-axis, and zeros are contours on the sphere parallel to the x-y plane. The functions with $l = |m|$ are called sectorial harmonics and the zeros define apple slice-esque regions.

l must be a non-negative integer, while $-l \leq m \leq l$. It is useful to reformulate these indices into a single parameter $i = l(l+1) + m$.

Alternative: On the sphere, the modes of the Helmholtz equation are used to represent functions $Y_l^m = P_l^m(\theta)e^{im\lambda}$ where θ is colatitude and λ is east longitude.

C.2 Properties

Any real-valued function, $f(\hat{s})$ may be approximated as $\tilde{f}(\hat{s})$, using a linear combination of SH basis functions.

$$f(\hat{s}) \approx \tilde{f}(\hat{s}) = \sum_{l=0}^{n-1} \sum_{m=-l}^{l} Y_l^m(\hat{s}) f_l^m, \tag{C.4}$$

The coefficients f_l^m are generated by projecting f onto each of the basis functions, Y_l^m, using the inner product, $< f, Y_l^m >$:

$$f(l,m) =< f.Y_l^m >= k_l^m \int_0^{\pi} \left[\int_0^{2\pi} e^{-im\theta} f(\theta,\phi) d\phi \right] P_l^m(\cos\theta) \sin\theta d\theta \tag{C.5}$$

Spherical harmonics are orthonormal ($\int f_i f_j = \delta_{ij}$) and rotation invariant, and, according to the Nyquist sampling theorem, a band-limited function, f, of bandwidth, B, can be reconstructed from at least $2B$ samples without aliasing. Therefore the above calculation can be reduced to finite length weighted sums using a simple sampling scheme over S^2 in the longitudinal and latitudinal directions. Thus,

$$f(l,m) = \frac{\sqrt{2\pi}}{2B} \sum_{j=0}^{2B-1} \sum_{k=0}^{2B-1} a_j^{(B)} f(\theta_j,\phi_k) e^{-im\phi_k} P_l^m(\cos\theta_j) \tag{C.6}$$

where, $(\theta_j, \phi_k) = (\pi(2j+1)/4B, k\pi/B)$ are chosen as the sample points, and $a_j^{(B)}$ are weights analogous to the $\sin\theta$ factor in the integrals. Obviously, the sampling near the poles will be different to that near the equator, but this should not affect anything.

C.3 Zonal harmonics

These are the functions that have $m = 0$. If they are oriented so that the axis around which they have rotation symmetry is the z-axis, then zeros of these functions form lines of constant latitude, and the functions only depend on θ. Therefore, the coefficient vector only has 1

non-zero coefficient per band, so it has n rather then n^2 coefficients. To rotate, all that is required is to evaluate the SH basis functions in the new direction, d.:

$$f(s) = \sum_l z_l \sqrt{\frac{4\pi}{sl+1}} \sum_m y_l^m(d) y_l^m(s) \tag{C.7}$$

With the resulting coefficients as:

$$f_l^m = \sqrt{\frac{4\pi}{sl+1}} z_l y_l^m(d) \tag{C.8}$$

C.4 Rotations

Given a function $g(\hat{s})$ (which is a function $f(\hat{s})$ rotated by a rotation matrix, Q, so that $g(\hat{s}) = f(Q(\hat{s}))$, the projection of g is identical to rotating f and re-projecting it.

C.5 Integration

Due to the orthonormality of the SH basis, given any two SH functions, a and b, the integral of the product is simply the dot product of the coefficient vectors:

$$\int \tilde{a}(s)\tilde{b}(s) = \sum_{i=0}^{n^2} a_i b_i \tag{C.9}$$

C.6 Frequency filtering on the sphere

The spectrum is represented as a collection of complex coefficients, $\{\hat{P}_l^m | 0 \leq l < B. |m| \leq l\}$, which can be adjusted by multiplying by a 'filtering function', $\hat{h}(l,m)$:

$$\hat{P}\prime(l,m) = \hat{P}(l,m) * \hat{h}(l,m) \tag{C.10}$$

A *low-pass filter* (ignores anything over a threshold frequency) can be performed by :

$$\hat{h}(l,m) = \left\{ \begin{array}{c} 0 \text{ if } \sqrt{l^2+m^2} > K_l \\ 1 \text{ else} \end{array} \right\} \tag{C.11}$$

Stop-band filtering can be achieved by:

$$\hat{h}(l,m) = \left\{ \begin{array}{c} 0 \text{ if } K_a < \sqrt{l^2+m^2} < K_b \\ 1 \text{ else} \end{array} \right\} \tag{C.12}$$

Enhance certain frequency ranges by:

$$\hat{h}(l,m) = \left\{ \begin{array}{c} \kappa \text{ if } K_a < \sqrt{l^2+m^2} < K_b \\ 1 \text{ else} \end{array} \right\} \tag{C.13}$$

where $\kappa > 1$

For example, to improve the smoothness of a function, modify the ideal low-pass filter to perform a trapeziform low-pass filter (instead of a straight line cutoff, it drops down to zero over a certain range:

$$\hat{h}(l,m) = \left\{ \begin{array}{l} 0 \text{ if } \sqrt{l^2+m^2} > K_b \\ 1 \text{ if} \sqrt{l^2+m^2} < K_a \\ (K_b - \sqrt{l^2+m^2})/(K_b - K_a) else \end{array} \right\} \tag{C.14}$$

In this case, the filter drops to zero at K_b

C.7 Aliasing

There are only a finite number of e^{ikx} that be be distinguished on a set of points $x_n = n\frac{2\pi}{N}$, as:

$$e^{inkx\frac{2\pi}{N}} = e^{-imn2\pi} e^{in(k+mN)\frac{2\pi}{N}} = e^{in(k+mN)\frac{2\pi}{N}} \tag{C.15}$$

In other words, for any k there exists $-\frac{N}{2} < k_1 \leq \frac{N}{2}$ such that:

$$e^{ikx_n} = e^{ik_1 x_n} \tag{C.16}$$

This means that e^{ikx} and $e^{ik_1 x}$ cannot be distinguished on x_n - they are alternate characterisations (aliases) of each other. High frequency components will alias onto the lower frequency components.

For a grid with N latitudinal points, l must be less than or equal to N, and, as $m \leq n$ by definition, this is the discrete basis of harmonic functions. This means that an analysis

is invariant under any rotation or translation of the spherical coordinate system - the same harmonic representation is obtained no matter where the pole is placed.

Also, if we imagine a grid with 2N longitudes and N latitudes, the discrete harmonic basis consists of the functions:

$$\cos m\lambda P_l^m(\theta) \text{ and } \sin m\lambda P_l^m(\theta) \tag{C.17}$$

with $m \leq n$ and $n \leq N$, which gives N^2 basis functions - which is only *half* the number of grid points $2N^2$. This means there is no solution to the interpolation problem on the sphere, so spectral analysis is given as the solution to the least squares problem.

C.8 Power spectrum in signal processing

In signal processing, the total power of a function, f, is defined as the integral of the function squared, divided by the area of its domain.

For real unit-power spherical harmonic functions, this means that the total power of a function defined on the unit sphere is related to its spectral coefficients by a generalisation of *Parseval's theorem*:

$$\frac{1}{4\pi}\int_\Omega f(\Omega)^2 d\Omega = \sum_{l=0}^{\inf} S_{ff}(l) \tag{C.18}$$

where:

$$s_{ff}(l) = \sum_{m=-l}^{l} f_{lm}^2 \tag{C.19}$$

is defined as the angular power spectrum.

This leads naturally into the definition of the cross-power of two functions as:

$$\frac{1}{4\pi}\int_\Omega f(\Omega)g(\Omega) d\Omega = \sum_{l=0}^{\inf} S_{fg}(l) \tag{C.20}$$

where:

$$s_{fg}(l) = \sum_{m=-l}^{l} f_{lm}g_{lm} \tag{C.21}$$

is defined as the cross-power spectrum.

If the functions f and g have zero mean (i.e. f_{00} and g_{00} are zero) then $S_{ff}(l)$ and $S_{fg}(l)$ represent the contributions to the function's variance and covariance for degree l respectively. It is common that the (cross-)power spectrum is approximated by a power law of the form:

$$S_{ff}(l) = Cl^{\beta} \tag{C.22}$$

When $\beta = 0$, the spectrum is "white" as each degree possesses equal power. If $\beta < 0$, the spectrum is termed "red" as there is more power in the low degrees with long wavelengths than the higher degrees. In $\beta > 0$ it is termed "blue".

C.9 Parseval's theorem

This usually refers to the result that the Fourier transform is unitary - that the sum (or integral) of the square of a function is equal to the sum (or integral) of the square of its transform. It originates from a 1799 theorem about series by Marc-Antoine Parseval, which was later applied to the Fourier series. It is also known as Rayleigh's energy theorem, or Rayleigh's Identity, after John William Strutt, Lord Rayleigh.[2] Although the term "Parseval's theorem" is often used to describe the unitarity of any Fourier transform, especially in physics and engineering, the most general form of this property is more properly called the Plancherel theorem.

Suppose we have two complex-valued functions, $A(x)$ and $B(x)$, with Fourier series:

$$A(x) = \sum_{n=-\inf}^{\inf} a_n e^{inx} \tag{C.23}$$

$$B(x) = \sum_{n=-\inf}^{\inf} b_n e^{inx} \tag{C.24}$$

$$\tag{C.25}$$

Then,

$$\sum_{n=-\inf}^{\inf} a_n \bar{b_n} = \frac{1}{2\pi} \int_{-\pi}^{\pi} A(x) \bar{B(x)} dx \tag{C.26}$$

From here, we can see that if $A = B$:

$$\sum_{n=-\inf}^{\inf} |a_n|^2 = \frac{1}{2\pi} \int_{-\pi}^{\pi} |A(x)|^2 dx \tag{C.27}$$

Also, if we are considering the Fourier series for real-valued functions A and B (which corresponds to a_0 real, $a_{-n} = \bar{a_n}$, b_0 real, and $b_{-n} = \bar{b_n}$). In this case:

$$a_0 b_0 + 2 * Real[\sum_{n=1}^{\inf} a_n \bar{b_n}] = \frac{1}{2\pi} \int_{-\pi}^{\pi} A(x)B(x)dx \tag{C.28}$$

C.9.1 Usage

C.9.1.1 Continuous Fourier transfom

$$\int_{-\inf}^{\inf} |x(t)|^2 dt = \int_{-\inf}^{\inf} |X(f)|^2 df \tag{C.29}$$

Here, we are stating that the total energy contained in a waveform $x(t)$ summed across all time, t, is equal to the total energy of the waveform's Fourier Transform, $X(f)$, summed across all of its frequency components, f.

C.9.1.2 Discrete Fourier transform

$$\sum_{n=0}^{N-1} |x[n]|^2 = \frac{1}{N} \sum_{k=0}^{N-1} |X[k]|^2 \tag{C.30}$$

$X[k]$ is the DFT of $x[n]$, both having length N.

C.9.2 Equivalence of norm and inner product forms

Definition of inner product form:

$$\int_{-\inf}^{\inf} x(t)\bar{y(t)}dt = \int_{-\inf}^{\inf} X(f)\bar{Y(f)}df \tag{C.31}$$

norm form:

$$\int_{-\inf}^{\inf} |x(t)|^2 dt = \int_{-\inf}^{\inf} |X(f)|^2 df \tag{C.32}$$

Using the polarisation identity (expressing the dot product of two vectors in terms of the Euclidean norm - $||x|| := \sqrt{x_1^2 + x_2^2 + ... + x_n^2}$ for Real, and $||z|| := \sqrt{|z_1|^2 + |z_2|^2 + ... + |z_n|^2} = \sqrt{z_1\bar{z_1} + z_2\bar{z_2} + ...z_n\bar{z_n}}$ for Complex numbers (also called absolute value or modulus)):

$$u.v = \frac{1}{2}(||u+v||^2 - ||u||^2 - ||v||^2) \tag{C.33}$$

$$u.v = \frac{1}{2}(||u||^2 + ||v||^2 - ||u-v||^2) \tag{C.34}$$

$$u.v = \frac{1}{4}(||u+v||^2 - ||u-v||^2) \tag{C.35}$$

It's possible to use the parallelogram law ($2||\mathbf{u}||^2 + 2||\mathbf{v}||^2 = ||\mathbf{u}+\mathbf{v}||^2 + ||\mathbf{u}-\mathbf{v}||^2$) to get other forms.

C.9.2.1 Derivation

Relation between norm and dot product is:

$$||v||^2 = v.v \tag{C.36}$$

So:

$$||u+v||^2 = (u+v).(u+v)// \quad = (u.u) + (u.v) + (v.u) + (v.v)// = ||u||^2 + ||v||^2 + 2(u.v) \tag{C.37}$$

and similarly:

$$||u-v||^2 = ||u||^2 + ||v||^2 - 2(u.v) \tag{C.38}$$

Thus, the first and second identities above follow by solving these equations for $||u.v||$, whilst the third follows from subtracting them (adding gives the parallelogram law).

Thus:

$$a\bar{b} = \frac{1}{4}(|a+b|^2 + i|a+ib|^2 + i^2|a+i^2b|^2 + i^3|a+i^3b|^2) \tag{C.39}$$

which is true for all complex numbers a and b.

C.10 Cross-correlator

$$\xi_{1,2}(\hat{s}) = \int_{t-(\tau/2)}^{t+(\tau/2)} g_1(t^{'} - \vec{r}_1 \cdot \hat{s}) g_2(t^{'} - \vec{r}_2 \cdot \hat{s}) dt^{'} \tag{C.40}$$

As the data has been discretely sampled at an interval of δt (sample frequency of f_s), the above integral can be written as an inner product:

$$\xi_{1,2}(\hat{s}) = \frac{1}{N}\vec{g}_1^T \cdot \mathbb{T}^T(\vec{r}_1 \cdot \hat{s}) \cdot \mathbb{T}(\vec{r}_2 \cdot \hat{s}) \cdot \vec{g}_2 \tag{C.41}$$

Appendix D

Analysis of S6B run

D.1 Science run S6B/VSR2

The "B" period began at GPS time 937958415 (Saturday Sept 26,2009 at 00:00 UTC) and ended at GPS 947023216 (Friday Jan.8, 2010 at 22:00 UTC). The total amount of background generated for S6B is 123745850 seconds, which works out at just under 4 years of analysed data, which produced 4723658 triggers. As these are split into two sets for the purposes of tuning the analysis, it leaves us 2361829 background triggers - potentially meaning we can get down to a false alarm rate of $0.5yr^{-1}$. As stated previously, the analysis requires tuning, so we split the triggers into two, train on half, then apply the cut to the other half.

D.1.1 Background distribution of the null energy

As with S6A, Figure D.1 shows how the null energy is distributed across the coherent/incoherent plane. It look very similar to the distribution for the S6A background and provides a similar border hugging line to use in rejecting background The rates plot is shown in Figure D.2, and demonstrates the effect of the null cut on the background. Once again, the number of significant triggers is reduced and there is no tail trailing to high Γ.

D.1.2 Background distribution of the energy in the spherical harmonic coefficients

Figure D.3 shows how the "coherent" and "incoherent" parts of the spherical harmonic coefficients are distribution on the plane. The shape is different to that produced for the background in S6A, Figure 5.10, with a more noticeable "hole", or bifurcation, as the plot tends towards high E_{sh} and I_{sh}. Still, the tuning produces a closely fitting cut that tracks the boundary well, with very few surviving triggers. Figure D.4 is the rate vs Γ plot similar to Figure D.2

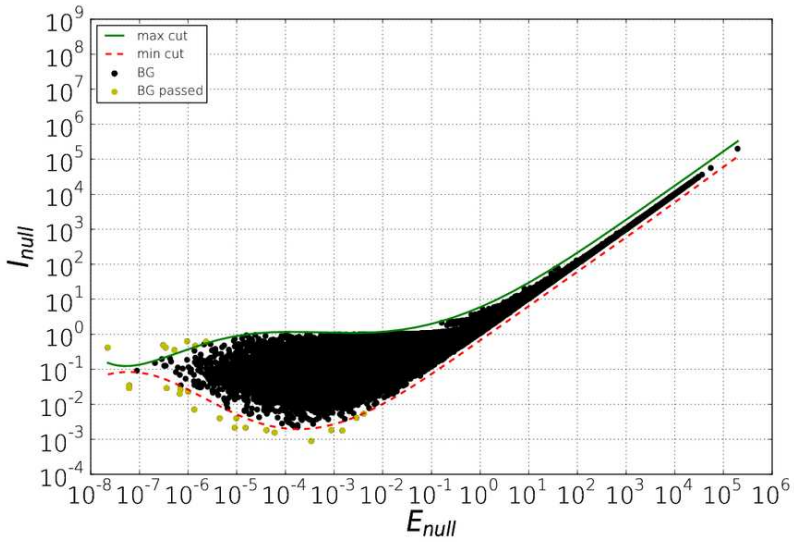

FIGURE D.1: Scatter plot showing the E/I distribution for the 2^{nd} set of background triggers and the corresponding cut for the null energy.

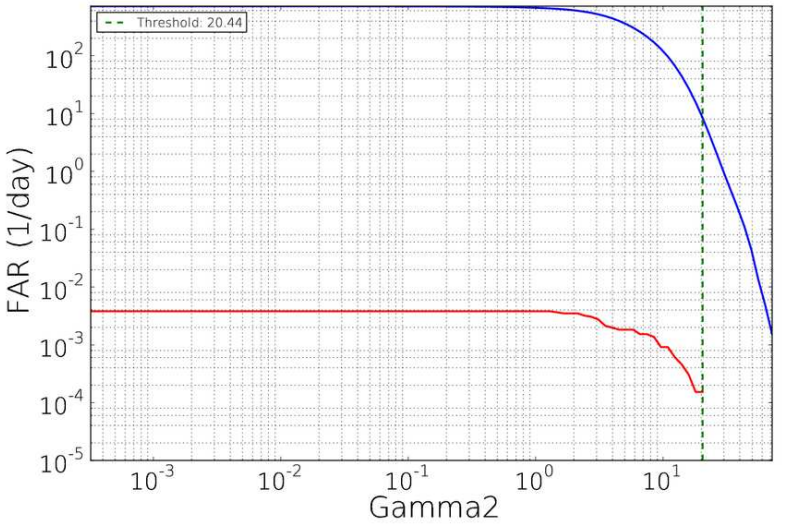

FIGURE D.2: Plot of cumulative rate vs Γ for tuning and null cut background

but the cut applied to the background is that generated by the tuning based on the spherical harmonic coefficients. In this case, both cuts leave a loud background trigger with a Γ value of 20.44.

D.1.3 Short duration simulations

Just as with S6A, we will present figures demonstrating how triggers from the different injection types are distributed in the null and spherical harmonic plane (see Table 5.1 for the parameters

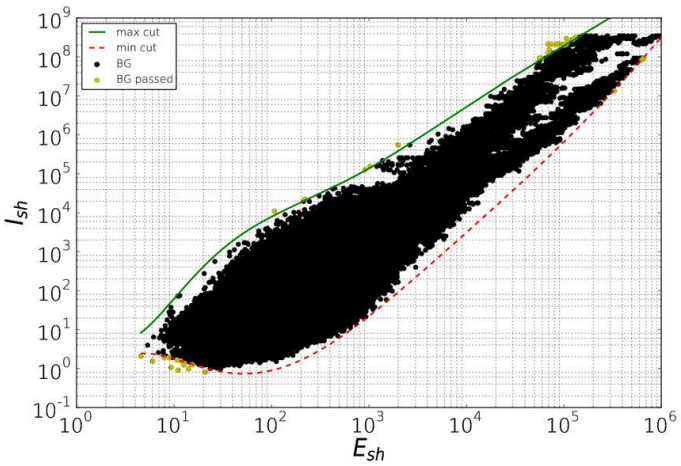

FIGURE D.3: Scatter plot showing the E/I distribution of background triggers and the corresponding cut for the spherical harmonic coefficients energy

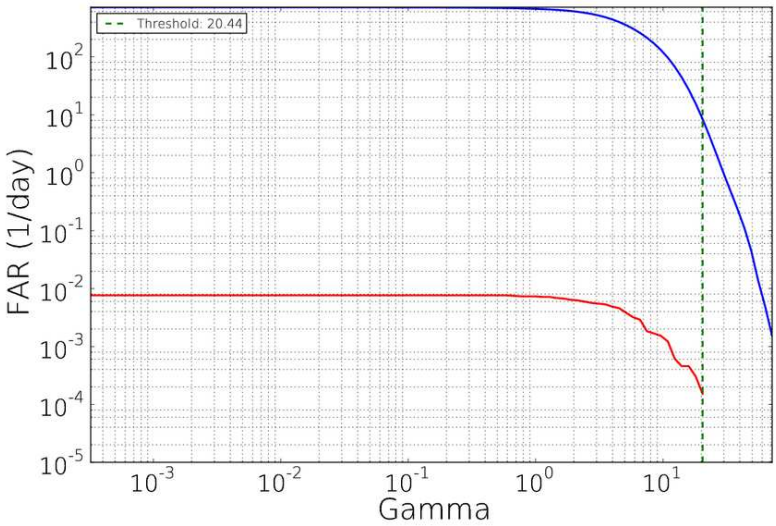

FIGURE D.4: Plot of cumulative rate vs Γ for tuning and spherical harmonic cut background

used) in the following sections. Note that the paired scatter plots share a common legend while the efficiency plots have their own legend giving the value of the 50% and 90% upper limits.

D.1.3.1 Behaviour of white noise bursts

The following figures demonstrate the distribution of triggers of white noise burst injections in both the null and spherical harmonic plane. Figure D.5, demonstrates the good separation between the loud injection triggers (plotted in red) and the background triggers (plotted in black). The efficiency curves show how the null cut does better for lower frequency injections, while the spherical harmonic cut just edges it for higher frequency injections. However, we do get a small amount of turn over at high signal amplitudes which was not present in the

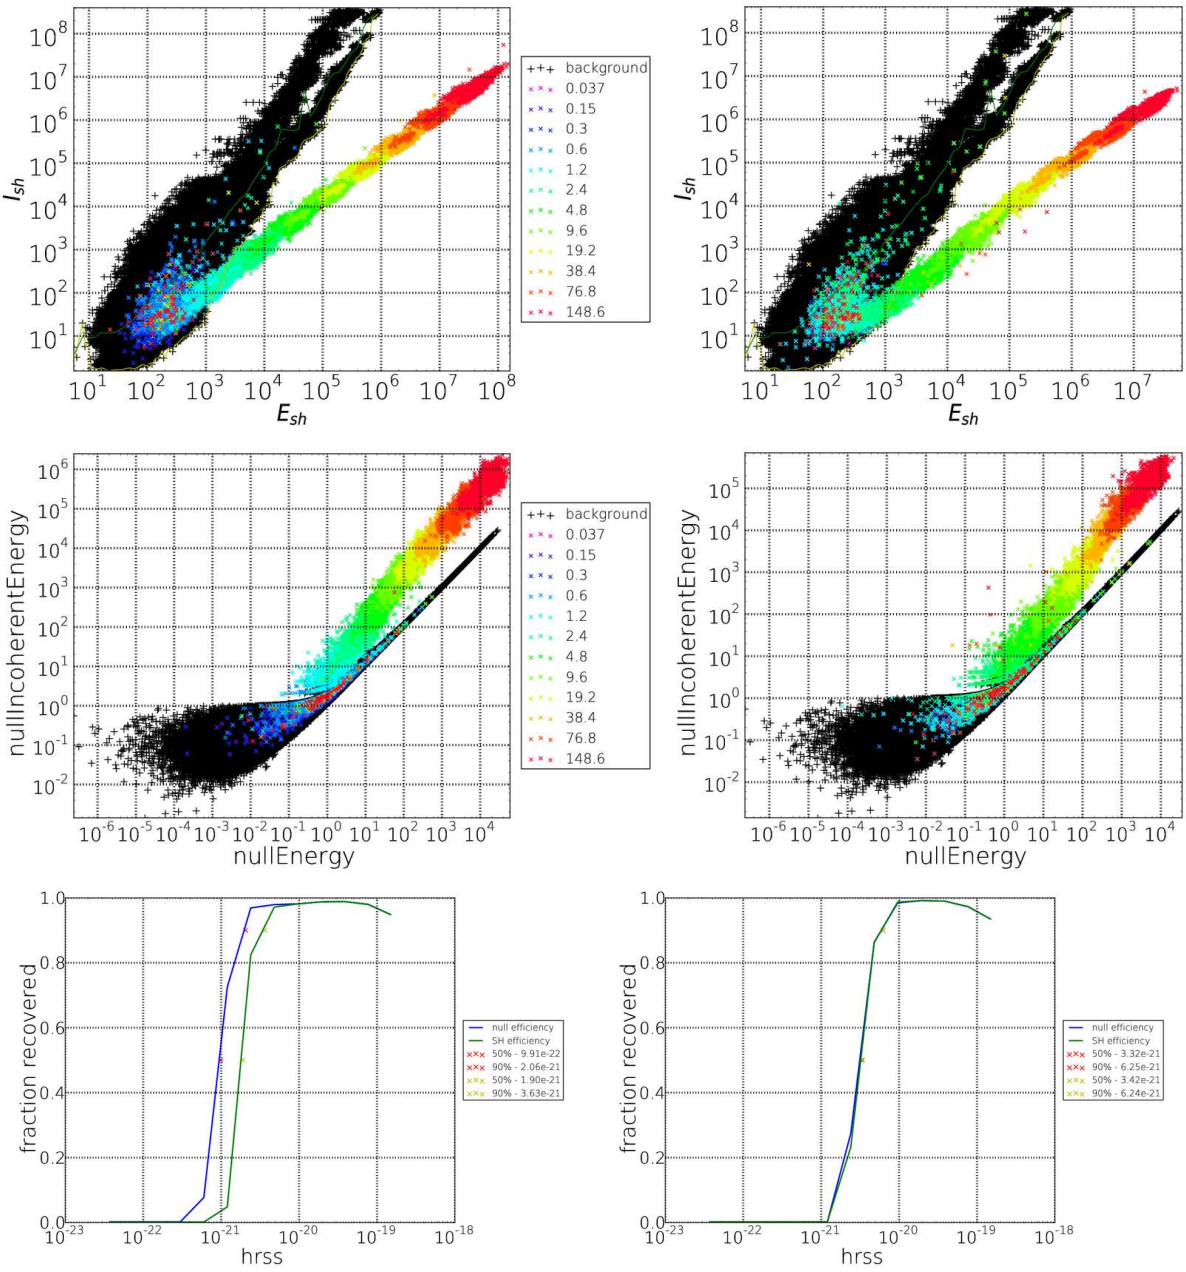

FIGURE D.5: The plots in this figure show the distribution of triggers in the null and spherical harmonic plane for S6B, and the corresponding efficient curves generated by applying the appropriate cut to the triggers. The top pair of figures correspond to spherical harmonic energy, while the centre pair show the same for null energy. The left most figures relate to WNBs with f_c 300 Hz and Δf=50 Hz, while the right most figures show WNBs with a f_c of 750 Hz and Δf of 250 Hz. The bottom pair of plots show the efficiency curves for both the null (in blue) and spherical harmonic (in green) cuts.

S6A results. While the exact cause of this is unknown, the fact that there are red crosses present in the background suggests that the wrong triggers are being associated with these loud injections.

D.1.3.2 Behaviour of linearly polarised sine-Gaussians

The following figures demonstrate the distribution of triggers of SGL injections in both the null and spherical harmonic plane. Once again, Figure D.6 demonstrates the good separation between loud injection triggers (plotted in red) and background triggers (plotted in black). The efficiency curves show that the null cut does a better job at lower frequencies, while the spherical harmonic cut is slightly better at higher frequencies. Generally the plots in Figure D.7 are similar to those seen for S6A (Figure D.7), although the efficiency of injection recovery is slightly worse.

D.1.3.3 Behaviour of circularly polarised sine-Gaussians

The following figures demonstrate the distribution of triggers for injected circularly polarised sine-Gaussians in both the null and spherical harmonic plane. Once again, the plots in Figure D.8 demonstrate that the circularly polarised SGs are recovered the best out the SGs. Even though there has been a general reduction in efficiency across the board for SG in S6B, it is a marginal decrease.

D.1.4 Long duration simulations

The figures in the following sections demonstrate the distribution of triggers from the different injection types in the null and spherical harmonic plane. See Tables 5.2, 5.3 and 5.4 for the parameters used for the POW, inspiral waveforms and the SGCs respectively. Note that the paired scatter plots share a common legend while the efficiency plots have their own legend giving the value of the 50% and 90% efficiency.

D.1.4.1 Behaviour of poly-oscillating waveforms

The following figures demonstrate the distribution of triggers for injected POW in both the null and spherical harmonic plane.

Regarding Figure D.9, it demonstrates how well the SphRad pipeline can recover complicated long duration waveforms. There is a large separation between the loud injection triggers (plotted in red) and the background triggers (plotted in black). Note however that there some of the loudest triggers (in red) appear to be present among the background, correspondingly

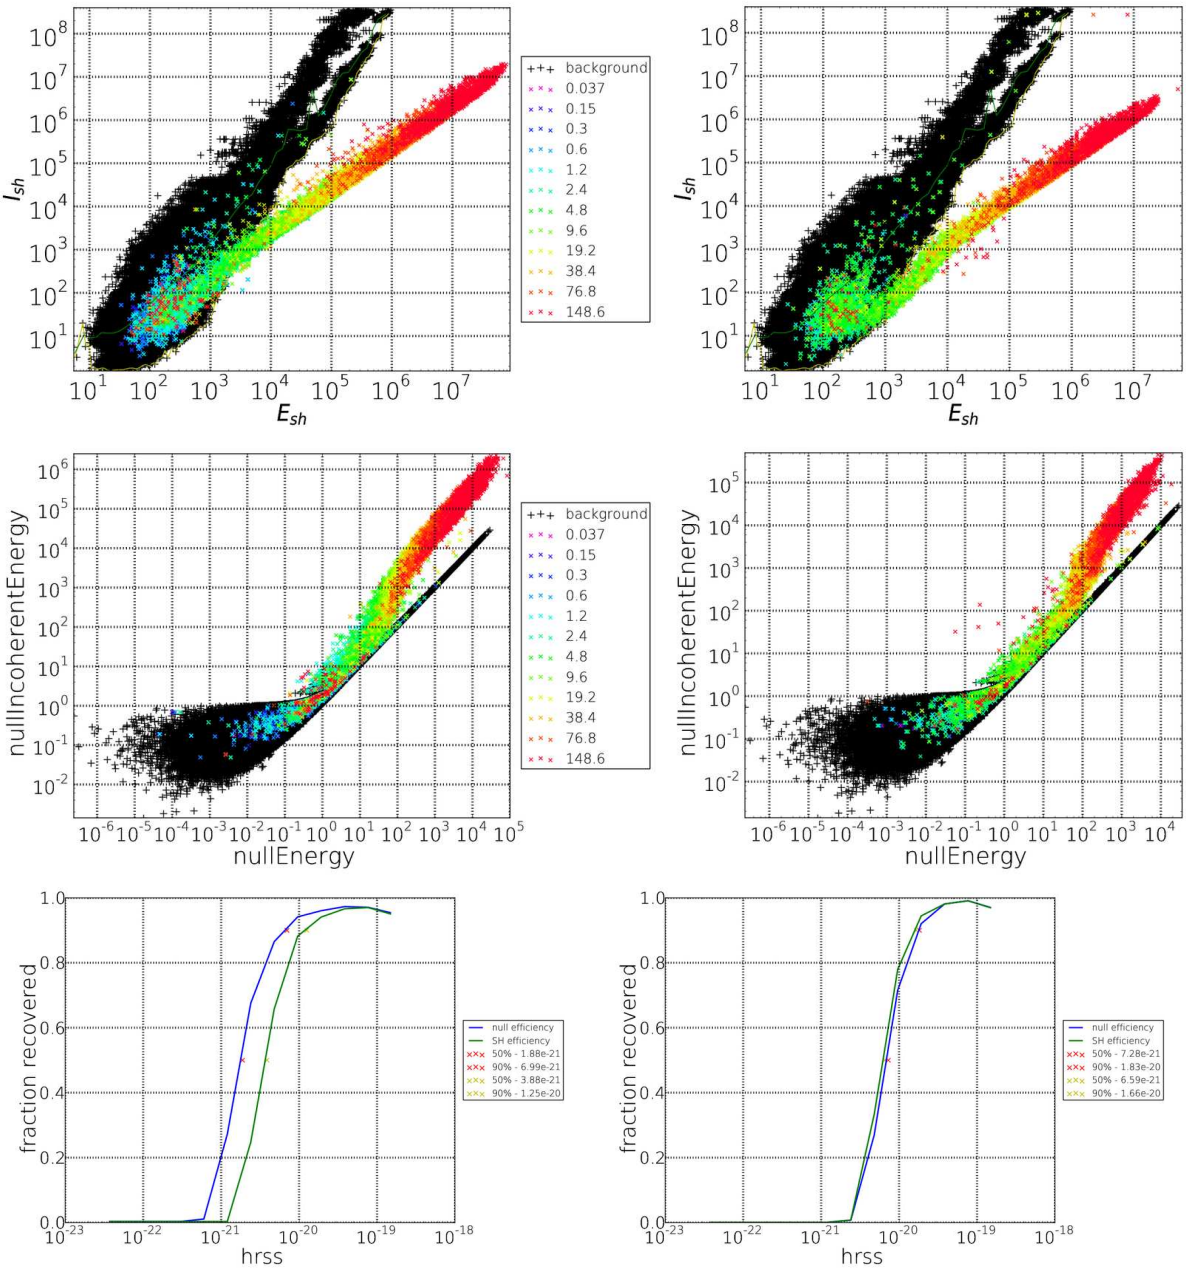

FIGURE D.6: The plots in this figure show the distribution of triggers in the null and spherical harmonic plane for S6B, and the corresponding efficient curves generated by applying the appropriate cut to the triggers. The top pair of figures correspond to spherical harmonic energy, while the centre pair show the same for null energy. The left most figures relate to SGLs with f_c 235 Hz and Q=3, while the right most figures show SGL with f_c = 853 Hz and Q=3. The bottom pair of plots show the efficiency curves for both the null (in blue) and spherical harmonic (in green) cuts.

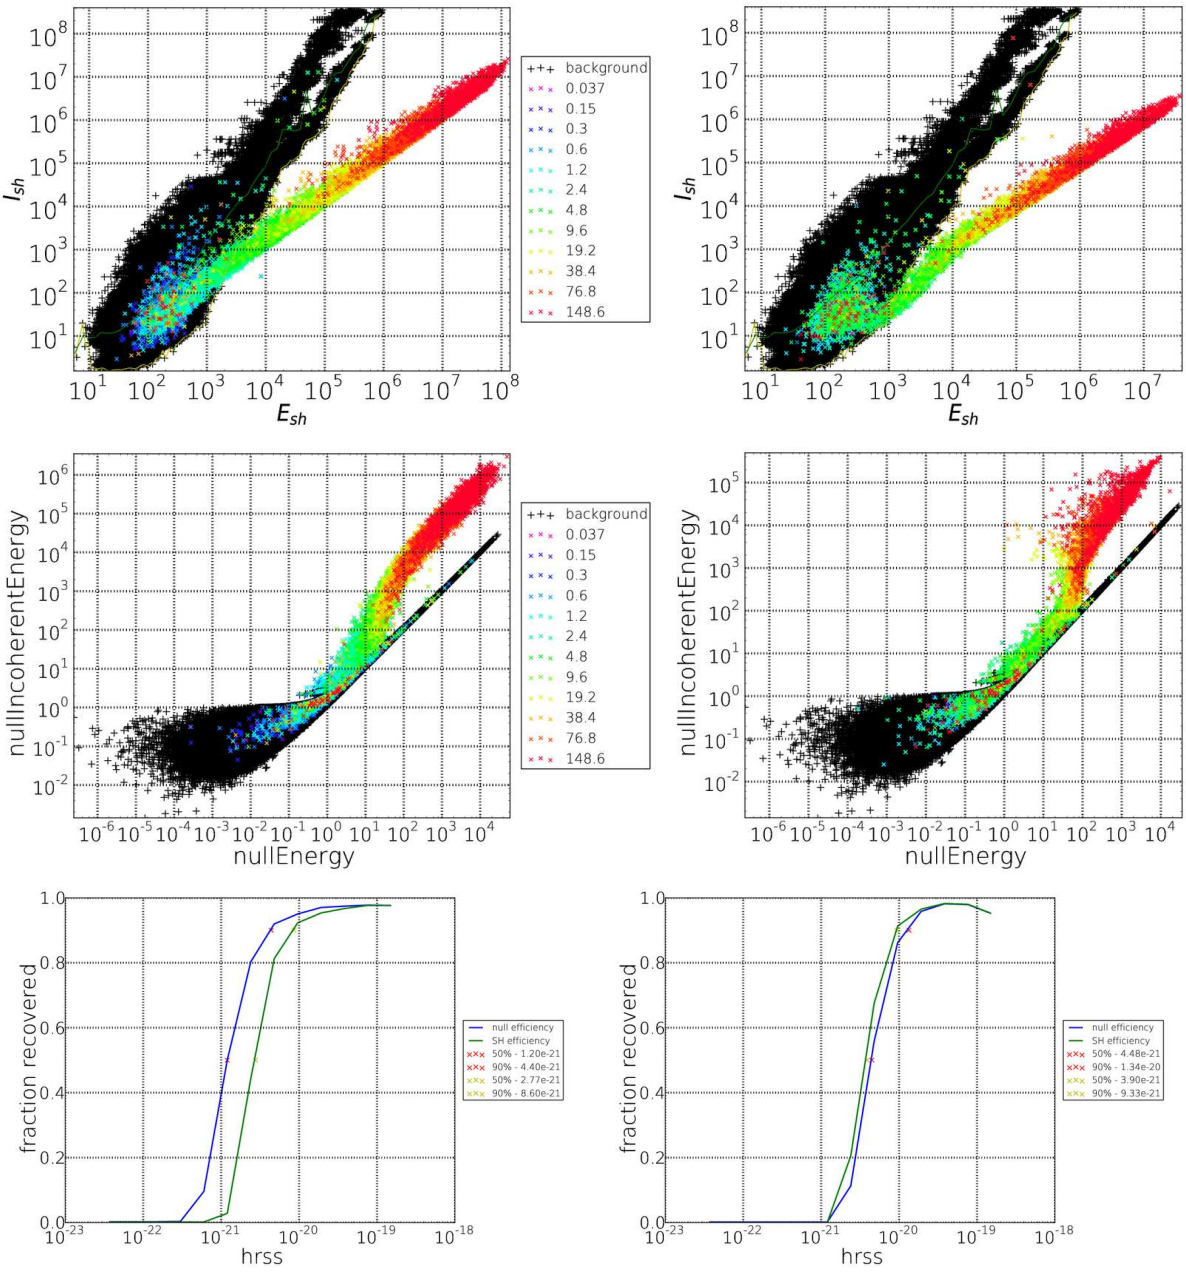

FIGURE D.7: This figure shows the distribution of triggers in the null and spherical harmonic plane for S6B, and the corresponding efficient curves generated by applying the appropriate cut to the triggers. The top pair of figures correspond to spherical harmonic energy, while the centre pair show the same for null energy. The left most figures relate to SGLs with f_c 235 Hz and Q=9, while the right most figures show SGL with $f_c = 853$ Hz and Q=9. The bottom pair of plots show the efficiency curves for both the null (in blue) and spherical harmonic (in green) cuts.

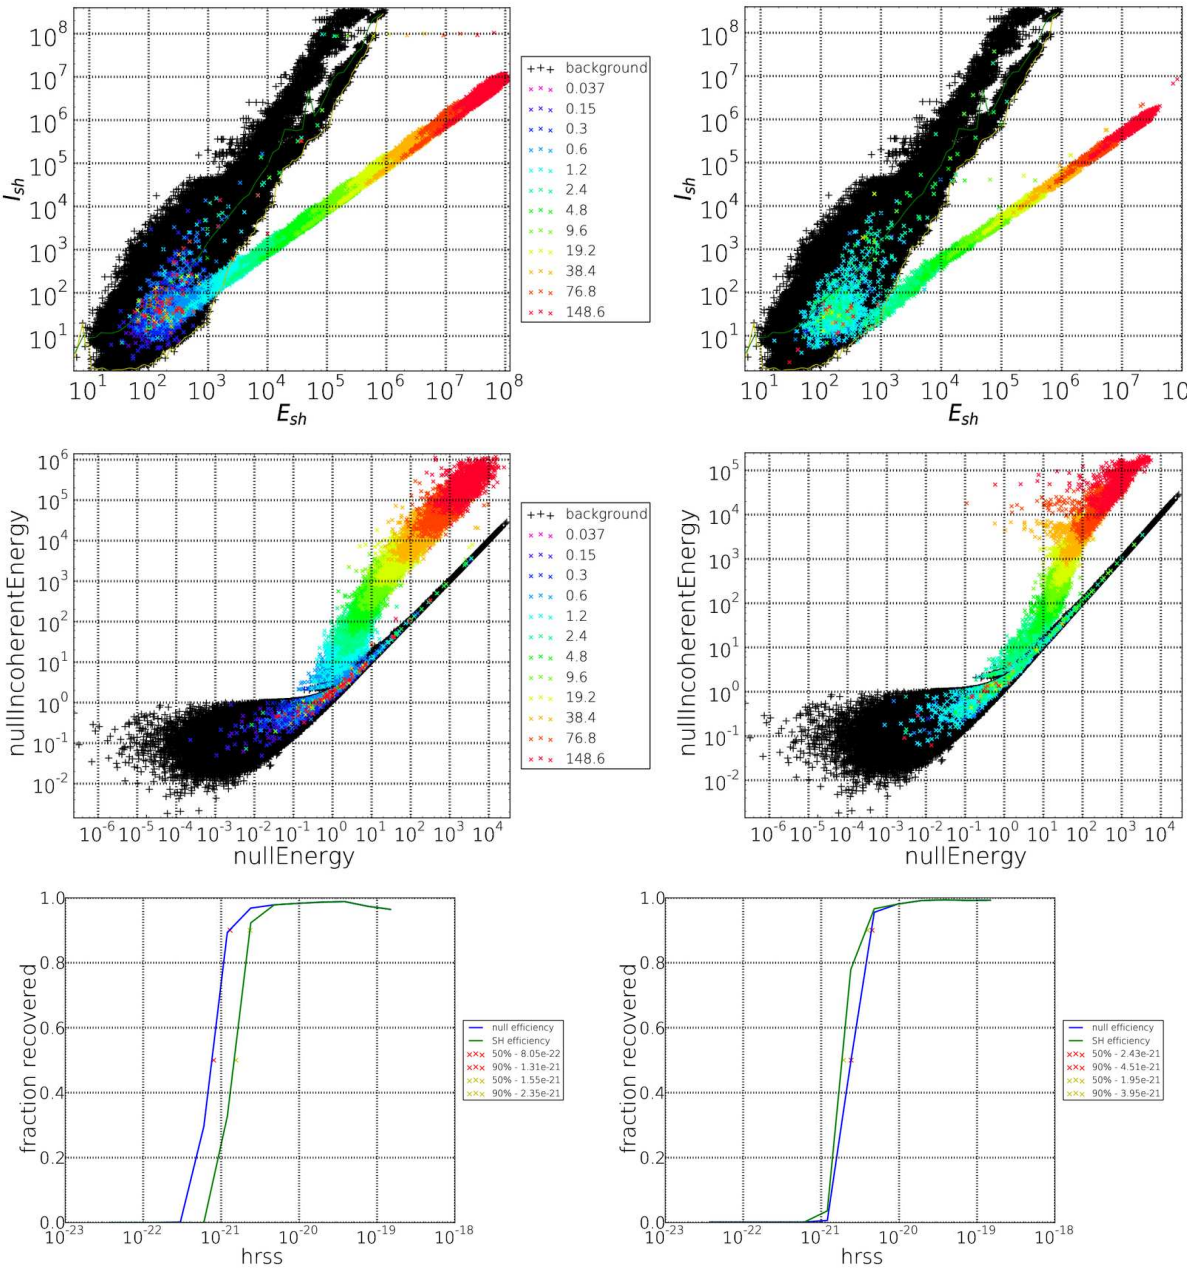

FIGURE D.8: The plots of this figure show the distribution of triggers in the null and spherical harmonic plane for S6B, and the corresponding efficient curves generated by applying the appropriate cut to the triggers. The top pair of figures correspond to spherical harmonic energy, while the centre pair show the same for null energy. The left most figures relate to SGCs with f_c 235 Hz and Q=9, while the right most figures show SGCs with f_c = 853 Hz and Q=9. The bottom pair of plots show the efficiency curves for both the null (in blue) and spherical harmonic (in green) cuts.

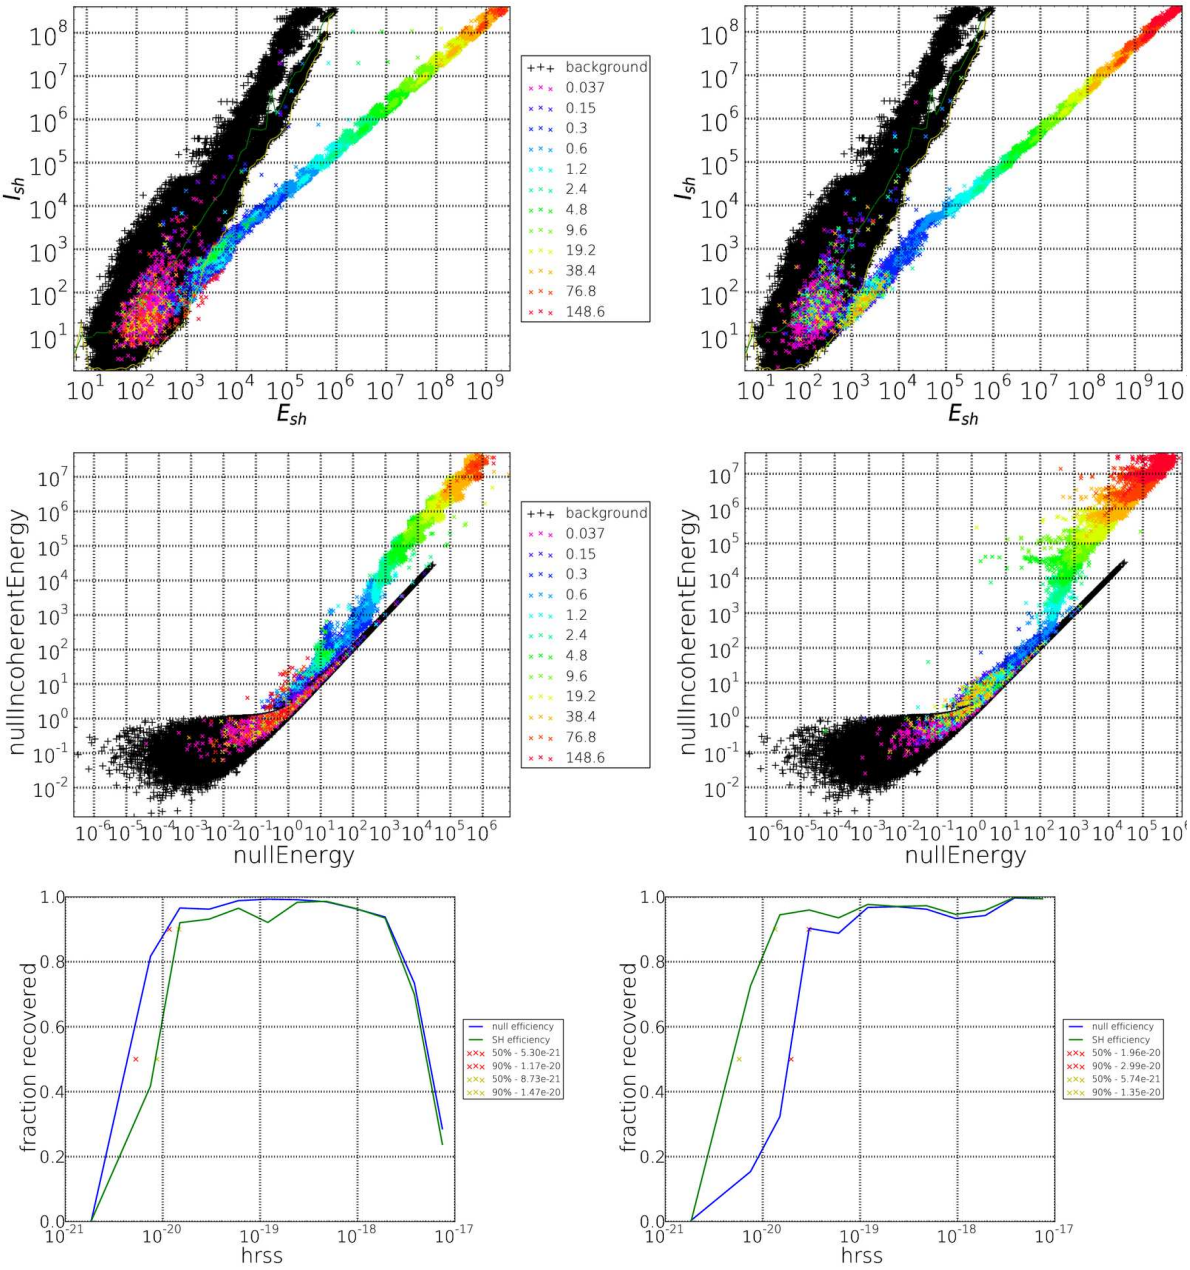

FIGURE D.9: The plots of this figure show the distribution of triggers in the null and spherical harmonic plane for S6B, and the corresponding efficient curves generated by applying the appropriate cut to the triggers for S6B. The top figures correspond to spherical harmonic energy, while the centre show the same for null energy. The left most figures relate to POWs with $f_c = 235$ Hz, $\Delta f = 100$ Hz, and $\Delta t = 12$ seconds, while the right most figures show POWs with $f_c = 852$, $\Delta f = 50$ Hz, and $\Delta t = 24$ seconds. The bottom plots show the efficiency curves for both the null (in blue) and spherical harmonic (in green) cuts.

pulling the efficiency at very large injection down. The cause of this is probably as simple as the wrong trigger being associated with the loud injections (due to timing mismatches etc), or just the fact that this is a unrealistically high amplitude for a real gravitational wave.

D.1.4.2 Behaviour of Inspiral simulations

The following figures demonstrate the distribution of triggers for injected inspiral waveforms in both the null and spherical harmonic plane.

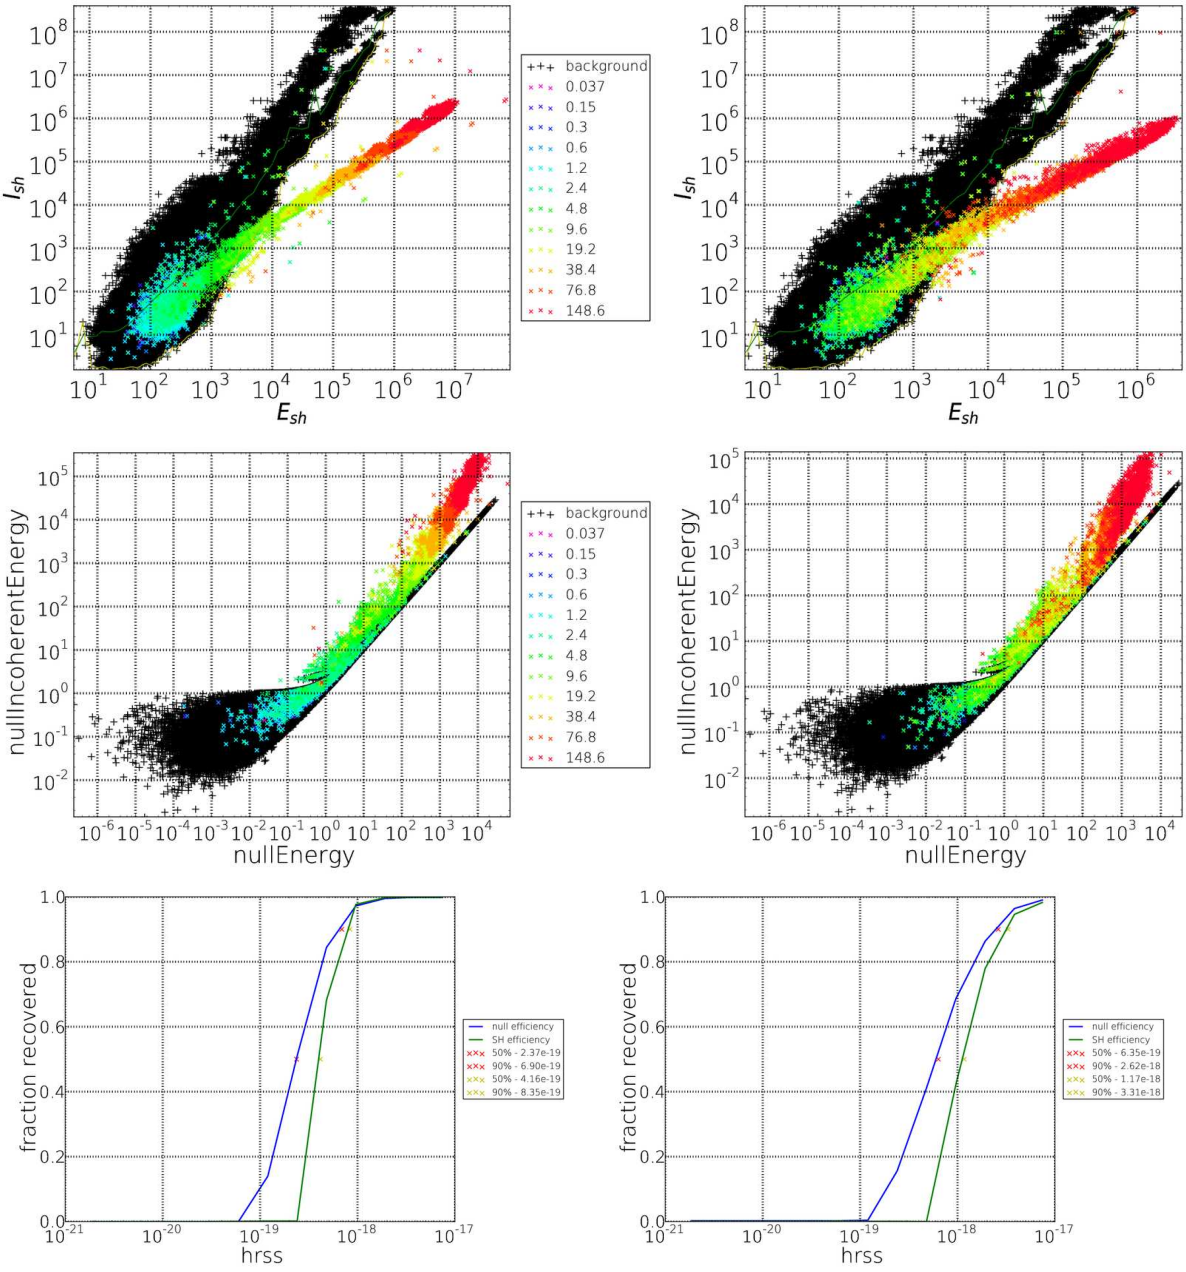

FIGURE D.10: The plots of this figure show the distribution of triggers in the null and spherical harmonic plane for S6B, and the corresponding efficient curves generated by applying the appropriate cut to the triggers. The top pair of figures correspond to spherical harmonic energy, while the centre pair show the same for null energy. The left most figures relate to inspirals with an inclination angle of $\pi/4$, while the right most figures show inspirals with an inclination angle of $\pi/2$ (corresponds to edge on). The bottom pair of plots show the efficiency curves for both the null (in blue) and spherical harmonic (in green) cuts.

Figure D.10 continues to shows that SphRad struggles to recover inspirals with any great efficiency. Could be timing issues, could be something to do with the whitening issue.

D.1.4.3 Behaviour of long duration circularly polarised sine-Gaussians

The following figures demonstrate the distribution of triggers for long duration circularly polarised sine-Gaussians in both the null and spherical harmonic plane. The plots shown in Figure D.11 display the same lack of performance as Figure 5.18, highlighting the whitening issue.

D.1.4.4 Behaviour of accretion-disk instabilities

The following figures demonstrate the distribution of triggers for very long duration ($\approx$ 220 seconds) accretion disk instability waveforms in both the null and spherical harmonic plane. Note that the ADI waveform is precomputed and so does not have any free parameters. The waveform is still injected with a range of injection factors however, so we can generate the efficiency plots. The plots relating to the recovery of the ADI waveform as shown in Figure D.12 are not substantially difference to their S6A counterparts, like most of the plots for S6B. However, there is a noticeable turning over of some of the efficiency curves at very high amplitudes, perhaps suggesting that the data from the S6B run is a little cleaning than S6A, so what we are seeing is some sort of saturation effect.

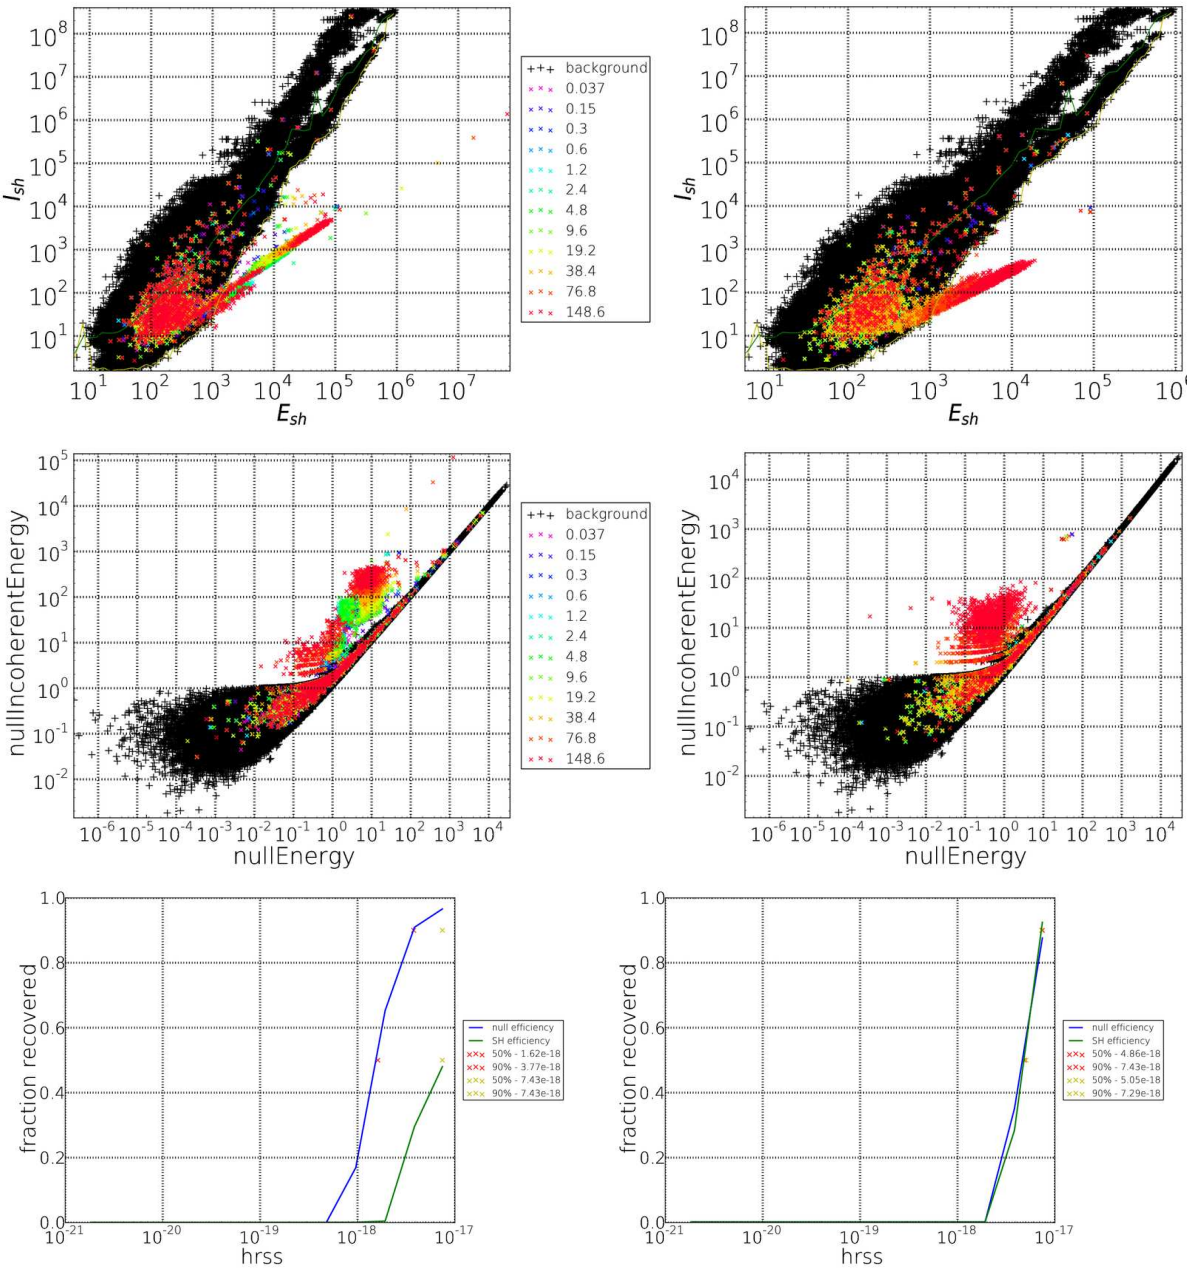

FIGURE D.11: The plots of this figure show the distribution of triggers in the null and spherical harmonic plane for S6B, and the corresponding efficient curves generated by applying the appropriate cut to the triggers. The top figures correspond to spherical harmonic energy, while the centre show the same for null energy. The left most figures relate to SGCs with f_c 235 Hz and Δt=24 seconds, while the figures on the right hand side show SGCs with f_c 512 Hz and Δt=96 seconds. The bottom plots show the efficiency curves for both the null (in blue) and spherical harmonic (in green) cuts.

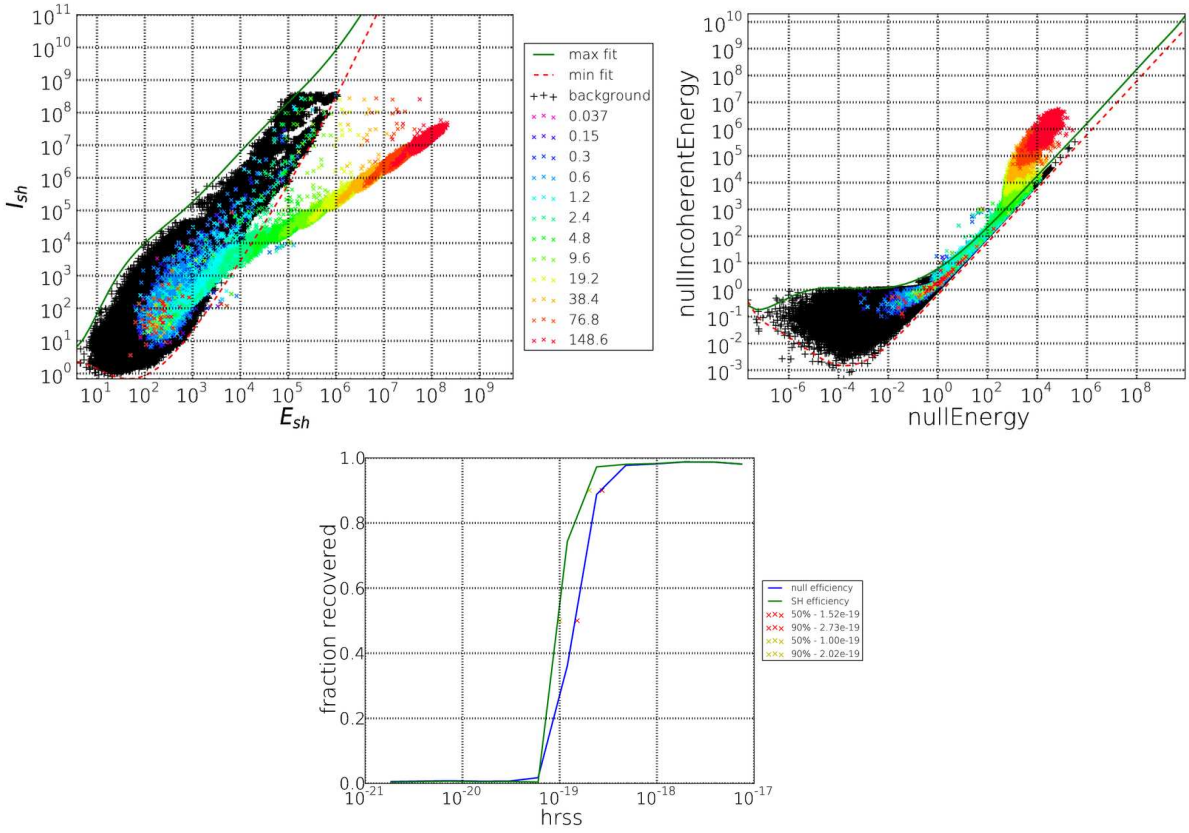

FIGURE D.12: The plots of this figure show the distribution of triggers in the null and spherical harmonic plane for S6B, and the corresponding efficient curves generated by applying the appropriate cut to the triggers. The top left figure shows the distribution for spherical harmonic energy, while the top right shows the same for null energy. The bottom plot show the efficiency curves for both the null (in blue) and spherical harmonic (in green) cuts.

Appendix E

Analysis of S6D run

E.1 Science run 6D/VSR3

The "d" period started at GPS 961545615 (Sat Jun 26 2010 00:00:00 UTC) and ended at GPS 971654415 (Thu Oct 21 2010 00:00:00 UTC).

The total amount of background generated for S6D is 245383628 seconds, which works out at just under 8 years of analysed data, which produced 8876284 triggers. As these are split into two sets for the purposes of tuning the analysis, it leaves us 4438142 background triggers - potentially meaning we can get down to a false alarm rate of $0.25yr^{-1}$.

E.1.1 Background distribution of the null energy

Figure E.1 shows how the null energy is distributed across the coherent/incoherent plane. This appears to be a very consistent plot, as all three Periods have looked similar in their distribution. The rates plot is shown in Figure E.2, and demonstrates the effect of the null cut on the background. Once again, the number of significant triggers is reduced and there is no tail trailing to high Γ.

E.1.2 Background distribution of the energy in the spherical harmonic coefficients

Figure E.3 shows how the spherical harmonic coefficients are distribution on the plane. Once again, the shape is different to that produced by the background in S6A and B, Figures 5.10 & D.3, with a more filled in appearance. Still, the tuning produces a closely fitting cut that tracks the boundary well, with very few surviving triggers. Figure E.4 is the rate vs Γ plot similar to Figure E.2 but the cut applied to the background is that generated by the tuning

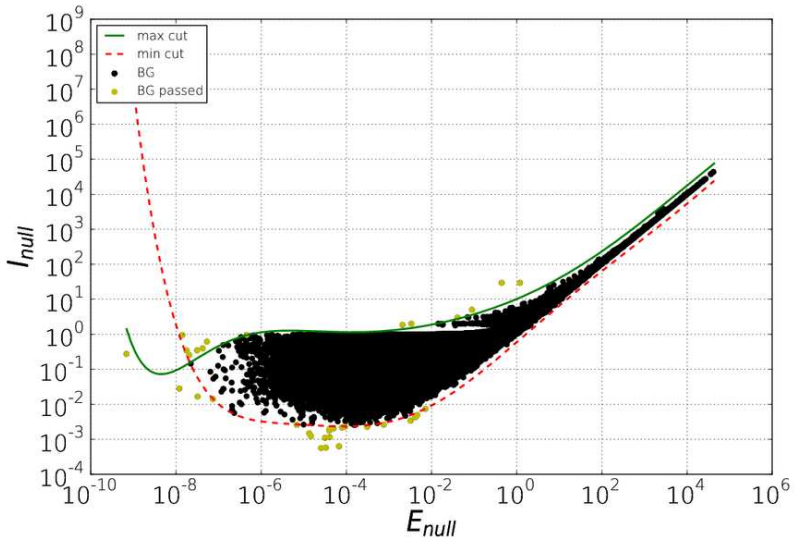

FIGURE E.1: Scatter plot showing the E/I distribution for the 2^{nd} set of background triggers and the corresponding cut for the null energy.

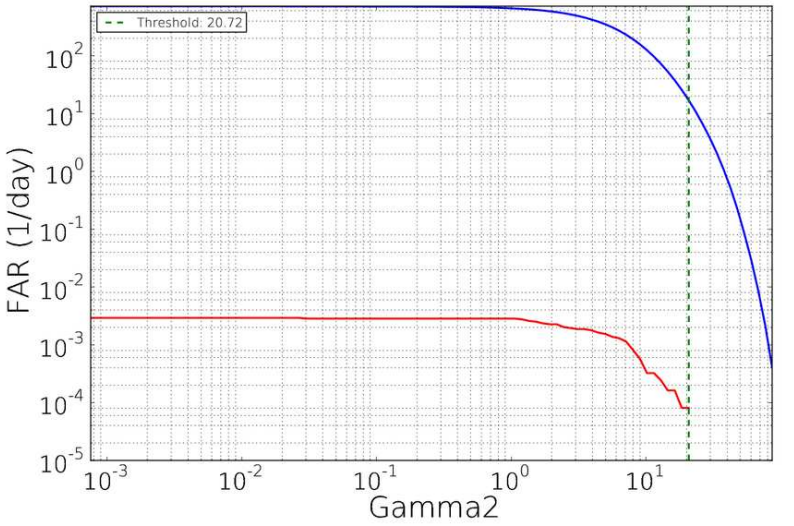

FIGURE E.2: Plot of cumulative rate vs Γ for tuning and null cut background

based on the spherical harmonic coefficients. In this case, both cuts leave a loud background trigger with a Γ value of 20.44.

E.1.3 Short duration simulations

In this section we will present example figure demonstrating how triggers from the different injection types are distributed in the null and spherical harmonic plane (see Table 5.1 for the parameters used). Note that the paired scatter plots share a common legend while the efficiency plots have their own legend giving the value of the 50% and 90% upper limits.

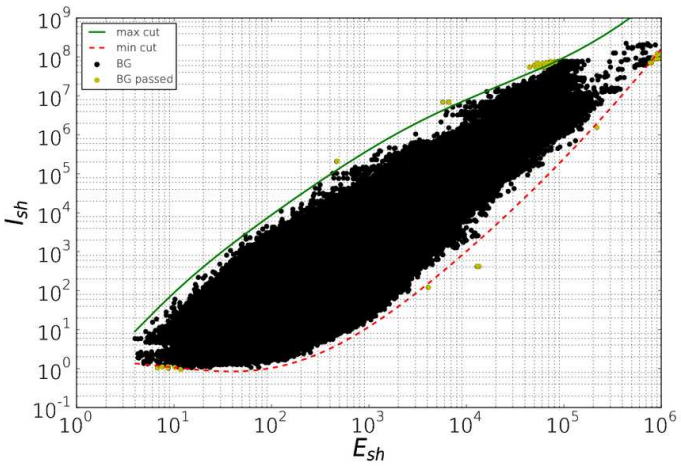

FIGURE E.3: Scatter plot showing the E/I distribution of background triggers and the corresponding cut for the spherical harmonic coefficients energy

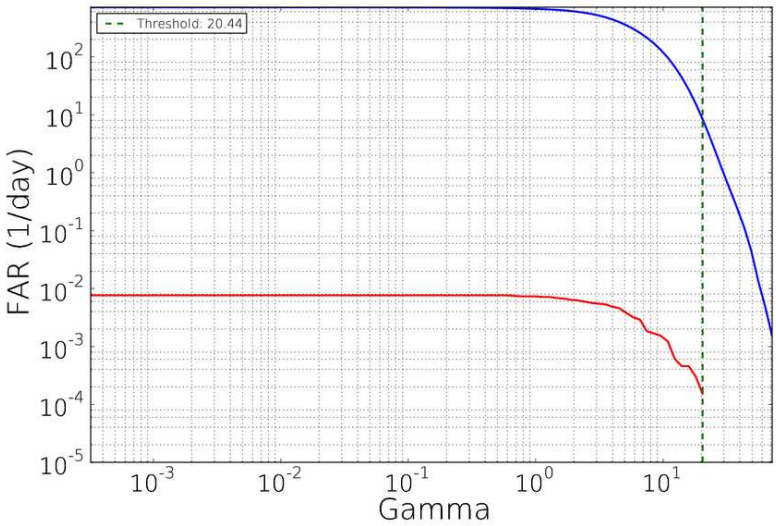

FIGURE E.4: Plot of cumulative rate vs Γ for tuning and spherical harmonic cut background

E.1.3.1 Behaviour of white noise bursts

The following figures demonstrate the distribution of triggers of white noise burst injections in both the null and spherical harmonic plane. No real surprises in Figure E.5. There is good separation between the loud injection triggers (plotted in red) and the background triggers (plotted in black), while the efficiency are on par with the others. However, we do get a small amount of turn over at high signal amplitudes which was not present in the S6A results.

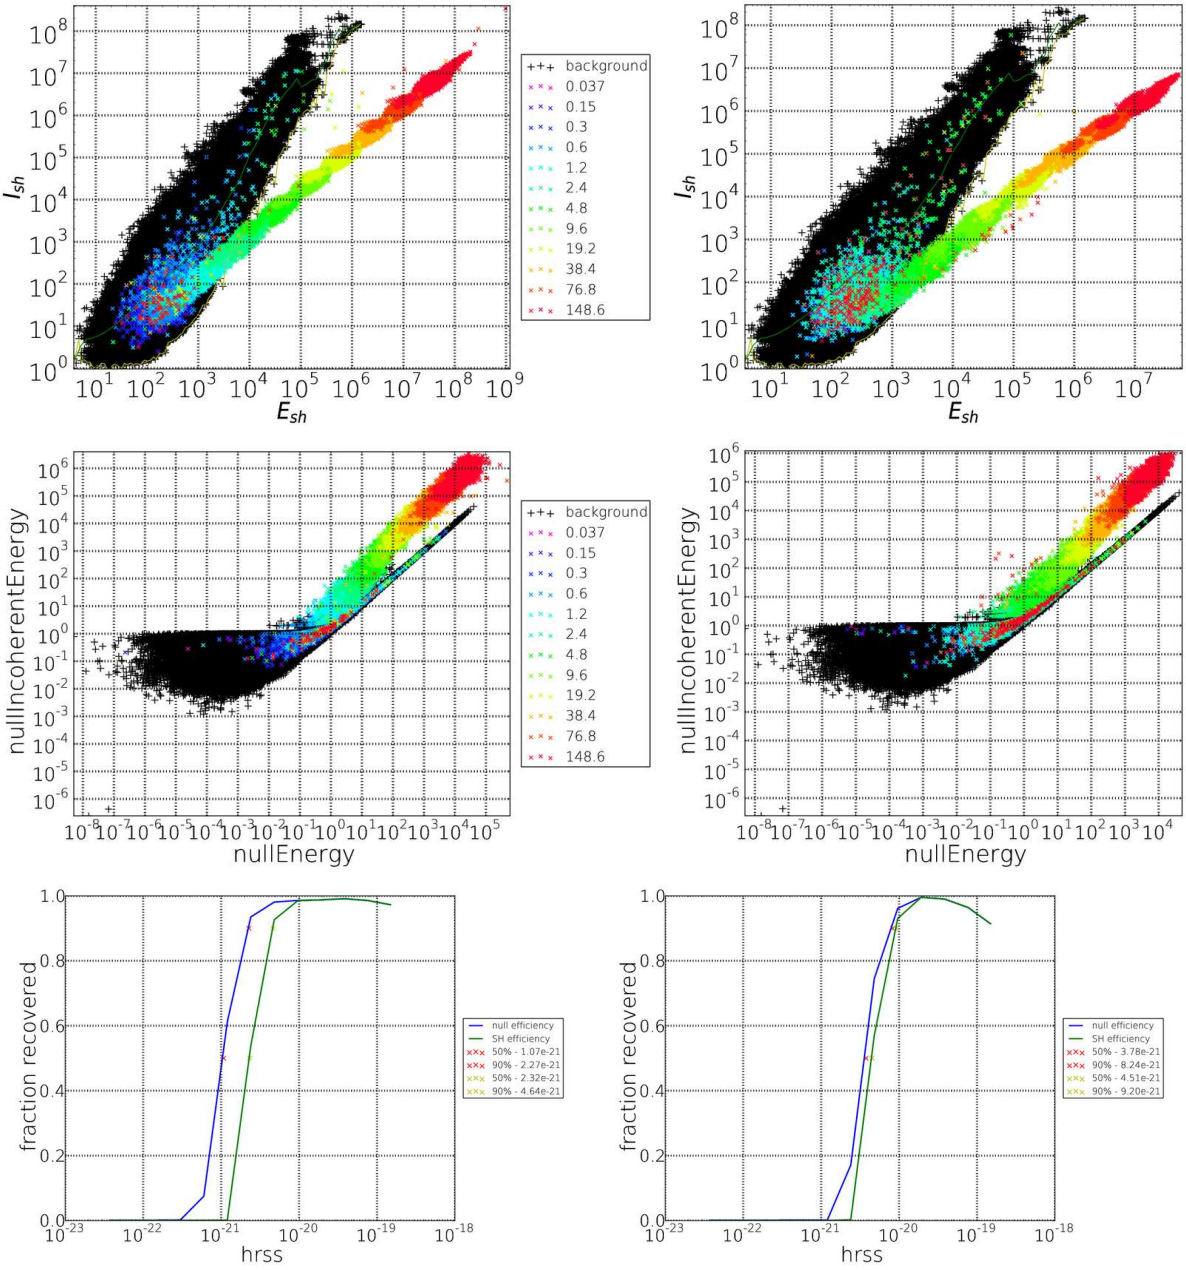

FIGURE E.5: The plots in this figure show the distribution of triggers in the null and spherical harmonic plane for S6D, and the corresponding efficient curves generated by applying the appropriate cut to the triggers. The top pair of figures correspond to spherical harmonic energy, while the centre pair show the same for null energy. The left most figures relate to WNBs with f_c 300 Hz and Δf=50 Hz, while the right most figures show WNBs with a f_c of 750 Hz and Δf of 250 Hz. The bottom pair of plots show the efficiency curves for both the null (in blue) and spherical harmonic (in green) cuts.

E.1.3.2 Behaviour of linearly polarised sine-Gaussians

The following figures demonstrate the distribution of triggers of SGL injections in both the null and spherical harmonic plane. Again, no real surprises, Figure E.6 demonstrates the

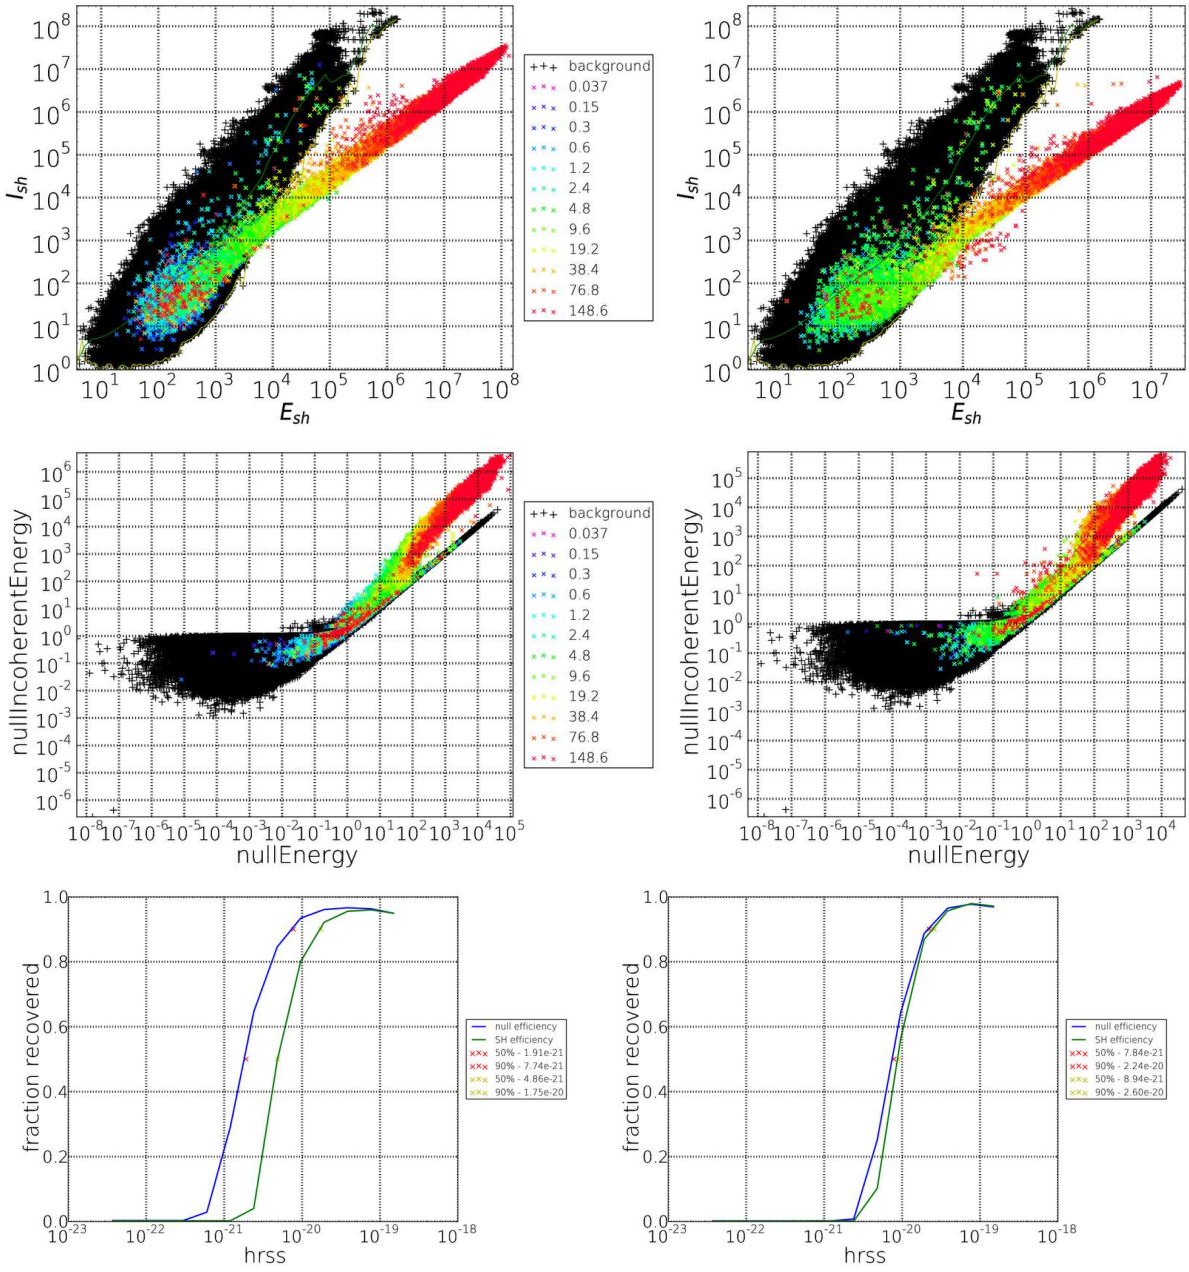

FIGURE E.6: The plots in this figure show the distribution of triggers in the null and spherical harmonic plane for S6B, and the corresponding efficient curves generated by applying the appropriate cut to the triggers. The top pair of figures correspond to spherical harmonic energy, while the centre pair show the same for null energy. The left most figures relate to SGLs with f_c 235 Hz and Q=3, while the right most figures show SGL with f_c = 853 Hz and Q=3. The bottom pair of plots show the efficiency curves for both the null (in blue) and spherical harmonic (in green) cuts.

same good separation between loud injection triggers (plotted in red) and background triggers

(plotted in black). While the efficiency curves persist in showing that the null cut does a better job at lower frequencies. Generally the plots in Figure E.7 are similar to those seen for S6A

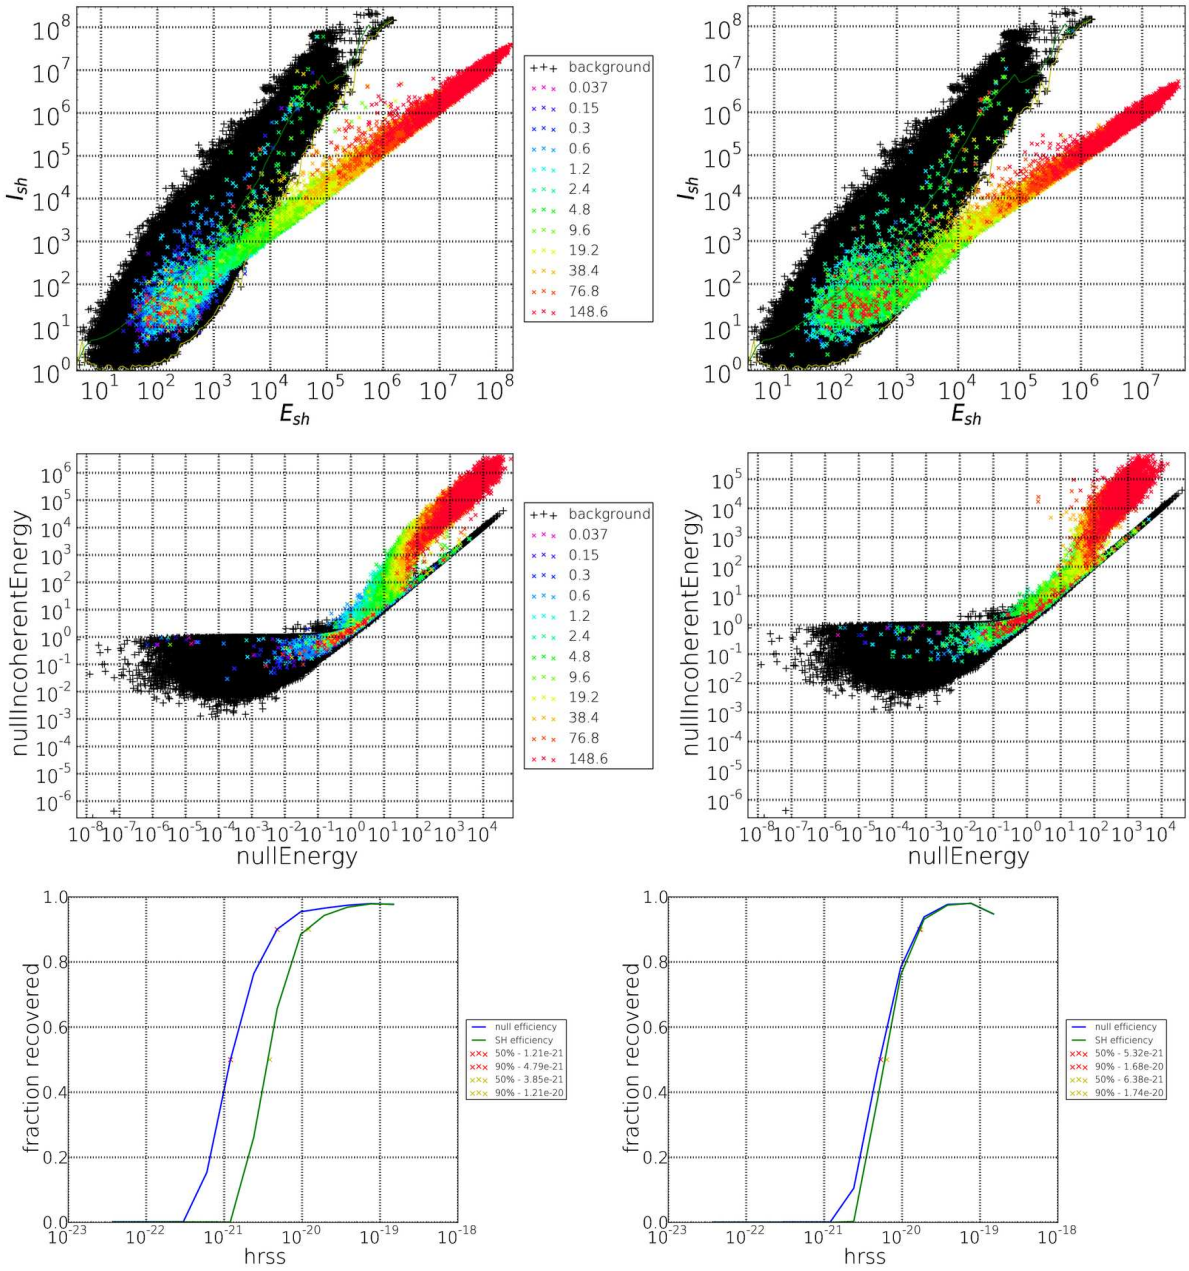

FIGURE E.7: This figure shows the distribution of triggers in the null and spherical harmonic plane for S6D, and the corresponding efficient curves generated by applying the appropriate cut to the triggers. The top pair of figures correspond to spherical harmonic energy, while the centre pair show the same for null energy. The left most figures relate to SGLs with f_c 235 Hz and Q=9, while the right most figures show SGL with $f_c = 853$ Hz and Q=9. The bottom pair of plots show the efficiency curves for both the null (in blue) and spherical harmonic (in green) cuts.

(Figure 5.14), although the efficiency curve based on the null cut appears to be pulling away from the curve based on the spherical harmonic cut. This could be a quirk with the tuning, where the spherical harmonic cut is tighter than it needs to be.

E.1.3.3 Behaviour of circularly polarised sine-Gaussians

The following figures demonstrate the distribution of triggers for injected circularly polarised sine-Gaussians in both the null and spherical harmonic plane. The plots in Figure E.8 show a

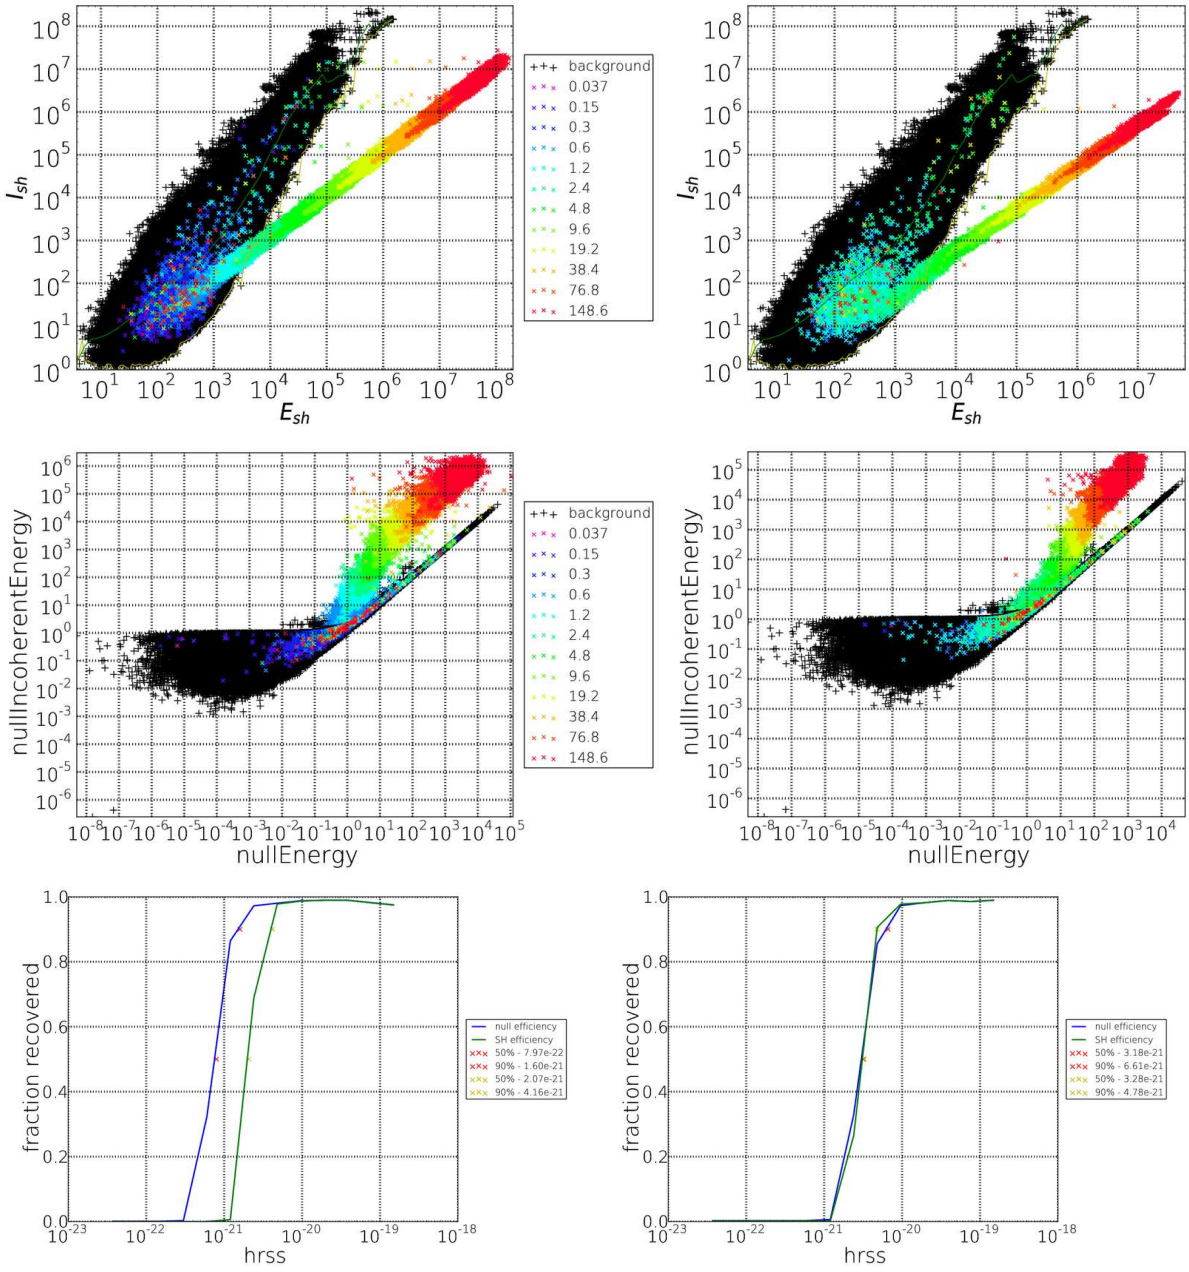

FIGURE E.8: The plots of this figure show the distribution of triggers in the null and spherical harmonic plane for S6D, and the corresponding efficient curves generated by applying the appropriate cut to the triggers. The top pair of figures correspond to spherical harmonic energy, while the centre pair show the same for null energy. The left most figures relate to SGCs with f_c 235 Hz and Q=9, while the right most figures show SGCs with f_c = 853 Hz and Q=9. The bottom pair of plots show the efficiency curves for both the null (in blue) and spherical harmonic (in green) cuts.

nice clean recovery of the circularly polarised SGs, with none of the oddness that appears to be affecting the other SGs.

E.1.4 Long duration simulations

The figures in the following sections demonstrate the distribution of triggers from the different injection types in the null and spherical harmonic plane. See Tables 5.2, 5.3 and 5.4 for the parameters used for the POW, inspiral waveforms and the SGCs respectively. Note that the paired scatter plots share a common legend while the efficiency plots have their own legend giving the value of the 50% and 90% efficiency.

E.1.4.1 Behaviour of poly-oscillating waveforms

The following figures demonstrate the distribution of triggers for injected POW in both the null and spherical harmonic plane. Figure E.9 shows that while there is generally a good separation between the loud injection triggers (plotted in red) and the background triggers (plotted in black), note that there some of the loudest triggers appear to be present among the background (as seen in the S6B results as shown in Figure D.9).

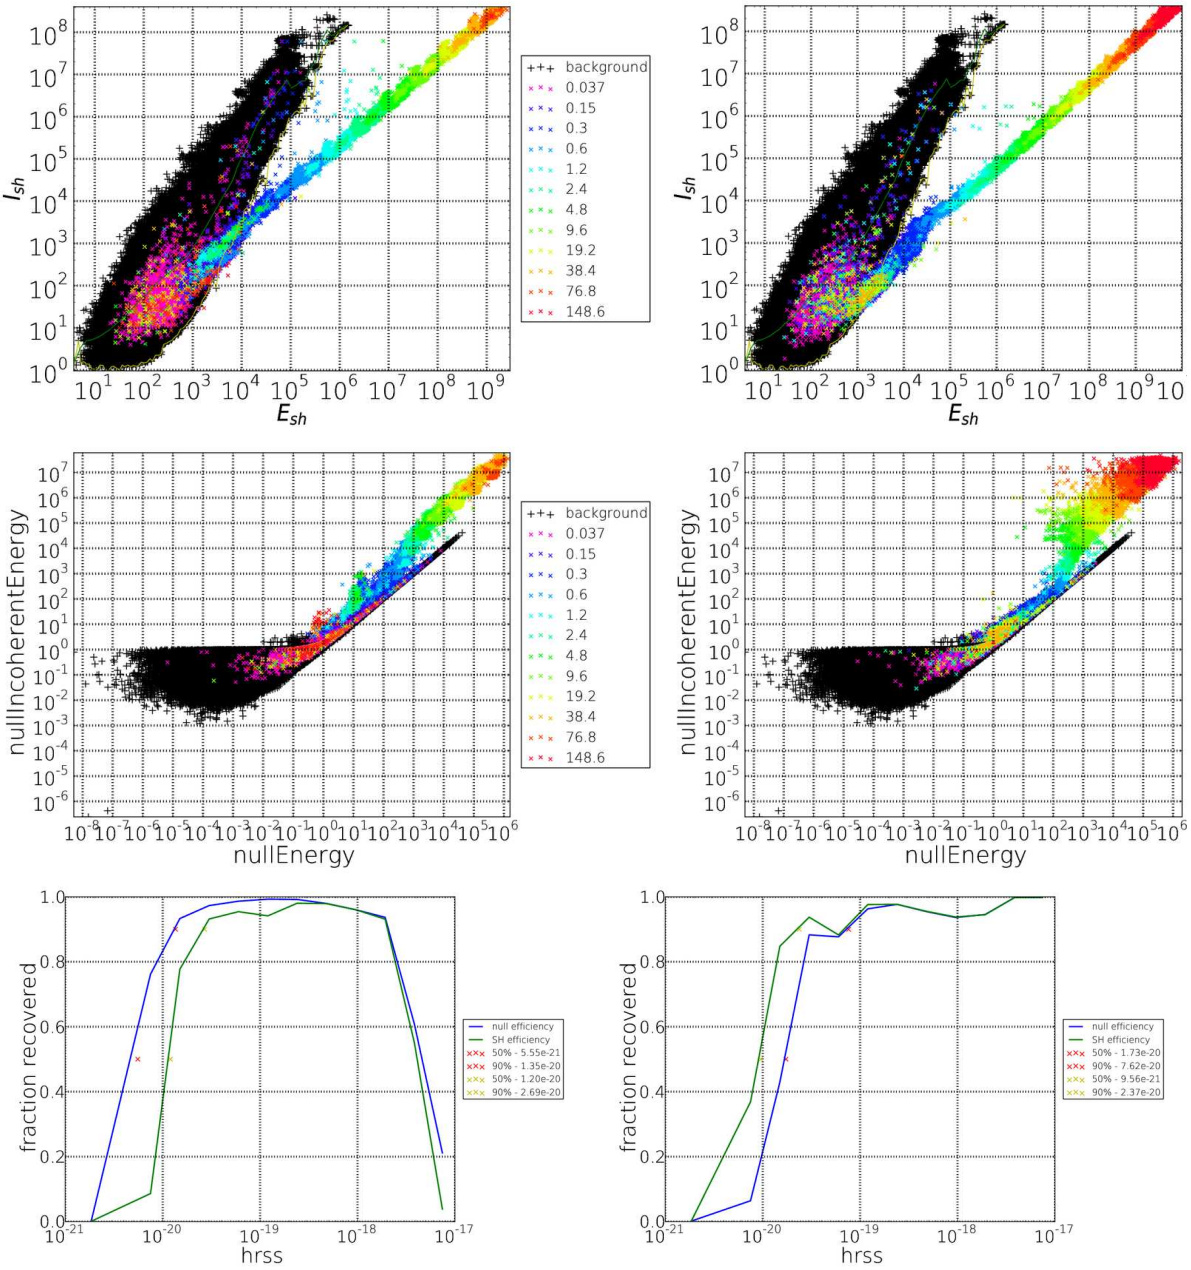

FIGURE E.9: The plots of this figure show the distribution of triggers in the null and spherical harmonic plane for S6D, and the corresponding efficient curves generated by applying the appropriate cut to the triggers for S6D. The top figures correspond to spherical harmonic energy, while the centre show the same for null energy. The left most figures relate to POWs with $f_c = 235$ Hz, $\Delta f = 100$ Hz, and $\Delta t = 12$ seconds, while the right most figures show POWs with $f_c = 852$, $\Delta f = 50$ Hz, and $\Delta t = 24$ seconds. The bottom plots show the efficiency curves for both the null (in blue) and spherical harmonic (in green) cuts.

E.1.4.2 Behaviour of Inspiral simulations

The following figures demonstrate the distribution of triggers for injected inspiral waveforms in both the null and spherical harmonic plane. Figure E.10 continues to shows that SPHRAD

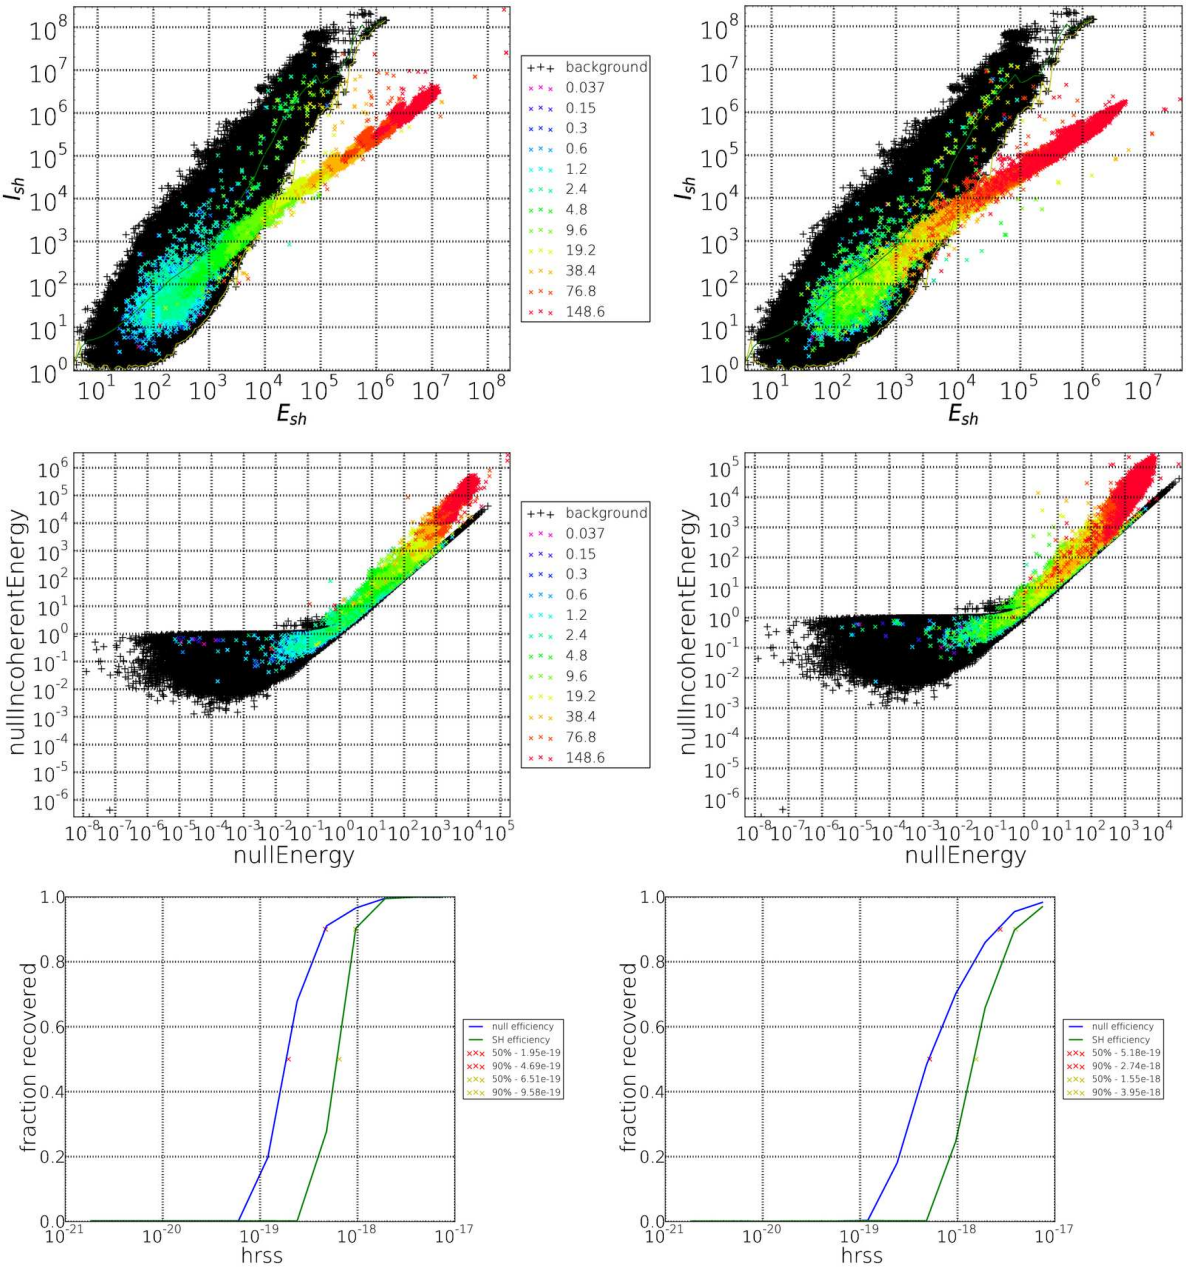

FIGURE E.10: The plots of this figure show the distribution of triggers in the null and spherical harmonic plane for S6D, and the corresponding efficient curves generated by applying the appropriate cut to the triggers. The top pair of figures correspond to spherical harmonic energy, while the centre pair show the same for null energy. The left most figures relate to inspirals with an inclination angle of $\pi/4$, while the right most figures show inspirals with an inclination angle of $\pi/2$ (corresponds to edge on). The bottom pair of plots show the efficiency curves for both the null (in blue) and spherical harmonic (in green) cuts.

struggles to recover inspirals with any great efficiency. Could be timing issues, could be something to do with the whitening issue.

E.1.4.3 Behaviour of long duration circularly polarised sine-Gaussians

The following figures demonstrate the distribution of triggers for long duration circularly polarised sine-Gaussians in both the null and spherical harmonic plane. The plots shown in

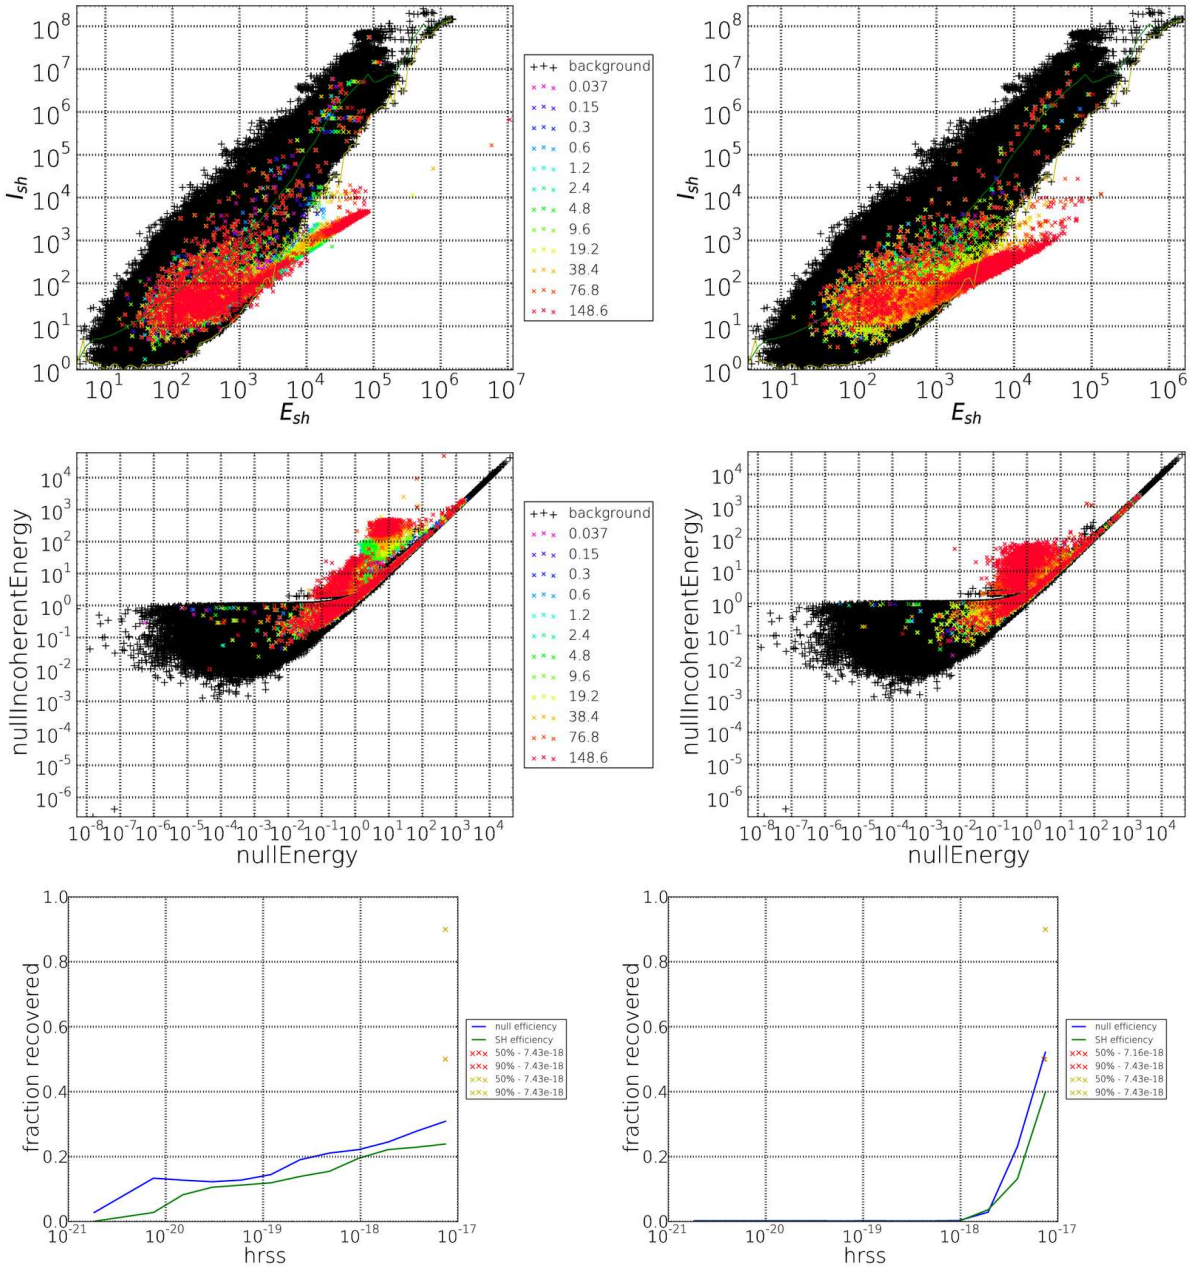

FIGURE E.11: The plots of this figure show the distribution of triggers in the null and spherical harmonic plane for S6D, and the corresponding efficient curves generated by applying the appropriate cut to the triggers. The top figures correspond to spherical harmonic energy, while the centre show the same for null energy. The left most figures relate to SGCs with f_c 235 Hz and Δt=24 seconds, while the figures on the right hand side show SGCs with f_c 512 Hz and Δt=96 seconds. The bottom plots show the efficiency curves for both the null (in blue) and spherical harmonic (in green) cuts.

Figure E.11 display the same lack of performance as Figure 5.18, emphasising the whitening issue.

E.1.4.4 Behaviour of accretion-disk instabilities

The following figures demonstrate the distribution of triggers for very long duration ($\approx$ 220 seconds) accretion disk instability waveforms in both the null and spherical harmonic plane. Note that the ADI waveform is precomputed and so does not have any free parameters. The waveform is still injected with a range of injection factors however, so we can generate the efficiency plots. Figure fig:S6Dl-ALL-ADI shows the recovery of the ADI waveform. This is

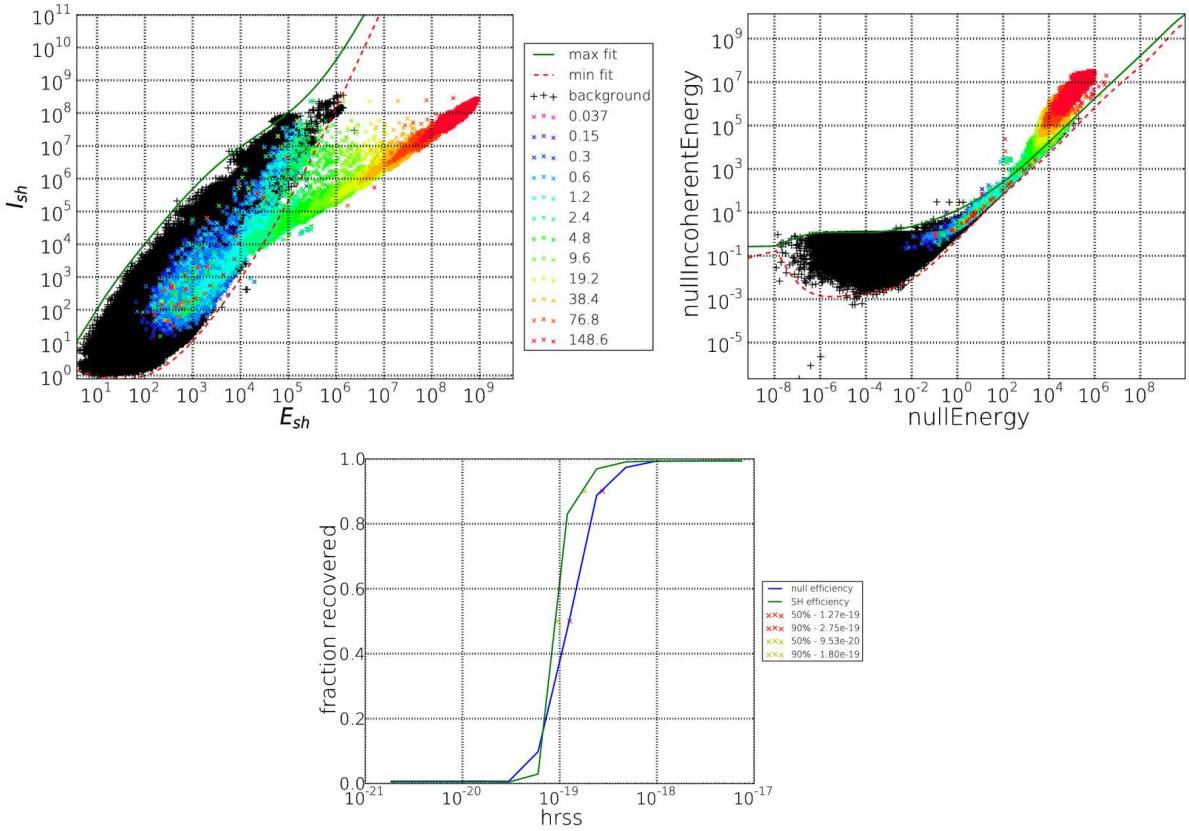

FIGURE E.12: The plots of this figure show the distribution of triggers in the null and spherical harmonic plane for S6D, and the corresponding efficient curves generated by applying the appropriate cut to the triggers. The top left figure shows the distribution for spherical harmonic energy, while the top right shows the same for null energy. The bottom plot show the efficiency curves for both the null (in blue) and spherical harmonic (in green) cuts.

a very long duration waveform and may span a couple of analysis segments, so it is good to see that the pipeline can cope with such a situation. However, there is some noticeably odd behaviour from some of the injections, especially the turning over of the efficiency curves at very high amplitudes.

Bibliography

[1] J. et al Abadie. All-sky search for gravitational-wave bursts in the second joint ligo-virgo run. *Phys. Rev. D*, 85:122007, Jun 2012. doi: 10.1103/PhysRevD.85.122007. URL http://link.aps.org/doi/10.1103/PhysRevD.85.122007.

[2] J. Aasi et al. (The LIGO Scientific Collaboration and the Virgo Collaboration). Prospects for Localization of Gravitational Wave Transients by the Advanced LIGO and Advanced Virgo Observatories. *(arXiv:1304.0670)*, 2013. URL http://xxx.lanl.gov/abs/arXiv:1304.0670.

[3] C.W. Misner, K.S. Thorne, and J.A. Wheeler. *Gravitation*. Gravitation. W. H. Freeman, 1973. ISBN 9780716703440.

[4] Bernard F. Schutz. *A first course in general relativity*. CAMBRIDGE, 1985.

[5] James B. Hartle. *Gravity - An Introduction to Einstein's General Relativity*. Addison Wesley, 2003.

[6] Michele Maggiore. *Gravitational Waves: Volume 1: Theory and Experiments*. Oxford University Press, 2008.

[7] Gregory M. Harry. Advanced LIGO: The next generation of gravitational wave detectors. *Class.Quant.Grav.*, 27:084006, 2010. doi: 10.1088/0264-9381/27/8/084006.

[8] Acernese et al. Virgo Technical Report VIR-0027A-09. 2009. URL https://tds.ego-gw.it/ql/?c=6589.

[9] Accadia et al. Virgo Document VIR-0128A-12. 2012. URL https://tds.ego-gw.it/ql/?c=8940.

[10] K. S. Thorne. *Gravitational radiation. In S. W. Hawking and W. Israel, editors, Three hundred years of gravitation*. Cambridge University Press, Cambridge, 1987.

[11] J. Abadie et al (The LIGO-Virgo Scientific Collaboration). All-sky search for gravitational-wave bursts in the second joint ligo-virgo run. *Phys. Rev. D*, 85:122007, Jun 2012. doi: 10.1103/PhysRevD.85.122007. URL http://link.aps.org/doi/10.1103/PhysRevD.85.122007.

[12] Abbott B et al. First upper limits from ligo on gravitational-wave bursts. *Phys. Rev. D*, 69:102001, 2004.

[13] Owen B. J. Maximum elastic deformations of compact stars with exotic equations of state. *Phys. Rev. Lett.*, 95:211101, 2005.

[14] Melatos A. and Payne D. J. B. Maximum elastic deformations of compact stars with exotic equations of state. *Astrophys. J.*, 623:1044, 2005.

[15] N. Andersson, Kokkotas K. D., and Stergioulas N. On the relevance of the r-mode instability for accreting neutron stars and white dwarfs. *Astrophys. J.*, 516:307–314, 1999.

[16] B. Owen et al. Gravitational waves from hot young rapidly rotating neutron stars. *Phys. Rev. D*, 58:084020, 1998.

[17] D.I. Jones and N. Andersson. Gravitational waves from freely precessing neutron stars. *Mon.Not.Roy.Astron.Soc.*, 331:203, 2002. doi: 10.1046/j.1365-8711.2002.05180.x.

[18] B. et al Abbott. Coherent searches for periodic gravitational waves from unknown isolated sources and Scorpius X-1: Results from the second LIGO science run. *Phys.Rev.*, D76: 082001, 2007. doi: 10.1103/PhysRevD.76.082001.

[19] et al. Abadie J. Directional limits on persistent gravitational waves using ligo s5 science data. *Phys. Rev. Lett.*, 107:271102, 2011.

[20] Patrick J. Sutton. A Rule of Thumb for the Detectability of Gravitational-Wave Bursts. *(arXiv:1304.0210)*, 2013.

[21] Curt Cutler and Kip S. Khorne. An overview of gravitational-wave sources. *preprint*, arXiv:gr-qc/0204090, 2002.

[22] I.F. Mirabel. Microquasars: summary and outlook. *Lect.Notes Phys.*, 794:1–15, 2010. doi: 10.1007/978-3-540-76937-8_1.

[23] S. Eikenberry, K. Matthews, M. Muno, P. Blanco, E. Morgan, et al. Faint infrared flares from the microquasar GRS 1915+105. *preprint*, arXiv.org:astro-ph/0001472, 2000.

[24] Arnon Dar and Alvaro De Rujula. A Cannonball model of gamma-ray bursts: Superluminal signatures. *Astron.Astrophys.*, 2000.

[25] Ehud B. Segalis and Amos Ori. Emission of gravitational radiation from ultrarelativistic sources. *Phys.Rev. D*, 64:064018, 2001. doi: 10.1103/PhysRevD.64.064018.

[26] Peter Meszaros. Gamma-Ray Bursts. *Rept.Prog.Phys.*, 69:2259–2322, 2006. doi: 10.1088/0034-4885/69/8/R01.

[27] Ehud Nakar. Short-Hard Gamma-Ray Bursts. *Phys.Rept.*, 442:166–236, 2007. doi: 10.1016/j.physrep.2007.02.005.

[28] S. E. Woosley. Gamma-ray bursts from stellar mass accretion disks around black holes. *Astrophys.J.*, 405:273, 1993. doi: 10.1086/172359.

[29] Bohdan Paczynski. Are gamma-ray bursts in star forming regions? *Astrophys.J.*, 494: L45, 1998. doi: 10.1086/311148.

[30] A. MacFadyen and S.E. Woosley. Collapsars: Gamma-ray bursts and explosions in 'failed supernovae'. *Astrophys.J.*, 524:262, 1999. doi: 10.1086/307790.

[31] T.J. Galama, Paul M. Vreeswijk, J. van Paradijs, C. Kouveliotou, T. Augusteijn, et al. Discovery of the peculiar supernova 1998bw in the error box of GRB 980425. *Nature*, 395: 670, 1998. doi: 10.1038/27150.

[32] J.S. Bloom, S.R. Kulkarni, Fiona A. Harrison, T. Prince, E.S. Phinney, et al. Expected characteristics of the subclass of supernova gamma-ray bursts (S-GRBs). *Astrophys.J.*, 506:L105–L108, 1998. doi: 10.1086/311655.

[33] Maurice van Putten. Gravitational radiation from a torus around a black hole. *preprint*, arXiv.org:astro-ph/010707, 2001.

[34] O. Bromberg, A. Levinson, and Maurice van Putten. The gravitational-wave spectrum of a non-axisymmetric torus around a rapidly spinning black hole. *New Astron.*, 11:619–62, 2006.

[35] Maurice van Putten et al. Gravitational radiation from gamma-ray bursts as observational opportunities for ligo and virgo. *Phys. Rev. D*, 69:044007, 2004.

[36] Nils Andersson. Gravitational waves from instabilities in relativistic stars. *Class. Quant. Grav.*, 20:R105, 2003.

[37] Nils Andersson and Kostas D. Kokkotas. The r-mode instability rotating neutron stars. *Int.J.Mod.Phys. D*, 10:381–442, 2001.

[38] Benjamin J. Owen and Lee Lindblom. Gravitational radiation from the r-mode instability. *Class.Quant.Grav.*, 19:1247–1254, 2002.

[39] A. Katrin Schenk, Phil Arras, Eanna E. Flanagan, Saul A. Teukolsky, and Ira Wasserman. Nonlinear mode coupling in rotating stars and the r mode instability in neutron stars. *Phys.Rev.*, D65:024001, 2002. doi: 10.1103/PhysRevD.65.024001.

[40] Lee Lindblom, Benjamin J. Owen, and Greg Ushomirsky. Effect of a neutron star crust on the r mode instability. *Phys.Rev.*, D62:084030, 2000. doi: 10.1103/PhysRevD.62.084030.

[41] Yanqin Wu, Christopher D. Matzner, and Phil Arras. R modes in neutron stars with crusts: Turbulent saturation, spindown, and crust melting. *Astrophys.J.*, 549:1011–1020, 2001. doi: 10.1086/319446.

[42] Lars Bildsten and Greg Ushomirsky. Viscous boundary layer damping of R modes in neutron stars. *Astrophys.J.Lett.*, 1999.

[43] P.B. Jones. Comment on 'Gravitational radiation instability in hot young neutron stars'. *Phys.Rev.Lett.*, 86:1384, 2001. doi: 10.1103/PhysRevLett.86.1384.

[44] P. B. Jones. Bulk viscosity of neutron star matter. *Phys.Rev.*, D64:084003, 2001. doi: 10.1103/PhysRevD.64.084003.

[45] Lee Lindblom and Benjamin J. Owen. Effect of hyperon bulk viscosity on neutron star r modes. *Phys.Rev.*, D65:063006, 2002. doi: 10.1103/PhysRevD.65.063006.

[46] Gregory Mendell. Magnetic effects on the viscous boundary layer damping of the R modes in neutron stars. *Phys.Rev.*, D64:044009, 2001. doi: 10.1103/PhysRevD.64.044009.

[47] Luciano Rezzolla, Frederick L. Lamb, Dragoljub Markovic, and Stuart L. Shapiro. Properties of r modes in rotating magnetic neutron stars. 1. Kinematic secular effects and magnetic evolution equations. *Phys.Rev.D*, 2001.

[48] Luciano Rezzolla, Frederick L. Lamb, Dragoljub Markovic, and Stuart L. Shapiro. Properties of r modes in rotating magnetic neutron stars. 2. Evolution of the r modes and stellar magnetic field. *Phys.Rev.D*, 2001.

[49] Sutton P., Jones G., Chatterji S., Kalmus P., Leonor I., Poprocki S., Rollins J., Searle A., Stein L., Tinto M., and Was M. X-pipeline: An analysis package for autonomous gravitational-wave burst searches. *New J. Phys.*, 12:053034, 2010.

[50] Sylvestre J. Time-frequency detection algorithm for gravitational wave bursts. *Phys. Rev. D*, 66:102004, 2002.

[51] Anderson W. G., Brady P. R., Creighton J. D. E., and Flanagan E .E. Excess power statistic for detection of burst sources of gravitational radiation. *Phys. Rev. D*, 63:042003, 2001.

[52] Klimenko S., Yakushin I., Mercer A., and Mitselmakher G. Coherent method for detection of gravitational wave bursts. *Class. Quant. Grav.*, 25:114029, 2008.

[53] S. Chatterji, L. Blackburn, G. Martin, and E. Katsavounidis. Multiresolution techniques for the detection of gravitational-wave bursts. *Class.Quant.Grav.*, 21:S1809–S1818, 2004. doi: 10.1088/0264-9381/21/20/024.

[54] M. Rakhmanov. Rank deficiency and tikhonov regularization in the inverse problem for gravitational-wave bursts. *Class. Quant. Grav.*, 21:S673–S685, 2006.

[55] C. W. Helstrom. *Statistical Theory of Signal Detection*. Permagon, London, 1968.

[56] Bruce Allen, Warren G. Anderson, Patrick R. Brady, Duncan A. Brown, and Jolien D.E. Creighton. FINDCHIRP: An Algorithm for detection of gravitational waves from inspiraling compact binaries. *Phys.Rev.*, D85:122006, 2012. doi: 10.1103/PhysRevD.85.122006.

[57] B. S. Sathyaprakash and S. V. Dhurandhar. Choice of filters for the detection of gravitational waves from coalescing binaries. *Phys. Rev. D*, 44:3819–3834, Dec 1991. doi: 10.1103/PhysRevD.44.3819. URL `http://link.aps.org/doi/10.1103/PhysRevD.44.3819`.

[58] R. Balasubramanian, B. S. Sathyaprakash, and S. V. Dhurandhar. Gravitational waves from coalescing binaries: Detection strategies and monte carlo estimation of parameters. *Phys. Rev. D*, 53:3033–3055, Mar 1996. doi: 10.1103/PhysRevD.53.3033. URL `http://link.aps.org/doi/10.1103/PhysRevD.53.3033`.

[59] Benjamin J. Owen. Search templates for gravitational waves from inspiraling binaries: Choice of template spacing. *Phys. Rev. D*, 53:6749–6761, Jun 1996. doi: 10.1103/PhysRevD.53.6749. URL `http://link.aps.org/doi/10.1103/PhysRevD.53.6749`.

[60] Benjamin J. Owen and B. S. Sathyaprakash. Matched filtering of gravitational waves from inspiraling compact binaries: Computational cost and template placement. *Phys. Rev. D*, 60:022002, Jun 1999. doi: 10.1103/PhysRevD.60.022002. URL `http://link.aps.org/doi/10.1103/PhysRevD.60.022002`.

[61] T. Cokelaer. Gravitational waves from inspiralling compact binaries: Hexagonal template placement and its efficiency in detecting physical signals. *Phys. Rev. D*, 76:102004, Nov 2007. doi: 10.1103/PhysRevD.76.102004. URL `http://link.aps.org/doi/10.1103/PhysRevD.76.102004`.

[62] S. Chatterji et al. Coherent network analysis technique for discriminating gravitational-wave bursts from instrument noise. *Phys. Rev. D*, 74:082005, 2006.

[63] Klimenko K., Mohanty S., Rakhmanov M., and Mitselmakher G. Constraint likelihood analysis for a network of gravitational wave detectors. *Phys. Rev. D*, 72:122002, 2005.

[64] Thrane E., Kandhasamy S., Ott C. D., Anderson W. G., Christensen N. L., Coughlin M. W., Dorsher S., Giampanis S., Mandic V., Mytidis A., Prestegard T., Raffai P., , and Whiting B. Long gravitational-wave transients and associated detection strategies for a network of terrestrial interferometers. *Phys. Rev. D*, 83:083004, 2011.

[65] J Abadie et al. All-sky search for gravitational wave bursts in the first joint ligo-geo-virgo run. *Physical Review D*, 81:10200, 2010.

[66] Kipp C. Cannon. An efficient algorithm for computing the time resolved full sky cross power in an interferometer with omnidirectional elements. *Phys. Rev. D*, 75:123003, 2007.

[67] Mark Edwards and Patrick Sutton. A glitch rejection method forged in spherical harmonic coefficients. *Journal of Physics: Conference Series*, 363(1):012025, 2012. URL `http://stacks.iop.org/1742-6596/363/i=1/a=012025`.

[68] Fast fourier analysis on groups, 2008. URL `www.cs.dartmouth.edu/~geelong/sphere`.

[69] R.C. Gonzalez and P Wintz. *Digital Image Processing*. Addison Wesley, 2003.

[70] Driscoll J R and Healy D. Computing fourier transforms and convolutions on the 2-sphere. *Adv. in Appl. Math.*, 15:202 – 250, 1994.

[71] Healy Jr. D., Rockmore D., Kostelec P., and Moore S. Ffts for the 2-sphere - improvements and variations. *The Journal of Fourier Analysis and Applications*, 9:4: 341 – 385, 2003.

[72] Stephen Fairhurst. Source localization with an advanced gravitational wave detector network. *Class. Quantum Grav.*, 28:105021, 2011.

[73] S. Klimenko et al. Localization of gravitational wave sources with networks of advanced detectors. *Phys. Rev. D*, 83:102001, 2011.

[74] Abbott B et al. Search for gravitational-wave bursts in ligo data from the fourth science run. *Class. Quant. Grav.*, 24:5343–5369, 2007.

[75] J. Aasi et al. Search for gravitational waves from binary black hole inspiral, merger, and ringdown in ligo-virgo data from 2009˘2010. *Phys. Rev. D*, 87:022002, Jan 2013. doi: 10.1103/PhysRevD.87.022002. URL `http://link.aps.org/doi/10.1103/PhysRevD.87.022002`.

[76] J. Abadie et al. Search for gravitational waves associated with gamma-ray bursts during LIGO science run 6 and Virgo science runs 2 and 3. *Astrophys.J.*, 760:12, 2012. doi: 10.1088/0004-637X/760/1/12.

[77] Christian D Ott. The gravitational-wave signature of core-collapse supernovae. *Class. Quantum Grav.*, 26(6):063001, 2009. URL `http://stacks.iop.org/0264-9381/26/i=6/a=063001`.

[78] J. Abadie et al. (The LIGO-Virgo Scientific Collaboration). Search for Gravitational Waves from Low Mass Compact Binary Coalescence in LIGO's Sixth Science Run and Virgo's Science Runs 2 and 3. *Phys. Rev. D*, 85:082002, 2012. URL `http://xxx.lanl.gov/abs/arXiv:1111.7314`.

[79] Maurice H. P. M. van Putten. Proposed source of gravitational radiation from a torus around a black hole. *Phys. Rev. Lett.*, 87:091101, Aug 2001. doi: 10.1103/PhysRevLett.87.091101. URL `http://link.aps.org/doi/10.1103/PhysRevLett.87.091101`.