Title

Black Holes as Precision Probes of Dark Sectors and Fundamental Physics

ISBN

9798190916348

Author

Ewasiuk, Christopher

Publication Date

2026-08-03

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA
SANTA CRUZ

BLACK HOLES AS PRECISION PROBES OF DARK SECTORS AND FUNDAMENTAL PHYSICS

A dissertation submitted in partial satisfaction of the
requirements for the degree of

DOCTOR OF PHILOSOPHY

in

PHYSICS

by

Christopher Ewasiuk

August 2026

The Dissertation of Christopher Ewasiuk
is approved:

Professor Stefano Profumo, Chair

Professor Stefania Gori

Professor Wolfgang Altmannshofer

Donald Smith
Interim Vice Provost and Dean of Graduate Studies

Contents

List of Figures

List of Tables

Abstract

Black Holes as Precision Probes of Dark Sectors and Fundamental Physics

by

Christopher Ewasiuk

This thesis investigates the observational and phenomenological consequences of large gravitationally coupled dark sectors, using black holes and high-frequency gravitational waves as the primary probes. The work spans four interconnected studies, progressing from broad model-independent constraints on dark sector size, through the physics of primordial black holes as a dark matter candidate, to a detailed re-examination of black hole evaporation dynamics, and culminating in a quantitative assessment of high-frequency gravitational-wave signals from superradiant boson clouds.

The first study derives comprehensive constraints on the number of dark degrees of freedom N from multiple astrophysical and collider observables. Hawking evaporation of stellar-mass black holes observed by LIGO/Virgo provides the dominant bound: using the merger event GW190924_021846, with chirp mass $5.7\,M_\odot$ and inferred age $(1.16\pm0.04)\times10^{10}$ yr, we constrain $N \lesssim 5\times10^{58}$ for massless spin-0 dark sectors and $N \lesssim 2.5\times10^{60}$ for spin-2, using full greybody-factor calculations with the BlackHawk code. Complementary limits from ultrahigh-energy cosmic-ray collisions ($N \lesssim 10^{71}$), LHC monojet searches ($N \lesssim$ few $\times\ 10^{62}$), and supernova cooling ($N \lesssim 6\times10^{68}$) are derived via detailed cross-section calculations, confirming that black hole evaporation

provides the strongest quasi-model-independent constraint on gravitationally coupled dark sectors across a wide range of dark sector particle masses.

The second study examines primordial black holes (PBHs) in the asteroid-mass window (10^{-16}–$10^{-11}\,M_\odot$), where they can constitute the entirety of the cosmological dark matter and are inaccessible to conventional gravitational-wave detectors. We update the calculation of the gravitational-wave sight distance for narrow-band resonant microwave cavities such as ADMX, showing that it becomes chirp-mass independent at large PBH masses, saturating at $d_{\rm sens} \simeq 0.03$ AU. Constructing the most optimistic possible merger event rate— incorporating early-time clustering, local dark matter overdensities, dichromatic mass functions, and $f_{\rm PBH} = 1$—we find that detectable merger events are extremely rare even under maximally favorable assumptions, motivating the search for alternative high-frequency gravitational-wave sources from PBH systems.

The third study extends the theory of black hole evaporation to Kerr and Reissner–Nordström geometries in the presence of large dark sectors, using the Page factor formalism. We demonstrate that the canonical evaporation hierarchy—charge neutralized first, followed by spin loss, then mass—can be qualitatively reversed when the number of uncharged dark sector species exceeds a critical threshold $N_{\rm crit} \approx 220$ (for a 10^{17} g black hole). Above this threshold, the mass-loss rate outpaces the charge-loss rate, causing the effective charge $Q_* = Q/M$ to grow during evaporation and potentially driving the black hole toward a near-extremal state. We further characterize the competing roles of Schwinger pair production, which restores rapid discharge for sufficiently light and highly charged black holes, and superradiance, which can extract angular momentum 10^8 times faster than Hawking radiation for optimal boson–black hole mass combinations.

The fourth study treats superradiance as a gravitational-wave source rather than merely a spin-loss mechanism, presenting a unified framework for gravitational-wave emission from both isolated and binary-perturbed gravitational atoms around primordial black holes. Ultralight bosons with masses $\mu \sim 10^{-8}$–10^{-5} eV form macroscopic coherent clouds around PBHs in the mass range 10^{-6}–$10^{-3}\,M_\odot$, emitting quasi-monochromatic gravitational waves in the MHz–GHz band at frequencies $f_{\rm ann} \simeq \mu/\pi$ and $f_{\rm tr} \simeq \mu\alpha^2/(4\pi)$, directly accessible to ADMX-like resonant cavities. We derive analytic frequency-domain strain templates for both the annihilation and level-transition channels, and apply the Landau–Zener formalism to binary-induced resonant transitions. Although binary-driven bursts satisfy detector ring-up criteria over a range of parameters, their characteristic strain falls orders of magnitude below current ADMX sensitivity at astrophysically plausible distances ($\gtrsim$ 1 kpc), with detectable events requiring source distances $\lesssim$ 1 AU. We quantify the improvements in strain sensitivity, bandwidth, and ring-up time required to render these signals observable, and identify PBH gravitational atoms as a theoretically well-motivated target for future high-frequency gravitational-wave searches. Taken together, these studies establish black holes as stellar-mass merger remnants, primordial dark matter candidates, and hosts of superradiant boson clouds—as uniquely powerful instruments for probing physics beyond the Standard Model through the gravitational window.

To my parents and brother, John, Glenda, and Jon

my partner Mana,

and to Stefano,

for believing in this work before I did.

Acknowledgments

First and foremost, I want to thank my advisor Stefano Profumo for his guidance, generosity, and genuine love of physics; qualities that made this work possible and made the journey worthwhile. I am grateful to my committee: Stefania Gori, Wolfgang Altmannshofer, and Jason Prochaska who pushed this work to be better than it would otherwise have been. Graduate school is not done alone, and I am lucky to have shared it with Aidan Morson, Phillip Popp, and Sam Christensen. Their friendship, humor, and solidarity made my time here pleasant, in addition to being productive. To my parents John and Glenda, and my brother Jon, thank you for your love and support throughout all of this. And to Mana, who has been my anchor from the beginning to end. Thank you for everything.

This work was supported in part by the U.S. Department of Energy under grant DE-SC0010107.

Chapter 1

Introduction

1.1 The Dark Matter Problem and the Case for a Dark Sector

The evidence for dark matter is overwhelming and multi-scale. Anomalously flat galactic rotation curves, the dynamics of galaxy clusters, the angular power spectrum of the cosmic microwave background (CMB), and the matter power spectrum of large-scale structure all point to the same conclusion: roughly 27% of the universe's energy budget is carried by matter that interacts gravitationally yet emits no detectable electromagnetic radiation [169]. Despite decades of increasingly refined searches—direct-detection experiments deep underground, indirect searches for annihilation and decay products across the electromagnetic spectrum, and dedicated production channels at the Large Hadron Collider—the microscopic identity of dark matter remains one of the most pressing open questions in fundamental physics.

The dominant paradigm for the past four decades has been the weakly interacting massive particle (WIMP), motivated by the remarkable coincidence that a particle with weak-scale mass and coupling automatically freezes out of thermal equilibrium with a relic abundance close to the observed dark matter density, a fact sometimes called the "WIMP miracle" [18]. The progressive exclusion of large regions of WIMP parameter space has broadened the theoretical landscape considerably. Among the well-motivated alternatives are ultralight axion-like particles [184, 4, 77], primordial black holes [96], and entirely gravitationally coupled *dark sectors*—the central subject of this thesis.

A dark sector is an extension of the Standard Model (SM) field content populated by degrees of freedom that carry no SM gauge charges and couple to the visible sector exclusively through gravity, or through an additional portal that is itself very weakly coupled. Such sectors arise naturally in a wide variety of theoretical frameworks. Twin Higgs models [56] contain a mirror sector of roughly SM-sized field content. The minimal supersymmetric SM introduces a superpartner for every SM degree of freedom [161]. Scenarios with large extra dimensions [20] predict a Kaluza–Klein tower of graviton modes whose multiplicity scales as $(M_*/M_P)^{d-4}$, potentially reaching astronomical numbers in the four-dimensional effective theory. "Clockwork" [115] and "stasis" [74] cosmological scenarios similarly contain large numbers of approximately degenerate species.

Even without any non-gravitational portal, a large number N of light dark-sector degrees of freedom modifies Newton's constant through radiative corrections,

$$\delta G^{-1}(\mu) = G^{-1}(\mu) - G^{-1}(0) \sim N m^2 \ln\left(\mu^2/m^2\right), \qquad m < \mu, \tag{1.1}$$

where m is the dark-sector particle mass and μ the energy scale [47]. For massive species with $m < \mu$, equation (1.1) governs the running; in the strictly massless limit the m^2 prefactor is replaced by μ^2, and the correction persists as a pure logarithmic running. The bound quoted below for massless species follows from this latter regime, as derived in detail in Chapter 2. Precision equivalence-principle tests at distances $d \sim 10^3$ km constrain $\delta G/G \lesssim 10^{-9}$ [101, 195], which translates at the scale $\mu \sim 1/d \sim 10^{-13}$ eV into the weak bound $N \lesssim 10^{73}$ for massless species. These renormalization-group limits are examined carefully in Chapter 2, but they carry important caveats: gauge bosons contribute with the opposite sign to scalars and Weyl fermions, and large cancellations involving the cosmological constant are not excluded. Qualitatively different, and observationally far more powerful, are constraints that arise from the direct effects of a large dark sector on macroscopic astrophysical objects—in particular, on the Hawking evaporation of stellar-mass and primordial black holes, and on the gravitational-wave signals those objects produce. It is these constraints, and the observational signatures they generate, that form the subject of this thesis.

1.2 Organization of This Thesis

This thesis is organized as a compilation of four research papers, each of which constitutes a self-contained chapter. The chapters are arranged to follow a coherent scientific narrative, progressing from precision laboratory constraints on large dark sectors through the phenomenology of primordial black holes as dark matter, then deepening into the theoretical structure of black hole evaporation in generic geometries, and culminating in a quantitative detectability study of high-frequency gravitational-wave signals from superradiant boson clouds around primordial black holes.

Chapter 2, "Precision gravity constraints on large dark sectors" (Ewasiuk & Profumo, arXiv:2509.02801 [97]), establishes a model-independent framework for bounding the number of gravitationally coupled dark degrees of freedom using precision tests of Newton's law. By treating general relativity as a low-energy effective field theory and translating loop-induced quantum corrections to the Newtonian potential into the Yukawa-type language of fifth-force experiments, the chapter derives quantitative multiplicity limits that reach $N \lesssim \mathcal{O}(10^{61})$ for truly massless species.

Chapter 3, "The maximal gravitational wave signal from asteroid-mass primordial black hole mergers at resonant microwave cavities" (Profumo, Brown, Ewasiuk, Ricarte & Su, arXiv:2410.15400 [187]), narrows the focus to primordial black holes as a specific dark matter candidate. The chapter updates the calculation of the gravitational-wave sight distance for narrow-band resonant-cavity detectors, derives the asymptotic chirp-mass independence of that distance in the $N_c/Q \ll 1$ limit, and constructs the maximum possible merger event rate under optimistic assumptions about PBH clustering, local dark matter density, and mass function shape. The conclusion that even the most favorable event rates fall far below current sensitivity motivates the search for alternative high-frequency gravitational-wave sources from the same PBH population.

Chapter 4, "Dark-sector modifications to Kerr and Reissner–Nordström black hole evaporation" (Ewasiuk & Profumo, arXiv:2505.04812 [98]), provides the theoretical foundation for understanding how a large dark sector alters black hole evaporation dynamics in generic geometries. Using the Page factor formalism extended to Kerr and Reissner–Nordström black holes, the chapter demonstrates that the canonical charge-first evaporation hierarchy can be reversed when the number of uncharged dark-sector species exceeds a critical threshold $N_{\mathrm{crit}} \approx 220$, producing effective charge growth and potentially near-extremal relics. The chapter also characterizes Schwinger pair production and superradiance as competing mechanisms that further modify the evaporation dynamics, providing the theoretical bridge to the final chapter.

Chapter 5, "High-frequency gravitational wave transients from superradiance" (Su, Brown, Ewasiuk & Profumo, arXiv:2604.01407 [203]), completes the program by treating superradiance as a gravitational-wave source rather than merely a spin-loss mechanism. Drawing on the PBH mass range and resonant-cavity detector framework developed in Chapter 3 and the superradiance formalism introduced in Chapter 4, the chapter derives analytic waveform templates for both isolated and binary-perturbed gravitational atoms, performs a rigorous detectability analysis at ADMX sensitivity, and quantifies the detector improvements needed to access this class of signals.

1.3 Precision Gravity Constraints on Dark Sectors

General relativity, treated as a low-energy effective field theory valid below the Planck scale, predicts calculable quantum corrections to Newton's law arising from loops of matter and graviton fields. For the Standard Model particle content these corrections are utterly negligible at laboratory distances. The situation changes qualitatively, however, when the theory is extended by a hidden sector containing a large number N of light degrees of freedom: loop-induced modifications to the Newtonian potential accumulate with N, and for sufficiently large dark sectors they can approach the sensitivity thresholds of current precision gravity experiments.

The leading quantum correction to the static potential takes different forms depending on the mass of the dark-sector particle. Massless species generate a long-range, nonanalytic $1/r^3$ power-law correction whose coefficient grows linearly with N. Massive species with mass m_X generate Yukawa-suppressed contributions with characteristic range $\lambda = 1/(2m_X)$ that collectively scale as $N\beta_X e^{-2m_X r}/r^2$, where β_X is a spin-dependent loop coefficient. Both contributions can be mapped onto the standard experimental Yukawa template

$$U_{\exp}(r) = -\frac{Gm_1 m_2}{r}\left[1 + \gamma\, e^{-r/\lambda}\right], \tag{1.2}$$

making precision fifth-force exclusion curves directly reinterpretable as bounds on the dark-sector multiplicity.

Chapter 2 carries out this program systematically, deriving the effective field theory predictions for loop-induced corrections and translating published inverse-square-law and Yukawa-type exclusion data into multiplicity bounds. Using the most precise available torsion-balance data [204], the per-spin limits for truly massless species are

$$N_s \lesssim 1.08 \times 10^{62} \quad \text{(real scalar)}, \tag{1.3}$$

$$N_f \lesssim 1.80 \times 10^{61} \quad \text{(Dirac fermion)}, \tag{1.4}$$

$$N_v \lesssim 9.02 \times 10^{60} \quad \text{(vector)}. \tag{1.5}$$

For massive particles the analysis provides a continuous exclusion curve in the (m, N) plane, with torsion-balance experiments dominating at masses $m \sim 10^{-16}$–10^{-13} eV and planetary dynamics and lunar laser ranging providing sensitivity at smaller masses. These constraints are complementary to cosmological bounds from BBN and the CMB, which limit relic abundances under specified production histories, whereas laboratory tests constrain the present-day spectrum of light states irrespective of their cosmological population.

1.4 Primordial Black Holes as Dark Matter: The Asteroid-Mass Window and the HFGW Detection Program

Among the scenarios in which black holes illuminate the dark sector, perhaps the most striking is the possibility that black holes themselves *are* the dark matter. Primordial black holes (PBHs), formed from the gravitational collapse of large-amplitude density perturbations in the early universe at the scale of the cosmic event horizon, are a well-motivated dark matter candidate whose mass spectrum depends sensitively on the primordial power spectrum and the thermal history at the time of formation [52, 96].

A well-known observational "asteroid-mass window" spanning roughly $10^{-16}\,M_{\odot}$ to $10^{-11}\,M_{\odot}$ permits PBHs to constitute the *entirety* of the cosmological dark matter abundance ($f_{\rm PBH} = 1$) without violating any existing bound. The lower edge of this window is set by Hawking evaporation: PBHs lighter than $\sim 10^{-16}\,M_{\odot}$ would have evaporated before the present epoch, with products constrained by extragalactic gamma-ray backgrounds and CMB spectral distortions [52]. The upper edge near $10^{-11}\,M_{\odot}$ is set by gravitational microlensing surveys [172]. Crucially, the entire five-decade-wide window is inaccessible to any currently planned gravitational-wave interferometer—LIGO, Virgo, KAGRA, and LISA all target mass ranges well outside it—making a direct detection of asteroid-mass PBHs one of the defining challenges for next-generation gravitational-wave astronomy.

High-frequency gravitational waves (HFGWs) offer a rare window into this mass range. A merger of two black holes with chirp mass $\mathcal{M}_c$ emits gravitational radiation at an instantaneous frequency $\nu \simeq 2/(6^{3/2}\pi M_{\rm tot})$ at the innermost stable circular orbit, which for $M_{\rm tot} \sim 10^{-11}\,M_{\odot}$ corresponds to $\nu \sim 1$ GHz. Resonant microwave cavities—narrow-band detectors originally designed for axion dark matter searches [199], sensitive to oscillatory metric perturbations via the inverse Gertsenshtein effect [178]—naturally operate in this frequency range. The strain sensitivity of an ADMX-like cavity is [105]

$$h_{\rm sens}(\nu) \simeq 3\times 10^{-22}\left(\frac{0.1}{\eta_n}\right)\left(\frac{8\ \rm T}{|\vec{B}|}\right)\left(\frac{0.1\ \rm m^3}{V_{\rm cav}}\right)^{5/6}\left(\frac{10^5}{Q}\right)^{1/2} \\ \times\left(\frac{T_{\rm sys}}{1\ \rm K}\right)^{1/2}\left(\frac{1\ \rm GHz}{\nu}\right)^{3/2}\left(\frac{\Delta\nu}{10\ \rm kHz}\right)^{1/4}\left(\frac{1\ \rm min}{\Delta t}\right)^{1/4}, \tag{1.6}$$

where $\eta_n \approx 0.1$ is the effective coupling coefficient for the TM_{010} mode, Q the quality factor, and $\Delta\nu \sim \nu/Q$ the bandwidth.

Chapter 3 updates the calculation of the sight distance for such detectors and derives the maximal possible event rate for PBH mergers at an ADMX-like cavity. A key technical result is that in the regime where the number of gravitational-wave cycles within the cavity bandwidth satisfies $N_c/Q \ll 1$, the cavity signal becomes independent of chirp mass; the sight distance therefore saturates at

$$d_{\rm sens} \simeq 0.03\ \mathrm{AU}\left(\frac{\Delta\nu}{10\ \rm kHz}\right)^{1/4}. \tag{1.7}$$

The chapter then constructs the most optimistic possible event rate, simultaneously incorporating early-time PBH clustering that enhances three-body binary formation, a local dark matter over-density, a dichromatic PBH mass function of the form $\psi(m) = f_1\,\delta(m-m_1)+f_2\,\delta(m-m_2)$ constrained by $\sum_i f_i/f_{\rm PBH}^{\rm max}(m_i) \le 1$, and a PBH abundance saturating the cosmological dark matter density. Even under these maximally optimistic conditions, detectable merger events are extraordinarily rare: the $d_{\rm sens} \simeq 0.03$ AU sight distance falls far below any characteristic distance implied by realistic PBH binary formation rates, with event rates below $10^{-17}\ \mathrm{yr}^{-1}$ even for the most favorable future detector configurations. This conclusion motivates the search for alternative sources of GHz gravitational waves from the same PBH population, which is the subject of

Chapter 5. But first, a detailed understanding of how a large dark sector alters the *dynamics* of black hole evaporation for generic geometries is necessary, which is the subject of Chapter 4.

1.5 Dark-Sector Modifications to Black Hole Evaporation Hierarchies in Kerr and Reissner–Nordström Geometries

The analysis of Chapter 2 treats dark-sector effects through their cumulative loop contributions to the Newtonian potential. A qualitatively different and complementary probe of large dark sectors arises from black hole evaporation itself: because black holes radiate every species lighter than their Hawking temperature regardless of gauge charge, they are sensitive to the full particle content of the theory in a way that laboratory or collider experiments are not. A generic black hole is characterized by three conserved quantities, mass M, angular momentum J, and charge Q. These quantities are dissipated on very different timescales that depend sensitively on the available particle spectrum.

Page's classic analysis [174, 175, 176] established a canonical hierarchy for SM-only black holes: charge Q is neutralized first, angular momentum J is shed next, and mass loss dominates last. This ordering reflects the fact that the number of particles needed to neutralize a charged black hole scales as M, while shedding all angular momentum J requires of order M^2 emissions, and mass loss operates on the longest timescale of all.

Chapter 4 demonstrates that this hierarchy is not universal: it depends sensitively on the particle content of the theory and can be qualitatively reversed by a large dark sector. The central tool is the Page factor formalism, in which scale-invariant loss rates

$$f(M, a^*) = -M^2 \frac{dM}{dt}, \tag{1.8}$$

$$g(M, a^*) = -\frac{M}{a^*} \frac{dJ}{dt}, \tag{1.9}$$

$$h(M, Q^*) = -\frac{M^2}{Q^*} \frac{dQ}{dt}, \tag{1.10}$$

are computed numerically using BlackHawk for the Kerr case and a custom charged-metric adaptation for the Reissner–Nordström (RN) case [98]. The key physical insight is that introducing N_{dof} massless, neutral dark-sector species augments the mass-loss Page factor f while leaving the charge-loss factor h unchanged, since only charged particles contribute to h. The effective charge $Q^* = Q/M$ therefore evolves as

$$\frac{dQ^*}{dt} = \frac{Q^* \left[f(M, Q^*) - h(M, Q^*) \right]}{M^3}, \tag{1.11}$$

which changes sign—from charge loss to effective charge growth—when f surpasses h. The critical threshold derived in Chapter 4 is $N_{\text{crit}} \approx 220$ for a 10^{17} g black hole with massless dark-sector particles. For $N_{dof} \gtrsim 270$, the effective charge grows throughout the entire evaporation, driving the black hole asymptotically toward extremality ($Q^* \to 1$) where the Hawking temperature

$$T_{\text{RN}} = \frac{r_+ - r_-}{4\pi r_+^2} \xrightarrow{Q^* \to 1} 0 \tag{1.12}$$

and evaporation halts. This is a non-trivial prediction: large hidden sectors can produce near-extremal charged black hole relics as a population-level outcome of Hawking evaporation, independent of initial conditions.

Two additional mechanisms are incorporated in Chapter 4. Schwinger pair production [197] describes the spontaneous creation of electron-positron pairs in the near-horizon electric field and becomes dominant for light ($M \lesssim 10^{17}$ g), highly charged ($Q^* \gtrsim 0.5$) black holes, introducing a competing discharge channel that can restore the conventional hierarchy and even transiently increase Q^* before the field weakens. For the Kerr geometry, the chapter confirms the expected result that spin is always shed before mass, but demonstrates that the timescale for spin-down depends sensitively on the dark-sector particle spin: scalar fields are anomalously inefficient at carrying away angular momentum in the low-spin limit, producing characteristic divergences in the log-normalized ratio $\frac{M}{a^*}\frac{da^*}{dt}/\frac{dM}{dt}$ near the end of the spin-down phase.

The chapter further introduces superradiance as an angular-momentum extraction mechanism that is orders of magnitude more efficient than Hawking radiation for appropriate combinations of black hole spin and boson mass. The superradiant spin-down timescale

$$t_{\rm SR} = \Gamma_{n\ell m}^{-1} \approx 5 \text{ s} \left(\frac{M_{\rm BH}}{10^{-5}\, M_\odot}\right)\left(\frac{0.2}{\alpha}\right)^9\left(\frac{0.1}{C_{211}}\right), \tag{1.13}$$

where $\alpha \equiv G M_{\rm BH}\mu_a$ is the gravitational fine-structure constant and μ_a the boson mass, can be 10^8 times shorter than the Hawking spin-down timescale. This observation provides the theoretical bridge to Chapter 5, which treats superradiance not merely as a spin-loss mechanism but as a gravitational-wave source in its own right.

1.6 High-Frequency Gravitational Wave Transients from Superradiance around Primordial Black Holes

Superradiance is the amplification of bosonic waves at the expense of the rotational energy of a Kerr black hole. A bosonic field of mass μ satisfying the superradiant condition

$$0 < \omega < m\,\Omega_H, \qquad \Omega_H = \frac{a^*}{2GM_{\rm BH}\left(1+\sqrt{1-a^{*2}}\right)}, \tag{1.14}$$

is exponentially amplified by repeated scattering from the angular-momentum barrier of the Kerr potential [219, 220, 185]. When the boson's Compton wavelength $\bar{\lambda}_C = \mu^{-1}$ is comparable to the gravitational radius $r_g = GM_{\rm BH}$—i.e., when the gravitational fine-structure constant $\alpha \equiv GM_{\rm BH}\mu \sim \mathcal{O}(0.1)$—the amplified modes are gravitationally trapped and grow into a macroscopic coherent cloud, the "gravitational atom" (GA) [71, 43]. The resulting quasi-bound states carry hydrogenic quantum numbers $\{n, \ell, m\}$ with eigenfrequencies $\omega_n \simeq \mu(1-\alpha^2/(2n^2))$, and the saturation occupation number of the dominant $\{2, 1, 1\}$ state is

$$N_{\rm sat} \approx \frac{0.1\,M_{\rm BH}}{\mu} \approx 10^{75}\left(\frac{M_{\rm BH}}{10^{-5}\,M_\odot}\right)\left(\frac{10^{-6}\ \mathrm{eV}}{\mu}\right). \tag{1.15}$$

Two gravitational-wave emission channels become active once the cloud saturates. In the annihilation channel, pairs of cloud bosons convert into gravitons at frequency

$$f_{\rm ann} \simeq \frac{\mu}{\pi} \approx 484\ \mathrm{MHz} \left(\frac{\mu}{10^{-6}\ \mathrm{eV}}\right). \tag{1.16}$$

In the level-transition channel, the quadrupole self-interaction drives bosons from principal quantum number n_e to $n_g < n_e$, emitting at the Bohr frequency

$$f_{\rm tr} \simeq \frac{\mu\alpha^2}{4\pi}\left(\frac{1}{n_g^2} - \frac{1}{n_e^2}\right) \approx 121\ \mathrm{MHz} \left(\frac{\mu}{10^{-6}\ \mathrm{eV}}\right)\left(\frac{\alpha}{0.2}\right)^2 \left(\frac{1}{n_g^2} - \frac{1}{n_e^2}\right). \tag{1.17}$$

The MHz–GHz frequency window targeted by resonant-cavity detectors therefore corresponds to boson masses $\mu \sim 10^{-8}$–10^{-5} eV. For peak superradiance efficiency at $\alpha \simeq 0.2$, the required black hole mass is

$$M_{\rm BH} \simeq \frac{\alpha}{\mu G} \approx 2.7 \times 10^{-5}\, M_\odot \left(\frac{\alpha}{0.2}\right)\left(\frac{10^{-6}\ \mathrm{eV}}{\mu}\right), \tag{1.18}$$

which falls squarely in the asteroid-mass PBH window studied in Chapter 3. This is the central connection: the same PBH population that is the target of the merger-detection program is a natural host for superradiant gravitational atoms. Boson masses in the range $\mu \sim 10^{-8}$–10^{-5} eV are independently motivated by the QCD axion solution to the strong-CP problem [179, 212, 213], further linking the gravitational-wave signal to a concrete particle-physics target.

Chapter 5 presents the first unified treatment of gravitational-wave emission from both isolated and binary-perturbed gravitational atoms in the high-frequency regime, incorporating relativistic superradiant instability rates [128, 198], numerically computed cloud masses [163], the binary-induced resonance formalism of Kyriazis & Yang [147], and realistic ADMX sensitivity curves [35]. For isolated gravitational atoms, analytic closed-form frequency-domain strain templates, derived via the exponential integral E_1, are obtained for both the annihilation and level-transition channels. Peak strains for the most optimistic annihilation configurations reach $h \sim 10^{-22}$ at 1 kpc, while level-transition signals peak near $h \sim 10^{-23}$.

For PBH binaries, a companion introduces a time-varying tidal potential that drives resonant level transitions whenever the orbital frequency sweeps through the Bohr frequency of the gravitational atom. Chapter 5 applies the Landau–Zener formalism [33, 34] to the hyperfine transition $\{2, 1, 1\} \to \{2, 1, -1\}$ and derives the conditions under which the resulting transient GW burst duration exceeds the ADMX ring-up time $\tau_{\rm ring} \sim Q/\nu_0 \sim 10^{-4}$ s. While the ring-up criterion can be satisfied for a range of binary parameters, the detectability analysis reveals a fundamental strain gap: at astrophysically plausible distances ($\gtrsim 1$ kpc, set by the characteristic PBH binary formation rate from early two-body channels), the characteristic strain falls orders of magnitude below the ADMX noise floor. A detectable binary-induced burst would require the source to be at $\lesssim 1$ AU, implying an event rate far too low to be observationally relevant. The chapter quantifies the improvements in strain sensitivity, bandwidth, and ring-up time that would render these signals observable, establishing a concrete design target for future high-frequency gravitational-wave experiments.

Chapter 2

Precision Gravity Constraints on Dark Sectors

2.1 Introduction

Chapter 1 established the central motivation: a large number N of gravitationally coupled dark degrees of freedom modifies Newton's constant through radiative corrections, and black hole evaporation arguments already place the strongest quasi-model-independent bounds on N. This chapter pursues a complementary and independent constraint channel, precision laboratory gravity, that probes a qualitatively different regime of the same parameter space.

Precision "fifth-force" experiments, such as torsion-balance measurements pioneered by the Eöt-Wash group, now probe gravitational interactions down to tens of microns [204]. These experiments are typically interpreted in terms of Yukawa-type deviations from the inverse-square law, parametrized by a strength α and range λ [204, 8]. Within the low-energy effective field theory (EFT) of gravity [86], quantum loops of matter and gravitons induce calculable corrections to the Newtonian potential, with leading terms scaling as $1/r^3$ at large distances [37, 126, 100]. A single new scalar, fermion, or vector state contributes negligibly to such deviations. However, a sector containing a large multiplicity N of light degrees of freedom generates cumulative corrections that become appreciable: massless dark particles produce enhanced $1/r^3$ power-law corrections, while massive states contribute Yukawa-suppressed tails whose collective strength grows with N [45, 106].

The central aim of this chapter is to systematically connect these EFT predictions for loop-induced gravitational corrections with precision experimental limits, deriving robust bounds on the multiplicity and mass spectrum of hidden sectors coupled only gravitationally. We show how existing torsion-balance data can be reinterpreted as constraints on the number of additional massless and light massive fields, yielding quantitative limits that in some cases reach $\mathcal{O}(10^{61})$ for truly massless species, complementing the black hole evaporation bounds derived in Chapter 4 and the cosmological probes of relativistic degrees of freedom.

The chapter is organized as follows. Section 2.2 reviews the EFT framework and presents the relevant nonlocal form factors. Section 2.7 discusses massless-limit consistency conditions. Section 2.10 applies the formalism to hidden sectors with many degrees of freedom, mapping loop-induced corrections onto the Yukawa template used in experiments. Section 2.14 translates published fifth-force bounds into constraints on dark sect multiplicities. We conclude in Section 2.19 with implications for dark sect phenomenology and prospects for future improvements.

2.2 Quantum Corrections to Newton's Constant from Massive and Massless Fields in Effective Field Theory

General relativity can be viewed as a low-energy effective field theory (EFT), valid far below the Planck scale. In this approach, quantum loops of matter fields and gravitons lead to scale-dependent modifications of Newton's constant. These effects can be described either as a running gravitational coupling $G(q^2)$ in momentum space, or equivalently through nonlocal form factors multiplying curvature terms in the effective action. At low energies the universal content is carried by nonanalytic terms, such as $q^2 \ln\left(1 + q^2/m^2\right)$ for massive particles or $q^2 \ln\left(-q^2\right)$ for massless fields. The Fourier transforms of these terms produce observable corrections to the Newtonian potential, modifying the familiar $\frac{1}{r}$ force law at short or long distances depending on the particle content.

The subsections that follow outline this framework in detail, first reviewing the structure of the one-loop effective action and then showing how different particle species contribute to the running of G and to corrections of the potential.

2.3 Effective Action and Nonlocal Form Factors

Quantum corrections to gravity can be systematically treated within the one-loop effective action, viewing general relativity as a low-energy EFT. For a matter field of mass m and spin s one writes [60]

$$\Gamma_s^{(1)} \;=\; \tfrac{1}{2}(-1)^{2s}\,\mathrm{Tr}\,\ln\bigl(\Delta_s + m^2\bigr), \tag{2.1}$$

where Δ_s is the kinetic operator appropriate to spin s. The nonlocal terms of interest appear as form factors multiplying curvature invariants. It is convenient and complete to organize the result in the quadratic curvature basis (see, e.g., [31, 86, 87, 45, 107]):

$$\Gamma^{(2)}_{\mathrm{matter}} = \int d^4x\,\sqrt{-g}\Bigl[R\,F_{1,s}\Bigl(\tfrac{\Box}{m^2}\Bigr)R + R_{\mu\nu}\,F_{2,s}\Bigl(\tfrac{\Box}{m^2}\Bigr)R^{\mu\nu}\Bigr], \tag{2.2}$$

with the dimensionless argument

$$\tau \;\equiv\; -\frac{\Box}{m^2}. \tag{2.3}$$

To leading nonlocal order, each spin-s species contributes the universal shape

$$F_{i,s}(\tau) = c_{i,s}\Bigl[\,1 - \tau\ln\Bigl(1+\frac{1}{\tau}\Bigr)\Bigr], \qquad i = 1,2, \tag{2.4}$$

which resums the infrared logarithms and smoothly interpolates between the massive and massless regimes.

Note that in the background–field quantization of a (Stueckelberg-regulated) Abelian vector in a covariant gauge, the one-loop effective action at quadratic order reads schematically

$$\Gamma^{(1)}_{(1)} \;=\; \tfrac{1}{2}\,\mathrm{Tr}\ln\bigl({\Delta_\mu}^{\nu} + m^2\bigr) \;-\; \mathrm{Tr}\ln\bigl(\Delta_{\mathrm{gh}} + m^2\bigr) \;+\; \tfrac{1}{2}\,\mathrm{Tr}\ln\bigl(\Delta_{\mathrm{gf}} + m^2\bigr), \tag{2.5}$$

where the first term comes from the gauge field, the second from the complex Grassmann Faddeev–Popov ghosts (hence the overall minus sign), and the last from the gauge-fixing Jacobian. Decomposing into spin projectors and rewriting the result in the curvature basis gives

$$\Gamma^{(2)}_{\text{matter}} \supset \int \sqrt{-g}\,\Big[R\,F_{1,1}\Big(\frac{\Box}{m^2}\Big)R \;+\; R_{\mu\nu}\,F_{2,1}\Big(\frac{\Box}{m^2}\Big)R^{\mu\nu}\Big],$$

with the same nonlocal shape for both channels, $F_{i,1}(\tau) = c_{i,1}\,[\,1 - \tau\ln(1+1/\tau)\,]$, but with spin-1 coefficients $c_{i,1}$ that include the ghost/gauge-fixing contributions. The static inverse–square–law combination probed by our fits is $F_1^{(\mathrm{ISL})}(\tau) = F_{2,1}(\tau) - \frac{1}{2}F_{1,1}(\tau)$, so the net coefficient is $c_1 \equiv c_{2,1} - \frac{1}{2}c_{1,1}$. Including the Faddeev–Popov sector yields a *negative* value for this combination, $c_1 = -1/(96\pi^2)$; see, e.g., [31, 45, 107] for explicit derivations in the covariant form-factor language.

It is important to emphasize that this curvature–squared form–factor coefficient c_1 is not identical to the coefficient of the spin–1 contribution in the long–distance potential written in Eq. (I.1) of Ref. [45], where only the *physical* vector degrees of freedom contribute (counted by N_1), with the ghost sector omitted from that parametrization. After translating between the physical field–counting basis used in Eq. (I.1) and the covariant form–factor basis used here (which automatically incorporates the ghost sector via Eq. (2.5)), both approaches yield the same long–distance $1/r^3$ correction to the Newtonian potential. The negative value (2.7) therefore reflects the standard background–field definition of the ISL coefficient for spin–1 matter.

At the linearized level about flat space, the static potential probes a definite combination of the spin-2 and spin-0 channels carried by the propagator. This can be expressed as

$$F_s^{(\mathrm{ISL})}(\tau) = F_{2,s}(\tau) - \tfrac{1}{2}F_{1,s}(\tau) = c_s\Big[1 - \tau\ln\Big(1+\frac{1}{\tau}\Big)\Big], \qquad c_s \equiv c_{2,s} - \tfrac{1}{2}c_{1,s}. \tag{2.6}$$

Equation (2.6) shows explicitly that both R^2 and $R_{\mu\nu}R^{\mu\nu}$ are generated at one loop, while clarifying why a single coefficient c_s suffices for the static potential: it is the short-distance combination entering inverse–square–law (ISL) observables.

For minimally coupled matter fields, the coefficients entering the ISL combination are

$$c_0 = \frac{1}{384\,\pi^2}\ \text{(real scalar)}, \quad c_{1/2} = \frac{1}{192\,\pi^2}\ \text{(Dirac fermion)}, \quad c_1 = -\frac{1}{96\,\pi^2}\ \text{(vector+ghosts)}. \tag{2.7}$$

Summing over all species gives

$$F^{(\mathrm{ISL})}(\tau) \ = \ \sum_s N_s\, c_s\left[1 - \tau\ln\Big(1+\frac{1}{\tau}\Big)\right], \tag{2.8}$$

where N_s counts the number of fields of spin s. If a non-minimal scalar coupling $\frac{1}{2}\xi R\phi^2$ is present, the trace (spin-0) channel carries the familiar $(1-6\xi)^2$ weight [31, 45]; in practice one may track this by $c_0 \to c_0(1-6\xi)^2$ in (2.8).

The long–distance $1/r^3$ correction to the Newtonian potential originates from the nonanalytic part of the covariant form factors $F_{i,s}(\tau)$. In the massless or high–momentum regime ($\tau \to \infty$), the form factors take the well–known limit $F_{i,s}(\tau) \simeq c_{i,s}(\ln\tau - 2)$, reproducing the $\ln\Big(-\Box/\mu^2\Big)$ running characteristic of massless loops [31, 45]. For massive fields with momenta well below threshold ($|\tau| \ll 1$), the same nonlocal structure admits the local curvature expansion

$$F_{i,s}(\tau) = c_{i,s}\Big(\tfrac{\tau}{6} - \tfrac{\tau^2}{60} + \cdots\Big), \tag{2.9}$$

which is the low–momentum expansion of the form factor that yields the universal $1/r^3$ tail upon resummation. At intermediate distances, this physics is conveniently interpolated by a Yukawa–envelope fit [107, 45]. Our mapping from ISL data to constraints on hidden fields therefore consistently incorporates both the curvature–squared operators arising at low momenta and their nonlocal running inherited from the full form factor.

For example, integrating out a real scalar yields a correction

$$U(r) = -\frac{Gm_1m_2}{r}\left[1+\frac{c_0}{G}\frac{1}{r^2}+\dots\right], \tag{2.10}$$

where the second term encodes the leading quantum effect. In general summing over all particle species (weighted by c_s) determines the net sign and magnitude of these corrections.

If the loop particle is massless, thus $\tau \to \infty$, the form factors reduce to the pure nonanalytic structure $F_s(\tau) \sim c_s \ln\left(-\Box/\mu^2\right)$, with μ a renormalization scale. This directly implies a scale-dependent gravitational coupling $G(q^2)$. We discuss the massless case in more detail in the following section. For particles with nonzero mass, the mass scale m acts as a threshold: at distances much larger than $1/m$, the quantum corrections are strongly suppressed, while at shorter distances the corrections gradually approach the same logarithmic running seen in the massless case. In this way, the massive case smoothly connects between $1/m^2$-suppressed effects in the low energy ($q^2 \ll m^2$) and universal logarithmic behavior in the high energy case ($q^2 \gg m^2$), a feature that plays an important role in precision tests of gravity across laboratory and astrophysical regimes.

2.4 Running and Threshold Effects in $G(q^2)$

The one-loop corrections can be interpreted as being momentum-dependent, or "running," Newton's constant. In flat-space, the running of Newton's constant can be parametrized by a correction to its inverse propagator,

$$G_{\text{eff}}^{-1}(q^2) \;=\; G_0^{-1}\left[1+\delta_s(q^2)\right], \tag{2.11}$$

where G_0 is the noncorrected gravitational constant and $\delta_s(q^2)$ encodes the quantum correction from a particle of mass m and spin s. In the case of massless fields, $m \to 0$, the expression reduces to a purely nonlocal form,

$$\delta_s(q^2) \;\sim\; c_s'\, G_0 q^2 \ln\left(\frac{-q^2}{\mu^2}\right), \tag{2.12}$$

where μ is a renormalization scale. This result is explicitly derived and discussed in detail later in Section 2.7. This universal logarithmic dependence is responsible for long-distance, nonlocal corrections to Newton's law [86, 37, 144].

For massive fields, the behavior is governed by threshold effects. At low momentum, $q^2 \ll m^2$, the logarithm can be expanded and the correction takes the form

$$\delta_s(q^2) \;\approx\; \frac{c_s}{2}\, G_0\, \frac{q^4}{m^2}, \tag{2.13}$$

which is analytic in q^2 and therefore indistinguishable from local higher-derivative operators. For graviton propagators with large momentum transfer, $q^2 \gg m^2$, the particle contributes fully to the running and the correction asymptotes to

$$\delta_s(q^2) \;\approx\; c_s\, G_0\, q^2 \ln\left(\frac{q^2}{m^2}\right). \tag{2.14}$$

The full nonlocal expression, $\ln\Big(1 + q^2/m^2\Big)$, thus interpolates smoothly between these two regimes, ensuring that heavy particles decouple at low energies while recovering the massless logarithmic running in the ultraviolet.

Summing over all particle species yields a compact expression for the effective coupling in Euclidean space,

$$G_{\rm eff}(-q^2) \;\simeq\; G_0 \left[1 + G_0 q^2 \sum_s c_s \ln\left(1 + \frac{q^2}{m_s^2}\right) + \dots \right], \tag{2.15}$$

where the sum runs over all fields with mass m_s. This interpolating behavior between the low-energy analytic regime and high-energy logarithmic running captures the essential physics of how massive fields contribute to gravitational quantum corrections. The threshold scale m determines whether particles contribute significantly to the running of Newton's constant: light particles ($q^2 \gg m^2$) participate fully in the quantum corrections, while heavy particles ($q^2 \ll m^2$) decouple and contribute only suppressed analytic terms. This mass-dependent decoupling will prove crucial when analyzing the cumulative effects of large hidden sectors, as it determines which mass ranges are most relevant for observable deviations from Newton's law at different distance scales.

2.5 Correction to the Newtonian Potential

As outlined above, the corrections to the Newtonian potential are obtained by integrating the nonlocal form factor above for all possible particle spins s. In the case of the massless fields, such corrections to the gravitational potential induce long range effects of the form $O(1/r^3)$ [86, 37, 144, 100, 126]:

$$U(r) = -\frac{Gm_1 m_2}{r}\left[1 + \alpha\,\frac{G}{r^2} + \dots\right] \tag{2.16}$$

where [37]

$$\alpha = \frac{41}{10\pi} \qquad \text{(graviton loops)}, \tag{2.17}$$

is obtained via graviton ($s = 2$) corrections and is modified via additional contributions from spins 0, 1/2, 1 to get:

$$\alpha = \frac{41}{10\pi} + \sum_f N_f \frac{1}{20\pi} + \sum_s N_s \frac{1}{120\pi} + \sum_v N_v \left(-\frac{1}{10\pi}\right) \tag{2.18}$$

where $N_{f,s,v}$ count Dirac fermions, real scalars, and vectors. For massive fields, corrections are exponentially suppressed beyond the Compton wavelength and interpolate between short-distance (massless) and long-distance (suppressed) regimes. Explicit analytic and numerical forms are given in [106, 45], including Bessel and Struve functions.

2.6 Validity of the EFT and species effects.

Loop-generated quadratic curvature terms can be reorganized into massive spin-0 and spin-2 propagating poles (the latter being a ghost in a purely local R^2+$R_{\mu\nu}R^{\mu\nu}$ truncation). In a Wilsonian picture the induced masses scale schematically like

$$M_{0,2} \sim \frac{M_{\mathrm{Pl}}}{\sqrt{\alpha_{0,2}}}, \qquad \alpha_{0,2} \propto \frac{N}{16\pi^2},$$

so extremely large N can lower $M_{0,2}$. Independently of N, collider searches for massive spin-2 resonances at the LHC place a lower bound $M_{0,2} \gtrsim \mathcal{O}(\text{few TeV})$ on any new state that couples to Standard Model matter with gravitational strength [2]; for moderate N

this bound is automatically satisfied, since the loop suppression factor $1/16\pi^2$ keeps $M_{0,2}$ well above the TeV scale, but in the large-N regime the two constraints become complementary. For moderate N this collider bound is not operative, since the loop suppression factor $1/16\pi^2$ keeps $M_{0,2}$ well above the TeV scale. However, in the large-N regime that is the primary focus of this work, $M_{0,2}$ can be driven toward the TeV scale, and the LHC bound then provides a complementary lower bound on $M_{0,2}$ that tightens the allowed parameter space independently of the laboratory gravity analysis. Our laboratory analysis is performed at distances $r_\star$ in the mm to μm range. If $M_{0,2}^{-1} \ll r_\star$, the massive poles merely renormalize the local curvature terms and the familiar nonlocal $1/r^3$ tail governs the signal. If instead $M_{0,2}^{-1} \gtrsim r_\star$, the same physics manifests as *additional Yukawa components* with ranges $\lambda_{0,2} \sim M_{0,2}^{-1}$, which are already encompassed by our envelope fit and lead to *tighter* constraints. Thus a conservative consistency condition for our nonlocal description is $M_{0,2}^{-1} \ll r_\star$; when it fails, one should include the corresponding Yukawa piece explicitly, which our mapping accommodates without changing the empirical logic of the bounds.

2.7 Massless Limit Consistency

A crucial consistency check for the quantum corrections derived in the preceding sections is that, in the limit of vanishing mass, the expressions for massive fields smoothly reproduce the well-known massless results both in momentum-space running and in the real-space potential. This property is explicitly realized in the referenced calculations [106, 45].

2.8 Momentum-Space Running

For a field of mass m and spin s, the nonlocal correction to Newton's constant in momentum space is obtained via integrating the quantum corrections given by F_s in Eq. 2.6 to obtain the form [106]:

$$\delta_s(q^2) = c_s\, G_0\, q^2 \ln\left(1 + \frac{q^2}{m^2}\right). \tag{2.19}$$

In the massless limit, we are interested in the $m \to 0$ limit. Doing so, the logarithm expands as:

$$\ln\left(1 + \frac{q^2}{m^2}\right) \longrightarrow \ln\left(\frac{q^2}{m^2}\right). \tag{2.20}$$

With this minor expansion, the overall expression for the massless correction to Newton's constant takes the form

$$\delta_s(q^2) \to c_s\, G_0\, q^2 \ln\left(\frac{q^2}{\mu^2}\right). \tag{2.21}$$

where μ has replaced the mass m and defines the renormalization scale in the massless case. This matches the standard massless result, where the running of G is governed by a $q^2 \ln\left(-q^2\right)$ form factor [86, 37]. In the limit $m \to 0$, the nonlocal form factor reduces to the universal massless expression $\sim q^2 ln(-q^2)$, in agreement with earlier EFT analyses of graviton and massless matter loops. The scale m in the massive case thus plays the role of a regulator, which is replaced by a renormalization scale μ in the massless theory.

2.9 Position-Space Potential

Massive Correction	**$m \to 0$ Limit (Massless)**	**Source**
Exponential/Yukawa: e^{-mr}/r^3	Nonanalytic $1/r^3$ tail	[106, 45]
Bessel/Struve function form	Expands to $1/r^3$	[45]

The effect of quantum corrections can also be understood directly in position space by examining their contribution to the Newtonian potential. For a particle of finite mass m, the correction takes a Yukawa-suppressed form

$$\delta U(r) \sim \frac{e^{-mr}}{r^3}, \tag{2.22}$$

so that the contribution is short-ranged and becomes negligible beyond the particle's Compton wavelength. More detailed treatments [45] show that the exact expression can be written in terms of special functions (e.g., Bessel and Struve), but the essential behavior is well captured by the exponential factor. Both [45, 106] express the full quantum-corrected potential $V(r)$ for massive loops using Bessel/Struve functions:

- For $mr \ll 1$, the expansion reproduces the $1/r^3$ behavior of massless loop corrections.
- For $mr \gg 1$, the correction is exponentially (Yukawa) suppressed.

In the massless limit, $m \to 0$, the exponential suppression goes to unity, leaving a long-range correction of the form

$$\delta U(r) \sim \frac{1}{r^3}. \tag{2.23}$$

This result matches the standard nonanalytic correction derived in effective field theory and demonstrates the smooth connection between the massive and massless cases. At short distances ($mr \ll 1$) the potential reduces to the $1/r^3$ form characteristic of massless loops, while at large distances ($mr \gg 1$) the correction is exponentially suppressed and therefore negligible. In this way, the position-space picture confirms the consistency of the massless limit and illustrates how heavy particles decouple from long-range gravitational physics.

2.10 Constraints on Hidden or Dark Sectors with Many Degrees of Freedom

The preceding analysis, based on the Standard Model particle content, shows that quantum corrections to Newton's law are utterly negligible at accessible length scales. However, the situation can change qualitatively in the presence of a dark sector containing a large number of light or very weakly coupled particles—for example, many real scalars, fermions, or vectors with masses m_i below or near the meV scale.

2.11 Quantum Corrections from Large Hidden Sectors.

The full quantum-corrected Newtonian potential arises from the combined effect of all light and heavy degrees of freedom running in loops [37, 106]. Each particle species contributes a correction determined by its spin, mass, and multiplicity, with massless states generating long-range, nonanalytic terms and massive states producing short-range Yukawa-suppressed effects. In practice, the net modification is obtained by summing over all particle types, weighted by the spin-dependent coefficients introduced in the previous section. This leads to the general expression for a modified gravitational potential:

$$U(r) = -\frac{Gm_1m_2}{r}\left[1 + \sum_{j\in\text{massless}} \alpha_j \frac{G}{r^2} + \sum_{i\in\text{massive}} \beta_i \frac{Ge^{-2m_i r}}{r^2} + \cdots\right] \tag{2.24}$$

with α_j and β_i determined by the spin-dependent loop coefficients c_s (see previous section), and the sums extend over all massless and massive dark sector states.

In the massive case with N nearly degenerate, weakly-coupled species with similar masses m_X, the Yukawa correction is enhanced by the total number of additional degrees of freedom N,

$$U_{\text{Yukawa}}(r) \sim -\frac{G^2 m_1 m_2}{r}\, N\, \beta_X \frac{e^{-2m_X r}}{r^2}. \tag{2.25}$$

This means that the quantum gravitational correction, while tiny for each individual state, can become collectively significant if N is large or if m_X is sufficiently small ($m_X \lesssim$ meV). For heavier states the Yukawa corrections are strongly suppressed, but ultralight particles can give rise to measurable deviations in precision experiments.

Note that our reinterpretation counts any *propagating* degree of freedom lighter than the experimental lever arm. Concretely, a state contributes to the potential if its mass satisfies $m \lesssim r_\star^{-1}$. In confining theories, heavy resonances above this scale are integrated out and do not affect our fits, whereas light PNGBs, glueballs, or other composites couple universally to gravity and contribute on the same footing as elementary fields. Our constraints apply to any light degree of freedom that behaves as an effectively pointlike scalar at the momentum transfers relevant for ISL experiments. For PNGBs, this condition is naturally satisfied because their masses can be parametrically smaller than their compositeness scale. For glueballs, this occurs only if the confinement scale of the hidden sector is sufficiently low (e.g. keV–MeV), so that the lightest glueball is both kinematically accessible and structureless at the probed momenta. In contrast, QCD glueballs are too heavy to satisfy this condition, and do not enter our bounds. The bounds therefore depend only on the infrared spectrum, not on UV compositeness.

2.12 Experimental Limits: Mapping onto the Yukawa Template.

A standard approach to constraining possible deviations from Newton's law is to express the potential in terms of a Yukawa-like correction, parameterized as

$$U_{\rm exp}(r) = -\frac{Gm_1m_2}{r}\left[1+\gamma\, e^{-r/\lambda}\right] \tag{2.26}$$

where γ is the dimensionless strength and λ the range of a hypothetical new Yukawa force. This form provides a convenient benchmark because it captures both the possibility of a new long-range force (if λ is large) and a short-range modification (if λ is small). Comparing to the Yukawa correction in Eq. 2.25, the collective quantum loop correction from a large hidden sector maps onto this template by identifying

$$\gamma = N\,\beta_X \frac{G}{2m_X\lambda} \qquad \text{with} \qquad \lambda = \frac{1}{2m_X}, \tag{2.27}$$

where the detailed normalization depends on the spin, couplings, and explicit computation of β_X for the field(s) in question [106, 45]. However, the mapping is not reliable as $m \to 0$; we return to a detailed discussion of this massless limit in Sec. 2.14.

2.13 Constraints and Projected Sensitivity

Experimental searches for deviations from Newton's law are commonly presented as exclusion plots in the (γ, λ) plane, where γ sets the relative strength and λ the range of a Yukawa-type correction. Within our framework, these bounds can be directly reinterpreted as limits on the number and properties of hidden-sector states that contribute to quantum loop corrections of gravity.

For instance, identifying $\lambda = 1/(2m_X)$ for a species of mass m_X, one finds that experiments typically constrain $|\gamma| \lesssim 10^{-2}$–$10^{-4}$ at micron length scales [45]. If the collective enhancement from a large number of species N or from very light masses m_X pushes the predicted quantum correction above these bounds, the corresponding models can be ruled out or tightly constrained. The overall size and sign of the effect depend on the weighted sum of coefficients c_s for the contributing fields, and can therefore be evaluated model by model using the master expressions for $U(r)$ derived above [106, 37]. If the hidden sector also contains massless states (such as scalars or dark photons), the cumulative correction is instead of the form

$$\alpha_{\text{tot}} = \sum_{j \in \text{massless}} N_j \, c_j, \tag{2.28}$$

leading to an additive $1/r^3$ term in the potential. While these corrections fall off faster than Yukawa-type terms and are generally harder to detect, their cumulative effect could still be constrained through precision fits to short-distance deviations from the $1/r$ force law. Still, the quantum loop corrections to the gravitational potential—Yukawa-suppressed for each individual massive state and $1/r^3$-suppressed for massless ones—can, if multiplied by a large number of hidden-sector degrees of freedom, become significant enough to be probed by laboratory experiments. The basic procedure is to map the theoretical prediction

$$\Delta U(r) \;=\; -\frac{Gm_1m_2}{r}\sum_X N_X\,\beta_X\,\frac{e^{-2m_Xr}}{r^2} \tag{2.29}$$

onto the experimental Yukawa form with parameters (γ,λ). Published exclusion plots in the (γ,λ) plane then translate directly into limits on the number N_X and mass m_X of new species [106, 45], as we will see below.

2.14 From fifth–force bounds to multiplicity limits

We now show how experimental fifth–force searches can be recast as direct bounds on the multiplicity of hidden-sector degrees of freedom that couple to gravity. This provides a framework between laboratory constraints on Yukawa or power–law deviations from Newton's law and theoretical scenarios with large numbers of light or massless states. Throughout this section we set $\hbar = c = 1$ and denote Newton's constant by G_N.

2.15 Set–up

To connect experimental fifth–force searches to hidden–sector physics, we consider a generic dark sector containing N_s real scalars, N_f Dirac fermions, and N_v vectors, all coupled only through gravity. The effect of these fields is to modify the Newtonian potential by loop exchange, producing a correction of the form

$$\frac{\Delta U}{U_{\rm Newt}}(r) = \sum_{X=s,f,v} N_X \frac{G_N}{r^2} \mathcal{K}_X(m_X r), \qquad U_{\rm Newt}(r) = -\frac{G_N m_1 m_2}{r}, \tag{2.30}$$

where $\mathcal{K}_X$ is a dimensionless loop kernel encoding the spin and mass dependence of the exchanged particle. Two limits of $\mathcal{K}_X$ are especially important:

$$\mathcal{K}_X(0) = \beta_X, \quad \beta_s = \frac{1}{120\pi}, \quad \beta_f = \frac{1}{20\pi}, \quad \beta_v = -\frac{1}{10\pi}, \tag{2.31}$$

for massless fields, and

$$\mathcal{K}_X(\xi) \xrightarrow[\xi\gg 1]{} \beta_X^{(\rm as)}(\xi)\, e^{-2\xi}, \qquad \xi \equiv m_X r, \tag{2.32}$$

The superscript (as) indicates the asymptotic form: $\beta_X^{(\rm as)}(\xi)$ gives the coefficient in the large-ξ expansion, which generally differs from the massless coefficient β_X due to additional corrections that emerge in the massive case at large distances. For massive fields at distances large compared to their Compton wavelength. Thus, massless species always generate a long–range $1/r^3$ correction, whereas massive species decouple exponentially at large r.

On the experimental side, deviations from Newton's law are typically reported in two standard parameterizations:

$$V_k(r) = -\frac{G_N m_1 m_2}{r}\left[1 + \beta_k \left(\frac{r_0}{r}\right)^{k-1}\right], \tag{2.33}$$

for power–law deformations of the inverse–square law, and

$$V_Y(r) = -\frac{G_N m_1 m_2}{r}\Big[1 + \alpha_{\rm exp}(\lambda)\, e^{-r/\lambda}\Big], \tag{2.34}$$

for Yukawa–type forces. The loop–induced $1/r^3$ correction maps onto the $k = 3$ case of Eq. (2.33), while finite–mass fields map onto the Yukawa form with range $\lambda = 1/(2m_X)$.

Our methodology for constraining multiplicity limits involves utilizing given experimental exclusion curves to obtain values of α and λ to set direct bounds on the number of degrees of freedom. In the following subsections we treat separately the *massless* case, which allows closed–form analytic bounds, and the *massive* case, which requires mapping onto Yukawa envelopes.

2.16 Massless dark degrees of freedom: closed–form bound with numbers

When $m_X = 0$, Eqs. (2.30) and (2.31) give

$$\frac{\Delta U}{U_{\rm Newt}}(r) = \Big(\frac{N_f}{20\pi} + \frac{N_s}{120\pi} - \frac{N_v}{10\pi}\Big)\frac{G_N}{r^2}. \tag{2.35}$$

Direct ISL tests at short range (torsion balances) often report the $k = 3$ power–law coefficient at $r_0 = 1$ mm. Writing the fractional residual as $\varepsilon(r) \equiv |\Delta U/U_{\rm Newt}|$, the $k = 3$ parameterization implies

$$\varepsilon(r)\, r^2 = \beta_3\, r_0^2 \qquad (k = 3,\ r_0 = 1\ \text{mm}), \tag{2.36}$$

independent of r. Combining Eqs. (2.35) and (2.36) yields the general multiplicity bound

$$\boxed{\left|\frac{N_f}{20\pi} + \frac{N_s}{120\pi} - \frac{N_v}{10\pi}\right| \le \frac{\beta_3\, r_0^2}{G_N}\,.} \tag{2.37}$$

Using the most precise published ISL *power–law* constraint (Eöt–Wash–style torsion balance),[1]

$$|\beta_3| \lesssim 7.5 \times 10^{-5} \quad (68\% \text{ CL}) \qquad \Rightarrow \qquad \frac{\beta_3 r_0^2}{G_N} \lesssim 2.87 \times 10^{59}, \tag{2.38}$$

where we used $r_0 = 1$ mm and $G_N = l_P^2 = 2.611 \times 10^{-70}\,\text{m}^2$. Equation (2.37) then gives the *per–spin* limits (assuming only one spin species is present):

$$\boxed{N_s \lesssim 1.08 \times 10^{62}} \qquad \text{(real scalar)}, \tag{2.39}$$

$$\boxed{N_f \lesssim 1.80 \times 10^{61}} \qquad \text{(Dirac fermion)}, \tag{2.40}$$

$$\boxed{N_v \lesssim 9.02 \times 10^{60}} \qquad \text{(vector; using } |\beta_v|\text{)}. \tag{2.41}$$

For Weyl/Majorana fermions, divide (2.40) by two. If multiple spins are present, apply (2.37) to the spin–weighted sum; one may also impose a robust (no–cancellation) limit by bounding the sum of absolute contributions. Equations (2.39)–(2.41) show that current short–range ISL data allow extremely large multiplicities of truly massless gravitationally–coupled fields; any improvement in $|\beta_3|$ translates linearly to these bounds.

[1] Tan *et al.* quote power–law bounds including $k = 3$; for convenience we take their tabulated $k = 3$ value at 68% CL, $|\beta_3| \lesssim 7.5 \times 10^{-5}$ [204]. Older analyses (e.g. [131]) give consistent though weaker values. At 95% CL, numbers inflate by a factor $\sim$1.6–2.

2.17 Massive dark degrees of freedom: mapping to Yukawa envelopes

For $m_X > 0$, the experimental summaries provide 95% CL limits $\alpha_{\rm exp}^{\rm max}(\lambda)$ in Eq. (2.34). Identify the range with particle mass as

$$\lambda = \frac{1}{2m_X}, \tag{2.42}$$

and match the leading long–range piece of (2.30) to the Yukawa form at the apparatus lever arm $r_\star$ (choosing $r_\star \simeq \lambda$ maximizes sensitivity). Defining a dimensionless shape factor $\mathcal{S}_X(\xi) \equiv \mathcal{K}_X(\xi)\, e^{2\xi}$ with $\mathcal{S}_X(0) = 1$ and $\mathcal{S}_X(\frac{1}{2}) = \mathcal{O}(1)$, one obtains

$$\alpha_{\rm th}(\lambda) \simeq N_X\, \beta_X\, \frac{G_N}{r_\star^2}\, \mathcal{S}_X(m_X r_\star) \quad \Rightarrow \quad \boxed{N_X^{\rm max}(\lambda) \simeq \frac{\alpha_{\rm exp}^{\rm max}(\lambda)}{\beta_X\, \mathcal{S}_X(1/2)}\, \frac{\lambda^2}{G_N}\,, \quad \lambda = \frac{1}{2m_X}.} \tag{2.43}$$

Setting $\mathcal{S}_X(1/2) = 1$ gives a slightly conservative estimate that is convenient for quick re-interpretations of published (α, λ) curves (torsion balances, LLR, planetary dynamics, *etc.*). Applying the results from precision experiments at varying length scales, we can obtain continuous limits on the multiplicities of dark sectors, as is illustrated in Fig. 2.1 below.

Including additional Yukawa poles when they enter the lever arm. If the higher-derivative poles discussed in Sec. 2.2 descend into the laboratory window, their effect is incorporated by *adding* explicit Yukawa pieces to the static potential, rather than relying solely on the single-envelope mapping. We write

$$U(r) \;=\; -\frac{G\,m_1 m_2}{r}\left[1+\sum_{k=1}^{n_{\rm Y}} \alpha_k\, e^{-r/\lambda_k}\right], \tag{2.44}$$

with (λ_k, α_k) the range and strength of the k-th Yukawa term. For the quadratic-curvature spin channels one may take $n_{\rm Y} = 2$, with

$$(\lambda_0, \alpha_0) \;\text{ and }\; (\lambda_2, \alpha_2), \qquad \lambda_{0,2} = \frac{1}{2M_{0,2}}, \tag{2.45}$$

mirroring the convention in Eq. (2.42). The data analysis then proceeds by profiling the likelihood (or χ^2) over (α_0, α_2) at fixed (λ_0, λ_2), yielding one–sided bounds $\alpha_{0,2}^{95}(\lambda_{0,2})$ (single–pole case) or joint contours for two poles. Because Eq. (2.44) fits the *sum* of deviations from $1/r$, any additional pole within the probed range can only increase the total deviation unless it is finely tuned to cancel other pieces; our envelope construction avoids such tuned cancellations and thus remains conservative.

2.18 Experimental Multiplicity limits in both the Massive and Massless Case

Utilizing values of α and λ from experimental exclusion curves, we construct analogous exclusion plots for the upper limits on the number of degrees of freedom based on Eq. (2.43) for massive particles in Figure 2.1. For concreteness, we assume the shape factor S_X is unity and model the gravitational correction as arising entirely from a large sector of fermionic particles with spin-dependent loop coefficient $\beta_f = \frac{1}{20\pi}$. The methodology readily extends to scalar and vector particles by substituting the

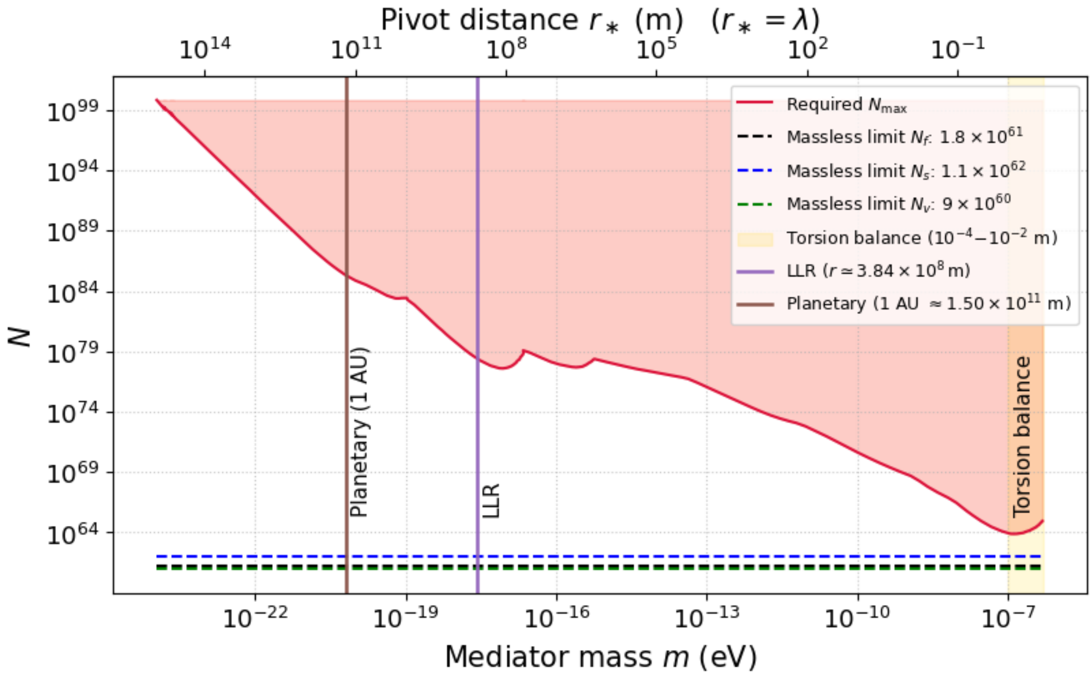

Figure 2.1: Degrees of freedom set by experimental constraints on G and the Yukawa parameters (α, λ). The shaded region indicates the exclusion set by experimental limits at each distance scale $r_\star$. The red exclusion curve is derived from the massive Yukawa mapping and loses sensitivity as $m \to 0$ because the pivot distance $\lambda = 1/(2m) \to \infty$ exceeds any laboratory lever arm; the dashed lines are computed independently from the massless $1/r^3$ power-law regime and do not represent the $m \to 0$ limit of the red curve, as discussed in Section 2.8.

appropriate coefficients $\beta_s = \frac{1}{120\pi}$ and $\beta_v = -\frac{1}{10\pi}$, respectively, which modify the bounds by at most an order of magnitude.

Figure 2.1 presents these constraints in the (m, N) plane, where the horizontal axis shows the dark sector particle mass m and the vertical axis displays the maximum allowed multiplicity N. The corresponding pivot distance $r_\star = \lambda = \frac{1}{2m}$ is indicated on the top axis to facilitate comparison with experimental length scales. The red exclusion curve represents the upper bound $N_{\max}$ derived from Yukawa envelope constraints using Eq. (2.43), with the shaded region above this curve excluded by current data. Horizontal dashed lines mark the massless limits from Eqs. (2.39)–(2.41), providing direct comparison between massive and massless scenarios.

Several key features emerge from this analysis. First, the massless limit provides the most stringent constraints across all particle types, with vector bosons offering the tightest bound at $N_v \lesssim 9 \times 10^{60}$ degrees of freedom. This enhanced sensitivity reflects the long-range nature of massless corrections, which scale as $1/r^3$ and remain unsuppressed at the distance scales probed by precision experiments.

For massive particles, the constraints exhibit strong scale dependence. Torsion balance experiments operating at millimeter to submillimeter scales ($r_\star \sim 10^{-4}$–10^{-2} m) provide the most restrictive bounds, limiting multiplicities to $N \lesssim 10^{64}$ for particles with masses in the 10^{-16}–10^{-13} eV range. This dominance arises because torsion balances probe precisely the distance scales where massive corrections transition from the short-range exponentially suppressed regime to the intermediate regime where quantum effects are maximally observable.

At larger distances (smaller masses), the constraints weaken considerably as only planetary dynamics and lunar laser ranging (LLR) provide sensitivity. These techniques, while remarkable in their precision over astronomical scales, cannot match the controlled laboratory environment of torsion balance experiments for constraining quantum gravitational effects. The vertical bands in Figure 2.1 clearly delineate these complementary experimental regimes, illustrating how different techniques carve out distinct regions of parameter space.

The steep rise in $N_{\max}$ toward smaller masses reflects the exponential suppression of Yukawa corrections: as particle masses decrease, larger multiplicities are required to generate observable deviations from Newton's law. However, this trend is ultimately cut off by the massless limit, where the analytic $1/r^3$ corrections provide the fundamental theoretical ceiling.

These results demonstrate that precision gravitational tests already impose non-trivial constraints on dark sector model building. While the numerical bounds may appear large—reaching $\mathcal{O}(10^{64})$ in some cases—they represent genuine, model-independent limits on the complexity of hidden sectors that can exist while remaining consistent with laboratory measurements. As experimental sensitivity continues to improve, particularly in the submillimeter regime where torsion balance techniques excel, these bounds will tighten proportionally, providing increasingly powerful probes of beyond-Standard-Model physics through purely gravitational interactions.

The complementarity between different experimental approaches highlights an underappreciated synergy in fundamental physics: laboratory-scale precision gravity experiments, originally conceived to test general relativity or search for extra dimensions, emerge as sensitive probes of dark sector multiplicity and structure. This connection opens new avenues for constraining theoretical models that predict large numbers of light degrees of freedom, from strongly coupled dark gauge theories to string theory compactifications with extensive moduli sectors.

2.19 Discussion

In this section we have systematically derived and analyzed quantum gravitational corrections to Newton's law arising from large dark sectors, translating effective field theory predictions into experimentally testable constraints. Our analysis demonstrates that while individual particles contribute negligibly to gravitational modifications, sectors containing many light degrees of freedom can generate cumulative effects that approach the sensitivity thresholds of current precision gravity experiments. For truly massless hidden sector particles, current inverse-square law tests constrain the multiplicities to remarkably large values: $N_s \lesssim 1.08 \times 10^{62}$ real scalars, $N_f \lesssim 1.80 \times 10^{61}$ Dirac fermions, and $N_v \lesssim 9.02 \times 10^{60}$ vectors. These bounds, while numerically enormous, represent the first direct laboratory constraints on the size of massless dark sectors coupled purely gravitationally. The $1/r^3$ power-law corrections from massless fields provide the strongest leverage for such constraints, as they are not exponentially suppressed at the distance scales probed by torsion balance experiments.

For massive dark particles, the quantum corrections take the form of Yukawa-suppressed modifications with characteristic range $\lambda = 1/(2m)$. By mapping these corrections onto the experimental Yukawa template, we have shown how published fifth-force exclusion curves can be directly translated into multiplicity bounds $N_{\max}(\lambda)$ for particles of mass $m = 1/(2\lambda)$. This approach provides a systematic framework for constraining the mass spectrum and abundance of hidden sector states using existing gravitational data. Beyond the specific numerical bounds, our work establishes a general methodology for connecting quantum field theory predictions in gravity with precision experimental constraints. The translation between effective field theory nonlocal form factors and the phenomenological parametrizations used by experimentalists provides a crucial bridge between theoretical dark sector model-building and laboratory tests.

The gravitational constraints derived here occupy a unique position in the landscape of dark sector searches. Unlike collider experiments, which require non-gravitational interactions, or cosmological observations, which are sensitive to the collective energy density, precision gravity tests can probe purely gravitationally-coupled sectors at the level of individual particle multiplicities. Cosmological constraints on relativistic degrees of freedom, typically expressed through effective neutrino number N_{eff}, provide complementary information but probe different physics. While cosmology constrains the total energy contribution of light species during nucleosynthesis and recombination, precision gravity tests are sensitive to the detailed spectrum and multiplicity structure of dark sectors at much later times and smaller scales. The mass ranges accessible to laboratory gravity experiments, corresponding to Compton wavelengths from micrometers to millimeters, complement both high-energy collider searches and astrophysical probes. This intermediate scale regime is particularly relevant for dark sectors arising from string theory compactifications or strongly-coupled hidden gauge theories, where light states with masses in the meV to eV range naturally emerge.

Note that our bounds are laboratory and model-independent: they constrain the *present-day* multiplicity of light states without assuming a cosmological population. Cosmology can provide stronger constraints when the hidden sector was populated and remained relativistic at BBN/CMB epochs, typically phrased as limits on ΔN_{eff}. However, low reheating temperatures, suppressed portals, late entropy injection, or purely gravitational/out-of-equilibrium production histories can substantially weaken cosmological sensitivity while leaving our reinterpretation intact. We therefore view the two approaches as complementary: CMB/BBN constrain *abundances* under specified histories, whereas precision gravity constrains the *spectrum* of light degrees of freedom irrespective of their relic density.

Our results have direct implications for several classes of theoretical models. Strongly coupled dark gauge sectors may contain multiple light states (e.g. PNGBs or glueballs) below the experimental lever arm; while each contributes negligibly, the cumulative effect scales with the number of such light states: any unexpectedly rich light spectrum is directly testable in the same, model-independent way. Our bounds provide quantitative guidance for the maximum complexity such sectors can achieve while remaining consistent with laboratory tests. Extra-dimensional theories often predict towers of Kaluza-Klein modes or large numbers of moduli fields, and the gravitational constraints complement existing bounds from fifth-force searches to help restrict the allowed parameter space of phenomenologically viable compactifications. Similar analyses could be extended to axion-like particles that couple to gravity through topological terms or induce effective modifications to gravitational couplings.

It is important to remark that Kaluza–Klein towers fit naturally into our framework *provided one sums only over modes lighter than the probed scale.* In the short-distance regime, massive KK gravitons generate a set of Yukawa deviations whose envelope is already captured by our mapping; the well-studied short-range gravity limits in ADD/RS scenarios can be rephrased in exactly this language. Our plots should thus be read as providing an efficient, model-agnostic translation between measured inverse-square-law tests and the cumulative effect of any tower of light modes.

We also note that for extremely large N, loop-induced higher-derivative poles can drift down toward the laboratory window. In that regime, our reinterpretation remains applicable and typically becomes *more* constraining, because such poles manifest as Yukawa deviations from $1/r$ that are bounded by the same precision data.

Several aspects of our analysis merit further investigation. We have focused on leading-order quantum corrections arising from one-loop effects, but higher-order contributions could modify the detailed form of the gravitational modifications, particularly in strongly-coupled dark sectors where perturbative calculations may break down. Our treatment assumes minimal gravitational coupling through the stress-energy tensor, but dark sector particles with non-minimal couplings to curvature could generate different signatures that might be more readily accessible to experiment. The analysis presented here applies to vacuum quantum corrections, though in cosmological or astrophysical contexts, thermal corrections from dark sector particles could modify the gravitational dynamics in observable ways. The constraints presented here are based on current experimental sensitivities, and ongoing improvements in torsion balance experiments, lunar laser ranging, and other precision gravity tests will tighten these bounds and potentially access new regions of parameter space.

The intersection of precision gravity experiments with dark sector phenomenology represents an underexplored frontier in fundamental physics. While the Standard Model contributions to gravitational quantum corrections are utterly negligible, the situation changes qualitatively in the presence of large hidden sectors. This work demonstrates that laboratory tests of Newton's law, originally conceived to search for extra dimensions or light scalar fields, can serve as powerful probes of the multiplicity and structure of dark matter sectors. The enormous numerical values of the bounds we derive, reaching $\mathcal{O}(10^{61})$ for some particle types, should not obscure their physical significance. These constraints represent genuine, model-independent limits on the complexity of dark sectors that could exist in nature while remaining consistent with precision tests of gravity. As experimental sensitivity improves, these bounds will tighten proportionally, providing increasingly stringent tests of theoretical models that predict large numbers of light degrees of freedom.

We have established a quantitative framework for constraining large dark sectors using precision tests of Newton's gravitational law. By systematically translating quantum field theory predictions for loop-induced gravitational corrections into the experimental language of Yukawa deviations and inverse-square law violations, we derive robust bounds on the multiplicities and mass spectra of hidden sector particles coupled only gravitationally. Our results demonstrate that current torsion balance experiments and related precision gravity tests already impose non-trivial constraints on dark sector model building, with bounds reaching $\mathcal{O}(10^{61})$ for massless degrees of freedom. As experimental techniques continue to improve, this approach will provide increasingly powerful tests of theories predicting large numbers of light particles beyond the Standard Model. The methodology developed here opens new avenues for connecting theoretical dark sector phenomenology with laboratory experiments, highlighting an unexpected synergy between precision gravity measurements and particle physics beyond the Standard Model, and establishing precision gravitational tests as a valuable complement to collider searches and cosmological observations in the quest to understand the hidden sectors that may populate our universe.

Chapter 3

The Maximal Gravitational Wave Signal from Asteroid-Mass Primordial Black Hole Mergers At Resonant Microwave Cavities

3.1 Introduction

Chapter 1 established that asteroid-mass primordial black holes (PBHs) in the mass window 10^{-16}–$10^{-11}\,M_{\odot}$ represent a well-motivated dark matter candidate that is largely inaccessible to conventional gravitational-wave detectors, and identified narrow-band resonant microwave cavities as the most promising detection avenue for gravitational waves (GWs) from PBH mergers in this regime. This chapter carries out that program in detail, deriving the maximum possible merger event rate and the corresponding detection prospects at ADMX-like cavities.

A distinctive feature of high-frequency gravitational waves (HFGWs) is the purported absence of known astrophysical backgrounds, making a detection a potent signal of new physics or a new class of GW sources. Among the well-motivated sources in the MHz–GHz band are light PBH mergers, which could shed light on early Universe conditions and the microscopic nature of cosmological dark matter [51, 96]. The detection channel we focus on exploits the inverse Gertsenshtein effect [123], wherein gravitons convert to photons in the presence of a static magnetic field, as in axion search experiments such as ADMX [178, 82]. The interested reader is referred to Refs. [9, 9] for an overview of the broader HFGW experimental landscape.

Prior work has examined relic GW energy densities from PBH mergers [79, 129, 186], stochastic backgrounds from early PBH coalescences [214, 108], and resonant-cavity detection of GWs from PBH coalescences and hyperbolic encounters [28, 29]. The present study extends this program by focusing on the conditions that maximize the detectable event rate: (i) considerable early-time clustering leading to enhanced binary formation, (ii) a large local overdensity of dark matter, (iii) a non-monochromatic PBH mass function, and (iv) a PBH abundance close to the totality of the cosmological dark matter. We additionally demonstrate that the sight distance at resonant cavities becomes chirp-mass independent at asymptotically large PBH masses; a result with direct implications for detector targeting.

The chapter is organized as follows. Section 3.2 describes the formalism and observational constraints on non-monochromatic PBH mass functions. Section 3.3 out-lines the sensitivity of microwave cavities to HFGWs from PBH mergers. Section 3.4 derives the dominant merger pathway and the corresponding rate of detectable events at an ADMX-like cavity; subdominant pathways are discussed in Appendix 7.1. Section 3.6 presents our conclusions.

3.2 Constraints for a generic mass function

We consider a mass distribution of PBHs (the so-called "mass function"), the mass-weighted differential number of PBHs, of the standard form

$$\psi(m) \equiv \frac{1}{m}\frac{dn_{\rm PBH}(m)}{dm}, \tag{3.1}$$

normalized so that the cosmological mass density of PBHs $\rho_{\rm PBH}$ relative to the cosmological dark matter abundance $\rho_{\rm DM}$

$$f_{\rm PBH} = \int \psi(m) dm. \tag{3.2}$$

We refer to monochromatic mass functions, corresponding to the form

$$\psi_0(m) = f_{\rm PBH}(m_0)\delta(m - m_0), \tag{3.3}$$

and to "dichromatic" mass functions,

$$\psi_{\rm DC}(m) = f_1\delta(m - m_1) + f_2\delta(m - m_2). \tag{3.4}$$

The latter is especially well motivated in the present context, as previous studies have numerically shown that the structure of mass functions that maximize black hole merger rates are generally rather close to the schematic form in Eq. (3.4) (see e.g. Ref. [151]).

Note that translating from the customarily shown constraints on monochromatic mass functions, representing the largest-possible value of $f_{\rm PBH}^{\rm max}(m_0)$ for Eq. (3.3) above compatible with experimental and observational constraints, to a generic mass function $\psi(m)$ amounts to requiring that [53]

$$\int \frac{\psi(m)}{f_{\rm PBH}^{\rm max}(m)} dm \leq 1. \tag{3.5}$$

Hence, the resulting constraint on f_1 and f_2 for a monochromatic mass function is (dropping the superscript "max" from now on):

$$\frac{f_1}{f_{\rm PBH}(m_1)} + \frac{f_2}{f_{\rm PBH}(m_2)} \leq 1. \tag{3.6}$$

For a given f_1, the value of f_2 that maximizes the mass-density of PBHs is then fixed by

$$f_2 = {\rm Max}\left(f_{\rm PBH}(m_2)\left(1 - \frac{f_1}{f_{\rm PBH}(m_1)}\right), 0\right) = {\rm Max}\left(\frac{f_{\rm PBH}(m_2)}{f_{\rm PBH}(m_1)}\left(f_{\rm PBH}(m_1) - f_1\right), 0\right). \tag{3.7}$$

3.3 Microwave cavities: sensitivity

Resonant cavities – first proposed to search for axions [199] – are sensitive to high-frequency GWs: a GW passing in a cavity endowed with a static magnetic field sources an effective electromagnetic current that, in turn, generates an electromagnetic field at the same frequency of the GW [105, 82, 36, 109]. Currently-operating and proposed detectors include ADMX [39, 40, 30], HAYSTAC [221], CAPP [149], and ORGAN [164]. Here, for definiteness, we consider an ADMX-type detector for our estimates. In particular, we adopt the following estimate for ADMX's sensitivity to the GW amplitude [105]:

$$\begin{aligned} h_{\rm sens}(\nu) &= 3\times 10^{-22}\left(\frac{0.1}{\eta_n}\right)\left(\frac{8\ \mathrm{T}}{|\vec{B}|}\right)\left(\frac{0.1\ \mathrm{m}^3}{V_{\rm cav}}\right)^{5/6}\left(\frac{10^5}{Q}\right)^{1/2} \\ & \left(\frac{T_{\rm sys}}{1\ \mathrm{K}}\right)^{1/2}\left(\frac{1\ \mathrm{GHz}}{\nu}\right)^{3/2}\left(\frac{\Delta\nu}{10\ \mathrm{kHz}}\right)^{1/4}\left(\frac{1\ \mathrm{min}}{\Delta t}\right)^{1/4}. \end{aligned} \tag{3.8}$$

In the equation above, η_n is the cavity's effective GW coupling; Note that in the most-optimal possible case, i.e. when the azimuthal direction of the relevant cavity mode resonates with the spin structure of the gravitational field, the cavity's effective coupling coefficient η_n is $\mathcal{O}(0.1)$ [35]. This includes the TM_{010} and TM_{020} modes in use for axion experiments like ADMX. It was also found in [109] that $\eta_n \approx 0.14$ for the TM_{012} mode. Note that while [105] presents the coupling coefficient as $\mathcal{O}(1)$, the authors used $\eta_n = 0.1$ in their calculations for the strain sensitivities of ADMX and SQMS. We use the same here, albeit with future optimization this may actually be an under-estimate.

Also in Eq. (3.8), $|\vec{B}|$ is the cavity's magnetic field; Q is the quality factor of the cavity; ν the operating resonant frequency; and $\Delta\nu \sim \nu/Q$ the bandwidth. The other parameters are the cavity's volume, operating temperature, and acquisition observation time. Unless otherwise specified, we will utilize the pivot values in Eq. (3.8) as our reference values for the sensitivity of an ADMX-like microwave cavity.

As customary, we define the chirp mass of the binary as

$$m_c \equiv \frac{(m_1 m_2)^{3/5}}{(m_1+m_2)^{1/5}}. \tag{3.9}$$

The gravitational wave amplitude for a merger of a binary with chirp mass m_c at a distance d reads (see e.g. Eq.(4.3) and (4.29) of Ref. [158])

$$h(m_c, d, \nu) = \frac{4}{d}\left(\frac{Gm_c}{c^2}\right)^{5/3}\left(\frac{\pi\nu}{c}\right)^{2/3}, \tag{3.10}$$

at frequency ν. The time derivative of the frequency reads (Eq. (4.18) in Ref. [158])

$$\dot{\nu} = \frac{96}{5}\pi^{8/3}\left(\frac{Gm_c}{c^3}\right)^{5/3}\nu^{11/3}. \tag{3.11}$$

The signal in the cavity depends on whether the cavity is fully "rung-up" or not; in turn, this depends on comparing the so-called *ring-up time* of the cavity, $t_r = Q/\nu_0 = 10^{-4}$ sec, the latter value referring to the specific case of ADMX, with the time spent at frequency ν, given by

$$t_g(M, \nu) = \int_{\nu_{\min}}^{\nu_{\max}} \frac{d\nu}{\dot{\nu}}, \tag{3.12}$$

where $\nu_{\mathrm{max,min}} = \nu_0 \pm \nu_0/Q$, with $\nu_0 = 1$ GHz the resonant frequency of the cavity. Carrying out the integral, and noting that $Q^2 - 1 \simeq Q^2$ for a narrow-band cavity, we find

$$t_g(M, \nu) = \frac{5}{48}\frac{1}{Q}(\pi\nu_0)^{-8/3}\left(\frac{GM}{c^3}\right)^{-5/3}. \tag{3.13}$$

Note that in fact one should use, above, the Lorentzian suppression of the cavity's response around the resonance frequency, and the proper integral is

$$t_{g'} = \int \frac{d\nu}{\dot{\nu}}\frac{1}{1+\left(\frac{\nu-\nu_0}{\nu_0/Q}\right)^2}; \tag{3.14}$$

using the expression above, however, we find that the result is very close, to a few percent, to the estimate that neglects the Lorentzian factor.

Note that the expressions above do not follow the behavior expected for broad-band detectors,

$$N_{\rm cycles}^{\rm broad} = \int_{t_{\rm min}}^{t_{\rm max}} \nu dt = \int_{\nu_{\rm min}}^{\nu_{\rm max}} d\nu \frac{\nu}{\dot{\nu}} \simeq \frac{\nu^2}{\dot{\nu}}, \tag{3.15}$$

with the characteristic time spent at a frequency ν is

$$t_g^{\rm broad}(m_c, \nu) = \frac{\nu}{\dot{\nu}}, \tag{3.16}$$

as quoted e.g. in Ref. [105], Eq. (2.33), and Ref. [109], Eq. (12).

The ring-up condition for a resonant cavity can also be expressed in terms of the quality factor and the number of cycles the signal spends at the resonant frequency. The latter, i.e. the number of cycles within the cavity's bandwidth is

$$N_c = \int_{\nu_{\rm min}}^{\nu_{\rm max}} \frac{\nu}{\dot{\nu}} d\nu \simeq \frac{5}{48} \frac{\nu_0}{(\pi\nu_0)^{8/3} Q} \left(\frac{GM}{c^3}\right)^{-5/3}, \tag{3.17}$$

and thus the condition for the cavity to be rung up is equivalently expressed as $N_c = Q$.

The signal in the cavity follows from the ring-up equation for resonant cavities

$$h_{\rm sig}(m_c, d, \nu) = h(m_c, d, \nu) \,\min\left[1 - \exp\left(N_c(m_c, \nu)/Q\right)\right]. \tag{3.18}$$

We define the sight distance, or distance sensitivity, for a chirp mass m_c at a frequency ν as the distance at which

$$h(m_c, d_{\rm sens}, \nu) = h_{\rm sens}(\nu). \tag{3.19}$$

Note that for $N_c/Q \ll 1$, $1 - \exp\left(N_c/Q\right) \simeq N_c/Q$; therefore, the signal in the cavity goes as

$$\lim_{N_c \ll Q} h_{\rm sig}(m_c, d, \nu) = \frac{4}{d}\left(\frac{Gm_c}{c^2}\right)^{5/3} \left(\frac{\pi\nu}{c}\right)^{2/3} \frac{5}{48} \frac{\nu}{(\pi\nu)^{8/3} Q} \left(\frac{Gm_c}{c^3}\right)^{-5/3} = \frac{5}{12\pi^2 d} \frac{c}{\nu_0 Q}, \tag{3.20}$$

and is thus independent of mass; the distance sight then simply goes as

$$d_{\rm sens} \sim 0.03\ {\rm AU} \left(\frac{\Delta\nu}{10\ {\rm kHz}} \right)^{1/4} = 0.03\ {\rm AU} \left(\frac{\nu}{1\ {\rm GHz}} \frac{10^5}{Q} \right)^{1/4} . \tag{3.21}$$

We show the ADMX distance sensitivity at a frequency ν for mergers involving two masses m_1 and m_2 in Fig. 3.1. The top left panel shows the monochromatic mass function case, $m_1 = m_2$ for three different frequencies, $\nu = 10^8,\ 10^9$ and 10^{10} Hz, while the top right panel assumes a di-chromatic mass function of mass ratio $m_1/m_2 = 0.01$ and the bottom left panel of $m_1/m_2 = 0.01$, again for the same three frequencies. Note that the shape of the curves as a function of mass reflects the mass dependence of the sensitivity function $d_{\rm sens}(m_1, m_2, \nu)$; note, in particular, the chirp-mass-independent sight distance at large masses ($N_c/Q \ll 1$) discussed above; also note that ADMX's sensitivity, corresponding, approximately, to the red lines, is typically around 0.03 AU only.

In the bottom right panel of Fig. 3.1 we show, for $\nu = 10^9$ Hz, the ratio of the distance sensitivity for a given mass ratio m_2/m_1 to that corresponding to a monochromatic mass function ($m_1 = m_2$), as a function of m_1. The plot illustrates how mass function ratios, $m_2/m_1 < 1$, can be detected at a greater distance than the monochromatic case, due to presence of lighter binaries that can ring up the cavity for larger values of m_1; similarly, for high mass ratios, $m_2/m_1 > 1$ the largest merger ratios occur at smaller m_1, again, that is, when one of the masses in the binary is sufficiently small. We note that in any case *a non-monochromatic (here, a dichromatic) mass function generically enables larger distance sensitivities than the monochromatic*

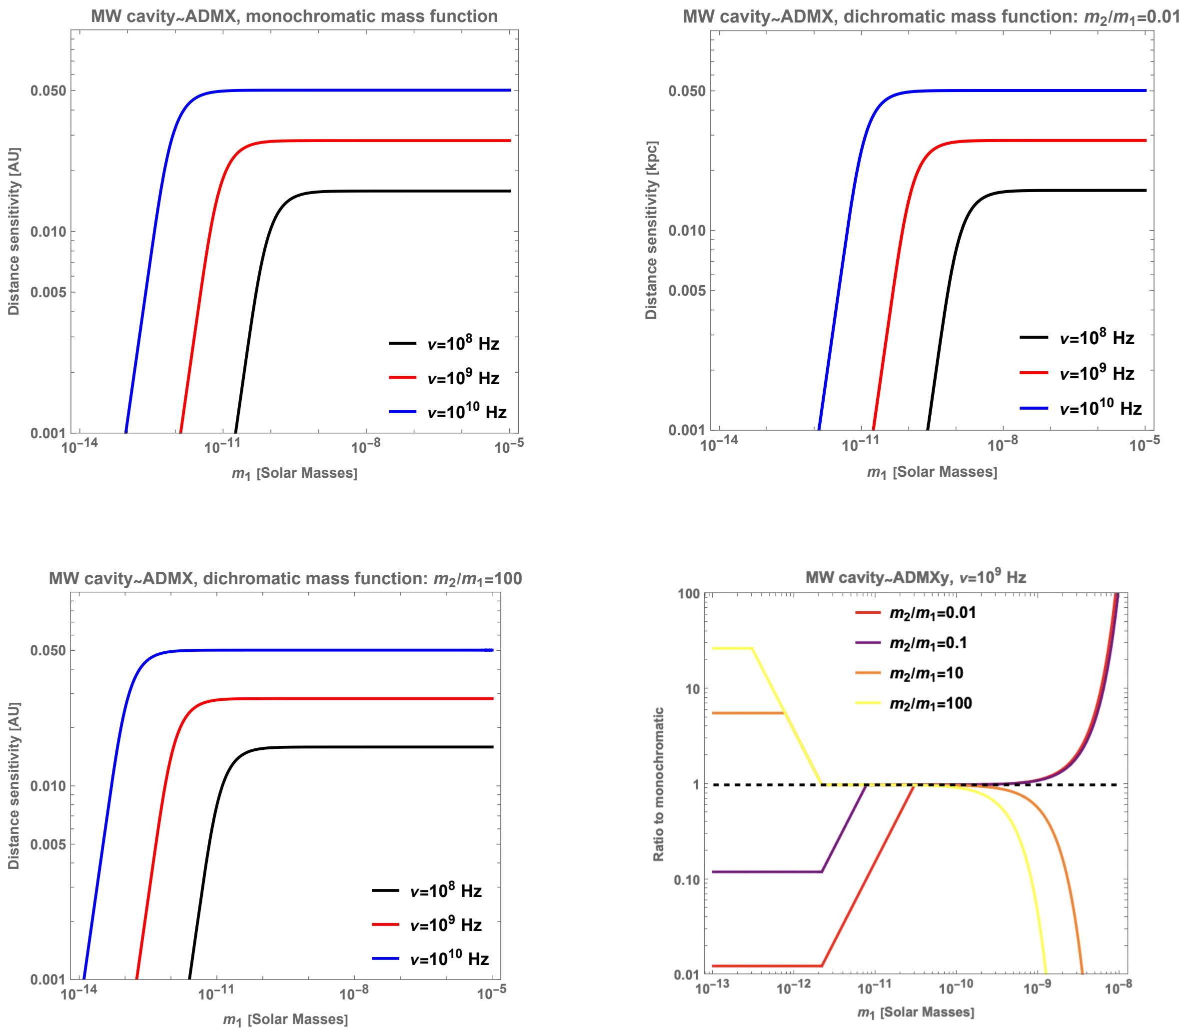

Figure 3.1: Distance sensitivity comparison for a microwave cavity with specifications close to ADMX's for three different frequencies, $\nu = 10^8,\ 10^9$ and 10^{10} Hz. In the top left corner we assume a monochromatic mass function, in the top right that $m_2/m_1 = 10^{-2}$, in the bottom left that $m_2/m_1 = 10^2$; in the bottom right we instead show the ratio, for $\nu = 10^9$ Hz, of the distance sensitivity for a given mass ratio m_2/m_1 as in the legend, divided by that for the monochromatic case, $m_1 = m_2$.

case, for sufficiently spaced-out masses.

3.4 Rate of detectable events from light black hole mergers

In what follows we are concerned with the calculation of the rate of detectable events from light PBH mergers, principally with a monochromatic or dichromatic mass function, at a microwave resonant cavity such as ADMX.

For a generic mass function $\psi(m)$, the *rate* of visible mergers at frequency ν, for a detector whose critical distance sensitivity is given by the function $d_{\rm sens}(m_1, m_2, \nu)$ in Eq. (3.21), reads

$$R[\psi](f) = \int_{m_{\rm min}}^{\infty} dm_1 \int_{m_{\rm min}}^{\infty} dm_2 \, \frac{dR[\psi](f)}{dm_1 dm_2} \frac{4\pi}{3} d_{\rm sens}^3(m_1, m_2, f), \tag{3.22}$$

where $m_{\rm min}$ corresponds to black holes long-lived enough not to have expired yet and R is the merger rate for the given mass function and frequency.

We demonstrate in the Appendix that the largest differential merger rate today corresponds to the early three-body binary formation pathway ($E3$ in what follows), discussed in detail in Ref. [190] (see their Eq. (2.13)):

$$\begin{aligned} \frac{dR_{E3}}{dm_1 dm_2 dm_3} &= \frac{9}{296\pi} \frac{1}{\tilde{\tau}} \left(\frac{t_0}{\tilde{\tau}}\right)^{-34/37} \left(\Gamma\left[\frac{58}{37}, \tilde{N}\left(\frac{t_0}{\tilde{\tau}}\right)^{3/16}\right] - \Gamma\left[\frac{58}{37}, \tilde{N}\left(\frac{t_0}{\tilde{\tau}}\right)^{-1/7}\right]\right) \\ &\quad \times \tilde{x}^{-3} \delta_{\rm dc}^{-1} \tilde{N}^{53/37} \bar{m}^3 \frac{\psi(m_1)}{m_1} \frac{\psi(m_2)}{m_2} \frac{\psi(m_3)}{m_3}. \end{aligned} \tag{3.23}$$

In the equation above (restoring factors of c suppressed in the original reference),

$$\tilde{\tau} \equiv c^5 \frac{384}{85} \frac{\alpha^4 \beta^7 a_{\rm eq}^4 m_3^7 \tilde{x}^4}{G^3 \eta M^{10}}, \tag{3.24}$$

where α and β are $\mathcal{O}(1)$ numerical constants parameterizing the semi-major and semi-minor axes of the forming binary in the three-body approximation [190], and we will hereafter assume $\alpha = \beta = 1$.

$$\eta \equiv \frac{m_1 m_2}{M^2}, \qquad M \equiv m_1 + m_2, \tag{3.25}$$

and

$$\tilde{x}^3 \equiv \frac{3}{4\pi} \frac{M}{a_{\mathrm{eq}}^3 \rho_{\mathrm{eq}}}. \tag{3.26}$$

Finally,

$$\tilde{N} \equiv \delta_{\mathrm{dc}} \Omega_{\mathrm{DM,eq}} \frac{M}{\bar{m}}, \quad \text{and} \quad \bar{m} \equiv \left(\int dm \frac{\psi(m)}{m} \right)^{-1}. \tag{3.27}$$

In the expressions above, the matter-radiation equality corresponds to the energy density relative to the critical density, and expansion factor

$$\Omega_{\mathrm{eq}} = 0.4, \quad \text{with} \quad a_{\mathrm{eq}} \simeq 1/3400. \tag{3.28}$$

A crucial input to the expressions above is the local density contrast δ_{dc} at the decoupling redshift corresponding to

$$a_{\mathrm{dc}} \approx a_{\mathrm{eq}} \left(\frac{x}{\tilde{x}} \right)^3, \tag{3.29}$$

where $x \equiv r/a$ is the comoving distance. Assuming that the PBH two-point function $\xi(x)$ is constant at comoving distances smaller than $\tilde{x}$ specified above in Eq. (3.26), i.e.

$$1 + \xi(x) \approx \delta_{\mathrm{dc}} \quad \text{for} \quad x < \tilde{x}, \tag{3.30}$$

strong clustering corresponds to $\delta_{\mathrm{dc}} \gg 1$ [190]. Following Ref. [105] and [188], in order to assess the maximal-possible merger rate we additionally consider a large local enhancement factor $\delta_{\rm local} \sim 2 \times 10^5$, linearly impacting the merger rate $R \to \delta_{\rm local} R$, reflecting the fact that the merger rate is sensitive to the *local* dark matter density – i.e. assuming that the binary density follows the dark matter density.

We note that the expression above can be *maximized* by expanding the Gamma function in the asymptotically large $\delta_{\rm dc}$ limit, by using the asymptotic expansion for the incomplete Gamma function

$$\Gamma[a,z] \simeq e^{-z} z^{a-1} \qquad z \to \infty. \tag{3.31}$$

Neglecting the second Gamma function, which is very suppressed compared to the first one for large values of the second argument, one obtains

$$\begin{aligned} \frac{dR_{E3}}{dm_1 dm_2 dm_3} &= \frac{9}{296\pi}\frac{1}{\tilde{\tau}}\left(\frac{t_0}{\tilde{\tau}}\right)^{-13/16}\left(\exp\left[-\tilde{N}\left(\frac{t_0}{\tilde{\tau}}\right)^{3/16}\right]\right)\tilde{N}^2 \\ &\quad \times \tilde{x}^{-3}\delta_{\rm dc}^{-1}\bar{m}^3\frac{\psi(m_1)}{m_1}\frac{\psi(m_2)}{m_2}\frac{\psi(m_3)}{m_3}. \end{aligned} \tag{3.32}$$

We will not use the expression above, which is, however, relevant to interpret our numerical results below.

Note that important physical effects can alter the quantitative values for the merger rates outlined above here. For instance, Ref. [189] finds, with dedicated numerical simulations, that the rates for early binaries can be importantly reduced when one excludes binaries ending up in a PBH cluster induced by Poisson fluctuations, a suppression factor that can be lower than 10^{-2} for large PBH fraction. We elaborate on this point in the Appendix, showing that even then the early 3-body pathway would dominate over all other pathways, even if the problem is amplified in scenarios with

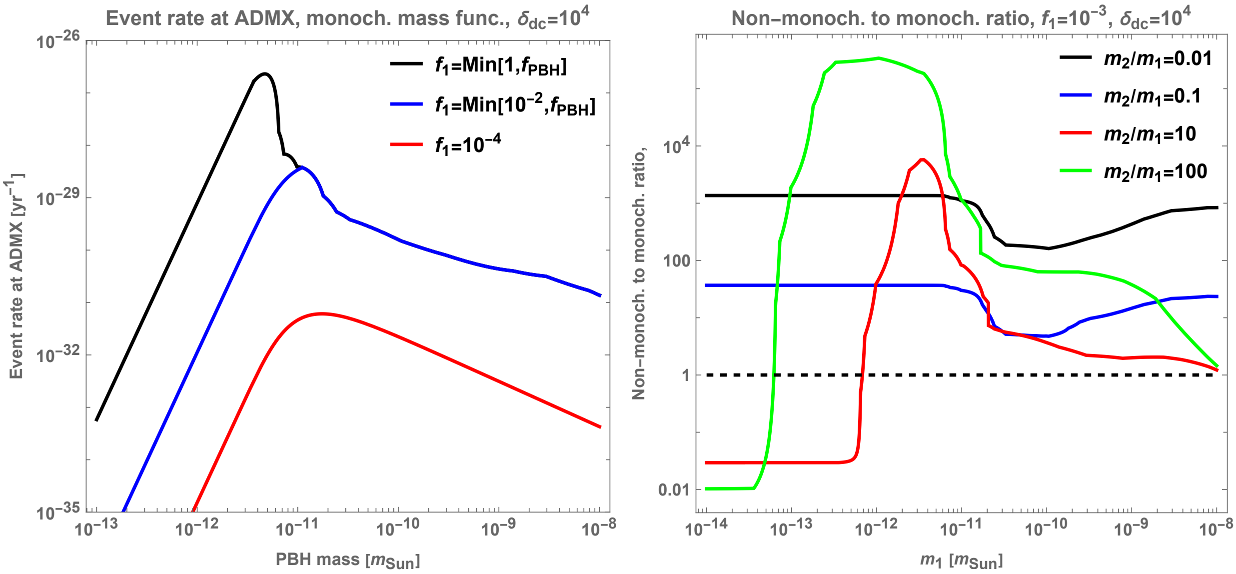

Figure 3.2: Left: rates (in events per year) for a monochromatic mass function with a fixed mass fraction f_1 (if $f_1 > f_{\rm PBH}$, i.e. when the assumed mass fraction violates observational limits, we set the mass fraction to its maximal possible value $f = f_{\rm PBH}$). Right: the ratio, as a function of m_1, of event rates for dichromatic mass functions with a given $0.01 \le m_2/m_1 \le 100$, to the monochromatic case, for $f_1 = 10^{-3}$ and for $\delta_{\rm dc} = 10^4$

initial clustering, whereby regular interactions in clusters would reduce the predicted rates. Additionally, we note here that the PBH halos induced by a strong initial clustering could be dynamically instable and expand, as this is the case for PBH clusters induced by Poisson fluctuations, albeit it is unclear if this effect would impact the early 3-body merger rates. We give a detailed discussion of these effects, as well as the issue of how clustering of PBHs is constrained by direct methods such as microlensing and GWs, in Appendixes 7.2 and 7.3 below.

Fig. 3.2, left, shows the merger rates for a monochromatic mass function with a fixed mass fraction f (if $f > f_{\rm PBH}$, i.e. when the assumed mass fraction violates observational limits, we set the mass fraction to its maximal possible value $f = f_{\rm PBH}$). The peak occurs at masses where (i) the constraints on $f_{\rm PBH}$ allow PBH to be the entirety of the dark matter, and (ii) the distance sensitivity plateaus at its maximal value, for the frequency under consideration.

The right panel of Fig. 3.2 shows the ratio, as a function of m_1 of event rates for dichromatic mass functions with a given mass ratio $0.01 \leq m_2/m_1 \leq 100$, to the monochromatic case, for $f_1 = 10^{-3}$ and for $\delta_{\rm dc} = 10^4$ (this latter parameter is largely uninfluential here). The plot bears out what commented on for the bottom, right panel of Fig. 3.1: large mass ratios produce enhanced rates for smaller m_1, and smaller mass ratios for larger m_1, and the event rates are generally larger for dichromatic mass functions than for monochromatic mass functions. The non-trivial features in the figure causing non-monotonic behaviors are due to a non-trivial combination of the sensitivity distance and constraints on the abundance of PBHs at a given mass.

The contour plots in Fig. 3.3 show, on a Log10 scale, the event rate per year for a putative ADMX-like cavity for a monochromatic mass function varying the binary merger mass and $f_{\rm PBH}$ (left) with $\delta_{\rm dc} = 10^4$ (left), and varying instead $\delta_{\rm dc}$ with $f = \min(f_{\rm PBH}, 1))$. The sharp decrease on the left side of the plots is related to increasingly tight constraints on the abundance of PBH at those masses; we also note that the region with the largest event rates extends over a relatively broad range of masses, $10^{-12} \lesssim m_{\rm PBH}/m_{\rm Sun} \lesssim 10^{-8}$, and down to PBH densities of about 0.1% of the dark

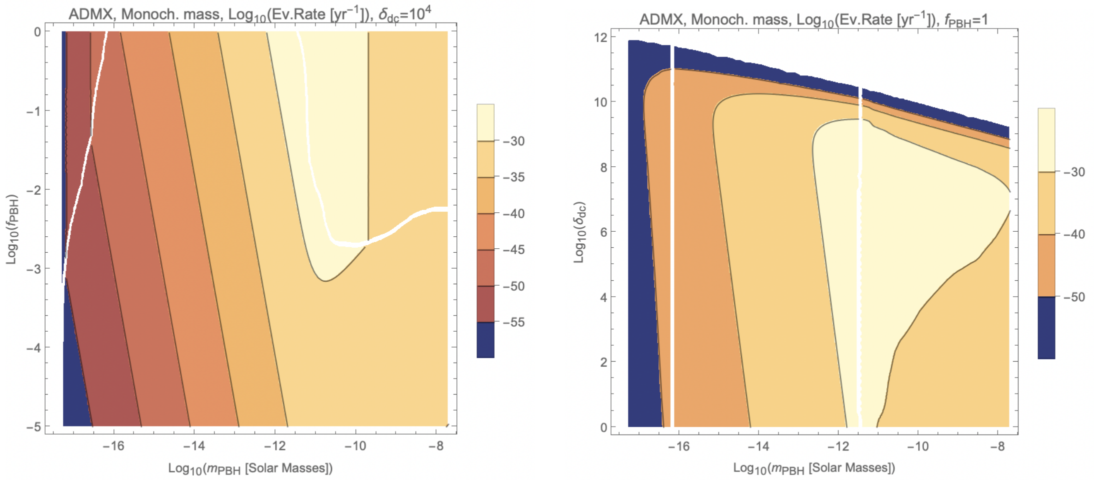

Figure 3.3: Contours of constant Log10 of the event rate per year for a putative ADMX-like cavity for a monochromatic mass function varying the binary merger mass $m_{\rm PBH}$ and f_1 with $\delta_{\rm dc} = 10^4$ (left), and varying instead $\delta_{\rm dc}$ with $f = \min(f_1, 1)$ (right).

matter density; additionally, the left panel highlights the relatively weak dependence of the event rate on $\delta_{\rm dc}$ outside of the exponential suppression at $\delta_{\rm dc} \gtrsim 10^9$, noted already e.g. in [105].

Fig. 3.4 relaxes the assumption of a monochromatic mass function, and considers, for three different values of the "weight" (relative abundance) f_1 of the population with mass m_1 (see Eq.(3.4)), the event rate as a function of the mass ratio m_2/m_1. Note that here and hereafter we assume the maximal f_2, fixed by Eq. (3.7) for $f_1 \leq_{\rm PBH} (m_1)$, otherwise the curves are interrupted as the abundance exceeds limits. In the figure, we assume $m_1 = 10^{-11}\ m_{\rm Sun}$, for $\delta_{\rm dc} = 10^4$. The right panel shows the same, but as a ratio to the monochromatic $m_2 = m_1$ case, and highlights how a non-monochromatic mass ratio predicts, for m_2 between $0.05 m_1$ and $0.9 m_1$, a larger event rate than the

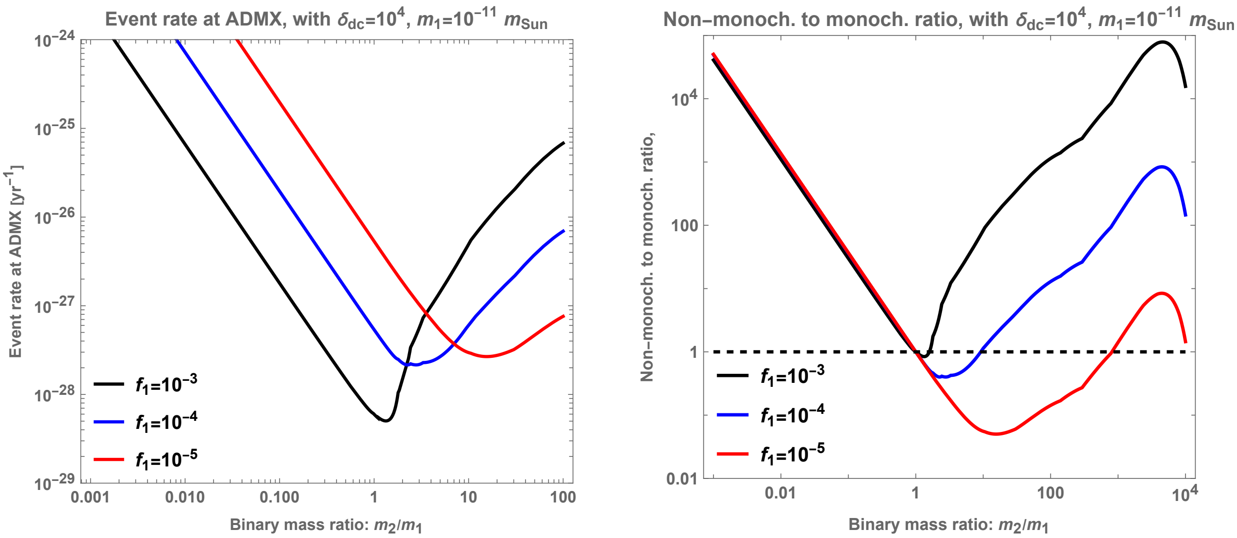

Figure 3.4: Left: Event rate at an ADMX-like cavity for different $f_1 = 10^{-3},\ 10^{-4},\ 10^{-5}$ as a function of the binary mass ratio m_2/m_1, with $m_1 = 10^{-11}\ m_{\rm Sun}$, for $\delta_{\rm dc} = 10^4$. Right: same as in the left panel, but as a ratio to the monochromatic $m_2 = m_1$ rate

monochromatic case.

Fig. 3.5 studies the effect of clustering, but for a dichromatic mass function, varying a number of parameters, and as a function of the clustering parameter $\delta_{\rm dc}$. Both panels assume $m_1 = 10^{-11} m_{\rm Sun}$; in the left panel we fix $f_1 = 0.01$ and vary the mass ratio $m_2/m_1 = 0.1,\ 1,\ 10$; in the right panel we fix, instead, $m_2/m_1 = 0.01$, and vary $f_1 = 0.01,\ 0.001$, and 10^{-4}. As usual, f_2 is maximized for a given f_1, as per Eq. (3.7). The first peak, to the left, corresponds to m_2, while the one to the right to m_1 in the dichromatic mass function.

The main takeaway point in this figure is that the effect of early clustering strongly depends on the mass ratio for a non-monochromatic mass function. Also,

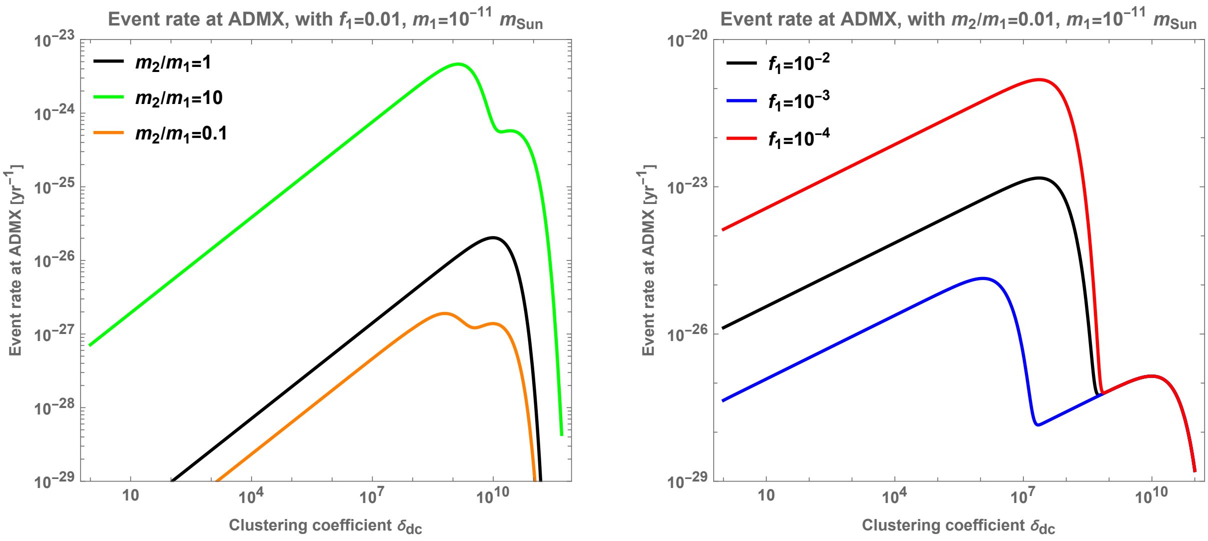

Figure 3.5: The dependence of the event rate per year on the clustering coefficient $\delta_{\rm dc}$, for three different mass ratios $m_2/m_1 = 1,\ 10,\ 0.1$ with $f_1 = 0.01$ and $m_1 = 10^{-11}\ m_{\rm Sun}$ (left) and for three different values of $f_1 = 10^{-2},\ 10^{-3},\ 10^{-4}$

notably, at very large clustering there exists an exponential suppression, deriving from the physical effect that mergers in this case have already occurred at very early times and the binaries, thus, have been exhausted into early mergers.

The contour plots of Fig. 3.6 illustrate, as above on a Log10 scale, but different color-coding, the event rate on the planes defined by the mass ratio m_2/m_1 versus clustering coefficient $\delta_{\rm dc}$ (left) and f_1 (right). In both plots we fix $m_1 = 10^{-12} m_{\rm Sun}$, while we set $f_1 = 10^{-5}$ in the left panel and $\delta_{\rm dc} = 10^4$ in the right panel. Note that as in previous plots, f_2 is maximized for a given f_1. The white regions in the left panel have vanishingly small event rates because of the exponential suppression discussed above.

The left panel highlights that the maximal rate, approximately independent

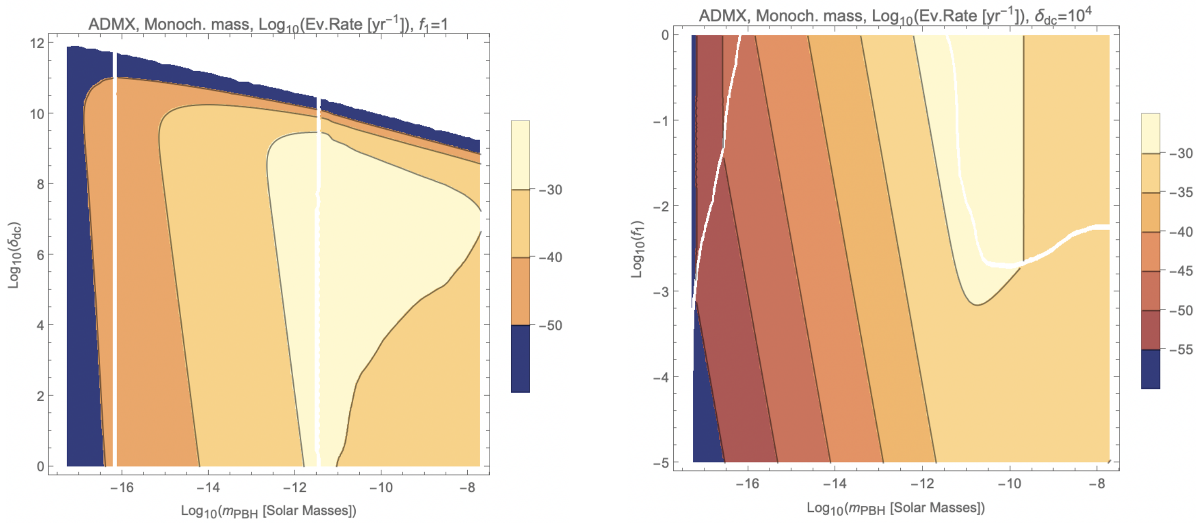

Figure 3.6: Contour plots of the Log10 of the event rate per year as a function of the dichromatic mass function ratio m_2/m_1, for $m_1 = 10^{-12}\, m_{\rm Sun}$, versus $\delta_{\rm dc}$ with $f_1 = 10^{-5}$ (and maximized f_2) (left) and f_1 (right) with $\delta_{\rm dc} = 10^4$.

of m_2/m_1, is achieved for $10^1 \lesssim \delta_{\rm dc} \lesssim 10^4$, and for the smallest mass ratios m_2/m_1. The right panel, instead, shows how the event rate is maximized at extreme mass ratios $m_2/m_1 \gg 1$ at large f_1, or $m_2/m_1 \ll 1$ at small f_1.

In the final four panels of Fig. 3.7 we study the event rate, on the plane defined by the masses in the merging binary m_1 and m_2. In all panels we fix $\delta_{\rm dc} = 10^4$, and, clockwise from the top left, we set $f_1 = 10^{-1}$, 10^{-3}, 10^{-5} and 10^{-10}. The white regions correspond to $f_1 > f_{\rm PBH}(m_1)$ and are therefore excluded by constraints on the PBH abundance at that mass.

In all cases, we find that rates are largest in the range around $10^{-13} \lesssim m_{1,2}/m_{\rm Sun} \lesssim 10^{-11}$, with peaks at values slightly in excess of $10^{-12} m_{\rm Sun}$. The details

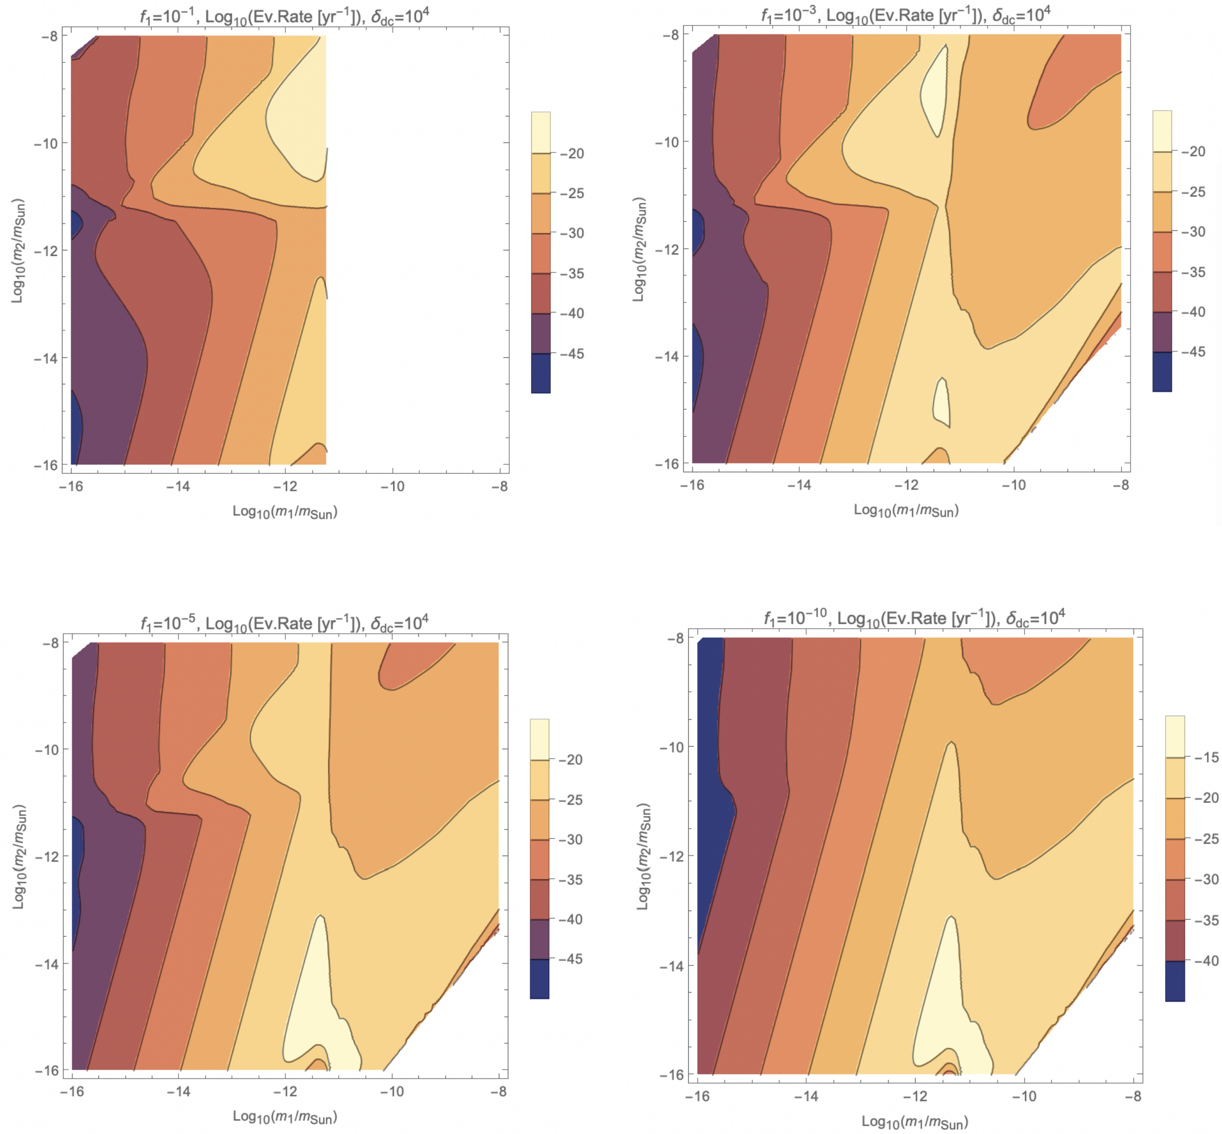

Figure 3.7: Four panels showing the Log10 contours for the yearly event rate on the (m_1, m_2) plane, all at $\delta_{\rm dc} = 10^4$, for $f_1 = 10^{-1},\ 10^{-3},\ 10^{-5},\ 10^{-10}$ in clockwise order from the top left

of each panel are interesting and straightforward to interpret based on the discussion of the plots above. Note that the color range is not the same for all plots, and that the event rate is increasingly large with decreasing f_1.

3.5 Comparison of different detectors and frequencies

We compare, here, the rates expected with improved, future detectors, sensitive to a broader range of frequencies. Specifically, we consider a version of ADMX with a significantly reduced resonant frequency of $\nu = 0.65$ GHz (corresponding to ADMX run1A; note that for instance the recent run1D is around 1.4GHz), with all other parameters set to the default values indicated above. For all detectors we assume 60 seconds integration times.

For the SQMS detector we employed the parameters indicated in Ref. [35]: $f \in (1-2)$ GHz, $Q \sim 10^6$, $|\mathbf{B}| = 5$ T, $V_{\rm cav} = 100$ L, $T_{\rm sys} = 1$ K, and $\eta_n = 0.1$ as for ADMX.

For the ADMX-EFR we utilize projections provided by the ADMX Collaboration[1], featuring a magnetic field – 9.4 T at central field (mean over cavities on the order of 9.1 T), a cavity volume consisting of an 18 cavity array with combined volume of approximately 218 liters in the frequency range 2-3 GHz and 182 liters (3-4 GHz), and a Cavity Unloaded Quality factor of 120k; the total system noise temperature was assumed to be 440 mK.

Finally - albeit the geometric structure of the detector may pose non-trivial coupling differences from the linear structure of ADMX - for DM Radio GUT we assume

[1] G.P. Carosi, private communication

the parameters given in[2] and also outlined in [82] (assuming the sensitivity is approximately reflected in our Eq.(3.8) above): a magnetic field of 16 T, a volume of 4,580 L, a quality factor of 10^6, a system noise temperature of 1 mK, and a frequency of 4 GHz.

We present a comparison of the expected, maximal yearly event rate with these various setups in Fig. 3.8. Generally, as expected, higher frequencies of operation extend the detectors sensitivity to lower masses, and increase the range of masses where a significant event rate should occur (see especially the right panel). The scaling of detectors performance in terms of the detector parameters is otherwise clearly evident from the expression in Eq. (3.8).

Unfortunately, even in the most optimistic setup under consideration, the event rate at the most promising future resonant cavity setup is found to be below 10^{-17} events per year.

3.6 Discussion

The realization that a number of devices originally conceived for different experimental purposes could be used in the quest for high-frequency GWs has spurred a resurgent interest in potential signals [9]. Parallel to this, as the search for weakly-interacting particle dark matter dawns [18] with no compelling evidence of it being realized in nature has also given rise to renewed interest in primordial black holes as dark matter candidates. Interestingly, a broad black hole mass range, between, approximately, masses of 10^{-16} and 10^{-10} solar masses (the "asteroid mass range") is

[2]https://indico.fnal.gov/event/63051/

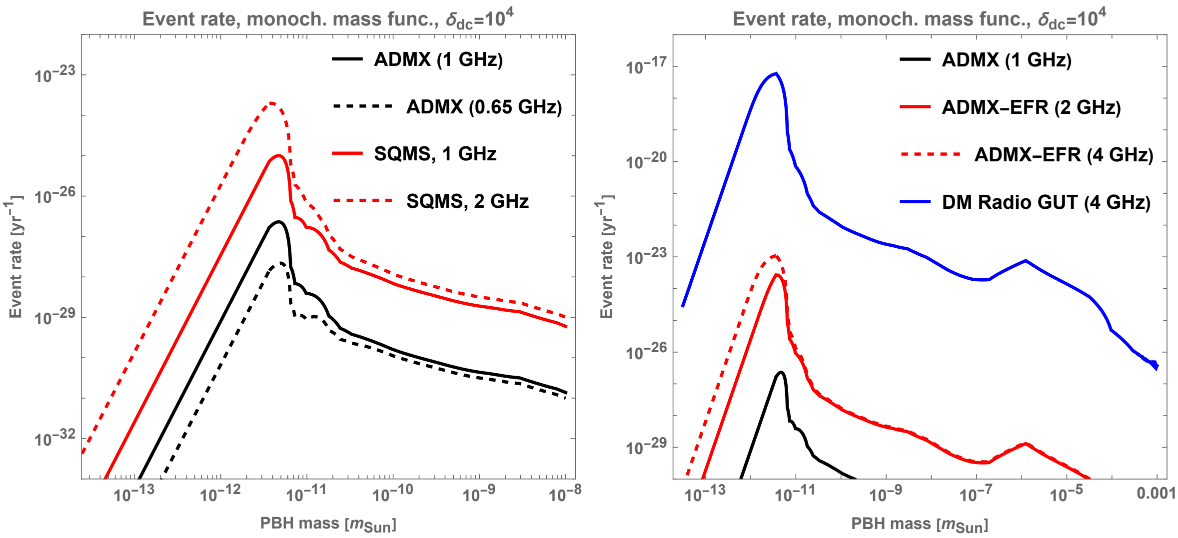

Figure 3.8: Left: A comparison of the yearly event rate at different frequencies for ADMX (1 and 0.65 GHz) and for SQMS (1 and 2 GHz); Right: The same as in the left panel, but for ADMX-EFR at 2 and 4 GHz, and for DM Radio GUT at 4 GHz.

compatible with these objects being the entirety of the cosmological, and Galactic, dark matter [53, 51].

The asteroid mass window is, however, extremely complicated to probe. At low masses, evaporation via Hawking-Bekenstein radiation is too slow; at high masses microlensing is not effective because of a number of technical reasons related to finite size source effects, wave optics, and the interplay of event duration and observational cadence. However, mergers of such light black holes could give rise to GHz frequency gravitational wave signals similar to those detected, at much lower frequency, by the interferometers LIGO/Virgo/KAGRA, and potentially detectable by the above-mentioned re-purposed detectors.

Here, we provided an in-depth examination of the expected event rate for light PBH mergers at experiments similar to the microwave resonant cavity ADMX. Our assumptions are purposely optimistic, stretching the event rates to their maximal possible values. This includes maximizing the effect of clustering upon binary formation, the local density of Galactic dark matter, and the black hole mass function, the latter being customarily and systematically assumed to be monochromatic in previous studies (see e.g. [105, 82, 109]).

We found that even with future experiments, and with the most-optimistic possible assumptions, the event rate at resonant cavities is vanishingly small, due to the experiments' sight distance being at most a fraction of an AU.

This study addressed closely the role of a non-monochromatic mass function, showing that event rates can be enhanced by orders of magnitude compared to the usually-assumed monochromatic case; everywhere, we took into consideration the relevant observational constraints on the black hole abundance, as they apply to non-trivial mass functions. We first assessed the most significant pathway to binary formation, and concluded that throughout the parameter space of interest it is associated with early, three-body binary formation. We then studied how event rates depend on a "dichromatic" mass function - one that features two subpopulations of different masses in different proportions, which has been shown elsewhere to maximize the merger event rate [151]. We also studied the effect of clustering at early (and late) times, and prospects for current and future detectors alike. Finally, we demonstrated that the sight distance to black hole merger events is mass-independent for any narrow-band detector; additionally, we demonstrated that the resonant cavity ring-up condition corresponds to the number of cycles spent by the signal within the bandwidth being equal to the cavity's quality factor.

Chapter 4

Dark sector modifications to Kerr and Reissner-Nordström black hole evaporation

4.1 Introduction

Chapters 2 and 3 established that a large dark sector accelerates black hole evaporation and that this effect provides the strongest quasi-model-independent bounds on gravitationally coupled dark degrees of freedom. Those analyses treated evaporation primarily as a mass-loss process, implicitly assuming Schwarzschild-like geometry. This chapter generalizes the framework to rotating (Kerr) and charged (Reissner–Nordström) black holes, where mass, angular momentum, and charge are all dissipated simultaneously, and where the presence of a large dark sector can qualitatively restructure the evaporation hierarchy.

Page factors [174, 176] provide the quantitative framework for describing these loss rates within the semi-classical approximation, which assumes a fixed background metric with quantum mechanical evolution of fields and successfully describes evaporation of sufficiently large black holes where back-reaction effects remain negligible [121]. The central question this chapter addresses is whether the canonical evaporation ordering—charge dissipated first, then angular momentum, then mass—is robust to the addition of a large dark sector, or whether a sufficiently large N can reverse this hierarchy and produce qualitatively different endpoint geometries.

4.2 Black Hole Types and Evaporation Mechanisms

Black holes are characterized by mass (M), angular momentum (J), and charge (Q). Schwarzschild black holes, being non-rotating and uncharged, lose only mass through isotropic thermal radiation [174]. Kerr black holes introduce rotational complexity: frame-dragging effects produce anisotropic emission, with angular momentum typically lost much faster than mass through higher-spin particle emission, driving a "spin-down" toward Schwarzschild-like states [174, 175, 205]. RN black holes, charged but non-rotating, preferentially emit oppositely charged particles, dissipating charge on much shorter timescales than mass when such particles are thermodynamically available (i.e., when their mass is comparable to or below the hole's temperature) [176, 125].

These evaporation patterns critically depend on the particle spectrum. The existence of degrees of freedom beyond the Standard Model (SM) can substantially modify the dissipation hierarchy described above. Furthermore, greybody factors—which quantify deviations from ideal blackbody spectra due to scattering in curved spacetime [171, 85]—play a significant role in determining emission rates and energy partitioning across particle types and modes.

4.3 Comparative Analysis and Extremal Effects

While individual black hole types have been studied extensively, systematic comparative analyses using Page factors remain sparse. Schwarzschild black holes evaporate slowest due to isotropic Hawking radiation [174, 162]. Kerr black holes lose angular momentum several times faster than mass, with superradiance effects enhancing energy emission from ergoregions [175, 162]. RN black holes rapidly shed charge, subsequently suppressing radiation through Coulomb potential barriers [176].

High charge-to-mass ratios (Q/M) in RN black holes induce distinctive phenomena including Schwinger pair production and quantum charge fluctuations near event horizons [85, 125]. Near-extremal configurations—whether Kerr or RN black holes approaching physical limits—exhibit vanishing surface gravity, suppressing thermal emission within semi-classical approximations, though quantum corrections modify this behavior [44, 122]. Additionally, superradiance enables light bosons to form "gravitational atoms" around rotating black holes, anomalously accelerating angular momentum loss for specific combinations of spin and boson mass.

4.4 Motivation and Scope

Despite detailed studies of individual black hole evaporation characteristics, significant gaps remain in understanding how mass, angular momentum, and charge evolve comparatively across black hole types. Key questions include:

- How do Page factors for energy, angular momentum, and charge loss differ quantitatively across Schwarzschild, Kerr, and RN black holes?
- What are the relative timescales for losing conserved quantities, and how do greybody factors influence these dynamics?
- How do near-extremal conditions (high a/M or Q/M) modify evaporation processes?
- How do degrees of freedom beyond the Standard Model alter dissipation patterns?
- What additional mechanisms can modify relative loss rates of charge, angular momentum, and mass?

This study quantitatively addresses these questions. Section 4.5 describes Page factor evolution for light black holes, outlining computational procedures for time-evolved factors with spin and charge evolution, including effects of dark degrees of freedom. Section 4.6 presents new results and discusses their broader implications. Section 4.7 examines the effect of a large dark sector on the evolution of black hole parameters during evaporation. Section 4.10 discusses how relative dissipation of mass, charge, and spin varies with additional particle species. Section 4.11 examines time evolution under Schwinger pair production, which significantly impacts light charged black holes. The penultimate section analyzes how superradiance and gravitational atom formation affect angular momentum extraction and Kerr black hole lifetimes. Section 4.13 presents final considerations and conclusions. Throughout, we work in natural units where $G = \hbar = c = 1$.

4.5 Page Factor Evolution

The evolution of intrinsic BH parameters such as spin and mass can be determined by the so-called Page factors, as outlined in [175]. Generally, the lifetime of a BH scales as the black hole mass to the third power, M^3, given that the evaporation rate follows the classical Stefan-Boltzmann equation, and area is proportional to $R_S^2 \sim M^2$, R_S being the black hole Schwarzschild radius, and $T_H \sim M^{-1}$, with T_H the Hawking temperature; as a result, the scale-invariant Page factor for the BH mass is defined as

$$f(M, x^*) = -M^3 \frac{d\ln(M)}{dt} = -M^2 \frac{dM}{dt} \tag{4.1}$$

where x^* corresponds to the intrinsic spin a^* or charge Q^*, depending on if we are describing the Kerr or RN metrics. We will use here the effective dimensionless quantities $a^* = \frac{a}{M} = \frac{J}{M^2}$ and $Q^* = \frac{Q}{M}$.

The Page factor for angular momentum is defined as

$$g(M, a^*) = -M^3 \frac{d\ln(J)}{dt} = -\frac{M}{a^*} \frac{dJ}{dt} \tag{4.2}$$

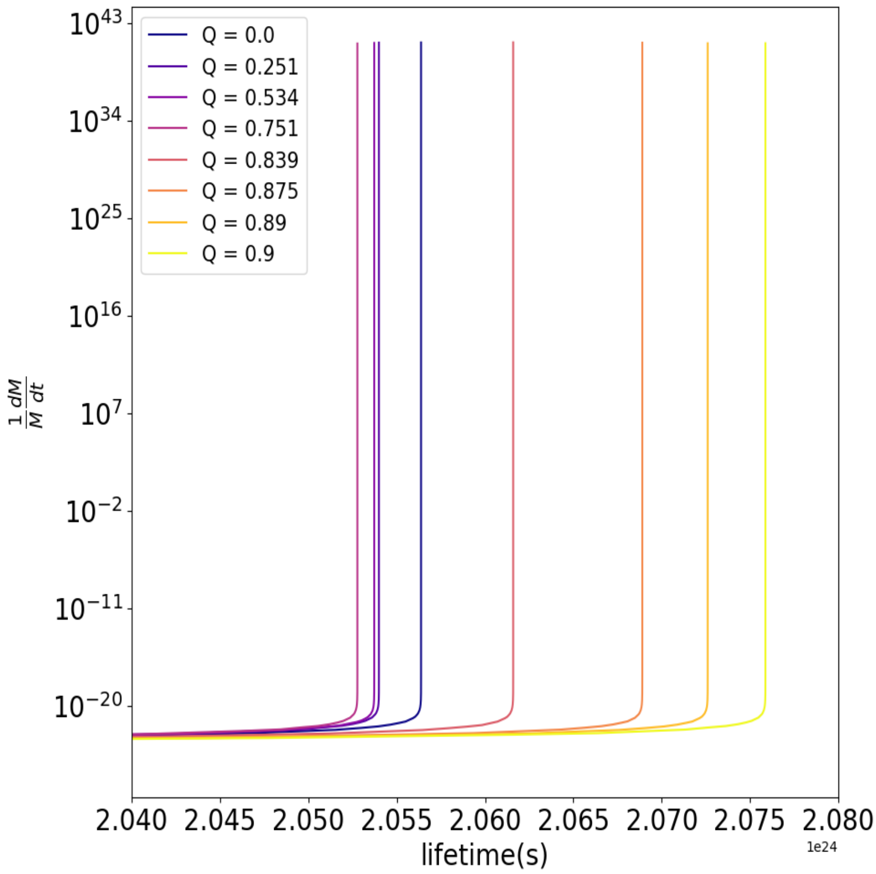

Figure 4.1: Lognormalized mass loss rates as a function of lifetime for different effective charge. Increasing effective charge lowers the overall temperature of the BH, slowing particle emission. The initial BH mass is set at 10^{17} g as larger masses will not significantly change the initial mass loss rate.

To describe charge evolution, we define the additional scale-invariant quantity

$$h(M, Q^*) = -M^3 \frac{d\ln(Q)}{dt} = -\frac{M^2}{Q^*}\frac{dQ}{dt}. \tag{4.3}$$

The rates of particle emission are described as in Ref. [120] and give

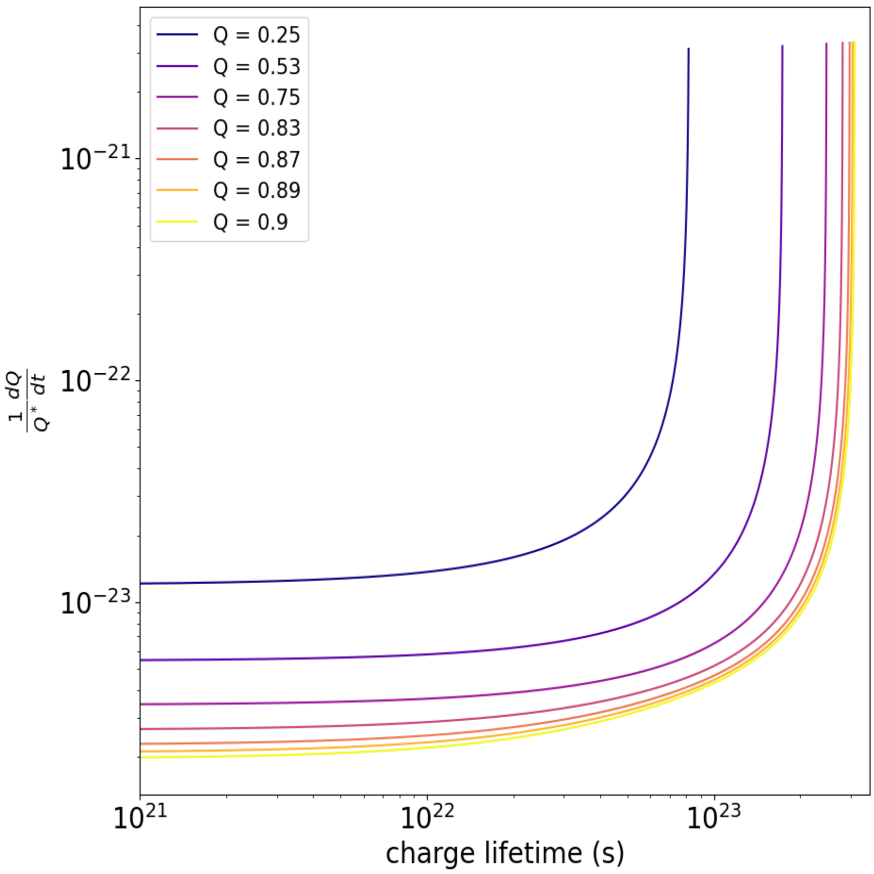

Figure 4.2: Charge loss rate due to Hawking radiation for various spins normalized by the initial effective charge Q^* of a BH.

$$f(M, x^*) = -M^2 \frac{dM}{dt} = -M^2 \int_0^\infty E \sum_i \frac{d^2 N_i}{dtdE} dE, \tag{4.4}$$

$$g(M, a^*) = -\frac{M}{a^*} \frac{dJ}{dt} = -\frac{M}{a^*} \int_0^\infty \sum_i g_i \sum_{l,m} \frac{d^2 N_{l,m}}{dtdE} dE, \tag{4.5}$$

where the terms

$$\frac{d^2 N_{l,m}}{dtdE} = \frac{1}{2\pi} \frac{\Gamma_{s,lm}(M, E, x^*)}{e^{E/T} - (-1)^{2s}}, \tag{4.6}$$

$$\frac{d^2N_i}{dtdE} = g_i \sum_{l,m} \frac{d^2N_{l,m}}{dtdE} \tag{4.7}$$

describe the total number of outgoing particles created within the BH horizon, as described in [120]. This leads us to an analogous Page factor describing charge evolution from Eq. (4.3):

$$h(M, Q^*) = -\frac{M^2}{Q^*}\frac{dQ}{dt} = -\frac{M^2}{Q^*}\int_0^\infty \sum_i q_i g_i \sum_{l,m} \frac{d^2N_{l,m}}{dtdE} dE. \tag{4.8}$$

The terms q_i are the individual particle charge and g_i are available particle degrees of freedom. Note that the term $h(M, Q^*)$ does not include a double counting of the number of degrees of freedom from including antiparticles, as antiparticles and particles included would have opposite charge. We define the effective charge as being from 0 to 1, as the process for shedding negative charge is identical to the scenario of shedding positive charge aside from a sign change.

The computations in the following section are done using *BlackHawk* [17] for the Kerr metric. However, *BlackHawk* does not account for time evolution of a RN BH and a charged metric adaptation of the code was written to overcome this issue. The analysis code used in this work is openly available online[1]. In our consideration of the charged metric, we have adopted the optical limit approximation to compute the particle's greybody factor:

$$\Gamma = 27M^2E^2, \tag{4.9}$$

[1] `https://github.com/cewasiuk/Spin-and-Charge-BH-Evolution`

where E is the total energy of the emitted particle, taken to be extremely large in comparison to an individual particles rest mass. This approximation becomes increasingly accurate in the high-frequency regime ($\omega M \gg 1$), and is therefore suited to the extremely light black holes considered in our study, whose high Hawking temperatures push the emission spectrum toward the geometric-optics limit. While not exact, this approach captures the qualitative behavior of RN evaporation in the presence of additional massless degrees of freedom. The optical-limit treatment leads to modest deviations in the Page factors and produces an overall shortening of the total lifetime relative to the fully tabulated greybody factors. Although our approximation slightly overestimates the evaporation rate, it accurately reproduces the qualitative trends relevant for exploring how large hidden sectors modify RN black hole evolution.

4.6 New Results and Physical Implications

In this work we identify several new features in the standard evaporation hierarchy of black holes—where charge is typically lost first, followed by spin and finally mass—when the evolution is extended to include large numbers of beyond–Standard Model particle species. Our principal results are as follows:

1. **Reversal of charge-loss hierarchy**

   For Reissner–Nordström black holes coupled to large, massless hidden sectors ($N_{\mathrm{dof}} \gtrsim 100$), the standard evaporation hierarchy is inverted. The mass-loss rate exceeds the charge-loss rate,

$$\left|\frac{\dot{M}}{M}\right| > \left|\frac{\dot{Q}^*}{Q^*}\right|,$$

leading to effective charge growth during evaporation. This contrasts with previous expectations that charged black holes always neutralize first.

2. **Critical threshold for reversal**

   We derive a quantitative condition for this inversion, showing that it occurs when

   $$N_{\mathrm{dof}} > N_{\mathrm{crit}} \approx 220,$$

   for beyond Standard model massless degrees of freedom, beyond which energy loss to uncharged species dominates over electromagnetic discharge.

3. **Delayed or halted neutralization**

   In the high-N_{dof} regime, black holes can approach near-extremal states ($Q^*/M \to 1$) without rapid charge loss. These configurations exhibit prolonged lifetimes due to suppressed Hawking temperatures and reduced Schwinger emission.

4. **Interplay with rotation**

   For Kerr configurations, spin-down remains dominant, but frame-dragging modifies the effective emission balance, coupling angular momentum loss to charge evolution. This coupling introduces a correlated decay track in the (a^*, Q^*) phase space not captured in prior treatments.

5. **Reemergence of discharge via Schwinger pair production**

   In the case of a large dark sector or light extremal black holes ($M \lesssim 10^{17}$ g), quantum electrodynamic discharge again dominates, restoring neutrality and terminating the reversed regime.

6. **Superradiant extraction of angular momentum**

   We incorporate the effects of black hole superradiance into the evaporation hierarchy, showing that massive bosonic fields can remove angular momentum far more efficiently than Hawking radiation.

7. **Numerical confirmation**

   Numerical solutions to the coupled Page-factor equations confirm these behaviors, delineating regions in (M, Q^*, a^*) space where charge growth occurs and illustrating the modified evaporation trajectories under large hidden-sector particle counts.

Together these results demonstrate that the canonical evaporation sequence, charge loss preceding mass loss, is not inherently universal but depends sensitively on the particle content of the theory. Large hidden sectors can qualitatively change black hole evolution, producing transiently charged or near-extremal relics with potential implications for early-universe black hole populations and dark sect phenomenology.

4.7 Effect of a large dark sector on parameter evaporation

Introducing a dark sector containing a large number of degrees of freedom results in significant changes in the evolution of charged or rotating BHs. The effective spin and charge of a BH are defined as $a^* = \frac{J}{M^2}$ and $Q^* = \frac{Q}{M}$, where an effective spin or charge equal to 1 corresponds to extremality where evaporation shuts off. The total

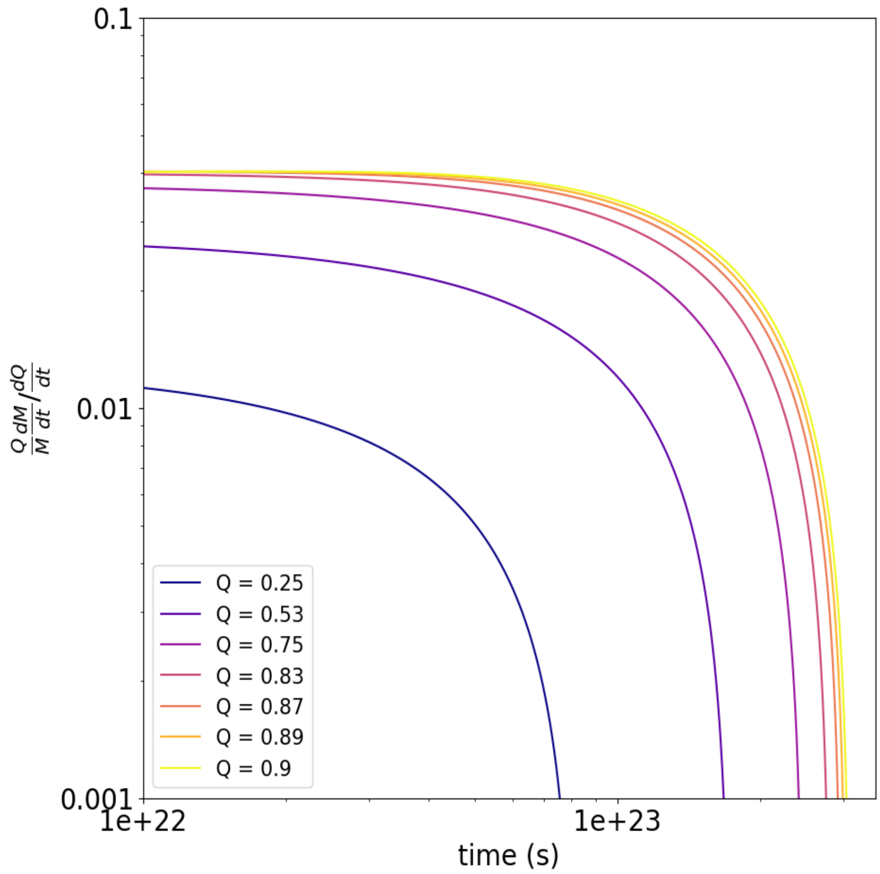

Figure 4.3: Lines of $\frac{\frac{1}{M}dM/dt}{\frac{1}{Q^*}dQ^*/dt}$ lognormalized for a RN BH for various effective charge values. Initial conditions are for a $M = 10^{17}g$ BH which evaporate only standard model particles.

number of particles needed to neutralize a charged black hole are of order M, whereas the number needed to carry off overall angular momenta J scales as M^2. Therefore the standard hierarchy of parameter loss, when considering purely evaporative processes, begins with charge neutralization, followed by loss of spin and finally mass evaporation which, as noted above, scales as M^{-2}.

Rather than comparing baseline cases, we target regimes that perturb the hierarchy itself. While introducing a large number of dark sect degrees of freedom may at first appear unconventional, such scenarios are not unconstrained. A variety of observational and experimental considerations already place bounds on relativistic or metastable hidden sectors, including limits from high-energy cosmic rays, collider searches, supernova cooling, precision tests of gravity, and the evaporation of stellar-mass black holes. The feasibility of large dark sectors consistent with these constraints has been examined in our previous work, and we refer the reader to Ref. [97] for a detailed discussion. In this context, black hole evaporation provides a complementary and largely model-independent probe of such sectors.

Starting with charged black holes, we include a substantial population of uncharged dark sect species and quantify the parameter ranges where $f(M, Q^*)$ surpasses $h(M, Q^*)$, signaling hierarchy reversal. The very large lifetimes shown in several figures reflect the suppression of Hawking emission for massive or near-extremal black holes and are intended to illustrate relative scaling and hierarchy modification; any potential cosmological relevance instead arises for sufficiently light black holes, for which the same mechanisms operate on much shorter, early-universe–relevant timescales.

4.8 Reissner-Nordström Evaporation Rates

The overall effective charge of a black hole (BH) shrinks both the horizon and Cauchy radii according to

$$r_{\pm} = r_{H,C} = r_S \frac{1 \pm \sqrt{1 - Q^{*2}}}{2}, \tag{4.10}$$

where $r_S = 2M$. This modifies the black hole temperature,

$$T_{RN} = \frac{r_+ - r_-}{4\pi r_+^2}, \tag{4.11}$$

and thus alters the emission spectrum. As T_{RN} increases, the probability of emitting particles of a given rest mass rises, as reflected in Eq. (4.6).

Figure 4.1 shows the extension of lifetime as effective charge increases for a given black hole mass. This occurs because a larger effective charge reduces the temperature, thereby suppressing massive particle emission. In this case only standard model particles are considered.

At the same time, stronger charge enhances the absolute rate of charge loss. When this rate is normalized by the effective charge, however, we find that black holes with smaller initial Q^* lose charge more efficiently. This behavior, illustrated in Figs. 4.2 and 4.3, explains why highly charged black holes require longer timescales to neutralize. For sufficiently light black holes capable of thermally emitting electrons, the charge-loss rate can exceed the mass-loss rate, indicating that charge depletion dominates the early evaporation phase.

To compare the relative rates of mass and charge loss, we define the log-

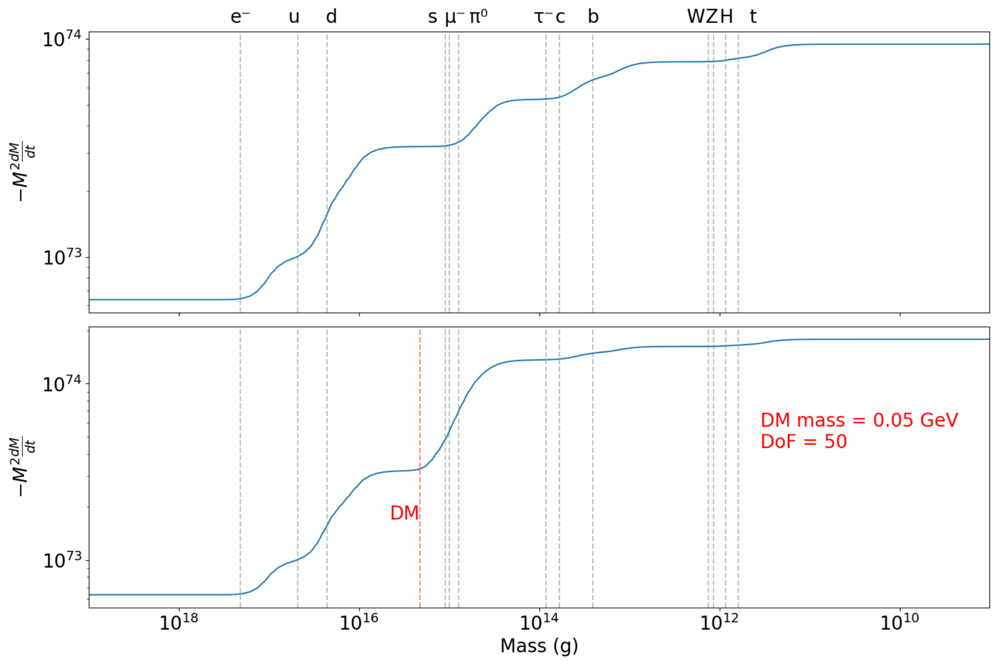

Figure 4.4: Approximate black hole mass threshold for massive particles to significantly contribute to the normalized mass loss rate. All massive particles have a nonzero emission rate at all times, but are exponentially suppressed by their respective greybody factor. Significant emission occurs when $T_{BH} \approx m_p$. Influence of a light, neutral dark sector on the normalized mass loss rate is included to illustrate its effect.

normalized ratio (LNR)

$$\frac{\frac{1}{M}\frac{dM}{dt}}{\frac{1}{Q^*}\frac{dQ^*}{dt}}, \tag{4.12}$$

which tracks the time evolution of both quantities on a dimensionless scale. The effective charge derivative is given by

$$\frac{dQ^*}{dt} = \frac{1}{M}\frac{dQ}{dt} - \frac{Q^*}{M}\frac{dM}{dt}, \tag{4.13}$$

taking care to distinguish between the true charge-loss rate $\frac{dQ}{dt}$ and the effective charge-loss rate $\frac{dQ^*}{dt}$. Written in terms of the Page factors $f(M, Q^*)$ and $h(M, Q^*)$ defined in Eqs. (4.4) and (4.8), Eq. (4.13) becomes

$$\frac{dQ^*}{dt} = Q^* \frac{f(M, Q^*) - h(M, Q^*)}{M^3}. \tag{4.14}$$

Our analysis focuses on the relative scaling between the mass and charge-loss rates. Both Page factors depend on the number of available degrees of freedom of the emitted species. However, the Page factor associated with charge, $h(M, Q^*)$, only receives contributions from charged particles. As a result, the emission of neutral species enhances the mass and spin Page factors while leaving $h(M, Q^*)$ unchanged.

In order to invert the typical evaporation sequence in which charge neutralization precedes mass loss, we introduce massless, neutral degrees of freedom to create a regime in which the black hole loses mass faster than charge. The introduction of massless particles is chosen so that the mass loss rate can be enhanced prior to the emission of any charged particles, such as the electron. The black hole mass is taken to be relatively light, at 10^{17} g—approximately the scale at which the emission of massive particles becomes significant. The corresponding mass thresholds at which all Standard Model species are emitted at non-negligible rates is included for reference in Fig. 4.4. We find that if the number of new degrees of freedom is less than $N \approx 200$, $\frac{dQ^*}{dt}$ remains negative and the black hole continues to evaporate normally without any effective charge growth. The impact of introducing dark sectors of this scale are illustrated in Fig. 4.5. At early times, the true charge-loss rate $\frac{dQ}{dt}$ stays nearly constant (extremely small, but nonzero due to a nonzero emission probability for any massive particle), while the mass-loss rate $\frac{dM}{dt}$ increases as new dark sect species contribute to evaporation. As the black hole evolves, charged particle emission becomes more frequent, enhancing $\frac{dQ}{dt}$ and producing the downward trend in the log-normalized ratio. For the modest number of additional degrees of freedom shown in Fig. 4.5, mass loss continues to dominate over charge loss throughout the evolution.

For the parameters considered, the LNR (Eq. (4.12)) approaches unity when the normalized mass and effective charge loss rates become comparable. In our simulation this occurs at the upper limit of the degrees of freedom considered above, or $N = 200$. Around this value charge loss and mass loss proceed at similar rates, leading the effective charge Q^* to temporarily plateau.

Degrees of freedom higher $N \simeq 200$ always result in some duration of a BH's lifetime where $\frac{dM}{dt}$ is dominant in comparison to $\frac{dQ}{dt}$. When this number of additional degrees of freedom are introduced, some portion of the BH's lifetime will be spent increasing effective charge. Additionally, we also find that the entirety of a BH's lifetime will be spent with effective charge increasing for degrees of freedom greater than $N \simeq 270$, which would result in a maximally charged BH with a temperature that approaches zero. We can verify this by looking at the limit as $N \to \infty$ for Eq. (4.13)

$$\lim_{N\to\infty} \big(\frac{1}{M}\frac{dQ}{dt} - \frac{Q^*}{M}\frac{dM}{dt}\big) \approx -\frac{Q^*}{M}\frac{dM}{dt}, \tag{4.15}$$

since increasing the number of massless, neutral degrees of freedom will only increase $\frac{dM}{dt}$. The LNR will then reduce to :

$$\frac{\frac{1}{M}\frac{dM}{dt}}{\frac{1}{Q^*}\frac{dQ^*}{dt}} = \frac{\frac{1}{M}\frac{dM}{dt}}{\frac{1}{Q^*}\big(-\frac{Q^*}{M}\frac{dM}{dt}\big)} = -1 \tag{4.16}$$

for large values of $\frac{dM}{dt}$. This effect can be seen within Fig. 4.6, where the values for $\frac{dQ^*}{dt}$ transition from negative to positive. This causes the effective charge of the BH to increase after the introduction of a large number of new, non-charged degrees of freedom. For introduced degrees of freedom that allow $f(M, Q^*)$ to become comparable to $h(M, Q^*)$, the effective charge neutralization of the BH stagnates and both true charge and mass evaporate at comparable rates.

For a BH of mass 10^{17}g, we estimate that number of massless degrees of freedom required hovers between roughly $220 \leq N_{dof} \leq 270$ and can be seen in Fig. 4.6. The number of degrees of freedom required for this transition is BH mass dependent, again seen through Eq. (4.13). Any introduced number of degrees of freedom larger than this, for a given BH mass, will result in a scenario where the effective charge of the BH will increase. This upsets the traditional hierarchy where charge will be evaporated away first. Varying the mass of the included dark matter serves to increase the BH lifetime; i.e. as the mass of the dark matter increases, the temperature threshold for dark matter evaporation increases, resulting in a longer duration spent radiating non-DM particles.

In this setup we have considered relatively light BHs that can thermally emit massive particles. For smaller masses, the Hawking temperature rises rapidly, allowing emission of progressively heavier species and broadening the emission spectrum. Conversely, heavier black holes have much lower temperatures, limiting emission to the lightest or massless particles—photons and neutrinos—while heavier ones are exponentially suppressed. Most of the lifetime of such massive black holes is therefore spent radiating massless degrees of freedom, which reduces mass but not charge. In this

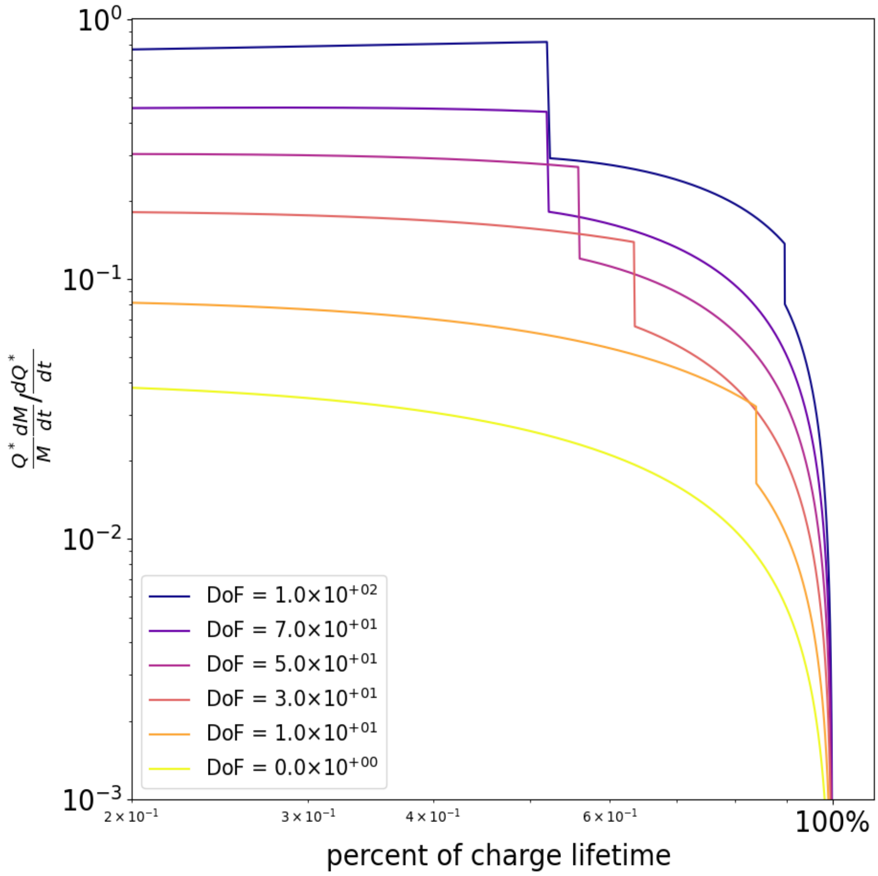

Figure 4.5: Illustration of the effects of adding spin 0, chargeless DM to the BH evaporation spectrum while varying the number of degrees of freedom. Initial conditions are for a BH with mass $10^{17}g$, effective charge $Q^* = 0.9$ and introduced dark matter with $m_{DM} = 0$ GeV. Adding larger degrees of freedom will result in a transition from a regime where $\frac{1}{Q}\frac{dQ}{dt}$ is dominant to a regime where $\frac{1}{M}\frac{dM}{dt}$ begins to dominate. Such a transition is illustrated for the same initial parameters in Fig. 4.6, and occurs when the lines reach a maximum of 1. Illustration only shows rates when the BH is considered charged, afterwards the BH is considered static.

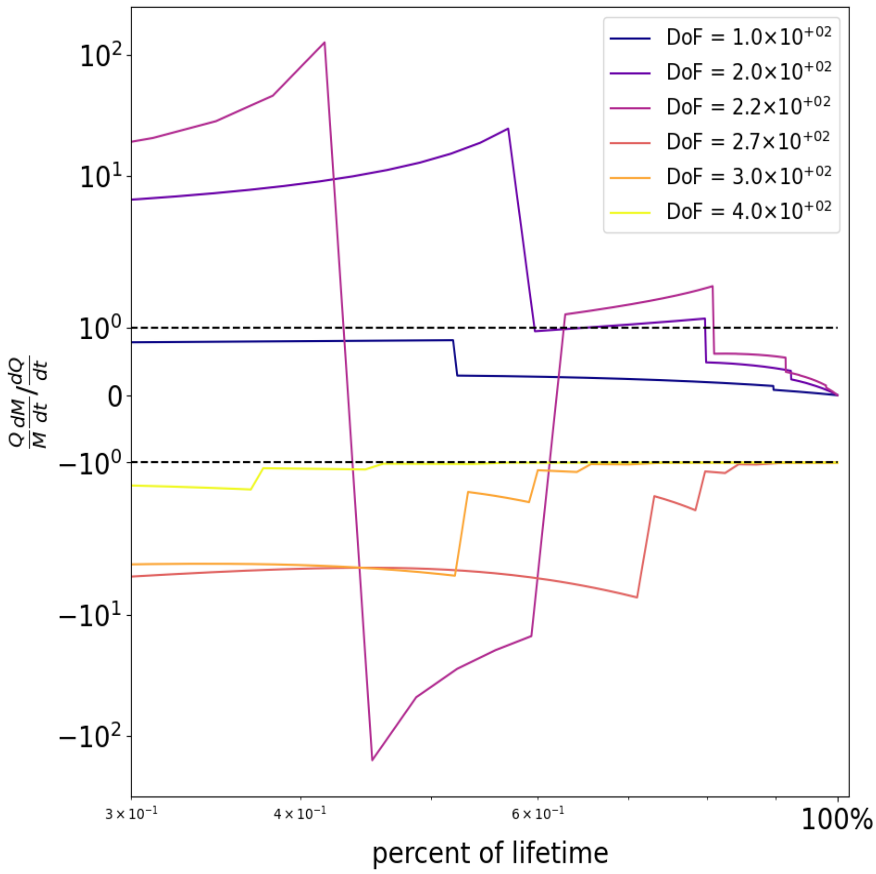

Figure 4.6: Lognormalized ratio of mass and charge loss for chargeless degrees of freedom starting when mass loss dominates and transitioning to a region where effective charge loss dominates. Effective charge loss rates close to zero result in the numerical divergences, and are moments when charge loss rates match mass loss rates.

regime, the effective charge $Q^* = Q/M$ increases steadily, always driving the black hole toward extremality *without* the need for an introduced dark sector. Neglecting other discharging mechanisms, this effect is also expected to increase the lifetime of the BH indefinitely. The increase of effective charge lowers the BH temperature, suppressing the mass loss page factor $f(M, Q^*)$. The end result is effectively an extremal black hole with a suppressed mass and charge loss mechanism and a Hawking temperature that approaches zero. We will see later that the tendency for a BH to approach more extremal can be negated in the presence of an extreme electric field when we include the effects of Schwinger pair production. The inclusion of this effect in addition to Hawking radiation is talked about in more detail within Section 4.11.

4.9 Kerr Evaporation Rates

We consider here BHs with intrinsic angular momentum J instead. The greybody factors used in this section were numerically evaluated for each configuration of the Kerr metric, providing more accurate predictions for spin evolution and black-hole lifetimes compared to the charged case. The parameter describing the evolution of a BH with respect to its angular momentum J is the effective spin parameter a^*. Similarly to Eq. (4.14), we can find the effective spin loss rate

$$\frac{da^*}{dt} = \frac{1}{M^2}\frac{dJ}{dt} - 2\frac{J}{M^3}\frac{dM}{dt} \tag{4.17}$$

In terms of Page factors defined in Eqs. (4.4), (4.5), this is

$$\frac{da^*}{dt} = a^* \frac{2f(M, a^*) - g(M, a^*)}{M^3} \tag{4.18}$$

which will dictate the evaporation of spin for a Kerr BH.

Contrary to the case of shedding charge, there is no scenario in which an emitted particle always carries zero angular momentum away from the black hole. In the case of charge loss, the black hole can, in principle, emit a neutral particle, meaning that not every emission necessarily contributes to charge depletion. However, for angular momentum, any emitted particle can carry away some amount of total angular momentum, ensuring that the black hole's angular momentum is always being affected by its radiation. Even if we consider a particle with a quantum spin number of zero (such as a scalar boson), the total angular momentum J of the emitted particle is given by the sum of its intrinsic spin S and its orbital angular momentum L,

$$J = L + S \tag{4.19}$$

For a spinless particle, $S = 0$, but the emitted particle will still have some nonzero orbital angular momentum L. Since angular momentum is quantized, even the lowest possible allowed orbital angular momentum state for a given emission process is nonzero, particularly when considering emissions from a rotating body. This results in a scenario where $g(M, a^*)$ will always be greater than $2f(M, a^*)$ for all greybody factors tabulated. Therefore, spin will always be shed before mass and will cause a BH to tend away from an extremal BH for all particles considered. Still, we can compare timescales of effective spin loss against the charge loss values obtained above. Looking at Fig. 4.7, we can see that the timescale for spin evaporation for a maximally spinning ($a^* = 0.9$), 10^{17} g mass BH is $\approx 1.8 \times 10^{24}$ seconds. Compared to the charge lifetime of $\approx 3 \times 10^{23}$ seconds, we find that spin is shed approximately 6 times slower in our analysis. Again, it is important to emphasize that our lifetime results for charge are an overestimation, however, the hierarchy coincides with what is to be expected.

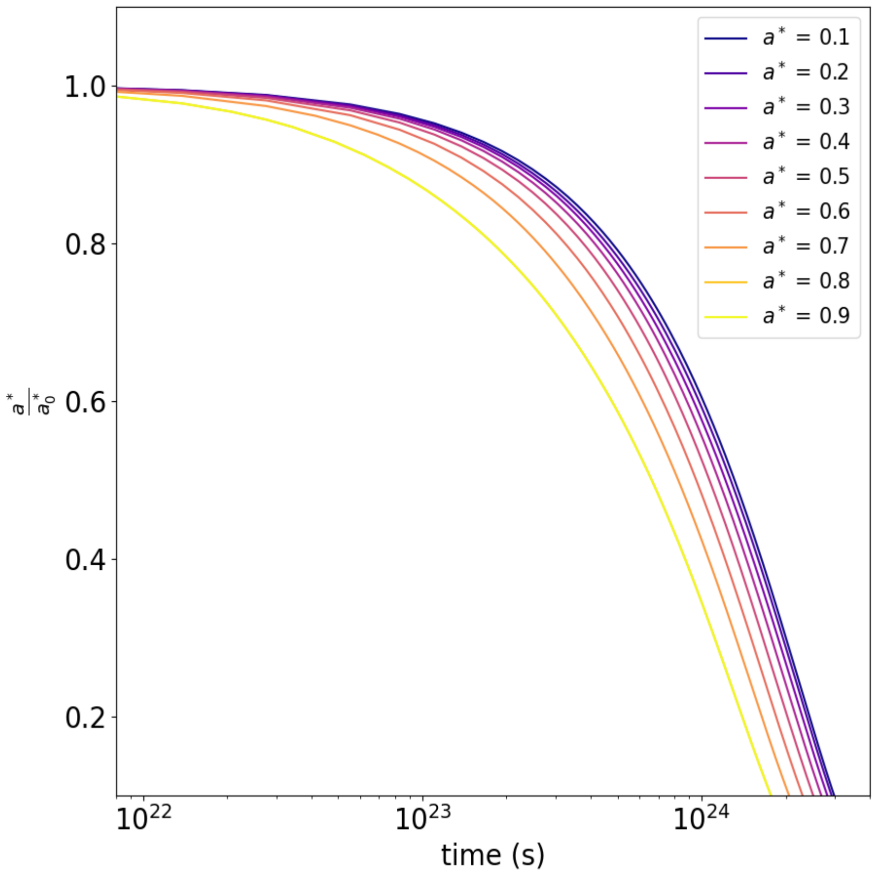

Figure 4.7: Overlay of varying effective spin normalized by initial spin value as a function of time for a Kerr BH. Lifetime is given for a Kerr BH of mass $10^{17}g$, and introduced DM is not considered.

We have also investigated the effects of introducing massless particles of varying degrees of freedom and spin on shedding angular momenta. The results can be seen in Fig. 4.8. The most interesting case is the far left panel, where we have considered spin zero particles. In this case, the overall spin carried away by the particle is entirely orbital angular momentum. While the page factor for spin $g(M, a^*)$ is always larger than

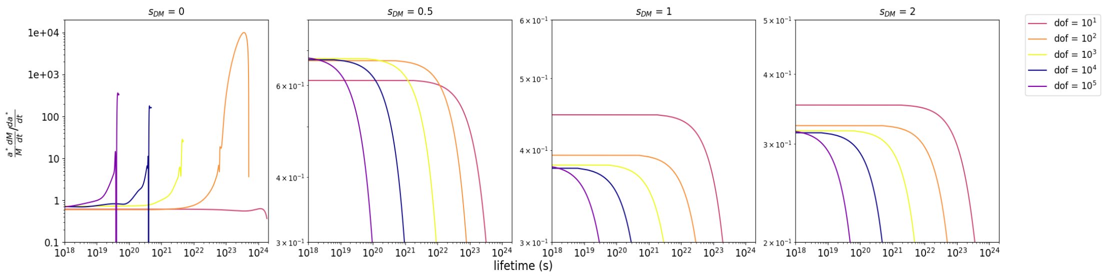

Figure 4.8: Effect of adding various spin dark matter with various masses to the evaporation model of a Kerr BH. BH initial parameters are $a^* = 0.8$, $m_{BH} = 10^{17}g$.

the mass loss term $2f(M, a^*)$, in this case the difference between the two becomes very small. This results in the denominator term $\frac{da^*}{dt}$ in the lognormalized ratio to become extremely small in comparison to the mass loss term $\frac{dM}{dt}$. For all other cases the intrinsic spin S of the particle acts to increase the spin loss term, resulting in a lognormalized ratio less than 1. We observe apparent divergences in the spin-zero case near the end of the black hole's spin-down phase. This behavior arises because a scalar field with a large number of degrees of freedom removes angular momentum increasingly inefficiently as the black hole approaches low spin. As a^* decreases, the effective spin-loss rate $\frac{da^*}{dt}$ falls sharply, since there is progressively less angular momentum per emitted quantum. When this declining spin-loss rate is combined with the continually increasing mass-loss rate $\frac{dM}{dt}$, the resulting ratio produces a sharp spike for the $S = 0$ case. We do not see this effect for the other cases as having some nonzero spin number S causes the effective spin loss rate to be much more stable even as the BH has a small amount of angular momentum left to radiate.

We have also illustrated how effective spin affects the lognormalized ratio

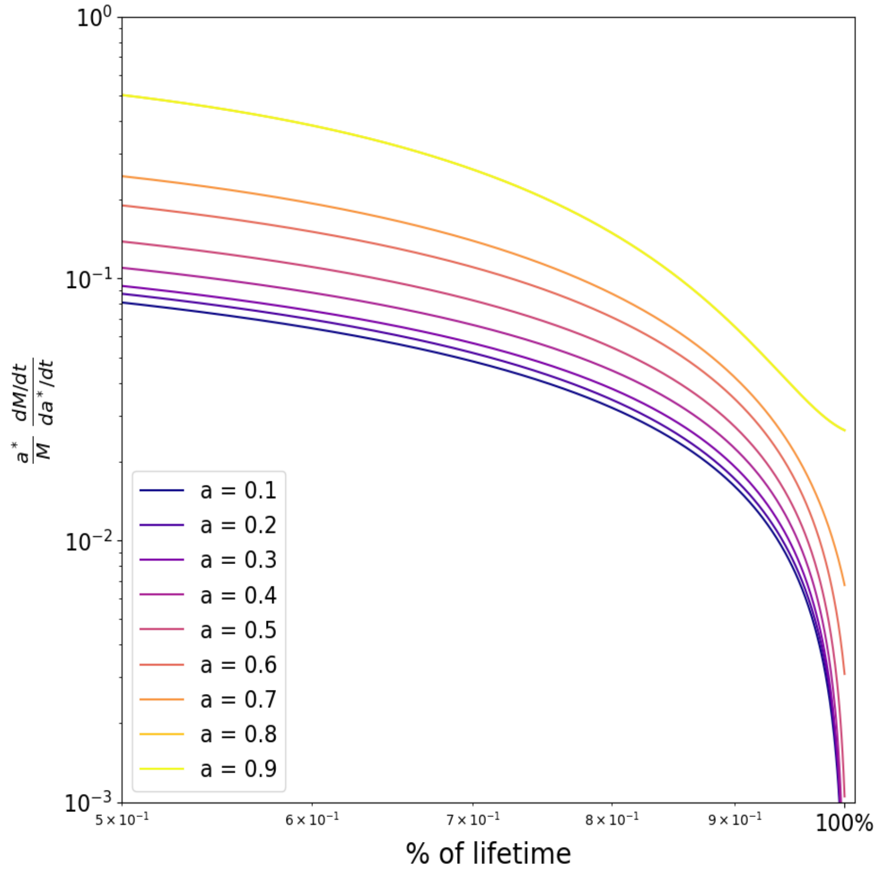

Figure 4.9: Lines of $\frac{\frac{1}{M}dM/dt}{\frac{1}{a^*}da^*/dt}$ lognormalized for a Kerr BH for various effective spin values. In this case, the emitted particle spectrum consists exclusively of Standard Model species.

$\frac{a^*}{M}\frac{dM/dt}{da^*/dt}$ in the absence of a large number of introduced degrees of freedom in Fig. 4.9. Again the lognormalized ratio expressed in the figure shows how the normalized mass loss and spin loss rates compare. Values greater than 1 illustrate regions where mass is shed faster than spin, and values less than 1 represent the inverse. In Fig. 4.9, the plot describes that as effective spin of a BH is increased the spin loss rate is decreased in comparison to the mass loss rate. Again this is a result of the dependence of a Kerr BH's temperature on the effective spin of the BH. As effective spin increases the overall BH temperature decreases, resulting in delayed particle evaporation thresholds and a higher exponential suppression of evaporated particles in comparison to lower effective spins. As we approach the end of a BH's lifetime the spin loss term begins to dominate until spin is completely shed.

4.10 Relative Hierarchies

The purpose of the section above is to illustrate the possibility of upsetting the hierarchy of charge, spin and mass loss as dictated by the standard model. Our findings indicate that under any circumstances, spin is lost before mass. However, we have found that through the introduction of a large number of massive or massless degrees of freedom we can create a scenario where we can increase the effective charge of a RN BH. The standard hierarchy we set out to upset is as follows; charge is neutralized first, followed by spin and finally mass will be depleted ($\frac{dM}{dt} < \frac{dJ}{dt} < \frac{dQ}{dt}$). There are two distinct ways in which we can upset this natural order.

The first possibility of changing the order of parameters lost is through the introduction of a select number of chargeless degrees of freedom. In our consideration of a BH of mass $M = 10^{17}$g, if we include massless scalar degrees of freedom within the range of $100 \lessapprox N_{dof} \lessapprox 220$, we find a scenario where the hierarchy changes and spin is shed first, followed by charge neutralization and then finally mass loss ($\frac{dM}{dt} < \frac{dQ}{dt} < \frac{dJ}{dt}$).

Since the introduction of chargeless degrees of freedom has no effect on the rate of charge loss, this introduction solely affects the rates of mass and spin loss. Degrees of freedom larger than ≈ 220 result in increasing effective charge, and would make mass evaporate faster than charge. However, degrees of freedom greater than ≈ 100 will result in spin being lost faster than charge being lost, with mass still evaporated away last. Again this is a direct result of these introduced degrees of freedom always being able to carry away some nonzero amount of angular momentum. The upper limit of the number of degrees of freedom is set by the BH mass and can be seen in the dependence of $\frac{dQ^*}{dt}$ on M in Eq. (4.13). The lower limit is also dependent on the intrinsic spin S of the introduced particle. Higher particle spins will require lower degrees of freedom due to the dependence of the intrinsic spin S in the total angular momentum J carried away from the BH.

The second possibility is when we consider degrees of freedom higher than ≈ 220. In this scenario, we find that the effective charge increases for a $M = 10^{17}$g BH. While the overall charge of the BH goes down, it is outpaced by mass. In this scenario we find that again spin is lost first, but we have the order of mass and charge loss reversed in comparison to the first scenario ($\frac{dQ}{dt} < \frac{dM}{dt} < \frac{dJ}{dt}$).

Note, again, that in the absence of Schwinger pair production extremely massive BHs act similar to blackbodies. In this case, even without the addition of any dark degrees of freedom, we have a scenario where the BH evaporates a significant portion of mass before having any noticeable amount of charge evaporated. Since any charged particle within the standard model has some nonzero mass, the production of such particles will be exponentially suppressed for the majority of the BH's lifetime. Once the BH reaches a mass of roughly 10^{17}g, it begins to emit a significant amount of charged particles. Prior to that mass threshold we find that the effective charge will increase and approach an extremal charged BH. For the majority of the lifetime of a BH, effective charge increases until charged particles can be produced. While this isn't necessarily upsetting the natural hierarchy discussed above since charge is eventually dissipated before mass near the end of the BH's lifetime, the majority of the lifetime is naturally spent increasing the effective charge for heavier BHs.

4.11 A Charged Caveat: Schwinger Pair Production

The above analysis for a RN BH only considers the production of charged particles through the computation of their greybody factors outlined in [17]. We did not consider, however, purely classical processes that could neutralize the BH. For example, if the BH were surrounded by charged interstellar media, we would see a fairly quick neutralization of the charge due to selective accretion of particles with charge opposite the BH. Still, even in the absence of such media we can still have an enhanced charge neutralization through quantum effects. A possible significant contribution comes from the Schwinger effect [197, 154, 57]. The Schwinger effect describes spontaneous pair creation of charged particles in the presence of an extremely strong electric field. In the majority of electric fields the production is exponentially suppressed, yet for electric fields higher than the threshold field this effect can become dominant. Such a strong field results in vacuum decay and the subsequent creation of electron-positron pairs, serving to lower the charge of the field by $2e$ for each pair created [197] [124]. The greybody factors for electron and positron production can then be significantly enhanced via Schwinger pair production.

We can estimate the effect that this has on the charge lifetime of a RN BH by modifying Eq. (4.13). Starting with the Schwinger pair production formula [124], the electron-positron creation rate per unit volume is

$$\Gamma = \frac{e^2}{4\pi^3}\frac{Q^2}{r_+^4}\exp\left(-\frac{\pi m_e^2 r_+^2}{eQ}\right)\left[1+\mathcal{O}\left(\frac{e^3 Q}{m_e^2 r_+^2}\right)+\cdots\right] \tag{4.20}$$

where m_e is the mass of the electron and Q is the BH charge defined as $Q = Q^* M$. We will require that the term $\frac{e^3 Q}{m^2 r^2} << 1$ and consider only the leading order term. If this assumption is violated the series is no longer convergent and can be assumed to approach infinity, resulting in a massive discharge of the RN BH until the effective charge is low enough to allow the condition to be satisfied. The charge loss rate per unit time can be obtained via integration

$$\frac{dQ_S}{dt} = -2e\int \Gamma dV, \tag{4.21}$$

where Q_S has been introduced to denote the difference between charge loss specifically due to Schwinger pair production. Integrating Eq. 4.20 rate per unit volume we find the charge loss rate to be

$$\frac{dQ_S}{dt} = \frac{-e^3}{\pi^2}\frac{1}{r_+}\exp\left\{\left[-\frac{r_+^2}{QQ_0}\right\}\right] - \frac{\pi}{(QQ_0)^{1/2}}\mathrm{erfc}\left[\frac{r_+}{(QQ_0)^{1/2}}\right]. \tag{4.22}$$

where $\mathrm{erfc}(x)$ is the complementary error function and $Q_0 = \frac{e}{\pi m_e^2}$. Combining the effects

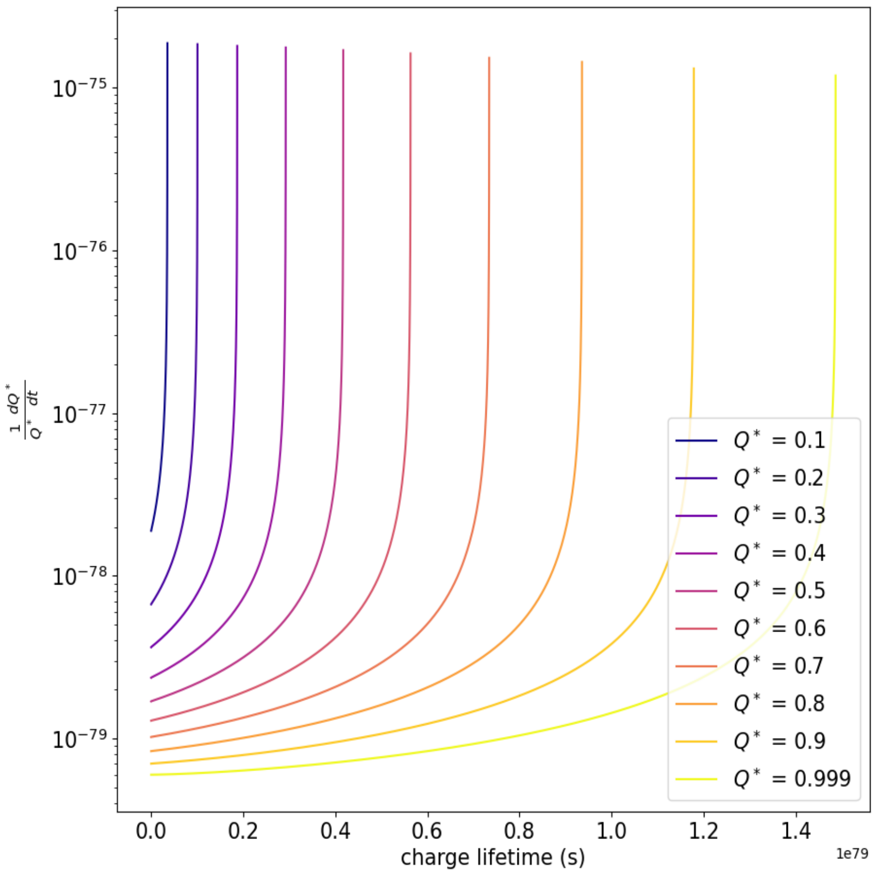

Figure 4.10: Effect of Schwinger pair production on $\frac{dQ^*}{dt}$ for various effective charges for a BH of initial mass $M = 10^4 M_\odot$. All charge is lost before any significant amount of mass is lost.

of Schwinger pair production, the effective charge loss in Eq. (4.13) becomes

$$\frac{dQ*}{dt} = \frac{1}{M}\frac{dQ}{dt} + \frac{1}{M}\frac{dQ_S}{dt} - \frac{Q^*}{M}\frac{dM}{dt} \tag{4.23}$$

where $\frac{dQ}{dt}$ describes the charge particle production at the horizon due to quantum vacuum fluctuations and $\frac{dQ_S}{dt}$ describes electron-positron production due to the presence of a large electric field at various $r > r_+$. The rapid discharge due to Schwinger

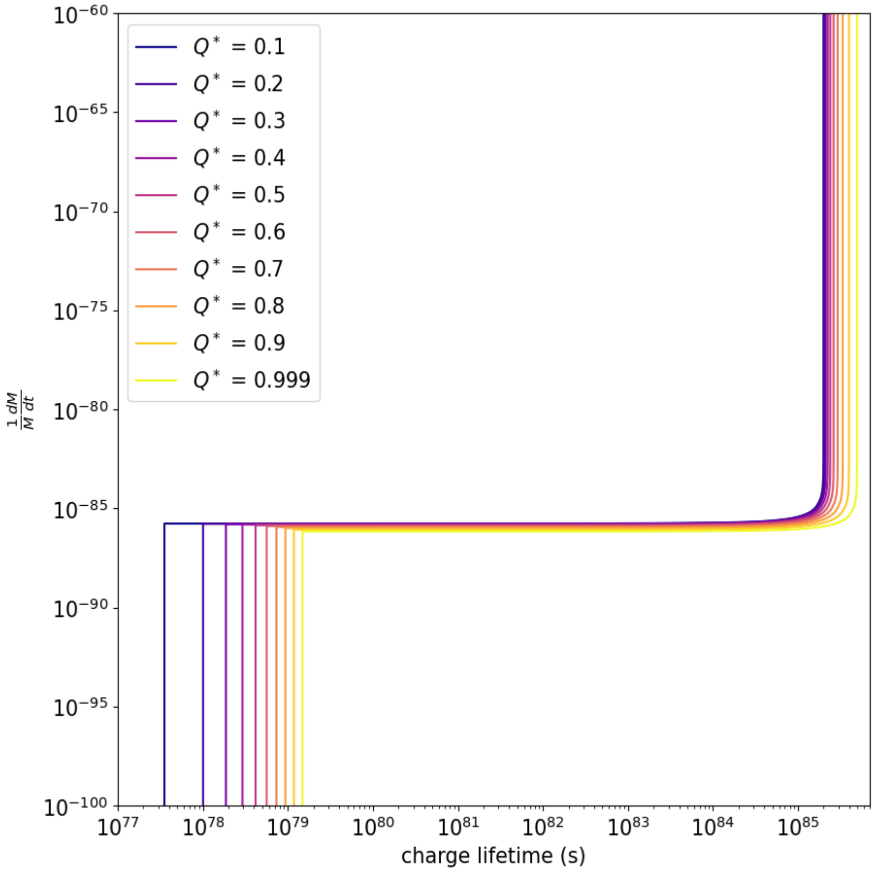

Figure 4.11: Effect of Schwinger pair production of $\frac{dM}{dt}$ for various effective charges. BH has initial mass $M = 10^4 M_\odot$. Sudden jump on the left hand side of the plot is when charge is neutralized. From then on BH behaves as if static.

pair production would also affect the mass loss rate $\frac{dM}{dt}$ as well. The mass loss rate can be obtained by looking at

$$\frac{dM}{dt} = \frac{\partial M}{\partial Q}\frac{\partial Q}{\partial t} = \frac{\partial M}{\partial r_+}\frac{\partial r_+}{\partial Q}\frac{\partial Q}{\partial t} \tag{4.24}$$

Working out the algebra we have the expression

$$\frac{dM_S}{dt} = \frac{Q^* M}{r_+}\frac{dQ_S}{dt} \tag{4.25}$$

which is analogous to the result from [124] in natural units. This additional term will serve to speed up the mass loss rate for a charged BH as long as $\frac{dQ_S}{dt}$ is nonzero.

We can calculate an order of magnitude estimate to see if this effect will have a significant impact on the discharging of a RN BH. Both the Schwinger and the Hawking discharge rate feature an exponential suppression; in the case of the Schwinger discharge rate this is

$$\Gamma_{\text{Sch}} \sim \exp\left(-\frac{\pi m_e^2 r_+^2}{eQ}\right), \tag{4.26}$$

with

$$r_+ = M\left(1 + \sqrt{1 + (Q/M)^2}\right)$$

and e the electron charge. Note that this is the discharge rate per unit volume and the Hawking rate is per unit energy. Expressing the above in natural units, neglecting order 1 factors, we get

$$\Gamma_{\text{Sch}} \sim \exp\left(-\frac{m_e^2 M^2}{eQ}\right). \tag{4.27}$$

The Hawking evaporation rate is Boltzmann suppressed by m_e/T_H, where

$$T_H \sim \frac{1}{M}\left(1 - \frac{Q^2}{M^2}\right),$$

thus

$$\Gamma_{\text{Hawk}} \sim \exp\left(-\frac{M m_e}{1 - \frac{Q^2}{M^2}}\right). \tag{4.28}$$

The Schwinger discharge rate dominates when

$$\Gamma_{\text{Sch}} \gtrsim \Gamma_{\text{Hawk}} \Rightarrow \frac{m_e^2 M^2}{eQ} \lesssim \frac{m_e M}{\left(1 - \frac{Q^2}{M^2}\right)}. \tag{4.29}$$

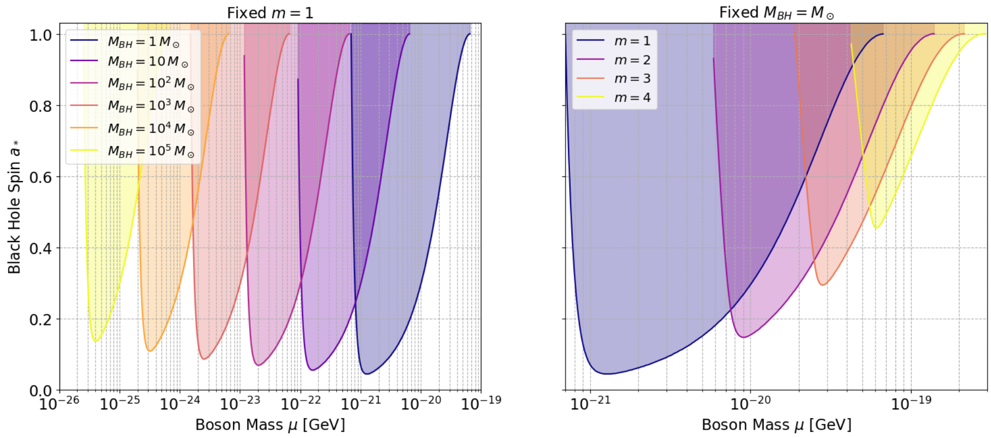

Figure 4.12: Parameter space for possible superradiance regions for a given black hole mass and light boson.

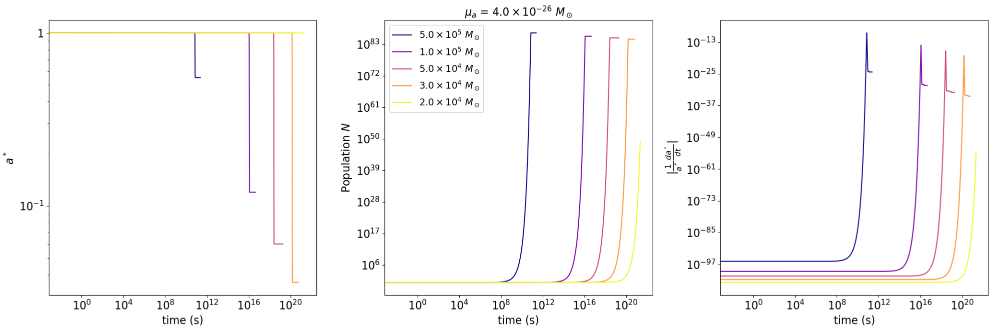

Figure 4.13: Effective spin loss due to superradiance of a Kerr BH. Extracting the majority of spin ($\mathcal{O}(0.1)$ remaining) occurs in approximately 10^{16}s in comparison to the Hawking lifetime for Kerr BHs of 1.8×10^{24} s.

We can rearrange this equation to find that in terms of mass-to-charge ratio $x_e = (e/m_e)$ and $x_{\rm BH} = Q/M$, Schwinger discharge dominates if

$$\frac{1 - x_{\rm BH}}{x_{\rm BH}} \lesssim x_e \simeq 6.4 \times 10^{21}. \tag{4.30}$$

In turn, this implies

$$x_{\rm BH} \gtrsim 10^{-22}. \tag{4.31}$$

Since the Planck charge is $11.7e$, we find that Schwinger discharge dominates for any black hole with any electric charge $Q \geq e$ when $M \leq 10^{18}$ grams.

Looking at Fig 4.10, we can see how Schwinger pair production enhances the number of charged particles in the Hawking spectra. In this case we have considered a much larger BH ($10^4 M_\odot$). The primary purpose of this choice is to consider a stage in a charged BH's life when there is essentially only blackbody radiation. This mechanism only serves to deplete mass from a charged BH and, by Eq. (4.13), would serve to increase the overall effective charge of such a large BH. Without Schwinger pair production, there would be a tendency for charged BHs to approach extremality in all scenarios. However, we can see that this effect serves to rapidly discharge the BH in comparison to the mass loss timescales. Looking at Fig. 4.11, we can see that the lifetime of a BH with various effective charges hovers around 10^{85}s for BH of this size. This is roughly 10^{16} times larger than the timescale for which a BH will discharge due to Schwinger pair production. It's important to note that this analysis only includes effects from the leading order term for Schwinger production and therefore indicates the largest lifetimes possible when including this effect. All higher order terms serve to increase the electron-positron production rate, and therefore will decrease the lifetime of a charged BH.

4.12 Superradiance and Spin: Angular Momentum Extraction

Recent work has shown that black hole superradiance can extract angular momentum on timescales much shorter than Hawking emission [207]. Although a black hole cannot be completely spun down through this mechanism, a substantial fraction of its angular momentum can be removed, and we include this effect alongside the Hawking analysis for Kerr black holes.

Superradiance arises from the existence of the ergoregion between the outer and inner horizons, where negative-energy states are permitted [42]. A particle entering the ergoregion may decay into two modes, one of which carries negative energy $E_1 < 0$. Energy conservation then requires the other mode to emerge with energy $E_2 = E_0 - E_1 > E_0$, so that the outgoing particle carries more energy than the particle originally entering the black hole. The negative-energy mode remains inside the ergoregion and effectively reduces the black hole's rotational energy.

Superradiance is the amplification of this process: bosonic waves satisfying the superradiant condition grow exponentially by repeatedly extracting rotational energy. The resulting accumulation of a scalar cloud outside the horizon forms a "gravitational atom," enabling rapid and efficient angular momentum extraction compared to Hawking radiation alone. In the following section we will outline the timescales associated with removing angular momentum from the system via the introduction of massive scalar bosons and the superradiant effect.

Superradiance of a BH is allowed so long as the superradiant condition [207]:

$$\omega < m\Omega_+ \tag{4.32}$$

is satisfied. In the above equation ω refers to the frequency of the incoming scalar wave, m the azimuthal quantum number of a given state and Ω_+ the angular velocity of the BH horizon. In general, we assume the population of states occurs with ω to be roughly the boson mass $\omega \approx \mu_a$. The overall angular momentum available to be extracted via superradiance is then proportional to the number of particles produced by this process. For each particle extracted from a given BH, the overall angular momentum of the BH is lowered by $\Delta J = m$, or an integer quanta of angular momentum proportional to the given state occupied. Therefore, the overall rate of angular momentum extraction from the BH is equivalent to the rate at which superradiance occurs Γ_{SR} for a given particle multiplied by the number of particles of a given state N_m:

$$\frac{dJ}{dt} = -m\Gamma_{SR}N_m. \tag{4.33}$$

Again using Eq. (4.17), we estimate the change in effective spin as long as we have the mass loss rate due to superradiance. The mass loss rate of the BH is obtained in a similar manner, except we can replace the quanta of spin m by the scalar particle mass μ_a to get

$$\frac{dM}{dt} = -\mu_a\Gamma_{SR}N_m. \tag{4.34}$$

All together, the expression for effective spin loss due to superradiance is given by;

$$\frac{da^*_{SR}}{dt} = -\frac{1}{M^2}m\Gamma_{SR}N_m + 2\frac{a^*}{M}\mu_a\Gamma_{SR}N_m. \tag{4.35}$$

This equation is expected to be valid so long as the superradiance condition in Eq. (4.32) is satisfied and the superradiance rate Γ_{SR} is faster than the Eddington accretion time. We have simulated the effective spin available to be extracted from the BH in the Regge trajectory plots in Fig. 4.12. For any BH mass we can see that, for quantum numbers $\{n, l, m\}$, the $m = 1$ state allows for the largest Δa^* in the right hand panel. Additionally, we simulate the parameter space for various BH and scalar boson masses that allow for superradiance. We can see as a general trend that as BH mass increases the overall spin available for extraction decreases. Additionally, we expect that constraints on light boson candidates to be another limiting factor in the possibility of this SR regime being reached for heavy BHs.

Still, we can estimate the spin down lifetime of a BH due to superradiance with Eq. 4.32. Utilizing the results of [207], the $\{n, l, m\} = \{2, 1, 1\}$ state is the most efficient at extracting angular momentum. In general the lifetime of a BH's superradiant event of a BH can be estimated as $\tau_{SR} = \frac{1}{\Gamma_{SR}}$ [21]. Using our superradiant rates derived from [207], we illustrate the expected lifetime of a given superradiant event for a solar mass BH and variable gravitational fine structure constant $\alpha = \mu_a M$ in Fig. 4.14. The shortest spin-down time we find is approximately 10^{16} s, corresponding to a reduction from $a^* = 0.9999$ to $a^* \approx 0.15$ for a $10^5 M_\odot$ black hole. This is approximately 10^8 times faster than it takes for a Kerr BH to lose spin due to purely radiative processes. Granted, this is still a rotating BH, but this process serves to greatly reduce the spin and overall lifetime due to the significant influence of effective spin on BH temperature and particle emission rate. For the $m = 1$ superradiant states, smaller values of the

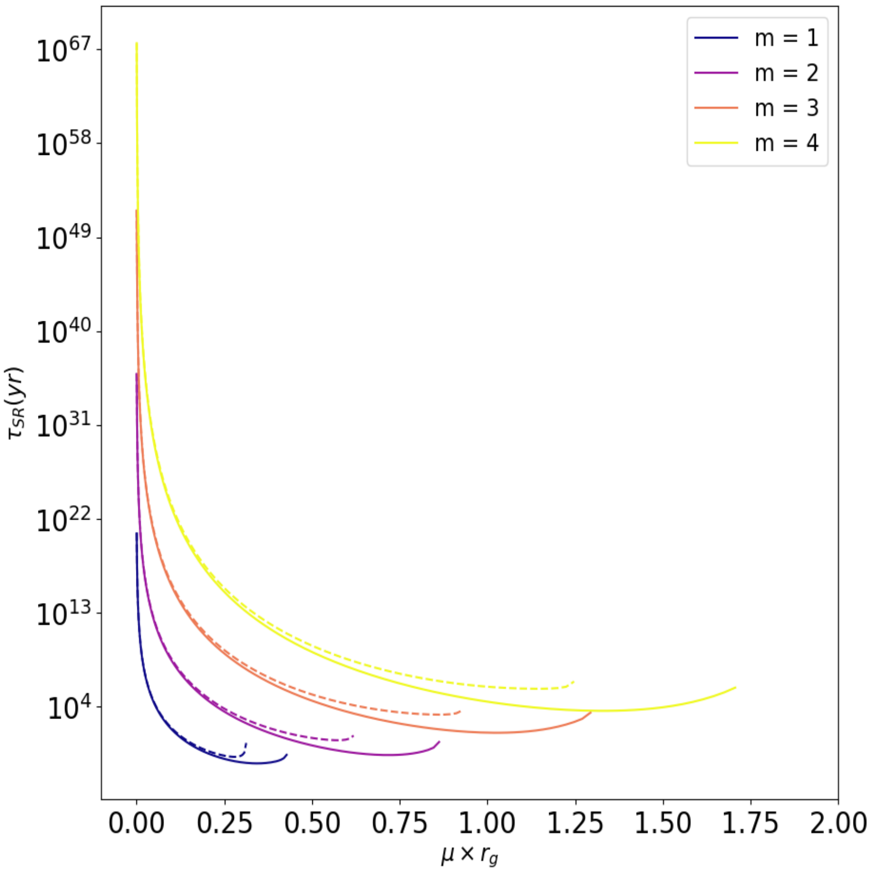

Figure 4.14: Superradiant event lifetime for a given BH and boson mass, set by the gravitational fine structure constant $\alpha = \mu_a \times r_g$. We use $M = 10^5 M_\odot$. Solid lines correspond to $a^* = 0.99$ and dotted to $a^* = 0.9$. Lifetimes are expected to increase proportionally as mass increases, as there will be more angular momentum available for a given effective spin.

gravitational fine-structure constant α correspond to more spin being extractable. This explains why the superradiant lifetime grows as $\alpha = \mu_a r_g$ approaches zero in the figure. Overall, a significant amount of the available spin can be lost from a superradiant event.

When considering the lowest superradiant state, all BHs shed spin under approximately 10^{20} seconds for the BH mass considered, which is much faster than the expected Kerr evaporation timescales.

4.13 Discussion

This comprehensive study systematically compares the evaporation dynamics of Schwarzschild, Kerr, and Reissner-Nordström black holes through the lens of Page factors, revealing critical insights into the interplay of mass, charge, and angular momentum loss. For **Schwarzschild black holes**, isotropic Hawking radiation drives mass loss on timescales governed solely by the Stefan-Boltzmann law, with evaporation rates scaling as M^{-2}. In contrast, **Kerr black holes** exhibit rapid angular momentum dissipation due to preferential emission of high-spin particles and superradiance effects, leading to a spin-down phase orders of magnitude faster than mass loss. This asymmetry arises from frame-dragging effects, which amplify emission from equatorial regions and enhance the coupling of angular momentum to radiation fields.

For **RN black holes**, charge loss dominates early-stage evaporation, mediated by the efficient emission of charged particles through both Hawking radiation and Schwinger pair production. The Coulomb potential barrier introduces a self-regulating mechanism: high charge-to-mass ratios (Q_*/M) suppress thermal emission by lowering the black hole temperature, while simultaneously accelerating charge loss via enhanced particle tunneling. Near-extremal configurations ($Q_* \to 1$) exhibit prolonged lifetimes due to suppressed Hawking temperatures, though quantum corrections destabilize these states over cosmological timescales. The figures with very long lifetimes are intended to illustrate the existence and scaling of hierarchy modification rather than to imply direct cosmological durations.

A pivotal finding is the **modification of evaporation hierarchies** by non-Standard Model particle species. Introducing large numbers of uncharged dark matter degrees of freedom shifts the balance between mass and charge loss rates. For example, a dark sector with $\mathcal{O}(100)$ massless particles can invert the canonical hierarchy, causing $\dot{M}/M$ to exceed $\dot{Q}_*/Q_*$ even for highly charged black holes. This challenges assumptions about "charge neutrality" as a universal endpoint of evaporation and suggests that astrophysical black holes in dark matter-rich environments may retain residual charges longer than predicted by SM-only models. Although the values of N_{dof} considered here may appear large, they are not unconstrained and are subject to existing astrophysical, collider, and gravitational bounds. Black hole evaporation thus emerges as a complementary and largely model-independent probe of such dark sectors.

The role of **greybody factors** further complicates these dynamics. For Kerr black holes, frame-dragging distorts emission spectra, favoring low-angular-momentum modes and accelerating spin loss. In RN systems, greybody corrections suppress charged particle emission at low energies but enhance high-energy discharges, creating a feedback loop that regulates charge loss. These effects are particularly pronounced in near-extremal regimes, where geometric optics approximations break down, necessitating full wave-equation treatments for accuracy.

Schwinger pair production emerges as a critical mechanism for light ($M \lesssim 10^{17}, \mathrm{g}$), highly charged black holes. At $Q^*/M \gtrsim 0.5$, this quantum electrodynamic process dominates over Hawking radiation, ejecting $e^{\pm}$ pairs and accelerating charge loss. For primordial black holes formed with near-extremal charges, the competition between Schwinger-driven discharge and mass loss can nevertheless lead to a transient increase in the effective charge $Q^* = Q/M$, even as the absolute charge decreases. Any possible connection to early-universe baryogenesis would therefore be indirect and restricted to sufficiently light black holes, while the extremely long lifetimes shown in some figures should be interpreted as illustrating scaling behavior in the near-extremal regime rather than cosmologically relevant durations.

Superradiance plays a role for any mass BH to extract significant amounts of angular momentum of a Kerr BH in the form of a light boson cloud. For a given BH mass, a range of axion masses allow for $\Delta a^* = \mathcal{O}(1)$ changes in the overall angular momentum on timescales much faster than that of Hawking radiation. The possibility of forming a gravitational atom allows for much more rapid increases in temperature of the BH and speeds up the overall mass loss of the system as well.

These results underscore the inadequacy of single-parameter evaporation models and highlight the need for **multivariate frameworks** incorporating particle physics beyond the SM. Future work should explore the interplay between evaporation and accretion in astrophysical environments, quantum gravity corrections to semi-classical approximations, and observational signatures of charged or spinning black holes in gravitational wave or cosmic ray data. Additionally, the discovery that dark sectors can dramatically alter evaporation hierarchies invites experimental tests via high-precision measurements of black hole merger remnants or gamma-ray burst spectra. By bridging thermodynamics, quantum field theory, and astrophysics, this study advances our understanding of black holes as multichannel dissipative systems and provides tools to probe fundamental physics through their evaporation.

Data Availability Statement

The data and code supporting the findings of this study are publicly available in Ref. [99].

Chapter 5

High-frequency gravitational wave transients from superradiance

5.1 Introduction

Chapter 3 established that asteroid-mass primordial black holes (PBHs) in the mass window 10^{-16}–$10^{-11}\,M_\odot$ are among the most promising sources of high-frequency gravitational waves (HFGWs) detectable at resonant microwave cavities, yet showed that even maximally optimistic merger event rates fall short of current ADMX sensitivity. Chapter 4 demonstrated that the spin dynamics of Kerr black holes are qualitatively modified by a large dark sector, motivating a closer look at angular-momentum–driven emission mechanisms. This chapter brings these threads together by treating superradiance not merely as a spin-loss mechanism but as a gravitational-wave source in its own right, and by performing the first unified detectability analysis of superradiant GW signals from PBH gravitational atoms at GHz-band detectors.

The basic mechanism is as follows. A bosonic field of mass μ satisfying the superradiant condition

$$0 < \omega < m\,\Omega_H, \tag{5.1}$$

where Ω_H is the angular velocity of the Kerr horizon and m the azimuthal quantum number, is exponentially amplified at the expense of the black hole's rotational energy [219, 220, 185]. When the boson's Compton wavelength is comparable to the gravitational radius, the amplified modes are gravitationally trapped and grow into a macroscopic quasi-bound cloud—a *gravitational atom* (GA) [71, 43]—with occupation numbers reaching $N \sim 10^{70}$–10^{80} [21]. The dynamics are controlled by the gravitational fine-structure constant

$$\alpha \equiv GM_{\rm BH}\,\mu, \tag{5.2}$$

which simultaneously sets the cloud radius $r_c \sim M_{\rm BH}/\alpha^2$, the superradiant instability rate, and the hydrogenic spectrum of bound levels [22, 21]. Efficient growth requires $\alpha \sim 0.1$–0.5, which for a given boson mass μ selects a preferred BH mass—or equivalently, for a given BH mass, a preferred range of boson masses. Notably, bosons in the mass range $\mu \sim 10^{-7}$–10^{-5} eV selected by sub-solar PBH masses are independently motivated by the axion solution to the strong CP problem [179, 212, 213] and the cosmology thereof [184, 4, 77, 72].

The GW frequency is set almost entirely by μ: annihilation of two cloud bosons into a graviton emits at $\omega_{\rm ann} \simeq 2\mu$, while level transitions between states with principal quantum numbers n_g and n_e emit at the Bohr frequency $\omega_{\rm tr} \propto \mu\alpha^2(n_g^{-2} - n_e^{-2})$. For sub-solar PBH masses this places signals in the MHz–GHz band targeted by resonant-cavity haloscopes such as ADMX. Once superradiance saturates, two emission channels become active in the isolated system:

1. *Level transitions*, in which the quadrupole self-interaction of two populated states drives bosons from an excited level to a lower one, emitting quasi-monochromatic GWs at the Bohr frequency $\omega_{\rm tr}$.

2. *Annihilations*, in which pairs of bosons convert into gravitons at $\omega_{\rm ann} \simeq 2\mu$, producing an extremely long-lived, slowly decaying quasi-monochromatic signal.

Bosenova collapses, in which the cloud implodes once its occupation number exceeds a critical value set by bosonic self-interactions, introduce additional burst-like GW emission and periodically reset the superradiant growth [215, 216]; we account for the resulting upper bound on cloud populations when deriving strain amplitudes.

In realistic astrophysical environments GAs are often not isolated. A binary companion introduces a time-varying tidal potential that drives resonant transitions between GA levels whenever the orbital frequency sweeps through a Bohr frequency of the gravitational atom. The cloud then behaves as a driven two-level system whose response is captured by a Landau–Zener analysis [34]: depending on the adiabaticity parameter $z \equiv \eta^2/(|\Delta m|\dot{\Omega})$, transitions range from partial population transfer ($z \ll 1$) to complete level inversion ($z \gg 1$). The resulting GW signal is a short-duration transient burst whose characteristic strain and spectral shape we compute following the formalism of Kyriazis & Yang [147], applied here for the first time to the PBH binary case.

This chapter brings together four ingredients not previously combined in the context of high-frequency GWs: (i) relativistic superradiant instability rates [128, 127], (ii) numerically computed cloud masses [198, 163], (iii) the binary-induced resonance formalism of Ref. [147], and (iv) realistic ADMX sensitivity curves [35]. Concretely, we make the following contributions:

- We map the superradiant parameter space in the PBH mass range using Regge trajectories (Sec. 5.2), identifying the $(\mu, M_{\rm BH}, \alpha)$ combinations that produce MHz–GHz emission and assessing the role of bosenova saturation and accretion in bounding the allowed cloud populations.

- For isolated GAs we derive time-domain GW strain envelopes and analytic closed-form frequency-domain templates, via the exponential integral E_1, for both the level-transition and annihilation signals (Sec. 5.3 and App. 7.4). Peak strains reach $h \sim 10^{-22}$ at 1 kpc for the most optimistic annihilation configurations, while level-transition signals typically peak at $h \sim 10^{-23}$.

- For PBH binaries we apply the Landau–Zener formalism to the hyperfine transition $\{2, 1, 1\} \to \{2, 1, -1\}$, compute the transient GW burst characteristic strain across the $M_{\rm BH}$–α parameter space (Sec. 5.4), and derive the conditions under which the signal duration exceeds the ADMX ring-up time.

- We provide a detectability analysis for ADMX (Sec. 5.5), finding that binary-induced burst strains fall orders of magnitude below the ADMX sensitivity threshold at astrophysically plausible distances. Using PBH merger rates from the early two-body formation channel [191], detectable events would require source distances $\lesssim 1\,\mathrm{AU}$, whereas the merger rate implies characteristic distances exceeding $\sim 9\,\mathrm{kpc}$.

Taken together, our results show that PBH gravitational atom systems are among the very few theoretically well-motivated sources of MHz–GHz gravitational waves, but that a significant improvement in detector strain sensitivity beyond current ADMX runs is required to make them observable. We identify improved sensitivity, faster ring-up times, and lower frequency coverage as the priorities for future high-frequency GW detector design, and provide analytic frequency-domain templates as concrete waveform targets for such searches. We use natural units $\hbar = c = 1$ throughout.

5.2 Superradiance and the Parameter Space for High-Frequency GWs

Rotating black holes interacting with ultralight bosonic fields define a structured parameter space in which GW emission frequencies are tied directly to the masses of both the black hole and the boson. This section establishes the parametric foundation for the rest of the paper. We derive the conditions for superradiant bound-state formation, introduce the Regge trajectories that delimit the allowed region, give the instability rate governing cloud growth, and identify the portion of parameter space that produces MHz–GHz gravitational waves.

The central result—summarized graphically in Fig. 5.1—is that the MHz–GHz window is populated by black holes in the primordial mass range $M_{\rm BH} \sim 10^{-6}$–$10^{-3}\,M_\odot$ for boson masses $\mu \sim 10^{-8} - 10^{-5}\,\text{eV}$ and gravitational coupling $\alpha \sim 0.1 - 0.3$, as appropriate for efficient superradiance.

5.2.1 Bound States, Regge Trajectories, and the GW Frequency Map

A massive scalar boson of mass μ on a Kerr background admits quasi-bound states whenever its Compton wavelength ${}_C = \mu^{-1}$ (in natural units $\hbar = c = 1$) is comparable to the gravitational radius $r_g \equiv GM_{\rm BH}$. The ratio of these two scales is the gravitational fine-structure constant

$$\alpha \;\equiv\; \frac{r_g}{{}_C} \;=\; GM_{\rm BH}\,\mu\,. \tag{5.3}$$

In the limit $\alpha \ll 1$ the quasi-bound spectrum is precisely hydrogenic [71, 78]:

$$\omega_n \;\simeq\; \mu\left(1 - \frac{\alpha^2}{2n^2}\right), \tag{5.4}$$

where n is the principal quantum number. The spatial extent of the cloud in state $\{n, \ell, m\}$ is

$$r_c \;=\; \frac{n^2}{\alpha^2}\,r_g \;=\; \frac{n^2}{\mu\alpha}\,, \tag{5.5}$$

so the cloud becomes more diffuse at smaller α. The regime $\alpha \sim 0.1$–0.5 simultaneously ensures that the cloud is compact enough for efficient angular-momentum extraction and energetic enough for detectable GW emission; both smaller and larger α values suppress the instability rate exponentially, as discussed in Sec. 5.2.2.

Superradiance is active when the bound-state frequency satisfies the *superradiant condition*

$$\omega < m\,\Omega_H\,, \tag{5.6}$$

where m is the azimuthal quantum number and $\Omega_H = a_*/\left[2GM_{\rm BH}(1+\sqrt{1-a_*^2}\,)\right]$ is the angular velocity of the Kerr horizon [43], with a_* the dimensionless BH spin. Substituting the leading-order binding energy $\omega \approx \mu$ into Eq. (5.6) and solving for the critical spin, one finds the threshold

$$a_*^{\rm crit}(\alpha, m) \;=\; \frac{4m\alpha}{m^2+4\alpha^2}\,, \tag{5.7}$$

above which superradiance operates ($a_* > a_*^{\rm crit}$) and below which it does not. This threshold defines the *Regge trajectory* of the mode: the curve $a_* = a_*^{\rm crit}(\alpha, m)$ in the (a_*, α) plane, or equivalently in the $(a_*, M_{\rm BH})$ or (a_*, μ) planes for fixed μ or $M_{\rm BH}$. A BH that forms with spin above a Regge trajectory will be driven toward it by superradiant spin-down; at the trajectory, growth saturates.

The spin requirement is an important model dependence of this construction. In the simplest radiation-dominated formation channel from nearly Gaussian curvature perturbations, PBHs are generally expected to be born with small dimensionless spin, because large overdensity peaks are close to spherical and pressure support suppresses the growth of angular momentum [58, 67]. However, larger spins can arise in less minimal formation histories. Examples include PBH formation during an early matter-dominated era, where tidal torques and the longer collapse time can generate substantial angular momentum [118]; strongly nonspherical or non-Gaussian collapse; delayed collapse associated with first-order phase transitions; and post-formation PBH–PBH mergers, whose remnants naturally acquire spins of order $a_* \simeq 0.7$ for comparable-mass, initially nonspinning binaries [206]. Accretion can also spin up black holes in sufficiently dense environments, although it is likely inefficient for the isolated low-mass PBHs emphasized here. We therefore do not assume that all PBHs are rapidly rotating. Rather, the Regge trajectories in Fig. 5.1 should be read as identifying the subset of PBHs whose spins exceed the relevant threshold and hence can host the clouds studied in the rest of the paper.

In the limit $a_* \to 1$ (maximally spinning BH), Eq. (5.7) gives $a_*^{\mathrm{crit}} \to 1$ when $\alpha \to m/2$, so the hard upper bound for superradiance to be possible at all is

$$\alpha < \frac{m}{2}. \tag{5.8}$$

This bound is visible in all four panels of Fig. 5.1 as the right-hand edge of each shaded superradiant region.

Figure 5.1 presents the Regge trajectories across four complementary slices of parameter space. In each panel the shaded area marks the region where a given mode can grow:

- *Top-left*: the (μ, a_*) plane for varying M_{BH} at fixed $m = 1$. The allowed region shifts to larger boson masses as M_{BH} decreases, illustrating that lighter PBHs couple to heavier bosons at the same α.

- *Top-right and bottom-left*: the (μ, a_*) and (M_{BH}, a_*) planes for varying m. Higher azimuthal number extends the region to larger α (Eq. (5.8)); at fixed α these modes have a *lower* critical spin but strongly suppressed growth rates, so $m = l = 1$ dominates in practice

- *Bottom-right*: the (α, a_*) plane at fixed $M_{\rm BH}$, showing the Regge trajectories in their most natural parametrization. The state with $m = \ell = 1$ has the *largest* critical spin of the four modes shown at given α, but its superradiant growth rate is the least suppressed in α (Sec. 5.2.2); wherever the spin threshold is cleared, $m = \ell = 1$ therefore grows fastest and extracts angular momentum first.

Throughout, quantum numbers are taken to satisfy $m = \ell = n - 1$ (the lowest radial mode of each angular sector), since these states have the largest angular momentum content per unit energy and therefore achieve the highest saturation occupation numbers [43, 22].

The gravitational-wave emission frequencies are set by μ through two distinct channels. For *annihilation* of two cloud bosons into a single graviton:

$$\omega_{\rm ann} \simeq \frac{\mu}{\pi} \approx 484\,{\rm MHz}\left(\frac{\mu}{10^{-6}\,{\rm eV}}\right). \tag{5.9}$$

For a *level transition* from principal quantum number n_e down to n_g:

$$\omega_{\rm tr} \simeq \frac{\mu\alpha^2}{4\pi}\left(\frac{1}{n_g^2} - \frac{1}{n_e^2}\right) \approx 4.84\,{\rm MHz}\left(\frac{\mu}{10^{-6}\,{\rm eV}}\right)\left(\frac{\alpha}{0.2}\right)^2\left(\frac{1}{n_g^2} - \frac{1}{n_e^2}\right). \tag{5.10}$$

The MHz–GHz frequency window therefore corresponds to boson masses $\mu \sim 10^{-8}$–10^{-5} eV via annihilation, or to slightly heavier bosons when transitions are the dominant channel. The black hole mass required for superradiance efficiency to peak at $\alpha \simeq 0.2$

is then

$$M_{\rm BH} \simeq \frac{\alpha}{\mu G} \approx 2.7\times 10^{-5}\, M_\odot \left(\frac{\alpha}{0.2}\right)\left(\frac{10^{-6}\,\mathrm{eV}}{\mu}\right), \tag{5.11}$$

Across the range $\mu \sim 10^{-8} - 10^{-5}\,\mathrm{eV}$ relevant for MHz–GHz emission, this corresponds to black hole masses $M_{\rm BH} \sim 10^{-6} - 10^{-3}\,M_\odot$ for $\alpha \sim 0.1 - 0.3$. This is the PBH mass range, establishing the central connection exploited throughout this paper: *high-frequency GW detectors such as ADMX are natural instruments for probing ultralight bosons through superradiance around PBHs.* For reference, Table 5.1 displays this mapping for frequencies spanning from low-RF band to the GHz range. This illustrates that the MHz–GHz frequency range associated with superradiant gravitational atoms spans and extends beyond the current ADMX sensitivity window, motivating broader frequency coverage in future high-frequency gravitational-wave searches.

5.2.2 Superradiant Instability Rates and Growth Timescales

The cloud occupation number grows as $\dot{N} = \Gamma_{n\ell m}\, N$, where $\Gamma_{n\ell m} = \mathrm{Im}(\omega)$ is the superradiant instability rate. In the non-relativistic limit ($\alpha \ll 1$) this rate can be computed analytically by matched-asymptotic methods [71, 43, 33]:

$$\Gamma_{n\ell m} = \frac{2r_+}{M} C_{n\ell} g_{\ell m}(m\Omega_H - \omega_{n\ell m})\alpha^{4\ell+5}, \tag{5.12}$$

with

$$C_{n\ell} = \frac{2^{4\ell+1}(n+\ell)!}{n^{2\ell+4}(n-\ell-1)!}\left(\frac{\ell!}{(2\ell)!(2\ell+1)!}\right)^2, g_{\ell m} = \prod_{k=1}^{\ell}\left[k^2(1-a^{*2}) + \left(a^* m - 2r_+\omega_{n\ell m}\right)^2\right]$$

where $C_{n\ell}g_{\ell m}(a_*)$ is a dimensionless, spin-dependent coefficient. Meanwhile, for a rotating Kerr black hole with angular momentum per unit mass $a = J/Mc$, the event

Table 5.1: Mapping between boson mass μ, host BH mass at $\alpha = 0.2$, and annihilation GW frequency ω_{ann}. The ADMX Run 1 data-collection band (0.645–1.4 GHz) lies between the third and fourth rows, with the corresponding frequencies bracketing the experimental sensitivity range.

μ	$M_{\text{BH}}/M_{\odot}$	ω_{ann}	Band
10^{-9} eV	2.7×10^{-2}	484 kHz	low-RF
10^{-7} eV	2.7×10^{-4}	48 MHz	HF (broadband RF)
10^{-6} eV	2.7×10^{-5}	484 MHz	∼ADMX Run 1 (lower)
3×10^{-6} eV	9×10^{-6}	1.45 GHz	ADMX Run 1 (upper)
10^{-5} eV	2.7×10^{-6}	4.8 GHz	Next-generation RF

horizon is located at:

$$r_{+} = \frac{GM}{c^2}\left(1 + \sqrt{1 - \frac{a^2}{M^2c^2}}\right).$$

The e-folding time for cloud growth is

$$\tau_{\text{SR}} \equiv \Gamma_{n\ell m}^{-1}, \tag{5.13}$$

Figure 5.2a shows τ_{SR} in years as a function of α for modes $\ell = m = 1, \ldots, 4$; Fig. 5.2b shows the same rate in geometrized units $r_g\,\Gamma_{n\ell m}$, making the ℓ-dependent peak structure visible.

Three features of these figures are particularly relevant to our analysis:

1. *Steep α-dependence at small α.* Because $\tau_{\rm SR} \propto \alpha^{-(4\ell+5)}$, the growth time rises by many orders of magnitude as α decreases below ~ 0.1. For the ADMX-relevant PBH mass range, $\tau_{\rm SR}$ ranges from seconds at $\alpha \sim 0.4$ to timescales exceeding the Hubble time at $\alpha \lesssim 0.05$. This sets a practical lower bound $\alpha \gtrsim 0.05$ for any detectable cloud to have formed.

2. *Cutoff at the superradiance threshold.* Both figures show $\Gamma_{n\ell m} \to 0$ as $\alpha \to m/2$, because the superradiant enhancement $C_{n\ell} g_{\ell m}$ vanishes there. The instability therefore operates within a finite window $\alpha \in (0.05, m/2)$, peaking at $\alpha \sim 0.2$–0.3.

3. *Multi-mode competition.* For sufficiently high initial BH spin, several m-modes can satisfy the superradiance condition simultaneously. Although the $m = 1$ mode dominates the growth rate, modes with $m \geq 2$ develop non-negligible populations over long timescales, contributing sub-dominant but non-zero GW channels at their respective transition and annihilation frequencies.

Throughout Sec. 5.3 we compute instability rates from the matched-asymptotic expression of Eq. (5.12), supplemented at larger α by the WKB envelope of Ref. [22] (their Eq. 27), since Eq. (5.12) alone is formally valid only for $\alpha \ll 1$.

5.2.3 Cloud Saturation and Bosenova Cycling

Superradiant growth drives the BH spin from its initial value toward the Regge trajectory $a_* = a_*^{\rm crit}$, at which point the instability saturates. The fraction of BH mass transferred to the cloud at saturation is determined by simultaneous conservation of energy and angular momentum. For the dominant $\{2, 1, 1\}$ state at $\alpha \simeq 0.2$ and $a_* \to 1$ this fraction is approximately 10.8% [208], consistent with our numerical calculation of 10.81% shown in Fig. 5.4. The corresponding saturation occupation number is

$$N_{\rm sat} \approx \frac{G_N M_{\rm BH}^2\, \Delta a}{m} \approx 10^{65} \left(\frac{\Delta a}{0.1}\right) \left(\frac{M_{\rm BH}}{10^{-5}\, M_\odot}\right)^2 . \tag{5.14}$$

The enormous value of $N_{\rm sat}$—a consequence of the tiny boson mass—is what enables the macroscopic, coherent GW emission discussed in subsequent sections.

For bosons with attractive self-interactions, parametrized by the axion decay constant f_a, a separate ceiling on the occupation number arises from the *bosenova* instability [22, 215]. When the occupation number exceeds

$$N_{\rm bosenova} \;\approx\; 5 \times 10^{78} \frac{n^4}{\alpha^3} \left(\frac{M}{M_\odot}\right)^2 \left(\frac{f_a}{M_P}\right)^2 , \tag{5.15}$$

attractive self-interactions overcome the gravitational binding energy and the cloud collapses [215, 216]. The implosion deposits most of the extracted angular momentum back into the BH in a burst lasting $\sim r_g/c$, after which superradiance restarts at reduced spin. This cycling behavior caps the effective occupation number at $N_{\rm eff} \equiv \min(N_{\rm sat}, N_{\rm bosenova})$.

For QCD-axion values $f_a \sim 10^{16}\,$GeV we find $f_a/M_P \sim 10^{-3}$, giving

$$\frac{N_{\rm bosenova}}{N_{\rm sat}} \;\sim\; 5 \times 10^6 \left(\frac{n^4/\alpha^3}{10}\right) \left(\frac{M_{\rm BH}}{10^{-5}\, M_\odot}\right) \left(\frac{\mu}{10^{-6}\,{\rm eV}}\right) \left(\frac{f_a/M_P}{10^{-3}}\right)^2 , \tag{5.16}$$

so the bosenova threshold is not reached before saturation in the bulk of our PBH parameter space. We include $N_{\rm bosenova}$ as an upper bound in all strain calculations, but note that for QCD-axion-like couplings it is non-restrictive. For very small f_a (e.g., light moduli with $f_a \ll M_P$) the bosenova may dominate and the cloud cycles; such scenarios, while physically interesting, fall outside the scope of this paper.

Finally, we note that accreting matter onto the BH can spin it back toward $a_* = 1$, potentially restarting superradiance after partial spin-down or counteracting the spin-down while the cloud is growing. For isolated PBHs in vacuum—the scenario relevant to the PBH mass range considered here—negligible accretion is expected and superradiance proceeds unimpeded. The visible imprint of accretion is the left-hand suppression of the allowed region in Fig. 5.1: as $M_{\rm BH}$ decreases at fixed μ, the condition $\alpha < m/2$ becomes harder to satisfy if spin-up by accretion is competitive with superradiant spin-down. We treat this effect as absent in what follows.

5.3 Gravitational-Wave Signatures from Isolated Gravitational Atoms

Before examining the modifications introduced by a binary companion (Sec. 5.4), we establish the GW signals from an *isolated* gravitational atom (GA) as the baseline against which binary-induced effects will be compared. Two physically distinct emission channels are active once superradiance saturates:

1. *Level transitions* (Sec. 5.3.1), in which the gravitational self-interaction of two simultaneously occupied superradiant levels drives bosons from the higher to the lower state, emitting quasi-monochromatic GWs at the Bohr frequency $\omega_{\rm tr}$.

2. *Boson annihilation* (Sec. 5.3.2), in which pairs of cloud bosons annihilate into a single graviton at $\omega_{\rm ann} \simeq 2\mu_a$, producing an extremely long-lived, slowly decaying signal.

In both cases the signal is quasi-monochromatic and persists on timescales far exceeding the ring-up time $\tau_{\rm ring}$ of a resonant-cavity detector; for these persistent sources the steady-state cavity-sensitivity comparison is therefore the appropriate limiting case. The time-domain waveforms derived here serve as inputs to the frequency-domain strain templates of App. 7.4.

Throughout this section we assume the BH spin has been reduced to the saturation value $a_*^{\rm crit}$ (Eq. (5.7)) and that the cloud is in a well-defined two-level configuration consisting of an "excited" state with quantum numbers (n_e, ℓ, m) and a "ground" state (n_g, ℓ, m), with $n_e > n_g$. All cloud masses are computed numerically using the results of Refs. [198, 163], and superradiance rates use Eq (5.12).

5.3.1 Gravitational Waves from Level Transitions

5.3.1.1 Two-level population dynamics

When the BH spin lies above the critical value for two levels simultaneously, both states are initially populated independently through superradiance. For occupation numbers well below saturation, where the gravitational self-interaction between levels is negligible, each level grows independently:

$$\frac{dN_g}{dt} = N_g \, \Gamma_{\rm SR}^{(n_g)}, \tag{5.17}$$

$$\frac{dN_e}{dt} = N_e \, \Gamma_{\rm SR}^{(n_e)}, \tag{5.18}$$

where $\Gamma_{\rm SR}^{(n)}$ is the superradiant instability rate for level n (Eq. (5.12)). The $\{6, 4, 4\}$ and $\{5, 4, 4\}$ states have nearly equal instability rates at the benchmark parameters of Fig. 5.3, so both grow on comparable timescales.

This picture changes once the occupation numbers become large enough that the gravitational self-interaction of the cloud is no longer negligible. A macroscopic cloud in state $|n_e, \ell, m\rangle$ generates a time-dependent potential $V_*(t, \boldsymbol{r})$ that mixes the two levels through the matrix element $\langle n_g, \ell, m|V_*(t, \boldsymbol{r})|n_e, \ell, m\rangle$. This introduces a coherent transition rate Γ_t, coupling the two occupation numbers:

$$\frac{dN_g}{dt} = N_g \, \Gamma_{\rm SR}^{(n_g)} + N_g N_e \, \Gamma_t, \tag{5.19}$$

$$\frac{dN_e}{dt} = N_e \, \Gamma_{\rm SR}^{(n_e)} - N_g N_e \, \Gamma_t. \tag{5.20}$$

The term $-N_g N_e \, \Gamma_t$ acts as a stimulated-emission drain on the excited state: as N_e grows, the transition rate accelerates, depopulating $|n_e\rangle$ and populating $|n_g\rangle$ while emitting a GW at the Bohr frequency. The per-boson transition rate is obtained from the gravitational quadrupole formula applied to the off-diagonal current between the two levels [22, 21]:

$$\Gamma_t \;\sim\; \frac{2\, G_N \, \omega_{\rm tr}^5}{5} \, \mu_a^2 \, r_c^4 \;=\; \mathcal{O}(10^{-6}\text{–}10^{-8}) \, \frac{G_N \, \alpha^9}{r_g^3} \,, \tag{5.21}$$

where $\omega_{\rm tr}$ is the transition frequency (Eq. (5.22)), $r_c \sim n^2 r_g/\alpha^2$ is the characteristic cloud radius, and $r_g = GM$ is the gravitational radius. The scaling $\Gamma_t \propto \alpha^9/r_g^3$ follows from substituting $r_c^4 \propto r_g^4/\alpha^8$ and $\omega_{\rm tr} \propto \mu_a \alpha^2 \propto \alpha^3/r_g$ into the first form. The numerical range reflects the variation across quantum numbers n and ℓ for the configurations of interest [22].

5.3.1.2 Time-domain gravitational-wave strain

The GW angular frequency for transitions between states (n_e, ℓ, m) and (n_g, ℓ, m) with $n_e > n_g$ is the Bohr frequency:

$$\omega_{\rm tr} \;=\; \omega_{n_g \ell m} - \omega_{n_e \ell m} \;\simeq\; \frac{\mu_a\, \alpha^2}{2} \left(\frac{1}{n_g^2} - \frac{1}{n_e^2} \right), \tag{5.22}$$

where we have used the leading-order hydrogenic approximation. For the $\{6, 4, 4\} \to \{5, 4, 4\}$ transition at the benchmark parameters of Fig. 5.3, this gives $\omega_{\rm tr} \approx 833.03\,\mathrm{MHz}$.

The oscillating quadrupole moment of the two-level system sources a GW whose strain envelope follows from the standard power–strain relation $h^2 = 4G_N P_{\rm tr}/(r^2\omega_{\rm tr}^2)$ [16]. Since each transition quantum carries energy $\omega_{\rm tr}$ and the quanta are emitted at rate $\Gamma_t N_g N_e$, the radiated power is $P_{\rm tr} = \omega_{\rm tr}\, \Gamma_t N_g(t) N_e(t)$; substituting and cancelling one power of $\omega_{\rm tr}$ gives: [22, 21]:

$$h_{0,{\rm tr}}(t) \;=\; \sqrt{\frac{4\, G_N}{r^2\, \omega_{\rm tr}}\, \Gamma_t\, N_g(t)\, N_e(t)}\,, \tag{5.23}$$

where $N_g(t)$ and $N_e(t)$ are solutions of Eqs. (5.19)–(5.20). This is the *slowly-varying envelope* of the GW signal: it evolves on timescales set by $\Gamma_{\rm SR}$ and Γ_t, which are much longer than the GW oscillation period $2\pi/\omega_{\rm tr}$. The full waveform, decomposed into the standard plus and cross polarizations, is:

$$h_+(t) = h_{0,\mathrm{tr}}(t)\,\frac{1+\cos^2\iota}{2}\,\cos\big(\omega_{\mathrm{tr}}\,t + \phi_0\big), \tag{5.24}$$

$$h_\times(t) = h_{0,\mathrm{tr}}(t)\,\cos\iota\,\sin\big(\omega_{\mathrm{tr}}\,t + \phi_0\big), \tag{5.25}$$

where ι is the inclination of the BH spin axis with respect to the line of sight and ϕ_0 is an arbitrary initial phase. The signal is quasi-monochromatic at ω_{tr} with a slowly drifting amplitude; it is the amplitude modulation through $h_{0,\mathrm{tr}}(t)$ that, in the frequency domain (App. 7.4), produces the characteristic narrow Lorentzian lineshape at $f = \omega_{\mathrm{tr}}/2\pi$.

5.3.1.3 Saturation occupation number and peak strain

The peak GW strain occurs at the moment when the coupled system (Eqs. (5.19)–(5.20)) transitions from the superradiance-dominated to the transition-dominated regime, i.e., when $\Gamma_t N_e \sim \Gamma_{\mathrm{SR}}^{(n_e)}$. This equal population condition corresponds to the moment of maximum stimulated emission and thus maximum GW power. Specifically, the maximum gravitational wave power occurs when the excited state equals the ground state

$$N_{tr} = \frac{\Gamma_{n\ell m}}{\Gamma_t}. \tag{5.26}$$

Note that this level transition always occurs before the cloud population becomes saturated. In Fig. 5.3, the transition occurs when $N_g = N_e = 10^{55}$, and the population continues to grow until $N_{sat} \approx 10^{62}$. Substituting $N_g \sim N_e \sim N_{\mathrm{tr}}/2$ into Eq. (5.23) gives the peak strain:

$$h_{0,\mathrm{tr}}^{\mathrm{peak}} \simeq \frac{1}{2}\sqrt{\frac{4\,G_N\,\Gamma_{n\ell m}^2}{r^2\,\omega_{\mathrm{tr}}\Gamma_t}} \simeq 8\times10^{-31}\left(\frac{1\,\mathrm{au}}{r}\right)\left(\frac{\alpha}{0.3}\right)^{15}\left(\frac{M}{10^{-7}\,M_\odot}\right). \tag{5.27}$$

where the numerical scaling absorbs G_N, r_g, and $N_{\mathrm{sat}}(\alpha, M)$, and is normalized to the $\ell = 4$ transition ($\{6,4,4\} \to \{5,4,4\}$) shown in Fig. 5.3. Higher-ℓ transitions are far weaker than the dominant $\ell = 1$ channel: the wavefunction overlap entering $\Gamma_{n\ell m}$ and Γ_t suppresses $h_{0,\mathrm{tr}}^{\mathrm{peak}}$ by roughly two additional orders of magnitude per unit of ℓ, and steepens the α-scaling from α^3 at $\ell = 1$ to α^{15} here. This scaling should be treated as an order-of-magnitude guide rather than a precise formula, since the coefficient also depends sensitively on the black-hole spin near the superradiant threshold.

Using the benchmark parameters of $M = 10^{-7}\,M_\odot$ and $\alpha = 0.75$ we obtain $h_{0,\mathrm{tr}}^{\mathrm{peak}} \approx 7 \times 10^{-27}$ at $r = 100$ au, in good agreement with the numerically computed peak $h_{\max} \approx 6 \times 10^{-27}$ seen in Fig. 5.3. The smallness of this value relative to the optimal $\ell = 1$ case reflects the $\ell = 4$ level structure adopted there, rather than the maximal achievable transition strain.

5.3.1.4 Signal lifetime

Once the transition-dominated phase begins, the excited state drains with a characteristic timescale

$$\tau_{\mathrm{tr}} \sim \frac{1}{\Gamma_t\, N_{\mathrm{sat}}} \sim \mathcal{O}(10^3\text{–}10^6)\ \mathrm{yr}, \tag{5.28}$$

for the parameters relevant to MHz–GHz emission. Since $\tau_{\mathrm{tr}} \gg \tau_{\mathrm{ring}} \sim 1/\Delta f_{\mathrm{band}} \sim 10^{-5}\,\mathrm{s}$ (for a cavity bandwidth $\Delta f_{\mathrm{band}} \sim 10\,\mathrm{kHz}$), the level-transition signal is effectively persistent on detector timescales. The extremely narrow spectral width $\Delta f \sim \tau_{\mathrm{tr}}^{-1}/(2\pi) \ll 1\,\mathrm{Hz}$ (compared to the carrier frequency $f_{\mathrm{tr}} \sim$ MHz–GHz) makes this signal highly coherent over any realistic observation window; the resulting frequency-domain Lorentzian template is developed in App. 7.4.

5.3.2 Gravitational Waves from Boson Annihilation

5.3.2.1 Annihilation frequency and rate

When two cloud bosons annihilate they produce a single graviton. Each boson in the bound state has energy $\omega_b = \omega_{n\ell m}$, so the emitted graviton carries

$$\omega_{\text{ann}} \;=\; 2\,\omega_{n\ell m} \;\simeq\; 2\mu_a\left(1 - \frac{\alpha^2}{2n^2}\right), \tag{5.29}$$

to leading order in α. The $\mathcal{O}(\alpha^2)$ correction shifts the emission frequency slightly below $2\mu_a$; for the benchmark parameters of Fig. 5.5 this gives $f_{\text{ann}} = \omega_{\text{ann}}/(2\pi) \approx 833.09\,\text{MHz}$.

The pair-annihilation rate per boson—the rate at which a single pair converts into a graviton—is obtained from the gravitational quadrupole formula applied to the $2\mu_a$ oscillation of the cloud's stress–energy tensor [21]:

$$\Gamma_a \;\simeq\; K_{n\ell m}\left(\frac{\alpha/\ell}{0.5}\right)^p \frac{G_N}{r_g^3}\,, \tag{5.30}$$

where the exponent $p = 17$ for $\ell = 1$ and $p = 4\ell + 11$ for $\ell \geq 2$ is determined by the power of α in the hydrogenic wavefunction overlap, and $K_{n\ell m} \approx 10^{-10}$ is a dimensionless coefficient that depends weakly on n for $n \lesssim \ell + 2$ [21, 43]. Since $\Gamma_a \propto \alpha^{4\ell+11}/r_g^3$ while $\Gamma_{\text{SR}} \propto \alpha^{4\ell+5}/r_g$ and $\Gamma_t \propto \alpha^9/r_g^3$, the hierarchy $\Gamma_{\text{SR}} \gg \Gamma_t \gg \Gamma_a$ holds over the entire parameter space of interest, confirming that annihilation is the slowest of the three processes.

5.3.2.2 Time-domain occupation number and strain

After superradiance saturates the cloud at $N_{\rm max}$ and the level-transition phase has depleted the excited state, the occupation number of the surviving ground level evolves under pure annihilation. Solving $dN/dt = -\Gamma_a N^2$ gives:

$$N(t) \;=\; \frac{N_{\rm max}}{1+\Gamma_a\, N_{\rm max}\, t}\,, \tag{5.31}$$

starting from $t=0$ at saturation. The GW strain envelope follows directly from the quadrupole luminosity:

$$h_{0,\rm ann}(t) \;=\; N(t)\sqrt{\frac{4\,G_N}{r^2\,\omega_{\rm ann}}\,\Gamma_a} \;=\; \frac{N_{\rm max}}{1+\Gamma_a N_{\rm max} t}\sqrt{\frac{4\,G_N\,\Gamma_a}{r^2\,\omega_{\rm ann}}}\,, \tag{5.32}$$

which is the slowly-varying amplitude of the wave [21]. The full waveform polarizations are:

$$h_+(t) = h_{0,\rm ann}(t)\,\frac{1+\cos^2\iota}{2}\,\cos\big(\omega_{\rm ann}\,t+\phi_0\big), \tag{5.33}$$

$$h_\times(t) = h_{0,\rm ann}(t)\,\cos\iota\,\sin\big(\omega_{\rm ann}\,t+\phi_0\big), \tag{5.34}$$

where the inclination ι and phase ϕ_0 carry the same meaning as in Eqs. (5.24)–(5.25). The envelope $h_{0,\rm ann}(t)$ includes the effects of BH mass and spin loss, which cause $N_{\rm max}$ and $\omega_{\rm ann}$ to decrease slowly over time; these corrections are numerically small ($\lesssim 10\%$ over cosmic timescales for the configurations considered) and are incorporated in our numerical calculations.

5.3.2.3 Peak strain scaling

The peak strain, attained at $t = 0$ (saturation), is:

$$h_{0,\text{ann}}^{\text{peak}} = N_{\max} \sqrt{\frac{4\, G_N\, \Gamma_a}{r^2\, \omega_{\text{ann}}}} \simeq 2 \times 10^{-20} \left(\frac{1\,\text{au}}{r}\right) \left(\frac{\alpha/\ell}{0.5}\right)^{p/2} \left(\frac{M}{10^{-6}\, M_\odot}\right)$$
$$p = 17\ (\ell = 1), \quad 4\ell + 11\ (\ell \geq 2), \tag{5.35}$$

where the numerical scaling uses $N_{\max} \sim M^2(\tilde{a}_*^{\text{init}} - a_*^{\text{crit}}) \cdot \omega_{n\ell m}/m$ and the Γ_a scaling of Eq. (5.30). For the benchmark parameters of Fig. 5.5, the simulated peak is $h_{0,\text{ann}}^{\max} \approx 10^{-21}$, consistent with this estimate.

5.3.2.4 Signal lifetime and detectability window

The characteristic decay time of the annihilation signal is:

$$\tau_{\text{ann}} \equiv \frac{1}{\Gamma_a N_{\max}} \sim \frac{r_g^3}{G_N K_{n\ell m}} \left(\frac{0.5}{\alpha/\ell}\right)^p \frac{1}{N_{\max}}, \tag{5.36}$$

which, for the parameters relevant to MHz–GHz emission, evaluates to $\tau_{\text{ann}} \sim 10^{25}$–$10^{60}$ yr. This vastly exceeds both the Hubble time and any conceivable observation period, confirming that the annihilation signal is effectively a *monochromatic steady source* for the purposes of detection. The extremely slow decay also means that $N(t) \approx N_{\max}$ throughout any observation, so $h_{0,\text{ann}}$ is well approximated by its peak value.

Like the transition signal, the annihilation signal is continuous on detector timescales ($\tau_{\text{ann}} \gg \tau_{\text{ring}}$), and its spectral width $\Delta f \sim \tau_{\text{ann}}^{-1}/(2\pi) \lesssim 10^{-33}$ Hz is extraordinarily narrow. The frequency-domain template, detailed in App. 7.4, takes the form of an exponential-integral function peaked at $f_{\text{ann}} = \omega_{\text{ann}}/(2\pi)$ with a width set by $\Gamma_a N_{\max}$.

5.3.3 Summary and comparison of isolated-GA signals

Table 5.2 collects the key properties of the two isolated-GA GW channels for the benchmark parameter sets introduced in Table 5.1. The level-transition and annihilation signals differ in amplitude, duration, and spectral width, but share the property that both persist for times enormously exceeding the ring-up time of any resonant-cavity detector. This motivates the use of characteristic strain $h_c(f) = 2f|\tilde{h}(f)|$ as the figure of merit (App. 7.4), since the large number of coherent cycles accumulated over $\tau_{\rm tr}$ or $\tau_{\rm ann}$ enhances the characteristic strain above the raw amplitude by a factor $\sqrt{f\,\tau}$.

The key conclusion of this section is that *both* isolated-GA emission channels produce continuous, narrowband GWs in the MHz–GHz band for PBH masses in the range $M_{\rm BH} \sim 10^{-16}$–$10^{-4}\,M_\odot$. Their detectability at a specific detector depends on the signal duration relative to $\tau_{\rm ring}$ (which is easily satisfied) and on whether the peak strain exceeds the detector noise floor.

5.4 Gravitational-Wave Signatures from Gravitational Atoms in Binaries

The emission channels studied in Sec. 5.3 treat the gravitational atom (GA) as evolving in isolation. In realistic astrophysical environments, however, a GA may reside in a binary system whose companion's tidal field drives resonant transitions between cloud levels. When the orbital frequency Ω_0 sweeps through a Bohr frequency of the gravitational atom, the system behaves as a driven two-level quantum system

Property	Level transition	Annihilation
Emission frequency	$\omega_{\rm tr}/(2\pi)$, Eq. (5.22)	$\omega_{\rm ann}/(2\pi)$, Eq. (5.29)
Rate controlling amplitude	Γ_t, Eq. (5.21)	Γ_a, Eq. (5.30)
Peak strain at 1 kpc	$\sim 10^{-23}$	$\sim 10^{-22}$
Signal lifetime	$\tau_{\rm tr} \sim 10^3$–$10^6\,$yr, Eq. (5.28)	$\tau_{\rm ann} \sim 10^{25}$–$10^{60}\,$yr, Eq. (5.36)
Spectral width Δf	$\sim \tau_{\rm tr}^{-1}/(2\pi) \ll 1\,$Hz	$\lesssim 10^{-33}\,$Hz
$\Delta f/f$	$\ll 10^{-6}$	$\ll 10^{-30}$
Signal type	Quasi-monochromatic transient	Persistent monochromatic

Table 5.2: Comparison of the two isolated-GA gravitational-wave channels for the PBH benchmark parameters of Table 5.1. Both signals satisfy $\tau \gg \tau_{\rm ring}$ and are candidates for resonant-cavity detection in the MHz–GHz band. The level-transition signal is stronger at early times but decays over $\tau_{\rm tr}$; the annihilation signal is weaker but essentially eternal.

and can undergo a Landau–Zener resonance [33, 34]. In this section we apply the binary-perturbation formalism of Kyriazis & Yang [147] to PBH binaries in natural units $c = \hbar = 1$, compute the characteristic GW strain of the transient burst produced by the resonant $\{2, 1, 1\} \to \{2, 1, -1\}$ hyperfine transition, and assess the prospects for detection with ADMX.

5.4.1 Time- and Frequency-Domain Waveforms

The gravitational-wave strain produced by tidal perturbations of a GA cloud from its binary companion can be expressed as [147]

$$h_{+,211}(t) = h_0 \, \frac{1+\cos^2\iota}{2} \, \mathrm{Re}\Big[e^{-2i\Delta m\,\varphi(t_{\mathrm{re}})}\, Q(t_{\mathrm{re}})\Big]\,, \tag{5.37}$$

$$h_{\times,211}(t) = h_0 \, \cos\iota \, \mathrm{Im}\Big[e^{-2i\Delta m\,\varphi(t_{\mathrm{re}})}\, Q(t_{\mathrm{re}})\Big]\,, \tag{5.38}$$

where ι is the inclination angle between the line of sight and the normal to the binary orbital plane, and t_{re} is the retarded time. The oscillating term $e^{-2i\Delta m\varphi(t_{\mathrm{re}})}$ represents the carrier wave, where $-2\Delta m\varphi(t_{\mathrm{re}}) = \Phi_{GW}(t)$ represents the corresponding GW phase. The binary orbital phase is $\varphi(t) = \Omega_0 t + \frac{1}{2}\gamma t^2 + \phi_0$, which assumes a linearly chirping orbit with rate $\dot{\Omega} = \gamma$ (see Eq. (5.42) below) within an arbitary phase factor of ϕ_0. The instantaneous GW carrier frequency is then defined by $\dot{\Phi}_{GW}(t) = -2\Delta m\Omega_0 - 2\Delta m\gamma t$, resembling a linear chirping signal. The quantity $\Delta m \equiv m_f - m_i = -2$ is the difference in azimuthal quantum numbers between the final ($m_f = -1$) and initial ($m_i = +1$) states of the $\{2,1,1\} \to \{2,1,-1\}$ transition. The complex modulation $Q(t)$ captures the dynamics of the level population transfer: it encodes the evolving occupation-number amplitudes of the two mixed states as the binary sweeps through resonance, and its derivation and explicit form are given in Ref. [147]. The exponential factor $e^{-2i\Delta m\,\varphi}$ carries the fast orbital-phase oscillation, while the variation of $Q(t)$ is slow compared to $1/\Omega_0$.

The overall GW amplitude is set by

$$h_0 = \frac{24\, q_c\, GM}{r}\, \frac{(GM\Omega_0)^2}{\alpha^4}\,, \tag{5.39}$$

where M is the mass of the host BH (the primary of the gravitational-atom–binary, or GAB, system), Ω_0 is the binary orbital frequency at resonance, r is the observer distance, and $q_c \equiv \frac{M_{\rm cloud}}{M - M_{\rm cloud}}$ is the cloud-to-BH mass ratio. The latter is determined by the saturation condition [147]:

$$q_c = \frac{8\alpha^2 \left(1 - \dfrac{\alpha a_*}{m}\right)}{m^2 \left(1 - \sqrt{1 - \left(\dfrac{4\alpha}{m}\left(1 - \dfrac{\alpha a_*}{m}\right)\right)^2}\,\right)} - 1\,, \tag{5.40}$$

where m is the azimuthal quantum number of the initial state and a_* is the dimensionless BH spin. Throughout we set $a_* = a_*^{\rm crit}(m_i)$ (Eq. (5.7)), i.e., the spin is taken at the Regge-trajectory saturation value that maximizes α for fixed boson mass.

For the rest of the section, we shall only concern ourselves with the "plus" polarization h_+, as the two strain polarizations only differ by a relative phase and inclination weight, which can be observed through the comparison between Eq. (5.37) and Eq. (5.38) for the time domain, and proved within Section 7.4.3 for the frequency domain.

In frequency space, the Fourier transform of $h_+(t)$ takes the analytic form [147]

$$\tilde{h}_+(f) = h_0(1+\cos^2\iota)\,\sqrt{\pi}\,|\Delta m|^2\, i\, e^{i\Psi_+(f)}\, \frac{\sqrt{z}}{|\Gamma| - i\pi(f - f_c)}\, e^{-\pi z}\, \exp\left[-2z \arctan\left(\frac{\pi(f - f_c)}{|\Gamma|}\right)\right] \tag{5.41}$$

where $f_c = 2\Omega_0/\pi$ is the carrier frequency (four times the orbital frequency divided by 2π), and the GW phase is $\Psi_+(f) = 2\pi f r + 4\pi^2(f - f_c)^2/(4|\Delta m|\gamma) - \pi/4$. The quantity

$\Gamma \equiv \Gamma_{\rm SR}^{(2,1,-1)}$ is the superradiant instability rate of the *final* state $\{2,1,-1\}$ (Eq. (5.12)), which governs how rapidly bosons that have transitioned into $m=-1$ fall back into the BH; it sets the spectral width of the signal. Equation (5.41) describes a Lorentzian-shaped frequency-domain burst centered at the peak frequency f_p (Eq. (5.45)), with width $\sim |\Gamma|/\pi$.

The orbital chirp rate, derived from the leading-order gravitational-wave energy loss of a circular binary with mass ratio $q \equiv M_c/M$, is

$$\frac{\gamma}{\Omega_0^2} = \frac{96}{5}\,\frac{q}{(1+q)^{1/3}}\,(GM\Omega_0)^{5/3}\,. \tag{5.42}$$

The adiabaticity parameter z governs the degree of population transfer during the resonance. It is defined as

$$z \equiv \frac{\eta^2}{|\Delta m|\,\gamma}\,, \tag{5.43}$$

where $\eta = \langle \psi_f | V_*(t,\boldsymbol{r}) | \psi_i \rangle$ is the tidal mixing amplitude between the $\{2,1,1\}$ and $\{2,1,-1\}$ states evaluated numerically from the companion's gravitational potential V_*. In the limit $z \ll 1$ the resonance is non-adiabatic: the orbital frequency sweeps through the Bohr frequency too rapidly for the levels to equilibrate, and only a fraction $\sim z$ of the cloud population is transferred. In the opposite limit $z \gg 1$ the resonance is adiabatic and the populations invert almost completely. In practice, however, the final state $\{2,1,-1\}$ has a non-zero instability rate $|\Gamma|$ that continuously drains bosons back into the BH, effectively truncating the transition and limiting the maximum transferred population regardless of z. This instability is fully accounted for in Eq. (5.41)

As the binary sweeps through the resonance band, the GW signal duration is

$$\Delta t = \frac{2|\Gamma|}{\gamma}(1+2z)\,. \tag{5.44}$$

The peak GW frequency is shifted from the carrier by the finite-z correction,

$$f_p = f_c - \frac{2z|\Gamma|}{\pi} = \frac{2\Omega_0}{\pi} - \frac{2z|\Gamma|}{\pi}\,. \tag{5.45}$$

For $z \ll 1$ the peak frequency approaches the orbital carrier, $f_p \to f_c$; for $z \gg 1$ the signal is redshifted from the carrier by a term proportional to $|\Gamma|$, quantifying the downward pull of the final-state decay.

5.4.2 Parameter Space Constraints and Cavity-Response Diagnostic

Several physical requirements simultaneously constrain the allowed (α, q) parameter space. The binary orbit is treated in the linear-chirp approximation (Eq. (5.42)), which is valid only when the fractional frequency change during the resonance is small: $\Delta\Omega/\Omega_0 = \gamma\,\Delta t/\Omega_0 \ll 1$. The annihilation timescale of the initial $\{2,1,1\}$ state (Eq. (5.36)) must also exceed the resonance band crossing time Δt, otherwise the cloud is depleted by annihilation before it can undergo the binary-driven transition. In practice, both of these are less restrictive than the most stringent condition, which requires the signal duration to be shorter than the binary merger timescale:

$$\Delta t < t_{\rm merge}\,, \tag{5.46}$$

where the Peters-formula merger time at orbital frequency Ω_0 is

$$t_{\rm merge} = \frac{5}{256}\,G^{-5/3}\,\frac{M}{(M\Omega_0)^{8/3}}\,\frac{(1+q)^{1/3}}{q}\,. \tag{5.47}$$

The condition (5.46) ensures that the resonant transition occurs while the binary is still intact—a prerequisite for the waveform model of Sec. 5.4.1 to apply. Figure 5.6 displays the ratio $\Delta t/t_{\rm merge}$ in the (α, q) plane, with contours at $\Delta t/t_{\rm merge} = 0.01, 0.1, 1$. The white region ($\Delta t/t_{\rm merge} > 1$) is excluded because the transition would not complete before merger. The figure shows that the allowed region is bounded by $\alpha \lesssim 0.26$ and $q \lesssim 0.01$. We enforce these constraints in all subsequent calculations in this section.

For resonant cavities such as ADMX, in order to implement the steady-state sensitivity curves within Sec. 5.4.3, a further parameter space constraint is introduced by the requirement of a fully rung up cavity. A useful steady-state diagnostic is whether the duration of the GW burst within the cavity bandwidth satisfies

$$\Delta t \gtrsim \tau_{\rm ring} \equiv \frac{1}{\Delta f_{\rm band}}\,, \tag{5.48}$$

where $\Delta f_{\rm band} \simeq f/Q$ is the cavity bandwidth and $\tau_{\rm ring}$ is the ring-up time required for the cavity to reach its steady-state response. A single cavity mode with resonance frequency $\omega_{\rm cav}$ and damping rate $\gamma_{\rm cav} = \omega_{\rm cav}/Q$ responds to a gravitational-wave strain $h(t)$ as a driven oscillator,

$$\dot{A}(t) + \left[\frac{\gamma_{\rm cav}}{2} + i\big(\omega_{\rm cav} - \omega_{\rm GW}(t)\big)\right] A(t) = \kappa h(t), \tag{5.49}$$

where $A(t)$ is the electromagnetic mode amplitude and κ encodes the magnetic field, cavity geometry, and mode-overlap factor. Equivalently,

$$A(t) = \kappa \int_{-\infty}^{t} dt'\, h(t') \exp\left[-\frac{\gamma_{\rm cav}}{2}(t-t') - i\int_{t'}^{t} dt''\big(\omega_{\rm cav} - \omega_{\rm GW}(t'')\big)\right]. \tag{5.50}$$

For a monochromatic signal with duration $\Delta t \gg \gamma_{\rm cav}^{-1}$ this expression reduces to the familiar steady-state response used in standard cavity sensitivity estimates.

For a chirping signal, which is the relevant case for this section as the binary sweeps through a resonant transition band that is much larger than the cavity bandwidth, the deposited energy is controlled instead by the time, $t_{\rm dwell} \sim \Delta\omega_{\rm cav}/|\dot{\omega}_{\rm GW}|$, the waveform spends within the cavity bandwidth, and the detailed evolution of the phase-difference between the chirping GW and the cavity mode across the cavity bandwidth. The latter determines the decoherence time, $t_{\rm decoherence}$, when the driving signal ceases to oscillate in-phase and coherent with the cavity mode to effectively deposit energy within the cavity, which limits the number of cavity mode oscillations that can be amplified by the driving force for a cavity ring up. For linear chirping signals $\omega(t) = \omega_{\rm cav} + \chi t$, an estimate for the decoherence time can be found as [194]:

$$t_{\rm decoherence} \sim \sqrt{\frac{\pi}{\chi}} \tag{5.51}$$

which can also be considered as the coherence time interval $t_{\rm coherence}$ that sets the maximum number of allowed oscillations, N_{max}, where the mode function amplitude grows.

Recall from Sec. 5.4.1, the chirping signal is: $\omega_{GW}(t) = 4\Omega_0 + 4\gamma t$. At the $z \ll 1$ limit where $f_p \to f_c$, we have $4\Omega_0 + 4\gamma t \sim \omega_{\rm cav} + \chi t$, such that $\chi = 4\gamma$. The decoherence time for our considerations then becomes:

$$t_{\rm decoherence} \sim \sqrt{\frac{\pi}{4\gamma}} \tag{5.52}$$

The criterion for determining whether a chirping signal is capable of satisfying the condition for applying the steady-state comparison is $\Delta t_{\rm eff} \gg \tau_{\rm ring}$, where $\Delta t_{\rm eff} \sim \min\{t_{\rm decoherence}, t_{\rm dwell}\}$, this ensures the GW signal is capable of fully ringing up the cavity within the allowed energy deposit time window.

To provide an estimation on binary source parameters that are capable of ringing up the cavity, in Figure 5.7 we plot the ratio of $\Delta t_{\rm eff}/\tau_{\rm ring}$ within the allowed parameter space defined by the binary merger time in Fig. 5.6 for an ADMX-like cavity with $Q \sim 8 \times 10^4$, but without restricting the cavity resonance frequency to the ADMX hardware band. Since $t_{\rm decoherence}$ (Eq. (5.52)) provides a tighter constraint against $\tau_{\rm ring}$ when compared to $t_{\rm dwell}$ within the parameter space, we take $\Delta t_{\rm eff} \sim t_{\rm decoherence}$. We identify GWs from the region where $t_{\rm decoherence}/\tau_{\rm ring} < 5$ as "Transiets", signals that would fail to ring up the cavity to apply the steady-state approximation.

However, we stress that Fig. 5.7 serves as an order-of-magnitude estimation, and should not be treated as a strict detection forecast. We also point out that generically, the ring-up condition should not be interpreted as a sharp detectability boundary, but only as a prerequisite for applying the steady-state sensitivity curves. Signals that fail to fully ring up the cavity could still deposit considerable amounts of energy to enable a measurable response. Nevertheless, for estimation purposes, in the remainder of the section we constraint ourselves within the region of Fig. 5.7 where the steady-state approximation remains valid.

5.4.3 Detectability Prospects for PBH Gravitational-Atom Binaries

We now apply the binary-perturbation formalism to PBH systems and assess detectability at ADMX. To avoid overestimating the signal we adopt conservative numerical inputs throughout: (i) relativistic instability rates from Refs. [128, 127] for the final-state decay rate Γ, rather than the leading-order analytic approximation of Eq. (5.12); (ii) numerically computed mixing amplitudes $\eta = \langle \psi_f | V_*(t, \boldsymbol{r}) | \psi_i \rangle$; and (iii) numerically computed cloud masses M_{cloud} from Refs. [198, 163].

5.4.3.1 Peak-frequency map in the PBH mass range

Figure 5.8 shows the peak GW frequency f_p (Eq. (5.45)) across the M_{BH}–α parameter space for several values of the boson mass μ_b, with the companion-to-host mass ratio fixed at $q = 10^{-4}$. We restrict to the PBH mass range $M_{\text{BH}} \in [10^{-16}, 10^{1}\, M_\odot]$ and exclude parameter combinations outside the allowed region of Fig. 5.6. For the hyperfine transition $\{2,1,1\} \to \{2,1,-1\}$, the superradiance condition $\alpha < m/2 = 0.5$ is automatically satisfied within the constrained region.

The figure demonstrates that GW signals from binary-perturbed GA systems naturally populate the ultra-high-frequency (UHF) band—including the GHz range targeted by ADMX—*if and only if* the host BH is of primordial origin. The ADMX Run-1 frequency scan range (0.645–1.4 GHz, shown in purple) is intersected by trajectories with $M_{\text{BH}} \sim 10^{-11}$–$10^{-12}\, M_\odot$ and $\mu_b \sim 2.5\,\text{eV}$ at $\alpha \sim 0.14$.

5.4.3.2 Benchmark waveform

Figure 5.9 shows the time-domain strain $h_+(t)$ and its Fourier transform $\tilde{h}(f)$ for the benchmark binary GA system with

$$\alpha = 0.14, \quad M = 6.68 \times 10^{-12}\, M_\odot, \quad q = 10^{-4}, \quad r = 10\,\mathrm{kpc}. \tag{5.53}$$

These parameters lie within the allowed region of Fig. 5.6 and Fig. 5.7, place the peak frequency $f_p = 1.001\,\mathrm{GHz}$ within the ADMX band, and correspond to the event rates computed in Sec. 5.5. The resulting signal is a quasi-monochromatic GW burst.

5.4.3.3 ADMX sensitivity comparison

The minimum GW strain detectable by a resonant cavity is [35]

$$\begin{aligned} h_{\min} = 3 \times 10^{-22} &\left(\frac{0.1}{\eta_n}\right)\left(\frac{8\,\mathrm{T}}{|\mathbf{B}|}\right)\left(\frac{0.1\,\mathrm{m}^3}{V_{\mathrm{cav}}}\right)^{5/6}\left(\frac{10^5}{Q}\right)^{1/2} \\ &\times \left(\frac{T_{\mathrm{sys}}}{1\,\mathrm{K}}\right)^{1/2}\left(\frac{1\,\mathrm{GHz}}{f}\right)^{3/2}\left(\frac{\Delta f}{10\,\mathrm{kHz}}\right)^{1/4}\left(\frac{1\,\mathrm{min}}{t_{int}}\right)^{1/4}, \end{aligned} \tag{5.54}$$

where η_n is the cavity–GW coupling coefficient, $|\mathbf{B}|$ is the axial magnetic field, V_{cav} is the cavity volume, Q is the quality factor, T_{sys} is the system noise temperature, f is the resonant frequency, $\Delta f \simeq f/Q$ is the cavity bandwidth, and t_{int} is the measurement integration time. For ADMX Run 1 the relevant parameters are $\eta_n \sim 0.1$, $|\mathbf{B}| = 0.75\,\mathrm{T}$, $V_{\mathrm{cav}} = 136\,\mathrm{L}$, $Q \sim 8 \times 10^4$, $T_{\mathrm{sys}} = 0.6\,\mathrm{K}$, , $f \in [0.65, 1.4]\,\mathrm{GHz}$, and $t_{int} = 2\,\mathrm{min}$ [35], giving a typical threshold $h_{\min} \sim 10^{-22}$ within the ADMX band. The GW characteristic strain for binary-induced transitions is $h_c(f) = 2f|\tilde{h}(f)|$ with $\tilde{h}(f)$ given by Eq. (5.41), Fig. 7.3 shows the characteristic strain for the benchmark parameters of Eq. (5.53).

Figure 5.10 displays both h_c (filled contours) and $h_{\rm min}$ (dashed curves) across the $M_{\rm BH}$–α plane at $q = 10^{-4}$ and $r = 10\,{\rm kpc}$. The dashed contours of $h_{\rm min}$ are evaluated at fixed values of the ADMX sensitivity parameter (Eq. (5.54)) but without restricting f to the ADMX hardware band; the orange solid curves then overlay the ADMX frequency range from Fig. 5.8 to delineate the region where the ADMX hardware could physically respond, during Run1 data, in turn corresponding to $h_{\rm min} \sim 10^{-22}$. The PBH mass range of primary interest, $10^{-16}\,M_\odot \lesssim M_{\rm PBH} \lesssim 1\,M_\odot$, is bounded by the black dotted lines.

For the benchmark parameters at $r = 10\,{\rm kpc}$, the characteristic strain h_c falls *several orders of magnitude* below the ADMX detection threshold throughout the accessible parameter space. No overlap between the h_c and $h_{\rm min}$ contours of equal magnitude is found within the PBH mass range, confirming that ADMX as currently configured cannot detect the binary-induced burst at astrophysically motivated distances.

The gap between h_c and the ADMX threshold persists even when one optimizes over the remaining free parameters. Increasing the mass ratio q toward its upper bound $q \lesssim 0.01$ (Sec. 5.4.2) enhances h_c only modestly, with the most direct path to improved sensitivity is reducing the source distance r: from Eq. (5.39), $h_0 \propto 1/r$, so ADMX threshold strains are reached only at $r \lesssim 1\,{\rm AU}$. Such distances are strongly disfavored by merger-rate arguments, as we demonstrate quantitatively in Sec. 5.5.

We reiterate that this comparison is intended as an order-of-magnitude benchmark rather than a complete chirping/transient-search forecast, as discussed in Sec. 5.4.2. For the isolated annihilation and level-transition signals of Sec. 5.3, whose lifetimes are vastly longer than any cavity response time, the steady-state sensitivity curve is the appropriate limiting case. For the binary-induced signals considered here, a rigorous analysis would require filtering the time-domain waveform through the cavity transfer function, or equivalently computing the total deposited electromagnetic energy in the relevant modes. Recent studies of electromagnetic cavities exposed to non-monochromatic high-frequency gravitational waves show that the loss of sensitivity for short or chirping signals need not be captured by a step-function ring-up criterion alone [14, 38, 194]. Regardless, our main conclusion for binary-induced gravitational-atom bursts is insensitive to this refinement: at astrophysically plausible distances the predicted strain lies many orders of magnitude below the current ADMX-equivalent sensitivity, and the event-rate argument requires detectable sources to occur at unrealistically small distances. A more complete detector-response calculation may change the precise numerical reach, but it cannot remove the combined strain and rate deficits identified here.

5.4.4 Doppler Frequency Shift for Solar-System Flybys

The only scenario in which a PBH binary could approach within $\sim 1\,\mathrm{AU}$ is a transient flyby through the solar system. In this case the relative velocity between Earth and the binary imprints a Doppler shift on the observed GW frequency. We estimate the maximum frequency shift by setting the flyby velocity equal to the local dark-matter virial velocity, $v_{\mathrm{DM}} \approx 200\,\mathrm{km\,s^{-1}}$:

$$\delta f = \frac{v_{\mathrm{DM}}}{c}\, f_p \;\approx\; 6.7 \times 10^{-4} \times f_p\,. \tag{5.55}$$

For the benchmark peak frequency $f_p = 1.001\,\mathrm{GHz}$, this gives $\delta f \approx 670\,\mathrm{kHz}$. Although this shift exceeds the instantaneous cavity bandwidth of ADMX ($\Delta f_{\mathrm{band}} \simeq f/Q \approx 13\,\mathrm{kHz}$) and would therefore displace the signal out of resonance during a naive fixed-frequency scan, it remains a small fraction of the full ADMX Run-1 frequency scan range $(1.4 - 0.645)\,\mathrm{GHz} = 755\,\mathrm{MHz}$. A flyby source would therefore be visible in some portion of the scan range, and a matched-filter search that accounts for a $\pm 1\,\mathrm{MHz}$ Doppler drift would recover the full signal duration.

We note, however, that the detectability problem identified in Sec. 5.4.3 is one of *strain* sensitivity, not frequency coverage: even at $r = 1\,\mathrm{AU}$ the strain falls roughly six orders of magnitude below the ADMX threshold. The Doppler correction is therefore not the obstacle to detection.

5.5 Prospects for Observation at High-Frequency Gravitational-Wave Detectors

The detectability analysis of Sec. 5.4.3 shows that the characteristic strain of the binary-induced burst falls several orders of magnitude below the ADMX threshold at $r = 10\,\mathrm{kpc}$, requiring systems which can produce binary-induced level transitions to exist at much closer distances, potentially $\ll 1$ AU. Whether such close encounters occur at any appreciable rate depends on the abundance and spatial distribution of PBHs, and on the statistics of PBH binary formation. Because the orbital separations that drive the $\{2,1,1\} \to \{2,1,-1\}$ resonant transition occur immediately prior to merger (Sec. 5.4.2), the question reduces to the local PBH merger rate filtered by the conditions on mass ratio and total binary mass identified in Sec. 5.4.3: $q \lesssim 10^{-4}$ and $M = 6.68 \times 10^{-12} M_\odot$.

In Ref. [191], merger rates were computed for PBHs formed in the early universe across four binary-formation channels and for generic mass functions. We focus on the *early two-body channel*, in which two neighboring PBHs form a gravitationally bound pair when their mutual attraction overcomes the Hubble flow. This channel is conceptually the cleanest, yields the most direct mapping to the (M, q) space, and dominates the total merger rate at low $f_{\rm PBH}$, where $f_{\rm PBH}$ is the PBH fraction of the dark-matter energy density. The competing *early three-body channel*, which can become dominant at large $f_{\rm PBH}$, requires a treatment of N-body perturbations to the initial binary orbit and carries larger theoretical uncertainties [191]; we discuss its potential impact briefly in Sec. 5.5.3 below.

5.5.1 Merger Rate from the Early Two-Body Formation Channel

The differential merger rate density per unit logarithmic mass for the early two-body channel is given in [191] as:

$$\frac{\mathrm{d}R_{E2}}{\mathrm{d}\ln m_1\,\mathrm{d}\ln m_2} \approx \frac{1.6\times 10^6}{\mathrm{Gpc}^3\,\mathrm{yr}}\, f_{\rm PBH}^{53/37}\,\eta_{\rm PBH}^{-34/37}\left(\frac{M_{\rm tot}}{M_\odot}\right)^{-32/37}\left(\frac{t}{t_0}\right)^{-34/37} S_L\,S_E\,\psi(m_1)\,\psi(m_2)\,, \tag{5.56}$$

where m_1 and m_2 are the masses of the two black holes in the binary, $\eta_{\rm PBH} \equiv \langle m\rangle^2/\langle m^2\rangle$ is the inverse normalized second moment of the mass function $\psi(m)$ (not to be confused with the tidal mixing amplitude η of Eq. (5.43)), t is the merger time, and t_0 is the age of the universe. The suppression factor

$$S_L \approx \min\left\{1,\ 0.01\left[\left(\frac{t}{t_0}\right)^{0.44} f_{\rm PBH}\right]^{-0.65} \exp\left(0.03\ln^2\left[\left(\frac{t}{t_0}\right)^{0.44} f_{\rm PBH}\right]\right)\right\} \tag{5.57}$$

accounts for the disruption of PBH binaries by late-time large-scale structure, and

$$S_E \approx \frac{\sqrt{\pi}\,(5/6)^{21/74}}{\Gamma(29/37)} \left[\frac{\langle m^2\rangle/\langle m\rangle^2}{\bar{N}(y)+C} + \frac{\sigma_M^2}{f_{\rm PBH}^2}\right]^{-21/74} e^{-\bar{N}(y)} \tag{5.58}$$

where

$$C = f_{\rm PBH}^2\,\frac{\langle m^2\rangle/\langle m\rangle^2}{\sigma_M^2}\left\{\left[\frac{\Gamma(29/37)}{\sqrt{\pi}}\,U\left(\frac{21}{74},\frac{1}{2},\frac{5f_{\rm PBH}^2}{6\sigma_M^2}\right)\right]^{-74/21} - 1\right\}^{-1} \tag{5.59}$$

captures the suppression from Poisson fluctuations in the number of neighboring PBHs within the relevant comoving volume. Here $\bar{N}(y)$ is the mean number of PBH neighbors within comoving separation y, and $\sigma_M^2/f_{\rm PBH}^2$ is the variance of the mass-weighted density contrast [191].

The rate in Eq. (5.56) is evaluated at $t = t_0$ (i.e., for mergers occurring today), which gives the present-day merger rate density $R_{E2}(t_0)$ relevant for ADMX observations. We integrate over (m_1, m_2) via Monte-Carlo to obtain the portion of the rate associated with binaries in any chosen (m_1, q) bin, using the relations $q = m_2/m_1$ (with $m_2 < m_1$).

5.5.2 Dichromatic Mass Function and Event Rates

To evaluate the merger rate for our benchmark system ($m_1 = 6.68\times10^{-12}\,M_\odot$, $m_2 = 6.68\times10^{-16}\,M_\odot$, so $q = 10^{-4}$); we examine two mass functions capable of producing large numbers of mixed-mass binary systems: a dichromatic mass function

$$\psi_{\rm di}(m) = f_1\,\delta(m - m_1) + f_2\,\delta(m - m_2)\,, \quad f_1 + f_2 = 1\,, \tag{5.60}$$

which places all PBH mass at two discrete values, and a double-lognormal mass function

$$\psi_{\rm dl}(m) = \sum_{i=1}^{2} \frac{f_i}{\sqrt{2\pi}\,\sigma\,m} \exp\left[-\frac{\ln^2(m/m_i)}{2\sigma^2}\right]\,, \quad f_1 + f_2 = 1\,. \tag{5.61}$$

Under Eq. (5.60) exactly three binary configurations are possible: equal-mass pairs at m_1, equal-mass pairs at m_2, and unequal-mass pairs (m_1, m_2) with $q = 10^{-4}$. Under Eq. (5.61), there is greater allowed variation in m_1 and q, but most systems obtain parameters close to the previous three scenarios. As $\sigma \to 0$, the double lognormal case should reduce to the dichromatic case.

The resulting present-day merger rate densities, computed by Monte-Carlo integration of Eq. (5.56) with $f_{\rm PBH} = 1$ (all dark matter in PBHs, a conservative upper bound on the rate), are summarized in Table 5.3. The equal-mass (m_2, m_2) configuration dominates at $R_{E2} = 4.9 \times 10^{17}\,{\rm Gpc}^{-3}\,{\rm yr}^{-1}$. The unequal-mass (m_1, m_2) configuration yields $R_{E2} \approx 2.3 \times 10^{17}\,{\rm Gpc}^{-3}\,{\rm yr}^{-1}$, and equal-mass (m_1, m_1) configurations yield $R_{E2} \approx 3.2 \times 10^{12}\,{\rm Gpc}^{-3}\,{\rm yr}^{-1}$. Only the (m_1, m_2) binaries satisfy the parameter space constraints $q \lesssim 10^{-2}$ bounded by Fig. 5.6; this is the rate we use for the observation probability estimate.

Binary type	q	R_{E2} [$\mathrm{Gpc}^{-3}\,\mathrm{yr}^{-1}$]	ADMX-relevant?
(m_1, m_1): $6.68 \times 10^{-12}, 6.68 \times 10^{-12}\,M_\odot$	1.0	3.2×10^{12}	No
(m_2, m_2): $6.68 \times 10^{-16}, 6.68 \times 10^{-16}\,M_\odot$	1.0	4.9×10^{17}	No
(m_1, m_2): $6.68 \times 10^{-12}, 6.68 \times 10^{-16}\,M_\odot$	10^{-4}	2.3×10^{17}	**Yes**

Table 5.3: Present-day merger rate densities from the early two-body channel (Eq. (5.56)) for the dichromatic mass function (Eq. (5.60)) with $m_1 = 6.68 \times 10^{-12}\,M_\odot$, $m_2 = 6.68 \times 10^{-16}\,M_\odot$, and $f_{\mathrm{PBH}} = 1$. Only (m_1, m_2) binaries satisfy the parameter space constraints $q \lesssim 10^{-2}$ in Fig. 5.6.

5.5.3 Characteristic Detection Distance and the Strain Gap

Given a volumetric merger rate R_{E2} [$\mathrm{Gpc}^{-3}\,\mathrm{yr}^{-1}$], the characteristic distance within which one event per year is expected is

$$d_{1\,\mathrm{yr}} = \left(\frac{3}{4\pi R_{E2}}\right)^{1/3}, \tag{5.62}$$

obtained by setting the expected number of events in a sphere of radius d equal to unity: $R_{E2} \cdot \frac{4}{3}\pi d^3 = 1$. For the ADMX-relevant unequal-mass (m_1, m_2) rate $R_{E2} = 2.3 \times 10^{17}\,\mathrm{Gpc}^{-3}\,\mathrm{yr}^{-1}$, Eq. (5.62) gives

$$d_{1\,\mathrm{yr}}^{(q=10^{-4})} \approx 1\,\mathrm{kpc}\,. \tag{5.63}$$

At this distance, the characteristic strain from Sec. 5.4.3 ($h_c \sim 10^{-37}$) falls $\mathcal{O}(10^{15})$ times below the ADMX sensitivity threshold $h_{\mathrm{min}} \sim 10^{-22}$. In order for this deficit to be made up, a source would need to be placed within approximately 100 kilometers of ADMX, which is wholly unrealistic both due to the extremely low event rates corresponding to mergers at this distance and the complication that systems within this distance to ADMX would likely be strongly perturbed by the presence of the Earth.

Figure 5.11 summarizes the characteristic distances $d_{1\,\mathrm{yr}}$ for all three binary classes in the dichromatic model. We find that a lognormal mass function centered on the same benchmark masses produces very similar results.

Effect of PBH clustering. Local overdensities of PBHs increase the effective merger rate by a factor δ relative to the homogeneous prediction, which rescales all merger rates in Eq. (5.56) by δ [191]. Since the event rate scales as $R \propto \delta$, the characteristic distance scales as $d_{1\,\mathrm{yr}} \propto \delta^{-1/3}$. The most optimistic clustering estimates in the literature give $\delta \lesssim 10^5$ [191], which reduces $d_{1\,\mathrm{yr}}$ by a factor of $(10^5)^{1/3} \approx 46$:

$$d_{1\,\mathrm{yr}}^{(q=10^{-4},\ \delta=10^5)} \approx \frac{1\,\mathrm{kpc}}{46} \approx 22\,\mathrm{pc}. \tag{5.64}$$

Even with this maximally optimistic clustering, the strain deficit from systems that merge at appreciable rates is improved by only one to two orders of magnitude—far from enough to make up the original $\mathcal{O}(10^{15})$ gap between the binary-induced level transition signal and the ADMX strain sensitivity.

Effect of 3-Body Channel. The prior analysis relies on estimating the merger rate of PBH systems generated via the "early 2-body" channel, where PBH binaries form shortly after PBH formation due to mutual gravitational attraction between close pairs.

However, as noted in [191], binaries can also form via encounters between early 2-body channels and isolated PBHs in what is known as the "3-body channel". These encounters lead to a population of PBH binaries with different orbital properties, which can increase their merger rates, especially when $f_{\rm PBH}$ approaches 1.0. We find that, while this *can* increase the PBH merger rate for the types of systems we are interested in, the increase in merger rate is typically not more than a few orders of magnitude, once again falling far below the level of enhancement we would need to bridge the detectability gap.

5.5.4 Implications for Future Detectors

The analysis above identifies two independent, multiplicative obstacles to detecting binary-driven GA transitions at present-day cavity experiments:

1. *Strain deficit*: At the natural merger-rate distance $d_{1\,\rm yr} \approx 1\,\rm kpc$, the signal strain is $\sim 10^{15}$ times below the ADMX threshold. A factor-of-10^{15} improvement in strain sensitivity would be required from this factor alone.

2. *Rate deficit*: Reducing the merger distance to a few hundred kilometers while maintaining an appreciate merger rate requires an overall rate enhancement of $(1\,\rm kpc/100\,km)^3 \sim 10^{38}$ over the homogeneous prediction—far beyond what any clustering model can provide.

These obstacles together imply that a qualitative change in experimental approach, rather than incremental improvements to existing cavities, is necessary for binary-GA signals to become observable. The specific detector properties that would help most are: (i) improved strain sensitivity, achievable in principle through larger effective volume, stronger magnetic fields, lower system temperature, or enhanced mode overlap; (ii) a response model and readout strategy optimized for transient and frequency-drifting sources, rather than only for steady-state monochromatic signals; and (iii) broader frequency coverage, especially at sub-GHz frequencies where heavier PBHs produce signals outside the current ADMX scan range. Recent proposals using deposited-energy calculations for non-monochromatic cavity signals, multimode cavities that can track frequency drift, and geographically separated detector networks provide promising directions for going beyond the single-mode steady-state approximation used as our baseline comparison [14, 38, 194].

The isolated-GA annihilation and level-transition signals studied in Sec. 5.3 are not subject to the same rate limitations, since they arise from persistent, continuous emission rather than merger events. Their detectability is governed instead by the peak strain at a given distance, with $h_{0,\mathrm{ann}}^{\mathrm{peak}} \sim 10^{-22}$ at $1\,\mathrm{kpc}$ (Eq. (5.35)). For sources within $\sim 1\,\mathrm{kpc}$—well within the Milky Way PBH halo—the annihilation signal of an isolated GA is a more promising target for current ADMX-band experiments than the binary-induced burst, and we highlight it as the primary observational target for near-term searches.

5.6 Discussion and Conclusions

We have presented a unified treatment of high-frequency gravitational-wave emission from superradiant boson clouds around primordial black holes, covering both isolated gravitational atoms and binary-perturbed systems. In this section we synthesize the main results, place them in the context of the broader superradiance and high-frequency GW literature, discuss the principal caveats and uncertainties, and identify the most promising directions for future work.

5.6.1 Summary and Key Conclusions

Parameter space and Regge trajectories (Sec. 5.2). The gravitational fine-structure constant $\alpha = GM_{\rm BH}\mu$ determines both the efficiency of superradiant cloud growth and the frequency of the emitted GWs. For efficient superradiance ($\alpha \in [0.05, 0.5]$) and GW emission in the MHz–GHz band, the superradiance condition selects boson masses $\mu \sim 10^{-8}$–10^{-5} eV and, through $M_{\rm BH} = \alpha/(G\mu)$, host BH masses $M_{\rm BH} \sim 10^{-6}$–$10^{-3}\, M_\odot$—precisely the primordial BH mass window (Table 1). The Regge trajectories make this mapping explicit: each allowed $(M_{\rm BH}, \mu, m)$ configuration traces a curve in the $(a_*, M_{\rm BH})$ plane along which the cloud grows, extracts angular momentum, and emits narrowband GWs. The maximum cloud mass fraction of 10.8%, obtained numerically for $\alpha \approx 0.2$ and $a_* \to 1$, is in close agreement with the analytic upper bound of Ref. [208].

Isolated gravitational atom (Sec. 5.3). For an isolated GA we computed time-domain waveforms for both the level-transition and boson-annihilation channels, together with analytic frequency-domain templates (Appendix 7.4). The key quantitative results are:

- *Annihilation*: peak strain $h_{0,\mathrm{ann}}^{\mathrm{peak}} \simeq 10^{-21}$ at $r = 145\,\mathrm{au}$, with emission frequency $\omega_{\mathrm{ann}} \approx 833.09\,\mathrm{MHz}$ for the benchmark system ($\mu_a = 2 \times 10^{-6}\,\mathrm{eV}$, $M_{\mathrm{BH}} = 1.7 \times 10^{-5}\,M_{\odot}$, $\alpha = 0.21$). The signal is effectively a monochromatic steady source with a decay timescale $\tau_{\mathrm{ann}} \sim 10^{25}$–$10^{60}\,\mathrm{yr}$, enormously exceeding both the Hubble time and any observation window.

- *Level transitions*: peak strain $h_{0,\mathrm{tr}}^{\mathrm{peak}} \simeq 6 \times 10^{-27}$ at $r = 100\,\mathrm{au}$, with emission frequency $f_{\mathrm{tr}} \simeq 833.03\,\mathrm{MHz}$ for the $\{6,4,4\} \to \{5,4,4\}$ benchmark. The signal lifetime $\tau_{\mathrm{tr}} \sim 10^{3}$–$10^{6}\,\mathrm{yr}$ is long compared to any observation period but finite, giving a spectral width $\Delta f \sim \tau_{\mathrm{tr}}^{-1}/(2\pi) \ll 1\,\mathrm{Hz}$.

- The frequency-domain templates take analytic closed forms: an exponential-integral E_1 expression for the annihilation signal, and a complex Lorentzian shifted by the carrier frequency ω_{ann} or ω_{tr} for the two polarizations (Appendix 7.4). Both persistent signals trivially satisfy the steady-state cavity-response condition $\tau_{\mathrm{signal}} \gg \tau_{\mathrm{ring}} \sim 10^{-5}\,\mathrm{s}$.

Binary-perturbed gravitational atom (Secs. 5.4 and 5.5). Applying the Landau–Zener resonance formalism of Ref. [147] to PBH binaries, we computed the characteristic strain of the transient GW burst produced when the binary sweeps through the $\{2,1,1\} \to \{2,1,-1\}$ hyperfine transition. The principal results are:

- The peak GW frequency f_p (Eq. (5.45)) lies naturally in the GHz band for host BH masses $M_{\rm BH} \sim 10^{-11}$–$10^{-12}\, M_\odot$ at $\alpha \approx 0.14$, confirming that PBH binaries are the only astrophysical system that can populate wthe ADMX frequency range through binary-induced superradiance.

- The conservative duration-versus-response diagnostic $\Delta t \gg \tau_{\rm ring}$ is satisfied throughout the allowed (α, q) parameter space of Fig. 5.7, regardless of resonant frequency. This is sufficient for applying the steady-state comparison within our conservative diagnostic model, but it is not a full detector-response calculation and is not by itself sufficient for detection.=

- The GW characteristic strain at $r = 1\,\text{kpc}$ falls *fifteen orders of magnitude* below the ADMX threshold $h_{\rm min} \sim 10^{-22}$. Detection would require sources within to be unrealistically close.

- The early two-body PBH merger rate [191] evaluated for the benchmark dichromatic mass function gives a characteristic event distance $d_{1\,\rm yr} \approx 1\,\text{kpc}$ for ADMX-relevant ($q = 10^{-4}$) binaries. Even with the maximally optimistic PBH clustering factor $\delta \sim 10^5$, this reduces only to $\approx 22\,\text{pc}$—still approximately thirteen orders of magnitude in strain below the ADMX threshold. We conclude that binary-driven level-transition events are not detectable with current ADMX sensitivity.

5.6.2 Comparison to Related Work

The GW signatures of superradiant clouds around *stellar-mass* BHs are well studied [22, 21, 43, 135], with constraints on ultralight bosons from LIGO/Virgo spin measurements and non-detections of continuous GWs establishing the most robust exclusion regions in the (μ, M_{BH}) plane for $\mu \sim 10^{-13}$–$10^{-12}\,\mathrm{eV}$ [170]. The present work extends this program into qualitatively new territory in two respects.

First, by focusing on PBH masses $M_{\mathrm{BH}} \lesssim 10^{-3}\,M_{\odot}$ we access boson masses $\mu \sim 10^{-7}$–$10^{-5}\,\mathrm{eV}$ for which no spin-measurement constraints currently exist, and GW frequencies in the MHz–GHz range that are inaccessible to LIGO-band searches. High-frequency GW experiments such as ADMX and the broadband detectors proposed in Refs. [9, 83] offer the only existing instrumental sensitivity in this band; our frequency-domain templates (Appendix 7.4) provide the waveform targets needed for a matched-filter search.

Second, the binary-perturbation analysis builds directly on the Landau–Zener framework within [34] and the recent extension by Kyriazis & Yang [147], which derived the analytic GW strain template for tidally driven transitions. Previous applications of the Landau–Zener formalism focused primarily on floating orbits, inspiral modifications, dephasing of the binary GW signal, and the depletion of clouds in the context of LISA-band sources [34], and on the resonance history in Ref. [208]. We are the first to apply the Kyriazis & Yang waveform template to PBH binaries in the MHz–GHz regime and to compare it directly to the sensitivity of an existing high-frequency detector, yielding a quantitative (negative) detectability result rather than a parametric estimate.

The bosenova bound on the cloud occupation number, following Refs. [215, 216], is found to be non-restrictive for QCD-axion-like couplings ($f_a \sim 10^{16}$ GeV) in the bulk of our PBH parameter space (Eq. (5.14)), so the strain amplitudes we derive are not suppressed by bosenova cycling. For very weakly coupled bosons ($f_a \ll M_{\rm Pl}$) the cloud would cycle and the strain would be reduced; we note this as a limitation and leave the quantitative treatment of bosenova-modified waveforms to future work.

5.6.3 Caveats and Limitations

Several idealizations warrant discussion.

Spin distribution of PBHs. Our strain calculations assume that the host PBH initially lies above the relevant Regge trajectory, and in several benchmark cases we take a near-extremal initial spin in order to maximize the extractable rotational energy and hence the saturation cloud occupation number. The existence of an abundant population of PBHs with such spins remains an open question. Standard radiation-era formation from nearly Gaussian perturbations tends to predict small natal spins [58, 67], while early matter domination, highly nonspherical or non-Gaussian collapse, phase-transition formation channels, subsequent PBH mergers, and environmental accretion can produce larger spins in more model-dependent scenarios [118, 206]. Consequently, our amplitudes should be interpreted as optimistic conditional estimates for the high-spin subpopulation, not as population-averaged predictions for all PBHs. If only a fraction f_{spin} of PBHs of the relevant mass satisfy $a_* > a_*^{\text{crit}}(\alpha, m)$, the corresponding source counts and event rates should be rescaled by this factor, or by the appropriate power of it when more than one spinning PBH is required. A full treatment of the joint PBH mass, abundance, binary, and spin distributions is beyond the scope of this work and is left for future study.

Isolated versus binary GAs. We have treated the isolated and binary channels independently, but in a realistic PBH population both may be active simultaneously or sequentially. Binary companions that sweep through resonance before the isolated cloud reaches saturation will partially deplete the cloud, suppressing the subsequent annihilation and level-transition signals. Conversely, an isolated cloud that has already undergone significant annihilation before the binary resonance is crossded will produce a weaker binary-induced burst than our estimates suggest. A self-consistent treatment that evolves both channels simultaneously remains to be carried out.

Detector response to chirps/transients. Our detectability estimates use the published ADMX-equivalent strain sensitivity as a baseline and supplement it with a simple duration-versus-ring-up diagnostic. This is adequate for the persistent isolated-GA signals, but it is not a substitute for a full detector-response analysis of binary-induced chirping/transient signals. For such sources, the relevant observable is the electromagnetic energy deposited in one or more cavity modes after convolution with the detector transfer function, and this response can differ from the naive expectation based solely on whether the source duration is longer or shorter than τ_{ring}. We have therefore interpreted the binary-burst sensitivity estimates conservatively and regard a dedicated matched-filter or deposited-energy analysis as an important next step.

Multi-level dynamics. Our binary waveform analysis follows Ref. [147] in treating the $\{2,1,1\} \to \{2,1,-1\}$ transition as a two-level system. In reality, multiple superradiant levels may be populated simultaneously (Fig. 5.4), and the binary may sweep through several resonances in succession before merger. The resonance history can be complex and depends on the level hierarchy and the chirp timescale [208]; multi-level corrections to the waveform template are beyond the scope of this paper.

5.6.4 Outlook and Future Directions

Despite the negative detectability conclusion for binary-induced bursts at current ADMX sensitivity, several physically well-motivated avenues remain open.

Isolated-GA annihilation as a near-term target. The isolated-GA annihilation signal ($h_{0,\mathrm{ann}}^{\mathrm{peak}} \sim 10^{-21}$ at 145 au, $f_{\mathrm{ann}} \sim$ MHz–GHz) is not subject to the event-rate limitations that render binary-driven bursts undetectable. For PBHs distributed throughout the Milky Way dark-matter halo, sources within ~ 200 au are plausibly present with an abundance proportional to the local PBH dark-matter fraction f_{PBH}. The continuous, monochromatic nature of the signal ($\tau_{\mathrm{ann}} \gg t_{\mathrm{Hubble}}$, $\Delta f \lesssim 10^{-33}$ Hz) makes it an ideal target for coherent integration strategies that accumulate signal-to-noise over extended observation periods. We identify this channel as the primary observational target within the current ADMX band, and provide the explicit E_1-based frequency-domain template

(Appendix A.1) needed to implement a matched-filter search.

Next-generation detector requirements for binary signals. For binary-induced bursts to become detectable, three improvements are needed in combination: (i) enhancement in strain sensitivity $h_{\rm min}$ (Eq. (5.54)) through larger effective volume, stronger magnetic field, reduced system temperature, or enhanced mode overlap; (ii) broader frequency coverage, particularly at sub-GHz frequencies where heavier and potentially more abundant PBHs contribute; and (iii) detector response and readout strategies optimized for transient and frequency-drifting sources, rather than only for steady-state monochromatic signals. Novel detector concepts—including plasma haloscopes, LC circuit detectors, broadband bulk acoustic resonators, multimode cavities, and correlated detector networks [9, 14, 38, 83, 194]—may provide complementary sensitivity in frequency ranges or signal classes not covered by current ADMX searches.

Constraints on the ultralight boson spectrum. Even without a detection, the non-observation of a persistent narrowband GW signal at a given frequency places an upper bound on the product $f_{\rm PBH} \cdot h_{0,\rm ann}^{\rm peak}$, which translates into a constraint on the PBH density in the local dark-matter halo as a function of boson mass. Deriving quantitative spin-down constraints analogous to those obtained from LIGO-band continuous-wave searches [170], but using ADMX non-detection data, is a direct application of the templates presented in this paper and represents a novel intersection of axion dark-matter searches with GW physics.

PBH spin distribution and formation constraints. The sensitivity of the GW strain to the initial BH spin (through $N_{\rm sat} \propto a_* - a_*^{\rm crit}$) makes PBH-GA systems a probe of PBH formation mechanisms, complementary to mass-function constraints from microlensing and CMB observations [52]. A statistical population analysis combining the GW signal distribution with the PBH spin distribution would allow simultaneous constraints on both the PBH abundance and their formation spin, and is facilitated by the analytic templates derived here.

In sum, in this paper we have systematically studied the MHz–GHz gravitational-wave signatures of superradiant boson clouds around primordial black holes. Our results establish PBH–gravitational atom systems as among the very few theoretically well-motivated sources of high-frequency gravitational waves, and provide the complete set of waveform templates needed for targeted searches at ADMX and successor experiments. The gap between current sensitivity and the predicted signal levels defines a concrete engineering target for next-generation high-frequency GW detectors, and the boson mass range $\mu \sim 10^{-7}$–$10^{-5}\,\mathrm{eV}$ probed by these signals remains largely unconstrained by any other observational technique.

While isolated systems can produce long-lived, narrow-band signals with potentially observable strain amplitudes under optimistic conditions, binary-induced signals are generically suppressed both in amplitude and event rate. Depending on binary parameters, some signals may even be transient compared to detector response times. This study highlights a clear experimental target: improved strain sensitivity, broader frequency coverage, and detector response/readout strategies optimized for transient and frequency-drifting signals are all essential to probe this class of sources. Finally, the analytic templates developed here provide concrete benchmarks for future high-frequency gravitational-wave searches.

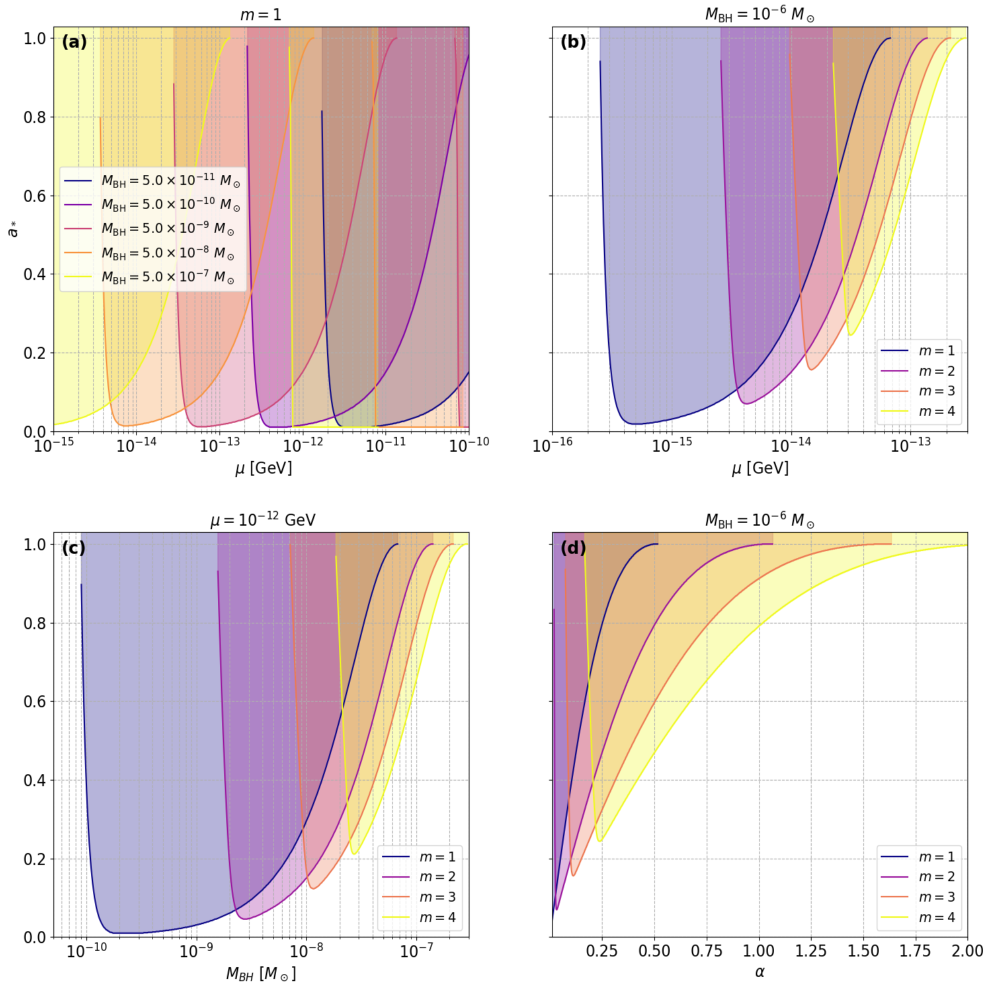

Figure 5.1: Regge trajectories showing the minimum black hole spin a_* required to sustain superradiant growth as a function of boson and black hole parameters, for azimuthal modes $m = 1, 2, 3, 4$. Shaded regions indicate parameter space where superradiance is active. **(a)** Superradiance threshold as a function of boson mass μ for fixed $m = 1$ and varying black hole mass M. **(b)** Threshold as a function of μ for fixed $M = 10^{-6}\ M_\odot$ and varying m. **(c)** Threshold as a function of $M_{\rm BH}$ for fixed boson mass $\mu = 10^{-3}$ eV and varying m. **(d)** Threshold as a function of the gravitational coupling $\alpha = M\mu$ for fixed $M = 10^{-6}\ M_\odot$ and varying m.

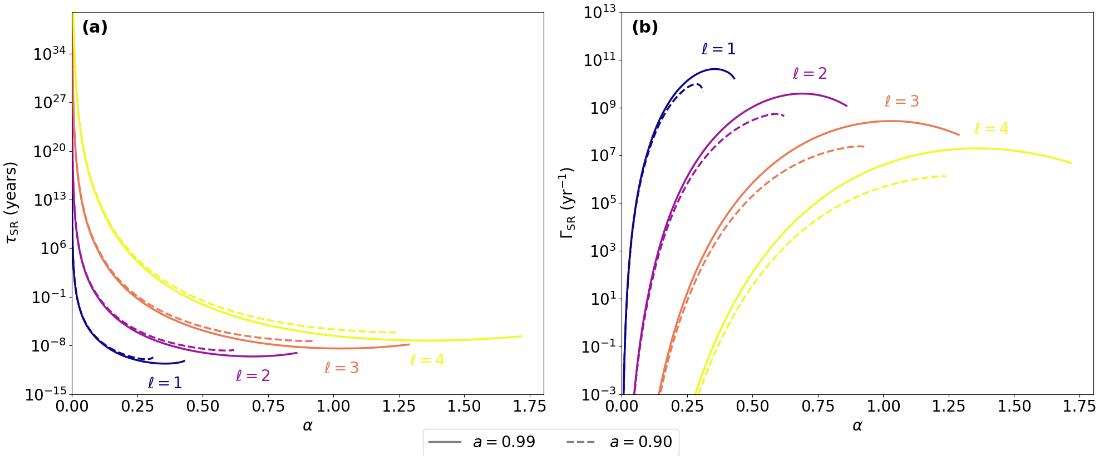

Figure 5.2: **(a)** Estimates of superradiance lifetime $\tau = \Gamma^{-1}$ as a function of α and azimuthal quantum number $\ell = m$ for a $10^{-6}\,M_\odot$ BH. Solid lines show $a = 0.99$, dashed lines show $a = 0.90$. **(b)** Superradiance rate $\Gamma_{\rm SR}$ as a function of α for the same parameter space. Superradiance continues until $m\Omega_{\rm H} \geq \omega$ is violated through spin extraction via the superradiant cloud.

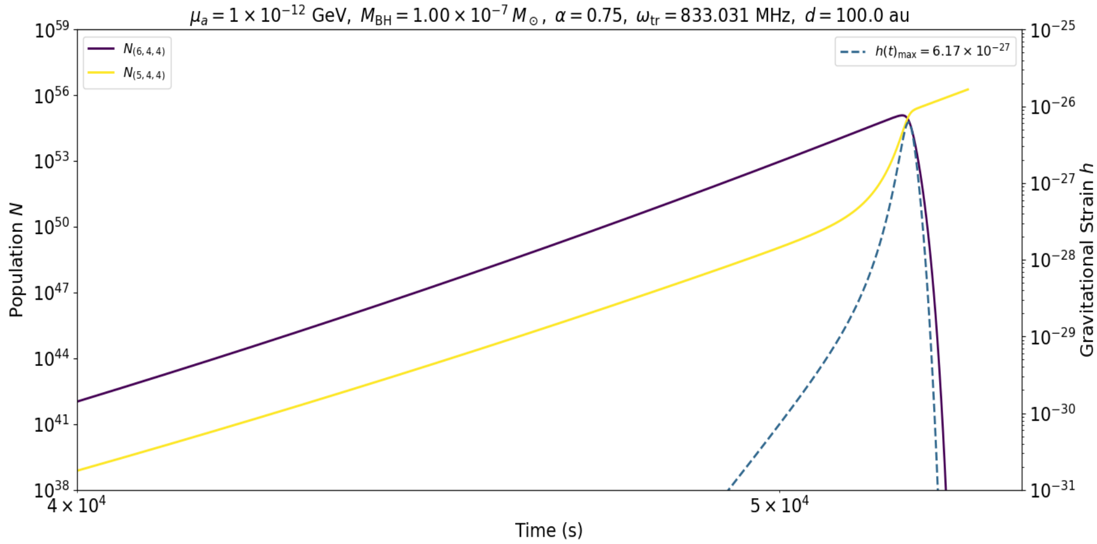

Figure 5.3: Time evolution of the occupation numbers $N_{\{6,4,4\}}$ and $N_{\{5,4,4\}}$ alongside the corresponding gravitational-wave strain envelope $h(t)$ (dashed), computed for $\mu_a = 10^{-3}$ eV, $M_{\rm BH} = 10^{-7}\,M_\odot$, $\alpha = 0.75$, and a source distance of $d = 100$ au. At early times both states grow independently under superradiance (Eqs. (5.17)–(5.18)). Once N_e is large enough, the gravitational self-interaction coupling $\Gamma_t N_g N_e$ (Eq. (5.21)) begins to dominate the excited-state dynamics, and a net flux of bosons cascades into the lower state, emitting a nearly monochromatic gravitational wave analogous to stimulated emission in atomic physics. The peak strain reflects a specific benchmark configuration and is not representative of the maximal transition strain, which can reach depends primarily on BH mass and source distance. This specific parameter configuration was chosen to illustrate a source within ADMX's detectable frequency band.

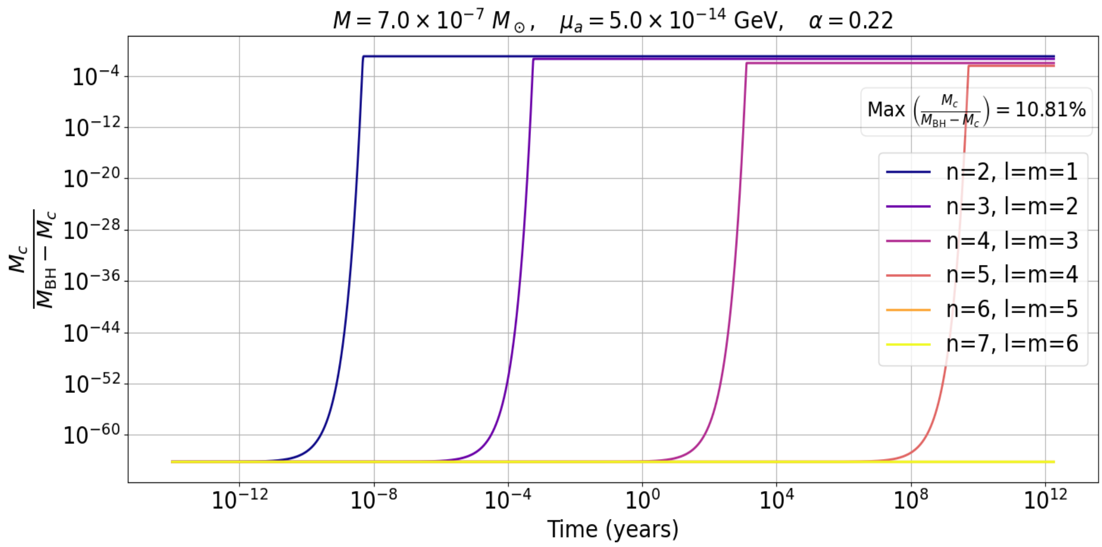

Figure 5.4: Fractional BH mass transferred to each superradiant level, $M_c/(M_{\rm BH} - M_c)$, as a function of time for states with $n = m + 1$, $\ell = m$, and $m = 1$–6, computed at $M_{\rm BH} = 6 \times 10^{-7}\, M_\odot$, $\mu_a = 5 \times 10^{-5}\,$eV, $\alpha = 0.22$. The $\{2, 1, 1\}$ level (darkest) reaches its saturation first and at the largest fractional mass. The maximum extracted mass for the most efficient state is 10.81%, in excellent agreement with the analytic upper bound of 10.8% derived in Ref. [208] for $\alpha \approx 0.2$ and $a_* \to 1$. Higher-m levels saturate later and accumulate correspondingly smaller clouds, as their slower growth is overtaken by the already-saturated lower levels.

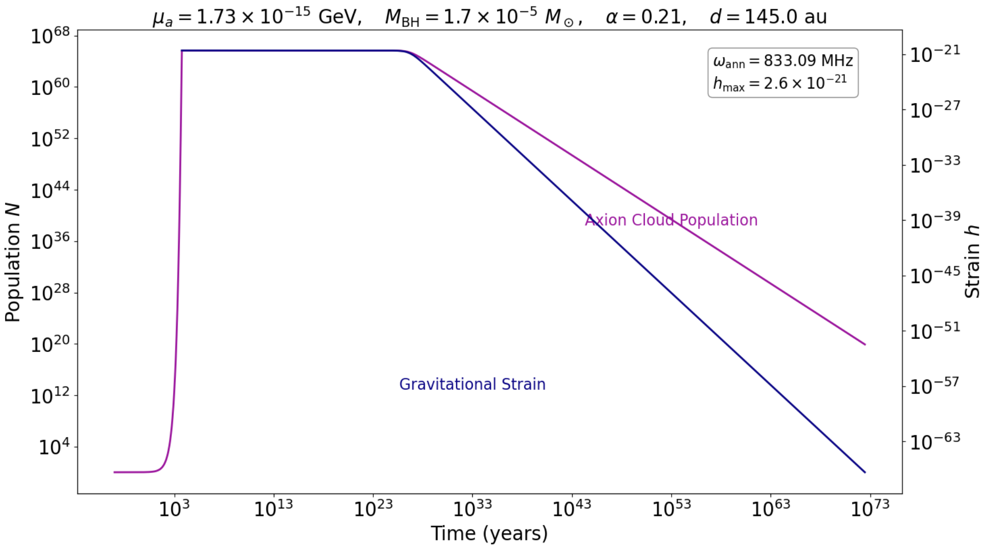

Figure 5.5: Time evolution of the cloud occupation number $N(t)$ (solid) and the corresponding annihilation strain envelope $h_{0,\rm ann}(t)$ (dashed), computed for $\mu_a \approx 2\times10^{-6}$ eV, $M_{\rm BH} = 1.7\times10^{-5}\, M_\odot$, $\alpha = 0.21$, and $d = 145$ au. The occupation number grows during the superradiant phase (left portion of the curve) until it saturates at $N_{\rm max}$, after which $N(t)$ decays as $[1 + \Gamma_a N_{\rm max}\, t]^{-1}$ (Eq. (5.31)) as boson pairs annihilate into gravitons. The corresponding strain peaks at $h_{0,\rm ann}^{\rm max} \approx 10^{-21}$ for nearby distances and decays on a timescale $\tau_{\rm ann} = (\Gamma_a N_{\rm max})^{-1}$ that can exceed the Hubble time, making the signal effectively monochromatic and persistent for any observation campaign. The annihilation frequency $\omega_{\rm ann}/(2\pi) \approx 833.09$ MHz was chosen to match the emitted signal from Fig. 5.3 and is within the ADMX scan band.

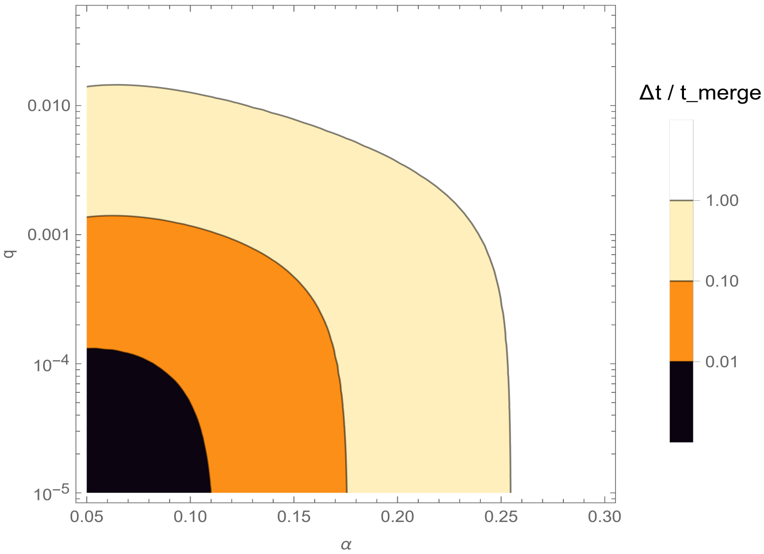

Figure 5.6: Ratio $\Delta t/t_{\mathrm{merge}}$ (Eqs. (5.44) and (5.47)) in the (α, q) parameter plane for the $\{2,1,1\} \to \{2,1,-1\}$ transition, evaluated at the resonant orbital frequency Ω_0. Contours correspond to $\Delta t/t_{\mathrm{merge}} = 0.01$ (inner), 0.1 (middle), and 1 (outer). The white region, where $\Delta t/t_{\mathrm{merge}} > 1$, is excluded: the signal duration would exceed the remaining merger time, violating the linear-orbit approximation, and the binary would not be intact. The allowed region is bounded by $\alpha \lesssim 0.26$ and $q \lesssim 0.01$; throughout Secs. 5.4.3 and 5.5 we adopt the benchmark values $\alpha = 0.14$ and $q = 10^{-4}$, which lie well within this allowed region.

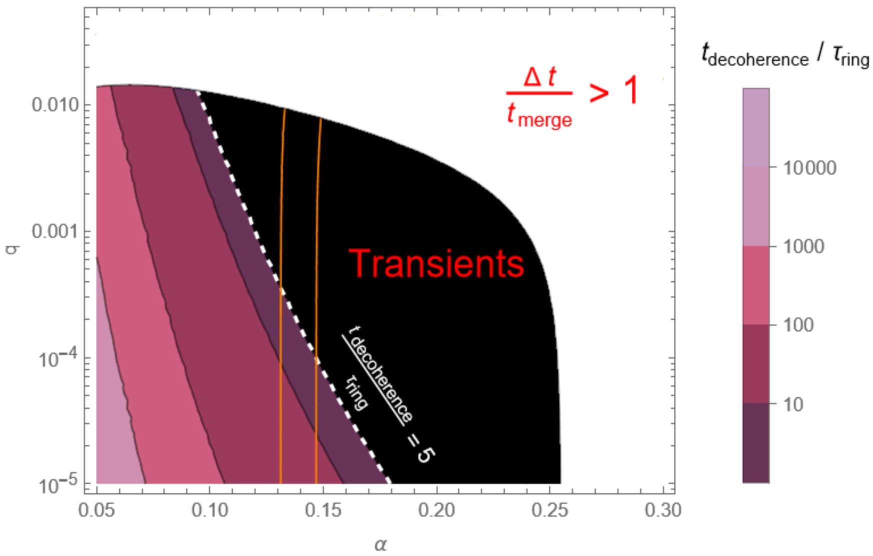

Figure 5.7: Ratio of $\Delta t_{\rm eff}/\tau_{\rm ring}$ for the conservative response diagnostic with $\Delta t_{\rm eff} \sim t_{\rm decoherence}$(Eq. (5.52)) within the allowed parameter space bounded by Fig. 5.6 for an ADMX-like cavity with $Q \sim 8 \times 10^4$, but without restricting the cavity resonance frequency to the ADMX hardware band contained between the orange solid curves. The white dash line corresponds to the boundary of $\Delta t_{\rm eff}/\tau_{\rm ring} = 5$, where binaries that produce smaller timescale ratios would correspond to transient signals incompatible with the steady-state approximation.

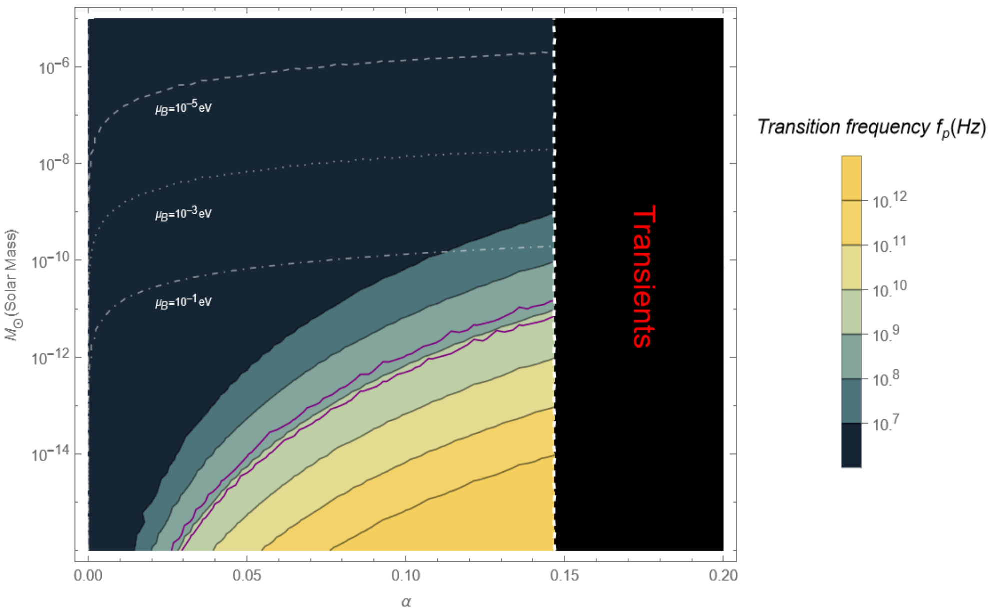

Figure 5.8: Peak transition frequency f_p (Eq. (5.45)) in the $M_{\rm BH}$–α parameter space for the $\{2,1,1\} \to \{2,1,-1\}$ hyperfine transition, for boson masses $\mu_b = 10^{-5}\,$eV (top), $10^{-3}\,$eV (middle), and $10^{-1}\,$eV (bottom), with $q = 10^{-4}$ fixed. The α value that defines the transient signal boundary shifts with different choices of q, as shown in Fig. 5.7. The PBH mass range cropped to $M_{\rm BH} \in [10^{-16}, 10^{-5}\,M_\odot]$ is shown, corresponding to the HF-band. The ADMX Run-1 scan band (0.645–1.4 GHz) is highlighted in purple; it is accessible for $M_{\rm BH} \sim 10^{-11}$–$10^{-12}\,M_\odot$ and $\mu_b \sim 2.5\,$eV at $\alpha \sim 0.14$. Sources lying in this band are the primary targets of the detectability analysis in Sec. 5.4.3.

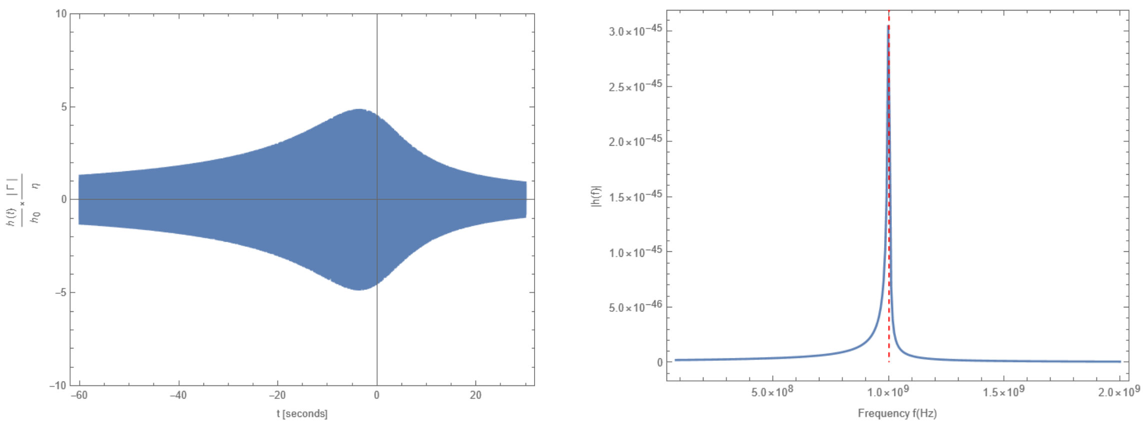

Figure 5.9: Time-domain strain $h_+(t)$ (left) and frequency-domain amplitude $|\tilde{h}(f)|$ (right) for the benchmark binary GA system of Eq. (5.53): $\alpha = 0.14$, $M = 6.68 \times 10^{-12}\, M_\odot$, $q = 10^{-4}$, $r = 10\,\text{kpc}$. In the time domain the signal is a short quasi-monochromatic burst with duration Δt (Eq. (5.44)). In the frequency domain the signal is a Lorentzian-shaped peak centered at the peak frequency $f_p = 1.001\,\text{GHz}$ (red dashed line), with spectral width $\sim |\Gamma|/\pi$. Both panels are computed using Eqs. (5.37) and (5.41) with the numerical inputs described in Sec. 5.4.3.

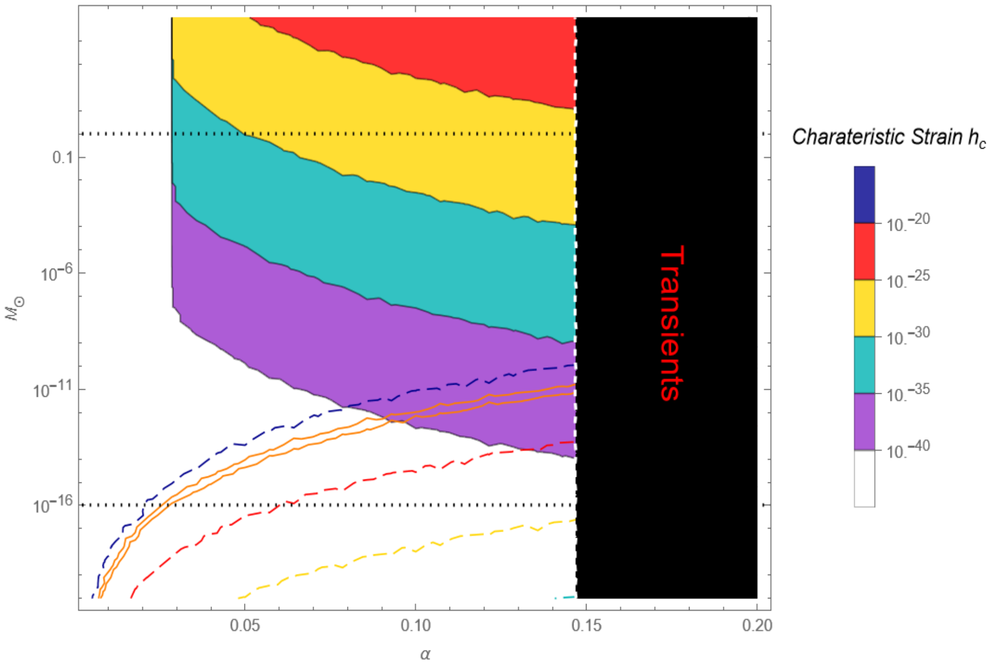

Figure 5.10: GW characteristic strain h_c (filled contours) and ADMX minimum detectable strain $h_{\rm min}$ (dashed contours, Eq. (5.54)) for binary GA systems in the $M_{\rm BH}$–α plane, with $q = 10^{-4}$ and $r = 10\,{\rm kpc}$. Filled contour levels span $h_c \in [10^{-40}, 10^{-20}]$; dashed contour levels show the same decades of $h_{\rm min}$, computed without restricting f to the ADMX hardware band. The orange solid curves overlay the ADMX frequency-band boundaries from Fig. 5.8, delimiting the region where ADMX can physically respond; within this region $h_{\rm min} \sim 10^{-22}$. The PBH mass range $10^{-16}\,M_\odot \lesssim M_{\rm PBH} \lesssim 1\,M_\odot$ is bounded by the black dotted lines. No h_c contour overlaps a $h_{\rm min}$ contour of equal magnitude within the PBH mass range at this distance, demonstrating that the signal falls orders of magnitude below the ADMX detection threshold.

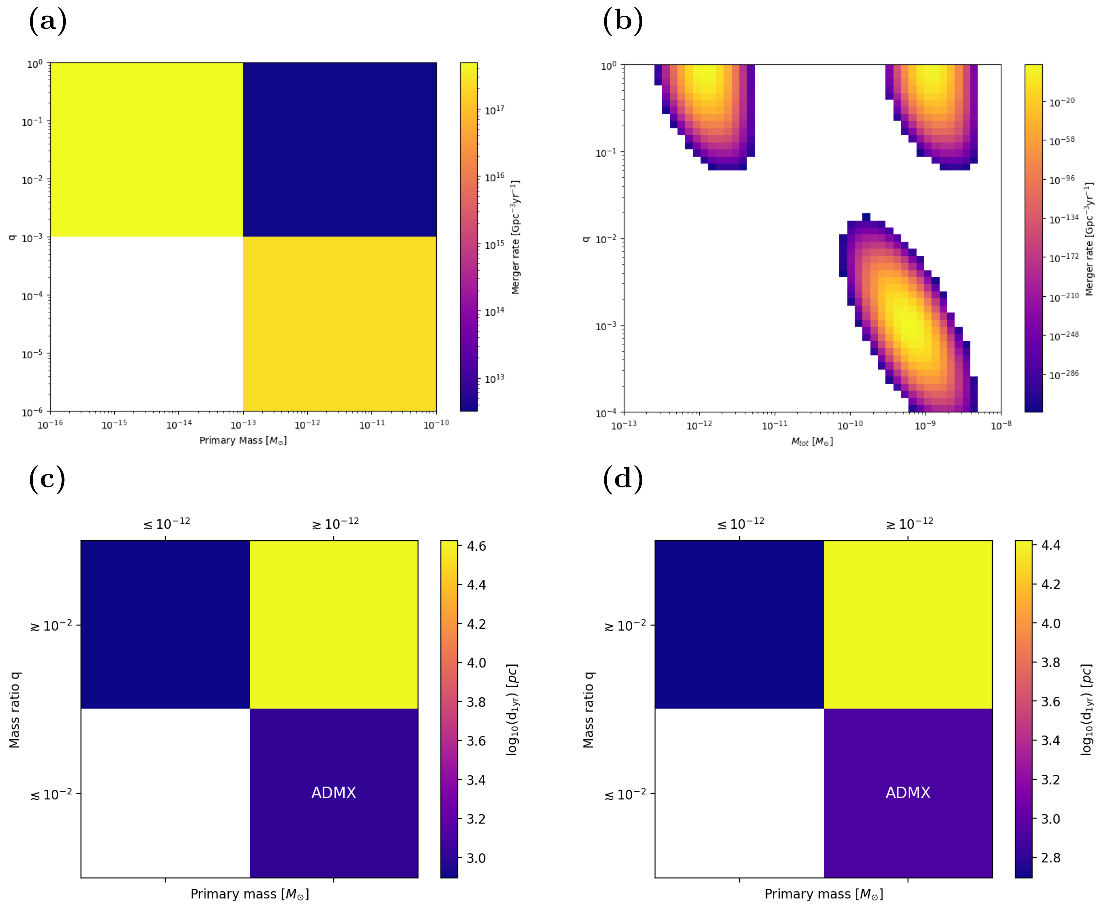

Figure 5.11: (a-b) Merger rates in $M_{tot} - q$ space for dichromatic (a) and double lognormal (b) PBH mass functions with $m_1 = 6.68 \times 10^{-12}\, M_\odot$, $m_2 = 6.68 \times 10^{-16}\, M_\odot$, and $f_1 = f_2 = 0.5$ for both functions and $\sigma = 0.05$ for the double lognormal. (c-d) Characteristic distances $d_{1\,\mathrm{yr}}$ (Eq. (5.62)) at which one merger per year is expected on average, for the dichromatic mass function (c) and for the double lognormal (d), subject to the binaries falling within a general range of benchmark parameters outlined in Sec .5.4. For the dichromatic case, the three binary types correspond to those shown in Table 5.3. The ADMX-relevant (m_1, m_2) binaries (labeled "ADMX") merge at appreciable rates only at $d \gtrsim 1\,\mathrm{kpc}$ for both mass functions, whereas ADMX requires $r \ll 1\,\mathrm{AU}$ to reach the detection threshold. The gap spans roughly fifteen orders of magnitude in strain.

Chapter 6

Conclusion

This thesis has used black holes as precision instruments for probing physics beyond the Standard Model, exploiting their unique property of interacting gravitationally with every species lighter than their Hawking temperature regardless of gauge charge. The four studies that comprise this work form a coherent progression: from broad model-independent constraints on the size of gravitationally coupled dark sectors, through the phenomenology of primordial black holes as a dark matter candidate, to a detailed re-examination of black hole evaporation in generic geometries, and culminating in a quantitative detectability analysis of high-frequency gravitational-wave transients from superradiant boson clouds. We summarize the principal results and their implications here before identifying the most important open questions.

Summary of Results

Chapter 2: Precision Gravity Constraints on Dark Sectors. The first study pursued two complementary routes to constraining gravitationally coupled dark sectors. On the laboratory side, we connected effective field theory predictions for loop-induced quantum corrections to Newton's constant with precision torsion-balance measurements, reinterpreting published fifth-force exclusion curves as constraints on the multiplicity N of hidden-sector fields. For truly massless dark species, this yields bounds as strong as $N \lesssim \mathcal{O}(10^{61})$. On the astrophysical side, the dominant constraint comes from the requirement that stellar-mass black holes observed by LIGO/Virgo survive to their inferred ages: using the merger event GW190924_021846, with chirp mass $5.7\,M_{\odot}$ and inferred age $(1.16 \pm 0.04) \times 10^{10}\,\mathrm{yr}$, we constrain $N \lesssim 5 \times 10^{58}$ for massless spin-0 dark sectors and $N \lesssim 2.5 \times 10^{60}$ for spin-2, using full greybody-factor calculations with BlackHawk. Complementary limits from ultrahigh-energy cosmic-ray collisions ($N \lesssim 10^{71}$), LHC monojet searches ($N \lesssim \text{few} \times 10^{62}$), and supernova cooling ($N \lesssim 6 \times 10^{68}$) confirm that black hole evaporation provides the strongest quasi-model-independent constraint on large gravitationally coupled dark sectors across a wide range of dark sect particle masses.

Chapter 3: Gravitational Waves from Asteroid-Mass PBH Mergers. The second study narrowed the focus to primordial black holes in the asteroid-mass window (10^{-16}–$10^{-11}\,M_{\odot}$), where they can constitute the entirety of the cosmological dark matter and are inaccessible to conventional interferometric detectors. We updated the gravitational-wave sight distance for narrow-band resonant microwave cavities such as ADMX, demonstrating analytically that it becomes chirp-mass independent at large PBH masses, saturating at $d_{\rm sens} \simeq 0.03\,\mathrm{AU}$ in the $N_c/Q \ll 1$ limit. Constructing the most optimistic possible merger event rate—incorporating early-time clustering with $\delta_{\rm dc} \sim 10^4$, local dark matter overdensities, dichromatic mass functions, and $f_{\rm PBH} = 1$—we found that detectable merger events are extremely rare even under maximally favorable assumptions. This sobering conclusion establishes a sensitivity floor for the resonant-cavity approach to PBH mergers and motivates the search for alternative high-frequency gravitational-wave sources from the same population.

Chapter 4: Dark Sector Modifications to Black Hole Evaporation Hierarchies. The third study generalized the evaporation framework from the Schwarzschild case to Kerr and Reissner–Nordström geometries, where mass, angular momentum, and charge are all dissipated simultaneously. The canonical Page hierarchy—charge neutralized first, angular momentum shed next, mass loss last—was shown to be non-universal. Using the Page factor formalism extended to charged and rotating metrics, we derived a critical threshold $N_{\rm dof} \gtrsim N_{\rm crit} \approx 220$ for massless dark species above which the mass-loss rate outpaces the charge-loss rate, producing effective charge *growth* during evaporation and driving the black hole asymptotically toward extremality. For $N_{\rm dof} \gtrsim 270$, charge grows throughout the entire evaporation history, independent of initial conditions. Two additional mechanisms were characterized: Schwinger pair production restores the conventional hierarchy for light ($M \lesssim 10^{17}\,\mathrm{g}$), highly charged black holes by introducing a competing QED discharge channel; and superradiance was shown to extract angular momentum on timescales far shorter than Hawking radiation, introducing a correlated decay track in the (a^*, Q_*) phase space not captured in prior treatments. Together, these results establish that large hidden sectors can produce near-extremal charged black hole relics as a population-level outcome of Hawking evaporation.

Chapter 5: High-Frequency Gravitational Wave Transients from Superradiance. The fourth and final study treated superradiance not merely as a spin-loss mechanism but as a gravitational-wave source in its own right, performing the first unified detectability analysis of superradiant signals from PBH gravitational atoms at GHz-band detectors. For isolated gravitational atoms, we derived analytic closed-form frequency-domain waveform templates via the exponential integral E_1 for both the level-transition and annihilation channels, finding peak strains of $h \sim 10^{-22}$ at $1\,\mathrm{kpc}$ for the annihilation channel and $h \sim 10^{-23}$ for level transitions. For PBH binaries, we applied the Landau–Zener formalism to the hyperfine transition $\{2, 1, 1\} \to \{2, 1, -1\}$ and found that while the ring-up criterion can be satisfied over a range of binary parameters, the characteristic strain falls orders of magnitude below the ADMX sensitivity threshold at astrophysically plausible distances. Using PBH merger rates from the early two-body formation channel, detectable binary-induced bursts would require source distances $\lesssim 1\,\mathrm{AU}$, whereas the merger rate implies characteristic event distances exceeding $\sim 9\,\mathrm{kpc}$. The isolated-GA annihilation signal—persistent, continuously emitting, and not subject to the same rate limitations—emerges as the more promising near-term observational target for ADMX-band experiments.

Overarching Themes and Lessons

Several cross-cutting conclusions emerge from the body of work as a whole.

Black holes are uniquely democratic probes. Unlike collider or direct-detection experiments, whose sensitivity depends on specific portal couplings, black holes couple to every species lighter than their Hawking temperature through gravity alone. This makes Hawking evaporation arguments the strongest *and* most model-independent constraints on large gravitationally coupled dark sectors, dominating over cosmic-ray, collider, and precision-gravity probes across most of the relevant parameter space.

The evaporation hierarchy is not universal. A recurring theme across Chapters 2 and 4 is that the standard Schwarzschild-based intuition about black hole evaporation breaks down in qualitatively new ways when the particle content of the theory is extended. The charge-first hierarchy can be reversed, the spin-down rate is dominated by superradiance rather than Hawking radiation in many regimes, and large hidden sectors can produce near-extremal relics as a population-level outcome. These are not fine-tuned exceptions—they are generic consequences of large N.

Sensitivity gaps motivate new detector technology. Both the PBH merger study and the superradiance study arrive at the same qualitative conclusion: the signals are theoretically well-motivated, but current ADMX sensitivity falls orders of magnitude short of what would be needed for detection at astrophysically plausible distances. This is not a reason for pessimism—it is a concrete design target. The analyses in Chapters 3 and 5 identify improved strain sensitivity, faster ring-up times, and broader frequency coverage at lower frequencies as the specific detector improvements that would open this class of signals to observation.

Future Directions

The natural extensions of this work span both theoretical and experimental directions.

On the theoretical side, the evaporation hierarchy analysis of Chapter 4 treated the Kerr and Reissner–Nordström cases largely independently. A fully self-consistent treatment of Kerr–Newman black holes—simultaneously rotating and charged—in the presence of a large dark sector remains to be carried out, and the coupled (a^*, Q_*) phase-space trajectories may reveal new phenomenology not captured by either limit alone. The superradiance analysis of Chapter 5 treated the isolated and binary channels independently; a self-consistent treatment that evolves both channels simultaneously, accounting for partial cloud depletion by binary resonances before isolated saturation is reached, would sharpen the detectability estimates considerably. Multi-level dynamics beyond the two-state Landau–Zener approximation—particularly in cases where multiple superradiant modes are populated simultaneously—also remain to be explored for the PBH mass range.

On the observational side, the LIGO/Virgo O3 constraints on dark sector size derived in Chapter 2 used a single benchmark event (GW190924_021846) as the most constraining known black hole. A systematic reanalysis of the full O3 and O4 merger catalogues, marginalizing over merger histories and mass uncertainties, could significantly tighten these bounds or reveal population-level signatures of dark sector particle emission. For the high-frequency gravitational-wave program, the analytic waveform templates derived in Chapter 5 provide concrete targets for matched-filter searches at existing and future ADMX runs; their implementation in a search pipeline is a natural next step.

More broadly, the program pursued in this thesis—using black holes of all masses, from stellar remnants to primordial relics, as gravitational spectrometers of the dark sector—remains in its early stages. The next decade's improvements in both gravitational-wave detector sensitivity and high-frequency cavity technology will determine how far down this road observation can take us.

Chapter 7

Appendix

7.1 Other merger pathways

In [191] the authors presented an update of their previous work [190] on early three-body mergers. Albeit an improvement in several aspects of the calculation, this latter discussion does not include a detailed discussion of the effect of early-time clustering. We verified that the differences in the merger rate compared to Ref. [190] are actually marginal and within a factor of a few, when $\delta_{\rm dc} = 1$. As such, we resort to the expression discussed above.

Other merger pathways include the early two-body (E2), late two-body (L2), and late three-body (L3) formation scenarios. For a generic mass function ψ, the E2 pathway merger rate for masses $m_1,\ m_2$ reads [191]

$$\frac{dR_{E2}}{dm_1 dm_2} \approx \frac{1.6 \times 10^6}{\text{Gpc}^3\ \text{yr}} f_{\rm PBH}^{\frac{53}{37}} \eta^{-\frac{34}{37}} \left(-\frac{32}{37}\right)^{-\frac{32}{37}} S_L S_E \frac{\psi(m_1)}{m_1} \frac{\psi(m_2)}{m_2} \tag{7.1}$$

where as usual $f_{\rm PBH} = \int \psi(m) dm$, $\eta = m_1 m_2/(m1 + m_2)^2$, $M = m_1 + m_2$, and S_L and

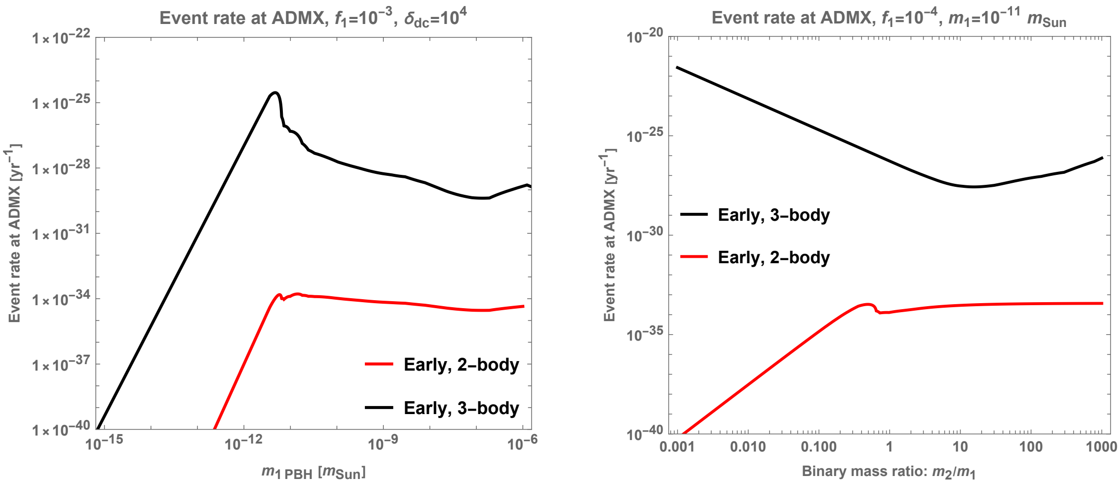

Figure 7.1: Pathways comparison – see the text for details.

S_E are suppression factors (which for the sake of comparison with the $E3$ rate we set to 1 below).

The late, 2-body pathway leads to a merger rate [191]

$$\frac{dR_{E2}}{dm_1 dm_2} \approx \frac{3.4 \times 10^{-6}}{\text{Gpc}^3\ \text{yr}} f_{\text{PBH}}^2 \delta_{\text{eff}} \left(\frac{\sigma_v}{\text{km/s}}\right)^{-\frac{11}{7}} \eta^{-\frac{5}{7}} \frac{\psi(m_1)}{m_1} \frac{\psi(m_2)}{m_2}, \tag{7.2}$$

where δ_{eff} is a local DM density contrast factor, and we assume a local velocity dispersion $\sigma_v \sim 100$ km/s.

Finally, the 3-body late mergers reads [191]

$$\frac{dR_{L3}}{dm_1 dm_2} \approx \frac{1.3 \times 10^{-16} e^{-6.0(\gamma-1)}}{\text{Gpc}^3\,\text{yr}} f_{\text{PBH}}^3 \delta_{\text{eff}}^2 \left(\frac{\sigma_v}{\text{km/s}}\right)^{-9+\frac{8\gamma}{7}} \eta^{-1+\frac{\gamma}{7}} \tilde{\mathcal{F}} \left(\frac{\langle m \rangle}{2\eta M} \kappa_{\min}\right) \frac{\psi(m_1)}{m_1} \frac{\psi(m_2)}{m_2},$$

$$\frac{dP}{dj} \sim \gamma j^{\gamma-1},$$

is assumed to be 1.1 (we find that the induced uncertainty from varying γ is mild).

We compare the four pathways first, for simplicity, using a monochromatic mass function in the left panel of Fig. 7.1. We assume a density contrast $\delta_{\text{dc}} = 10^4$ for all pathways (we implement it as described at the end of Ref. [191] for the pathways E2, L2 and L3) and we use for each mass the maximal possible value of $f_{\text{PBH}}(m_{\text{PBH}}$ for a given mass. The figure shows that the E3 pathway dominates everywhere between 6 and 7 orders of magnitude to the next most significant pathway. The late 2- and 3-body pathways are several orders of magnitude below the range shown in the figure.

The right panel of Fig. 7.1 shows the results for a dichromatic mass function, with the mass function mass ratio m_2/m_1 shown in the x axis. Here again it is fully apparent how the E3 pathway dominates other pathways by several orders of magnitude, justifying in full the approach we take in this analysis. Again here, the late 2- and 3-body pathways are several orders of magnitude below the range shown in the figure.

7.2 Clustering Effects on Black Hole Binary Formation Pathways

PBH clustering plays a dual and complex role in influencing binary formation pathways, with distinct effects on early and late binary formation channels. Clustering suppresses early 3-body binary formation in a radiation-dominated Universe. This suppression occurs due to perturbations induced by nearby PBHs, tidal forces, angular momentum exchanges, and hierarchical dynamics, which destabilize nascent binaries and reduce merger rates. Studies consistently show that clustering significantly lowers early binary merger rates, particularly at high PBH dark matter fractions, with nearest-neighbor distances playing a critical role in suppressing binaries [191, 95, 68, 137, 27].

However, a subset of early binaries may survive in regions of low clustering density or when initial conditions are favorable to stable configurations [68]. Studies find that, generally, such effects amount to a rate suppression by at most two-three orders of magnitude [191, 95, 68, 137, 27].

In contrast, PBH clustering enhances late binary formation, particularly through 2-body dynamical capture within dense environments such as dark matter halos or small PBH clusters. These late pathways arise from increased interaction rates, clustering-driven density amplification, and gravitational harmonics that promote orbital binding. Studies highlight the role of nested overdensities and evolving structure formation in facilitating such interactions and highlight the resulting high-eccentricity late binaries detectable through gravitational wave signals [191, 202, 68, 70, 12]. While this enhancement boosts late-time merger rates, competing effects like cluster heating and evaporation in dense PBH clusters could partially inhibit late binary formation in some scenarios [145]. In the present case, late-time merger rates are suppressed by over 30 orders of magnitude, so we do not expect the hierarchy of merger pathways we assume here to change.

Clustering's influence evolves over cosmic time, transitioning from suppression of early (post-recombination) binaries to amplification of late-stage binary formation during structure formation epochs dominated by dark matter halos. Redshift-dependent analyses indicate that suppression mechanisms are most prominent at high redshift, while clustering-enhanced late pathways dominate in lower redshift environments [191, 202, 68, 12]. This dual role of clustering raises important questions about the exact thresholds and transitions between suppression and enhancement mechanisms, as well as the downstream effects of disrupted early binaries on late-time formation pathways [191, 202, 68, 70].

7.3 Microlensing and Gravitational Wave Constraints

PBHs within the mass range of Earth mass to $10^{-16}M_{\odot}$ present unique challenges and opportunities for observational detection, primarily through microlensing and GWs methods. Microlensing surveys, such as OGLE and Subaru/HSC, have robustly constrained Earth-mass PBHs, with strong restrictions imposed on their dark matter fraction. However, for smaller masses (below $10^{-12}M_{\odot}$), finite source size effects significantly reduce microlensing sensitivity. Femtolensing of gamma-ray bursts can, in principle, probe PBHs as small as $10^{-16}M_{\odot}$, but this method's efficacy is limited by the extended size of gamma-ray burst sources, calling into question earlier constraints [172, 140, 181, 180]. Additionally, clustered PBHs marginally impact microlensing event rates, as compactness requirements for clustering are often unmet, rendering practical deviations insignificant in most models [209, 180].

Continuous gravitational wave (GW) detection methods show promise for identifying binaries of planetary- and asteroid-mass PBHs, especially for masses in the $10^{-12}M_{\odot}$ to $10^{-5}M_{\odot}$ range. Studies using data from LIGO/Virgo's O3a observing run, such as those by Miller et al., have constrained PBH binaries through continuous wave (CW) searches, targeting signals from tightly bound compact binaries. While transient GW signals from mergers dominate for larger masses, continuous and transient continuous-wave (tCW) methods are more relevant for PBHs in this lower mass regime. Nevertheless, sub-$10^{-12}M_{\odot}$ PBHs remain largely undetectable with current GW observatories due to sensitivity limitations and the lack of strain signatures in the accessible frequency range. Future high-frequency GW detectors like ADMX and space-based observatories like LISA may be able to probe these smaller masses in clustered or binary configurations [166, 165, 84].

PBH clustering and binary formation are central to enhancing GW detection prospects, as clustering increases the likelihood of binary interactions and GW emission. While clustering may not substantially affect microlensing constraints, it has been proposed as a mechanism to boost GW merger rates, particularly for asteroid-mass PBHs forming in dense environments, as noted and explored above. Despite these theoretical advancements, the observational integration of microlensing and GW detection methods remains underdeveloped. Overlapping sensitivity windows (e.g., microlensing at Earth mass and GW constraints at planetary mass) could jointly refine PBH abundance and clustering properties. However, the challenges introduced by finite source size effects, computational limitations in GW template searches, and uncertainties in clustered PBH models hinder the realization of such synergies [209, 166, 165, 172].

In summary, microlensing and GW methods offer complementary, but not entirely overlapping, approaches to constraining PBHs across this mass range. While microlensing excels at Earth-mass PBH detection, its sensitivity diminishes at the lower end of the spectrum. In contrast, GW techniques are increasingly effective for clustered and binary systems in the planetary-mass regime but remain largely impractical for detecting individual small PBHs below $10^{-12} M_{\odot}$. To respond to these challenges, future research must refine clustering models, consider next-generation GW facilities, and develop joint observational frameworks to bridge the gaps in PBH detectability [165, 140, 180].

7.4 Frequency Domain Strain Templates

The detectability of a gravitational-wave signal is most cleanly expressed in terms of the characteristic strain, which is defined from the Fourier transform of the metric perturbation as $h_c(f) = 2f|\tilde{h}(f)|$. While the raw Fourier amplitude $\tilde{h}(f)$ depends on the details of the observation, such as the total duration of the time series and the specific sampling, the combination $h_c(f)$ isolates the physically relevant quantity: the accumulated signal power at a given frequency, including the contribution from the number of coherent wave cycles. For a quasi-monochromatic source observed over an observation time T, one has approximately $h_c(f_0) \sim h_0\sqrt{f_0 T}$, so that the characteristic strain increases with the square root of the effective number of cycles. This makes $h_c(f)$ the appropriate quantity for comparing theoretical signals to the sensitivity of gravitational-wave detectors.

A detector is characterized by its one-sided noise power spectral density $S_n(f)$. The corresponding noise amplitude is given by $h_n(f) = \sqrt{f S_n(f)}$. Signal power and detector noise are compared at the same level: both $h_c(f)$ and $h_n(f)$ have the same dimensions and represent amplitudes per logarithmic frequency interval. A source is detectable when its characteristic strain exceeds the noise curve, and the expected signal-to-noise ratio is

$$SNR^2 = \int_0^\infty \frac{h_c^2(f)}{f^2 S_n(f)}\, df = \int_0^\infty \frac{h_c^2(f)}{h_n^2(f)}\, d(\ln f). \tag{7.3}$$

This expression shows explicitly why $h_c(f)$ is the correct quantity to compare with detector sensitivity. Any direct comparison using $\tilde{h}(f)$ alone would be misleading, because $\tilde{h}(f)$ scales with the chosen Fourier window and does not represent a physical amplitude.

In the context of black hole superradiance and axion annihilation or level transitions, the characteristic strain provides a unified language for evaluating detectability against detector strain curves. The signals are intrinsically narrow-band and can be long-lived, so $h_c(f)$ reflects both the instantaneous gravitational-wave amplitude and the coherence of the emission process. Even if the raw strain is extremely small, the large number of coherent cycles at the transition or annihilation frequency can raise $h_c(f)$ above detector noise. Once expressed in terms of $h_c(f)$, the comparison to detector sensitivity curves is immediate: gravitational-wave observatories are sensitive to signals whenever $h_c(f)$ lies above (or not far below) $h_n(f)$ at the relevant frequency. This is the fundamental criterion for detection prospects.

7.4.1 Annihilation Strain

The bulk of this paper has shown that there are multiple potentially detectable signals in a gravitational atom system. These signals can arise from both an isolated gravitational atom or due to the presence of a binary companion. In the first case, we have shown that detectable signals arise from two separate phenomenon: level transitions and boson annihilations.

The first frequency domain strain that we consider arises due to the axion annihilation within a specified level. Again, the time domain strain is specified in Eq (5.32) and the only time dependence arises from the decay of the state population in Eq. (5.31):

$$h_{0,\mathrm{ann}}(t) = \frac{N_{\max}}{1+\Gamma_a N_{\max} t}\sqrt{\frac{4G_N}{r^2\omega_{\mathrm{ann}}}\Gamma_a} \tag{7.4}$$

where again we have made the assumptions that α and the superradiance rates are unchanging in time to obtain an approximate form of the annihilation strain amplitude. Again, this allows us to take a fourier transform to obtain an analytic form of the annihilation strain as a function of frequency. We begin with the one-sided Fourier transform,

$$\tilde{h}_{0,\mathrm{ann}}(\omega) = \frac{1}{\sqrt{2\pi}}\int_0^\infty \frac{A\,e^{-i\omega t}}{1+\beta t}\,dt, \qquad \beta > 0, \tag{7.5}$$

where we will consistently use the convention of $f = \frac{\omega}{2\pi}$.

$$A = N_{\max}\sqrt{\frac{4G_N\,\Gamma_a}{r^2\omega_{\mathrm{ann}}}}, \qquad \beta = \Gamma_a N_{max}. \tag{7.6}$$

Let $u = 1+\beta t$, so that $t = (u-1)/\beta$ and $dt = du/\beta$. Substituting, we find

$$\tilde{h}_{0,\mathrm{ann}} = \frac{A}{\sqrt{2\pi}\,\beta}e^{i\omega/\beta}\int_1^\infty \frac{e^{-i(\omega/\beta)u}}{u}\,du. \tag{7.7}$$

To evaluate the integral, we introduce a causal regulator $+i\epsilon$ to ensure convergence and define the exponential integral $E_1(x)$ in its equivalent forms:

$$E_1(x) \equiv \int_1^\infty \frac{e^{-tx}}{t}\, dt = \int_x^\infty \frac{e^{-u}}{u}\, du, \qquad (|\arg x| < \pi). \tag{7.8}$$

Comparing with Eq. (7.7), we immediately identify

$$x = i\frac{\omega}{\beta} + i\epsilon,$$

where the infinitesimal imaginary term $+i\epsilon$ enforces the causal $(+i\epsilon)$ prescription. Thus,

$$\int_1^\infty \frac{e^{-i(\omega/\beta)u}}{u}\, du = E_1\left(i\frac{\omega}{\beta} + i\epsilon\right). \tag{7.9}$$

Substituting Eq. (7.9) into Eq. (7.7), we obtain the compact result

$$\tilde{h}_{0,\mathrm{ann}} = \frac{A}{\sqrt{2\pi}\,\beta} e^{\,i\omega/\beta} E_1\left(i\frac{\omega}{\beta} + i\epsilon\right), \tag{7.10}$$

where $i\epsilon$ denotes the same causal regulator used in the definition of E_1. To rewrite Eq. (7.10) in terms of the exponential integral Ei and more familiar functions, we use the relation

$$E_1(z) = -\mathrm{Ei}(-z), \tag{7.11}$$

and continue it to complex arguments along the causal branch. For imaginary arguments, the continuation gives

$$\mathrm{Ei}(ix + i\epsilon) = \mathrm{PV}\big[\mathrm{Ei}(ix)\big] + i\pi\mathrm{Sign}(x), \tag{7.12}$$

where $\mathrm{PV}\big[\mathrm{Ei}(ix)\big]$ refers to the principal value of the exponential integral and the second term represents the usual $i\pi$ phase picked up when crossing the branch. Similarly, the

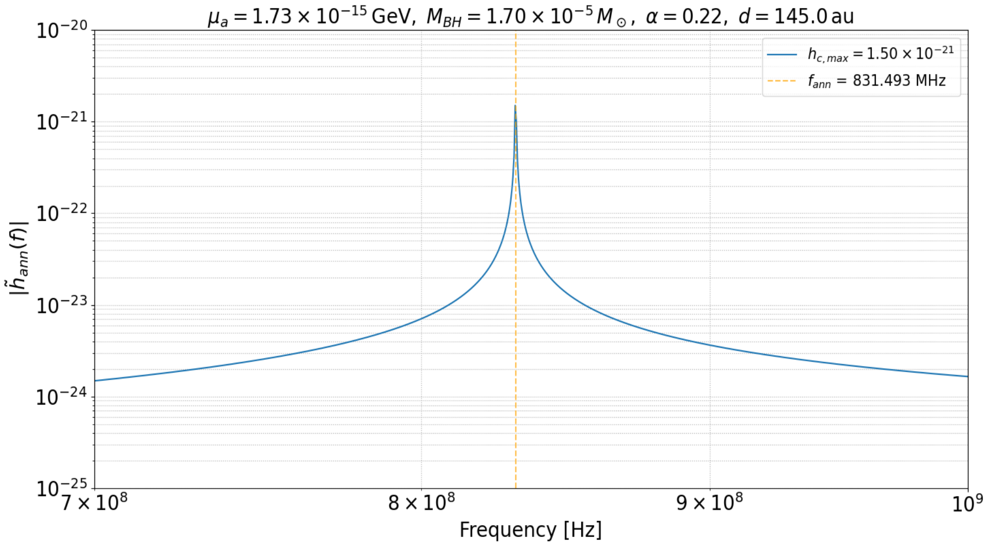

Figure 7.2: Characteristic strain of the Fourier transform for the annihilation strain. Input parameters match those from Fig. 5.5 for the $n = 2, l = m = 1$ state to obtain the strongest annihilation strain amplitude. The peak strain matches those expected in Ref [143].

logarithm acquires the standard phase shift across its branch cut,

$$\log(i\omega) - \log(-i\omega) = i\pi \mathrm{Sign}(\omega), \tag{7.13}$$

corresponding to the same causal choice of branch.

Using Equations (7.11)-(7.13), we can rewrite equation (7.10) as:

$$\tilde{h}_{0,\mathrm{ann}} = \frac{A}{\sqrt{2\pi}\,\beta}\, e^{i\omega/\beta} \left[-\,\mathrm{Ei}\left(-\frac{i\omega}{\beta} \right) - i\pi\, \mathrm{sgn}(\omega) \right]. \tag{7.14}$$

Finally, restoring A and $\beta = \Gamma_a N_{\max}$ gives

$$\tilde{h}_{0,\mathrm{ann}}(\omega) = \frac{1}{\sqrt{2\pi}} \sqrt{\frac{4G_N}{r^2 \omega_{\mathrm{ann}} \Gamma_a}}\, e^{i\omega/(\Gamma_a N_{\max})} \left[-\,\mathrm{Ei}\left(-\frac{i\omega}{\Gamma_a N_{\max}} \right) - i\pi\, \mathrm{sgn}(\omega) \right]. \tag{7.15}$$

which is identical to (7.14) upon using (7.13) and restoring all constants.

In order to return the oscillatory term to our frequency domain template, we need to multiply the envelope by $e^{\pm i\omega_{\rm ann} t}$. This has the effect of shifting our frequency, so we can describe the Fourier transform of our envelope as:

$$\mathcal{F}\Big[h_0(t)\, e^{\pm i\omega_{\rm ann} t}\Big] = \tilde{h}(\omega \mp \omega_{\rm ann}) \tag{7.16}$$

With $\cos = \frac{1}{2}(e^{i\phi_0} e^{i\omega_{\rm ann} t} + e^{-i\phi_0} e^{-i\omega_{\rm ann} t})$, $\sin = \frac{1}{2i}(e^{i\phi_0} e^{i\omega_{\rm ann} t} - e^{-i\phi_0} e^{-i\omega_{\rm ann} t})$, the polarization spectra are

$$\tilde{h}_+(\omega) = \frac{1+\cos^2(\iota)}{4}\Big[e^{i\phi_0}\, \tilde{h}_(\omega - \omega_{\rm ann}) + e^{-i\phi_0}\, \tilde{h}(\omega + \omega_{\rm ann})\Big], \tag{7.17}$$

$$\tilde{h}_\times(\omega) = \frac{\cos(\iota)}{2i}\Big[e^{i\phi_0}\, \tilde{h}(\omega - \omega_{\rm ann}) - e^{-i\phi_0}\, \tilde{h}(\omega + \omega_{\rm ann})\Big]. \tag{7.18}$$

Substituting $\tilde{h}$ gives the explicit closed forms:

$$\begin{aligned}\tilde{h}_+(\omega) = {} & \frac{N_{\rm max}}{\sqrt{2\pi}}\sqrt{\frac{4G_N}{r^2\omega_{\rm ann}}}\frac{1+\cos^2(\iota)}{4}\Big[e^{i\phi_0}\, e^{i(\omega-\omega_{\rm ann})/\beta}\Big(-2\,\mathrm{Ei}(\tfrac{i(\omega-\omega_{\rm ann})}{\beta}) + i\pi\,\mathrm{Sign}(\omega-\omega_{\rm ann})\Big) \\ & + e^{-i\phi_0}\, e^{i(\omega+\omega_{\rm ann})/\beta}\Big(-2\,\mathrm{Ei}(\tfrac{i(\omega+\omega_{\rm ann})}{\beta}) + i\pi\,\mathrm{Sign}(\omega+\omega_{\rm ann})\Big)\Big]\,,\end{aligned} \tag{7.19}$$

$$\begin{aligned}\tilde{h}_\times(\omega) = {} & \frac{N_{\rm max}}{\sqrt{2\pi}}\sqrt{\frac{4G_N}{r^2\omega_{\rm ann}}}\frac{\cos(\iota)}{2i}\Big[e^{i\phi_0}\, e^{i(\omega-\omega_{\rm ann})/\beta}\Big(-2\,\mathrm{Ei}(\tfrac{i(\omega-\omega_{\rm ann})}{\beta}) + i\pi\,\mathrm{Sign}(\omega-\omega_{\rm ann})\Big) \\ & - e^{-i\phi_0}\, e^{i(\omega+\omega_{\rm ann})/\beta}\Big(-2\,\mathrm{Ei}(\tfrac{i(\omega+\omega_{\rm ann})}{\beta}) + i\pi\,\mathrm{Sign}(\omega+\omega_{\rm ann})\Big)\Big]\,.\end{aligned} \tag{7.20}$$

7.4.2 Isolated Level Transition Strain

We next develop a frequency domain template of the transition amplitude gravitational strain outlined in Eq. (5.23) by taking a Fourier transform of the approximate solution. Since both the states grow exponentially and approximately independently for the majority of the lifetime of the excited state, we will be approximating the form of both the ground and excited states as $\frac{dN_g}{dt} = \Gamma_g^{sr} N_g + \Gamma_t N_g N_e$ and $\frac{dN_e}{dt} = \Gamma_e^{sr} N_e - \Gamma_t N_g N_e$. The resulting time dependence of each state is then an exponential of the form $N_g(t) = N_0 \exp\Big\{(\Gamma_g^{SR} + \Gamma_t N_e)t\Big\}$ and $N_e(t) = N_0 \exp\Big\{(\Gamma_e^{SR} - \Gamma_t N_g)t\Big\}$. The transition amplitude as a function of time, assuming a constant in time superradiant transition rate Γ_t, gives us a transition strain of the form

$$h_{0,tr} \approx \sqrt{\frac{4G_N}{r^2 \omega_{tr}} \Gamma_t N_0^g N_0^e} \exp\Big\{[\frac{1}{2}(\Gamma_e^{sr} + \Gamma_g^{sr} + \Gamma_t N_e - \Gamma_t N_g)t]\Big\} \tag{7.21}$$

Under these approximations we can take a Fourier transform to find the functional form of our strain. Doing so, we obtain the following expression:

$$h_{0,tr}(\omega) = \frac{1}{2\pi} \frac{\sqrt{\frac{4G_N}{r^2 \omega_{\rm tr}} \Gamma_t N_0^g N_0^e}}{\left(\Gamma_e^{sr} + \Gamma_g^{sr} + \Gamma_t N_g - \Gamma_t N_e - i\omega\right)} \left(-1 + [\frac{1}{2}(\Gamma_e^{sr} + \Gamma_g^{sr} + \Gamma_t N_e - \Gamma_t N_g - i\omega)T]\right) \tag{7.22}$$

with T being the time taken from the start of superradiance to when the excited state has transitioned completely. This is the approximate form of our gravitational wave signal in the frequency domain due to level transitions within the gravitational atom. Since detectors look at the frequency decomposition of these signals, this is ultimately the desired form for the gravitational strain. Again, to include the physical oscillations of the gravitational wave and its orientation, we can treat the transition signal as an amplitude (or envelope) modulated by a carrier frequency corresponding to the transition energy. In other words, the total strain is just our smooth envelope multiplied by an oscillatory factor at $\omega_{\rm tr}$:

$$h_{\rm tr}(t) = h_{0,\rm tr}(t)\cos(\omega_{\rm tr}t + \phi_0) = \tfrac{1}{2}\, h_{0,\rm tr}(t)\left(e^{i(\omega_{\rm tr}t+\phi_0)} + e^{-i(\omega_{\rm tr}t+\phi_0)}\right).$$

When we Fourier transform this, the oscillatory pieces simply shift the spectrum to $\pm\omega_{\rm tr}$, giving

$$\mathcal{F}\Big[h_{\rm tr}^{\rm env}(t)e^{\pm i\omega_{\rm tr}t}\Big] = \tilde{h}_{\rm tr}^{\rm env}(\omega \mp \omega_{\rm tr}). \tag{7.23}$$

So the frequency-domain strain becomes two terms centered on the transition frequency, each weighted by the phase ϕ_0:

$$\tilde{h}_{\rm tr}(\omega) = \frac{1}{2}\left[e^{i\phi_0}\,\tilde{h}_{0,\rm tr}(\omega-\omega_{\rm tr}) + e^{-i\phi_0}\,\tilde{h}_{0,\rm tr}(\omega+\omega_{\rm tr})\right]. \tag{7.24}$$

Finally, we can include the detector-frame polarization dependence. The observed strain depends on the inclination angle ι of the system relative to the line of sight. The two standard polarizations then take the form

$$\tilde{h}_{\rm tr}^{(+)}(\omega) = \frac{1+\cos^2\iota}{4}\Big[e^{i\phi_0}\,\tilde{h}_{\rm tr}^{\rm env}(\omega-\omega_{\rm tr}) + e^{-i\phi_0}\,\tilde{h}_{\rm tr}^{\rm env}(\omega+\omega_{\rm tr})\Big], \tag{7.25}$$

$$\tilde{h}_{\rm tr}^{(\times)}(\omega) = \frac{\cos\iota}{2i}\Big[e^{i\phi_0}\,\tilde{h}_{\rm tr}^{\rm env}(\omega-\omega_{\rm tr}) - e^{-i\phi_0}\,\tilde{h}_{\rm tr}^{\rm env}(\omega+\omega_{\rm tr})\Big]. \tag{7.26}$$

These two polarizations represent the characteristic “plus” and “cross” patterns of the transition signal, differing only by a relative phase and the inclination weighting. In practice, both appear as narrow peaks in the frequency spectrum at $\omega_{\rm tr}$, modulated by the growth and decay encoded in $\tilde{h}_{\rm tr}^{\rm env}(\omega)$.

The spectral width $\Delta\omega$ of the template is controlled entirely by the characteristic timescale over which the envelope $A(t)$ varies (i.e. by the relevant superradiant and transition rates),

$$\Delta\omega \sim \Gamma_{\rm tot} \sim \mathcal{O}\left(\mathrm{yr}^{-1}\right) \;\; \Rightarrow \;\; \Delta f \sim \frac{\Delta\omega}{2\pi} \ll 1\ \mathrm{Hz}, \tag{7.27}$$

even though the line is centered at $f_{\rm tr} = \omega_{\rm tr}/2\pi \sim$ MHz. Consequently, the frequency-domain signal appears as an extremely narrow (sub-Hz) Lorentzian line located at the MHz transition frequency. In general, the carrier frequency fixes the centering of the line position, while the slow population evolution ($\mathcal{O}(10^3)$ years) sets the linewidth. Faster population evolutions will therefore result in Lorentzian signals with much wider frequency resolution, whereas slowly evolving states will appear as increasingly narrow, near-monochromatic lines in the frequency spectrum.

7.4.3 Binary Level Transition Strain

The frequency domain template for the "plus" polarization for binary level transition strain given within Eq. (5.37) is already defined by Eq. (5.41), where a detailed derivation using the stationary phase approximation can be found within [147]. The frequency domain template for the "cross" polarization can also be easily obtained using the same method: Rewriting Eq. (5.38) into the form of:

$$h_{\times,211}(t) = h_0 \cos\iota \left[\frac{e^{-2i\Delta m\varphi(t_{\rm re})}Q(t_{\rm re}) - e^{2i\Delta m\varphi(t_{\rm re})}Q^*(t_{\rm re})}{2i}\right] \tag{7.28}$$

by expanding the imaginary component, then plugging into the Fourier transform:

$$\tilde{h}_{+,\times}(\omega) = \int dt\, h_{+,\times}(t)\, e^{i\omega t}. \tag{7.29}$$

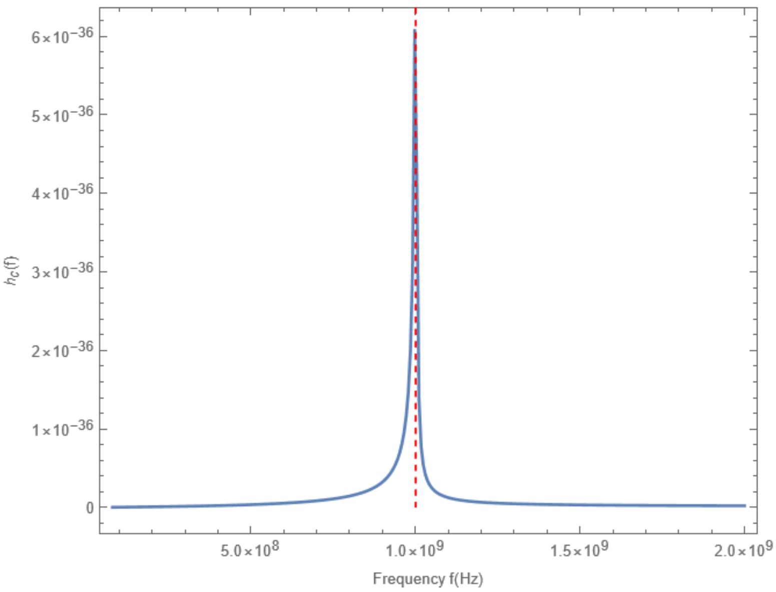

Figure 7.3: Computation of the binary perturbation characteristic strain $h_c(f) = 2f|\tilde{h}(f)|$ with $\tilde{h}(f)$ given by Eq. (5.41). Benchmark parameters from Sec. 5.4 are used to match the high frequency detection range of resonant cavity based detectors, such as ADMX. The above plot shows the characteristic strain as a function of frequency for the benchmark binary GA system, with the the peak frequency $f_p = 1.001\,\mathrm{GHz}$ visualized by the red dashed line.

we obtain the following expression:

$$\tilde{h}_\times(\omega) = \frac{h_0 \cos\iota}{2i} e^{i\omega r} \int dt \left[e^{i\omega t - 2i\Delta m\varphi(t)} Q(t) - e^{i\omega t + 2i\Delta m\varphi(t)} Q^*(t) \right] \tag{7.30}$$

where we renamed $t_{\mathrm{re}} \to t$.

Evaluating the integral through the same method of stationary phase approximation within [147] with the same arguments so that the first integral vanishes, we obtain:

$$\tilde{h}_\times(\omega) = -\frac{\sqrt{\pi}h_0 \cos\iota}{2i\sqrt{\gamma|\Delta m|}} Q^*(t_+(\omega))e^{i\Psi_+(\omega)} \tag{7.31}$$

where $t_+(\omega) = \frac{\omega-\omega_c}{2|\Delta m|\gamma}$ donates the non-zero stationary point and the phase argument $\Psi_+(\omega)$ is given by $\Psi_+(\omega) = \omega r + \frac{(\omega-\omega_c)^2}{4|\Delta m|\gamma} - \frac{\pi}{4}$, identical to the form we defined within Eq. (5.41). Substituting in $Q(t)$ as detailed within [147], we obtain:

$$\tilde{h}_\times(f) = -\frac{2}{i}h_0 \cos\iota\,\sqrt{\pi}\,|\Delta m|^2\, i\, e^{i\Psi_+(f)}\, \frac{\sqrt{z}}{|\Gamma| - i\pi(f-f_c)}\, e^{-\pi z}\, \exp\left[-2z \arctan\left(\frac{\pi(f-f_c)}{|\Gamma|}\right)\right] \tag{7.32}$$

where we rewrote the frequency as $f = \omega/2\pi$. Comparing to Eq. (5.41):

$$\tilde{h}_+(f) = h_0(1+\cos^2\iota)\,\sqrt{\pi}\,|\Delta m|^2\, i\, e^{i\Psi_+(f)}\, \frac{\sqrt{z}}{|\Gamma| - i\pi(f-f_c)}\, e^{-\pi z}\, \exp\left[-2z \arctan\left(\frac{\pi(f-f_c)}{|\Gamma|}\right)\right]$$

we notice that similar to isolated level transitions described within Sec. 7.4.2, the two polarizations differ only by the inclination weighting and a relative phase of $\pi/2$ given by the factor of $-2/i$. Figure 7.3 showcases the characteristic strain for the benchmark parameters that produce UHF-GWs with frequency $\sim$1 GHz.

Bibliography

[1] M. Aaboud et al. Search for dark matter and other new phenomena in events with an energetic jet and large missing transverse momentum using the atlas detector. *Journal of High Energy Physics*, 2018(1), January 2018.

[2] Georges Aad et al. Search for resonances decaying into photon pairs in 139 fb^{-1} of pp collisions at $\sqrt{s}$=13 TeV with the ATLAS detector. *Phys. Lett. B*, 822:136651, 2021.

[3] B. P. Abbott and month=oct others, year=2016. The basic physics of the binary black hole merger gw150914. *Annalen der Physik*, 529(1–2).

[4] L. F. Abbott and Pierre Sikivie. A cosmological bound on the invisible axion. *Phys. Lett. B*, 120:133–136, 1983.

[5] R. Abuter et al. Mass distribution in the Galactic Center based on interferometric astrometry of multiple stellar orbits. *Astron. Astrophys.*, 657:L12, 2022.

[6] M. Ackermann et al. The large area telescope on thefermi gamma-ray space telescopemission. *The Astrophysical Journal*, 697(2):1071–1102, May 2009.

[7] M. Ackermann et al. Search for gamma-ray spectral lines with the fermi large area telescope and dark matter implications. *Physical Review D*, 88(8), October 2013.

[8] E. G. Adelberger, B. R. Heckel, and A. E. Nelson. Particle-physics implications of a recent test of the gravitational inverse-square law. *Phys. Rev. Lett.*, 98:131104, 2007.

[9] Nancy Aggarwal et al. Challenges and opportunities of gravitational-wave searches at MHz to GHz frequencies. *Living Rev. Rel.*, 24(1):4, 2021.

[10] S. O. Alexeyev, X. Calmet, and B. N. Latosh. Gravity induced non-local effects in the standard model. *Phys. Lett. B*, 776:111–114, 2018.

[11] Yacine Ali-Haïmoud and Marc Kamionkowski. Cosmic microwave background limits on accreting primordial black holes. *Phys. Rev. D*, 95(4):043534, 2017.

[12] Muhsin Aljaf and Ilias Cholis. Simulating binary primordial black hole mergers in dark matter halos. *Phys. Rev. D*, 111:063020, Mar 2025.

[13] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H.-S. Shao, T. Stelzer, P. Torrielli, and M. Zaro. The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations. *Journal of High Energy Physics*, 2014(7), July 2014.

[14] Dorian Amaral et al. Global detector network to search for high-frequency gravitational waves (GravNet): conceptual design. 3 2026.

[15] Ana, Gia Dvali, and Emmanouil Koutsangelas. New mass window for primordial black holes as dark matter from memory burden effect. 2024.

[16] Alexandre Arbey and Jérémy Auffinger. Physics Beyond the Standard Model with BlackHawk v2.0. *Eur. Phys. J. C*, 81:910, 2021.

[17] Alexandre Arbey and Jérémy Auffinger. Blackhawk: a public code for calculating the hawking evaporation spectra of any black hole distribution. *The European Physical Journal C*, 79(8), August 2019.

[18] Giorgio Arcadi, Maíra Dutra, Pradipta Ghosh, Manfred Lindner, Yann Mambrini, Mathias Pierre, Stefano Profumo, and Farinaldo S. Queiroz. The waning of the WIMP? A review of models, searches, and constraints. *Eur. Phys. J. C*, 78(3):203, 2018.

[19] Nima Arkani-Hamed, Timothy Cohen, Raffaele Tito D'Agnolo, Anson Hook, Hyung Do Kim, and David Pinner. Solving the hierarchy problem at reheating with a large number of degrees of freedom. *Physical Review Letters*, 117(25), December 2016.

[20] Nima Arkani-Hamed, Savas Dimopoulos, and G. R. Dvali. The Hierarchy problem and new dimensions at a millimeter. *Phys. Lett. B*, 429:263–272, 1998.

[21] Asimina Arvanitaki, Masha Baryakhtar, and Xinlu Huang. Discovering the qcd axion with black holes and gravitational waves. *Physical Review D*, 91(8), April 2015.

[22] Asimina Arvanitaki and Sergei Dubovsky. Exploring the string axiverse with precision black hole physics. *Phys. Rev. D*, 83:044026, 2011.

[23] Zaven Arzoumanian et al. The NANOGrav 12.5 yr Data Set: Search for an Isotropic Stochastic Gravitational-wave Background. *Astrophys. J. Lett.*, 905(2):L34, 2020.

[24] E. Bagui, S. Clesse, V. D. Luca, J. Ezquiaga, G. Franciolini, J. Garc'ia-Bellido, Cristian Joana, R. K. Jain, S. Kuroyanagi, I. Musco, T. Papanikolaou, A. Raccanelli, S. Renaux-Petel, A. Riotto, E. R. Morales, M. Scalisi, O. Sergijenko, C. Unal, V. Vennin, and David Wands. Primordial black holes and their gravitational-wave signatures, 2023.

[25] Michael J. Baker and Andrea Thamm. Probing the particle spectrum of nature with evaporating black holes. *SciPost Phys.*, 12(5):150, 2022.

[26] Michael J. Baker and Andrea Thamm. Black hole evaporation beyond the Standard Model of particle physics. *JHEP*, 01:063, 2023.

[27] G Ballesteros and M Taoso. On the merger rate of primordial black holes: effects of nearest neighbors distribution and clustering. *Journal of Cosmology and Astroparticle Physics*, 2018.

[28] Aurélien Barrau, Juan García-Bellido, Thierry Grenet, and Killian Martineau. Prospects for detection of ultra high frequency gravitational waves from compact binary coalescenses with resonant cavities. 3 2023.

[29] Aurélien Barrau, Juan García-Bellido, Killian Martineau, and Martin Teuscher. Prospects for detection of ultra high frequency gravitational waves from hyperbolic encounters with resonant cavities. 4 2024.

[30] C. Bartram et al. Search for Invisible Axion Dark Matter in the 3.3–4.2 μeV Mass Range. *Phys. Rev. Lett.*, 127(26):261803, 2021.

[31] A. O. Barvinsky and G. A. Vilkovisky. Covariant perturbation theory (i). nonlocal effective action. *Nucl. Phys. B*, 333:471–511, 1990.

[32] Masha Baryakhtar, Robert Lasenby, and Matthew Teo. Black hole superradiance signatures of ultralight vectors. *Phys. Rev. D*, 96(3):035019, 2017.

[33] Daniel Baumann, Horng Sheng Chia, and Rafael A. Porto. Probing ultralight bosons with binary black holes. *Phys. Rev. D*, 99(4):044001, 2019.

[34] Daniel Baumann, Horng Sheng Chia, Rafael A. Porto, and Jason Stout. Gravitational collider physics. *Phys. Rev. D*, 101(8):083019, 2020.

[35] Asher Berlin, Diego Blas, Raffaele Tito D'Agnolo, Sebastian A. R. Ellis, Roni Harnik, Yonatan Kahn, and Jan Schütte-Engel. Detecting high-frequency gravitational waves with microwave cavities. *Phys. Rev. D*, 105(11):116011, 2022.

[36] Asher Berlin et al. Searches for New Particles, Dark Matter, and Gravitational Waves with SRF Cavities. 3 2022.

[37] Niels E. J. Bjerrum-Bohr, John F. Donoghue, and Barry R. Holstein. Quantum

gravitational corrections to the nonrelativistic scattering potential of two masses. *Phys. Rev. D*, 67:084033, 2003.

[38] Diego Blas, Yifan Chen, Yuxin Liu, Yanfei Shang, and Jing Shu. Cavity Multimodes as an Array for High-Frequency Gravitational Waves. 1 2026.

[39] C. Boutan et al. Piezoelectrically Tuned Multimode Cavity Search for Axion Dark Matter. *Phys. Rev. Lett.*, 121(26):261302, 2018.

[40] T. Braine et al. Extended Search for the Invisible Axion with the Axion Dark Matter Experiment. *Phys. Rev. Lett.*, 124(10):101303, 2020.

[41] Torsten Bringmann, Valerie Domcke, Elina Fuchs, and Joachim Kopp. High-frequency gravitational wave detection via optical frequency modulation. *Phys. Rev. D*, 108(6):L061303, 2023.

[42] R. Brito, V. Cardoso, and P. Pani. *Superradiance: Energy Extraction, Black-Hole Bombs and Implications for Astrophysics and Particle Physics.* Lecture Notes in Physics. Springer International Publishing, 2015.

[43] Richard Brito, Vitor Cardoso, and Paolo Pani. Superradiance. *Lect. Notes Phys.*, 906:1–237, 2015.

[44] Adam R. Brown, Luca V. Iliesiu, Geoff Penington, and Mykhaylo Usatyuk. The evaporation of charged black holes. 2024.

[45] Daniel Burns and Apostolos Pilaftsis. Matter Quantum Corrections to the Graviton Self-Energy and the Newtonian Potential. *Phys. Rev. D*, 90:024033, 2014.

[46] C. Bìhm, A. Kobakhidze, Ciaran A.J. O'Hare, Zachary S.C. Picker, and M. Sakellariadou. Eliminating the ligo bounds on primordial black hole dark matter, 2020.

[47] Xavier Calmet, Stephen D. H. Hsu, and David Reeb. Quantum gravity at a TeV and the renormalization of Newton's constant. *Phys. Rev. D*, 77:125015, 2008.

[48] Antonino Cangialosi. Cosmological constraints on primordial black holes, 2019.

[49] Chiara Caprini and Daniel G. Figueroa. Cosmological backgrounds of gravitational waves. *Class. Quant. Grav.*, 35(16):163001, 2018.

[50] Vitor Cardoso and Shijun Yoshida. Testing the nature of dark compact objects: a status report. *Living Rev. Relativ.*, 22(1):4, 2019.

[51] Bernard Carr and Florian Kuhnel. Primordial Black Holes as Dark Matter: Recent Developments. *Ann. Rev. Nucl. Part. Sci.*, 70:355–394, 2020.

[52] Bernard Carr, Florian Kühnel, and Marit Sandstad. Primordial black holes as dark matter. *Phys. Rev. D*, 94(8):083504, 2016.

[53] Bernard Carr, Martti Raidal, Tommi Tenkanen, Ville Vaskonen, and Hardi Veermäe. Primordial black hole constraints for extended mass functions. *Phys. Rev. D*, 96(2):023514, 2017.

[54] Sean M. Carroll. *Spacetime and Geometry: An Introduction to General Relativity.* Cambridge University Press, 7 2019.

[55] J. Casares, J. A. Orosz, C. Zurita, T. Shahbaz, J. M. Corral-Santana, J. E. McClintock, M. R. Garcia, I. G. Martínez-Pais, P. A. Charles, R. P. Fender, and R. A. Remillard. Refined Orbital Solution and Quiescent Variability in the Black Hole Transient GS 1354-64 (= BW Cir). *Ap. J. S.*, 181(1):238–243, March 2009.

[56] Z. Chacko, Hock-Seng Goh, and Roni Harnik. The Twin Higgs: Natural electroweak breaking from mirror symmetry. *Phys. Rev. Lett.*, 96:231802, 2006.

[57] Chiang-Mei Chen and Ming-Fan Wu. Spontaneous Pair Production in Reissner-Nordström Black Holes. *Phys. Rev. D*, 85:124041, 2012.

[58] Takeshi Chiba and Shuichiro Yokoyama. Spin distribution of primordial black holes. *Progress of Theoretical and Experimental Physics*, 2017(8):083E01, 2017.

[59] D.B. Clark, E. Godat, and F.I. Olness. Maneparse: A mathematica reader for parton distribution functions. *Computer Physics Communications*, 216:126–137, July 2017.

[60] Alessandro Codello and Rajeev Kumar Jain. On the covariant formalism of the effective field theory of gravity and leading order corrections. *Class. Quant. Grav.*, 33(22):225006, 2016.

[61] HAWC Collaboration, A. U. Abeysekara, R. Alfaro, C. Alvarez, J. D. Álvarez, R. Arceo, J. C. Arteaga-Velázquez, H. A. Ayala Solares, A. S. Barber, B. M. Baughman, N. Bautista-Elivar, E. Belmont, S. Y. BenZvi, D. Berley, M. Bonilla Rosales, J. Braun, R. A. Caballero-Lopez, K. S. Caballero-Mora, A. Carramiñana,

M. Castillo, U. Cotti, J. Cotzomi, E. de la Fuente, C. De León, T. DeYoung, R. Diaz Hernandez, J. C. Díaz-Vélez, B. L. Dingus, M. A. DuVernois, R. W. Ellsworth, A. Fernandez, D. W. Fiorino, N. Fraija, A. Galindo, F. Garfias, L. X. González, M. M. González, J. A. Goodman, V. Grabski, M. Gussert, Z. Hampel-Arias, C. M. Hui, P. Hüntemeyer, A. Imran, A. Iriarte, P. Karn, D. Kieda, G. J. Kunde, A. Lara, R. J. Lauer, W. H. Lee, D. Lennarz, H. León Vargas, E. C. Linares, J. T. Linnemann, M. Longo, R. Luna-GarcIa, A. Marinelli, H. Martinez, O. Martinez, J. Martínez-Castro, J. A. J. Matthews, P. Miranda-Romagnoli, E. Moreno, M. Mostafá, J. Nava, L. Nellen, M. Newbold, R. Noriega-Papaqui, T. Oceguera-Becerra, B. Patricelli, R. Pelayo, E. G. Pérez-Pérez, J. Pretz, C. Rivière, D. Rosa-González, H. Salazar, F. Salesa, F. E. Sanchez, A. Sandoval, E. Santos, M. Schneider, S. Silich, G. Sinnis, A. J. Smith, K. Sparks, R. W. Springer, I. Taboada, P. A. Toale, K. Tollefson, I. Torres, T. N. Ukwatta, L. Villaseñor, T. Weisgarber, S. Westerhoff, I. G. Wisher, J. Wood, G. B. Yodh, P. W. Younk, D. Zaborov, A. Zepeda, and H. Zhou. The hawc gamma-ray observatory: Dark matter, cosmology, and fundamental physics, 2013.

[62] J. Conrad, J. Cohen-Tanugi, and L. E. Strigari. Wimp searches with gamma rays in the fermi era: Challenges, methods and results. *Journal of Experimental and Theoretical Physics*, 121(6):1104–1135, December 2015.

[63] Adam Coogan, Logan Morrison, and Stefano Profumo. Hazma: a python toolkit for studying indirect detection of sub-gev dark matter. *Journal of Cosmology and*

Astroparticle Physics, 2020(01):056–056, January 2020.

[64] A. M. Cruise and R. M. Ingley. A prototype gravitational wave detector for 100-mhz. *Class. Quant. Grav.*, 23:6185–6193, 2006.

[65] Hooman Davoudiasl, Peter B. Denton, and David A. McGady. Ultralight fermionic dark matter. *Phys. Rev. D*, 103(5):055014, 2021.

[66] José Antonio de Freitas Pacheco. Evaporation of Primordial Charged Black Holes: Timescale and Evolution of Thermodynamic Parameters. *Symmetry*, 16(7):895, 2024.

[67] V. De Luca, V. Desjacques, G. Franciolini, A. Malhotra, and A. Riotto. The initial spin probability distribution of primordial black holes. *JCAP*, 05:018, 2019.

[68] V De Luca and A Riotto. The clustering evolution of primordial black holes. *Journal of Cosmology and Astroparticle Physics*, 2020.

[69] Adrián del Rio, Ruth Durrer, and Subodh P. Patil. Tensor Bounds on the Hidden Universe. *JHEP*, 12:094, 2018.

[70] Sten M Delos and Fabian Schmidt. Structure formation with primordial black holes: collisional dynamics, binaries, and gravitational waves. *Journal of Cosmology and Astroparticle Physics*, 2024.

[71] Steven Detweiler. Klein-gordon equation and rotating black holes. *Phys. Rev. D*, 22:2323–2326, Nov 1980.

[72] Giovanni di Cortona, Edward Hardy, Javier Pardo Vega, and Giovanni Villadoro. The qcd axion, precisely. *JHEP*, 01:034, 2016.

[73] Sofia Di Gennaro and Yen Chin Ong. How Not to Extract Information from Black Holes: Cosmic Censorship as a Guiding Principle. *Phys. Lett. B*, 819:136432, 2021.

[74] Keith R. Dienes, Lucien Heurtier, Fei Huang, Doojin Kim, Tim M. P. Tait, and Brooks Thomas. Stasis in an expanding universe: A recipe for stable mixed-component cosmological eras. *Phys. Rev. D*, 105(2):023530, 2022.

[75] Keith R. Dienes and Brooks Thomas. Dynamical Dark Matter: I. Theoretical Overview. *Phys. Rev. D*, 85:083523, 2012.

[76] Keith R. Dienes and Brooks Thomas. Dynamical Dark Matter: II. An Explicit Model. *Phys. Rev. D*, 85:083524, 2012.

[77] Michael Dine and Willy Fischler. The not-so-harmless axion. *Phys. Lett. B*, 120:137–141, 1983.

[78] Sam R. Dolan. Instability of the massive Klein-Gordon field on the Kerr spacetime. *Phys. Rev. D*, 76:084001, 2007.

[79] Alexander D. Dolgov and Damian Ejlli. Relic gravitational waves from light primordial black holes. *Phys. Rev. D*, 84:024028, 2011.

[80] V. Domcke, Sebastian A. R. Ellis, and Joachim Kopp. Dielectric haloscopes as gravitational wave detectors, 2024.

[81] V. Domcke and C. Garcia-Cely. Potential of radio telescopes as high-frequency gravitational wave detectors., 2021.

[82] Valerie Domcke, Camilo Garcia-Cely, and Nicholas L. Rodd. Novel Search for High-Frequency Gravitational Waves with Low-Mass Axion Haloscopes. *Phys. Rev. Lett.*, 129(4):041101, 2022.

[83] Valerie Domcke, Camilo Garcia-Cely, and Nicholas L. Rodd. A novel search for high-frequency gravitational waves with low-mass axion haloscopes. *Phys. Rev. Lett.*, 129(4):041101, 2022.

[84] Guillem Domènech and Misao Sasaki. Probing primordial black hole scenarios with terrestrial gravitational wave detectors. *Classical and Quantum Gravity*, 2024.

[85] R. Dong and D. Stojkovic. Gravitational wave production by hawking radiation from rotating primordial black holes. *Journal of Cosmology and Astroparticle Physics*, 2015.

[86] John F. Donoghue. General relativity as an effective field theory: The leading quantum corrections. *Phys. Rev. D*, 50:3874–3888, 1994.

[87] John F. Donoghue and Barry R. Holstein. Quantum corrections to the newtonian potential. *J. Phys. G*, 44(4):04LT01, 2017.

[88] GIA DVALI. Nature of microscopic black holes and gravity in theories with

particle species. *International Journal of Modern Physics A*, 25(02n03):602–615, January 2010.

[89] Gia Dvali. A microscopic model of holography: Survival by the burden of memory, 2018.

[90] Gia Dvali, Lukas Eisemann, Marco Michel, and Sebastian Zell. Black hole metamorphosis and stabilization by memory burden. *Physical Review D*, 102(10), November 2020.

[91] Gia Dvali, Juan Sebastián Valbuena-Bermúdez, and Michael Zantedeschi. Memory burden effect in black holes and solitons: Implications for pbh, 2024.

[92] Yohei Ema, Kazunori Nakayama, and Yong Tang. Production of purely gravitational dark matter. *Journal of High Energy Physics*, 2018(9), September 2018.

[93] Razieh Emami and Abraham Loeb. Formation redshift of the massive black holes detected by ligo, 2019.

[94] Solomon Endlich and Riccardo Penco. A modern approach to superradiance. *JHEP*, 05:052, 2017.

[95] A Escrivá and S Clesse. Simulations of pbh formation at the qcd epoch and comparison with the gwtc-3 catalog. *Journal of Cosmology and Astroparticle Physics*, 2022.

[96] Albert Escrivà, Florian Kuhnel, and Yuichiro Tada. Primordial Black Holes. 11 2022.

[97] Christopher Ewasiuk and Stefano Profumo. Constraints on the maximal number of dark degrees of freedom from black hole evaporation, cosmic rays, colliders, and supernovae. *Phys. Rev. D*, 111:015008, 2025.

[98] Christopher Ewasiuk and Stefano Profumo. Dark-sector modifications to Kerr and Reissner–Nordström black hole evaporation. 2025.

[99] Christopher Ewasiuk and Stefano Profumo. Data and code for "dark-sector modifications to kerr and reissner–nordström black hole evaporation". `https://github.com/cewasiuk/Spin-and-Charge-BH-Evolution`, 2025. GitHub repository.

[100] S. Faller. Effective field theory of gravity: Leading quantum gravitational corrections to Newton's and Coulomb's laws. *Phys. Rev. D*, 77:124039, 2008.

[101] Pierre Fayet. MICROSCOPE limits for new long-range forces and implications for unified theories. *Phys. Rev. D*, 97(5):055039, 2018.

[102] Nicolas Fernandez and Stefano Profumo. Unraveling the origin of black holes from effective spin measurements with ligo-virgo. *Journal of Cosmology and Astroparticle Physics*, 2019(08):022–022, August 2019.

[103] A. V. Filippenko, D. C. Leonard, T. Matheson, W. Li, E. C. Moran, and A. G. Riess. A black hole in the x-ray nova velorum 1993. *Publ. Astron. Soc. Pac.*, 111:969–979, 1999.

[104] Robert Foot, H. Lew, and R. R. Volkas. A Model with fundamental improper space-time symmetries. *Phys. Lett. B*, 272:67–70, 1991.

[105] Gabriele Franciolini, Anshuman Maharana, and Francesco Muia. Hunt for light primordial black hole dark matter with ultrahigh-frequency gravitational waves. *Phys. Rev. D*, 106(10):103520, 2022.

[106] Markus B. Fröb. Quantum gravitational corrections for spinning particles. *JHEP*, 10:051, 2016. [Erratum: JHEP 11, 176 (2016)].

[107] Markus B. Fröb, William C. C. Lima, and Elizabeth Winstanley. Quantum corrections to the newtonian potential in nonlocal gravity. *Phys. Rev. D*, 92:064013, 2015.

[108] Juan Garcia-Bellido, Marco Peloso, and Caner Unal. Gravitational Wave signatures of inflationary models from Primordial Black Hole Dark Matter. *JCAP*, 09:013, 2017.

[109] Claudio Gatti, Luca Visinelli, and Michael Zantedeschi. Cavity detection of gravitational waves: Where do we stand? *Phys. Rev. D*, 110(2):023018, 2024.

[110] Thomas C. Gehrman, Barmak Shams Es Haghi, K. Sinha, and Tao Xu. Baryogenesis, primordial black holes and mhz–ghz gravitational waves, 2022.

[111] Thomas C. Gehrman, Barmak Shams Es Haghi, K. Sinha, and Tao Xu. The primordial black holes that disappeared: connections to dark matter and mhz-ghz gravitational waves, 2023.

[112] G. W. Gibbons. Vacuum Polarization and the Spontaneous Loss of Charge by Black Holes. *Commun. Math. Phys.*, 44:245–264, 1975.

[113] G. W. Gibbons and S. W. Hawking. Cosmological event horizons, thermodynamics, and particle creation. *Phys. Rev. D*, 15:2738–2751, May 1977.

[114] Massimo Giovannini. Relic gravitons and high-frequency detectors. *JCAP*, 05:056, 2023.

[115] Gian F. Giudice and Matthew McCullough. A Clockwork Theory. *JHEP*, 02:036, 2017.

[116] Andrew D. Gow, C. Byrnes, A. Hall, and J. Peacock. Primordial black hole merger rates: distributions for multiple ligo observables, 2019.

[117] A. Green and B. Kavanagh. Primordial black holes as a dark matter candidate, 2020.

[118] Tomohiro Harada, Chul-Moon Yoo, Kazunori Kohri, Ken-ichi Nakao, and Sanjay Jhingan. Spins of primordial black holes formed in the matter-dominated phase of the Universe. *Phys. Rev. D*, 96:083517, 2017.

[119] Gregory M. Harry, Thomas R. Stevenson, and Ho Jung Paik. Detectability of gravitational wave events by spherical resonant mass antennas. *Phys. Rev. D*, 54:2409–2420, 1996.

[120] S. W. HAWKING. Black hole explosions? *Nature*, 248(5443):30–31, 1974.

[121] S. W. Hawking. Particle creation by black holes. *Communications in Mathematical Physics*, 43(3):199–220, 1975.

[122] S. W. Hawking and M. Taylor-Robinson. Evolution of near extremal black holes. *Physical Review D*, 1997.

[123] Nicolas Herman, André Füzfa, Léonard Lehoucq, and Sébastien Clesse. Detecting planetary-mass primordial black holes with resonant electromagnetic gravitational-wave detectors. *Phys. Rev. D*, 104(2):023524, 2021.

[124] W. A. Hiscock and L. D. Weems. Evolution of Charged Black Holes. *Phys. Rev. D*, 41:1142, 1990.

[125] Shahar Hod. Hawking radiation may violate the penrose cosmic censorship conjecture. *International Journal of Modern Physics D*, 2019.

[126] Barry R. Holstein. On-shell calculation of mixed electromagnetic and gravitational scattering. *Phys. Rev. D*, 108(6):066013, 2023.

[127] Sebastian Hoof. bhsr: Black hole superradiance code. `https://github.com/sebhoof/bhsr`, 2025. GitHub repository.

[128] Sebastian Hoof, David J. E. Marsh, Júlia Sisk-Reynés, James H. Matthews, and Christopher Reynolds. Getting more out of black hole superradiance: a statistically rigorous approach to ultralight boson constraints from black hole spin measurements. *Monthly Notices of the Royal Astronomical Society*, 546(2):staf1564, September 2025.

[129] D. Hooper, G. Krnjaic, J. March-Russell, Samuel D. McDermott, and

R. Petrossian-Byrne. Hot gravitons and gravitational waves from kerr black holes in the early universe., 2020.

[130] Dan Hooper and Stefano Profumo. Dark Matter and Collider Phenomenology of Universal Extra Dimensions. *Phys. Rept.*, 453:29–115, 2007.

[131] C. D. Hoyle, D. J. Kapner, B. R. Heckel, E. G. Adelberger, J. H. Gundlach, U. Schmidt, and H. E. Swanson. Sub-millimeter tests of the gravitational inverse-square law. *Phys. Rev. D*, 70:042004, 2004.

[132] Keisuke Inomata, M. Kawasaki, K. Mukaida, Y. Tada, and T. Yanagida. Inflationary primordial black holes for the ligo gravitational wave events and pulsar timing array experiments, 2016.

[133] Aurora Ireland, Stefano Profumo, and Jordan Scharnhorst. Primordial gravitational waves from black hole evaporation in standard and nonstandard cosmologies. *Phys. Rev. D*, 107(10):104021, 2023.

[134] Aurora Ireland, Stefano Profumo, and Jordan Scharnhorst. Gravitational waves from primordial black hole evaporation with large extra dimensions. *JCAP*, 08:033, 2024.

[135] Maximiliano Isi, Ling Sun, Richard Brito, and Andrew Melatos. Directed searches for gravitational waves from ultralight bosons. *Physical Review D*, 99(8), apr 2019.

[136] H. Th. Janka. Neutrino Emission from Supernovae. 2 2017.

[137] K Jedamzik. Consistency of primordial black hole dark matter with ligo/virgo merger rates. *Physical Review Letters*, 2020.

[138] Vassiliki Kalogera and Gordon Baym. The maximum mass of a neutron star. *Astrophys. J. Lett.*, 470:L61–L64, 1996.

[139] D. J. Kapner, T. S. Cook, E. G. Adelberger, J. H. Gundlach, B. R. Heckel, C. D. Hoyle, and H. E. Swanson. Tests of the gravitational inverse-square law below the dark-energy length scale. *Phys. Rev. Lett.*, 98:021101, 2007.

[140] Andrey Katz et al. Femtolensing by dark matter revisited. *Journal of Cosmology and Astroparticle Physics*, 2018(12):005, 2018.

[141] Roy P. Kerr. Gravitational field of a spinning mass as an example of algebraically special metrics. *Phys. Rev. Lett.*, 11:237–238, Sep 1963.

[142] I. B. Khriplovich. Particle creation by charged black holes. *Phys. Rept.*, 465:1–18, 2008.

[143] Younggeun Kim, Jordan Gué, Changhao Xu, Diego Blas, Dmitry Budker, Sungjae Bae, Claudio Gatti, Junu Jeong, Jihn E. Kim, Kiwoong Lee, Arjan F. van Loo, Yasunobu Nakamura, Seonjeong Oh, Wolfram Ratzinger, Taehyeon Seong, Yannis K. Semertzidis, Kristof Schmieden, Mattias Schott, Sergey Uchaikin, and SungWoo Youn. Search for high-frequency gravitational waves via reanalysis of cavity axion data. *Phys. Rev. D*, 113:072015, Apr 2026.

[144] G. G. Kirilin and I. B. Khriplovich. Quantum power correction to the Newton law. *JETP*, 95:981–986, 2002.

[145] Valeriya Korol and Melvyn B Davies. Merger rates in primordial black hole clusters without initial binaries. *Monthly Notices of the Royal Astronomical Society*, 2019.

[146] Mrunal Korwar and Stefano Profumo. Late-forming black holes and the antiproton, gamma-ray, and antihelium excesses. *Phys. Rev. D*, 111(2):023032, 2025.

[147] Alexandros Kyriazis and Fan Yang. Gravitational waves from resonant transitions of tidally perturbed gravitational atoms. *JHEP*, 11:062, 2025.

[148] B. Latosh. Feyngrav 2.0. *Computer Physics Communications*, 292:108871, November 2023.

[149] S. Lee, S. Ahn, J. Choi, B. R. Ko, and Y. K. Semertzidis. Axion Dark Matter Search around 6.7 μeV. *Phys. Rev. Lett.*, 124(10):101802, 2020.

[150] B. Lehmann, C. Johnson, and S. Profumo. Indirect detection of gravitino dark matter including its three-body decays. *JCAP*, 08:033, 2019.

[151] Benjamin V. Lehmann, Stefano Profumo, and Jackson Yant. Model-independent discovery prospects for primordial black holes at LIGO. *Mon. Not. Roy. Astron. Soc.*, 501(3):3727–3740, 2021.

[152] Ulf Leonhardt. Cosmological horizons radiate. *Europhysics Letters*, 135(1):10002, July 2021.

[153] Antoine Letessier-Selvon and Todor Stanev. Ultrahigh Energy Cosmic Rays. *Rev. Mod. Phys.*, 83:907–942, 2011.

[154] Puxin Lin and Gary Shiu. Schwinger Effect of Extremal Reissner-Nordström Black Holes. *Not Provided*, 2024.

[155] V. D. Luca. Signals from the early universe: Black holes, gravitational waves and particle physics, 2022.

[156] Jane H. MacGibbon, B. J. Carr, and Don N. Page. Do evaporating black holes form photospheres? *Physical Review D*, 78(6), September 2008.

[157] Michele Maggiore. Gravitational wave experiments and early universe cosmology. *Phys. Rept.*, 331:283–367, 2000.

[158] Michele Maggiore. *Gravitational Waves. Vol. 1: Theory and Experiments*. Oxford University Press, 2007.

[159] Yann Mambrini and Keith A. Olive. Gravitational production of dark matter during reheating. *Physical Review D*, 103(11), June 2021.

[160] V. Mandic, Simeon Bird, and I. Cholis. Stochastic gravitational-wave background due to primordial binary black hole mergers., 2016.

[161] Stephen P. Martin. A Supersymmetry primer. *Adv. Ser. Direct. High Energy Phys.*, 18:1–98, 1998.

[162] Isabella Masina. Dark matter and dark radiation from evaporating kerr primordial black holes. *Gravitation and Cosmology*, 27:134–147, 2021.

[163] Taillte May, William E. East, and Nils Siemonsen. Self-gravity effects of ultralight boson clouds formed by black hole superradiance. *Phys. Rev. D*, 111:044062, Feb 2025.

[164] Ben T. McAllister, Graeme Flower, Justin Kruger, Eugene N. Ivanov, Maxim Goryachev, Jeremy Bourhill, and Michael E. Tobar. The ORGAN Experiment: An axion haloscope above 15 GHz. *Phys. Dark Univ.*, 18:67–72, 2017.

[165] Andrew L. Miller. Gravitational waves from sub-solar mass primordial black holes. 2024.

[166] Andrew L. Miller et al. Constraints on planetary and asteroid-mass primordial black holes from continuous gravitational-wave searches. *Physical Review D*, 104(10):104029, 2021.

[167] Christopher J. Moore, Robert H. Cole, and Christopher P. L. Berry. Gravitational-wave sensitivity curves. *Class. Quant. Grav.*, 32(1):015014, 2015.

[168] Logan Morrison, Stefano Profumo, and Yan Yu. Melanopogenesis: Dark Matter of (almost) any Mass and Baryonic Matter from the Evaporation of Primordial Black Holes weighing a Ton (or less). *JCAP*, 05:005, 2019.

[169] S. Navas et al. Review of particle physics. *Phys. Rev. D*, 110(3):030001, 2024.

[170] Ken K. Y. Ng, Otto A. Hannuksela, Salvatore Vitale, and Tjonnie G. F. Li. Searching for ultralight bosons within spin measurements of a population of binary black hole mergers. *Phys. Rev. D*, 103:063010, Mar 2021.

[171] T. Ngampitipan and P. Boonserm. Bounding the greybody factors for non-rotating black holes. *Classical and Quantum Gravity*, 29, 2012.

[172] Hiroko Niikura et al. Constraints on earth-mass primordial black holes from ogle 5-year microlensing events. *Physical Review D*, 99(8):083503, 2019.

[173] Alexander H. Nitz, Sumit Kumar, Yi-Fan Wang, Shilpa Kastha, Shichao Wu, Marlin Schäfer, Rahul Dhurkunde, and Collin D. Capano. 4-OGC: Catalog of Gravitational Waves from Compact Binary Mergers. *Astrophys. J.*, 946(2):59, 2023.

[174] Don N. Page. Particle emission rates from a black hole. i. massless particles from uncharged black holes. *Physical Review D*, 13(2):198–206, 1976.

[175] Don N. Page. Particle emission rates from a black hole. ii. massless particles from a rotating hole. *Physical Review D*, 14(12):3260–3273, 1976.

[176] Don N. Page. Particle emission rates from a black hole. iii. charged leptons from a nonrotating hole. *Physical Review D*, 16(8):2402–2411, 1977.

[177] K. Pal, L. V. Sales, and J. Wudka. Ultralight Thomas-Fermi dark matter. *Phys. Rev. D*, 100(8):083007, 2019.

[178] Andrea Palessandro and Tony Rothman. A simple derivation of the Gertsenshtein effect. *Phys. Dark Univ.*, 40:101187, 2023.

[179] R. D. Peccei and Helen R. Quinn. Cp conservation in the presence of pseudoparticles. *Phys. Rev. Lett.*, 38:1440–1443, 1977.

[180] Mihael Petač and Karsten Jedamzik. Microlensing constraints on clustered primordial black holes. *Physical Review D*, 105(4):043507, 2022.

[181] Konrad Pfrang et al. Optical microlensing by primordial black holes with iacts. In *Proceedings of 37th International Cosmic Ray Conference — PoS(ICRC2021)*, 2021.

[182] Zachary S. C. Picker and Alexander Kusenko. Explaining the GeV excess with exploding black holes. *Phys. Lett. B*, 845:138175, 2023.

[183] R. M. Plotkin, A. Bahramian, J. C. A. Miller-Jones, M. T. Reynolds, P. Atri, T. J. Maccarone, A. W. Shaw, and P. Gandhi. Towards a larger sample of radio jets from quiescent black hole X-ray binaries. *Mon. Not. Roy. Astron. Soc.*, 503(3):3784–3795, 2021.

[184] John Preskill, Mark B. Wise, and Frank Wilczek. Cosmology of the invisible axion. *Phys. Lett. B*, 120:127–132, 1983.

[185] William H. Press and Saul A. Teukolsky. Floating orbits, superradiant scattering and the black-hole bomb. *Nature*, 238:211–212, 1972.

[186] S. Profumo and Fengwei Yang. On the anisotropy of the stochastic gravitational wave background from sub-horizon-collapsed primordial black hole mergers, 2023.

[187] Stefano Profumo, Lucas Brown, Christopher Ewasiuk, Angelo Ricarte, and Yifan Su. The maximal gravitational wave signal from asteroid-mass primordial black hole mergers at resonant microwave cavities. 2024.

[188] Oriol Pujolas, Ville Vaskonen, and Hardi Veermae. Prospects for probing gravitational waves from primordial black hole binaries. *Phys. Rev. D*, 104(8):083521, 2021.

[189] Martti Raidal, Christian Spethmann, Ville Vaskonen, and Hardi Veermäe. Formation and Evolution of Primordial Black Hole Binaries in the Early Universe. *JCAP*, 02:018, 2019.

[190] Martti Raidal, Ville Vaskonen, and Hardi Veermae. Gravitational Waves from Primordial Black Hole Mergers. *JCAP*, 09:037, 2017.

[191] Martti Raidal, Ville Vaskonen, and Hardi Veermae. Formation of primordial black hole binaries and their merger rates. 4 2024.

[192] Andreas Ringwald. Exploring the role of axions and other wisps in the dark universe. *Phys. Dark Univ.*, 1:116–135, 2012.

[193] Remo Ruffini, Gregory Vereshchagin, and She-Sheng Xue. Electron–positron pairs in physics and astrophysics: From heavy nuclei to black holes. *Physics Reports*, 487(1–4):1–140, February 2010.

[194] Sebastian Schenk, Kristof Schmieden, and Pedro Schwaller. Signatures of High-Frequency Gravitational Waves in Electromagnetic Cavities. 12 2025.

[195] Stephan Schlamminger, K. Y. Choi, T. A. Wagner, J. H. Gundlach, and E. G. Adelberger. Test of the equivalence principle using a rotating torsion balance. *Phys. Rev. Lett.*, 100:041101, 2008.

[196] K. Schmieden and M. Schott. The global network of cavities to search for gravitational waves (gravnet): A novel scheme to hunt gravitational waves signatures from the early universe, 2023.

[197] Julian Schwinger. On gauge invariance and vacuum polarization. *Phys. Rev.*, 82:664–679, Jun 1951.

[198] Nils Siemonsen, Taillte May, and William E. East. Modeling the black hole superradiance gravitational waveform. *Phys. Rev. D*, 107(10):104003, 2023.

[199] P. Sikivie. Experimental Tests of the Invisible Axion. *Phys. Rev. Lett.*, 51:1415–1417, 1983. [Erratum: Phys.Rev.Lett. 52, 695 (1984)].

[200] Torbjörn Sjöstrand, Stefan Ask, Jesper R. Christiansen, Richard Corke, Nishita Desai, Philip Ilten, Stephen Mrenna, Stefan Prestel, Christine O. Rasmussen, and Peter Z. Skands. An introduction to pythia 8.2. *Computer Physics Communications*, 191:159–177, June 2015.

[201] Evgeny Sorkin and Tsvi Piran. Formation and Evaporation of Charged Black Holes. *Phys. Rev. D*, 63, 2001.

[202] V Stasenko. Redshift evolution of primordial black hole merger rate. *Physical Review D*, 2024.

[203] Henry Su, Lucas Brown, Christopher Ewasiuk, and Stefano Profumo. High-frequency gravitational wave transients from superradiance. 2026.

[204] W. H. Tan, J. C. Long, H. Chan, C. G. Wang, S. Q. Yang, J. Luo, C. G. Shao, L. C. Tu, W. H. Zhao, and J. Luo. Improvement for testing the gravitational inverse-square law at the submillimeter range. *Phys. Rev. Lett.*, 124(5):051301, 2020.

[205] B. Taylor and W. Hiscock. Evaporation of a kerr black hole by emission of scalar and higher spin particles. *Physical Review D*, 55:6116–6125, 1998.

[206] Wolfgang Tichy and Pedro Marronetti. The Final mass and spin of black hole mergers. *Phys. Rev. D*, 78:081501, 2008.

[207] Giovanni Maria Tomaselli. Gravitational atoms and black hole binaries, 2024.

[208] Giovanni Maria Tomaselli, Thomas F. M. Spieksma, and Gianfranco Bertone. Resonant history of gravitational atoms in black hole binaries. *Phys. Rev. D*, 110:064048, Sep 2024.

[209] KA Toshchenko and KE Belotsky. Studying method of microlensing effect estimation for a cluster of primordial black holes. In *Journal of Physics: Conference Series*, volume 1390, page 012108, 2019.

[210] T. N. Ukwatta, D. R. Stump, J. T. Linnemann, J. H. MacGibbon, S. S. Marinelli, T. Yapici, and K. Tollefson. Primordial Black Holes: Observational characteristics of the final evaporation. *Astroparticle Physics*, 80:90–114, July 2016.

[211] Yi-Fan Wang and Andreas Nitz. Prospects for detecting gravitational waves from eccentric subsolar mass compact binaries. *The Astrophysical Journal*, 2021.

[212] Steven Weinberg. A new light boson? *Phys. Rev. Lett.*, 40:223–226, 1978.

[213] Frank Wilczek. Problem of strong p and t invariance in the presence of instantons. *Phys. Rev. Lett.*, 40:279–282, 1978.

[214] Zhu Yi, Qing Gao, Y. Gong, and Zong hong Zhu. Primordial black holes and secondary gravitational waves from inflationary model with a non-canonical kinetic term, 2020.

[215] Hirotaka Yoshino and Hideo Kodama. Gravitational radiation from an axion cloud around a black hole: Superradiant phase. *Prog. Theor. Exp. Phys.*, page 043E02, 2014.

[216] Hirotaka Yoshino and Hideo Kodama. The bosenova and axiverse. *Class. Quant. Grav.*, 32(21):214001, 2015.

[217] Hirotaka Yoshino and Hideo Kodama. Gravitational radiation from an axion cloud around a black hole: Superradiant phase. *Prog. Theor. Exp. Phys.*, page 061E01, 2015.

[218] Fang yu Li, Meng xi Tang, and D. Shi. Electromagnetic response of a gaussian beam to high frequency relic gravitational waves in quintessential inflationary models, 2003.

[219] Ya. B. Zel'dovich. Generation of waves by a rotating body. *JETP Lett.*, 14:180, 1971.

[220] Ya. B. Zel'dovich. Amplification of cylindrical electromagnetic waves reflected from a rotating body. *Sov. Phys. JETP*, 35:1085, 1972.

[221] L. Zhong et al. Results from phase 1 of the HAYSTAC microwave cavity axion experiment. *Phys. Rev. D*, 97(9):092001, 2018.

[222] T. J. M. Zouros and Douglas M. Eardley. Instabilities of massive scalar perturbations of a rotating black hole. *Annals Phys.*, 118:139–155, 1979.