

# TOWARD THE STUDY OF STARS WITH LSST

by

**Troy J. Raen**

B.S., University of Utah, 2002, 2016

M.S., University of Pittsburgh, 2018

Submitted to the Graduate Faculty of  
the Dietrich School of Arts and Sciences in partial fulfillment  
of the requirements for the degree of

**Doctor of Philosophy**

University of Pittsburgh

2022

UNIVERSITY OF PITTSBURGH  
DIETRICH SCHOOL OF ARTS AND SCIENCES

This dissertation was presented

by

Troy J. Raen

It was defended on

April 27th, 2022

and approved by

Michael Wood-Vasey, Department of Physics and Astronomy

Andrew Zentner, Department of Physics and Astronomy

Ann Lee, Department of Statistics & Data Science / Machine Learning Department,

Carnegie Mellon University

Brett Andrews, Department of Physics and Astronomy

Brian Batell, Department of Physics and Astronomy

Jeffrey Newman, Department of Physics and Astronomy

Dissertation Director: Michael Wood-Vasey, Department of Physics and Astronomy

Copyright © by Troy J. Raen  
2022

## ABSTRACT

### TOWARD THE STUDY OF STARS WITH LSST

Troy J. Raen, PhD

University of Pittsburgh, 2022

In this work I present two contributions towards the study of stars: computational modeling of the impact of asymmetric dark matter (ADM) on stellar evolution, and the Pitt-Google community alert broker. Both serve as preparation for the Vera C. Rubin Observatory’s upcoming Legacy Survey of Space and Time (LSST). LSST will revolutionize astronomy by revealing massive numbers of stars. Its nightly “alert” stream will be dominated by the dynamic life events of 7 million variable stars – caught in the act of flickering, pulsating, and erupting – and it will broadcast the spectacular deaths of another 0.2 million stars in supernovae. Its final catalog will contain 17 billion resolved stars in active and quiescent phases. These unprecedented sample sizes will both enable long-term population studies and contain individual events that are fleeting and exotic for which follow-up observations will need to be triggered rapidly. I will first discuss the potential impact of ADM. I wrote a plugin for the code-base Modules for Experiments in Stellar Astrophysics (MESA) that models the capture of ADM in stars (via scattering events) and the energy transport that ensues. Using this code, I ran a grid of stellar models which show that ADM can significantly alter a star’s internal structure and the timing of its evolution through the standard phases. I further show that this may have observable effects on the isochrones of LSST star clusters. Next, I will discuss the Pitt-Google broker which will serve the astronomy community by ingesting LSST’s torrential alert stream, adding technological and scientific value, and distributing useful subsets through convenient access methods. Our broker adopts cloud-based technologies and methods that are unique in the space. Our motivation is to enable broad

public access to, and scientific analysis of, LSST's alert stream with a low barrier to entry. I have developed the infrastructure necessary to run the end-to-end broker, from stream ingestion through user access tools. This framework currently processes and serves the nightly Zwicky Transient Facility (ZTF) alert stream. I will discuss the system design, user access, and ongoing development work.

**Keywords:** stars: evolution; stars: interiors; stars: low-mass; galaxies: dwarf; dark matter; asymmetric dark matter; software: data analysis; software: development; software: documentation; software: public release; alert streams; alert brokers; data services; cloud computing; .

## TABLE OF CONTENTS

<b>1.0 INTRODUCTION</b>	1
1.1 Stellar Evolution	2
1.1.1 Modeling Main Sequence Stars	4
1.2 Dark Matter in Stars	8
1.3 Alert Brokers	11
1.4 Scope of the Dissertation	13
<b>2.0 THE EFFECTS OF ASYMMETRIC DARK MATTER ON STELLAR EVOLUTION I: SPIN-DEPENDENT SCATTERING</b>	16
2.1 Introduction	16
2.2 Dark Matter Properties and Capture in Stars	19
2.3 Methods	22
2.3.1 Energy Transport by Dark Matter	23
2.4 Results	25
2.4.1 Low-Mass Stars	26
2.4.2 High-Mass Stars	29
2.4.3 Main Sequence Lifetimes	31
2.4.4 Stellar Evolutionary Tracks	34
2.4.5 Stellar Population Isochrones	39
2.5 Discussion and Conclusions	44
<b>3.0 THE PITT-GOOGLE LSST COMMUNITY ALERT BROKER</b>	46
3.1 Introduction	46
3.2 Google Cloud	48

3.2.1	Costs	49
3.2.2	Pub/Sub	50
3.2.3	BigQuery	52
3.2.4	Cloud Storage	52
3.2.5	Cloud Functions and Cloud Run	53
3.2.6	Access Methods	54
3.2.7	New User Setup	55
3.3	Broker Software	56
3.3.1	Conductor	57
3.3.2	Storage	61
3.3.3	Science	63
3.3.4	Admin	65
3.3.5	Periphery	67
3.3.6	Work in Progress and Future Plans	67
3.3.7	Pitt-Google's Budget	73
3.4	Data Access	74
3.4.1	<code>pittgoogle-client</code>	76
3.4.2	<code>tom-pittgoogle</code>	79
3.5	Summary	80
<b>4.0</b>	<b>CONCLUDING REMARKS</b>	<b>82</b>
	<b>APPENDIX A. OBSERVING STELLAR EVOLUTION WITH LSST</b>	<b>85</b>
A.1	Isochrones	85
A.2	Light Curves and the Alert Stream	86
A.2.1	Binary systems	89
A.2.2	Post-Main Sequence Stars	90
A.2.2.1	Very low mass stars - Red giants and white dwarfs.	91
A.2.2.2	Low mass stars - Red giants and white dwarfs.	91
A.2.2.3	Intermediate mass stars - Red giants and white dwarfs.	91
A.2.2.4	White dwarfs that exceed $1.4 M_{\odot}$ - Type Ia supernovae.	92

A.2.2.5 High mass stars - Red and blue supergiants, white dwarfs, supernovae, neutron stars, and black holes. . . . .	92
A.2.3 Transients . . . . .	93
A.2.3.1 Gap transients. . . . .	94
A.2.3.2 Kilonovae. . . . .	94
A.2.4 Variable Stars . . . . .	95
A.2.4.1 Pulsations - Cepheids. . . . .	95
A.2.4.2 Eclipsing - Binary systems viewed nearly edge-on. . . . .	95
A.2.4.3 Magnetic activity - Flare stars. . . . .	95
A.2.4.4 Accretion from a companion - Cataclysmic variables. . . . .	96
<b>APPENDIX B. ALERT BROKERS: COMMON SERVICES . . . . .</b>	<b>98</b>
<b>APPENDIX C. BAYESIAN NETWORKS . . . . .</b>	<b>101</b>
<b>BIBLIOGRAPHY . . . . .</b>	<b>103</b>

## LIST OF FIGURES

1	<p>Mass-luminosity relation of 190 stars. These stars are all in binaries or triples, which allows their parameters to be well-measured. For most of the mass range, the slope is well-approximated with <math>L \propto M_{\star}^{\sim 3.5}</math>, but it becomes noticeably shallower at both ends, and there is a kink around <math>10 M_{\odot}</math>. This is mostly due to changes in opacity as a function of temperature. The scatter within segments of constant slope is caused primarily by age and metallicity. Open circles represent giant stars, which have recently evolved off the MS (see § A.2.2). Note that the mass range is biased because of the difficulty detecting both low mass stars, which are very faint, and very high mass stars, which live very short lives. Low mass stars dominate the Universe, both because the ZAMS mass distributions of star clusters are heavily weighted toward the low mass end and because these stars live much longer. Reproduced from Torres et al. [2010]. . . . .</p>	5
2	<p><math>1.0 M_{\odot}</math> profiles for NoDM (grey) and <math>\Gamma_B = 10^6</math> (dark blue) models. Each set of 3 panels shows stellar profiles of the stars at different evolutionary phases indicated by the fraction of hydrogen in the center, <math>X_c</math>, which decreases as the star evolves (ZAMS is "zero-age main sequence"). The profiles in each panel are: 1) <math>\epsilon_{\text{nuc}}</math>, the nuclear burning rate in [erg/g/s]; 2) <math>\epsilon_{\text{DM}}</math>, the rate at which DM transports energy (negative values indicate that energy is being removed), also in [erg/g/s]; 3) <math>\log(T/\text{K})</math>, <math>\log_{10}</math> of the temperature in [K]. ADM energy transport decreases the temperature and burning rate in the center and increases them in a shell at <math>m(&lt; r) \sim 0.1 M_{\odot}</math>. . . . .</p>	27

- 3 Same as Figure 2 for  $3.5 M_{\odot}$  models, except that the 3rd panel in each set shows  $\log(D)$ , where  $D$  is the diffusion coefficient for convective mixing in  $[\text{cm}^2/\text{s}]$ . In the NoDM model the convective core retreats *toward* the center over time, and the burning rate peaks at the center until the end of the main sequence when the burning rate drops dramatically and a shell of strong burning appears suddenly. In the  $\Gamma_B = 10^6$  model, convection at the very center shuts off relatively early in the MS and a convective shell retreats *away* from the center over time. The peak burning rate shifts gradually outward, following the inner edge of the convective shell. The  $\Gamma_B = 10^6$  model reaches the  $X_c$  evolutionary markers at younger ages, relative to the NoDM model, since convection cannot replenish the fuel at the center. . . . . 30
- 4 Changes in MS lifetimes, relative to a star of the same mass with no dark matter, seen in our simulations. Diamonds mark the transition from radiative to convective cores (left to right). For the purposes of this figure this is defined as the lowest  $M_{\star}$  for which the average (over the MS) mass of the convective core is greater than  $0.01 M_{\text{star}}$ . Stars to the right of the NoDM marker (grey diamond) have decreased lifetimes due to a reduction in the size of the convective core, which reduces the amount of hydrogen available for burning. The effect abruptly disappears as stellar lifetimes become shorter than the time required to build up a sufficient amount of ADM. Stars to the left of this marker show mixed behavior due to the competing effects of decreased central burning rates and higher temperatures around  $m(< r) \gtrsim 0.1 M_{\odot}$  which give the star access to more fuel. In addition to these trends, there are several abrupt dips (e.g., at  $2.4 M_{\odot}$ ) and spikes (e.g., at  $2.55 M_{\odot}$ ). This is due to rotational mixing that turns on part-way through the MS and funnels fresh hydrogen fuel to the center, which increases the lifetimes (spikes). The dips result when the NoDM model exhibits this feature, but the ADM model of the same mass does not. 32

- 5 Stellar evolution tracks, from ZAMS to core helium depletion ( $Y_c = 10^{-3}$ ), of select masses with  $\Gamma_B = 10^4$  (left) and  $\Gamma_B = 10^6$  (right). Tracks of NoDM models of the same mass are overplotted with a higher transparency. The chosen masses highlight some of the most dramatic changes seen in Figure 4. Triangles mark the location where stars leave the MS, which we define here as core hydrogen depletion below  $X_c = 10^{-3}$ . The location of the ZAMS and core hydrogen depletion for NoDM models are plotted as dotted and solid black lines respectively. The spike in the  $X_{c, NoDM} < 10^{-3}$  line near the  $2.5 M_\odot$  track is due to rotational mixing in the  $2.4 M_\odot$  model, which is discussed in § 2.4.3. The main effect of ADM on a star’s surface properties is to move the star through roughly the same sequence of events at a faster or slower pace, causing the offset of the  $X_c = 10^{-3}$  milestone relative to NoDM. High mass stars with sufficient ADM skip the convective hook because ADM shuts off convection in the core. These stars transition into shell burning, and therefore the sub-giant branch, more smoothly, similar to low mass stars. . . . . 35
- 6 Same as Figure 5, except that here we plot effective temperature as a function of stellar age. The sub-giant branch phase in a star’s evolution is seen here as a sharp drop in  $T_{\text{eff}}$ , and is a result of structural changes in the star that are triggered by large reductions in the core burning rate due to hydrogen depletion at the end of the MS. The duration of the transition period between the MS and the sub-giant branch is seen here as the temporal difference between the locations of  $X_c = 10^{-3}$  (triangles) and the drop in  $T_{\text{eff}}$ . ADM alters the duration of the transition period, tending to increase it in high mass stars and decrease it in low mass stars. This is opposite of ADM’s effect on MS lifetimes. The net effect is that the feature in  $T_{\text{eff}}$  always occurs either concurrently or earlier in the ADM model than in its NoDM counterpart. . . . . 36

- 7 Isochrones for  $\Gamma_B = 10^4$  (left) and  $\Gamma_B = 10^6$  (right) models, marked by circles, with NoDM models overplotted at higher transparency and marked by crosses. The data points are interpolated from the stellar tracks to a common set of times (isochrone ages). We connect the interpolated data points of a single isochrone with straight lines to guide the eye. The lowest data point on every isochrone is a  $0.9 M_\odot$  star. As stars leave the main sequence they evolve rapidly, and therefore subsequent phases are less well sampled due to the mass resolution. Isochrone ages have been chosen to maximize the sampling around the MS turnoff and sub-giant branch, and are not the same in each panel. We do not show the giant branches because we do not have enough data points there to be representative. The MS turnoff of isochrones around 1 Gyr happens at a higher effective temperature and a lower luminosity and skips the convective hook. This happens in a wider range of ages in populations that live in richer ADM environments (right panel). The oldest isochrones contain only low mass stars, and the ADM and NoDM populations look very similar except that populations in high  $\Gamma_B$  environments appear slightly older due to their decreased luminosity and temperature. . . . . 40
- 8 Effective temperature of the main-sequence-turnoff star, defined here as the hottest MS star, at a given age. We show only the NoDM,  $\Gamma_B = 10^4$ , and  $\Gamma_B = 10^6$  models to allow the reader to see them more clearly. The MS turnoff temperature of the  $\Gamma_B = 10^4$  models starts to become hotter than in the NoDM isochrones at  $\sim 0.7$  Gyr when stars of  $\sim 2 M_\odot$  begin leaving the MS, while those of the  $\Gamma_B = 10^6$  models become hotter at  $\sim 0.1$  Gyr when stars of  $\sim 3.5 M_\odot$  begin leaving the MS. As we move to older isochrones the effect is reversed, and ADM models have lower effective temperatures at the turnoff. The lines terminate when there are no more stars on the main sequence (the lowest mass in our set of models is  $0.9 M_\odot$ ). For the purposes of this figure, we exclude stars in the convective hook because this feature is not well-resolved in our isochrones (see Fig. 7). . . . . 41

9	Isochrones generated by the MIST code for NoDM (grey) and $\Gamma_B = 10^6$ (dark blue) populations. The lowest mass star in both isochrones has $M_\star = 2.25 M_\odot$ (the interpolation was not successful for lower masses), and we show data through core helium depletion ( $Y_c = 10^{-3}$ ) which corresponds to $M_\star = 3.2 M_\odot$ in both cases. The $\Gamma_B = 10^6$ isochrone skips the convective hook and crosses the sub-giant branch at a lower luminosity, consistent with Figure 7. Additionally, the tip of its red giant branch occurs at a lower luminosity, a trend which appears in most of the $\Gamma_B = 10^6$ MIST isochrones. . . . .	42
10	Pitt-Google broker software as of May 2022. Nodes shaded yellow are data resources with public access rights. The resource names are given in the text. Each module is deployed independently to the indicated compute service, shaded red. Conceptual groupings are outlined and labeled in grey. Modules running on Compute Engine each run on a single VM with specs given in the text. Modules running on Cloud Functions are triggered automatically by messages published to their input stream. Black arrows indicate data flow through the main pipelines and grey is through peripheral modules. . . . .	58
11	Consumer performance ingesting the ZTF alert stream on a typical night (top panel) and at high incoming rates (bottom panel, explained further in the text). The y-axis is the time elapsed between alert creation by ZTF and the consumer’s publish timestamp. The black line is the median latency and is labeled on the right. The marginal histogram at the top of each panel is binned to show the per-minute rate of alerts published by the consumer. Typical latencies are about 0.3 seconds regardless of the incoming rate. The latencies on 1 April 2022 were due to the delay in connecting to the Kafka stream. This panel’s marginal histogram shows that the consumer published alerts at LSST-like rates. . . . .	60

- 12 Same as Fig. 11 but for the Cloud Storage module. The y-axis is the difference between publish times of the input and output streams. This is the same as the previous figure but labeled differently to reflect that the delay is now entirely internal and due to the module’s message delivery and compute time. Median processing times are about 0.6 seconds and are essentially unaffected by the incoming rate (which is given in the top marginal histograms of Fig. 11). . . . 62
- 13 Same as Fig. 12 but for the BigQuery module. At typical ZTF rates the performance is similar to the Cloud Storage module. However, at higher incoming alert rates a Google Cloud limit on the rate of BigQuery API requests is exceeded. This results in Google throttling requests and the module is prevented from completing for a period of time. Once the wait period is over the module successfully processes a relatively small number of alerts before exceeding the limit again. This cycle causes the pattern seen in the figure. The limit does not apply to streaming inserts directly and the broker code can be updated to avoid the offending request and bypass the limit. . . . . 64
- 14 Same as Fig. 12 but for the SuperNNova module. The total number of processed alerts is lower because this module only classifies those that are likely extragalactic transients. At incoming rates typical of ZTF, performance is similar to previous modules. At higher incoming rates (bottom panel) the processing times are longer because this module sends its results to BigQuery, and so it experiences similar issues as seen in Fig. 13. . . . . 66
- 15 Diagram of the envisioned process by which researchers could hook into Pitt-Google’s alert streams and catalogs and use them to hunt for electromagnetic counterparts of gravitational waves, both forward and backward in time. The user module (gold) on the left listens to the GW stream, choses an interesting event, then does the following: deploy the module on the right (injecting the GW data), query historical LSST data (all the way up to the given moment), analyze it, and send results to a user-owned Pub/Sub topic. The module on the right listens to the LSST live stream, analyzes it, and sends results to the same topic. . . . . 72

16	Pitt-Google’s total charges incurred on Google Cloud in 2021, by month and service. This is a screenshot of our billing page and includes all GCP projects, services, and SKUs associated with our billing account. The work represented here includes running and storing the full ZTF alert stream plus all development and testing. The total cost was \$6154.87, and was offset 100% by Google Cloud research credits. The BigQuery charge in September is mostly due to a two-day test of a cross-match module, which queried a BigQuery table for every alert. Our transition from Dataflow to Cloud Functions can be seen in their respective trends, as well as the overall cost difference between the two.	75
17	Gaia isochrones of 32 well-measured star clusters, colored by age. The evolution of several distinctive and useful features can be seen including the slope of the MS, the location of the MS turnoff, and the tip of the red giant branch. Credit: Gaia Collaboration et al., A&A, 616, A10 6, 2018, reproduced with permission © ESO. . . . .	87
18	Gaia isochrone of the M67 cluster overlaid with a model generated by PARSEC [Marigo et al., 2017] for comparison. The model reproduces the MS slope and turnoff location reasonably well. Note that the MS is well-populated underneath the line representing the model; the scatter above the line is caused by the combined light of stars in unresolved binaries. The MS turnoff mass is about 1.28 $M_{\odot}$ [Mowlavi et al., 2012], right near the boundary where cores become convective. The discrepancy in the hook features likely indicates that the model’s treatment of overshooting (discussed in § 1.1.1) needs better calibration, as the quality of Gaia’s data places the feature very precisely [Gaia Collaboration et al., 2018]. Credit: Gaia Collaboration et al., A&A, 616, A10 15, 2018, reproduced with permission © ESO. . . . .	88

## 1.0 INTRODUCTION

The dynamic lives and spectacular deaths of many billions of stars will soon be on display in unprecedented detail. Starting in 2024, the Vera C. Rubin Observatory’s Legacy Survey of Space and Time (LSST) will image the southern night sky repeatedly for 10 years through 6 optical filters. It will cover 1,000 times more time-volume space than previous surveys, thanks to its large étendue and faint limiting magnitudes [LSST Science Collaboration et al., 2009]. In fact, it will surpass current leaders in a large number of categories – from survey volume and speed, to data rates and catalog sizes, to astrophysical types and sample sizes – many by an order of magnitude or more. The array of science it will enable is extremely broad. Goals throughout four major science themes are driving the design of the system: (1) Taking an Inventory of the Solar System; (2) Mapping the Milky Way; (3) Exploring the Transient Optical Sky; and (4) Probing Dark Energy and Dark Matter [LSST Science Collaboration et al., 2009]. The study of stars will contribute significantly to the last three.

This introductory chapter is organized as follows. In § 1.1, I give a brief overview of the lifecycles of stars and then discuss stellar structure and evolution<sup>1</sup> during the main sequence (i.e., core hydrogen burning). In § 1.2, I discuss the potential for certain types of dark matter to be captured by a star and alter its evolution. This is followed by an introduction to LSST’s alert stream in § 1.3, the community alert brokers selected to serve it, and the opportunity for cloud computing. Finally, § 1.4 lays out the scope of the dissertation.

---

<sup>1</sup>In this text, “evolution” refers to the lifecycle of a single star, unless explicitly stated otherwise.

## 1.1 STELLAR EVOLUTION

A star is “born” when a core of dense gas within a collapsing molecular cloud becomes hot enough that the pressure provided by sustained hydrogen fusion balances the inward pull from gravity, stabilizing it against further collapse. The star will spend most of its life in this pseudo-equilibrium state, known as the main sequence (MS), burning<sup>2</sup> hydrogen fuel into helium ash in its core and emitting an approximately constant flux of photons from its surface. When the fuel supply in the core runs out, fusion ends and stops providing pressure support against gravitational collapse. After this, stellar fates diverge. Some stars cycle through periods of core contraction followed by fusion-powered stability, burning successively heavier elements into new ones, each more quickly than the last. Others don’t. None can sustain the cycle forever. When they run out of fuel that they can compress enough to burn, they “die”. Some ultimately find a new balance from quantum mechanics, with their extremely dense cores supported by electron- or neutron degeneracy pressure, radiating their leftover heat into space for eons. Others go out with a bang as supernovae; they partially collapse and then explode, ejecting large amounts of matter and energy into space. Still others may simply continue collapsing, all the way down to black holes. If a star has a companion, and most do, they may be distant or close. If they are close enough to interact (other than gravitationally) the results run the gamut from symbiotic to catastrophic.

*Mass:* The single biggest determinant of a star’s evolutionary path is its mass at “zero-age main sequence” (ZAMS), when the pressure supplied by fusion first balances gravity and the star stabilizes. All stars experience the MS but their stellar structures and luminosities during this time show considerable variation according to their ZAMS mass. Higher masses result in larger gravitational forces and therefore higher temperatures in the core. In turn, this determines the rates and types of nuclear burning that can occur, and the stellar structure required to carry the newly released energy away. All stars experience at least some mass loss through stellar winds in their atmospheres<sup>3</sup>; high mass stars lose considerably more. Stars in successively higher mass ranges can burn successively heavier elements. Lifespan

---

<sup>2</sup>In this text, “burning” always refers to nuclear fusion.

<sup>3</sup>This paragraph discusses evolution as predicted by the ZAMS mass specifically (i.e., not including mass loss), though “ZAMS” is often omitted for brevity.

monotonically decreases with increasing mass. Higher mass stars will contract farther over the course of their lifetimes. Low- and very high mass stars do not explode. Other inherent or environmental characteristics can alter this basic picture, but the ZAMS mass is the dominant driver.

Stellar masses can range from  $0.08 M_{\odot} \lesssim M_{\star} \lesssim 150 M_{\odot}$  (see below), with considerable uncertainty in the high-mass limit. The distribution of ZAMS masses is continuous and monotonically decreases with increasing mass [Kroupa and Jerabkova, 2021]. Low mass stars comprise about 70% of the stars in the galaxy [LSST Science Collaboration et al., 2009], but they are very faint and thus difficult to observe. We see high mass stars and those in giant phases (see below) disproportionately often because they are much brighter than low mass stars on the MS.

*Metallicity:* Chemical composition is (arguably) the next leading cause of variation in stellar evolution. “Metals” (i.e., everything heavier than helium) tend to increase nuclear reaction rates and radiative opacities. Higher fractions of carbon, nitrogen, and oxygen increase reaction rates in the CNO cycle (see § 1.1.1). Low metallicity stars tend to be hotter at the surface [Pols, 2011, Mowlavi et al., 2012]. At the surface, high metallicity tends to increase mass-loss rates, especially in high mass stars [Mowlavi et al., 1998]. Stars in the very early Universe could have been much bigger than today because they had extremely low metallicities.

Chemical compositions are often written as mass fractions, with typical ZAMS values of:  $X = 0.60 - 0.75$  (hydrogen),  $Y = 0.25 - 0.35$  (helium), and  $Z = 0 - 0.05$  (metals). Our sun has a metallicity of  $Z \approx 0.014$ . Roughly,  $Z > 0.02$  is considered high metallicity. Some stars, e.g., in the central bulge of the galaxy have been observed with very high metallicities of  $Z \sim 0.1$  [Mowlavi et al., 1998, Cinquegrana and Karakas, 2022]. Stars born earlier in cosmic time are observed to have lower metallicities, and hydrogen (helium) fractions closer to 0.75 (0.25), consistent with the predictions from Big Bang nucleosynthesis of the early Universe’s chemical makeup. The overwhelming majority of metals in the Universe were *produced* by stars, either through the fusion that supports their lives or more energetic processes when they explode. Thus, the Universe is enriched with metals over time as successive generations of stars create and then release them through stellar winds or explosions.

### 1.1.1 Modeling Main Sequence Stars

We cannot observe stellar interiors, but models of their structure and evolution predict their surface properties quite well over long portions of their lives [see [Pols, 2011](#), [Kippenhahn et al., 2012](#), for comprehensive reviews of stellar evolution]. These models are typically formed from a set of coupled differential equations for the temperature, pressure, radius, luminosity, and chemical composition, with time and a Lagrangian mass coordinate as the independent variables. The mass coordinate tracks the radial direction; spherical symmetry is assumed. The equations can be derived from thermonuclear reaction rates, and the conservation of mass, energy, and momentum. Their solutions typically require expensive numerical calculations.

*Homology relations:* For stars on the MS, solutions to the structure equations are approximately homologous (e.g., two stars with different masses have approximately the same relative mass distribution). This results in the following relations:

$$L \propto M_{\star}^{\sim 3.5} \tag{1.1}$$

$$R \propto M_{\star}^{\sim 0.7} \tag{1.2}$$

$$\tau_{\text{MS}} \propto M_{\star}^{\sim -2.5} \tag{1.3}$$

where  $L$  is the luminosity at the surface,  $R$  is the radius at the surface, and  $\tau_{\text{MS}}$  is the MS lifetime. To remain in equilibrium, the star must radiate as much energy from its surface as is released by nuclear burning in its interior. Thus, the luminosity relation shows that nuclear burning is a strong function of mass, with higher mass stars burning fuel at much faster rates. This results in an inverse relationship between mass and MS lifetime; higher mass stars spend less time on the MS, even though they burn more fuel. A surface temperature,  $T_{\text{eff}}$ , can be defined using the luminosity and radius with a black body approximation, with the surface defined by a radius at, e.g., 1 optical depth. This is usually a good approximation for the photosphere, the deepest layer we can directly observe. While the numbers quoted above are typical, there is observed variation, particularly in different mass ranges. See [Fig. 1](#).

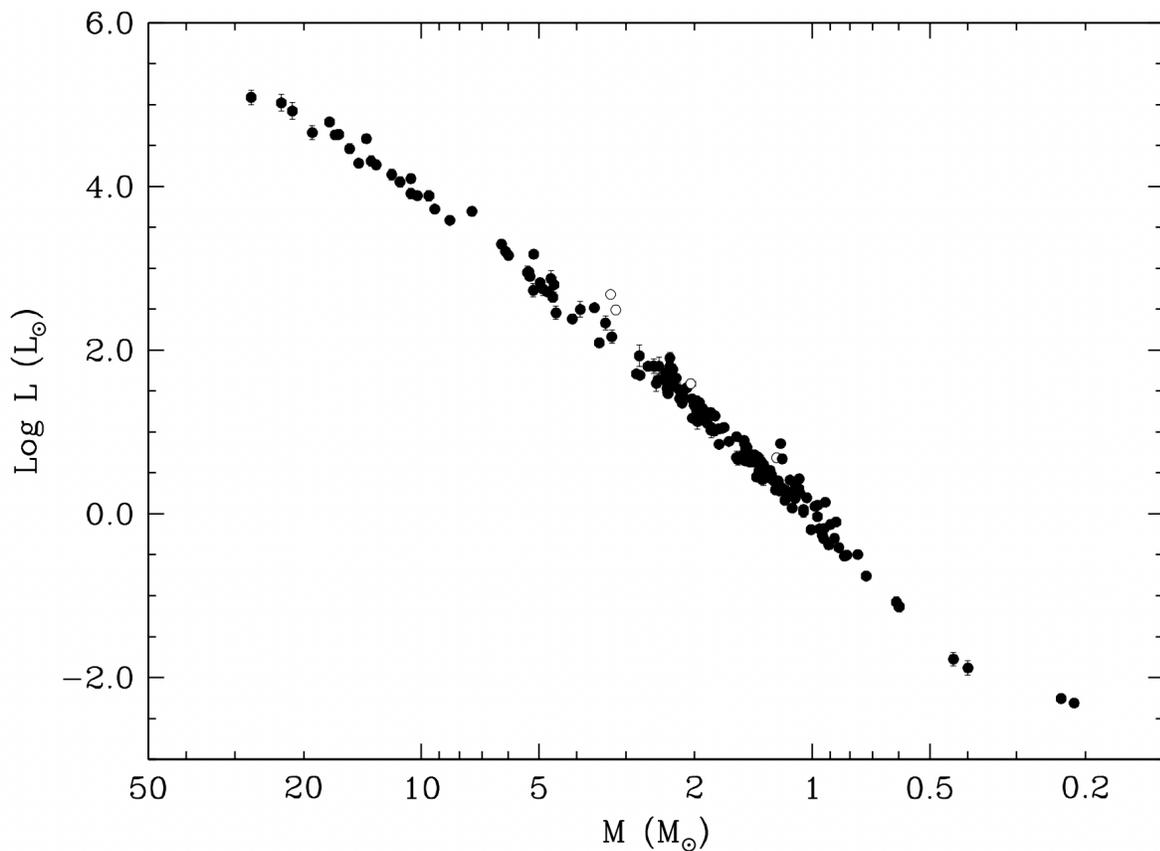


Figure 1: Mass-luminosity relation of 190 stars. These stars are all in binaries or triples, which allows their parameters to be well-measured. For most of the mass range, the slope is well-approximated with  $L \propto M_{\star}^{\sim 3.5}$ , but it becomes noticeably shallower at both ends, and there is a kink around  $10 M_{\odot}$ . This is mostly due to changes in opacity as a function of temperature. The scatter within segments of constant slope is caused primarily by age and metallicity. Open circles represent giant stars, which have recently evolved off the MS (see § A.2.2). Note that the mass range is biased because of the difficulty detecting both low mass stars, which are very faint, and very high mass stars, which live very short lives. Low mass stars dominate the Universe, both because the ZAMS mass distributions of star clusters are heavily weighted toward the low mass end and because these stars live much longer. Reproduced from [Torres et al. \[2010\]](#).

*Basic structure:* The temperature, density, and pressure in a MS star are peaked at the center and (almost always) decrease monotonically. This results in two basic regions: a core undergoing sustained fusion, surrounded by an atmosphere of increasingly thin and tepid gas.

*Burning channels and rates:* There are two main channels through which hydrogen fusion occurs in stars: the pp-chain and the CNO cycle. The pp-chain is a series of reactions whereby four protons fuse to form helium. Stable burning through this channel requires a minimum temperature of  $T_{min} \sim 8 \times 10^6$  K, and the burning rate scales as  $X^2 \rho T^{\sim 4}$  (where  $X$  is hydrogen mass fraction,  $\rho$  is density, and  $T$  is temperature). In the CNO cycle, carbon, nitrogen, and oxygen are used as catalysts in converting hydrogen to helium. Stable burning requires a higher minimum temperature,  $T_{min} \sim 1.5 \times 10^7$ , and the burning rate is a much stronger function of temperature, scaling as  $X Z_{CNO} \rho T^{\sim 18}$  (where  $Z_{CNO}$  is the mass fraction of catalysts). Thus, burning in lower mass stars is dominated by the pp-chain, while the CNO cycle dominates at higher masses. The transition occurs at  $M_{\star} \sim 1.3 M_{\odot}$  (where  $M_{\odot} \approx 1.989 \times 10^{30}$  kg is the mass of the sun) for stars similar in composition to the sun. Stars with lower metallicities contain fewer CNO catalysts, and so their cores must be hotter in order to release the same amount of energy (which they must do to support themselves against gravity). This causes lower metallicity stars to have convective cores at lower masses [Mowlavi et al., 2012]. Stars with masses near the transition point experience significant contribution from both channels, and the CNO contribution increases over time due to rising temperatures.

*Minimum mass:* Since a “star” is defined by the ability to sustain fusion, the pp-chain’s minimum temperature results in a minimum stellar mass of  $\sim 0.08 M_{\odot}$ .

*Energy transport and resulting structure:* The two most important modes of energy transport in a MS star are radiation (carried by photons) and convection (carried by the bulk motion of matter). Conduction (carried by the random motion of individual gas particles) also occurs, but contributes much less because the mean free path of nucleons is much smaller than that of photons. Radiative transport is always occurring. The amount of energy that can be carried away by radiation is proportional to the temperature gradient and inversely proportional to the opacity ( $\kappa$ , the ability of matter to absorb radiation).

If the total, local, energy flux is too large to be carried by radiation alone the region will become unstable. In most regions of the star, this results in convection. The details of convection are relatively uncertain, but the general process is a bulk, cyclic motion: large pockets of material which are slightly hotter than their surroundings rise a significant distance, expand, and cool; slightly colder pockets contract, sink, and absorb heat. This is difficult to model for at least two reasons. First is that it requires solving hydrodynamic equations in three dimensions across large regions of the star where the temperature, etc., change drastically. (Recall that models already require expensive numerical calculations but they assume spherical symmetry.) Second, while the boundary of a convection zone can be defined as the radius at which radiation can carry the full energy flux, rising (or sinking) pockets of gas have inertia. Thus, they can overshoot this boundary before dissolving. The interactions in this region are highly non-linear. Both of these complications are typically handled by parameterizing simpler models. Comparing the results with observations is an active area of research (see § A.1).

While convection is a dynamical instability within the convective zone, the star as a whole remains stable. The result is twofold: a net transport of heat (but not matter) towards the surface and mixing of chemical species between zones. The mixing often results in an influx of fresh hydrogen fuel into the burning region. Thus, a star with a convective core would have access to more fuel than a non-convective star of the same mass and would therefore experience a longer MS lifetime.

Because of the strong temperature dependence of the CNO cycle, the burning in higher mass stars peaks strongly at the center and releases large amounts of energy in a small region. These stars have convective cores surrounded by radiative envelopes, with the core size increasing with mass. The burning rates of lower mass stars are less strongly peaked in the center. However, the relatively low temperatures in their atmospheres result in larger opacities. As a result, lower mass stars tend to have radiative cores surrounded by convective envelopes, with the envelope size increasing with decreasing mass. The transition mass between these two groups is at or just below that of the burning channels ( $M_{\star} \sim 1.3 M_{\odot}$ ), and similarly, the convective cores of stars around this mass grow over time. Below  $\sim 0.35M_{\star}$ , the high opacities dominate and stars are completely convective. Thus, very low mass stars

can burn nearly all of their hydrogen, which extends their MS lifetimes.

*Maximum mass:* Close to the surface, the gas is too thin for energy transport by convection to be efficient (though convective *mixing* is still very efficient here). If the outward photon pressure dominates over gravitational attraction, layers of the atmosphere are blown off into space. This implies that there is a maximum surface luminosity,  $L_{\max} \propto M_{\star}$ , at which stars of a given mass will be stable. If stellar luminosity at the high-mass end continues to scale as  $M_{\star}$  to some power greater than 1, this implies a maximum mass for stars. Stars born above this limit would gain equilibrium by expelling the additional mass. This limit has historically been estimated at  $M_{\star} \approx 100 - 150 M_{\odot}$  for stars with solar composition, but there is evidence for stars of approximately twice this limit [Yusof et al., 2010]. Stars in the early Universe could have been much larger as the low opacities due to extremely low metal content would have impacted both the fragmentation of collapsing molecular clouds and the mass loss rates.

## 1.2 DARK MATTER IN STARS

The evidence for dark matter (DM) is abundant and nearly irrefutable, but its identity and fundamental characteristics remain unknown [e.g., Jungman et al., 1996, Bertone et al., 2005, Battaglieri et al., 2017, Profumo et al., 2019]. Stars in the outskirts of galaxies are observed to orbit much faster than is reasonable if the galaxy’s gravitational potential is provided only by the matter that we can detect from the light emanating from its baryons in the form of stars and gas. This can be seen by comparing the orbital speeds with those predicted by the virial theorem, which relates the kinetic and potential energies of a stable system. The same is true for galaxies themselves, orbiting in clusters. Evidence for DM also appears in the gravitational lensing of distant galaxies and the anisotropy of the microwave background radiation.

While research in recent decades has investigated wide ranging possibilities – from particle- to wave-like, hot to cold, primordial black holes to gravitational solitons – the search has focused on a class of particle candidates known as Weakly-Interacting Massive

Particles (WIMPs) [Jungman et al. \[1996\]](#). However, Earth-based direct detection experiments have failed to provide confirmative evidence and the parameter space (mass and scattering cross section) is rapidly shrinking [[Amole et al., 2017](#)]. This has spawned research on alternatives to the standard WIMP paradigm. Asymmetric dark matter (ADM) is a class of models which has some advantages over classic WIMPs. There are many variations of ADM but most predict masses in the range of  $m_{\text{DM}} \sim 1 - 10$  GeV, lighter than WIMPs and within a region of parameter space not yet ruled out (though it too is becoming increasingly constrained). As the name implies, the relic abundance of ADM is set by a primordial asymmetry with an antiparticle, rather than via thermal freeze-out as predicted for WIMPs.

DM with these properties could have interesting and potentially observable effects on the structures and evolution of stars, which has been recognized since at least the 1980's [[Press and Spergel, 1985](#), [Krauss et al., 1985](#), [Gaisser et al., 1986](#), [Griest and Seckel, 1987](#)]. As stars orbit through the dark matter halos that are home to their host galaxies, DM particles can scatter with stellar nucleons and lose enough energy to become gravitationally bound to the star. The capture rate depends on many properties, as follows (valid for  $m_{\text{DM}} \lesssim 20$  GeV) [[Zentner and Hearin, 2011](#)]:

$$C_{\text{DM}} \propto \left( \frac{\rho_{\text{DM}}}{\bar{v}} \right) \times \left( \frac{\sigma_{\text{n}}}{m_{\text{DM}}} \right) \times \left( M_{\star} v_{\text{esc}}^2 \right) \quad (1.4)$$

The first term in parentheses describes the star's local DM environment (density and velocity dispersion). The capture rates increase with larger ambient DM density and/or smaller velocity dispersion. The second term contains properties of the DM itself (cross section for scattering with nucleons and mass). The third term contains properties of the star (mass and escape velocity). Note that the capture rate scales linearly with stellar mass, a weaker dependence than that of either the luminosity (indicating the amount of energy generated) or the inverse lifetime. Therefore higher mass stars will have a difficult time building up sufficient DM within their lifetimes to affect the star.

Once captured, DM transports energy within the star through continued scattering events. The impact of this additional energy transport depends on the DM mass and cross section, the total amount of DM captured, and stellar properties such as its mass distribution and structure. To have a significant impact, the combination of the DM's mass and cross

section must be balanced, and the star must be able to collect and retain a relatively large number of particles. The DM must be heavy enough not to evaporate, but light enough so that it is not confined to the center of the star (i.e., it must experience an appreciable temperature gradient so that collisions transport energy). For a solar-mass star, this requires DM in the range of  $5 \text{ GeV} \lesssim m_{\text{DM}} \lesssim 20 \text{ GeV}$  [Petraki and Volkas, 2013]. Similarly, its cross section must be large enough to result in a significant number of collisions with stellar nucleons, but small enough so that its mean free path allows it to cross a significant temperature gradient between scattering events. ADM models tend to predict combinations of these parameters that place it within this “goldilocks” zone. In addition, ADM’s asymmetric property means that it can build up in large quantities without annihilating (as thermal-relic WIMPs would) since a negligible number of antiparticles would have survived beyond the early Universe.

An additional important property is whether the scattering interactions require that the stellar nucleons have a net spin. Both types of ADM models are predicted. Spin-dependent ADM in stellar interiors would primarily interact with protons (ionized hydrogen), since these are the only nucleons with a net spin that are present in significant amounts. Thus it would affect the star primarily during the MS, before it burns through its central supply of hydrogen. Spin-independent ADM would interact with all stellar nucleons. However the cross sections for this type of DM have been constrained by experiments to smaller values, resulting in relatively few interaction with protons. Thus its effect on the star would likely be insignificant until the end of the MS when the central helium fraction increases substantially.

Significant energy transport by DM can affect a star in many ways as it attempts to maintain or regain equilibrium. The immediate effect of DM removing (depositing) energy in a particular region will be a decrease (increase) in the local temperature. As a result, the local burning rate may be reduced (increased). This can result in less (more) radiation pressure, and so the local density of the star will likely adjust. If the energy transported by DM is sufficiently high it can obviate other energy-transport mechanisms such as convection, thus altering the star’s core structure. These changes are interrelated and complex and can propagate through the star’s outer layers, even if the DM itself does not reach that far. The combined result may include changes in the star’s surface luminosity and/or effective temperature, total amount of fuel burned, time spent on the main sequence, or even the

evolutionary phases it experiences (e.g., due to structural changes).

### 1.3 ALERT BROKERS

The size of LSST’s nightly alert stream will be equivalent to about 650 hours of HD video:  $\sim 800$  GB/night. This is far too much for every interested researcher to sort through on their own. This stream will announce of order  $10^7$  events each night, rushing by at an average rate of  $\sim 300$  alerts per second. Their sources will be very diverse: flare stars, CVs, Cepheids, merging neutron stars, thermonuclear and core collapse supernovae alongside dozens more manifestations of the dynamic lives of stars and a large number of non-stellar objects like near-Earth asteroids and the powerful quasars and blazars of distant active galactic nuclei. These are the subjects of wide-ranging research areas. The LSST Science Book says it best: “*a single data set [will serve] most science programs* (instead of science-specific surveys executed in series). One can think of this as *massively parallel astrophysics*.”<sup>4</sup> [LSST Science Collaboration et al., 2009]. However, significant challenges await in the parallelization, distribution, and analysis of this unprecedented dataset. LSST’s many strengths will double as sources of difficulty, both technical and scientific. Diverse efforts at many levels will be required to fully actualize LSST’s potential, and alert brokers will play a central role. Brokers will serve as the primary channels for LSST’s alert stream with the jobs of enriching and curating the full dataset into something digestible by individuals & collaborations, and of facilitating connections with external datasets and follow-up facilities. (See Appendix B for a more complete description of the role of brokers and the services they offer.)

LSST has selected 7 community alert brokers to whom they will deliver the full stream. They are ALerCE<sup>5</sup>, AMPEL<sup>6</sup>, ANTARES<sup>7</sup>, BABAMUL, Fink<sup>8</sup>, Lasair<sup>9</sup>, and Pitt-Google<sup>10</sup>.

---

<sup>4</sup>While this was said in reference to LSST’s data on the whole, a significant fraction will appear in the alert stream and it is my opinion that the sentiment applies to it as well. In further support of my usage here, alert brokers aggregate datasets and thus are likely to facilitate access to LSST’s catalogs as well.

<sup>5</sup><http://alerce.science/>

<sup>6</sup><https://ampelproject.github.io/>

<sup>7</sup><https://antares.noirlab.edu/>

<sup>8</sup><https://fink-broker.org/>

<sup>9</sup><https://lasair.roe.ac.uk/>

<sup>10</sup><https://pitt-broker.readthedocs.io/en/latest/>

These brokers are currently using the Zwicky Transient Facility (ZTF<sup>11</sup>) alert stream to prepare for LSST and in the process are demonstrating their unique strengths and approaches. I have been the lead developer of the Pitt-Google broker since Fall 2020, including taking the lead in writing our successful LSST proposal. My work on the broker is the subject of Chapter 3. We are using a cloud-based model that is unique among brokers and provides access in new ways.

The large amount of data LSST will provide in relatively unexplored regions of parameter space will necessitate rapid analysis and quick evolution of models and classifiers if we are to best take advantage, especially early on. The alert stream can provide the data to support this. And if we are to efficiently identify, e.g., M-dwarf flares, rare kilonovae amongst common CVs, and pure samples of Type Ia SNe, there will be plenty of room for many individuals with wide-ranging interests to test and evolve their own ideas. This will be a community effort and much preparation work is already taking place. However, most researchers interested in alert data are not accustomed to analyzing the streams directly, even though alert streams are now a major data product of astronomical surveys. Instead, many opt to receive a relatively small number of personal notifications and/or query catalogs & download the data at a later time for analysis.

Cloud computing can provide an ecosystem that changes this. Cloud-managed message delivery and compute services can allow individuals to accept, process, and store small or large streams in a very lightweight way without managing the client connection or compute resources, only their own analysis code. This has the potential to open up access to LSST's alert stream to a much broader range of people, across skill sets and access to resources. In turn this can enable individuals to deploy and evolve their own algorithms quickly, on live data streams, and to collaborate within and beyond sub-fields. This is the pace that will be required for the community to best realize the revolution that LSST will soon broadcast. In the past, the biggest technical barrier preventing community adoption of this processing method has been that the streams have not been provided in a conducive format, either by surveys or by brokers. While many brokers do operate in the cloud, either partially or fully, all except Pitt-Google do so using a more traditional model which does not expose cloud

---

<sup>11</sup><https://www.ztf.caltech.edu/>

services or facilitate their use by the community.

Pitt-Google is the only LSST broker to adopt an explicitly cloud-based model. As a result of my work, we are offering alert streams in a format that is unique in the space and with features that make them conducive to easy and direct analysis by students, astronomy professionals, and the public. Further, our pipeline demonstrates the efficient use of cloud services in a way that is easily replicable at both large and small scales, which can ease the adoption of this new approach by the community. In addition to live stream processing, we are developing and providing the usual services that are common to alert brokers, but doing so via Cloud-based methods which offer unique benefits.

## 1.4 SCOPE OF THE DISSERTATION

*Asymmetric dark matter in stars:* In Chapter 2 I model stellar evolution under the effect of ADM and show that its signature may be observable in the isochrones of star clusters. Specifically, I consider spin-dependent ADM with a fixed mass and scattering cross section and study its effect on sets of stars with ZAMS masses between  $0.9 \leq M_\star/M_\odot \leq 5.0$  in different ADM environments, covering the range of values expected to exist. With the chosen mass and cross section, ADM orbits have a characteristic radius of  $r_{\text{DM}} \sim 0.1R_\star$  and a mean free path of  $l_{\text{DM}} = (\sigma_p n_p)^{-1} \sim 1R_\star$ . Thus the ADM experiences substantial changes in local ambient temperature and completes several orbits between scattering events, resulting in significant energy transport once a sufficient amount of ADM has been captured. I wrote a module which calculates the amount of ADM captured and the resulting energy transport in the stellar interior at a given time. It is a plugin for MESA [Modules for Experiments in Stellar Astrophysics, Paxton et al., 2011, 2013, 2015, 2018, 2019] which I used to model the stellar evolution. I find that stars in ADM-rich environments can experience changes in their burning rates and core structures that are sufficient to alter their surface properties and MS lifetime. Finally, I produced isochrones for each set of stars. These show that ADM may alter isochrone features that are commonly used to measure distances and evaluate stellar evolution models and consider the implications for future observations. This chapter first

appeared as [Raen et al. \[2021\]](#).

*Pitt-Google alert broker:* In Chapter 3 I discuss my work on the Pitt-Google alert broker. Our broker was motivated by the desire to provide broad public access to the LSST alert stream and to facilitate scientific collaboration through a cloud-based solution that utilizes Google’s computational expertise. I have been the lead developer since Fall 2020 and a core member since early 2019. I led the writing of our successful proposal to become a full-stream LSST community alert broker; we will serve as one of the primary channels through which the community will access the nightly stream for the coming decade. I have been developing our broker in preparation, including a stream consumer, processing (e.g., categorizing/classifying) and storage pipelines, stream curation and publishing, a catalog with SQL access, and user access tools and tutorials. We have been processing the Zwicky Transient Facility (ZTF) alert stream consistently since late 2020 after I redesigned the broker’s core components. I have structured our pipeline to provide open access to the data at every level of processing along with features that facilitate custom filtering and efficient processing methods. I have either implemented solutions for, or done significant preparatory work towards, nearly all services typically provided by alert brokers.

I have been exposed to the diversity of stellar (and other) activity that will show up in the alert stream, as well as some complications that will arise in their identification and subsequent analysis. While there are many aspects to our broker, my focus in the previous section on enabling individuals to directly process the alert stream in real time is a reflection of my belief that this scale and rate of analysis can revolutionize the science, and that Pitt-Google can facilitate a unique and useful route.

My work touches a broad range of technical and scientific areas, and I document some of these in the appendixes. In Appendix A I outline some of the opportunities LSST will provide for studying stars. [A.1](#) discusses isochrones, which give us insight into stellar structure and access to the long timescales needed to study their full lifecycles. [A.2](#) details some of the dynamic stellar phases and events that will appear in the alert stream, as well as some complications in identifying them. This motivates my belief in the utility of rapid and widespread analysis, and my efforts to facilitate it. Appendix B describes alert brokers, outlining the features and services they typically provide. Appendix C discusses Bayesian

networks that I have proposed for both large-scale classification of the alert stream and targeted searches for exotic transients.

## 2.0 THE EFFECTS OF ASYMMETRIC DARK MATTER ON STELLAR EVOLUTION I: SPIN-DEPENDENT SCATTERING

### 2.1 INTRODUCTION

A preponderance of the evidence suggests that approximately 84% of the matter budget of the universe consists of a form of non-baryonic dark matter that has yet to be identified [e.g., [Jungman et al., 1996](#), [Bertone et al., 2005](#), [Battaglieri et al., 2017](#), [Profumo et al., 2019](#)]. In the standard picture of cosmological structure formation, galaxies form within the potential wells of large, nearly virialized halos of dark matter [[White and Rees, 1978](#), [Blumenthal et al., 1984](#)]. If the dark matter interacts with standard model particles, it can be captured by stars moving through dark matter halos [[Press and Spergel, 1985](#), [Krauss et al., 1985](#), [Gaisser et al., 1986](#), [Griest and Seckel, 1987](#)]. Once captured, continued scattering within the stellar interior contributes to energy transport within the star, potentially altering its evolution [[Spergel and Press, 1985](#), [Taoso et al., 2010](#), [Zentner and Hearin, 2011](#), [Iocco et al., 2012](#), [Lopes and Silk, 2012](#), [Casanelas and Lopes, 2013](#), [Casanelas et al., 2015](#), [Vincent et al., 2015](#), [Murase and Shoemaker, 2016](#), [Lopes and Silk, 2019](#), [Vincent, 2020](#)]. The significance of this energy transport depends on the following properties (in addition to the properties of the star): (1) the DM mass,  $m_{\text{DM}}$ ; (2) the DM-nucleon scattering cross section,  $\sigma_{\text{n}}$ ; and (3) the total number of DM particles captured by a star,  $N_{\text{DM}}$ , which itself depends on  $m_{\text{DM}}$  and  $\sigma_{\text{n}}$  as well as the local DM environment from which the particles are captured (see § 2.2). We study the effects of energy transport by asymmetric dark matter (ADM, see below) in stars of mass  $0.9 \leq M_{\star}/M_{\odot} \leq 5.0$  living within a variety of dark matter environments using the publicly-available code Modules for Experiments in Stellar Astrophysics [MESA, [Paxton et al., 2011, 2013, 2015, 2018, 2019](#)].

Evidence supporting the claim that  $\sim 84\%$  of the matter in the universe is in an unknown form of dark matter is abundant and varied, ranging from the anisotropy of the microwave background radiation to formation and structures of galaxies [e.g., [Jungman et al., 1996](#), [Bertone et al., 2005](#), [Planck Collaboration et al., 2018](#)]. For several decades, the leading candidate has been the so called Weakly-interacting massive particle (WIMP). The classic WIMP is a heavy ( $m_{\text{DM}} \sim 10^2 - 10^3$  GeV) thermal relic whose contemporary abundance is set by its annihilation rate in the early universe [e.g., [Kolb and Turner, 1990](#)]. Therefore, WIMPs are thought to have a fairly well established “standard” annihilation cross section [e.g., [Steigman et al., 2012](#)], which is comparable to typical weak-scale cross sections,  $\langle\sigma v\rangle \sim 10^{-26}$  cm<sup>3</sup>/s. This annihilation of WIMPs, which is so critical to guaranteeing that the correct abundance of dark matter in the contemporary Universe, in turn, limits the number of particles that can accumulate within a star. The rate of capture of new dark matter particles comes to equilibrium with dark matter particle annihilation in the stellar interior [[Krauss et al., 1985](#)]. Despite numerous ongoing terrestrial direct detection experiments [see [Schumann, 2019](#), for a review] and efforts to detect dark matter indirectly through its annihilation products [reviewed in [Slatyer, 2017](#)], dark matter has not been observed non-gravitationally. The available parameter space for relatively light ( $m_{\text{DM}} \lesssim 10^2$  GeV) DM is rapidly shrinking, which has triggered a surge in research into alternatives to the long-favored WIMP.

Asymmetric dark matter (ADM) is an alternative to the classic WIMP in which the relic abundance of the dark matter particle is set by a primordial asymmetry rather than via annihilation [for a review, see [Zurek, 2014](#), [Petraki and Volkas, 2013](#), and references therein]. If the baryon and dark matter asymmetries are related, then such models have the appealing property that they explain the fact that the contemporary dark matter and baryon abundances are of the same order of magnitude, which is otherwise surprising because these relic abundances are determined by unrelated physics in the WIMP scenario. Indeed, this was one of the early motivations for ADM-like models [e.g., [Nussinov, 1985](#), [Barr et al., 1990](#), [Chivukula and Walker, 1990](#), [Kaplan, 1992](#)]. The variety of specific incarnations of ADM is broad, but ADM models often predict particle masses smaller than the classic WIMP ( $m_{\text{DM}} \sim 1 - 10$  GeV) and little or no contemporary dark matter annihilation for lack of relic

dark matter anti-particles.

These predictions motivate studies to constrain ADM indirectly through stellar astrophysics. The lack of annihilation means that ADM may build up to very large quantities within stars because the capture of ADM is never countered by annihilation. Meanwhile, the relatively low masses compared to the classic WIMP mean captured ADM particles orbit within a significant volume of the star, out to  $r_{\text{DM}} \sim 0.1R_{\odot}$  for a Sun-like star, which means that they experience large differences in ambient temperature throughout their orbits and can thus transport energy outward from the stellar core extremely efficiently [Spergel and Press, 1985]. These features of ADM have already motivated research into the possibility that ADM may alter stellar evolution [e.g. Taoso et al., 2010, Zentner and Hearin, 2011, Iocco et al., 2012, Lopes and Silk, 2012, Casanellas and Lopes, 2013, Casanellas et al., 2015, Vincent et al., 2015, Murase and Shoemaker, 2016, Lopes and Silk, 2019, Lopes and Lopes, 2019, Vincent, 2020]. Our results are generally in agreement with previous works, insofar as they can be compared considering variations in the chosen parameters of each study which can include ADM properties (e.g., mass and cross section), halo environments (e.g., ADM density and velocity dispersion in the stellar neighborhood), and stellar mass. In this paper we undertake a study of the properties and evolution of stellar populations within halos of ADM. We fix the ADM mass and cross-section and study the effects of varying halo environments across a wider range of parameter space than has been done previously. (See § 2.2 for a discussion of ADM properties and environments.) The effects of stellar cooling are particularly large in environments in which the ambient dark matter density is high and velocity dispersion is low, such that the capture of dark matter is extremely efficient. Thus, these effects will be largest in dwarf satellite galaxies and high-redshift galaxies. In this first paper on the topic, we further narrow our study to spin-dependent ADM-nucleus scattering. Spin-independent ADM-nucleus scattering leads to behaviors that are qualitatively distinct from spin-dependent scattering; therefore, we will present results for the former case in a forthcoming manuscript.

In the following section, we summarize the dependence of the capture rate of dark matter within stars on both dark matter and stellar properties. In § 2.3, we describe our simulations of stellar evolution including cooling due to ADM. We present our results in § 2.4. We discuss

our results and draw conclusions in § 2.5.

## 2.2 DARK MATTER PROPERTIES AND CAPTURE IN STARS

Probing the parameter space of ADM with simulations of stellar evolution is computationally expensive. Consequently, we show results for an illustrative set of ADM parameters that we initially chose in order to: (1) make the effects of ADM on stellar evolution significant; (2) narrowly-evade the evaporation threshold; and (3) remain consistent with contemporary constraints on dark matter properties (but see the discussion in the next paragraph). We choose  $m_{\text{DM}} = 5 \text{ GeV}$  and a spin-dependent dark matter–nucleon scattering cross section of  $\sigma_{\text{p}} = 10^{-37} \text{ cm}^2$ . Hereafter, we will discuss ADM–proton scattering since protons are the only nuclei in MS stars with both a significant abundance and a net spin. With these parameters, ADM evaporation is negligible. We find that the largest evaporation mass in our models is  $\simeq 4.6 \text{ GeV}$  at a solar mass of  $\simeq 1.4 M_{\odot}$ , consistent with the literature on this topic (see, for example, the classic papers by [Gould 1987](#) and [Gould 1990](#) and the recent update by [Busoni et al. 2017](#)). We note that evaporation masses are slightly higher for spin-independent scattering, because the helium nucleus is more closely matched to the mass of the dark matter particle, so this must be considered in any extension of this work to spin-independent interactions. We assume that ADM self-interactions are negligible throughout; however, it is likely that self-interactions, if they existed, would lead to enhanced cooling [e.g., [Zentner, 2009](#)]. Exploring such models would constitute a potentially interesting follow-up to this work.

During the course of this work, the PICO collaboration was able to reduce its thresholds, unambiguously excluding dark matter with the specific parameter values listed above [[Amole et al., 2019](#)]. While the simulations that we have carried out are computationally intensive and it is impractical to repeat each of the  $\sim 600$  stellar evolution simulations, we believe our results are a meaningful exploration of the deviations in stellar evolution due to ADM and our energy transport module provides a basis for further exploration in future work. Consequently, we choose to present these results as a qualitative indication of the

effects that ADM can have upon stellar evolution. As we will see below, there is significant uncertainty involved in associating a particular stellar effect with a particular dark matter-nucleon scattering cross section due to a variety of model uncertainties. Consequently, while the effects of dark matter within stars will likely be milder than those that we describe here, it is possible that the same, qualitative effects could be realized in nature.

The amount of energy transported by dark matter (see § 2.3.1) is proportional to the amount of ADM within the star. In ADM models, in which annihilation of dark matter is negligible, the number of dark matter particles within the star at any given time,  $t$ , is determined by  $dN_{\text{DM}}/dt = C_{\text{DM}}$ , where  $C_{\text{DM}}$  is the instantaneous ADM capture rate. We use the capture rate from [Zentner and Hearin \[2011\]](#), which is a simplified form valid for dark matter particle masses  $m_{\text{DM}} \lesssim 20$  GeV [see [Gould, 1992](#), [Zentner, 2009](#), for more complete capture rates]:

$$\begin{aligned}
C_{\text{DM}} = C_{\odot} & \left( \frac{\rho_{\text{DM}}}{0.4 \text{ GeV cm}^{-3}} \right) \left( \frac{270 \text{ km s}^{-1}}{\bar{v}} \right) \\
& \times \left( \frac{\sigma_{\text{p}}}{10^{-43} \text{ cm}^2} \right) \left( \frac{5 \text{ GeV}}{m_{\text{DM}}} \right) \\
& \times \left( \frac{v_{\text{esc}}}{618 \text{ km s}^{-1}} \right)^2 \left( \frac{M_{\star}}{M_{\odot}} \right)
\end{aligned} \tag{2.1}$$

where  $C_{\odot} = 5 \times 10^{21} \text{ s}^{-1}$ ,  $\rho_{\text{DM}}$  is the DM density in the stellar environment,  $\bar{v}$  is the velocity dispersion of dark matter particles in the stellar neighborhood, and  $v_{\text{esc}}$  is the escape speed from the surface of the star.

The first line of Eq. (2.1) gives the dependence of the capture rate on the stellar environment. Both  $\rho_{\text{DM}}$  and  $\bar{v}$  are properties of the local stellar environment and are degenerate with one another in Eq. (2.1). A higher ambient density of dark matter leads to a higher rate of dark matter capture, while a lower relative velocity between the star and the infalling dark matter leads to a higher probability for capture. Because of this degeneracy, coupled with the fact that these parameters carry considerable uncertainty themselves, it is convenient to parameterize a star's local dark matter environment by an overall factor [[Zentner and Hearin, 2011](#), [Hurst et al., 2015](#)],

$$\Gamma_{\text{B}} = \left( \frac{\rho_{\text{DM}}}{0.4 \text{ GeV cm}^{-3}} \right) \left( \frac{270 \text{ km s}^{-1}}{\bar{v}} \right). \tag{2.2}$$

Normalized in this way,  $\Gamma_B$  specifies the capture rate,  $C_{DM}$ , relative to the rate that would be realized in the solar neighborhood for the same star. From this point on, we will characterize a star’s dark matter environment using  $\Gamma_B$ . In general, we will be most interested in values of  $\Gamma_B > 1$ . A value of  $\Gamma_B = 0$  describes a stellar environment with no dark matter (hereafter referred to as ‘standard models’ and labeled ‘NoDM’), and  $\Gamma_B = 1$  describes the solar neighborhood. A value of  $\Gamma_B = 10^2$  may specify an environment in which the dark matter density is 100 times that in the solar neighborhood at the same velocity dispersion, an environment in which the velocity dispersion is 1/100 that of the solar neighborhood at the same density, or any of an infinite number of other possible combinations.

It is interesting to consider the range of  $\Gamma_B$  values that would be considered reasonable. If the distribution of dark matter within galaxies, such as the Milky Way, follows a profile that diverges as the [Navarro et al. \[1997, NFW\]](#) density profile, then one might expect to find dark matter environments near the centers of galaxies with densities significantly higher than the local value and velocity dispersions significantly lower than the local value, giving  $\Gamma_B \gg 1$ . Such scenarios were explored in [Bertone and Fairbairn \[2008\]](#), [Fairbairn et al. \[2008\]](#), and [Scott et al. \[2009\]](#). While such large values of  $\Gamma_B$  may well lead to large effects on stellar structure, stellar populations near the Galactic Center are difficult to observe and any assumption about the dark matter density profile in the inner regions of any galaxy must be considered speculative. Interestingly, Local Group dwarf galaxies are extremely dark matter-dominated and have well-constrained dark matter profiles and velocity dispersions. In some cases, the Local Group dwarfs have densities  $\sim 3$  orders of magnitude higher than the dark matter density in the Solar neighborhood and velocity dispersions that are at least  $\sim 2$  orders of magnitude smaller than the local value. This suggests that values of  $\Gamma_B \sim 10^{5-6}$  could be realized within Local Group dwarf galaxies and has the further merit that  $\Gamma_B$  within Local Group dwarfs can be measured more precisely in the future. A third possibility for large values of the environmental factor are early-forming, very high-redshift galaxies. These galaxies begin forming in small, dense halos where the environmental boost factor can reach  $\Gamma_B \gtrsim 10^6$  at redshifts  $z \gtrsim 10$  [[Koushiappas et al., 2004](#)]. Of course, these stars will not be directly observable, but it is interesting to speculate that such stars could be detected as remnants of early mergers with the proto-Milky Way and/or that changes to the structure

and evolution of these stars could be detected indirectly in the chemical evolution of the larger, lower-redshift galaxies in which they will be found today.

Finally, while we have focused on the environmental parameter,  $\Gamma_{\text{B}}$ , as a proxy for the dark matter environment in which a star is embedded, we note that values of  $\Gamma_{\text{B}} \neq 1$  can also be mimicked through dark matter physics. In particular, dark matter self-interactions can greatly enhance the capture rates of dark matter within stars [Zentner, 2009]. This effect of dark matter self-capture itself grows with increasing ambient density and decreasing ambient velocity dispersion, so the two effects reinforce one another. For example, a value of  $\Gamma_{\text{B}} \sim 10^4$  may be realized by increasing the ambient dark matter density by a factor of  $\sim 10^3$ , while simultaneously introducing a dark matter self-interaction that boosts the number of captured dark matter particles by a factor of  $\sim 10$ . We relegate the separation of these effects to future work and encapsulate all of this uncertainty into the single parameter  $\Gamma_{\text{B}}$ .

### 2.3 METHODS

We study the impact of dark matter on the evolution of  $0.9 \leq M_{\star}/M_{\odot} \leq 5.0$  stars (with a mass step of  $0.05 M_{\odot}$ ) through core helium depletion ( $Y_c = 10^{-3}$ ) or a maximum age of 10 Gyr, whichever comes first, using the publicly-available code Modules for Experiments in Stellar Astrophysics [MESA, Paxton et al., 2011, 2013, 2015, 2018, 2019], release 12115. We used the MESA SDK version 20190830<sup>1</sup> to compile MESA. We base our stellar parameter inlist on the MESA test suite `1M_pre_ms_to_wd` inlists and use a metallicity of  $Z = 0.0142$ . Twelve of the models we ran did not complete (e.g., due to requiring unreasonably small timesteps), and we have excluded them from our final data set. Of these, 2 were  $\Gamma_{\text{B}} = 10^4$  models, and none of them were either NoDM or  $\Gamma_{\text{B}} = 10^6$  models (these are the three  $\Gamma_{\text{B}}$  values we highlight below). We wrote a module that calculates dark matter capture and energy transport (see § 2.3.1) and connects to MESA simulations via the provided `extra_energy_implicit` hook. For the reader interested in either examining the underlying

---

<sup>1</sup><https://zenodo.org/record/3560834>

data or reproducing our results, we make the following MESA input/output files available on Zenodo<sup>2</sup> and through the MESA Marketplace<sup>3</sup>: 1) inlist templates and `src` files; 2)  $1.0 M_{\odot}$  and  $3.5 M_{\odot}$  data for NoDM,  $\Gamma_B = 10^4$ , and  $\Gamma_B = 10^6$  models. Additional data will be shared on reasonable request to the corresponding author. For the reader interested in utilizing our module to explore the effects of varying parameters beyond the scope of this paper, our code is available on GitHub<sup>4</sup>.

We take advantage of the significantly improved numerical energy conservation capabilities in recent MESA versions [introduced in Paxton et al., 2019] as we find this to be crucial to a proper accounting of the effects of ADM. Energy transport by significant amounts of ADM alters the core structure of a star such that small changes in the temperature profile due to poor energy conservation can lead to a reversal in the direction of ADM energy transport which destabilizes the star. See § 2.4.1 for further discussion.

To generate isochrones (see § 2.4.5) from MESA’s stellar models we perform a linear interpolation of our stellar tracks to a uniform set of ages and choose isochrones that are as well-sampled as possible in regions of interest. MESA’s adaptive time steps resolve dynamic phases of evolution quite well, and this interpolation is not problematic. We also show one isochrone generated from the MIST code [MESA Isochrones and Stellar Tracks, Dotter, 2016, Choi et al., 2016], which uses a multi-step process to interpolate both the stellar tracks and the mass grid and therefore resolves otherwise sparsely populated regions of the isochrones. We are unable to show more MIST isochrones because the interpolation failed in most of our parameter space; most notably, it did not produce any isochrones older than 1 Gyr. This is likely due to non-monotonicity in the mass-age relation at fixed evolutionary phase (see § 2.4.3) which violates assumptions of the code.

### 2.3.1 Energy Transport by Dark Matter

The energy transported by captured ADM can, in principle, be computed by solving the Boltzmann equation; however, this strategy is too computationally intensive to combine with

---

<sup>2</sup><https://zenodo.org/record/4064115>

<sup>3</sup>[http://cococubed.asu.edu/mesa\\_market/add-ons.html](http://cococubed.asu.edu/mesa_market/add-ons.html)

<sup>4</sup><https://github.com/troyraen/DM-in-Stars>

a full-scale simulation of the evolution of stellar structure. To reduce the computational costs of our simulations, we estimate ADM energy transport using the approximations of [Spergel and Press \[1985\]](#). In particular, we assume a Maxwellian phase-space distribution for the ADM and calculate an orbit-averaged temperature,  $T_{\text{DM}}$ , by requiring that the distribution satisfy the first moment of the Boltzmann equation. This amounts to a requirement on energy conservation: ADM should neither inject nor remove a net energy from the star. The rate of energy transfer (per unit mass) from dark matter to protons is then

$$\begin{aligned} \epsilon_{\text{DM}}(r) = & 8\sqrt{\frac{2}{\pi}} \frac{n_{\text{DM}}(r)n_{\text{p}}(r)}{\rho(r)} \frac{m_{\text{DM}}m_{\text{p}}}{(m_{\text{DM}} + m_{\text{p}})^2} \sigma_{\text{p}} \\ & \times \left( \frac{m_{\text{p}}kT_{\text{DM}} + m_{\text{DM}}kT(r)}{m_{\text{DM}}m_{\text{p}}} \right)^{1/2} k[T_{\text{DM}} - T(r)], \end{aligned} \quad (2.3)$$

where  $n(r)$  is a number density,  $\rho(r)$  is the mass density,  $k$  is Boltzmann's constant, and the subscript p refers to protons. [See [Spergel and Press, 1985](#), for a detailed derivation].

Generally,  $n_{\text{p}}$ ,  $n_{\text{DM}}$ , and  $T$  all peak at the center, so the energy transport is most efficient here. The number density of dark matter particles,  $n_{\text{DM}}$  increases in proportion with  $N_{\text{DM}}$ , so we can expect the effects to increase with both  $\Gamma_{\text{B}}$  and stellar age through the MS, while hydrogen is abundant. As a star leaves the MS,  $n_{\text{p}}$  drops in the core and spin-dependent ADM energy transport is greatly diminished because there are relatively few protons left with which dark matter may scatter<sup>5</sup>.

The sign of  $\epsilon_{\text{DM}}(r)$  is given by the final term in (2.3),  $T_{\text{DM}} - T(r)$ , which is used to define an ADM characteristic radius,  $r_{\text{DM}}$ , implicitly as

$$T(r = r_{\text{DM}}) = T_{\text{DM}}. \quad (2.4)$$

Then dark matter takes energy from  $r < r_{\text{DM}}$  and deposits it at  $r > r_{\text{DM}}$  for a standard MS temperature profile (monotonically decreasing from the center outward). With our chosen ADM parameters we see typical values:

$$r_{\text{DM}} \sim \mathcal{O}(0.1R_{\star}) \quad (2.5)$$

$$l_{\text{DM}} = (\sigma_{\text{p}}n_{\text{p}})^{-1} \sim \mathcal{O}(1R_{\star}) \quad (2.6)$$

---

<sup>5</sup>This is one of the primary reasons that spin-dependent and spin-independent scattering give qualitatively distinct results. As the star burns H on the MS, the number of protons is reduced, reducing the importance of spin-dependent scattering processes. In the case of spin-independent scattering, the effect gets more important as helium is produced from H burning during the MS.

where  $l_{\text{DM}}$  is the ADM mean free path (implying that it completes several orbits between scattering events). These values allow dark matter to travel much larger distances than photons or ions within the star (which have  $l \lesssim 10^{-10} R_*$ ) and to traverse qualitatively distinct regions of the star. This large mean free path is what enables dark matter to serve as such an effective coolant despite being far less numerous than either photons or ions [Spergel and Press, 1985].

## 2.4 RESULTS

In standard stellar evolution, with no influence from dark matter, stars with  $0.9 \leq M_*/M_\odot \leq 5.0$  naturally split into two groups with qualitatively different structures, based on the dominant channel through which they burn hydrogen. Spin-dependent asymmetric dark matter (ADM) affects core hydrogen burning, mainly by flattening the temperature gradient. In this section, we will review standard stellar astrophysics [Kippenhahn et al., 2012] and then describe the effects of asymmetric dark matter seen in our MESA simulations.

The dominant burning channel is determined by the core temperature, with the transition happening at  $T_c \sim 2 \times 10^7$  K, which corresponds to a stellar mass of  $M_* \sim 1.3 M_\odot$ . Stars with  $M_* \lesssim 1.3 M_\odot$ , which we call low-mass stars, burn hydrogen primarily through the proton-proton (pp) chain for which the burning rate scales with temperature very roughly as  $\epsilon_{\text{pp}} \propto T^4$ . For stars in the mass range  $0.4 \lesssim M_*/M_\odot \lesssim 1.3$ , the transport of energy away from the core burning region is dominated by photon diffusion. Energy transport in the cores of such stars is said to be radiative.

Stars with  $M_* \gtrsim 1.3 M_\odot$ , which we call high-mass stars, are dominated by the carbon-nitrogen-oxygen (CNO) cycle, for which the burning rate scales much more strongly with temperature,  $\epsilon_{\text{CNO}} \propto T^{16-20}$ . In CNO-dominated stars, radiative energy transport is insufficient to carry away the energy produced by hydrogen burning; consequently, they have convective cores.

In § 2.4.1 and § 2.4.2, we will consider results for low-mass stars and high-mass stars separately and we will demonstrate that ADM has distinct effects on the evolution of the two

groups. § 2.4.3 details changes in MS lifetimes due to ADM energy transport. We discuss the effects on surface properties of individual stars in § 2.4.4, and on the isochrones of stellar populations in § 2.4.5. Note that all logarithms in this paper are base 10.

### 2.4.1 Low-Mass Stars

Standard model stars in the mass range  $0.9 \leq M_*/M_\odot \lesssim 1.3$  have relatively low central temperatures and so are powered primarily by the pp chain, which is much less sensitive to the temperature than burning via the CNO cycle. This means the burning does not peak as strongly at the center and radiative transport is sufficient to carry the energy flux, so the core is radiative. Without the mixing provided by convection, hydrogen depletes first at the very center and the burning shifts gradually outward into a shell.

As seen in Figure 2, energy transport by large amounts of ADM causes flatter temperature gradients in the center than those seen in the NoDM model. This reduces the burning rate in the center (as long as the local density does not get too high), where ADM is removing energy, and increases it in a shell, where ADM deposits energy. Note that in Figure 2 the burning rate is not significantly reduced at the center; this is due to a significant increase in density (with which the burning rate per gram scales linearly) as the star’s structure adjusts to cooling in the core. These results are generally in agreement with the results of previous papers studying ADM with similar properties. Taoso et al. [2010] found decreased core temperatures in the models of the sun affected by spin dependent ADM with  $\sigma_p = [1, 2, 3] \times 10^{-36} \text{ cm}^2$  and  $m_{\text{DM}} = 7 \text{ GeV}$ . Iocco et al. [2012] found similarly altered temperature and burning profiles in solar mass models affected by ADM with  $\sigma_p = 10^{-37} \text{ cm}^2$  and  $m_{\text{DM}} = 10 \text{ GeV}$  in  $\sim \Gamma_B = 10^3$  environments. Our work uses similar or slightly lower values and considers the full range of environments expected to exist in nature.

The increased burning at larger radii ( $m(< r) \sim 0.1 M_\odot$ ) causes a small net increase in the total luminosity of the star. Because ADM probes temperature differences over large portions of the star, energy transport by ADM remains efficient despite the shallower temperature gradient. The increased temperature at  $m(< r) \gtrsim 0.1 M_\odot$  means that more hydrogen burns during the MS in low-mass ADM models than in their NoDM counterparts. These two

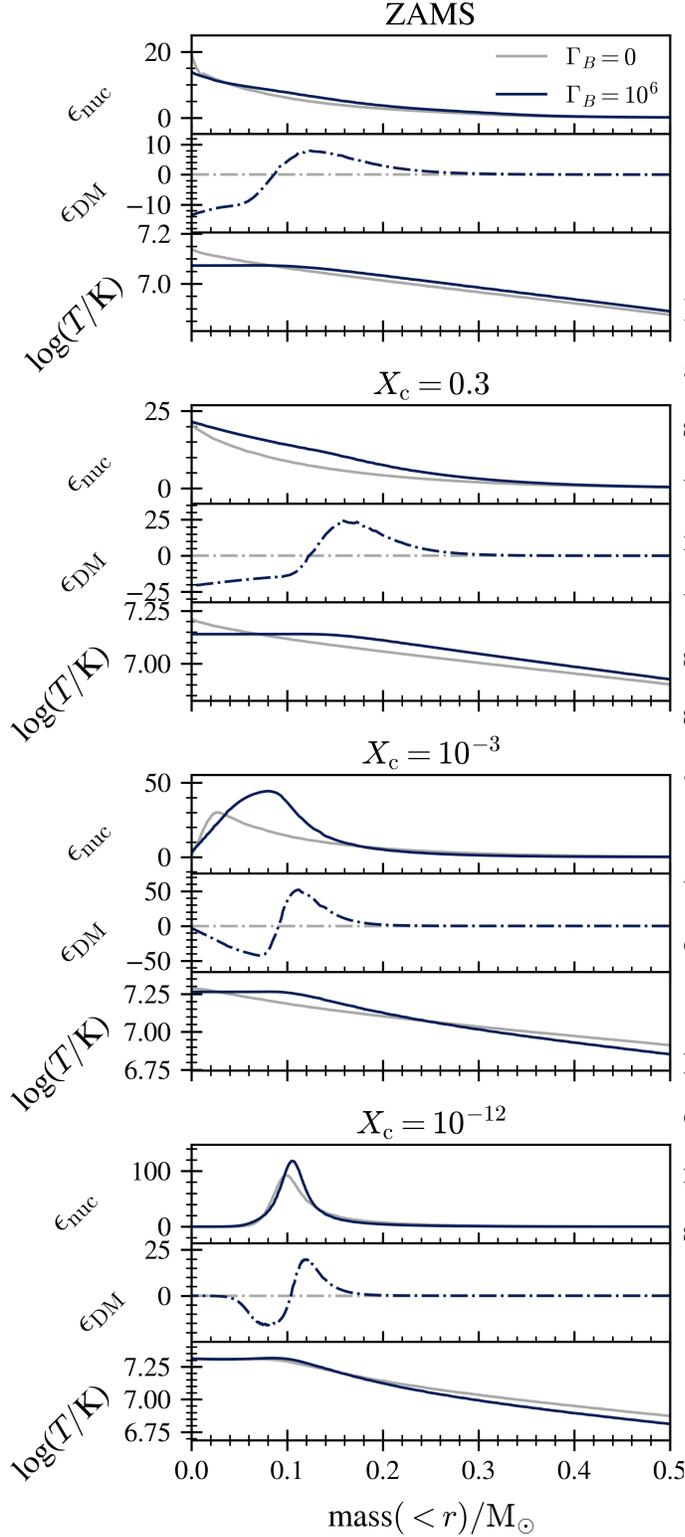


Figure 2:  $1.0 M_{\odot}$  profiles for NoDM (grey) and  $\Gamma_B = 10^6$  (dark blue) models. Each set of 3 panels shows stellar profiles of the stars at different evolutionary phases indicated by the fraction of hydrogen in the center,  $X_c$ , which decreases as the star evolves (ZAMS is "zero-age main sequence"). The profiles in each panel are: 1)  $\epsilon_{\text{nuc}}$ , the nuclear burning rate in [erg/g/s]; 2)  $\epsilon_{\text{DM}}$ , the rate at which DM transports energy (negative values indicate that energy is being removed), also in [erg/g/s]; 3)  $\log(T/\text{K})$ ,  $\log_{10}$  of the temperature in [K]. ADM energy transport decreases the temperature and burning rate in the center and increases them in a shell at  $m(< r) \sim 0.1 M_{\odot}$ .

competing effects dictate the gross evolution of the star. For low values of  $\Gamma_B$  (weaker ADM influence), the increase in the total amount of fuel wins and the stellar MS lifetime is an increasing function of  $\Gamma_B$ . At values of  $\Gamma_B \gtrsim 10^3$ , the burning rate continues to increase, but little additional fuel is burned because of the precipitous drop in temperature at larger radii, which cannot be overcome by the energy transport. The result is that the increase in MS lifetimes peaks at  $\Gamma_B \sim 10^3$  after which MS lifetimes decrease with increasing environmental factor, eventually declining below the lifetimes of NoDM stars at a point which depends on both the stellar mass and  $\Gamma_B$ . (We will discuss stellar lifetimes further in § 2.4.3). Stars in this regime have a higher surface luminosity than their NoDM counterparts, at fixed central hydrogen fraction. However, at fixed luminosity we should expect their effective temperatures to be roughly the same, because these are equilibrium points dictated by the equations of stellar structure. Indeed, this is what we see in our models. The result is that stars of a given mass but different ADM content move along roughly the same tracks in an HR diagram, but they do so at different rates. (We will discuss stellar tracks in § 2.4.4 and population isochrones in § 2.4.5.)

In our initial simulations we used an earlier version of MESA (release 10398) which included less rigorous energy conservation requirements and produced very different results in this low-mass regime. In those models, energy transport by large amounts of ADM caused the temperature profile to become inverted, with the central temperature falling below the ADM temperature. Once this occurred, ADM energy transport reversed direction and began moving energy *toward the center* of the star. This raised the central temperature until it surpassed the ADM temperature again, causing ADM energy transport to again reverse direction and move energy *away* from the center. This cycle was self-reinforcing and resulted in large oscillations in the core temperature, density, etc., which propagated outward and resulted in large oscillations in surface properties as well. Previous work by [Iocco et al. \[2012\]](#) found similar dramatic oscillations in solar mass stars in dense ADM environments, and noted that they were unable to determine whether it was a physical effect or a numerical artifact. Upon further investigation of our initial models, we found that they had poor energy conservation. A new MESA version had been released since we had begun this work that included improved energy conservation schemes. When we updated to MESA release 12115

and ran the models again, the energy conservation was much improved, and the central temperature was reduced such that it was very close to the ADM temperature but never dropped below it. Since  $T_c > T_{\text{DM}}$  throughout the star’s lifetime, ADM energy transport never reverses direction and the oscillations seen previously are now absent. We conclude that the oscillations in our initial simulations were a numerical artifact, and that strict energy conservation requirements are necessary for a proper accounting of ADM effects.

### 2.4.2 High-Mass Stars

In standard models, MS stars with  $M_\star \gtrsim 1.3 M_\odot$  are powered primarily by the CNO cycle. This has several important consequences: (1) the burning rate is much higher than in pp-dominated stars; (2) the burning rate is extremely sensitive to core temperature; (3) temperature gradients in the stellar core are relatively steep; and (4) stellar cores must be convective in order to carry away the energy produced by core hydrogen burning. Convective energy transport in the star also replenishes the core with unburnt hydrogen as the star evolves. Once hydrogen throughout the convective zone is depleted, the burning rate rapidly decreases and the star loses more energy at its surface than is being generated by burning. Gravity temporarily overcomes pressure support and the star contracts until the internal temperature increase is sufficient to ignite hydrogen in a shell outside the depleted core. This restructuring produces the so-called "convective hook" in an HR diagram as the star leaves the main sequence [Kippenhahn et al., 2012].

If a star captures enough ADM, the combination of dark matter + radiative energy transport becomes sufficient to carry the flux from nuclear burning at a temperature gradient that is insufficient to support convection. In other words, the additional energy transport by ADM can turn off convection within the stellar core. This can be seen in Figure 3 for a  $3.5 M_\odot$ ,  $\Gamma_B = 10^6$  star. Convection disappears from the center first (where ADM energy transport is most efficient) and retreats away from the core, into a narrowing shell. Without convective mixing, the hydrogen fuel supply depletes first at the very center (instead of simultaneously throughout the core) and the burning concomitantly shifts gradually into a shell, following the lower boundary of the convective zone. This can be seen in the time

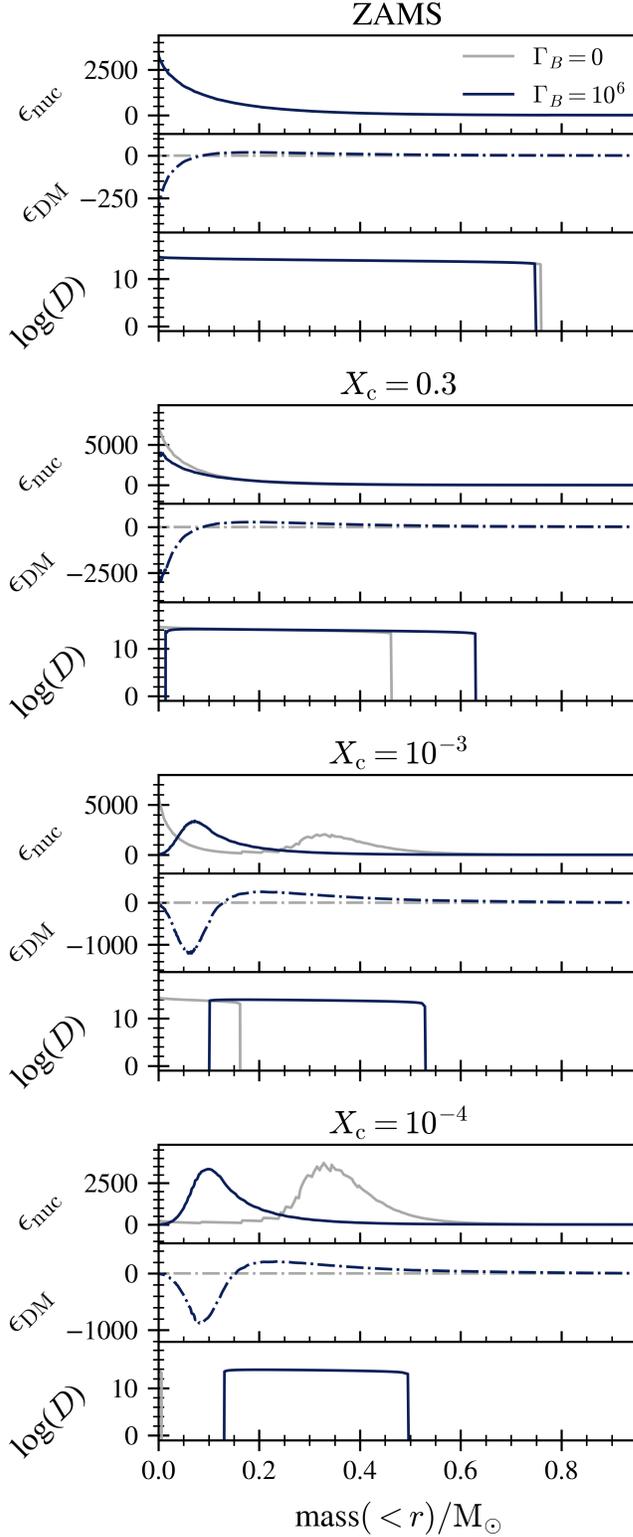


Figure 3: Same as Figure 2 for  $3.5 M_{\odot}$  models, except that the 3rd panel in each set shows  $\log(D)$ , where  $D$  is the diffusion coefficient for convective mixing in  $[\text{cm}^2/\text{s}]$ . In the NoDM model the convective core retreats *toward* the center over time, and the burning rate peaks at the center until the end of the main sequence when the burning rate drops dramatically and a shell of strong burning appears suddenly. In the  $\Gamma_B = 10^6$  model, convection at the very center shuts off relatively early in the MS and a convective shell retreats *away* from the center over time. The peak burning rate shifts gradually outward, following the inner edge of the convective shell. The  $\Gamma_B = 10^6$  model reaches the  $X_c$  evolutionary markers at younger ages, relative to the NoDM model, since convection cannot replenish the fuel at the center.

progression (down the page) of the  $\Gamma_B = 10^6$  (dark blue) model in Figure 3. The shift to shell burning is similar to the behavior of standard *low mass* stars.

The suppression of convection in the cores of stars affected by ADM has been noted in previous work. It was first predicted to occur in horizontal branch stars by [Renzini \[1987\]](#). Both [Casanellas and Lopes \[2013\]](#) and [Casanellas et al. \[2015\]](#) reported the effect in stars with  $\sim 1.3 M_\odot$  affected by ADM with varying  $\sigma_p$  and  $m_{\text{DM}}$  in a solar-like environment. The fixed values of ADM parameters used in this work are at the lower limits of the ranges considered there. [Lopes and Lopes \[2019\]](#) found similar suppression of convection in models of stars with  $M_\star \lesssim 2 M_\odot$  in the Milky Way’s nuclear star cluster ( $\Gamma_B \sim 10^3$ ), with ADM properties  $\sigma_p = 10^{-37} \text{ cm}^2$  and  $m_{\text{DM}} = 4 \text{ GeV}$ . Here we have explored the full range of ADM environments likely to be found in nature, and the full range of stellar masses that would be impacted. We show that not only is convection suppressed in the core, it moves into a shell that retreats away from the center as the star evolves. In the most extreme environments, stars with masses up to  $\sim 4 M_\odot$  are affected. See § 2.4.3 for further discussion of convective cores (and the impact on stellar lifetimes) with respect to varying stellar masses and environments.

### 2.4.3 Main Sequence Lifetimes

In Figure 4 we summarize the effects of ADM on main sequence (MS) lifetimes relative to a standard NoDM star of the same mass. For the purposes of this paper, we have *defined* the MS to end when the fractional abundance of hydrogen in the center,  $X_c$ , falls below  $10^{-3}$ . Once  $X_c < 10^{-3}$  the hydrogen burning rate is greatly reduced and the star transitions out of the MS and onto the sub-giant branch. This transition period is marked by relatively sudden and dramatic changes to the star’s structure. Stars that capture large amounts of ADM can have significantly different core structures at the end of the MS than their standard model counterparts, and these differences affect the stars’ transition out of the MS, including the duration, in ways that are qualitatively different than ADM’s effect on the MS itself. Therefore, different choices in the definition of when a star leaves the MS can affect the results. Our relatively moderate choice of  $10^{-3}$  highlights changes in the core of the star

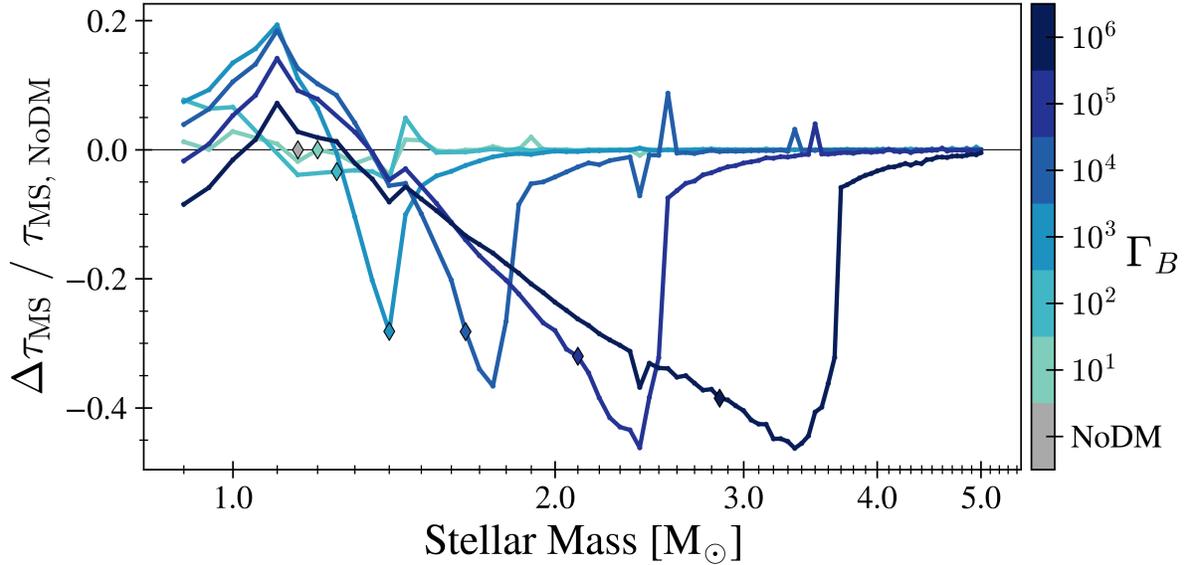


Figure 4: Changes in MS lifetimes, relative to a star of the same mass with no dark matter, seen in our simulations. Diamonds mark the transition from radiative to convective cores (left to right). For the purposes of this figure this is defined as the lowest  $M_*$  for which the average (over the MS) mass of the convective core is greater than  $0.01 M_{\text{star}}$ . Stars to the right of the NoDM marker (grey diamond) have decreased lifetimes due to a reduction in the size of the convective core, which reduces the amount of hydrogen available for burning. The effect abruptly disappears as stellar lifetimes become shorter than the time required to build up a sufficient amount of ADM. Stars to the left of this marker show mixed behavior due to the competing effects of decreased central burning rates and higher temperatures around  $m(< r) \gtrsim 0.1 M_{\odot}$  which give the star access to more fuel. In addition to these trends, there are several abrupt dips (e.g., at  $2.4 M_{\odot}$ ) and spikes (e.g., at  $2.55 M_{\odot}$ ). This is due to rotational mixing that turns on part-way through the MS and funnels fresh hydrogen fuel to the center, which increases the lifetimes (spikes). The dips result when the NoDM model exhibits this feature, but the ADM model of the same mass does not.

during the bulk of the MS, rather than changes in the transition period between the MS and the sub-giant branch. We discuss ADM's effects on this transition period in § 2.4.4.

The MS lifetimes of relatively low-mass stars (near  $\sim 1 M_\odot$ ) can be altered by up to 20%, however the sense and degree of the shift is not a monotonic function of the strength of the dark matter effect, parameterized by  $\Gamma_B$ . This complicated dependence on the amount of captured ADM is due to the competition between increased burning rates and increased availability of burnable hydrogen fuel as discussed in § 2.4.1.

At higher masses, the influence of ADM on stellar lifetimes is clearer. ADM shortens the lifetimes of high mass stars ( $M_\star \gtrsim 1.3 M_\odot$ ). In NoDM models, the central convection zone extends beyond the burning region, giving the star a source of fresh nuclear fuel as hydrogen from outside of the core is mixed into the center. Since ADM shuts off convection in the center, the star no longer gets this influx of fresh hydrogen. Consequently, the star has less fuel available to burn, and so it leaves the MS faster than the NoDM models. Note that the appearance of a convective core (diamonds, Fig 4) shifts to higher masses with increasing  $\Gamma_B$  due to larger amounts of ADM which can carry larger energy fluxes. The effect disappears abruptly as  $M_\star$  increases because stellar lifetimes scale as  $M_\star^{-2.5}$  and quickly become too short for a sufficient amount of ADM to build up (recall that the ADM capture rate scales roughly linearly with  $M_\star$ ), while the luminosity of the star increases rapidly with mass (roughly,  $L \propto M_\star^{3.5}$ ), meaning that more energy must be transported in order to alter the stellar structure.

These changes in MS lifetimes are consistent with those seen in Lopes and Lopes [2019] for stars in an environment similar to the Milky Way Nuclear Star Cluster (roughly  $\Gamma_B = 10^3$ ). Here we extend the study to a wider range of environments, including the most extreme environments likely to be found in nature, and therefore we see effects over a wider range of stellar masses. We will further extend the analysis to the period of transition out of the MS in § 2.4.4.

In addition to the MS lifetime trends we have discussed, Figure 4 has several abrupt dips (e.g., at  $2.4 M_\odot$ ) and spikes (e.g., at  $2.55 M_\odot$ ). This is due to rotational mixing that turns on part-way through the MS and funnels fresh hydrogen to the center, which increases the lifetimes (spikes). The dips result when the NoDM model exhibits this feature, but the

ADM model of the same mass does not. This rotational mixing occurs sporadically (i.e. at isolated masses, not in a continuous range of masses) in our models and may or may not be physical. However, it cannot be a bug introduced by our ADM energy transport module since some NoDM models display the feature, but our module is not called in this case. This phenomenon where a star of a given mass receives an influx of hydrogen to the center due to the onset of mixing, while stars of bracketing masses experience no such mixing, has also been reported previously in the MESA mailing lists<sup>6</sup>.

#### 2.4.4 Stellar Evolutionary Tracks

One of the goals of this work is to determine whether or not ADM can cause any gross changes to the properties of stars. We begin to answer this question with Figure 5, which shows evolutionary tracks on the HR diagram for many of our models. The tracks begin on the zero-age main sequence (ZAMS), delineated by the dotted black lines at the lower left of each panel. Stars evolve off of the ZAMS in a mass-dependent manner that is familiar from well-known aspects of standard stellar evolution. The tracks that we show leave the MS, defined as  $X_c < 10^{-3}$ , at the points marked by triangles. The points at which stars exit the MS in a standard model with no dark matter are indicated by the solid black lines in each panel. (The spike near the  $2.5 M_\odot$  track is due to rotational mixing in the  $2.4 M_\odot$  model, which is discussed in § 2.4.3.) Stars spend the majority of their lives on the MS and move more rapidly through the subsequent phases of stellar evolution. Our evolutionary tracks terminate when the core helium fraction falls below  $10^{-3}$ .

As is evident in Fig. 5, the effects of ADM on the evolutionary track of any individual star are generally subtle. Roughly speaking, this is not surprising. At a fixed central hydrogen fraction, stars containing ADM have different surface luminosities than their standard model counterparts, which is a result of ADM altering the structure of the star. This can be seen in the difference in location between the  $X_c = 10^{-3}$  markers in Fig. 5 (x's for ADM models, solid black line for NoDM models). However, if we instead consider stars at fixed luminosity, the temperature profiles, chemical compositions, opacities, and other properties of the overlying

---

<sup>6</sup><https://lists.mesastar.org/pipermail/mesa-users/>

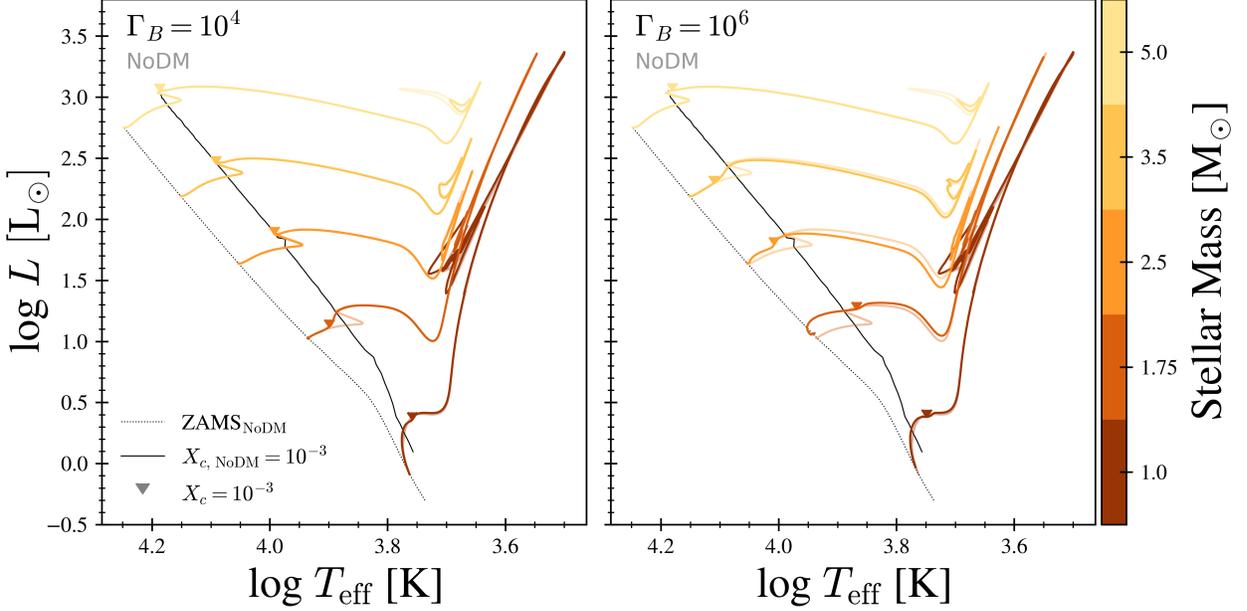


Figure 5: Stellar evolution tracks, from ZAMS to core helium depletion ( $Y_c = 10^{-3}$ ), of select masses with  $\Gamma_B = 10^4$  (left) and  $\Gamma_B = 10^6$  (right). Tracks of NoDM models of the same mass are overplotted with a higher transparency. The chosen masses highlight some of the most dramatic changes seen in Figure 4. Triangles mark the location where stars leave the MS, which we define here as core hydrogen depletion below  $X_c = 10^{-3}$ . The location of the ZAMS and core hydrogen depletion for NoDM models are plotted as dotted and solid black lines respectively. The spike in the  $X_{c, \text{NoDM}} < 10^{-3}$  line near the  $2.5 M_\odot$  track is due to rotational mixing in the  $2.4 M_\odot$  model, which is discussed in § 2.4.3. The main effect of ADM on a star’s surface properties is to move the star through roughly the same sequence of events at a faster or slower pace, causing the offset of the  $X_c = 10^{-3}$  milestone relative to NoDM. High mass stars with sufficient ADM skip the convective hook because ADM shuts off convection in the core. These stars transition into shell burning, and therefore the sub-giant branch, more smoothly, similar to low mass stars.

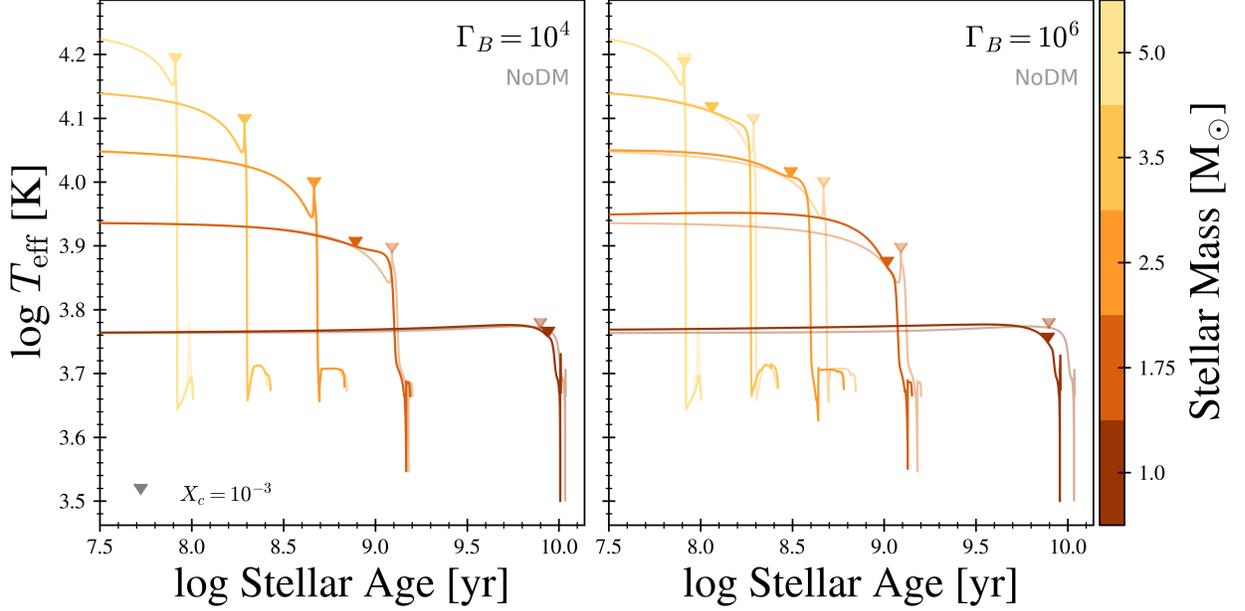


Figure 6: Same as Figure 5, except that here we plot effective temperature as a function of stellar age. The sub-giant branch phase in a star’s evolution is seen here as a sharp drop in  $T_{\text{eff}}$ , and is a result of structural changes in the star that are triggered by large reductions in the core burning rate due to hydrogen depletion at the end of the MS. The duration of the transition period between the MS and the sub-giant branch is seen here as the temporal difference between the locations of  $X_c = 10^{-3}$  (triangles) and the drop in  $T_{\text{eff}}$ . ADM alters the duration of the transition period, tending to increase it in high mass stars and decrease it in low mass stars. This is opposite of ADM’s effect on MS lifetimes. The net effect is that the feature in  $T_{\text{eff}}$  always occurs either concurrently or earlier in the ADM model than in its NoDM counterpart.

zones are approximately unaltered by ADM cooling. Consequently, the gross properties of the stellar photosphere, which are determined via the equations of stellar structure, are approximately fixed, at fixed luminosity. The result is that both ADM and standard model stars of a given mass follow roughly the same tracks in the HR diagram, but they do so at different rates.

Nonetheless, there are some small differences between standard evolutionary tracks and the tracks of stars with ADM. For example, consider the track of the  $1.75 M_{\odot}$  star in the left panel, corresponding to an environmental factor of  $\Gamma_B = 10^4$ . The standard model of stellar evolution shows a kink in the evolutionary track as the star exits the MS. This kink is known as the *convective hook*. The convective hook is caused by an overall contraction of the convective cores of the stars after hydrogen depletion. During this phase,  $T_{\text{eff}}$  increases. Eventually, at the hottest point on the hook, contraction of the former convective core is sufficient to ignite burning in a shell. After this point, shell burning ensues and the star continues to evolve along the sub-giant branch. What is clear from the evolutionary track of the  $1.75 M_{\odot}$  star in the left panel of Fig. 5 is that this evolutionary track exhibits *no convective hook*. This is because convection within the stellar core has been shut off by the ADM in these models. Instead of going through a phase of core contract followed by shell burning, such stars make a smooth transition to shell burning and, thus, a smooth transition to the sub-giant branch. The absence of convective hooks is evident for a wider range of masses in the right hand panel of Fig. 5, which corresponds to a larger environmental factor of  $\Gamma_B = 10^6$ .

The convective hook feature has been clearly seen in many open clusters for which the main sequence turn off lies between  $\sim 1.3 M_{\odot}$  and  $\sim 2 M_{\odot}$ , corresponding to stellar ages of  $\sim 1$  Gyr to  $\sim 4$  Gyr [e.g., see Fig. 18 in Gaia, 2018]. However, this does not yet provide any strong statement about the nature of dark matter because none of those environments are thought to contain significant amounts of dark matter. If a stellar population were identified within the appropriate age range, associated with a significant amount of dark matter, and containing a sufficiently large number of stars, the presence or absence of a convective hook should be immediately clear on an HR diagram.

In Figure 6 we plot the effective temperatures with respect to stellar age to better un-

derstand the transition from the MS to the sub-giant branch, seen here as the difference between  $X_c = 10^{-3}$  (triangles) and the sharp drop in  $T_{\text{eff}}$ . A star’s exit off of the main sequence is triggered when the core hydrogen fuel supply is depleted and the burning rate decreases such that it no longer provides sufficient pressure support and the star begins to contract. Densities and temperatures increase until the bottom layer of hydrogen, now in a shell surrounding the core, ignites. The outward pressure resulting from increased shell burning causes the star’s outer envelope to expand and cool at roughly constant luminosity, seen in Figure 6 as a large, sudden drop in  $T_{\text{eff}}$ . In models with no ADM, this transition period is more abrupt in high mass stars due to the mixing induced by their convective cores, which causes hydrogen to become depleted throughout the core simultaneously and shell burning to appear suddenly (see § 2.4.2). In standard model low mass stars the cores are not convective, and so hydrogen depletes first at the very center and the burning shifts into a shell more gradually (see § 2.4.1).

The temporal difference between ADM and standard models in the location of the sudden drop in  $T_{\text{eff}}$  (when increased shell burning causes the envelope to expand) is another indicator of the change in MS lifetime. Unlike our definition of the end of the main sequence ( $X_c = 10^{-3}$ ) this indicator is based on surface properties and occurs towards the end of the large structural changes that happen during the transition period. In some cases (e.g., in the  $1.75 M_{\odot}$ ,  $\Gamma_B = 10^4$  model) this temporal difference is much smaller than the change in MS lifetime given by our definition of leaving the MS (seen here as the difference between the triangle markers of the ADM model and its standard model counterpart), and in other cases (e.g., in the  $1.0 M_{\odot}$ ,  $\Gamma_B = 10^6$  model) it is larger.

ADM can affect both burning rates and stellar structure (e.g., convection), and therefore it is not surprising that ADM affects the timescale of a star’s transition off of the main sequence. High mass stars that skip the convective hook due to ADM energy transport take longer to move through this transition period because the burning shifts gradually into a shell (see § 2.4.2 for details). This behavior is very similar to standard model low mass stars. Conversely, the  $1.0 M_{\odot}$  ADM models move through this period more quickly than their NoDM counterpart. This is likely due to the fact that ADM has caused higher burning rates at the outer edge of the core during the MS, so that mixing during this transition period brings

more helium into the center than in the standard model star. These shifts are opposite the shifts in MS lifetimes, and the net result is that this feature in  $T_{\text{eff}}$  always occurs either concurrently or earlier in the ADM model than in its standard model counterpart.

#### 2.4.5 Stellar Population Isochrones

Though the evolutionary tracks are quite similar across all models, ADM changes the rate of evolution and, hence, the stellar ages at which stars reach particular evolutionary stages. To convey some of the information that is obscured in an evolutionary track, we present stellar isochrones in Figure 7. Each isochrone is a line on this plot that represents the locus of points that would be occupied by a population of stars of fixed age, but a wide range of masses.

The changes caused by ADM to individual stars' MS evolution is seen in these populations as a shift in the location of the MS turn-off, where the isochrones take a hard right turn. We have chosen these particular ages to maximize the sampling around this period and the subsequent crossing of the sub-giant branch. We do not show the giant branches because we do not have enough data points there to be representative, but we discuss the red giant branch further below. Due to the fact that stars move through the MS turnoff, sub-giant, and giant branches rather quickly, our mass sampling limits our ability to resolve these phases.

At around 1 Gyr, stars of  $M_{\star} \sim 1.75 M_{\odot}$  are leaving the MS. In Figure 7, we see that the MS turnoffs around this time occur at a higher effective temperature and skip the convective hook [consistent with Lopes and Lopes, 2019], reflecting ADM's affects on these stars, discussed in § 2.4.2. These isochrones also tend to cross the sub-giant branch at a lower luminosity. This reflects the fact that ADM speeds up the evolution of these stars, and so stars with smaller initial masses, which have lower effective temperatures, are crossing the sub-giant branch earlier than they otherwise would. This happens in a wider range of ages in populations that live in richer ADM environments (right panel). Stars move through this phase very quickly, meaning isochrones of real stellar populations are very sparsely populated in this region (known as the Hertzsprung Gap), however the location of the gap itself may contribute to discerning between models.

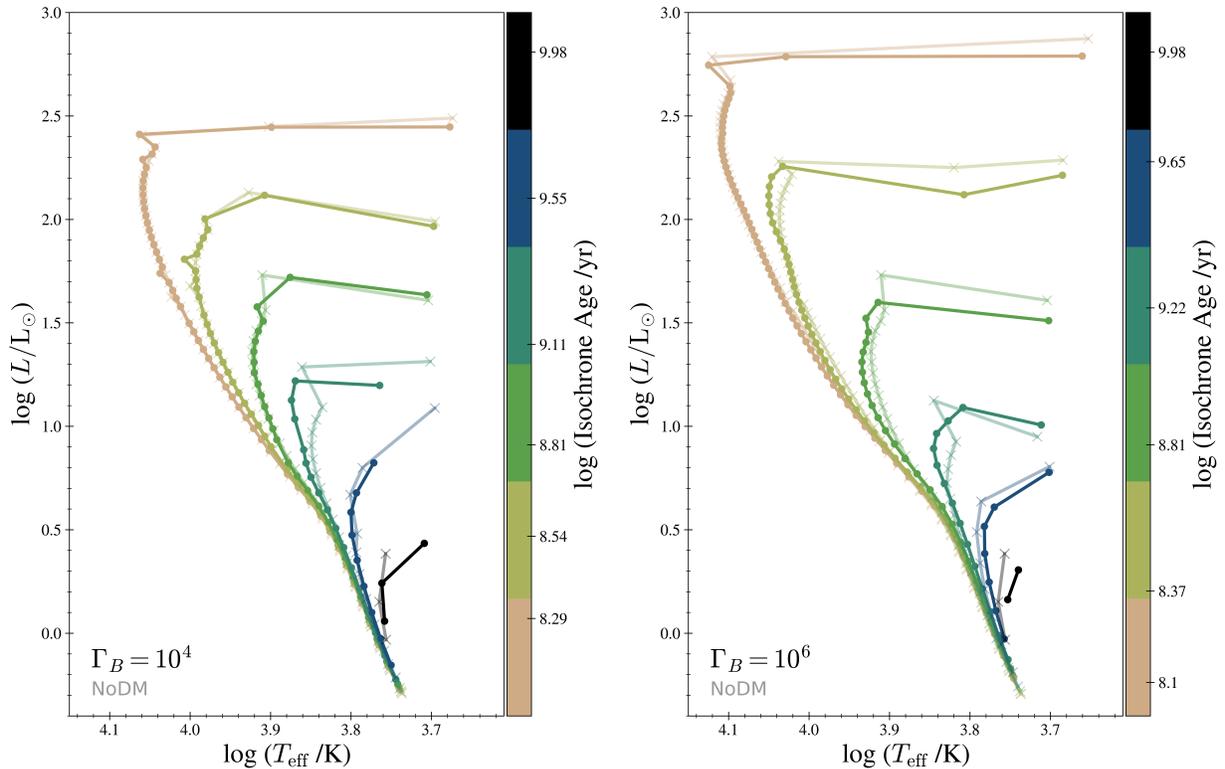


Figure 7: Isochrones for  $\Gamma_B = 10^4$  (left) and  $\Gamma_B = 10^6$  (right) models, marked by circles, with NoDM models overplotted at higher transparency and marked by crosses. The data points are interpolated from the stellar tracks to a common set of times (isochrone ages). We connect the interpolated data points of a single isochrone with straight lines to guide the eye. The lowest data point on every isochrone is a  $0.9 M_{\odot}$  star. As stars leave the main sequence they evolve rapidly, and therefore subsequent phases are less well sampled due to the mass resolution. Isochrone ages have been chosen to maximize the sampling around the MS turnoff and sub-giant branch, and are not the same in each panel. We do not show the giant branches because we do not have enough data points there to be representative. The MS turnoff of isochrones around 1 Gyr happens at a higher effective temperature and a lower luminosity and skips the convective hook. This happens in a wider range of ages in populations that live in richer ADM environments (right panel). The oldest isochrones contain only low mass stars, and the ADM and NoDM populations look very similar except that populations in high  $\Gamma_B$  environments appear slightly older due to their decreased luminosity and temperature.

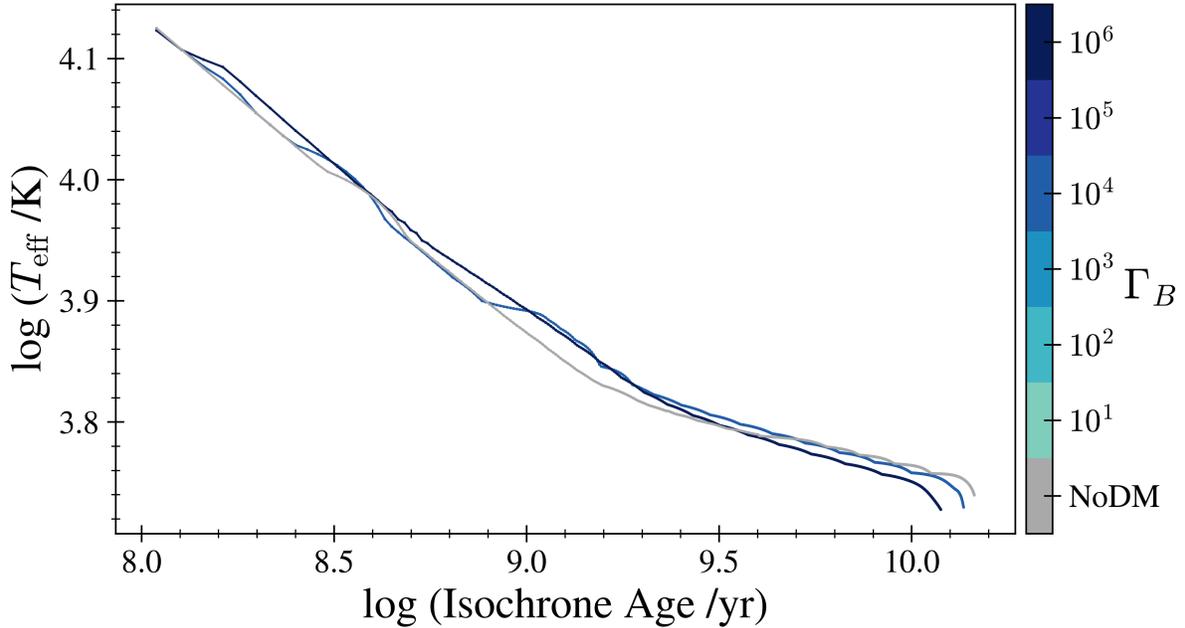


Figure 8: Effective temperature of the main-sequence-turnoff star, defined here as the hottest MS star, at a given age. We show only the NoDM,  $\Gamma_B = 10^4$ , and  $\Gamma_B = 10^6$  models to allow the reader to see them more clearly. The MS turnoff temperature of the  $\Gamma_B = 10^4$  models starts to become hotter than in the NoDM isochrones at  $\sim 0.7$  Gyr when stars of  $\sim 2 M_\odot$  begin leaving the MS, while those of the  $\Gamma_B = 10^6$  models become hotter at  $\sim 0.1$  Gyr when stars of  $\sim 3.5 M_\odot$  begin leaving the MS. As we move to older isochrones the effect is reversed, and ADM models have lower effective temperatures at the turnoff. The lines terminate when there are no more stars on the main sequence (the lowest mass in our set of models is  $0.9 M_\odot$ ). For the purposes of this figure, we exclude stars in the convective hook because this feature is not well-resolved in our isochrones (see Fig. 7).

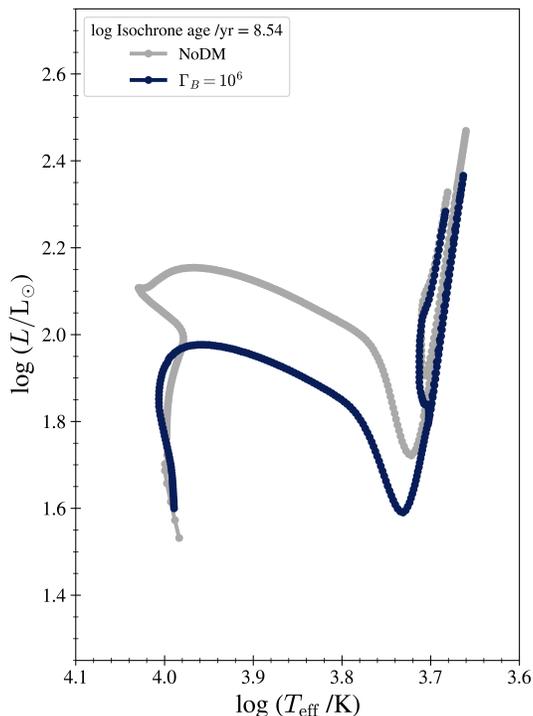


Figure 9: Isochrones generated by the MIST code for NoDM (grey) and  $\Gamma_B = 10^6$  (dark blue) populations. The lowest mass star in both isochrones has  $M_{\star} = 2.25 M_{\odot}$  (the interpolation was not successful for lower masses), and we show data through core helium depletion ( $Y_c = 10^{-3}$ ) which corresponds to  $M_{\star} = 3.2 M_{\odot}$  in both cases. The  $\Gamma_B = 10^6$  isochrone skips the convective hook and crosses the sub-giant branch at a lower luminosity, consistent with Figure 7. Additionally, the tip of its red giant branch occurs at a lower luminosity, a trend which appears in most of the  $\Gamma_B = 10^6$  MIST isochrones.

The oldest isochrones contain only low mass stars, since higher mass stars have already evolved into the giant branches and beyond. Here, the ADM isochrones look very similar to their standard model counterparts, except that populations in high  $\Gamma_B$  environments appear slightly older due to their decreased luminosity and temperature. This indicates that ADM causes the stars’ surface properties to evolve more quickly, likely due to the increase in shell burning.

To better resolve the isochrone’s MS turnoffs, Figure 8 shows the effective temperature of the MS turnoff star, which we define as the hottest MS star at a given age. At younger ages there is no difference between the ADM and standard models because the stars have not yet captured enough ADM to be significantly affected. Around 0.15 Gyr, isochrones of the  $\Gamma_B = 10^6$  model start to display higher temperatures, remaining high until  $\sim 3$  Gyr, after which their temperatures are cooler than their NoDM counterparts. The  $\Gamma_B = 10^4$  isochrones show similar trends, but they occur at later times, since it takes longer for stars in lower  $\Gamma_B$  environments to build up sufficient ADM. The waviness in the lines at older ages is a result of limited mass resolution.

The trends seen in our isochrones are consistent with what we were able to see in the isochrones generated by the MIST code, in regions where that code’s interpolation was successful. (The reader is reminded that this method did not produce any isochrones older than 1 Gyr; see § 2.3 for more information.) In addition, we noticed from MIST isochrones that the tip of the red giant branch tends to occur at a lower luminosity in populations with large amounts of ADM. The tip of the red giant branch is commonly used as a distance indicator, particularly for older populations. If the trend continues in isochrones older than 1 Gyr, ADM may add a source of uncertainty to these studies. To give the reader a sense of this shift, and to show a more well resolved MS turnoff and sub-giant branch, we show one particularly successful MIST isochrone for NoDM and  $\Gamma_B = 10^6$  models in Figure 9. The lowest mass star in both isochrones (lower left) has  $M_\star = 2.25 M_\odot$  (the interpolation was not successful for lower masses), and we show data through core helium depletion ( $Y_c = 10^{-3}$ ) which corresponds to  $M_\star = 3.2 M_\odot$  in both cases.

## 2.5 DISCUSSION AND CONCLUSIONS

We have studied the potential impact of asymmetric dark matter interacting with nucleons through a spin-dependent coupling on the gross evolution of stars. We accomplished this by incorporating a module that approximates heat transport by dark matter into the MESA stellar evolution software. We have identified several interesting qualitative distinctions between the standard evolution of stars and the evolution of stars in environments with a very high dark matter content. These include:

- 1) Flattened core temperature gradients, which alters the burning rates and stellar structures of low and high mass stars (where low mass stars have radiative cores and high mass stars have convective cores in models with no DM) in qualitatively different ways (§ 2.4.1 and § 2.4.2 respectively).
- 2) Convection is suppressed in the cores of high mass stars and pushed into a shell that retreats from the center (§ 2.4.2), resulting in the absence of a convective hook in stellar tracks (§ 2.4.4) and population isochrones (§ 2.4.5).
- 3) Changes to MS lifetimes (defined here as  $X_c > 10^{-3}$ , § 2.4.3). Lifetimes of low mass stars are *increased* by as much as 20%. Lifetimes of high mass stars ( $1.3 \lesssim M_*/M_\odot \leq 5.0$ ) are *reduced* by as much as 40%;
- 4) Stars in both mass regimes cross the sub-giant branch at younger ages (§ 2.4.4) and may reach the tip of the red giant branch at lower luminosities (§ 2.4.5).

Our results are consistent with previous work that considered similar ADM properties using a variety of stellar evolution codes, and we extend the field by considering the full range of ADM environments likely to exist in the universe and the range of affected stellar masses.

Finally, we find that strict energy conservation criteria in the stellar simulation code is crucial for a proper accounting of the effects on low mass stars with large amounts of ADM so as not to trigger large, nonphysical, self-reinforcing oscillations throughout the star (§ 2.4.1).

It is interesting to speculate on ways in which these effects could be used to identify and/or constrain dark matter or ways in which these effects may, at least, serve as an element of uncertainty in the analysis of stellar populations. Any constraint on dark matter

arising from these effects requires very high-quality observations of a stellar population residing in an environment with a large ambient dark matter density and thus there will be a significant element of serendipity involved. Such a population could potentially be observed by the Rubin Observatory’s Legacy Survey of Space and Time [LSST, Ivezić et al., 2019], which is expected to observe hundreds of dwarf galaxies with very high mass/light ratios (spectroscopic followup would be required). The population would need to contain enough stars that the isochrone features we have identified are observable, and stars in the parameter space of interest that are bright enough for quality spectroscopic measurements. If such a population is observed, our models suggest that if it is  $\sim 1$  Gyr old, the hottest MS star should be slightly *hotter* than expected for a population without ADM, and it should be slightly *cooler* at  $\sim 10$  Gyr (differences between NoDM and  $\Gamma_B = 10^6$  models are of  $\sim 5\%$  in both regimes). Contemporary measurements of  $T_{\text{eff}}$  regularly achieve precisions of a few percent, and can be as low as 1.5% or lower with high resolution, high signal-to-noise [see, for example, Soubiran et al., 2010, and references therein]. In addition, the tip of the red giant branch may occur at a lower luminosity. The tip of the red giant branch is commonly used as a distance indicator so ADM may add a source of uncertainty to these studies. Finally, the metallicity is known to affect many of the properties we have discussed (e.g., the locations of various phases in the HR diagram), so ADM may be a contaminant here as well.

Future work along these lines includes 1) chemical abundance studies exploring the effects of altered core burning; 2) asteroseismology of Sun-like MS and red giant branch stars, which could be seen in, e.g., the small frequency separation—a diagnostic that is sensitive to the core structure of the star. 3) spin-independent ADM-nucleus scattering, which should have a larger effect during later phases when stars are burning helium.

## 3.0 THE PITT-GOOGLE LSST COMMUNITY ALERT BROKER

### 3.1 INTRODUCTION

The Vera Rubin Observatory’s Legacy Survey of Space and Time [LSST [Krabbendam and Sweeney, 2010](#), [LSST Science Collaboration et al., 2009](#)] will revolutionize astronomy by imaging the entire visible night sky every  $\sim 3$  nights, probing deeper and faster into the Universe than any previous optical survey. It is projected to discover 100s of millions of variable and transient objects, for which it will publish a real-time data stream containing an estimated 10 million alerts per night for 10 years [[Graham et al., 2019](#)]. The objects in the stream will be of many dozens of different types, falling into several broad categories such as near-Earth asteroids, variable stars, supernovae, and active galactic nuclei. Researchers across astronomy will be looking to obtain and analyze subsets of the diverse data stream that will be small compared to the flood coming from LSST, but often very large compared to the datasets currently available to them. I am developing the Pitt-Google alert broker which will serve the community by processing all 10 million alerts per night and distributing the data in forms and through access methods that are convenient for working with information at this scale.

Alert brokers will provide the main channels for researchers, and the public, to access and study the data. They will take on the roles of ingesting the full stream, enriching alerts with cross match and classification information, curating subsets, distributing data, and providing access tools that are useful for the community. LSST has selected 7 community alert brokers to whom they will deliver the full stream. They are ALerCE<sup>1</sup>, AMPEL<sup>2</sup>,

---

<sup>1</sup><http://alerce.science/>

<sup>2</sup><https://ampelproject.github.io/>

ANTARES<sup>3</sup>, BABAMUL, Fink<sup>4</sup>, Lasair<sup>5</sup>, and Pitt-Google<sup>6</sup>. All of these brokers are currently using the Zwicky Transient Facility (ZTF<sup>7</sup>) alert stream to prepare for LSST and in the process are demonstrating their unique strengths and approaches. The Pitt-Google broker has adopted a cloud-based model that makes the data available through Google’s data services like Pub/Sub, BigQuery, and Cloud Storage. This gives users the opportunity to work with the data using all of the tools Google Cloud has developed to work with data at this scale.

The data coming from LSST will be abundant and rich. But along with it will come new challenges for the community, scientific and technical. LSST will image 1000 times the time-volume space of any other survey, and will do it in 6 broadband filters, where other surveys have used 2-5. It will see at least 2 magnitudes fainter than most other surveys, and it will observe faint objects in a much greater volume per time as others which reach similar limiting magnitudes [Bellm, 2016]. This will provide both much larger sample sizes across dozens of classes and many more headaches in differentiating the interesting events from the common ones. Especially to those seeking early identification in order to trigger follow up of interesting events, LSST’s 6 colors and wide area of coverage will show up looking like sparse light curve data that is scattered across filters. This will be especially true in the early years when there is still a dearth of training data available in the parameter space. There will be plenty of opportunity for many different individuals and collaborations to develop their own algorithms in these pursuits.

Cloud computing and cloud-managed data services have the potential to open up access to LSST’s alert stream to a broad range of people, across skill sets and access to resources. They can allow individuals to process large or small alert streams in a very lightweight way without managing the client connection or compute resources, only their own analysis code. These are technological solutions that are already being used in many industries to overcome very similar problems as those that will be presented by LSST – including real-time processing of high-volume message streams and all related compute and data management

---

<sup>3</sup><https://antares.noirlab.edu/>

<sup>4</sup><https://fink-broker.org/>

<sup>5</sup><https://lasair.roe.ac.uk/>

<sup>6</sup><https://pitt-broker.readthedocs.io/en/latest/>

<sup>7</sup><https://www.ztf.caltech.edu/>

challenges. These cloud-based solutions can enable researchers to deploy and evolve their own algorithms quickly, on the live alert streams, and to share information and collaborate within and beyond sub-fields.

In this chapter I discuss the Pitt-Google alert broker and our efforts to provide enhanced and distilled astronomical alert data through Google Cloud services. In particular, I detail the software that runs the broker’s alert stream pipelines and supporting infrastructure, and our custom APIs that facilitate data access. All software and related plans are the result of my own work, unless otherwise indicated. The chapter is structured as follows. In § 3.2, I give a somewhat detailed overview of the Google Cloud Platform<sup>8</sup>, including their services, tools, and related costs. This is intended as both an introduction and a reference for potential users, since data access will go directly through Google but many in the astronomy community are unfamiliar with the Google Cloud environment. In § 3.3, I present the software that our broker currently uses to run the ZTF alert stream, and then discuss work that is in progress and future development plans. The section concludes with details of Pitt-Google’s budget, which is intended to give potential users a better sense of cloud computing costs and what they might expect for various types of data access and/or service usage. In § 3.4, I present our broker’s custom APIs which facilitate access to the data resources that we maintain in Google Cloud services. This section includes example code that demonstrates data access. A summary is provided in § 3.5.

## 3.2 GOOGLE CLOUD

In this section I give an introduction to Google Cloud, including relevant services, access methods, and costs. The primary purpose is to provide potential users with an understanding of the environment within which their interactions with our broker will occur, as it may be unfamiliar to many. A secondary purpose of this section is to compile a reference of details about each service that will be relevant to future users/developers, but may not be immediately obvious in the online documentation. I will start by briefly discussing the cloud

---

<sup>8</sup><https://cloud.google.com>

model, outlining the major differences between our approach and that of other brokers and what this means for the end user. § 3.2.1 outlines relevant Google Cloud costs and options for offsetting them. In the next 4 subsections, I overview the specific Google Cloud services that are relevant to data access. § 3.2.6 discusses some of the data access methods that Google provides, and § 3.2.7 describes the basic setup process for a new user.

Our broker stores data in Pub/Sub, BigQuery, and Cloud Storage. This is a very different model from that used by other alert brokers, and results in significant differences for end users. Most brokers use Apache Kafka and Apache Cassandra, or similar solutions, to store data on virtual machines (VMs) that they manage themselves. As a result users access the data through the *broker* – even if that broker’s VMs happen to run in Google Cloud. In the model our broker has adopted, users will access the data through *Google Cloud*. This is because Pub/Sub, BigQuery, etc. are managed services, meaning that Google handles the entire process of data storage and delivery, including user authentication. With this approach, users have access to the data through Google’s native tools, like APIs and web consoles, and the option to run their own analyses in the Cloud, next to the data, with minimal setup. One consequence of this model is that each user is responsible to pay Google Cloud for their own data access and/or service usage. In this context, a “user” includes both the Pitt-Google broker and anyone who accesses our data resources, i.e., an “end-user”.

### 3.2.1 Costs

The basic fees associated with each service are described in the following subsections corresponding to the respective services. In general, there is a baseline level of usage that is free and renews monthly. Beyond this, there are typically charges for data storage and for any operations (load/publish data, query a table, pull a subscription, etc.). These charges tend to be small but frequent. The user that makes the request is usually the one who pays the fee. Here is one example: our broker pays to *send* data to BigQuery and to *store* it there, and the end-user pays to *query* it. The first 1 TB per month of data queried from BigQuery is free. After that the charge<sup>9</sup> is \$5.00 per TB. As this cost model is likely to be unfamiliar

---

<sup>9</sup>In this work, all monetary amounts are given in USD.

to many in the community, § 3.3.7 presents details of Pitt-Google’s budget to give readers a better sense of how it works.

There are several options for offsetting costs and accessing resources for free:

- Most services have a Free Tier<sup>10</sup> quota which renews monthly. For example, the first 10 GB per month of data pulled from a Pub/Sub subscription is free, and the first 1 TB of data queried per month from BigQuery is free.
- A 90-day, \$300 free trial<sup>11</sup> is available to everyone.
- Research credits<sup>12</sup> are available in many countries. PhD students can apply for \$1000, postdocs and faculty up to \$5000.
- Academic institutions can apply for an exemption from egress.

### 3.2.2 Pub/Sub

Google Pub/Sub<sup>13</sup> is a messaging service that uses the publish-subscribe pattern. Publishers and subscribers communicate asynchronously and are (mostly) disconnected, with the Pub/Sub service handling all message storage and delivery. Publishers send messages to a topic, and Pub/Sub immediately delivers them to all attached subscriptions. Subscriptions can be configured to either push messages to a client automatically or to wait for a client to make a pull request. The owner of the topic sets the access rights that determine who is allowed to attach a subscription. Messages published to a topic prior to a subscription being created will not be available to the subscriber. By default, Pub/Sub messages are not ordered.

“Push” subscriptions can be used to attach the stream to an HTTP endpoint, to which Pub/Sub immediately delivers the messages. These can be used to trigger automatic stream processing through Cloud Run and Cloud Functions, discussed below. Messages in a “pull” subscription remain there until they are fetched by a client, or for a maximum of 7 days.

This is because every Pub/Sub message is delivered and tracked independently. This allows Pub/Sub to support essentially unlimited client parallelization. The basic delivery

---

<sup>10</sup><https://cloud.google.com/free/docs/gcp-free-tier#free-tier>

<sup>11</sup><https://cloud.google.com/free/docs/gcp-free-tier/#free-trial>

<sup>12</sup>[https://edu.google.com/intl/ALL\\_us/programs/credits/research/](https://edu.google.com/intl/ALL_us/programs/credits/research/)

<sup>13</sup><https://cloud.google.com/pubsub/docs/overview>

process is as follows. Pub/Sub will lease a message to a subscriber client. If the client acknowledges the message before the lease ends, Pub/Sub marks it as successfully delivered. If not, Pub/Sub will re-deliver (re-lease) the message to a client at some arbitrary time in the future. This is very different than Apache Kafka, which is more familiar to the astronomy community. Kafka stores messages within a predetermined set of partitions on a topic. The messages are ordered within a given partition – they are delivered in order, and the delivery tracking mechanism relies on this ordering. Thus Kafka supports parallelization of subscriber clients up to the number of partitions on a topic. The Pub/Sub service is designed to be highly scalable and to handle very large-scale throughput. Google itself uses Pub/Sub to send over  $5 \times 10^8$  messages/sec for Gmail and other products<sup>14</sup>. The advertised, end-to-end latency is “on the order of 100 milliseconds”.

A filter<sup>15</sup> can be placed directly on a Pub/Sub subscription when it is created. The filter will have access to the custom attributes in message metadata (see § 3.3.6). These attributes are required to be strings, and so the filter supports a limited set of operations. For example, it does support the comparison operator = and the function `hasPrefix` (matching values that start with a given substring), but it does not support the operators > or <. With this method, all messages in the stream will be delivered to the *subscription*, but messages that do not pass the filter will be automatically acknowledged without being delivered to the *client*. This means that the subscriber will incur message delivery fees for all messages, but will not incur egress fees (if otherwise applicable) for messages that don’t pass the filter.

Pub/Sub costs<sup>16</sup> fall into three basic categories: delivery, egress, and storage. Fees apply to both publishers and subscribers. Publishers pay a delivery fee when they send a message to Pub/Sub. The message is automatically delivered to all attached subscriptions, for which the subscriber pays a delivery fee. Egress applies when the client is outside the Google Cloud region where the topic/subscription lives. The owner of the *topic* chooses the region; attached subscriptions live in the same region. Using the default settings, there is no charge for message storage. There is a free tier.

---

<sup>14</sup><https://cloud.google.com/pubsub/architecture>

<sup>15</sup><https://cloud.google.com/pubsub/docs/filtering>

<sup>16</sup><https://cloud.google.com/pubsub/pricing>

### 3.2.3 BigQuery

BigQuery<sup>17</sup> is a fully managed data warehouse with a SQL-based analytics engine. It is optimized for complex analytical queries on large datasets. It uses a columnar storage format and relational table structure with support for nested and repeated fields. It does not use or understand indexes, and so single row lookups can be expensive because they still require a full column scan. Tables can be partitioned and clustered for better query performance. Data can be loaded via batch jobs or streaming inserts<sup>18</sup>. Streamed data is typically available to queries immediately. It has built-in geospatial analysis and machine learning (BigQuery ML).

Costs<sup>19</sup> generally fall into one of two categories: storage or analysis. The owner of the dataset pays for storage. The user submitting a query pays for analysis, based on the amount of data scanned by the query. Egress may apply if data crosses the boundary of a Google Cloud region. There is a free tier.

### 3.2.4 Cloud Storage

Cloud Storage<sup>20</sup> is Google's object storage service. Objects are stored in buckets. Buckets have a flat namespace (meaning there is no such thing as a directory or folder), but folder-style functionality is provided by most of the access tools (e.g., console and APIs) which interpret folder hierarchies from slashes in the object name. (Note that in this text I often refer to this as “file” storage, though it is technically object storage.)

Cost<sup>21</sup> categories include storage and data processing (which includes things like GET requests for files). The owner of the bucket pays storage costs. The requestor pays data processing costs. Egress may apply if data crosses the boundary of a Google Cloud region. There is a free tier.

---

<sup>17</sup><https://cloud.google.com/bigquery/docs/introduction>

<sup>18</sup><https://cloud.google.com/blog/products/bigquery/life-of-a-bigquery-streaming-insert>

<sup>19</sup><https://cloud.google.com/bigquery/pricing>

<sup>20</sup><https://cloud.google.com/storage/docs/introduction>

<sup>21</sup><https://cloud.google.com/storage/pricing>

### 3.2.5 Cloud Functions and Cloud Run

Cloud Functions<sup>22</sup> and Cloud Run<sup>23</sup> are managed-compute services run by Google Cloud. They both run containers that are configured as HTTP endpoints. They can be used to process live message streams by attaching Pub/Sub push subscriptions. Incoming requests (i.e., messages) are processed in parallel. The number of container instances scales automatically and nearly instantaneously to meet incoming demand. Code to deploy a Python Cloud Function and attach a push subscription are given in § 3.2.6. Differences between the services are essentially tradeoffs between efficiency (at large scale) and ease of use.

Both services perform per-message processing – the user’s code should accept and process a single alert. This has some advantages including being simple to implement and reason about. Processing errors are isolated and a failure on one message does not impact any others. Perhaps the biggest tradeoff is that the efficiency of database queries or calls to external APIs may be greatly reduced since multiple messages cannot be batched for a single query/call. (Google Dataflow is a managed-compute service that does support batching.)

Cloud Functions are the simplest – the user writes and deploys a “bare function” (i.e., non-containerized). The function can call any other code/data files that reside in the same directory, plus any libraries that can be pip-installed (assuming Python). Various runtime environments are available (e.g., Python and Java but not shell), but GPUs are not available<sup>24</sup>. Upon deployment, Google will containerize the function and assign it an HTTP endpoint. A single container instance can process one request at a time.

With Cloud Run, the user is responsible for containerizing the code and any supporting files. This means the user has full control over the runtime environment. GPUs are available, but only through Cloud Run for Anthos<sup>25</sup>. A single container instance can process up to 1000 requests concurrently.

Costs<sup>26</sup> are based on the amount of memory, etc. a container is configured to use, the

---

<sup>22</sup><https://cloud.google.com/functions/docs/concepts/overview>

<sup>23</sup><https://cloud.google.com/run/docs/overview/what-is-cloud-run>

<sup>24</sup>Warning: Cloud Functions do not have the necessary system libraries to run Python packages that ship with GPU support (e.g., standard `pytorch` distribution), *even if your code doesn't use it*. Your deployment will fail (and worse, the actual cause may not be reported in error logs). To get around this, try installing a distribution (e.g., from a wheel) that does not support GPUs.

<sup>25</sup><https://cloud.google.com/anthos/docs/concepts/overview>

<sup>26</sup>Functions:<https://cloud.google.com/functions/pricing>. Run:<https://cloud.google.com/run/>

number of times it is invoked, and the amount of time it actually runs. There is no charge if no container instance is running. Both services include a free tier.

### 3.2.6 Access Methods

Google Cloud provides many methods to access data from Pub/Sub, BigQuery, and Cloud Storage. Here I provide basic details and references for those most likely to be of use to the reader.

The Console<sup>27</sup> is Google Cloud’s web interface. Most tasks can be completed here. This includes: creating/deleting projects, obtaining authentication credentials, previewing data in catalogs and streams, creating/deleting Pub/Sub subscriptions, running SQL queries in BigQuery, deploying Cloud Functions or Run, etc.

Google provides one or more APIs for all of its services. These can be used to access data from within or outside of Google Cloud. The general command-line interface tool is `gcloud`<sup>28</sup>. Client libraries for individual services are available in many languages, typically including Python<sup>29</sup>, Java, Go, and others. Python libraries can be pip-installed with, for example:

---

```
1 pip install google-cloud-pubsub
2 pip install google-cloud-bigquery
3 pip install google-cloud-storage
```

---

Managed compute services like Cloud Functions and Cloud Run (§ 3.2.5) can be used to process live message streams with minimal setup or resource management. Essentially, the user writes a function (or module) that accepts and processes a single alert, deploys it to a compute service, and attaches a Pub/Sub push subscription. Once this is done, messages will be automatically delivered and processed through the user’s code, in real time. The basic `gcloud` commands needed to deploy a Python Cloud Function and attach a push

---

pricing.

<sup>27</sup><https://console.cloud.google.com>

<sup>28</sup><https://cloud.google.com/sdk/gcloud>

<sup>29</sup><https://cloud.google.com/python/docs/reference>

subscription are given below. Angle brackets, `<>`, and everything between them should be replaced with the user's options. The function should be saved in a file called `main.py`, and the `deploy` command should be run from within the same directory. Commands are documented at <https://cloud.google.com/sdk/gcloud/reference>.

Deploy the function:

```
1 gcloud functions deploy <my-cloud-function> \  
2   --entry-point=<entry-point-name> \  
3   --runtime=<python37> \  
4   --trigger-http \  
5
```

Take note of the function's HTTP endpoint URL, which will be displayed after deployment. Then, create a subscription and configure it to automatically push messages to the endpoint. The following example connects the function to Pitt-Google's heartbeat stream, `ztf-loop`.

```
1 gcloud pubsub subscriptions create <ztf-loop-subscription> \  
2   --push-endpoint=<function-endpoint-url> \  
3   --topic=ztf-loop \  
4   --topic-project=ardent-cycling-243415 # Pitt-Google's project ID  
5
```

To delete the function and the subscription, use:

```
1 gcloud functions delete <my-cloud-function>  
2 gcloud pubsub subscriptions delete <ztf-loop-subscription>  
3
```

### 3.2.7 New User Setup

To access services like Pub/Sub and BigQuery, the user must be authenticated to a Google Cloud project. Projects are free and easy to create and delete.

The recommended authentication method for API calls is through a service account. To set this up starting from scratch, the basic steps are:

1. log into Google Cloud (e.g., using a Gmail address)
2. create a project
3. create a service account and download a key file containing credentials
4. set environment variables for your project ID and the path to your credentials

After this one-time setup, API calls are authenticated automatically. The process is described in more detail in our `pittgoogle-client` documentation, referenced in § 3.4.1.

### 3.3 BROKER SOFTWARE

The Pitt-Google broker facilitates public access to alert streams from astronomical surveys. The software is under active development in preparation for LSST, and has been running the ZTF alert stream consistently since Fall 2020. The broker implementation is explicitly Google Cloud-based. It employs Google-managed services for all related tasks including compute/processing, stream distribution, and data storage. It exposes the data at every stage of processing in ways that enable end-users' direct use of Google Cloud tools and services.

This section is structured as follows. Below, I give an overview of the broker software. The next 5 subsections describe the software that is currently (July 2022) running the ZTF alert stream. These include data resource names, technical details, and performance metrics. § 3.3.6 describes current work in progress and future plans for development. I give more technical detail than is customary in module descriptions because it can both serve as a template for future users who want to process in the cloud, and give context for our budget, which is presented in § 3.3.7.

Figure 10 is a diagram of the software architecture as of July 2022. The modules are grouped into five areas by functionality, outlined in grey. The groupings are for conceptual convenience only. Each module operates independently. There are no clusters of VMs.

The consumer is an Apache Kafka to Google Pub/Sub connector. It simply republishes the survey’s Kafka alert stream as a Pub/Sub stream. The storage and science pipelines are both triggered by the consumer’s Pub/Sub stream.

The storage and science pipelines are implemented using a microservices model, which results in per-alert processing through independent modules. In other words, there is no batching of alerts or fusion of modules. Every module is independently deployed to a managed compute service where it runs in an isolated container configured as an HTTP endpoint. The modules are loosely connected through their input/output Pub/Sub streams. Each module publishes messages to a dedicated topic containing the results of its processing alongside the original alert data (except for the storage notification streams). The next module in the pipeline connects to this topic using a push subscription, which automatically delivers new messages to the subscribing module’s HTTP endpoint. Every request (i.e., message or alert) received by an endpoint is processed independently through an instance of the module’s container. Modules also stream data to BigQuery and/or Cloud Storage as appropriate.

The highly modular pipeline provides open access to the alert and related data at every stage of processing via public Pub/Sub streams, BigQuery catalogs, and Cloud Storage buckets. Users can process a real-time alert stream, hooking into a stream that has been pre-processed and filtered to a level of their choosing. They can either process the alerts through their own analysis code on a Google Cloud compute services, or ingest the data into their own system outside of the Cloud. § 3.4 discusses data access.

### 3.3.1 Conductor

The conductor group contains modules that orchestrate the pipeline. The consumer is the single module in the group.

#### **Consumer module**

*Input* Apache Kafka alert stream produced by an astronomical survey

(current: ZTF; previous: DECAT; future: LSST)

*Output* Pub/Sub topic: `ztf-alerts`

The consumer ingests a survey’s alert stream and produces a Pub/Sub stream that is a

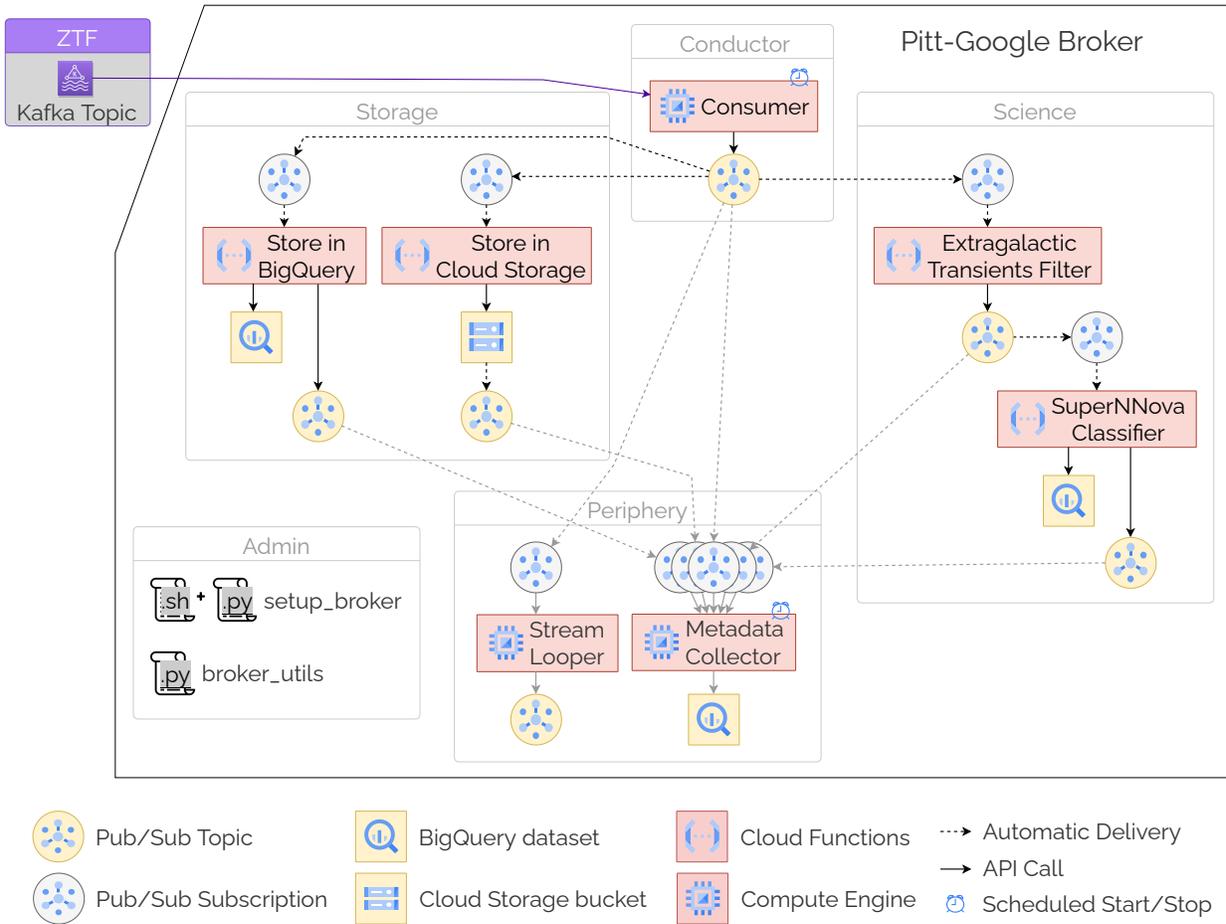


Figure 10: Pitt-Google broker software as of May 2022. Nodes shaded yellow are data resources with public access rights. The resource names are given in the text. Each module is deployed independently to the indicated compute service, shaded red. Conceptual groupings are outlined and labeled in grey. Modules running on Compute Engine each run on a single VM with specs given in the text. Modules running on Cloud Functions are triggered automatically by messages published to their input stream. Black arrows indicate data flow through the main pipelines and grey is through peripheral modules.

nearly exact replica. It runs an Apache Kafka to Google Pub/Sub connector<sup>30</sup> on a Compute Engine virtual machine (VM). ZTF alerts are sent via Kafka as binary encoded Avro data; LSST alerts will be very similar. The consumer is configured to pass the alert bytes directly through without decoding or transforming them in any way. It also attaches the original metadata, such as the Kafka topic and publish timestamp, as attributes on the corresponding Pub/Sub message. The simplicity makes it effective and efficient. Alerts received through Pub/Sub can be handled with the survey’s original tools and instructions, as expected. The software runs in “standalone” mode on a single “g1-small” VM (0.5 vCPU and 1.70 GB memory) and has been ingesting the ZTF stream for a year and a half with no issues. Even so, we may switch to “distributed” mode in the future to ensure fault tolerance.

Fig. 11 shows the consumer’s performance ingesting the ZTF stream on a typical night (top panel, 29 April 2022) and at high incoming rates (bottom panel, 1 April 2022). The high rates were the result of connecting to the ZTF topic after the night was over. The topic contained a backlog of nearly 200,000 alerts which the consumer ingested in just under 14 minutes, as fast as it could given the combined performance of ZTF’s Kafka broker and the Pitt-Google consumer. The main plot in each panel is a scatter plot of the alert latencies (elapsed time between publish timestamps from ZTF and the consumer). The top panel shows that typical latencies are about 0.3 seconds regardless of the incoming rate. The latencies in the bottom panel are due to the fact that the consumer began ingesting well after ZTF had published the alerts. Data has been sigma-clipped to  $4\sigma$ . The maximum latency on 29 April 2022 before clipping was 76.5 seconds. The top subplot of each panel is a marginal histogram that is binned to show the per-minute rate of alerts published by the consumer module. This histogram in the bottom panel shows that the consumer published alerts at rates approximately equal to the average rate expected from LSST. Subsequent figures will show the performance of various modules subject to these rates.

---

<sup>30</sup><https://github.com/GoogleCloudPlatform/pubsub/tree/master/kafka-connector>

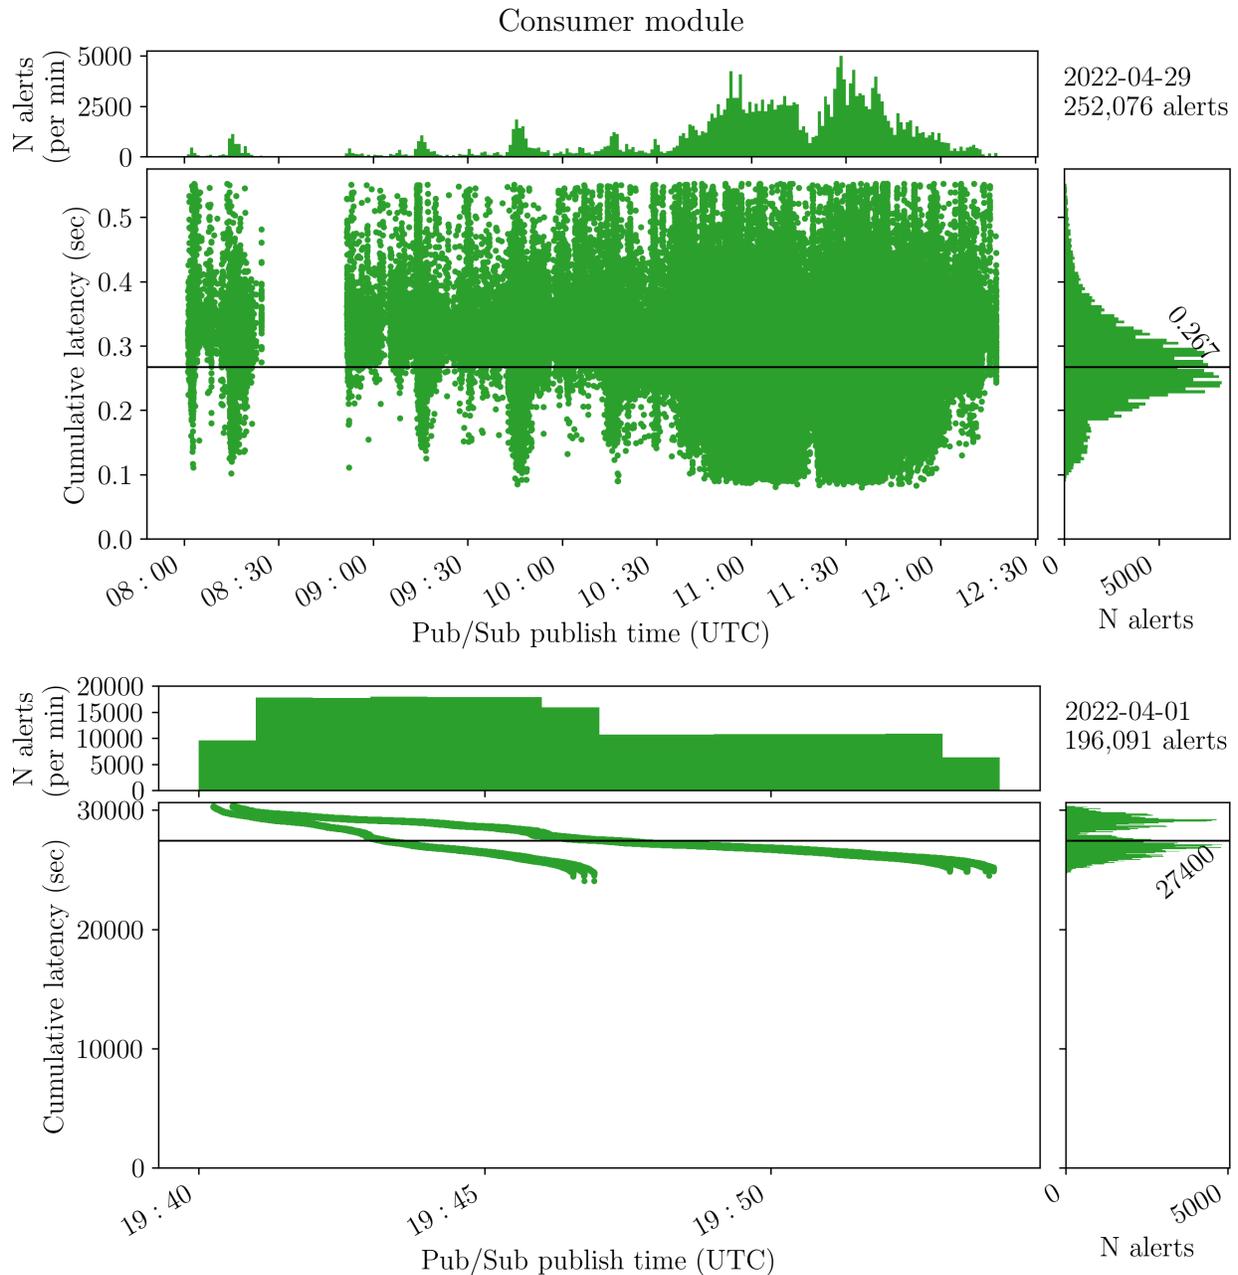


Figure 11: Consumer performance ingesting the ZTF alert stream on a typical night (top panel) and at high incoming rates (bottom panel, explained further in the text). The y-axis is the time elapsed between alert creation by ZTF and the consumer’s publish timestamp. The black line is the median latency and is labeled on the right. The marginal histogram at the top of each panel is binned to show the per-minute rate of alerts published by the consumer. Typical latencies are about 0.3 seconds regardless of the incoming rate. The latencies on 1 April 2022 were due to the delay in connecting to the Kafka stream. This panel’s marginal histogram shows that the consumer published alerts at LSST-like rates.

### 3.3.2 Storage

The storage group contains two modules that send incoming alert data to BigQuery tables and a Cloud Storage bucket respectively. Both run on Cloud Functions and use default settings: 256MB of memory, 60 second timeout. Both produce notification streams through Pub/Sub in addition to storing data.

#### Cloud Storage module

*Input* Pub/Sub stream: `ztf-alerts`

*Outputs* Cloud Storage bucket: `ardent-cyclimg-243415-ztf-alert_avros`

Pub/Sub topic: `ztf-alert_avros`

This module stores the complete incoming alert as an Avro file in a Cloud Storage bucket. The image cutouts can be obtained from these files. (They can also be obtained from the consumer's `ztf-alerts` stream but not from any other Pub/Sub stream or from BigQuery.) The module may alter the alert's Avro schema slightly, if necessary, to conform to strict standards in the ordering of type definitions. This is done to support a use-case where Avro files from Cloud Storage are streamed into BigQuery; this will fail if the schema is not formatted correctly. The bucket itself is configured to produce a Pub/Sub stream, emitting a message upon file upload. The message contains the bucket name, file name, IDs of the alert object and source, and its coordinates as ra and dec. The filename follows the syntax `objectid.sourceid.topic.avro` where `topic` is the name of Kafka topic that originated the alert. Fig. 12 shows the performance of the Cloud Storage module. The bottom panel demonstrates that Cloud Functions can process at LSST-like rates with very little change in performance.

#### BigQuery module

*Input* Pub/Sub stream: `ztf-alerts`

*Outputs* BigQuery tables: `ztf_alerts.alerts`; `ztf_alerts.DIASource`

Pub/Sub topic: `ztf-BigQuery`

This module loads alert data to BigQuery using streaming inserts, which makes the data immediately queryable. The full alert packet<sup>31</sup> is sent to the `alerts` table as a single row

---

<sup>31</sup>The image cutouts are not stored in BigQuery. In the context of BigQuery storage, all references to the

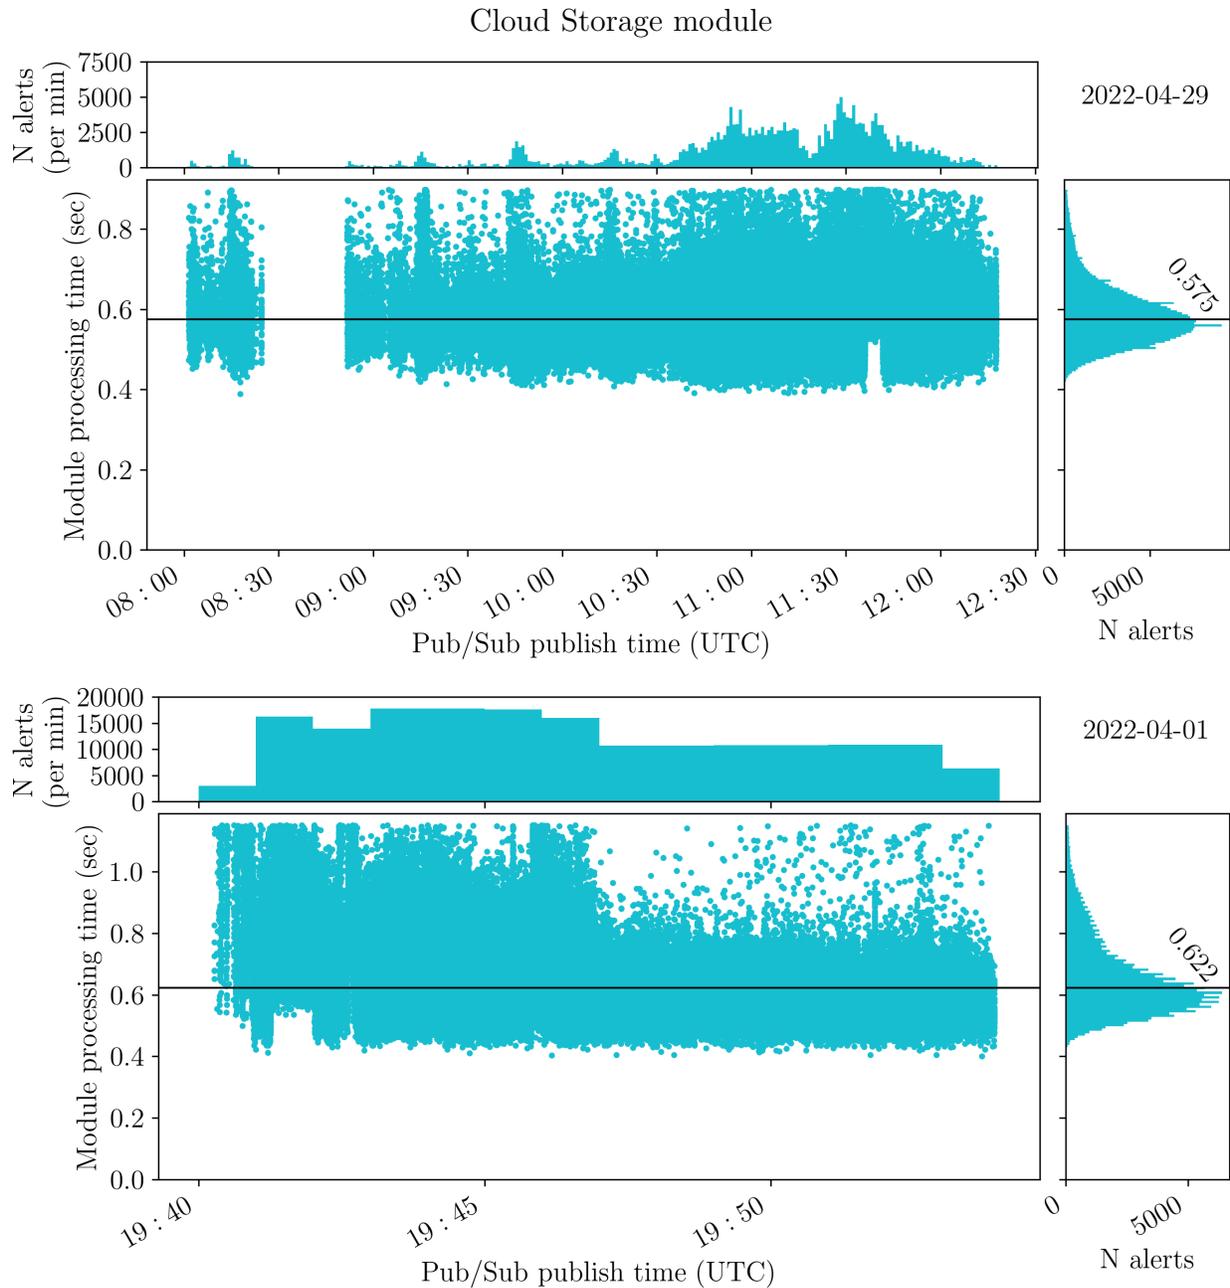


Figure 12: Same as Fig. 11 but for the Cloud Storage module. The y-axis is the difference between publish times of the input and output streams. This is the same as the previous figure but labeled differently to reflect that the delay is now entirely internal and due to the module’s message delivery and compute time. Median processing times are about 0.6 seconds and are essentially unaffected by the incoming rate (which is given in the top marginal histograms of Fig. 11).

with the same schema as the alert’s Avro schema. This provides immediate and persistent SQL access to a replica of the original alert stream. The tradeoffs are that the schema is complex (containing both nested and repeated columns) and the data is highly redundant (since each alert includes data from previous detections which themselves triggered alerts). To support a variety of access patterns, the module also streams data to the `DIASource` table. In this table, a row only contains the object ID, the source detection that triggered the alert, and a list of source IDs for previous detections. This is enough information to recreate the original ZTF alert but allows the table to have a flat schema and avoid data redundancy.

Fig. 13 shows the performance of the BigQuery module. It is similar to the Cloud Storage module when subject to incoming alert rates typical of ZTF, but markedly different at LSST-like rates. This is due to a Google Cloud limit on the rate of BigQuery API requests, beyond which throttling occurs. This limit does not apply to streaming inserts, however the broker makes a `GET` request for the table prior to data insertion which triggers the limit. The broker code can be updated to avoid the request and bypass the limit.

### 3.3.3 Science

The science group currently contains two modules: a filter for likely extragalactic transients followed by a classifier for Type Ia supernovae. Both run on Cloud Functions and use the default settings, with the exception that the classifier requires 512 MB of memory. Both produce Pub/Sub streams as output, and the classifier also streams its results into BigQuery.

There is ongoing work to this pipeline, which is described in § 3.3.6. In addition, this pipeline is ready to be extended by the inclusion of additional classifiers and other modules that can annotate streams with information useful to end users.

#### **Extragalactic Transients module**

*Input* Pub/Sub stream: `ztf-alerts`

*Output* Pub/Sub topic: `ztf-exgalac_trans_cf`

This module filters the alert stream, republishing only the alerts likely to represent ex-  

---

complete alert packet should be understood to exclude the cutouts.

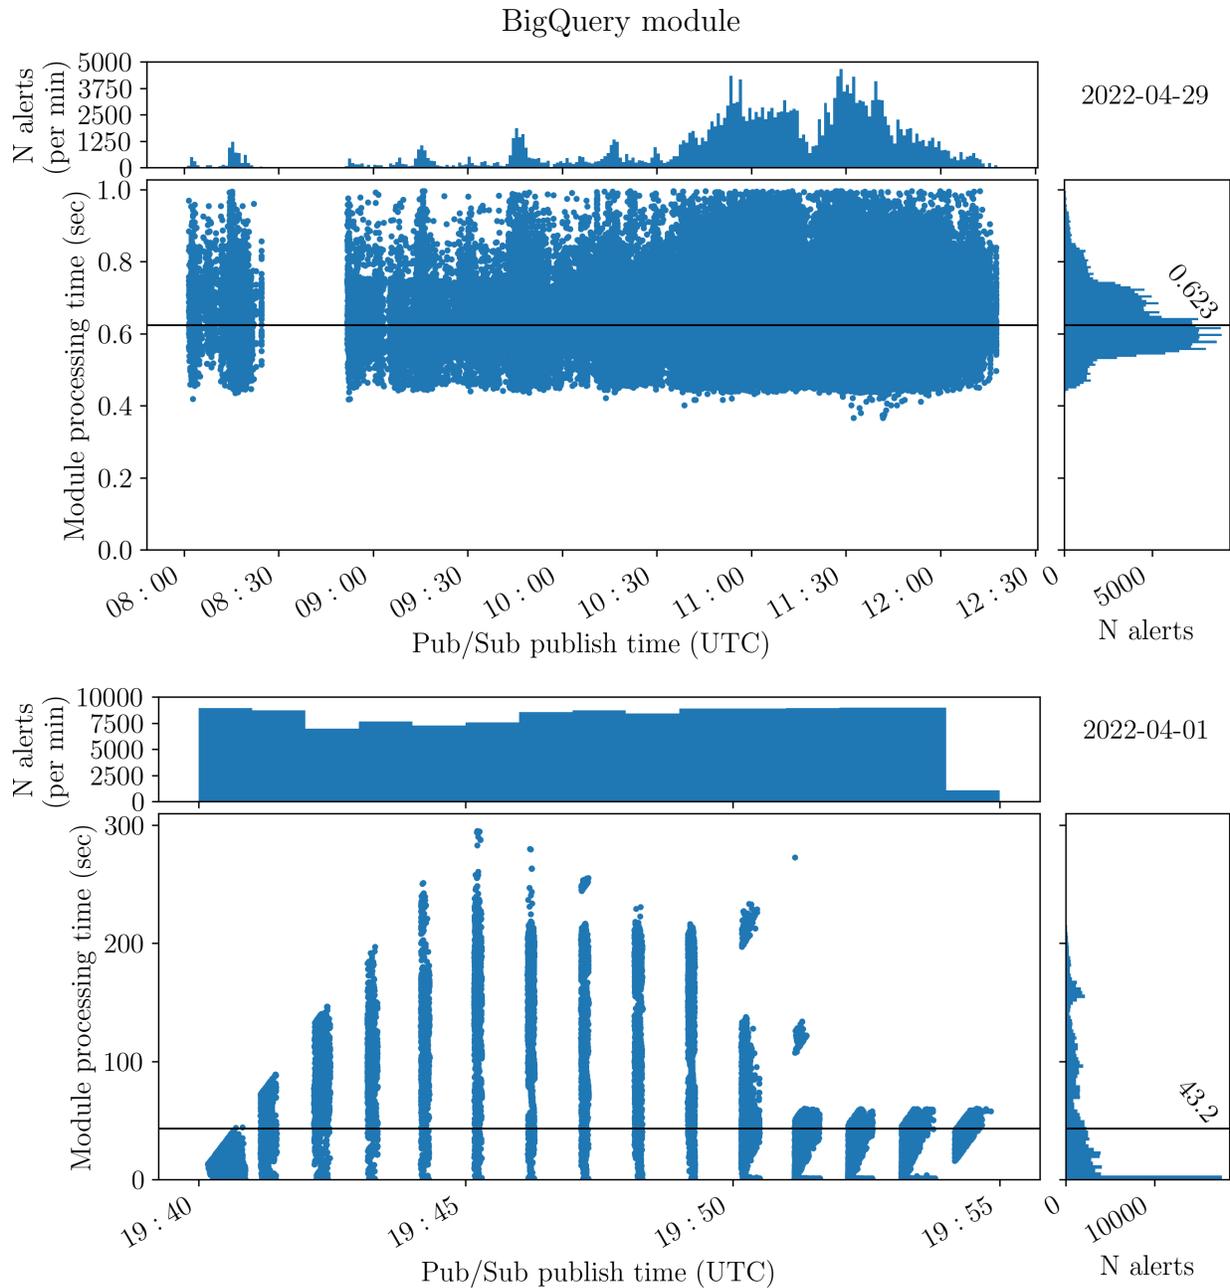


Figure 13: Same as Fig. 12 but for the BigQuery module. At typical ZTF rates the performance is similar to the Cloud Storage module. However, at higher incoming alert rates a Google Cloud limit on the rate of BigQuery API requests is exceeded. This results in Google throttling requests and the module is prevented from completing for a period of time. Once the wait period is over the module successfully processes a relatively small number of alerts before exceeding the limit again. This cycle causes the pattern seen in the figure. The limit does not apply to streaming inserts directly and the broker code can be updated to avoid the offending request and bypass the limit.

tragalactic transients. The determination is made using basic cuts on the data, for example, requirements that the difference image detection is positive and the object is not moving. It is an adaptation of a transients filter developed by the ZTF team<sup>32</sup>. The performance is similar to the SuperNNova module (below) but without the delay caused by BigQuery API requests at high incoming rates.

### SuperNNova module

*Input* Pub/Sub stream: `ztf-exgalac_trans_cf`

*Output* BigQuery table: `ztf_alerts.SuperNNova`

Pub/Sub topic: `ztf-SuperNNova`

This module computes a binary classification for Type Ia supernovae using SuperNNova [Möller and de Boissière, 2020]. This is a machine learning classifier that uses a recurrent neural network (RNN) with 2 hidden layers to predict the class, given the light curves. The model was trained by the SuperNNova developers on ZTF simulations.

### 3.3.4 Admin

Admin contains tools that support the broker pipelines. This includes a set of scripts that setup a new broker instance in Google Cloud and a library of python functions that are used throughout the broker. Both are under active development, which is described further in § 3.3.6.

The setup scripts are currently split between bash and python and use standard Google API calls to create Cloud resources and deploy code to compute services. The scripts assign unique names to all resources following a predetermined syntax. We are in the process of rewriting the setup module to use Terraform.

`broker-utils` is a pip-installable python package with a variety of tools used throughout the broker to transform data, connect to GCP resources, etc. I plan to move this package to the `pittgoogle-client` repo where it will form the basis for the next version of Pitt-Google's public client API.

---

<sup>32</sup>[https://github.com/ZwickyTransientFacility/ztf-avro-alert/blob/master/notebooks/Filtering\\_alerts.ipynb](https://github.com/ZwickyTransientFacility/ztf-avro-alert/blob/master/notebooks/Filtering_alerts.ipynb)

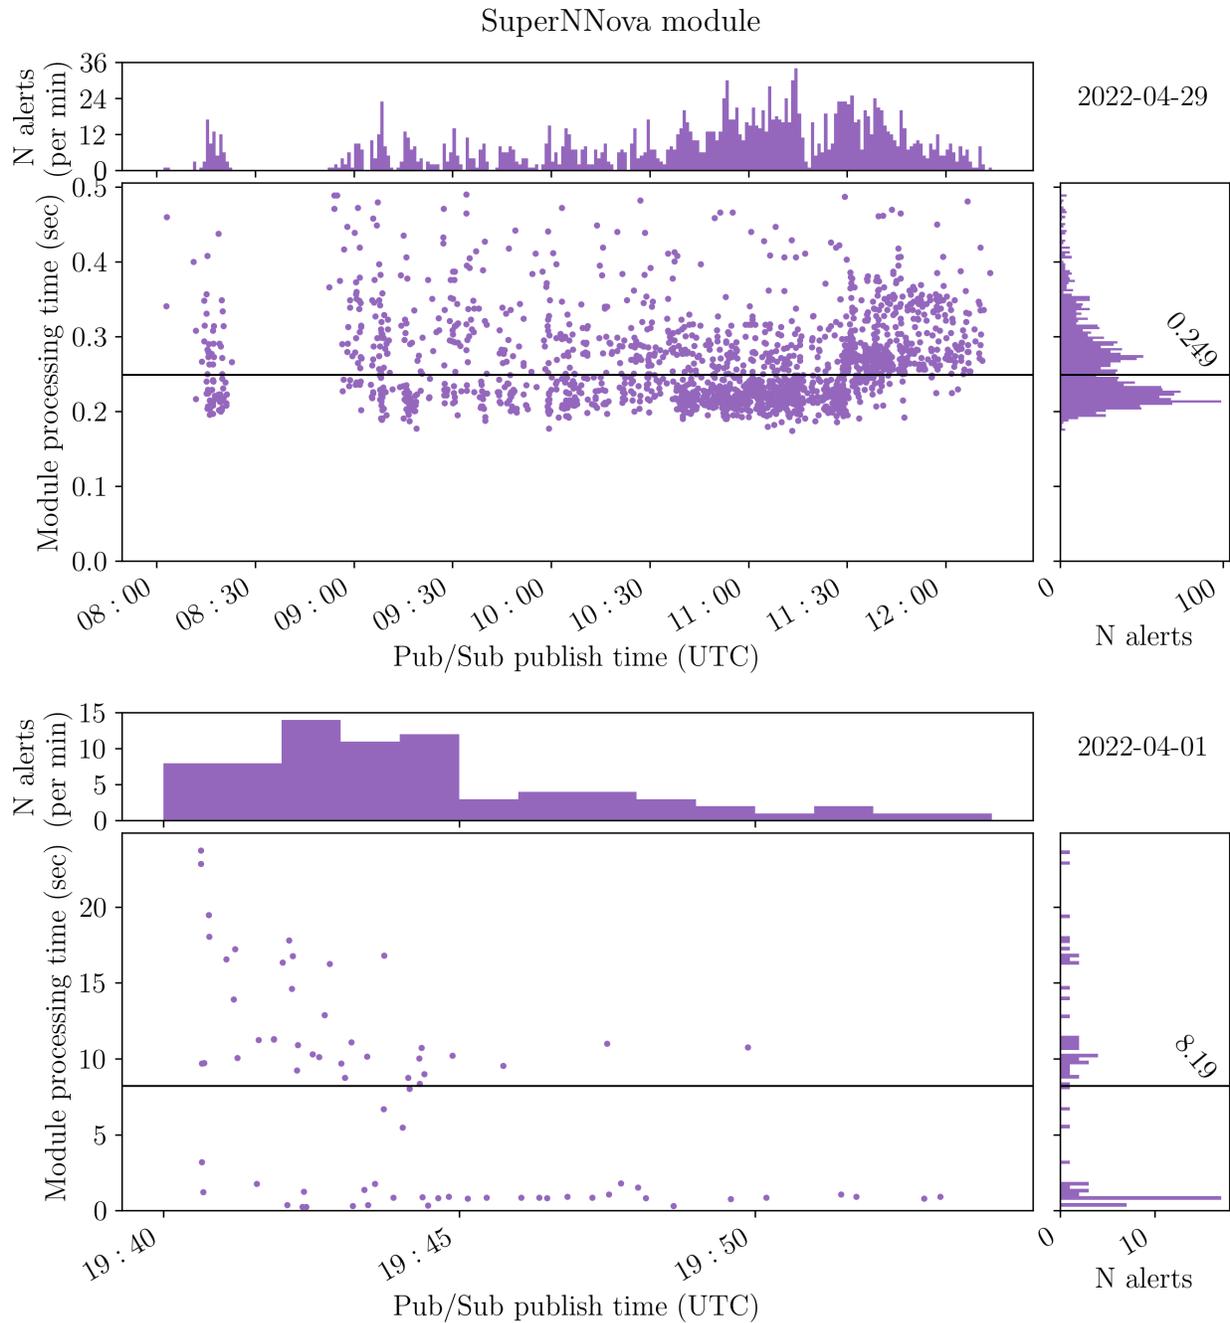


Figure 14: Same as Fig. 12 but for the SuperNNova module. The total number of processed alerts is lower because this module only classifies those that are likely extragalactic transients. At incoming rates typical of ZTF, performance is similar to previous modules. At higher incoming rates (bottom panel) the processing times are longer because this module sends its results to BigQuery, and so it experiences similar issues as seen in Fig. 13.

### 3.3.5 Periphery

The stream looper publishes a “heartbeat” stream, `ztf-loop`, intended for testing the connection. It is a VM that runs 24 hours a day and simply pulls and republishes recent ZTF alerts at a rate of 1/s.

The metadata collector pulls the output streams of all modules in the main pipelines. It extracts metadata from the messages and uploads it to BigQuery. This VM runs at the end of the night (i.e., every morning), it does not process streams in real-time.

### 3.3.6 Work in Progress and Future Plans

The broker is under active development. Here I describe some of the plans and prototypes that I have developed and work that is currently in progress. The numbers referenced throughout this section refer to open GitHub Issues detailing the plans, and Pull Requests (PRs) where the work is occurring. The Issues and PRs are in the main broker repo (<https://github.com/mwvgroup/Pitt-Google-Broker>) unless otherwise noted.

#### Design of the science pipeline

PR #154 is work in progress by an undergraduate student I am working with (Aaron Palmer) that will tag messages with categorizations (e.g., real/bogus, likely extragalactic transient, etc.) and other key data (e.g., binned magnitude, passband, etc.). These tags will be attached as custom attributes in the message metadata rather than to the message payload itself. This is intended to aid stream filtering (see § 3.2.2 and the discussion of metadata below). Issue #126 proposes to add additional tags.

We intend to process every alert through every module in the science pipeline. The tagger module (above) will replace the current extragalactic transient filter (see § 3.3.3) and the purity filter referenced by Issue #112. The tagger will not function as a filter – it will simply tag and republish every alert. Any future filters should be placed at the end of the pipeline, acting only to provide a stream containing a curated subset of alerts for end users.

Issue #83 proposes to put a module at the top of the science pipeline that publishes a lite, or “semantically compressed” version of the stream. Aaron Palmer recently began work on this. The lite stream will contain a message for every incoming alert but it will include

only a subset of the original data. This will help make it feasible to run every alert through every module in the pipeline, and will provide a stream for end users that can be more easily and cheaply listened to. The obvious tradeoff is that the end user will have to explicitly fetch any data needed for their analysis that was dropped by this by module. At a minimum, the lite stream must include all fields of the original alert that are accessed by modules the science pipeline. Determining a good balance of additional fields to include/exclude that will satisfy a broad range of use cases will likely be an ongoing process.

### **Access Rights**

Permissions in Google Cloud can be set with a very fine level of detail. PR #147 is work in progress that increases both our security and public availability. It improves the security of our projects by creating custom roles, each of which defines a set of permissions that is limited to the minimum necessary for the intended purpose. These can be assigned to developers and users as appropriate.

PR #147 also standardizes the configuration of public access rights for our data resources (Pub/Sub topics, BigQuery datasets, and Cloud Storage buckets). Specifically, it grants data access to `AllUsers` (defined as everyone on the internet), including permissions to attach subscriptions to topics, query datasets, make `GET` requests to buckets, etc<sup>33</sup>. We intend for all of our data resources to be public. However, this has not yet been standardized in our projects and some individual resources are not currently configured appropriately. This pull request will correct this.

### **ELAsTiCC challenge**

The Dark Energy Science Collaboration (DESC) is running a challenge for alert brokers called ELAsTiCC. They will be publishing a stream containing simulated LSST alerts. Brokers will process this stream and publish their classification results. DESC will then ingest the brokers' output streams containing classifications. This is largely intended as an opportunity for both brokers and DESC to get setup technologically for LSST and for brokers to develop classifiers with/for LSST-like data. PR #153 is work in progress by a graduate student I am working with (Christopher Hernandez) that updates the broker software to

---

<sup>33</sup>Note that many activities in Google Cloud (e.g., most API calls) require authentication, and granting permissions to `AllUsers` does not override this.

accommodate the ELAsTiCC schema, ingest the Kafka stream, and produce the requested classifications stream for DESC.

### **pittgoogle-client**

I have been developing a python API called `pittgoogle-client` (see § 3.4.1) to facilitate user access to our data resources. Various work is ongoing. PR #6 (in the client repo) adds a robust Pub/Sub client that performs a streaming pull and processes messages through a user-provided callback function. This is a general tool, but is being developed now so that DESC can use it in the ELAsTiCC challenge. This pull request includes a near-total rewrite of the documentation. See also, `broker-utils` (below).

### **broker-utils**

`broker-utils`<sup>34</sup> is a python library I have developed that is used internally by the broker. PR #150 contains a major rewrite of this library that updates dependencies and adds modules (e.g., `testing` and `types`). Issue #158 proposes to combine this library with the `pittgoogle-client`. This will simplify code management and provide end users with a more complete, mature, and robust set of tools.

### **tom-pittgoogle**

`tom-pittgoogle`<sup>35</sup> is a Pitt-Google plugin for the TOM Toolkit [Street et al., 2018]. It currently contains proof-of-concept implementations for three different access methods: Pub/Sub access via Python, Pub/Sub access via REST, and BigQuery access via Python. All three work, but the REST implementation has inherent limitations and doesn't add anything significant; it can be dropped. These all rely on consumers that were written specifically for the implementation and reside in the `tom-pittgoogle` repo. Issue #2 (in the `tom-pittgoogle` repo) proposes to replace these with `pittgoogle-client` consumers.

### **Integration testing**

PR #148 adds tools to perform integration testing. The test's workflow is to trigger a chain of one or more deployed modules by publishing alerts to the Pub/Sub topic at the top of the chain, then pull messages from a subscription attached at the end of the chain

---

<sup>34</sup>The current docs are at [https://pitt-broker.readthedocs.io/en/u-tjr-update\\_broker\\_utils/api/broker-utils/consumer-sim.html](https://pitt-broker.readthedocs.io/en/u-tjr-update_broker_utils/api/broker-utils/consumer-sim.html). Notably, they are not in the main branch; a problem that should be fixed by Issue #158.

<sup>35</sup>Repo: [https://github.com/mwvgroup/tom\\_pittgoogle](https://github.com/mwvgroup/tom_pittgoogle). Docs: <https://tom-pittgoogle.rtfid.io>

and validate that the input messages made it all the way through the chain. The test will validate that modules are functioning at a baseline level and are properly integrated with the rest of the pipeline.

### **Pipeline administration**

Issue #159 proposes to port all modules that currently run on Cloud Functions to Cloud Run. This should make code management easier and provide better performance at LSST scale.

PR #152 is work in progress that will port the broker's setup scripts to a Terraform<sup>36</sup> implementation. Currently, the setup uses a combination of bash and python scripts that setup Cloud resources and deploy modules by calling Google Cloud command-line and python APIs. Using Terraform will greatly simplify the deployment process, making it much easier to pick and choose which resources/modules are created/deployed. It will also greatly simplify the ongoing management of Cloud resources.

### **Rubin test stream**

PR #123 adds a connection to Rubin Observatory's test alert stream. Rubin will use a Kafka broker that is very similar to the one in use by ZTF. One major difference will be that Rubin's Avro-serialized alert packets will be schema-less. Brokers will need to cache the schema themselves. Ideally, our consumer would integrate directly with the schema registry used by Rubin. However, there are technical implementation issues (which are described in the pull request), and the pull request uses a work-around.

### **Cross matching**

Cross matching with large catalogs is a particular challenge in our pipeline because of the per-alert processing. I have developed three prototype cross-matching modules that each experiment with a different approach.

The first was implemented in Dataflow and attempted to batch alerts and then call the CDS cross-match service<sup>37</sup>. The batching never worked properly and we no longer use Dataflow. However, this method is feasible and we may want to revive it in the future, batching alerts for this module only and then returning to per-alert processing for downstream

---

<sup>36</sup><https://www.terraform.io/>

<sup>37</sup><http://cdsxmatch.u-strasbg.fr/>

modules.

PR #100 performs a cross match with the AllWISE<sup>38</sup> catalog that is hosted as a BigQuery public dataset<sup>39</sup>. This was very expensive to do on a per-alert basis (seen in 3.3.7). The discussion in pull request outlines some options to try, including optimizing the table structure (e.g., through partitioning and clustering) or using BigTable, which is likely to be better suited to this access pattern.

PR #106 uses a catalog of cataclysmic variables [Abril et al., 2020] that is small enough to be packaged and deployed with the module. It runs on Cloud Run and performs the cross match internally, avoiding the issues associated with calling external services on a per-alert basis. This module works, but would require some cleanup before being put in production. This approach could be extended to other small catalogs, but is not scalable to the multiple, large catalogs with which we would like to cross match.

### **Multi-messenger astronomy**

Issue #126 describes a use case for multi-messenger astronomy and outlines the necessary work. Fig. 15 diagrams the end-to-end process. The idea is for Pitt-Google to federate the LIGO-Virgo-KAGRA and LSST alert streams in a way that allows a researcher to listen for interesting gravitational wave (GW) events in the LIGO-Virgo-KAGRA stream and then trigger a custom search of LSST data for an electromagnetic (EM) counterpart. The search can monitor the LSST live stream going forward and also has immediate access to the alert stream history due to the streaming inserts to BigQuery. Issue #126 details the envisioned implementation, including both the functionality Pitt-Google needs to implement and the anticipated workflow and available options for the researcher. At a minimum, Pitt-Google would need to ingest a GW stream. Issue #121 proposes to implement access to LIGO/Virgo streams through a connection to the Gamma-ray Coordinates Network<sup>40</sup>.

### **BigQuery tables**

PR #141 restructures the way the results of classifiers/categorizers (science pipeline modules) are stored in BigQuery. It creates a new table called `classifications` that will hold a summary of the results from every module, and individual tables for each module

---

<sup>38</sup><https://wise2.ipac.caltech.edu/docs/release/allwise/>

<sup>39</sup><https://cloud.google.com/bigquery/public-data>

<sup>40</sup><https://gcn.gsfc.nasa.gov/>

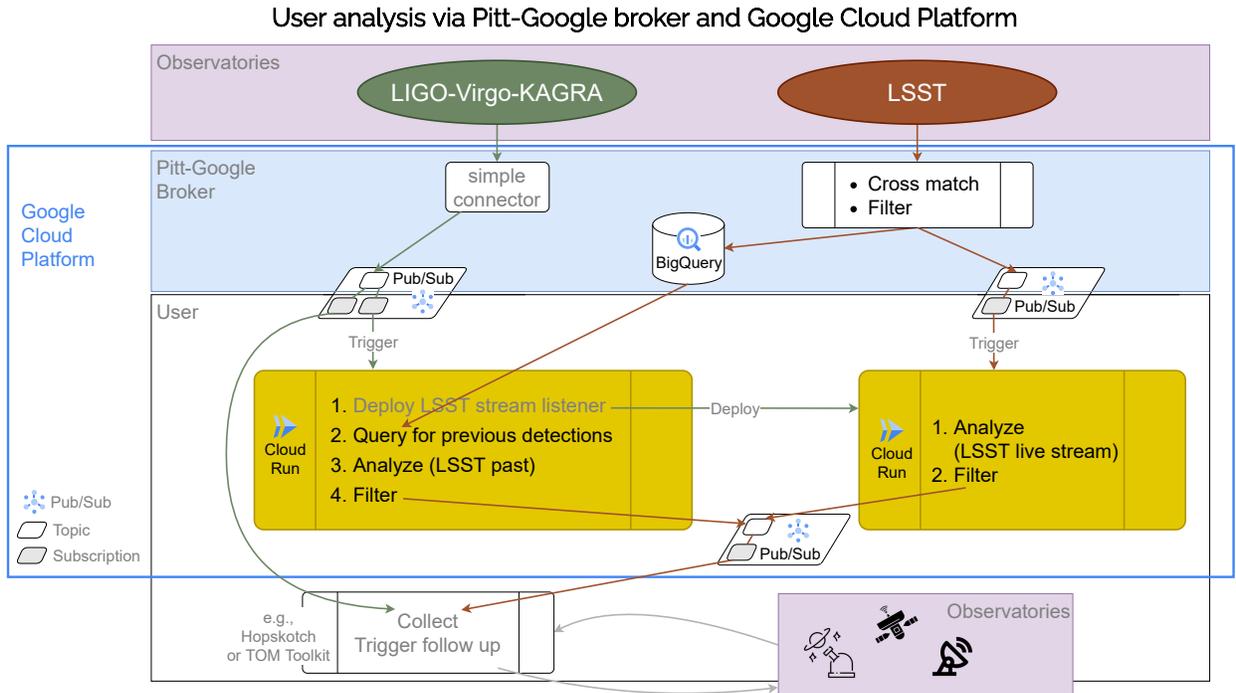


Figure 15: Diagram of the envisioned process by which researchers could hook into Pitt-Google’s alert streams and catalogs and use them to hunt for electromagnetic counterparts of gravitational waves, both forward and backward in time. The user module (gold) on the left listens to the GW stream, chooses an interesting event, then does the following: deploy the module on the right (injecting the GW data), query historical LSST data (all the way up to the given moment), analyze it, and send results to a user-owned Pub/Sub topic. The module on the right listens to the LSST live stream, analyzes it, and sends results to the same topic.

that will store detailed results.

Issue #102 proposes to calculate object statistics using DIA sources and store them in either BigQuery or BigTable. BigTable may provide a more efficient means to update on a regular basis. Aaron Palmer has recently picked this up and is attempting to create the table as a materialized view on one of the existing tables (`alerts` or `DIASource`).

### Metadata

Custom attributes that we attach in the message metadata can be used to filter subscriptions (§ 3.2.2). In the streams produced by the storage pipeline, these attributes include identifying information that will allow the listener to query/download the stored row/file. In the science pipeline, this currently includes only the object and source IDs, though PR #154 (discussed above) adds several more (e.g., binned magnitude) and Issue #126 (discussed above) proposes to add a HEALPix index to aid GW searches. A HEALPix index can also aid cross matching more generally. The magnitude is binned because these attributes must be strings, and thus the filters support a limited set of operations (discussed further in § 3.2.2). Issue #79 proposes to track code versions for classifiers and the overall broker version by adding them to metadata.

### 3.3.7 Pitt-Google’s Budget

In the LSST era, it may become common for individual researchers and collaborations to at least *want* to process ZTF-sized alert streams. Pitt-Google’s adoption of the cloud model offers a way to do this that has low overhead and low barrier to entry by providing researchers with direct access to cloud computing for their own purposes. However, the costs associated with cloud computing (discussed in § 3.2) can be a source of confusion and apprehension at first glance, for at least two reasons. First, the charges from the cloud provider are based on usage, and they tend to be very small but frequent. People may find it difficult to predict what their actual usage – and thus total cost – will be. Second, the overall budget for a project that runs in the cloud tends to breakdown along different categories than the budget of a project using more traditional processing methods on local machines, department clusters, etc. This may include percentages that go to FTEs (humans),

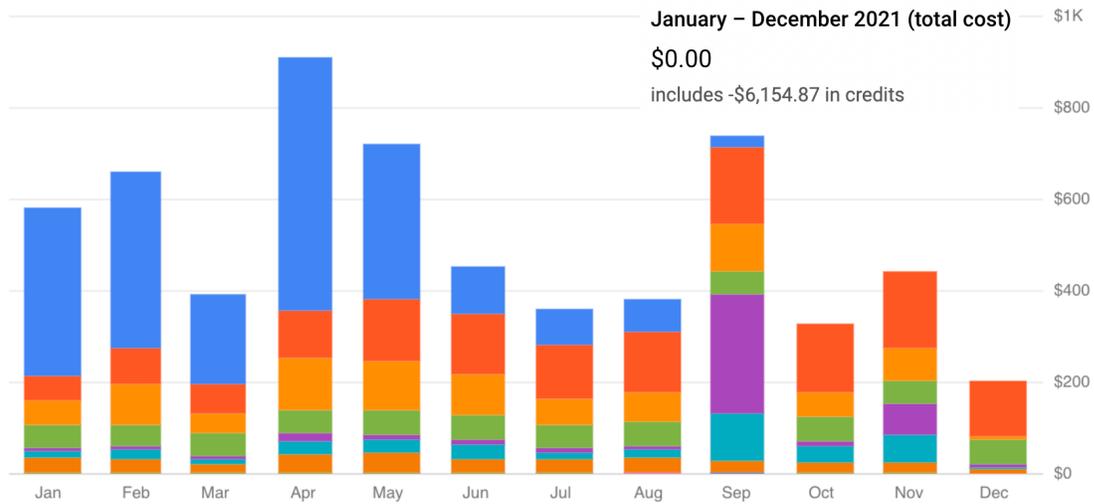
equipment/networking and related maintenance, a cloud provider, etc. This breakdown may also be difficult to predict. Here, I provide details of the Pitt-Google broker’s budget. It is intended to give people a sense of how the cloud computing costs break down, and what the total amounts end up being to develop the software described in this paper and run the ZTF alert stream.

We currently pay approximately \$500 per month for Google Cloud services. This includes running the full ZTF stream through the pipelines described in § 3.3 and all related testing and development. To date (1 March 2020 through 25 April 2022), we have incurred a total of \$8,548 in Google Cloud charges, of which \$7,527 was covered by research credits and other options (described above). All code development has been done by students at the University of Pittsburgh. Fall 2021 was the first time anyone was paid to work on the Pitt-Google broker.

We began processing the nightly ZTF alert stream on a consistent basis in late 2020. Here I describe our total budget for 2021. Fig. 16 shows a breakdown of our Google Cloud charges by month and service. We incurred \$6155 in total cost, but it was covered entirely by research credits (referenced above). This included running the full ZTF stream plus all development and testing. Our only other budget item in 2021 was one FTE (graduate student) for one semester.

### 3.4 DATA ACCESS

Pitt-Google broker is currently providing ZTF alert data in streams, catalogs, and Avro files, as described in § 3.3. Many access methods are available. This section discusses Pitt-Google’s custom APIs. Google Cloud’s native tools were discussed in § 3.2.6. For any access method, users need 1) the name of the Pitt-Google resource, and 2) authentication to Google Cloud. Resource names are given in § 3.3. Authentication is obtained directly from Google, and is discussed in § 3.2.7. Instructions are also provided in the online documentation for the `pittgoogle-client`, discussed below.



Service	↓ Cost	Discounts	Promotions and others	Subtotal
Cloud Dataflow	\$2,110.66	\$0.00	-\$2,110.66	\$0.00
Cloud Storage	\$1,432.20	\$0.00	-\$1,432.20	\$0.00
Cloud Pub/Sub	\$853.48	\$0.00	-\$853.48	\$0.00
Cloud SQL	\$612.91	\$0.00	-\$612.91	\$0.00
BigQuery	\$429.26	\$0.00	-\$429.26	\$0.00
Cloud Functions	\$371.90	-\$35.44	-\$336.46	\$0.00
Compute Engine	\$337.70	-\$6.90	-\$330.80	\$0.00
Cloud Scheduler	\$4.61	\$0.00	-\$4.61	\$0.00
Cloud Run	\$1.86	-\$1.86	\$0.00	\$0.00
Cloud Build	\$0.29	\$0.00	-\$0.29	\$0.00

Figure 16: Pitt-Google’s total charges incurred on Google Cloud in 2021, by month and service. This is a screenshot of our billing page and includes all GCP projects, services, and SKUs associated with our billing account. The work represented here includes running and storing the full ZTF alert stream plus all development and testing. The total cost was \$6154.87, and was offset 100% by Google Cloud research credits. The BigQuery charge in September is mostly due to a two-day test of a cross-match module, which queried a BigQuery table for every alert. Our transition from Dataflow to Cloud Functions can be seen in their respective trends, as well as the overall cost difference between the two.

### 3.4.1 `pittgoogle-client`

`pittgoogle-client`<sup>41</sup> is a Python library of custom tools for accessing Pitt-Google’s Cloud data resources and working with the data. Tools that access our Cloud resources are wrappers for Google’s APIs, discussed in § 3.2.6. These tools simplify the connections to data resources while also exposing the underlying Google API so that the user can call it directly at any point where they want more control. An example of this is given below.

The library can be installed using:

```
1 pip install pittgoogle-client
```

It is then imported in Python as:

```
1 import pittgoogle
```

The online documentation includes instructions for obtaining authentication credentials from Google, a description of the available data, names of the data resources, and instructions for using each module and class. Work on a major upgrade to this library is nearly complete (referenced in § 3.3.6). Many of the features and functionality described in the remainder of this section will be introduced in this upgrade.

The library supports authentication via two methods: a service account and OAuth 2.0. Current modules include `auth`, `bigquery`, `figures`, `pubsub`, `types`, and `utils`. Below I describe how to use this library to process an alert stream. The example assumes that setup has been completed using the recommended authentication method in § 3.2.7.

The `pittgoogle.pubsub` module includes a consumer that will execute what Pub/Sub calls a “streaming pull”. This is the recommended way to process a live alert stream. It is an asynchronous method that runs in a background thread and processes messages through a user-supplied callback, allowing the user to run a listener indefinitely without requiring the application to block while waiting for new messages. As demonstrated below, the user

---

<sup>41</sup>Repo: <https://github.com/mwvgroup/pittgoogle-client>. Docs: <https://pittgoogle-client.rtfid.io>

provides a callback and then calls the consumer's `stream()` method. The callback should accept a `pittgoogle.types.Alert` as input, process it and save results as desired, then return a `pittgoogle.types.Response`. Here is an example of a basic callback that saves the alert data to a file if the `objectId` matches a given string:

```
1 def my_callback(alert):
2     # alert is a pittgoogle.types.Alert
3     # with attributes determined by the consumer's unpack keyword
4     metadata_dict = alert.metadata
5     alert_dict = alert.dict
6
7     try:
8         # process the alert and store needed results
9         if metadata_dict["objectId"] == "ZTF18acvgznv":
10            with open("ZTF18acvgznv.txt", "w") as fout:
11                fout.write(str(alert_dict))
12
13    except:
14        # processing failed, so do not acknowledge the message
15        ack = False
16
17    else:
18        # processing succeeded, so acknowledge the message
19        ack = True
20
21    # return a Response
22    return pittgoogle.types.Response(ack=ack)
```

If processing was successful, the message should be acknowledged using `ack = True`. If it is not, Pub/Sub will redeliver the message at some arbitrary time in the future.

Here is the code to execute a streaming pull using this callback:

---

```

1 from pittgoogle.pubsub import Consumer
2
3 consumer = Consumer(
4     user_callback=my_callback,      # callback that will process alerts
5     unpack=["dict", "metadata"],    # data available to the callback
6 )
7
8 consumer.stream()                  # blocks the main thread by default

```

---

Instantiating a consumer will authenticate the user to Google Cloud, create a subscriber client, and make sure the client can connect to the subscription, creating it if necessary. There are various options that can be sent as keyword arguments to, for example, control which subscription the client connects to, how many messages it pulls before closing the stream, etc. By default, the call to `consumer.stream()` will block the main thread while it pulls and processes a small number of messages through the callback, and then close the background thread and return. These defaults are intended to support the user while developing the callback, especially to prevent the streaming pull from running out of control. *Ctrl + C* can be used at any time to regain control of the terminal – the consumer will attempt to close the stream and exit the background thread gracefully before returning.

Once the callback is working as desired, the streaming pull can be run indefinitely without blocking the main thread by using:

---

```

1 consumer = Consumer(
2     user_callback=my_callback,      # callback that will process alerts
3     unpack=["dict", "metadata"],    # data available to the callback
4     max_results=None,              # pull and process alerts indefinitely
5 )
6
7 consumer.stream(block=False)        # don't block the main thread

```

---

In this case, the consumer's `stop()` method can be used to close the stream and back-

ground thread:

---

```
1 consumer.stop()
```

---

The module is written in a way that allows the user to mix and match `pittgoogle` calls with direct calls to the underlying Google Cloud API. This allows the use of Pitt-Google’s wrappers when convenient without sacrificing the detailed control allowed by direct use of Google’s API. For example, the consumer created in the previous example uses a `google.cloud.pubsub_v1.subscriber.client.Client()` to manage the subscription, streaming pull, etc. This Pub/Sub client is exposed as `consumer.client.client`. So, if the user wants `pittgoogle` to handle the setup tasks like creating the client and making sure the subscription exists, but wants more control over the details of the streaming pull (e.g., parallelization), this can be done using the Pub/Sub client’s `subscribe()`<sup>42</sup> method, through the consumer:

---

```
1 consumer.streaming_pull_future = consumer.client.client.subscribe()
```

---

Appropriate arguments should be sent with `subscribe()`. Assigning the result to `consumer.streaming_pull_future` allows the use of Pitt-Google’s `consumer.stop()` to close the stream.

### 3.4.2 tom-pittgoogle

`tom-pittgoogle`<sup>43</sup> is a plugin for the TOM Toolkit [Street et al., 2018]. The TOM Toolkit is a software package that facilitates researchers setting up a database-driven system with a web front-end or GUI that allows them to aggregate and analyze data and connect with follow-up facilities. The `tom-pittgoogle` repo currently contains 3 working proof-of-concept implementations that connect our streams and catalogs into the system. Plans for continued development are discussed in § 3.3.6.

---

<sup>42</sup>[https://googleapis.dev/python/pubsub/latest/subscriber/api/client.html#google.cloud.pubsub\\_v1.subscriber.client.Client.subscribe](https://googleapis.dev/python/pubsub/latest/subscriber/api/client.html#google.cloud.pubsub_v1.subscriber.client.Client.subscribe)

<sup>43</sup>Repo: [https://github.com/mwvgroup/tom\\_pittgoogle](https://github.com/mwvgroup/tom_pittgoogle), Docs: <https://tom-pittgoogle.rtfid.io>

### 3.5 SUMMARY

Pitt-Google has been selected as a full-stream community alert broker for LSST. The software is under active development using the ZTF alert stream. Our founding motivations are to provide broad public access to the full LSST alert stream and to facilitate scientific collaboration with a low barrier to entry through Google Cloud’s advanced, big-data solutions. We have been running the ZTF alert stream consistently since late 2020.

We are providing ZTF alerts and value-added data through the Google Cloud services Pub/Sub, BigQuery, and Cloud Storage. These services, and the Google Cloud model in general, were described in § 3.2. There are many ways to access the data, both within and outside of the Cloud. Google’s own web console and APIs, and the process of using managed-compute services for real-time stream analysis, was discussed in § 3.2.6. Pitt-Google’s custom APIs were presented in § 3.4, along with code examples and links to our online documentation and tutorials. The custom APIs are Python libraries and include `pittgoogle-client`, our core set of user tools, and `tom-pittgoogle`, our plugin for the TOM Toolkit [Street et al., 2018]. They provide the user with data access via convenience wrappers around the relevant Google Cloud APIs, alongside other tools to unpack, visualize, and work with the alert data. Regardless of the access method, users will obtain authentication directly from Google. The process is detailed in our online documentation.

The current status of the broker software, plus ongoing work and future development plans, was presented in § 3.3. There is work in progress on nearly all aspects of the broker. The software runs on Google Cloud compute services, and exposes the data publicly at each stage of processing. The consumer ingests ZTF’s alert stream and republishes it to a public Pub/Sub topic. It publishes a byte-for-byte copy of the message payload (i.e., the alert packet/data), and attaches the Kafka metadata (e.g., publish time and topic) as custom attributes in the Pub/Sub metadata. The storage and science pipelines are triggered via push subscriptions on the consumer’s topic. The storage pipeline streams alert data to BigQuery and Cloud Storage where it is available immediately for queries and/or download.

Both pipelines use what is known as a “microservices” model. Each module in a pipeline is independently deployed to a managed-compute service like Cloud Functions and configured

(via Pub/Sub push subscriptions) to trigger automatically on messages published by the previous module in the pipeline. All processing is per-alert; there is no batching. The service automatically scales to meet the incoming demand and alert processing occurs in parallel for a given module, when needed. This model can be replicated by users who wish to process the live alert streams, and can be easily scaled for an individual running a single module to a full-scale downstream broker.

The science pipeline currently filters for extragalactic transients and then computes a binary classification for Type Ia supernovae using SuperNNova [Möller and de Boissière, 2020]. Notably, the science pipeline is being redesigned to include more modules and to run every alert through every module (rather than filtering before classifying). This work is discussed in § 3.3.6. A module is being added that will publish a “lite” version of the alert stream containing every alert but with a subset of the original data, allowing it to be listened to more cheaply. Another module is being added that will tag each alert with basic data like a binned magnitude and categorizations like real/bogus, attaching them as custom attributes in the message metadata where they are easily accessible to filters. There is also ongoing work toward adding cross matching modules.

## 4.0 CONCLUDING REMARKS

LSST’s data will be abundant and rich, and will throw the door open to revolutions across astronomy. But along with the multitude of opportunities, LSST’s unprecedented datasets will also bring new technical and scientific challenges for the community.

As shown in Chapter 2, the effect of DM in stars may be observable in the isochrones of star clusters residing in rich DM environments. These may show differences to key features such as the location of the MS turnoff, the presence or absence of a “convective hook”, or the luminosity of the tip of the red giant branch. As discussed in § 1.1, these features are also affected by, e.g., the age and metallicity of the cluster and are often used to compare with the results of stellar evolution models. Thus, at a minimum, care should be taken in the interpretation of isochrones when the cluster’s DM environment is either uncertain or amenable to large capture rates. Combined with independent measurements of a cluster’s age and/or metallicity, these isochrones may reveal the nature of DM.

Conditions in both high-redshift galaxies and the centers of larger galaxies (e.g., the Milky Way) are thought to be such that DM capture rates would be high. However, the effects will not be directly observable for the former, and dark matter density profiles are difficult to constrain for the latter. Dwarf satellite galaxies may offer the best chance to observe these effects. Dwarfs in the Local Group are known to have large mass-to-light ratios and their DM is well-constrained to have high densities and low velocity dispersions, amenable to large capture rates. LSST, with its wide survey area and faint limiting magnitudes, will observe hundreds of dwarf galaxies. Further, its excellent astrometric precision will support measurements of their internal motions, and by extension constraints on their density profiles [Ivezić et al., 2019]. Coupled with spectroscopic follow up (e.g., to estimate metallicity), this has the potential to reveal the nature of dark matter. Even barring such definitive results,

ADM has the potential to contaminate the interpretation of cluster isochrones from LSST by mimicking effects caused by distance or metallicity.

The Pitt-Google alert broker, presented in Chapter 3, is preparing to serve the astronomy community by processing the full LSST alert stream, enriching it with additional data, and redistributing it through channels and formats that simplify its further analysis by individual researchers and collaborations. The broker pipelines are designed to expose the data at every step, allowing researchers to hook into an alert stream that has been preprocessed to a level of their choosing. Our custom Python library, `pittgoogle-client`, simplifies data access, and the documentation provides tutorials. The `tom-pittgoogle` library integrates our data resources with the TOM Toolkit, a popular tool in astronomy for data collection and follow-up management. Many of the broker’s implementation details are designed to make the data easily accessible. Attaching summary data to the message metadata provides easy access for stream filters. The immediate SQL access to alert stream history, made possible by the broker’s streaming inserts to BigQuery, may prove helpful for uses cases like multi-messenger astronomy and searches for the electromagnetic counterparts of gravitational waves. This will allow researchers to immediately look both forward (via Pub/Sub streams) and backward (via BigQuery tables) in time with no gap in the data. Future development, including the addition of new modules that classify the objects in the alerts and cross match them with static catalogs, will further increase the usefulness of our broker’s services.

Cloud computing can expand access to astronomy’s emerging, big-data frontier, bringing it to those who would not otherwise have the resources or technical know-how to process data at this scale. The Pitt-Google broker’s unique approach is the first to directly enable this type of access. By providing the data through Google Cloud’s services like Pub/Sub and BigQuery, our broker is opening access to the alert data through all of the tools and services that Google has developed for large-scale data analysis. Individuals will be able to use Pub/Sub in combination with managed compute services like Cloud Functions to process LSST’s live data stream without having to manage the hardware, runtime environment, client connection, scaling to meet demand, etc. This can allow researchers to focus on developing their analysis code and evolving it quickly by testing on the live alert stream each night. This is the pace that will be required to best take advantage of the coming opportunities,

especially early in the survey when LSST begins to flood the unexplored parameter space with rich new data.

## APPENDIX A

### OBSERVING STELLAR EVOLUTION WITH LSST

#### A.1 ISOCHRONES

A single molecular cloud produces a *cluster* of stars, numbering dozens to hundreds of thousands, with essentially the same age and chemical composition. This allows us to study stellar evolution using a snapshot of a cluster, by plotting all of the member stars together on an HRD and comparing to the predictions of models. Since all stars in a cluster are the same age, their HRDs are called isochrones.

LSST will conduct the first systematic survey of star clusters in the southern sky in a volume that is unprecedented anywhere. The isochrones produced from the thousands of clusters it is expected to observe will span the spaces of age and metallicity for the first time [LSST Science Collaboration et al., 2009]. This will allow for crucial comparisons while fixing key drivers of the observed variation.

Well-populated isochrones can be used to study how stars evolve on timescales that are much longer than human life. This is because (a) the underlying (i.e., in a cluster with an infinite number of stars) distribution of ZAMS masses is continuous, and (b) stars within relatively large mass ranges undergo nearly the same evolution at different rates, and the rates scale smoothly with the ZAMS mass (see § 1.1.1). And so, with a large enough sample of stars of the same age, we can infer the evolution of a single star over time periods up to billions of years.

Fig. 17 shows the isochrones of 32 well-observed clusters colored by age, demonstrating

the evolution of populations over time. Stars on the MS line up diagonally in order of decreasing mass (top-left to bottom-right) with kinks that track those of the homology relations (as seen in Fig. 1). This is a well-populated feature on an HRD, since stars spend most of their lives here. An individual star leaving the main sequence would migrate to the right, and then perform a series of trips through the upper right portion of the diagram as it evolves through the giant branches (see § A.2.2). The shapes of isochrones track this evolution and contain many useful features. The mass of stars that are just turning off the main sequence indicates the cluster’s age. The morphology of the turnoff is affected by atomic diffusion [Gaia Collaboration et al., 2018], and the “convective hook” seen in Fig. 18 tracks the contraction of intermediate (§ A.2.2.3) and high mass stars (§ A.2.2.5) at the end of the MS – both revealing important clues about stellar structure. The hook’s location and morphology can be compared to model predictions and used to tune the parameters that encapsulate uncertain processes like overshooting (discussed in § 1.1.1). The location of the tip of the red giant branch is relatively insensitive to mass or metallicity and can be used as a “standard candle” to measure the cluster’s distance.

## A.2 LIGHT CURVES AND THE ALERT STREAM

*Light curves:* A light curve is a time series of photometric data, specifically the apparent brightness within a small range of wavelengths, determined by the telescope’s filter. Contrary to isochrones, these require repeated observations and reveal the evolution of an individual star over a very short timescale, relative to the star’s life. They are commonly used in the study of transients and variable stars for both classification of individual objects and statistical studies of populations. These objects often evolve differently in different filters, thus light curves in multiple filters (such as will be provided by LSST) can give important clues about the astrophysical sources.

*Alert stream:* The evolution seen in light curves is happening on human timescales which offers a unique opportunity for detailed study of particularly interesting events. Researchers can identify them as they happen and trigger additional follow-up observations to gather

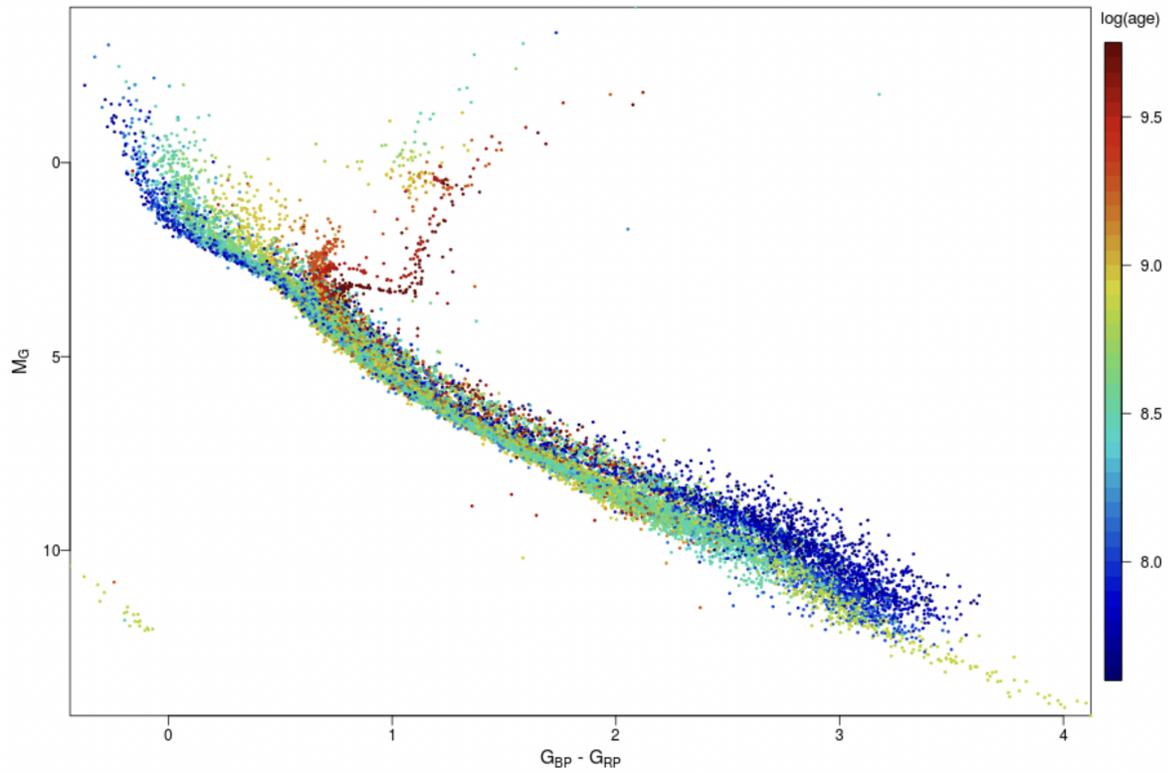


Figure 17: Gaia isochrones of 32 well-measured star clusters, colored by age. The evolution of several distinctive and useful features can be seen including the slope of the MS, the location of the MS turnoff, and the tip of the red giant branch. Credit: Gaia Collaboration et al., A&A, 616, A10 6, 2018, reproduced with permission © ESO.

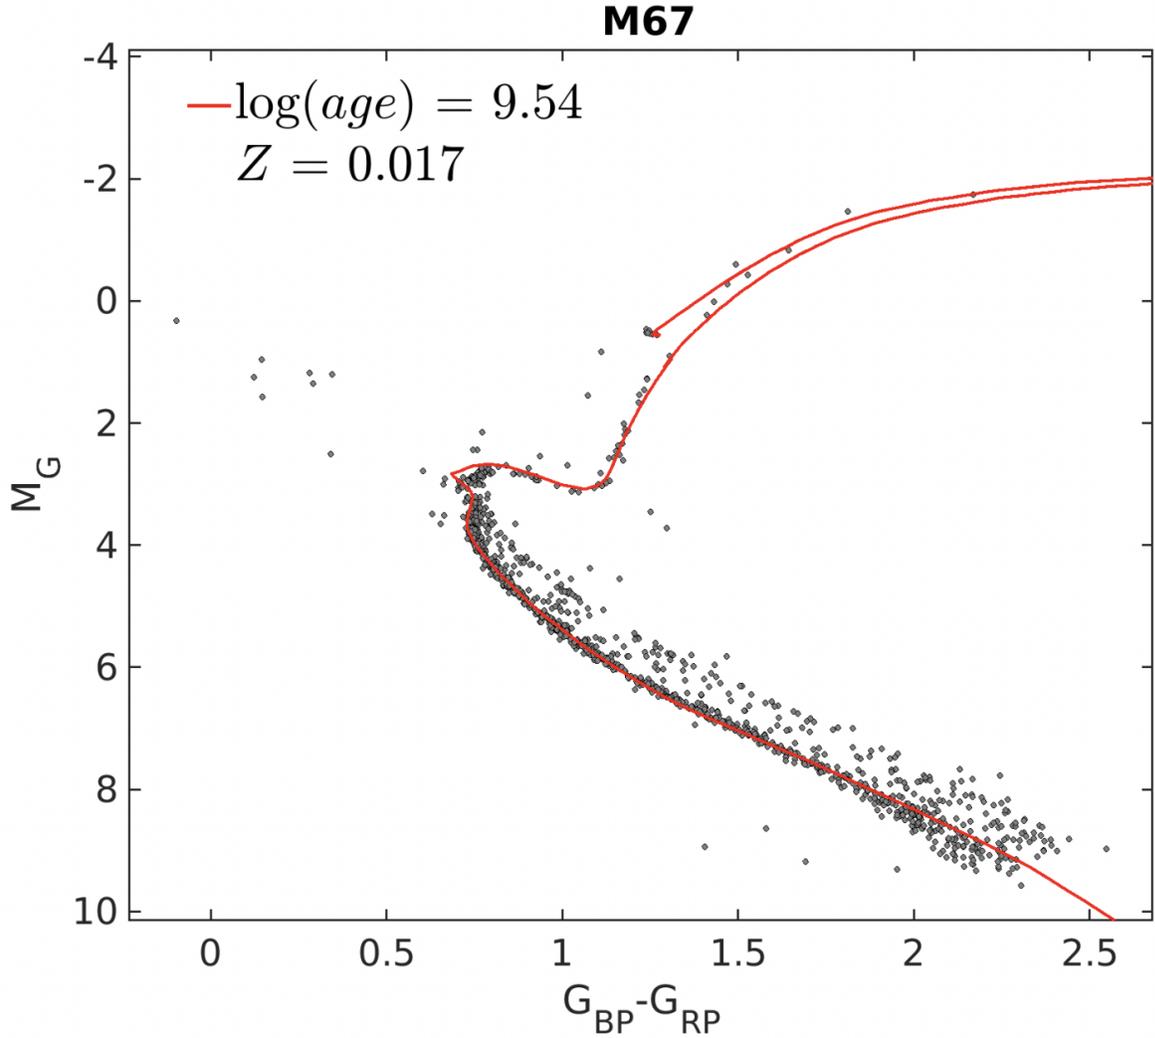


Figure 18: Gaia isochrone of the M67 cluster overlaid with a model generated by PARSEC [Marigo et al., 2017] for comparison. The model reproduces the MS slope and turnoff location reasonably well. Note that the MS is well-populated underneath the line representing the model; the scatter above the line is caused by the combined light of stars in unresolved binaries. The MS turnoff mass is about  $1.28 M_{\odot}$  [Mowlavi et al., 2012], right near the boundary where cores become convective. The discrepancy in the hook features likely indicates that the model’s treatment of overshooting (discussed in § 1.1.1) needs better calibration, as the quality of Gaia’s data places the feature very precisely [Gaia Collaboration et al., 2018]. Credit: Gaia Collaboration et al., A&A, 616, A10 15, 2018, reproduced with permission © ESO.

valuable information with different filters, spectral resolution, and/or cadence. For this reason, surveys like LSST are motivated to release data on all actively changing sources as a real-time message stream, called an alert stream, in addition to the catalogs they publish in standard data releases (approximately yearly). Each message is an “alert” about a single observation of a single object/event (called a source). These sources are found by performing image subtraction, i.e., subtracting a template of the same region of sky to reveal all sources which have changed in either brightness or position.

LSST’s alert stream will be more than an order of magnitude larger than the current leader, ZTF. Each night, it is projected to include  $\sim 7 \times 10^6$  variable stars and  $\sim 2 \times 10^5$  supernovae alongside (non-stellar objects)  $\sim 6 \times 10^4$  active galactic nuclei and  $\sim 2 \times 10^6$  solar system objects [Graham et al., 2019, Bellm et al., 2019].

The remainder of this section is devoted to the stellar activity that will appear in the alert stream and is structured as follows. In § A.2.1 and § A.2.2 I discuss binary systems and post-MS evolution respectively as these are dominant causes of stellar events that are revealed by image differencing. In § A.2.3 and § A.2.4 I discuss transients and variable stars respectively and give specific examples of each that will appear in the stream.

### A.2.1 Binary systems

Most stars live in binaries, gravitationally bound to a companion and orbiting their common center of mass. If they are separated by a relatively large distance this does not affect their evolution. But if they are closer together, their lives can become quite dynamic. When one of the stars evolves into a giant (via normal evolution, § A.2.2) its atmosphere may puff out so large as to envelop its companion, setting off a chain of (not-well-understood) events that leads to inspiral of the stars. This can end with a much closer binary system that has ejected the common envelope, or a dramatic merging of the two stars. A similar process which does not reach the common envelope configuration can instead lead to mass transfer as material from one star experiences more gravitational force from the other. Some examples are X-ray binaries and cataclysmic variables (CVs, § A.2.4.4). Very close binaries can produce gravitational waves as they orbit at high speeds in a small space, sending

out ripples through space. This removes energy from the system and causes further inspiral, likely resulting in a merger. Some examples are kilonovae (§ A.2.3.2) and Type Ia supernovae (§ A.2.2.4). Many of these events are difficult to observe and relatively poorly understood. But they will show up frequently in LSST’s alert stream, ready to be found and studied.

## A.2.2 Post-Main Sequence Stars

The gravitational collapse of a star can be halted only by nuclear burning or the degeneracy pressure provided by electrons or neutrons. When hydrogen runs out in the core, the star begins to contract, and it leaves the MS. The details of what follows vary, mostly depending on the star’s mass, but the general process is that it will go through cycles of core contraction interrupted by burning, either in the core or in a shell between the core and the atmosphere. The atmosphere responds by expanding and cooling, with intermittent phases of contraction and heating. It becomes very bright and we call it a “giant” or a “supergiant”, often qualified with “red” or “blue” indicating that the surface is relatively cool or hot respectively. This is often a very active phase and the star may pulsate (§ A.2.4.1), ignite fuel unstably in a flash, and/or lose a significant amount of mass via winds in its extended atmosphere.

When it can’t burn any more fuel, it ends its life in one of the following ways:

1. as a *white dwarf*, with  $M_{\star} \lesssim 1.4 M_{\odot}$ , supported by electron degeneracy pressure, slowly radiating away the heat left over from nuclear burning.
2. as a *neutron star* – the result of an extremely bright explosion (a supernova) that leaves behind a core supported by neutron degeneracy pressure, which slowly radiates its leftover heat away.
3. as a *black hole* – the result of either direct and complete collapse or an explosion that left behind a core too massive to be supported by neutron degeneracy pressure and thus collapsed.
4. in an explosion that tears the star apart completely, leaving *no remnant*.

The first three are remnants which can interact further with a companion, usually by accreting mass or merging, often with explosive results.

**A.2.2.1 Very low mass stars - Red giants and white dwarfs.** Stars with ( $M_\star \lesssim 0.3 M_\odot$ ) are very faint and are also the most numerous in the Universe. Since they are dominated by convection they can burn most of their hydrogen, and they burn it very slowly. The convection also stirs up the star’s magnetic fields, and so very low mass stars are more likely to flare (§ A.2.4.3). They live for  $\gtrsim 10^{12}$  years, and so none will have evolved off of the MS at the Universe’s current age. They are expected to transition slowly into red giants as their hydrogen burning slowly ceases. They will not reach temperatures high enough to fuse helium. They will slowly become white dwarfs and continue radiating leftover thermal energy for a very long time.

**A.2.2.2 Low mass stars - Red giants and white dwarfs.** Stars with ( $0.3 M_\odot \lesssim M_\star \lesssim 2 M_\odot$ ) transition smoothly from core to shell hydrogen burning. Their cores slowly contract until they become dominated by electron degeneracy pressure. They are not massive enough to contract any further unless they accrete matter from a companion. If they are heavy enough to burn helium [ $M_\star \gtrsim 0.8 M_\odot$ ; Pols, 2011] it ignites in a flash at the so-called tip of the red giant branch. They lose mass through winds from their extended atmospheres. This is mostly hydrogen but can include heavier elements that were produced during fusion and dredged up by deep convective envelopes. Towards the end of their lives they blow off the rest of their atmospheres, leaving their (mostly carbon and oxygen) cores exposed. They become white dwarfs.

**A.2.2.3 Intermediate mass stars - Red giants and white dwarfs.** ( $2 M_\odot \lesssim M_\star \lesssim 8 M_\odot$ ) have strongly convective cores on the MS which means that they are well-mixed and hydrogen burning ends fairly rapidly throughout the entire core. The entire star contracts in response until the temperature and density in a shell which is still hydrogen-rich is hot enough to ignite. This causes a feature on an HRD<sup>1</sup> called the convective hook, discussed

---

<sup>1</sup>A Hertzsprung–Russell diagram (HRD) is one of the most widely used visualizations in astronomy and comes in two forms. Theorists plot the (surface) luminosity and temperature, while observers plot magnitude and color. Though the mapping between the two is non-trivial it is monotonic. They convey essentially the same information and both display the features discussed in the text. Star catalogs are often sliced in various ways and plotted on HRDs to study, e.g., the effects of metallicity or the statistical features of stars in a given mass range.

in § A.1. Their evolution is similar to low mass stars except that burn helium stably. Their higher central temperatures and burning rates kept their cores less dense than those of low mass stars for most of their lives, but eventually they contract enough to be supported by electron degeneracy pressure and they also become white dwarfs.

**A.2.2.4 White dwarfs that exceed  $1.4 M_{\odot}$  - Type Ia supernovae.** The limit to the amount of mass that can be supported by electron degeneracy pressure is  $\sim 1.4 M_{\odot}$ , called the Chandrasekhar limit. White dwarfs cannot be more massive than this. If the white dwarf is in a binary and accretes enough mass from its companion to pass the Chandrasekhar limit, it is thought that runaway fusion would cause a thermonuclear explosion and result in a Type Ia SNe [where SNe is the plural abbreviation of supernova; see [Alsabti and Murdin](#), for a comprehensive review]. The merging of two white dwarfs is also expected to cause a Type Ia SNe. Both channels are thought to contribute, but their progenitors are elusive in observations. The explosion completely destroys the white dwarf, leaving no remnant. The extremely high energies released in all SNe trigger fusion in the ejected matter, and Type Ia SNe are the primary producers of iron in the Universe. Since all Type Ia SNe are produced by stars of  $\sim 1.4 M_{\odot}$  they all release approximately the same amount of energy. This makes them excellent distance indicators since their apparent luminosities can be compared with the expected absolute luminosity. Since they are also extremely bright (often briefly outshining their entire host galaxies), they are very important to cosmology where they are used to measure the expansion history of the Universe.

**A.2.2.5 High mass stars - Red and blue supergiants, white dwarfs, supernovae, neutron stars, and black holes.** Stars with ( $M_{\star} \gtrsim 8 M_{\odot}$ ) experience significant mass loss through stellar winds, especially post-MS. This tends to increase with increasing mass and decreasing metallicity and can be important even during the MS, though the uncertainties involved make it difficult to predict how much mass a given star will lose and when. They have sufficiently massive cores to ignite carbon, but beyond this their fates depend strongly on when and how much mass is lost. The lowest mass stars in this range are not massive enough to burn elements heavier than carbon. If their core mass at this stage is  $\lesssim 1.4 M_{\odot}$  they

become (mostly oxygen and neon) white dwarfs. If they are massive enough to ignite elements heavier than carbon then they are too massive to be supported by electron degeneracy pressure. They continue through cycles of contraction and fusion, burning neon, oxygen, and silicon until their cores are composed mostly of iron. Iron burning is endothermic and so stars cannot support themselves against gravity by burning it.

A large body of ongoing research is devoted to the very rapid and complicated processes that happen when fusion stops for good [see [Alsabti and Murdin](#), for a comprehensive review], but in short: Most of these stars likely contract until they surpass the density of an atomic nucleus and neutron degeneracy pressure causes the inner core to bounce back and collide with the outer core which is still in free fall. This causes an explosion called a core collapse supernova. There are many subtypes of core collapse SNe, mostly classified according to their composition (determined from spectral measurements) and the shapes of their light curves (§ [A.2](#)). If the remaining core has a mass low enough to be supported by neutron degeneracy pressure, it will stabilize as a neutron star. Otherwise, it will collapse into a black hole. The most massive stars at ZAMS may still be too massive (even after severe mass loss) for neutron degeneracy pressure to be strong enough to cause a core bounce and explosion. If so, they collapse directly into black holes without exploding. Metallicity can strongly impact the fates of massive stars, predominantly by altering mass loss rates. Very low metallicity stars lose significantly less mass, and so are much more likely to collapse directly into black holes.

### **A.2.3 Transients**

Transients, by definition, appear in observations for a (usually brief) period of time and then disappear forever. These are often catastrophically explosive events which are intrinsically very bright and can be seen at large distances, too far away for the progenitor star to have been detected. We can observe them in a much larger volume than variable stars but their inherent rates are much lower, the combined result being that LSST will observe about 35 times more variable stars than transients each night. The vast majority of transients are supernovae (§ [A.2.2.4](#) and § [A.2.2.5](#)). But there is a fairly unexplored gap in

the parameter space to which LSST will throw the door wide open.

**A.2.3.1 Gap transients.** “Gap transient” refers to an event with a peak absolute luminosity that falls in the gap between the brightest novae and the dimmest supernovae. We rarely observe events in this gap, which spans more than 2 orders of magnitude, but theory predicts that more should occur here. Some examples include the explosion of massive stars ( $> 25 M_{\odot}$ , § A.2.2.5) whose cores collapse into black holes (called fallback supernovae), kilonovae (§ A.2.3.2), common envelope mergers, and two types of accreting binaries [Kasliwal, 2011, LSST Science Collaboration et al., 2009]. LSST is expected to make the previously rare events in this space rather common. It is also expected to detect phenomena here that we have never seen before, and possibly some that we haven’t yet imagined [LSST Science Collaboration et al., 2009].

**A.2.3.2 Kilonovae.** Kilonovae are the result of a merger between a neutron star (§ A.2.2.5) and either a black hole or another neutron star. They evolve rapidly and are likely to be visible for only a few days. They produce rare, heavy elements like gold and platinum, and they generate gravitational waves. They are the subject of a particularly active subfield in astronomy’s emerging “multi-messenger” era. This has been particularly true since 2017, when GW170817 [Abbott et al., 2017] was discovered and became the first, and still only, event ever seen by both gravitational wave and electromagnetic detectors [see Metzger, 2019, for a review of the observations]. Scolnic et al. [2018] ran simulations of different surveys (e.g., observing cadence and sky coverage) and found that LSST is likely to have the capability of detecting  $\sim 75$  in total, almost a factor of 5 more than the next highest, the Roman Space telescope (formerly called WFIRST<sup>2</sup>). The fact that the study relied on kilonova models and estimated volumetric rates based on the single known event, plus searches that resulted in non-detections, only underscores the challenge that awaits in identifying these rare events in LSST’s stream, as this is the same paucity of data available to machine learning classifiers. Common “contaminants” include CVs [§ A.2.4.4; Andreoni et al., 2020].

---

<sup>2</sup><https://roman.gsfc.nasa.gov/>

## A.2.4 Variable Stars

Variable stars appear in observations consistently, or at least repeatedly. This category contains stars spread across most of the ranges of mass and life stage. Their variability can be caused by many things. Here I list some common causes and describe an example of each.

**A.2.4.1 Pulsations - Cepheids.** Cepheids are stars between  $\sim 4 - 12 M_{\odot}$  that are burning helium (§ A.2.2) and have surface temperatures near a boundary where opacity changes, causing cycles of expansion and contraction with a well-defined period-luminosity relationship [Pols, 2011]. This relationship allows estimation of the absolute (rather than apparent) luminosity which can be used to determine the (luminosity) distance. They are in a giant or supergiant phase and are therefore bright and can be seen from larger distances. Thus, Cepheids are used to map the expansion history of the local Universe and anchor the Type Ia SNe measurements that extend the “distance ladder” [LSST Science Collaboration et al., 2009].

**A.2.4.2 Eclipsing - Binary systems viewed nearly edge-on.** Stars in binaries (§ A.2.1) will be seen to eclipse each other when our line of sight is nearly edge-on to their orbital plane. This will enable direct measure of their component masses and radii, and by extension, accurate luminosities and distances. These provide stringent test cases when compared to models.

**A.2.4.3 Magnetic activity - Flare stars.** Flares are an explosion in the stellar atmosphere and are triggered by magnetic field lines when they rearrange suddenly. Many stars flare, including our sun, but the phenomenon is more common in very low mass stars (§ A.2.2.1), likely because they are dominated by convection which stresses the magnetic fields [Allred et al., 2015, Dzombeta and Percy, 2019]. These stars are numerous and very faint, thus LSST will detect many more than previous surveys [see LSST Science Collaboration et al., 2009, for an overview, including details about the statements that follow]. Their large numbers may be used to infer the star formation history of the Milky Way disk. A

flare can increase the stellar luminosity by 100 times or more [Allred et al., 2015], and since they are also the most numerous stars likely to host planets, understanding the irradiation impact is important to astrobiology. However, they can masquerade as transients and will contribute to the foreground fog<sup>3</sup> in searches for electromagnetic counterparts of gravitational waves. This can happen (in part) because the star may be too faint to be detected in quiescence; the brightness rises and falls within a few days, so that a given flare is likely to be captured by LSST only once or maybe twice, and other times missed altogether.

**A.2.4.4 Accretion from a companion - Cataclysmic variables.** Cataclysmic variables (CVs) are binary systems (§ A.2.1) composed of a white dwarf (§ A.2.2) accreting matter from a companion which is usually nearing or in its giant phase [for a review, see Warner, 1995, Zorotovic and Schreiber, 2019, Chomiuk et al., 2021]. They are among the most diverse and intriguing classes of variable stars. They are also very common, a result of the twin facts that most stars are in binaries and the most common endpoint of a single star is a white dwarf. However we don't see them as often as we think we should, and this is particularly true at the faint end, where LSST will excel.

They exhibit a very wide range of behavior due to accretion and a range of magnetic field strengths. If the recipient has an appreciable magnetic field, the accreting material will follow the field lines and impact the poles. The process can produce cyclotron emission and make them look red. If magnetic fields are weak, the material leaving the donor star will stream into a disk with the inner edge accreting onto the recipient. Accretion disks with high and low mass transfer rates are relatively stable, but accretion rates change over time. High rates tend to cause high luminosities and blue colors, and they are prone to frequent but non-periodic flickering. Intermediate mass transfer rates onto a disk are thermally unstable, due to opacity peaking near  $10^4$  K [Warner, 1995], and result in outbursting on the accretion disk. These are called dwarf novae and they repeat on timescales of tens of days. The brightest events produced by CVs are novae, which occur when the layer of hydrogen accreted onto the surface of the white dwarf reaches a critical mass, undergoes thermonuclear runaway,

---

<sup>3</sup>“Foreground fog” is composed of nearby (and usually very common) objects that can look similar to more distant and rare extragalactic events, and thus contribute contamination and confusion in searches for the latter.

and erupts. This is likely to reoccur on the order of thousands of years, but a few have been observed to do so within a few hundred years. It is currently unclear whether the eruption tends to expel all of the accreted matter. If it does not, the white dwarf can grow over time and eventually cross the Chandrasekhar limit. Thus CVs may be elusive progenitors of Type Ia supernovae (§ [A.2.2.4](#)).

## APPENDIX B

### ALERT BROKERS: COMMON SERVICES

Astronomical alert brokers are quickly becoming more common, driven in large part by the outsize need anticipated for LSST, but they have existed in a similar form for more than a decade [see [Narayan et al., 2018](#), for a review]. By 2009, the prototype Skyalert [[Williams et al., 2009](#)] was being built with the following functions and features, which still describe brokers today<sup>1</sup>:

#### **Pipeline features:**

*Generality* - ability to ingest and redistribute multiple alert streams from different surveys through a single framework.

*Aggregation* - identification and retrieval of additional data from various catalogs, surveys, and/or facilities that is related to an individual alert. This is often obtained by performing a spatial cross match with a catalog at the position of the alert's source. This can bring in information that is extremely valuable for interpreting, and taking action on, the alert such as observations in different wavelength regimes, spectra, and prior characterizations (e.g., of variability) and/or classifications. Possible host galaxies are also identified in this way and can bring in useful environmental information (such as star formation history and the source's projected location in the galaxy) and perhaps the most valuable bit of information, a redshift. Redshifts are used for accurate distance estimates, which can be used to break degeneracies and clearly discriminate the foreground fog (e.g., flare stars) from more distant

---

<sup>1</sup>I have rearranged and restated the original descriptions to more closely reflect the present-day models and terminology.

transients.

*Computation of new features* (e.g., periodicity) and/or classifications for an alert. Today this usually includes machine learning classifiers and is often done hierarchically, e.g., first separate transients from variable stars, then send the transients through a supernova classifier.

*Annotation, publication, and storage* - the original alert is combined with the aggregated and computed data, pushed to subscribers, and stored in a database & file storage for later access.

### **Data access services:**

*Website* - most commonly includes a “portfolio” page for each object/event that displays the light curves, images, other alert data and all related information the broker has collected and/or calculated. Today, broker websites are primary access channels for alert data and usually offer services like watchlists, pages presenting groups of similarly classified objects, and cone searches & other queries on the broker’s database. Additionally, users may be able to contribute their own comments and data.

*Subscriptions* - users register at a website and receive access to real time alerts, usually with some option to create custom filters. Alerts may be delivered in a variety of ways. I find it useful to distinguish two basic categories: notifications and streams (though in practice they don’t always separate this cleanly). *Notifications* are intended to alert humans and may be sent via tools like email, text, and (more recently) Slack. Recipients may act on notifications by, e.g., visiting the broker’s website to view the data and/or querying and downloading it through an API. *Streams* are intended to trigger automated processing and tend to have much higher throughput (alert rate). ZTF, LSST, and most of their community alert brokers use Apache Kafka to deliver message streams, while Pitt-Google uses Google Pub/Sub. Historically, it has been more common for researchers to receive and act on notifications, but this is becoming increasingly challenging as alert rates grow. LSST will necessitate another great leap forward in the automation of analysis and decision making, and community interest in processing streams appears to be growing. Pitt-Google’s use of Pub/Sub turns out to have important consequences here, discussed throughout § 3.

*API access* to catalogs, streams, and file storage is very commonly provided today (and

appears to have been offered in some form by Skyalert, though the specific term was not used).

**Follow-up services:**

These services can take a variety of forms, but the motivation is to connect scientists with the facilities/telescopes that perform follow-up observations upon request. One commonly used option is to integrate with tools or services which have been purpose-built for this task by others in the community. For example, brokers often develop a plugin for the TOM Toolkit [[Street et al., 2018](#)], which is a software framework for researchers to aggregate and analyze data and connect with follow-up facilities, all using a web front-end or GUI. Broker services may include the ability to automatically trigger a follow-up request on events that pass a filter.

## APPENDIX C

### BAYESIAN NETWORKS

LSST’s unprecedented reach has exciting implications for new discoveries of stellar variability at faint limiting magnitudes and exotic events in the luminosity gap between novae and supernovae. Many of these were discussed in § A.2. A lack of sufficient training data for machine learning (ML) classifiers, due to lack of previous observations in the space, will present challenges for early science. This will be compounded by sparse light curve data scattered across filters. Environmental information (e.g., host galaxy redshift, star formation rate, etc.) will be crucial but also subject to considerable uncertainty. Cross matches will be complicated by crowded fields and the possibility that the true match was not previously cataloged. Expert domain knowledge will be particularly important in bridging these gaps, but the enormous data rates will require it to be applied algorithmically like never before.

I have proposed to apply Bayesian networks to these problems, as they provide a powerful and natural framework for programmatically applying reasoning, learning from the available data, and making interpretable predictions. Bayesian networks are probabilistic graphical models of the relationships between observable and/or latent variables. I developed prototype Bayesian networks for both large-scale classification of the alert stream and targeted searches for exotic transients. In each case, the observables include both light curve information and contextual information, for example, offsets from likely host galaxies or local star formation rates. The object’s true class is a latent variable, and some of my proposed networks included the true redshift as a latent variable as well. The true class is then predicted by applying Bayesian statistics, and various solution methods exist.

As a broker, we can have the opportunity to consider classification of the alert stream at-large, and we also have access to a wealth of information about individual objects in the stream. In pursuit of these classifications, I developed the framework for a Bayesian network meta-classifier. The idea was to combine and augment the results of several more specialized classifiers that were (and are) being developed by various groups in the community who are searching for supernovae, different types of variable stars, etc. These classifiers typically consider the object’s light curves, or statistics derived from it, and possibly the redshift of a likely host galaxy. They do not typically consider other contextual information that will be available to the broker, such as the object’s offset from the likely host galaxy, the local star formation rate, or even the total number and types of nearby galaxies that are potential hosts (recall the crowding will becoming more prevalent in LSST data). This type of information can be extremely valuable in differentiating, for example, a common CV from a rare kilonova (see § A.2), and is often used by humans determining whether to pursue specific objects with valuable time on follow-up telescopes. The Bayesian meta-classifier would provide a programatic way to consider all of the information available and apply human-like reasoning at LSST scale, through the use of prior probabilities and Bayesian statistics.

I also proposed Bayesian networks that would perform targeted searches for specific types of objects. For these, I started with much simpler network designs. For example, I proposed a search for the peculiar calcium-rich gap transients using a “naive Bayes” network structure. This network’s observables included the host galaxy offset [since those of known calcium-rich transients have notably large offsets, [De et al., 2020](#)], a rise-time light curve statistic, and a redshift estimate. This is similar to [Djorgovski et al. \[2014\]](#) who found Bayesian networks to be an excellent method for incorporating spatial contextual data (e.g., host offset) and showed encouraging results using a small number of input variables.

## BIBLIOGRAPHY

B. P. Abbott, R. Abbott, T. D. Abbott, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, V. B. Adya, C. Affeldt, M. Afrough, B. Agarwal, M. Agathos, K. Agatsuma, N. Aggarwal, O. D. Aguiar, L. Aiello, A. Ain, P. Ajith, B. Allen, G. Allen, A. Allocca, P. A. Altin, A. Amato, A. Ananyeva, S. B. Anderson, W. G. Anderson, S. V. Angelova, S. Antier, S. Appert, K. Arai, M. C. Araya, J. S. Areeda, N. Arnaud, K. G. Arun, S. Ascenzi, G. Ashton, M. Ast, S. M. Aston, P. Astone, D. V. Atallah, P. Aufmuth, C. Aulbert, K. AultONeal, C. Austin, A. Avila-Alvarez, S. Babak, P. Bacon, M. K. M. Bader, S. Bae, M. Bailes, P. T. Baker, F. Baldaccini, G. Ballardini, S. W. Ballmer, S. Banagiri, J. C. Barayoga, S. E. Barclay, B. C. Barish, D. Barker, K. Barkett, F. Barone, B. Barr, L. Barsotti, M. Barsuglia, D. Barta, S. D. Barthelmy, J. Bartlett, I. Bartos, R. Bassiri, A. Basti, J. C. Batch, M. Bawaj, J. C. Bayley, M. Bazzan, B. Bécsy, C. Beer, M. Bejger, I. Belahcene, A. S. Bell, B. K. Berger, G. Bergmann, S. Bernuzzi, J. J. Bero, C. P. L. Berry, D. Bersanetti, A. Bertolini, J. Betzwieser, S. Bhagwat, R. Bhandare, I. A. Bilenko, G. Billingsley, C. R. Billman, J. Birch, R. Birney, O. Birnholtz, S. Biscans, S. Biscoveanu, A. Bisht, M. Bitossi, C. Biwer, M. A. Bizouard, J. K. Blackburn, J. Blackman, C. D. Blair, D. G. Blair, R. M. Blair, S. Bloemen, O. Bock, N. Bode, M. Boer, G. Bogaert, A. Bohe, F. Bondu, E. Bonilla, R. Bonnand, B. A. Boom, R. Bork, V. Boschi, S. Bose, K. Bossie, Y. Bouffanais, A. Bozzi, C. Bradaschia, P. R. Brady, M. Branchesi, J. E. Brau, T. Briant, A. Brillet, M. Brinkmann, V. Brisson, P. Brockill, J. E. Broida, A. F. Brooks, D. A. Brown, D. D. Brown, S. Brunett, C. C. Buchanan, A. Buikema, T. Bulik, H. J. Bulten, A. Buonanno, D. Buskulic, C. Buy, R. L. Byer, M. Cabero, L. Cadonati, G. Cagnoli, C. Cahillane, J. Calderón Bustillo, T. A. Callister, E. Calloni, J. B. Camp, M. Canepa, P. Canizares, K. C. Cannon, H. Cao, J. Cao, C. D. Capano, E. Capocasa, F. Carbognani, S. Caride, M. F. Carney, G. Carullo, J. Casanueva Diaz, C. Casentini, S. Caudill, M. Cavaglià, F. Cavalier, R. Cavalieri, G. Cella, C. B. Cepeda, P. Cerdá-Durán, G. Cerretani, E. Cesarini, S. J. Chamberlin, M. Chan, S. Chao, P. Charlton, E. Chase, E. Chassande-Mottin, D. Chatterjee, K. Chatziioannou, B. D. Cheeseboro, H. Y. Chen, X. Chen, Y. Chen, H. P. Cheng, H. Chia, A. Chincarini, A. Chiummo, T. Chmiel, H. S. Cho, M. Cho, J. H. Chow, N. Christensen, Q. Chu, A. J. K. Chua, S. Chua, A. K. W. Chung, S. Chung, G. Ciani, R. Ciolfi, C. E. Cirelli, A. Cirone, F. Clara, J. A. Clark, P. Clearwater, F. Cleva, C. Cocchieri, E. Coccia, P. F. Cohadon, D. Cohen, A. Colla, C. G. Collette, L. R. Cominsky, M. Constancio, L. Conti, S. J. Cooper, P. Corban, T. R. Corbitt, I. Cordero-Carrión, K. R. Corley, N. Cornish, A. Corsi, S. Cortese, C. A. Costa, M. W. Coughlin, S. B. Coughlin, J. P. Coulon, S. T. Countryman, P. Couvares, P. B. Co-

vas, E. E. Cowan, D. M. Coward, M. J. Cowart, D. C. Coyne, R. Coyne, J. D. E. Creighton, T. D. Creighton, J. Cripe, S. G. Crowder, T. J. Cullen, A. Cumming, L. Cunningham, E. Cuoco, T. Dal Canton, G. D'Álya, S. L. Danilishin, S. D'Antonio, K. Danzmann, A. Dasgupta, C. F. Da Silva Costa, V. Dattilo, I. Dave, M. Davier, D. Davis, E. J. Daw, B. Day, S. De, D. DeBra, J. Degallaix, M. De Laurentis, S. Deléglise, W. Del Pozzo, N. Demos, T. Denker, T. Dent, R. De Pietri, V. Dergachev, R. De Rosa, R. T. DeRosa, C. De Rossi, R. DeSalvo, O. de Varona, J. Devenson, S. Dhurandhar, M. C. Díaz, T. Dietrich, L. Di Fiore, M. Di Giovanni, T. Di Girolamo, A. Di Lieto, S. Di Pace, I. Di Palma, F. Di Renzo, Z. Doctor, V. Dolique, F. Donovan, K. L. Dooley, S. Doravari, I. Dorrington, R. Douglas, M. Dovale Álvarez, T. P. Downes, M. Drago, C. Dreissigacker, J. C. Driggers, Z. Du, M. Ducrot, R. Dudi, P. Dupej, S. E. Dwyer, T. B. Edo, M. C. Edwards, A. Effler, H. B. Eggenstein, P. Ehrens, J. Eichholz, S. S. Eikenberry, R. A. Eisenstein, R. C. Essick, D. Estevez, Z. B. Etienne, T. Etzel, M. Evans, T. M. Evans, M. Factourovich, V. Fafone, H. Fair, S. Fairhurst, X. Fan, S. Farinon, B. Farr, W. M. Farr, E. J. Fauchon-Jones, M. Favata, M. Fays, C. Fee, H. Fehrmann, J. Feicht, M. M. Fejer, A. Fernandez-Galiana, I. Ferrante, E. C. Ferreira, F. Ferrini, F. Fidocar, D. Finstad, I. Fiori, D. Fiorucci, M. Fishbach, R. P. Fisher, M. Fitz-Axen, R. Flaminio, M. Fletcher, H. Fong, J. A. Font, P. W. F. Forsyth, S. S. Forsyth, J. D. Fournier, S. Frasca, F. Frasconi, Z. Frei, A. Freise, R. Frey, V. Frey, E. M. Fries, P. Fritschel, V. V. Frolov, P. Fulda, M. Fyffe, H. Gabbard, B. U. Gadre, S. M. Gaebel, J. R. Gair, L. Gammaitoni, M. R. Ganija, S. G. Gaonkar, C. Garcia-Quiros, F. Garufi, B. Gateley, S. Gaudio, G. Gaur, V. Gayathri, N. Gehrels, G. Gemme, E. Genin, A. Genai, D. George, J. George, L. Gergely, V. Germain, S. Ghonge, Abhirup Ghosh, Archisman Ghosh, S. Ghosh, J. A. Giaime, K. D. Giardino, A. Giazotto, K. Gill, L. Glover, E. Goetz, R. Goetz, S. Gomes, B. Goncharov, G. González, J. M. Gonzalez Castro, A. Gopakumar, M. L. Gorodetsky, S. E. Gossan, M. Gosselin, R. Gouaty, A. Grado, C. Graef, M. Granata, A. Grant, S. Gras, C. Gray, G. Greco, A. C. Green, E. M. Gretarsson, P. Groot, H. Grote, S. Grunewald, P. Gruning, G. M. Guidi, X. Guo, A. Gupta, M. K. Gupta, K. E. Gushwa, E. K. Gustafson, R. Gustafson, O. Halim, B. R. Hall, E. D. Hall, E. Z. Hamilton, G. Hammond, M. Haney, M. M. Hanke, J. Hanks, C. Hanna, M. D. Hannam, O. A. Hannuksela, J. Hanson, T. Hardwick, J. Harms, G. M. Harry, I. W. Harry, M. J. Hart, C. J. Haster, K. Haughian, J. Healy, A. Heidmann, M. C. Heintze, H. Heitmann, P. Hello, G. Hemming, M. Hendry, I. S. Heng, J. Hennig, A. W. Heptonstall, M. Heurs, S. Hild, T. Hinderer, W. C. G. Ho, D. Hoak, D. Hofman, K. Holt, D. E. Holz, P. Hopkins, C. Horst, J. Hough, E. A. Houston, E. J. Howell, A. Hreibi, Y. M. Hu, E. A. Huerta, D. Huet, B. Hughey, S. Husa, S. H. Huttner, T. Huynh-Dinh, N. Indik, R. Inta, G. Intini, H. N. Isa, J. M. Isac, M. Isi, B. R. Iyer, K. Izumi, T. Jacqmin, K. Jani, P. Jaranowski, S. Jawahar, F. Jiménez-Forteza, W. W. Johnson, N. K. Johnson-McDaniel, D. I. Jones, R. Jones, R. J. G. Jonker, L. Ju, J. Junker, C. V. Kalaghatgi, V. Kalogera, B. Kamai, S. Kandhasamy, G. Kang, J. B. Kanner, S. J. Kapadia, S. Karki, K. S. Karvinen, M. Kasprzack, W. Kastaun, M. Katolik, E. Katsavounidis, W. Katzman, S. Kaufer, K. Kawabe, F. Kéfélian, D. Keitel, A. J. Kemball, R. Kennedy, C. Kent, J. S. Key, F. Y. Khalili, I. Khan, S. Khan, Z. Khan, E. A. Khazanov, N. Kijbunchoo, Chunglee Kim, J. C. Kim, K. Kim, W. Kim, W. S. Kim, Y. M. Kim, S. J. Kimbrell, E. J. King, P. J. King, M. Kinley-Hanlon, R. Kirchhoff, J. S. Kissel, L. Kleybolte, S. Klimenko, T. D. Knowles, P. Koch, S. M. Koehlenbeck, S. Koley, V. Kondrashov, A. Kontos, M. Korobko, W. Z. Korth, I. Kowalska, D. B. Kozak,

C. Krämer, V. Kringel, B. Krishnan, A. Królak, G. Kuehn, P. Kumar, R. Kumar, S. Kumar, L. Kuo, A. Kutynia, S. Kwang, B. D. Lackey, K. H. Lai, M. Landry, R. N. Lang, J. Lange, B. Lantz, R. K. Lanza, S. L. Larson, A. Lartaux-Vollard, P. D. Lasky, M. Laxen, A. Lazzarini, C. Lazzaro, P. Leaci, S. Leavey, C. H. Lee, H. K. Lee, H. M. Lee, H. W. Lee, K. Lee, J. Lehmann, A. Lenon, E. Leon, M. Leonardi, N. Leroy, N. Letendre, Y. Levin, T. G. F. Li, S. D. Linker, T. B. Littenberg, J. Liu, X. Liu, R. K. L. Lo, N. A. Lockerbie, L. T. London, J. E. Lord, M. Lorenzini, V. Lorette, M. Lormand, G. Losurdo, J. D. Lough, C. O. Lousto, G. Lovelace, H. Lück, D. Lumaca, A. P. Lundgren, R. Lynch, Y. Ma, R. Macas, S. Macfoy, B. Machenschalk, M. MacInnis, D. M. Macleod, I. Magaña Hernandez, F. Magaña-Sandoval, L. Magaña Zertuche, R. M. Magee, E. Majorana, I. Maksimovic, N. Man, V. Mandic, V. Mangano, G. L. Mansell, M. Manske, M. Mantovani, F. Marchesoni, F. Marion, S. Márka, Z. Márka, C. Markakis, A. S. Markosyan, A. Markowitz, E. Maros, A. Marquina, P. Marsh, F. Martelli, L. Martellini, I. W. Martin, R. M. Martin, D. V. Martynov, J. N. Marx, K. Mason, E. Massera, A. Masserot, T. J. Massinger, M. Masso-Reid, S. Mastrogiovanni, A. Matas, F. Matichard, L. Matone, N. Mavalvala, N. Mazumder, R. McCarthy, D. E. McClelland, S. McCormick, L. McCuller, S. C. McGuire, G. McIntyre, J. McIver, D. J. McManus, L. McNeill, T. McRae, S. T. McWilliams, D. Meacher, G. D. Meadors, M. Mehmet, J. Meidam, E. Mejuto-Villa, A. Melatos, G. Mendell, R. A. Mercer, E. L. Merilh, M. Merzougui, S. Meshkov, C. Messenger, C. Messick, R. Metzдорff, P. M. Meyers, H. Miao, C. Michel, H. Middleton, E. E. Mikhailov, L. Milano, A. L. Miller, B. B. Miller, J. Miller, M. Millhouse, M. C. Milovich-Goff, O. Minazzoli, Y. Minenkov, J. Ming, C. Mishra, S. Mitra, V. P. Mitrofanov, G. Mitselmakher, R. Mittleman, D. Moffa, A. Moggi, K. Mogushi, M. Mohan, S. R. P. Mohapatra, I. Molina, M. Montani, C. J. Moore, D. Moraru, G. Moreno, S. Morisaki, S. R. Morriss, B. Mours, C. M. Mow-Lowry, G. Mueller, A. W. Muir, Arunava Mukherjee, D. Mukherjee, S. Mukherjee, N. Mukund, A. Mullavey, J. Munch, E. A. Muñiz, M. Muratore, P. G. Murray, A. Nagar, K. Napier, I. Nardecchia, L. Naticchioni, R. K. Nayak, J. Neilson, G. Nelemans, T. J. N. Nelson, M. Nery, A. Neunzert, L. Nevin, J. M. Newport, G. Newton, K. K. Y. Ng, P. Nguyen, T. T. Nguyen, D. Nichols, A. B. Nielsen, S. Nissanke, A. Nitz, A. Noack, F. Nocera, D. Nolting, C. North, L. K. Nuttall, J. Oberling, G. D. O’Dea, G. H. Ogin, J. J. Oh, S. H. Oh, F. Ohme, M. A. Okada, M. Oliver, P. Oppermann, Richard J. Oram, B. O’Reilly, R. Ormiston, L. F. Ortega, R. O’Shaughnessy, S. Ossokine, D. J. Ottaway, H. Overmier, B. J. Owen, A. E. Pace, J. Page, M. A. Page, A. Pai, S. A. Pai, J. R. Palamos, O. Palashov, C. Palomba, A. Pal-Singh, Howard Pan, Huang-Wei Pan, B. Pang, P. T. H. Pang, C. Pankow, F. Panarale, B. C. Pant, F. Paoletti, A. Paoli, M. A. Papa, A. Parida, W. Parker, D. Pascucci, A. Pasqualetti, R. Passaquieti, D. Passuello, M. Patil, B. Patricelli, B. L. Pearlstone, M. Pedraza, R. Pedurand, L. Pekowsky, A. Pele, S. Penn, C. J. Perez, A. Perreca, L. M. Perri, H. P. Pfeiffer, M. Phelps, O. J. Piccinni, M. Pichot, F. Piergiovanni, V. Pierro, G. Pillant, L. Pinard, I. M. Pinto, M. Pirello, M. Pitkin, M. Poe, R. Poggiani, P. Popolizio, E. K. Porter, A. Post, J. Powell, J. Prasad, J. W. W. Pratt, G. Pratten, V. Predoi, T. Prestegard, M. Prijatelj, M. Principe, S. Privitera, R. Prix, G. A. Prodi, L. G. Prokhorov, O. Puncken, M. Punturo, P. Puppato, M. Pürrer, H. Qi, V. Quetschke, E. A. Quintero, R. Quitzow-James, F. J. Raab, D. S. Rabeling, H. Radkins, P. Raffai, S. Raja, C. Rajan, B. Rajbhandari, M. Rakhmanov, K. E. Ramirez, A. Ramos-Buades, P. Rapagnani, V. Raymond, M. Razzano, J. Read, T. Regimbau, L. Rei, S. Reid, D. H. Reitze, W. Ren, S. D. Reyes, F. Ricci,

P. M. Ricker, S. Rieger, K. Riles, M. Rizzo, N. A. Robertson, R. Robie, F. Robinet, A. Rocchi, L. Rolland, J. G. Rollins, V. J. Roma, J. D. Romano, R. Romano, C. L. Romel, J. H. Romie, D. Rosińska, M. P. Ross, S. Rowan, A. Rüdiger, P. Ruggi, G. Rutins, K. Ryan, S. Sachdev, T. Sadecki, L. Sadeghian, M. Sakellariadou, L. Salconi, M. Saleem, F. Salemi, A. Samajdar, L. Sammut, L. M. Sampson, E. J. Sanchez, L. E. Sanchez, N. Sanchis-Gual, V. Sandberg, J. R. Sanders, B. Sassolas, B. S. Sathyaprakash, P. R. Saulson, O. Sauter, R. L. Savage, A. Sawadsky, P. Schale, M. Scheel, J. Scheuer, J. Schmidt, P. Schmidt, R. Schnabel, R. M. S. Schofield, A. Schönbeck, E. Schreiber, D. Schuette, B. W. Schulte, B. F. Schutz, S. G. Schwalbe, J. Scott, S. M. Scott, E. Seidel, D. Sellers, A. S. Sengupta, D. Sentenac, V. Sequino, A. Sergeev, D. A. Shaddock, T. J. Shaffer, A. A. Shah, M. S. Shahriar, M. B. Shaner, L. Shao, B. Shapiro, P. Shawhan, A. Sheperd, D. H. Shoemaker, D. M. Shoemaker, K. Siellez, X. Siemens, M. Sieniawska, D. Sigg, A. D. Silva, L. P. Singer, A. Singh, A. Singhal, A. M. Sintes, B. J. J. Slagmolen, B. Smith, J. R. Smith, R. J. E. Smith, S. Somala, E. J. Son, J. A. Sonnenberg, B. Sorazu, F. Sorrentino, T. Souradeep, A. P. Spencer, A. K. Srivastava, K. Staats, A. Staley, M. Steinke, J. Steinlechner, S. Steinlechner, D. Steinmeyer, S. P. Stevenson, R. Stone, D. J. Stops, K. A. Strain, G. Stratta, S. E. Strigin, A. Strunk, R. Sturani, A. L. Stuver, T. Z. Summerscales, L. Sun, S. Sunil, J. Suresh, P. J. Sutton, B. L. Swinkels, M. J. Szczepańczyk, M. Tacca, S. C. Tait, C. Talbot, D. Talukder, D. B. Tanner, M. Tápai, A. Taracchini, J. D. Tasson, J. A. Taylor, R. Taylor, S. V. Tewari, T. Theeg, F. Thies, E. G. Thomas, M. Thomas, P. Thomas, K. A. Thorne, K. S. Thorne, E. Thrane, S. Tiwari, V. Tiwari, K. V. Tokmakov, K. Toland, M. Tonelli, Z. Tornasi, A. Torres-Forné, C. I. Torrie, D. Töyrä, F. Travasso, G. Traylor, J. Trinastic, M. C. Tringali, L. Trozzo, K. W. Tsang, M. Tse, R. Tso, L. Tsukada, D. Tsuna, D. Tuyenbayev, K. Ueno, D. Ugolini, C. S. Unnikrishnan, A. L. Urban, S. A. Usman, H. Vahlbruch, G. Vajente, G. Valdes, M. Vallisneri, N. van Bakel, M. van Beuzekom, J. F. J. van den Brand, C. Van Den Broeck, D. C. Vander-Hyde, L. van der Schaaf, J. V. van Heijningen, A. A. van Veggel, M. Vardaro, V. Varma, S. Vass, M. Vasúth, A. Vecchio, G. Vedovato, J. Veitch, P. J. Veitch, K. Venkateswara, G. Venugopalan, D. Verkindt, F. Vetranò, A. Viceré, A. D. Viets, S. Vinciguerra, D. J. Vine, J. Y. Vinet, S. Vitale, T. Vo, H. Vocca, C. Vorvick, S. P. Vyatchanin, A. R. Wade, L. E. Wade, M. Wade, R. Walet, M. Walker, L. Wallace, S. Walsh, G. Wang, H. Wang, J. Z. Wang, W. H. Wang, Y. F. Wang, R. L. Ward, J. Warner, M. Was, J. Watchi, B. Weaver, L. W. Wei, M. Weinert, A. J. Weinstein, R. Weiss, L. Wen, E. K. Wessel, P. Weßels, J. Westerweck, T. Westphal, K. Wette, J. T. Whelan, S. E. Whitcomb, B. F. Whiting, C. Whittle, D. Wilken, D. Williams, R. D. Williams, A. R. Williamson, J. L. Willis, B. Willke, M. H. Wimmer, W. Winkler, C. C. Wipf, H. Wittel, G. Woan, J. Woehler, J. Wofford, K. W. K. Wong, J. Worden, J. L. Wright, D. S. Wu, D. M. Wysocki, S. Xiao, H. Yamamoto, C. C. Yancey, L. Yang, M. J. Yap, M. Yazback, Hang Yu, Haocun Yu, M. Yvert, A. Zadrožny, M. Zanolin, T. Zelenova, J. P. Zendri, M. Zevin, L. Zhang, M. Zhang, T. Zhang, Y. H. Zhang, C. Zhao, M. Zhou, Z. Zhou, S. J. Zhu, X. J. Zhu, A. B. Zimmerman, M. E. Zucker, J. Zweizig, LIGO Scientific Collaboration, and Virgo Collaboration. GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral. *Phys. Rev. Lett.*, 119(16):161101, October 2017. doi: 10.1103/PhysRevLett.119.161101.

Javier Abril, Linda Schmidtbreick, Alessandro Ederoclite, and Carlos López-Sanjuan. Dis-

- entangling cataclysmic variables in Gaia's HR diagram. *MNRAS*, 492(1):L40–L44, February 2020. doi: 10.1093/mnrasl/slz181.
- Joel C. Allred, Adam F. Kowalski, and Mats Carlsson. A Unified Computational Model for Solar and Stellar Flares. *ApJ*, 809(1):104, August 2015. doi: 10.1088/0004-637X/809/1/104.
- Athem W. Alsabti and Paul Murdin, editors. *Handbook of Supernovae*. Springer, Cham. ISBN 978-3-319-20794-0. doi: 10.1007/978-3-319-20794-0. URL <https://link.springer.com/referencework/10.1007/978-3-319-21846-5>.
- C. Amole, M. Ardid, I. J. Arnquist, D. M. Asner, D. Baxter, E. Behnke, P. Bhattacharjee, H. Borsodi, M. Bou-Cabo, P. Champion, G. Cao, C. J. Chen, U. Chowdhury, K. Clark, J. I. Collar, P. S. Cooper, M. Crisler, G. Crowder, C. E. Dahl, M. Das, S. Fallows, J. Farine, I. Felis, R. Filgas, F. Girard, G. Giroux, J. Hall, O. Harris, E. W. Hoppe, M. Jin, C. B. Krauss, M. Laurin, I. Lawson, A. Leblanc, I. Levine, W. H. Lippincott, F. Mamedov, D. Maurya, P. Mitra, T. Nania, R. Neilson, A. J. Noble, S. Olson, A. Ortega, A. Plante, R. Podvianuk, S. Priya, A. E. Robinson, A. Roeder, R. Rucinski, O. Scallon, S. Seth, A. Sonnenschein, N. Starinski, I. Štekl, F. Tardif, E. Vázquez-Jáuregui, J. Wells, U. Wichoski, Y. Yan, V. Zacek, and J. Zhang. Dark Matter Search Results from the PICO-60 C<sub>3</sub>F<sub>8</sub> Bubble Chamber. *Physical Review Letters*, 118(25):1–14, 2017. ISSN 10797114. doi: 10.1103/PhysRevLett.118.251301.
- C. Amole, M. Ardid, I. J. Arnquist, D. M. Asner, D. Baxter, E. Behnke, M. Bressler, B. Broerman, G. Cao, C. J. Chen, U. Chowdhury, K. Clark, J. I. Collar, P. S. Cooper, C. B. Coutu, C. Cowles, M. Crisler, G. Crowder, N. A. Cruz-Venegas, C. E. Dahl, M. Das, S. Fallows, J. Farine, I. Felis, R. Filgas, F. Girard, G. Giroux, J. Hall, C. Hardy, O. Harris, T. Hillier, E. W. Hoppe, C. M. Jackson, M. Jin, L. Klopfenstein, T. Kozynets, C. B. Krauss, M. Laurin, I. Lawson, A. Leblanc, I. Levine, C. Licciardi, W. H. Lippincott, B. Loer, F. Mamedov, P. Mitra, C. Moore, T. Nania, R. Neilson, A. J. Noble, P. Oedeker, A. Ortega, M. C. Piro, A. Plante, R. Podvianuk, S. Priya, A. E. Robinson, S. Sahoo, O. Scallon, S. Seth, A. Sonnenschein, N. Starinski, I. Štekl, T. Sullivan, F. Tardif, E. Vázquez-Jáuregui, N. Walkowski, E. Weima, U. Wichoski, K. Wierman, Y. Yan, V. Zacek, J. Zhang, and PICO Collaboration. Dark matter search results from the complete exposure of the PICO-60 C<sub>3</sub>F<sub>8</sub> bubble chamber. *Phys. Rev. D*, 100(2):022001, July 2019. doi: 10.1103/PhysRevD.100.022001.
- Igor Andreoni, Erik C. Kool, Ana Sagués Carracedo, Mansi M. Kasliwal, Mattia Bulla, Tomás Ahumada, Michael W. Coughlin, Shreya Anand, Jesper Sollerman, Ariel Goobar, David L. Kaplan, Tegan T. Loveridge, Viraj Karambelkar, Jeff Cooke, Ashot Bagdasaryan, Eric C. Bellm, S. Bradley Cenko, David O. Cook, Kishalay De, Richard Dekany, Alexandre Delacroix, Andrew Drake, Dmitry A. Duev, Christoffer Fremling, V. Zach Golkhou, Matthew J. Graham, David Hale, S. R. Kulkarni, Thomas Kupfer, Russ R. Laher, Ashish A. Mahabal, Frank J. Masci, Ben Rusholme, Roger M. Smith, Anastasios Tzanidakis, Angela Van Sistine, and Yuhan Yao. Constraining the Kilonova Rate with

- Zwicky Transient Facility Searches Independent of Gravitational Wave and Short Gamma-Ray Burst Triggers. *ApJ*, 904(2):155, December 2020. doi: 10.3847/1538-4357/abfb4c.
- S. M. Barr, R. Sekhar Chivukula, and E. Farhi. Electroweak fermion number violation and the production of stable particles in the early universe. *Physics Letters B*, 241:387–391, May 1990. doi: 10.1016/0370-2693(90)91661-T.
- Marco Battaglieri, Alberto Belloni, Aaron Chou, Priscilla Cushman, Bertrand Echenard, Rouven Essig, Juan Estrada, Jonathan L. Feng, Brenna Flaugher, and Patrick J. Fox. US Cosmic Visions: New Ideas in Dark Matter 2017: Community Report. *arXiv e-prints*, art. arXiv:1707.04591, Jul 2017.
- Eric Bellm, Robert Blum, Melissa Graham, Leanne Guy, Željko Ivezić, William O Mullane, Maria Patterson, John Swinbank, Beth Willman, and Lsst Project. LDM-612: Large Synoptic Survey Telescope ( LSST ) Plans and Policies for LSST Alert Distribution. 2019.
- Eric C. Bellm. Volumetric survey speed: A figure of merit for transient surveys. *Publications of the Astronomical Society of the Pacific*, 128(966):1–22, 2016. ISSN 00046280. doi: 10.1088/1538-3873/128/966/084501.
- Gianfranco Bertone and Malcolm Fairbairn. Compact stars as dark matter probes. *Phys. Rev. D*, 77(4):043515, February 2008. doi: 10.1103/PhysRevD.77.043515.
- Gianfranco Bertone, Dan Hooper, and Joseph Silk. Particle dark matter: evidence, candidates and constraints. *Phys. Rep.*, 405(5-6):279–390, Jan 2005. doi: 10.1016/j.physrep.2004.08.031.
- G. R. Blumenthal, S. M. Faber, J. R. Primack, and M. J. Rees. Formation of galaxies and large-scale structure with cold dark matter. *Nature*, 311:517–525, Oct 1984. doi: 10.1038/311517a0.
- Giorgio Busoni, Andrea De Simone, Pat Scott, and Aaron C. Vincent. Evaporation and scattering of momentum- and velocity-dependent dark matter in the Sun. *J. Cosmology Astropart. Phys.*, 2017(10):037, October 2017. doi: 10.1088/1475-7516/2017/10/037.
- Jordi Casanellas and Ilídio Lopes. First Asteroseismic Limits on the Nature of Dark Matter. *ApJ*, 765(1):L21, March 2013. doi: 10.1088/2041-8205/765/1/L21.
- Jordi Casanellas, Isa Brandão, and Yveline Lebreton. Stellar convective cores as dark matter probes. *Phys. Rev. D*, 91(10):103535, May 2015. doi: 10.1103/PhysRevD.91.103535.
- R. S. Chivukula and T. P. Walker. Technicolor cosmology. *Nuclear Physics B*, 329:445–463, January 1990. doi: 10.1016/0550-3213(90)90151-3.
- Jieun Choi, Aaron Dotter, Charlie Conroy, Matteo Cantiello, Bill Paxton, and Benjamin D. Johnson. MESA Isochrones and Stellar Tracks (MIST). I: Solar-Scaled Models. *The Astrophysical Journal*, 823(2):1–48, 2016. ISSN 1538-4357. doi: 10.3847/0004-637X/823/

2/102. URL <http://arxiv.org/abs/1604.08592><http://dx.doi.org/10.3847/0004-637X/823/2/102>.

Laura Chomiuk, Brian D. Metzger, and Ken J. Shen. New Insights into Classical Novae. *ARA&A*, 59, September 2021. doi: 10.1146/annurev-astro-112420-114502.

Giulia C. Cinquegrana and Amanda I. Karakas. The most metal-rich stars in the universe: chemical contributions of low- and intermediate-mass asymptotic giant branch stars with metallicities within  $0.04 \leq Z \leq 0.10$ . *MNRAS*, 510(2):1557–1576, February 2022. doi: 10.1093/mnras/stab3379.

Kishalay De, Mansi M. Kasliwal, Anastasios Tzanidakis, U. Christoffer Fremling, Scott Adams, Robert Aloisi, Igor Andreoni, Ashot Bagdasaryan, Eric C. Bellm, Lars Bildsten, Christopher Cannella, David O. Cook, Alexandre Delacroix, Andrew Drake, Dmitry Duev, Alison Dugas, Sara Frederick, Avishay Gal-Yam, Daniel Goldstein, V. Zach Golkhou, Matthew J. Graham, David Hale, Matthew Hankins, George Helou, Anna Y. Q. Ho, Ido Irani, Jacob E. Jencson, David L. Kaplan, Stephen Kaye, S. R. Kulkarni, Thomas Kupfer, Russ R. Laher, Robin Leadbeater, Ragnhild Lunnan, Frank J. Masci, Adam A. Miller, James D. Neill, Eran O. Ofek, Daniel A. Perley, Abigail Polin, Thomas A. Prince, Eliot Quataert, Dan Reiley, Reed L. Riddle, Ben Rusholme, Yashvi Sharma, David L. Shupe, Jesper Sollerman, Leonardo Tartaglia, Richard Walters, Lin Yan, and Yuhan Yao. The Zwicky Transient Facility Census of the Local Universe. I. Systematic Search for Calcium-rich Gap Transients Reveals Three Related Spectroscopic Subclasses. *ApJ*, 905(1):58, December 2020. doi: 10.3847/1538-4357/abb45c.

S. G. Djorgovski, A. A. Mahabal, C. Donalek, M. J. Graham, A. J. Drake, M. Turmon, and T. Fuchs. Automated Real-Time Classification and Decision Making in Massive Data Streams from Synoptic Sky Surveys. *arXiv e-prints*, art. arXiv:1407.3502, July 2014.

Aaron Dotter. Mesa Isochrones and Stellar Tracks (Mist) 0: Methods for the Construction of Stellar Isochrones. *The Astrophysical Journal Supplement Series*, 222(1):8, 2016. ISSN 1538-4365. doi: 10.3847/0067-0049/222/1/8. URL <http://stacks.iop.org/0067-0049/222/i=1/a=8?key=crossref.80eafb1cc332e11191ee19714c861a18>.

K. Dzombeta and J. R. Percy. Flare Stars: A Short Review (Abstract). *Journal of the American Association of Variable Star Observers*, 47(2):282, December 2019.

Malcolm Fairbairn, Pat Scott, and Joakim Edsjö. The zero age main sequence of WIMP burners. *Phys. Rev. D*, 77(4):047301, February 2008. doi: 10.1103/PhysRevD.77.047301.

Collaboration Gaia. Gaia Data Release 2: Observational Hertzsprung-Russell diagrams. *Astronomy and Astrophysics*, 616, 2018. ISSN 14320746. doi: 10.1051/0004-6361/201832843. URL <http://cdsarc.u-strasbg.fr/viz-bin/qcat?J/A+A/616/A10>.

Gaia Collaboration, C. Babusiaux, F. van Leeuwen, M. A. Barstow, C. Jordi, A. Vallenari, D. Bossini, A. Bressan, T. Cantat-Gaudin, M. van Leeuwen, A. G. A. Brown, T. Prusti, J. H. J. de Bruijne, C. A. L. Bailer-Jones, M. Biermann, D. W. Evans, L. Eyer, F. Jansen,

S. A. Klioner, U. Lammers, L. Lindegren, X. Luri, F. Mignard, C. Panem, D. Pourbaix, S. Randich, P. Sartoretti, H. I. Siddiqui, C. Soubiran, N. A. Walton, F. Arenou, U. Bastian, M. Cropper, R. Drimmel, D. Katz, M. G. Lattanzi, J. Bakker, C. Cacciari, J. Castañeda, L. Chaoul, N. Cheek, F. De Angeli, C. Fabricius, R. Guerra, B. Holl, E. Masana, R. Messineo, N. Mowlavi, K. Nienartowicz, P. Panuzzo, J. Portell, M. Riello, G. M. Seabroke, P. Tanga, F. Thévenin, G. Gracia-Abril, G. Comoretto, M. Garcia-Reinaldos, D. Teyssier, M. Altmann, R. Andrae, M. Audard, I. Bellas-Velidis, K. Benson, J. Berthier, R. Blomme, P. Burgess, G. Busso, B. Carry, A. Cellino, G. Clementini, M. Clotet, O. Creevey, M. Davidson, J. De Ridder, L. Delchambre, A. Dell’Oro, C. Ducourant, J. Fernández-Hernández, M. Fouesneau, Y. Frémat, L. Galluccio, M. García-Torres, J. González-Núñez, J. J. González-Vidal, E. Gosset, L. P. Guy, J. L. Halbwachs, N. C. Hambly, D. L. Harrison, J. Hernández, D. Hestroffer, S. T. Hodgkin, A. Hutton, G. Jasiewicz, A. Jean-Antoine-Piccolo, S. Jordan, A. J. Korn, A. Krone-Martins, A. C. Lanzafame, T. Lebzelter, W. Löffler, M. Manteiga, P. M. Marrese, J. M. Martín-Fleitas, A. Moitinho, A. Mora, K. Muinonen, J. Osinde, E. Pancino, T. Pauwels, J. M. Petit, A. Recio-Blanco, P. J. Richards, L. Rimoldini, A. C. Robin, L. M. Sarro, C. Siopis, M. Smith, A. Sozzetti, M. Süveges, J. Torra, W. van Reeve, U. Abbas, A. Abreu Aramburu, S. Accart, C. Aerts, G. Altavilla, M. A. Álvarez, R. Alvarez, J. Alves, R. I. Anderson, A. H. Andrei, E. Anglada Varela, E. Antiche, T. Antoja, B. Arcay, T. L. Astraatmadja, N. Bach, S. G. Baker, L. Balaguer-Núñez, P. Balm, C. Barache, C. Barata, D. Barbato, F. Barblan, P. S. Barklem, D. Barrado, M. Barros, L. Bartholomé Muñoz, J. L. Bassilana, U. Becciani, M. Bellazzini, A. Berihuete, S. Bertone, L. Bianchi, O. Bienaymé, S. Blanco-Cuaresma, T. Boch, C. Boeche, A. Bombrun, R. Borrachero, S. Bouquillon, G. Bourda, A. Bragaglia, L. Bramante, M. A. Breddels, N. Brouillet, T. Brüsemeister, E. Brugaletta, B. Bucciarelli, A. Burlacu, D. Busonero, A. G. Butkevich, R. Buzzzi, E. Caffau, R. Cancelliere, G. Cannizzaro, R. Carballo, T. Carlucci, J. M. Carrasco, L. Casamiquela, M. Castellani, A. Castro-Ginard, P. Charlot, L. Chemin, A. Chiavassa, G. Cocozza, G. Costigan, S. Cowell, F. Crifo, M. Crosta, C. Crowley, J. Cuypers, C. Dafonte, Y. Damerджи, A. Dapergolas, P. David, M. David, P. de Laverny, F. De Luise, R. De March, D. de Martino, R. de Souza, A. de Torres, J. Debosscher, E. del Pozo, M. Delbo, A. Delgado, H. E. Delgado, S. Diakite, C. Diener, E. Distefano, C. Dolding, P. Drazinos, J. Durán, B. Edvardsson, H. Enke, K. Eriksson, P. Esquej, G. Eynard Bontemps, C. Fabre, M. Fabrizio, S. Faigler, A. J. Falcão, M. Farràs Casas, L. Federici, G. Fedorets, P. Fernique, F. Figueras, F. Filippi, K. Findeisen, A. Fonti, E. Fraile, M. Fraser, B. Frézouls, M. Gai, S. Galleti, D. Garabato, F. García-Sedano, A. Garofalo, N. Garralda, A. Gavel, P. Gavras, J. Gerssen, R. Geyer, P. Giacobbe, G. Gilmore, S. Girona, G. Giuffrida, F. Glass, M. Gomes, M. Granvik, A. Gueguen, A. Guerrier, J. Guiraud, R. Gutiérrez, R. Haigron, D. Hatzidimitriou, M. Hauser, M. Haywood, U. Heiter, A. Helmi, J. Heu, T. Hilger, D. Hobbs, W. Hofmann, G. Holland, H. E. Huckle, A. Hypki, V. Icardi, K. Janßen, G. Jevardat de Fombelle, P. G. Jonker, Á. L. Juhász, F. Julbe, A. Karamelas, A. Kewley, J. Klar, A. Kochoska, R. Kohley, K. Kolenberg, M. Kontizas, E. Kontizas, S. E. Kuposov, G. Kordopatis, Z. Kostrzewa-Rutkowska, P. Koubsky, S. Lambert, A. F. Lanza, Y. Lasne, J. B. Lavigne, Y. Le Fustec, C. Le Poncin-Lafitte, Y. Lebreton, S. Leccia, N. Leclerc, I. Lecoeur-Taibi, H. Lenhardt, F. Leroux, S. Liao, E. Licata, H. E. P. Lindstrøm, T. A. Lister, E. Livanou, A. Lobel, M. López, S. Managau, R. G. Mann, G. Mantelet, O. Marchal, J. M. Marchant, M. Mar-

- coni, S. Marinoni, G. Marschalkó, D. J. Marshall, M. Martino, G. Marton, N. Mary, D. Massari, G. Matijević, T. Mazeh, P. J. McMillan, S. Messina, D. Michalik, N. R. Millar, D. Molina, R. Molinaro, L. Molnár, P. Montegriffo, R. Mor, R. Morbidelli, T. Morel, D. Morris, A. F. Mulone, T. Muraveva, I. Musella, G. Nelemans, L. Nicastro, L. Noval, W. O’Mullane, C. Ordénovic, D. Ordóñez-Blanco, P. Osborne, C. Pagani, I. Pagano, F. Pailler, H. Palacin, L. Palaversa, A. Panahi, M. Pawlak, A. M. Piersimoni, F. X. Pineau, E. Plachy, G. Plum, E. Poggio, E. Poujoulet, A. Prša, L. Pulone, E. Racero, S. Ragaini, N. Rambaux, M. Ramos-Lerate, S. Regibo, C. Reylé, F. Riclet, V. Ripepi, A. Riva, A. Rivard, G. Rixon, T. Roegiers, M. Roelens, M. Romero-Gómez, N. Rowell, F. Royer, L. Ruiz-Dern, G. Sadowski, T. Sagristà Sellés, J. Sahlmann, J. Salgado, E. Salguero, N. Sanna, T. Santana-Ros, M. Sarasso, H. Saviotto, M. Schultheis, E. Sciacca, M. Segol, J. C. Segovia, D. Ségransan, I. C. Shih, L. Siltala, A. F. Silva, R. L. Smart, K. W. Smith, E. Solano, F. Solitro, R. Sordo, S. Soria Nieto, J. Souchay, A. Spagna, F. Spoto, U. Stampa, I. A. Steele, H. Steidelmüller, C. A. Stephenson, H. Stoev, F. F. Suess, J. Surdej, L. Szabados, E. Szegedi-Elek, D. Tapiador, F. Taris, G. Tauran, M. B. Taylor, R. Teixeira, D. Terrett, P. Teyssandier, W. Thuillot, A. Titarenko, F. Torra Clotet, C. Turon, A. Ulla, E. Utrilla, S. Uzzi, M. Vaillant, G. Valentini, V. Valette, A. van Elteren, E. Van Hemelryck, M. Vaschetto, A. Vecchiato, J. Veljanoski, Y. Viala, D. Vicente, S. Vogt, C. von Essen, H. Voss, V. Votruba, S. Voutsinas, G. Walmsley, M. Weiler, O. Wertz, T. Wevers, L. Wyrzykowski, A. Yoldas, M. Žerjal, H. Ziaee pour, J. Zorec, S. Zschocke, S. Zucker, C. Zurbach, and T. Zwitter. Gaia Data Release 2. Observational Hertzsprung-Russell diagrams. *A&A*, 616:A10, August 2018. doi: 10.1051/0004-6361/201832843.
- T. K. Gaisser, G. Steigman, and S. Tilav. Limits on cold-dark-matter candidates from deep underground detectors. *Phys. Rev. D*, 34(8):2206–2222, Oct 1986. doi: 10.1103/PhysRevD.34.2206.
- Andrew Gould. Weakly Interacting Massive Particle Distribution in and Evaporation from the Sun. *ApJ*, 321:560, October 1987. doi: 10.1086/165652.
- Andrew Gould. Evaporation of WIMPs with Arbitrary Cross Sections. *ApJ*, 356:302, June 1990. doi: 10.1086/168840.
- Andrew Gould. Cosmological density of WIMPs from solar and terrestrial annihilations. *The Astrophysical Journal*, 388:338, 1992. ISSN 0004-637X. doi: 10.1086/171156. URL <http://adsabs.harvard.edu/doi/10.1086/171156>.
- M L Graham, E Bellm, L Guy, C T Slater, G Dubois-Felsmann, and Melissa Graham. Large Synoptic Survey Telescope (LSST) Data Management LSST Alerts: Key Numbers Change Record. 2019. URL <https://github.com/lsst-dm/dmtn-102>.
- Kim Griest and David Seckel. Cosmic asymmetry, neutrinos and the sun. *Nuclear Physics B*, 283:681–705, Jan 1987. doi: 10.1016/0550-3213(87)90293-8.
- Travis J. Hurst, Andrew R. Zentner, Aravind Natarajan, and Carles Badenes. Indirect probes of dark matter and globular cluster properties from dark matter annihilation within the

- coolest white dwarfs. *Physical Review D - Particles, Fields, Gravitation and Cosmology*, 91(10), 2015. ISSN 15502368. doi: 10.1103/PhysRevD.91.103514.
- Fabio Iocco, Marco Taoso, Florent Leclercq, and Georges Meynet. Main Sequence Stars with Asymmetric Dark Matter. *Phys. Rev. Lett.*, 108(6):061301, February 2012. doi: 10.1103/PhysRevLett.108.061301.
- Ž. Ivezić, S. M. Kahn, J. A. Tyson, B. Abel, E. Acosta, R. Allsman, D. Alonso, Y. AlSayyad, S. F. Anderson, J. Andrew, and et al. LSST: From Science Drivers to Reference Design and Anticipated Data Products. *ApJ*, 873:111, March 2019. doi: 10.3847/1538-4357/ab042c.
- G. Jungman, M. Kamionkowski, and K. Griest. Supersymmetric Dark Matter. *Physics Reports*, 267(5):195–373, 1996. URL <http://arxiv.org/abs/hepph/9506380>{%}5Cnpapers2://publication/uuid/BAC1F598-6461-4943-B5A9-83F7D62754C1.
- D. B. Kaplan. Single explanation for both baryon and dark matter densities. *Physical Review Letters*, 68:741–743, February 1992. doi: 10.1103/PhysRevLett.68.741.
- Mansi M. Kasliwal. Transients in the local Universe: systematically bridging the gap between novae and supernovae. *Bulletin of the Astronomical Society of India*, 39(3):375–385, September 2011.
- Rudolf Kippenhahn, Alfred Weigert, and Achim Weiss. *Stellar Structure and Evolution*. Springer-Verlag Berlin Heidelberg, 2012. ISBN 978-3-642-61523-8.
- Edward W. Kolb and Michael S. Turner. *The early universe*, volume 69 of 1. Addison-Wesley Publishing Company, 1990.
- Savvas M. Koushiappas, James S. Bullock, and Avishai Dekel. Massive black hole seeds from low angular momentum material. *MNRAS*, 354(1):292–304, October 2004. doi: 10.1111/j.1365-2966.2004.08190.x.
- V. L. Krabbendam and D. Sweeney. The Large Synoptic Survey Telescope preliminary design overview. In Larry M. Stepp, Roberto Gilmozzi, and Helen J. Hall, editors, *Ground-based and Airborne Telescopes III*, volume 7733 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 77330D, July 2010. doi: 10.1117/12.857942.
- L. M. Krauss, K. Freese, D. N. Spergel, and W. H. Press. Cold dark matter candidates and the solar neutrino problem. *ApJ*, 299:1001–1006, Dec 1985. doi: 10.1086/163767.
- Pavel Kroupa and Tereza Jerabkova. The initial mass function of stars and the star-formation rates of galaxies. *arXiv e-prints*, art. arXiv:2112.10788, December 2021.
- Ilídio Lopes and Joseph Silk. Solar Constraints on Asymmetric Dark Matter. *ApJ*, 757(2):130, October 2012. doi: 10.1088/0004-637X/757/2/130.

Ilídio Lopes and Joseph Silk. Dark matter imprint on  $^8\text{B}$  neutrino spectrum. *Phys. Rev. D*, 99(2):023008, January 2019. doi: 10.1103/PhysRevD.99.023008.

José Lopes and Ilídio Lopes. Asymmetric dark matter imprint on low-mass main-sequence stars in the milky way nuclear star cluster. *The Astrophysical Journal*, 879(1):50, Jul 2019. ISSN 1538-4357. doi: 10.3847/1538-4357/ab2392. URL <http://dx.doi.org/10.3847/1538-4357/ab2392>.

LSST Science Collaboration, Paul A. Abell, Julius Allison, Scott F. Anderson, John R. Andrew, J. Roger P. Angel, Lee Armus, David Arnett, S. J. Asztalos, Tim S. Axelrod, Stephen Bailey, D. R. Ballantyne, Justin R. Bankert, Wayne A. Barkhouse, Jeffrey D. Barr, L. Felipe Barrientos, Aaron J. Barth, James G. Bartlett, Andrew C. Becker, Jacek Becla, Timothy C. Beers, Joseph P. Bernstein, Rahul Biswas, Michael R. Blanton, Joshua S. Bloom, John J. Bochanski, Pat Boeshaar, Kirk D. Borne, Marusa Bradac, W. N. Brandt, Carrie R. Bridge, Michael E. Brown, Robert J. Brunner, James S. Bullock, Adam J. Burgasser, James H. Burge, David L. Burke, Phillip A. Cargile, Srinivasan Chandrasekharan, George Chartas, Steven R. Chesley, You-Hua Chu, David Cinabro, Mark W. Claire, Charles F. Claver, Douglas Clowe, A. J. Connolly, Kem H. Cook, Jeff Cooke, Asantha Cooray, Kevin R. Covey, Christopher S. Culliton, Roelof de Jong, Willem H. de Vries, Victor P. Debattista, Francisco Delgado, Ian P. Dell’Antonio, Saurav Dhital, Rosanne Di Stefano, Mark Dickinson, Benjamin Dilday, S. G. Djorgovski, Gregory Dobler, Ciro Donalek, Gregory Dubois-Felsmann, Josef Durech, Ardis Eliasdottir, Michael Eracleous, Laurent Eyer, Emilio E. Falco, Xiaohui Fan, Christopher D. Fassnacht, Harry C. Ferguson, Yanga R. Fernandez, Brian D. Fields, Douglas Finkbeiner, Eduardo E. Figuera, Derek B. Fox, Harold Francke, James S. Frank, Josh Frieman, Sebastien Fromenteau, Muhammad Furqan, Gaspar Galaz, A. Gal-Yam, Peter Garnavich, Eric Gawiser, John Geary, Perry Gee, Robert R. Gibson, Kirk Gilmore, Emily A. Grace, Richard F. Green, William J. Gressler, Carl J. Grillmair, Salman Habib, J. S. Haggerty, Mario Hamuy, Alan W. Harris, Suzanne L. Hawley, Alan F. Heavens, Leslie Hebb, Todd J. Henry, Edward Hileman, Eric J. Hilton, Keri Hoadley, J. B. Holberg, Matt J. Holman, Steve B. Howell, Leopoldo Infante, Zeljko Ivezic, Suzanne H. Jacoby, Bhuvnesh Jain, R. Jedicke, M. James Jee, J. Garrett Jernigan, Saurabh W. Jha, Kathryn V. Johnston, R. Lynne Jones, Mario Juric, Mikko Kaasalainen, Styliani, Kafka, Steven M. Kahn, Nathan A. Kaib, Jason Kalirai, Jeff Kantor, Mansi M. Kasliwal, Charles R. Keeton, Richard Kessler, Zoran Knezevic, Adam Kowalski, Victor L. Krabbendam, K. Simon Krughoff, Shrinivas Kulkarni, Stephen Kuhlman, Mark Lacy, Sebastien Lepine, Ming Liang, Amy Lien, Paulina Lira, Knox S. Long, Suzanne Lorenz, Jennifer M. Lotz, R. H. Lupton, Julie Lutz, Lucas M. Macri, Ashish A. Mahabal, Rachel Mandelbaum, Phil Marshall, Morgan May, Peregrine M. McGehee, Brian T. Meadows, Alan Meert, Andrea Milani, Christopher J. Miller, Michelle Miller, David Mills, Dante Minniti, David Monet, Anjum S. Mukadam, Ehud Nakar, Douglas R. Neill, Jeffrey A. Newman, Sergei Nikolaev, Martin Nordby, Paul O’Connor, Masamune Oguri, John Oliver, Scot S. Olivier, Julia K. Olsen, Knut Olsen, Edward W. Olszewski, Hakeem Oluseyi, Nelson D. Padilla, Alex Parker, Joshua Pepper, John R. Peterson, Catherine Petry, Philip A. Pinto, James L. Pizagno, Bogdan Popescu, Andrej Prsa, Veljko Radcka, M. Jordan Rad-dick, Andrew Rasmussen, Arne Rau, Jeonghee Rho, James E. Rhoads, Gordon T. Richards,

- Stephen T. Ridgway, Brant E. Robertson, Rok Roskar, Abhijit Saha, Ata Sarajedini, Evan Scannapieco, Terry Schalk, Rafe Schindler, Samuel Schmidt, Sarah Schmidt, Donald P. Schneider, German Schumacher, Ryan Scranton, Jacques Sebag, Lynn G. Seppala, Ohad Shemmer, Joshua D. Simon, M. Sivertz, Howard A. Smith, J. Allyn Smith, Nathan Smith, Anna H. Spitz, Adam Stanford, Keivan G. Stassun, Jay Strader, Michael A. Strauss, Christopher W. Stubbs, Donald W. Sweeney, Alex Szalay, Paula Szkody, Masahiro Takada, Paul Thorman, David E. Trilling, Virginia Trimble, Anthony Tyson, Richard Van Berg, Daniel Vanden Berk, Jake VanderPlas, Licia Verde, Bojan Vrsnak, Lucianne M. Walkowicz, Benjamin D. Wandelt, Sheng Wang, Yun Wang, Michael Warner, Risa H. Wechsler, Andrew A. West, Oliver Wiecha, Benjamin F. Williams, Beth Willman, David Wittman, Sidney C. Wolff, W. Michael Wood-Vasey, Przemek Wozniak, Patrick Young, Andrew Zentner, and Hu Zhan. LSST Science Book, Version 2.0. *Science*, (November):596, 2009. URL <http://arxiv.org/abs/0912.0201>.
- Paola Marigo, Léo Girardi, Alessandro Bressan, Philip Rosenfield, Bernhard Aringer, Yang Chen, Marco Dussin, Ambra Nanni, Giada Pastorelli, Thaíse S. Rodrigues, Michele Trabucchi, Sara Bladh, Julianne Dalcanton, Martin A. T. Groenewegen, Josefina Montalbán, and Peter R. Wood. A New Generation of PARSEC-COLIBRI Stellar Isochrones Including the TP-AGB Phase. *ApJ*, 835(1):77, January 2017. doi: 10.3847/1538-4357/835/1/77.
- Brian D. Metzger. Kilonovae. *Living Reviews in Relativity*, 23(1):1, December 2019. doi: 10.1007/s41114-019-0024-0.
- A. Möller and T. de Boissière. SuperNNova: an open-source framework for Bayesian, neural network-based supernova classification. *MNRAS*, 491(3):4277–4293, January 2020. doi: 10.1093/mnras/stz3312.
- N. Mowlavi, G. Meynet, A. Maeder, D. Schaerer, and C. Charbonnel. On some properties of very metal-rich stars. *A&A*, 335:573–582, July 1998.
- N. Mowlavi, P. Eggenberger, G. Meynet, S. Ekström, C. Georgy, A. Maeder, C. Charbonnel, and L. Eyer. Stellar mass and age determinations . I. Grids of stellar models from  $Z = 0.006$  to  $0.04$  and  $M = 0.5$  to  $3.5 M_{\odot}$ . *A&A*, 541:A41, May 2012. doi: 10.1051/0004-6361/201117749.
- Kohta Murase and Ian M. Shoemaker. Detecting asymmetric dark matter in the Sun with neutrinos. *Phys. Rev. D*, 94(6):063512, September 2016. doi: 10.1103/PhysRevD.94.063512.
- Gautham Narayan, Tayeb Zaidi, Monika D. Soraisam, Zhe Wang, Michelle Lochner, Thomas Matheson, Abhijit Saha, Shuo Yang, Zhenge Zhao, John Kececioglu, Carlos Scheidegger, Richard T. Snodgrass, Tim Axelrod, Tim Jenness, Robert S. Maier, Stephen T. Ridgway, Robert L. Seaman, Eric Michael Evans, Navdeep Singh, Clark Taylor, Jackson Toeniskoetter, Eric Welch, Songzhe Zhu, and ANTARES Collaboration. Machine-learning-based Brokers for Real-time Classification of the LSST Alert Stream. *ApJS*, 236(1):9, May 2018. doi: 10.3847/1538-4365/aab781.

- Julio F. Navarro, Carlos S. Frenk, and Simon D. M. White. A Universal Density Profile from Hierarchical Clustering. *ApJ*, 490(2):493–508, Dec 1997. doi: 10.1086/304888.
- S. Nussinov. Technocosmology - could a technibaryon excess provide a “natural” missing mass candidate? *Physics Letters B*, 165(1-3):55–58, December 1985. doi: 10.1016/0370-2693(85)90689-6.
- Bill Paxton, Lars Bildsten, Aaron Dotter, Falk Herwig, Pierre Lesaffre, and Frank Timmes. Modules for Experiments in Stellar Astrophysics (MESA). *Astrophysical Journal, Supplement Series*, 192(1), 2011. ISSN 00670049. doi: 10.1088/0067-0049/192/1/3.
- Bill Paxton, Matteo Cantiello, Phil Arras, Lars Bildsten, Edward F. Brown, Aaron Dotter, Christopher Mankovich, M. H. Montgomery, Dennis Stello, F. X. Timmes, and Richard Townsend. Modules for experiments in stellar astrophysics (MESA): Planets, oscillations, rotation, and massive stars. *Astrophysical Journal, Supplement Series*, 208(1), 2013. ISSN 00670049. doi: 10.1088/0067-0049/208/1/4.
- Bill Paxton, Pablo Marchant, Josiah Schwab, Evan B. Bauer, Lars Bildsten, Matteo Cantiello, Luc Dessart, R. Farmer, H. Hu, N. Langer, R. H.D. Townsend, Dean M. Townsley, and F. X. Timmes. Modules for Experiments in Stellar Astrophysics (MESA): Binaries, pulsations, and explosions. *Astrophysical Journal, Supplement Series*, 220(1), 2015. ISSN 00670049. doi: 10.1088/0067-0049/220/1/15.
- Bill Paxton, Josiah Schwab, Evan B. Bauer, Lars Bildsten, Sergei Blinnikov, Paul Duffell, R. Farmer, Jared A. Goldberg, Pablo Marchant, Elena Sorokina, Anne Thoul, Richard H. D. Townsend, and F. X. Timmes. Modules for Experiments in Stellar Astrophysics (MESA): Convective Boundaries, Element Diffusion, and Massive Star Explosions. *The Astrophysical Journal Supplement Series*, 234(2):34, 2018. ISSN 0067-0049. doi: 10.3847/1538-4365/aaa5a8. URL <http://dx.doi.org/10.3847/1538-4365/aaa5a8>.
- Bill Paxton, R. Smolec, Josiah Schwab, A. Gaudsichy, Lars Bildsten, Matteo Cantiello, Aaron Dotter, R. Farmer, Jared A. Goldberg, Adam S. Jermyn, S. M. Kanbur, Pablo Marchant, Anne Thoul, Richard H. D. Townsend, William M. Wolf, Michael Zhang, and F. X. Timmes. Modules for Experiments in Stellar Astrophysics (MESA): Pulsating Variable Stars, Rotation, Convective Boundaries, and Energy Conservation. *The Astrophysical Journal Supplement Series*, 243(1):10, 2019. ISSN 0067-0049. doi: 10.3847/1538-4365/ab2241. URL <http://dx.doi.org/10.3847/1538-4365/ab2241>.
- Kalliopi Petraki and Raymond R. Volkas. Review of asymmetric dark matter. *Int. J. Mod. Phys. A*, 28:1330028, 2013. doi: 10.1142/S0217751X13300287.
- Planck Collaboration, N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo, S. Basak, R. Battye, K. Benabed, J. P. Bernard, M. Bersanelli, P. Bielewicz, J. J. Bock, J. R. Bond, J. Borrill, F. R. Bouchet, F. Boulanger, M. Bucher, C. Burigana, R. C. Butler, E. Calabrese, J. F. Cardoso, J. Carron, A. Challinor, H. C. Chiang, J. Chluba, L. P. L. Colombo, C. Combet, D. Contreras, B. P. Crill, F. Cuttaia, P. de Bernardis, G. de Zotti, J. Delabrouille,

J. M. Delouis, E. Di Valentino, J. M. Diego, O. Doré, M. Douspis, A. Ducout, X. Dupac, S. Dusini, G. Efstathiou, F. Elsner, T. A. Enßlin, H. K. Eriksen, Y. Fantaye, M. Farhang, J. Fergusson, R. Fernandez-Cobos, F. Finelli, F. Forastieri, M. Frailis, A. A. Fraisse, E. Franceschi, A. Frolov, S. Galeotta, S. Galli, K. Ganga, R. T. Génova-Santos, M. Gerbino, T. Ghosh, J. González-Nuevo, K. M. Górski, S. Gratton, A. Gruppuso, J. E. Gudmundsson, J. Hamann, W. Handley, F. K. Hansen, D. Herranz, S. R. Hildebrandt, E. Hivon, Z. Huang, A. H. Jaffe, W. C. Jones, A. Karakci, E. Keihänen, R. Keskitalo, K. Kiiveri, J. Kim, T. S. Kisner, L. Knox, N. Krachmalnicoff, M. Kunz, H. Kurki-Suonio, G. Lagache, J. M. Lamarre, A. Lasenby, M. Lattanzi, C. R. Lawrence, M. Le Jeune, P. Lemos, J. Lesgourgues, F. Levrier, A. Lewis, M. Liguori, P. B. Lilje, M. Lilley, V. Lindholm, M. López-Caniego, P. M. Lubin, Y. Z. Ma, J. F. Macías-Pérez, G. Maggio, D. Maino, N. Mandolesi, A. Mangilli, A. Marcos-Caballero, M. Maris, P. G. Martin, M. Martinelli, E. Martínez-González, S. Matarrese, N. Mauri, J. D. McEwen, P. R. Meinhold, A. Melchiorri, A. Mennella, M. Migliaccio, M. Millea, S. Mitra, M. A. Miville-Deschênes, D. Molinari, L. Montier, G. Morgante, A. Moss, P. Natoli, H. U. Nørgaard-Nielsen, L. Pagano, D. Paoletti, B. Partridge, G. Patanchon, H. V. Peiris, F. Perrotta, V. Pettorino, F. Piacentini, L. Polastri, G. Polenta, J. L. Puget, J. P. Rachen, M. Reinecke, M. Remazeilles, A. Renzi, G. Rocha, C. Rosset, G. Roudier, J. A. Rubiño-Martín, B. Ruiz-Granados, L. Salvati, M. Sandri, M. Savelainen, D. Scott, E. P. S. Shellard, C. Sirignano, G. Sirri, L. D. Spencer, R. Sunyaev, A. S. Suur-Uski, J. A. Tauber, D. Tavagnacco, M. Tenti, L. Toffolatti, M. Tomasi, T. Trombetti, L. Valenziano, J. Valiviita, B. Van Tent, L. Vibert, P. Vielva, F. Villa, N. Vittorio, B. D. Wandelt, I. K. Wehus, M. White, S. D. M. White, A. Zacchei, and A. Zonca. Planck 2018 results. VI. Cosmological parameters. *arXiv e-prints*, art. arXiv:1807.06209, July 2018.

Onno Pols. *Utrecht University, Lecture Notes: Stellar Structure and Evolution*. Number September. 2011. URL [http://www.astro.ru.nl/~onnop/education/stev/\\_utrecht/\\_notes/](http://www.astro.ru.nl/~onnop/education/stev/_utrecht/_notes/).

W. H. Press and D. N. Spergel. Capture by the sun of a galactic population of weakly interacting, massive particles. *ApJ*, 296:679–684, Sep 1985. doi: 10.1086/163485.

Stefano Profumo, Leonardo Giani, and Oliver F. Piattella. An Introduction to Particle Dark Matter. *Universe*, 5(10):213, October 2019. doi: 10.3390/universe5100213.

Troy J. Raen, Héctor Martínez-Rodríguez, Travis J. Hurst, Andrew R. Zentner, Carles Badenes, and Rachel Tao. The effects of asymmetric dark matter on stellar evolution - I. Spin-dependent scattering. *MNRAS*, 503(4):5611–5623, June 2021. doi: 10.1093/mnras/stab865.

A. Renzi. Effects of cosmions in the sun and in globular cluster stars. *A&A*, 171(1-2):121, January 1987.

Marc Schumann. Direct detection of WIMP dark matter: concepts and status. *Journal of Physics G Nuclear Physics*, 46(10):103003, October 2019. doi: 10.1088/1361-6471/ab2ea5.

- D. Scolnic, R. Kessler, D. Brout, P. S. Cowperthwaite, M. Soares-Santos, J. Annis, K. Herner, H. Y. Chen, M. Sako, Z. Doctor, R. E. Butler, A. Palmese, H. T. Diehl, J. Frieman, D. E. Holz, E. Berger, R. Chornock, V. A. Villar, M. Nicholl, R. Biswas, R. Hounsell, R. J. Foley, J. Metzger, A. Rest, J. García-Bellido, A. Möller, P. Nugent, T. M. C. Abbott, F. B. Abdalla, S. Allam, K. Bechtol, A. Benoit-Lévy, E. Bertin, D. Brooks, E. Buckley-Geer, A. Carnero Rosell, M. Carrasco Kind, J. Carretero, F. J. Castander, C. E. Cunha, C. B. D'Andrea, L. N. da Costa, C. Davis, P. Doel, A. Drlica-Wagner, T. F. Eifler, B. Flaugher, P. Fosalba, E. Gaztanaga, D. W. Gerdes, D. Gruen, R. A. Gruendl, J. Gschwend, G. Gutierrez, W. G. Hartley, K. Honscheid, D. J. James, M. W. G. Johnson, M. D. Johnson, E. Krause, K. Kuehn, S. Kuhlmann, O. Lahav, T. S. Li, M. Lima, M. A. G. Maia, M. March, J. L. Marshall, F. Menanteau, R. Miquel, E. Neilsen, A. A. Plazas, E. Sanchez, V. Scarpine, M. Schubnell, I. Sevilla-Noarbe, M. Smith, R. C. Smith, F. Sobreira, E. Suchyta, M. E. C. Swanson, G. Tarle, R. C. Thomas, D. L. Tucker, A. R. Walker, and DES Collaboration. How Many Kilonovae Can Be Found in Past, Present, and Future Survey Data Sets? *ApJ*, 852(1):L3, January 2018. doi: 10.3847/2041-8213/aa9d82.
- Pat Scott, Malcolm Fairbairn, and Joakim Edsjö. Dark stars at the Galactic Centre - the main sequence. *MNRAS*, 394(1):82–104, March 2009. doi: 10.1111/j.1365-2966.2008.14282.x.
- Tracy R. Slatyer. TASI Lectures on Indirect Detection of Dark Matter. *arXiv e-prints*, art. arXiv:1710.05137, October 2017.
- C. Soubiran, J. F. Le Campion, G. Cayrel de Strobel, and A. Caillo. The PASTEL catalogue of stellar parameters. *A&A*, 515:A111, June 2010. doi: 10.1051/0004-6361/201014247.
- D. N. Spergel and W. H. Press. Effect of hypothetical, weakly interacting, massive particles on energy transport in the solar interior. *The Astrophysical Journal*, 294:663, 1985. ISSN 0004-637X. doi: 10.1086/163336. URL <http://adsabs.harvard.edu/doi/10.1086/163336>.
- Gary Steigman, Basudeb Dasgupta, and John F. Beacom. Precise relic WIMP abundance and its impact on searches for dark matter annihilation. *Phys. Rev. D*, 86(2):023506, Jul 2012. doi: 10.1103/PhysRevD.86.023506.
- R. A. Street, M. Bowman, E. S. Saunders, and T. Boroson. General-purpose software for managing astronomical observing programs in the LSST era. In Juan C. Guzman and Jorge Ibsen, editors, *Software and Cyberinfrastructure for Astronomy V*, volume 10707 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, page 1070711, July 2018. doi: 10.1117/12.2312293.
- Marco Taoso, Fabio Iocco, Georges Meynet, Gianfranco Bertone, and Patrick Eggenberger. Effect of low mass dark matter particles on the Sun. *Phys. Rev. D*, 82(8):083509, October 2010. doi: 10.1103/PhysRevD.82.083509.

- G. Torres, J. Andersen, and A. Giménez. Accurate masses and radii of normal stars: modern results and applications. *A&A Rev.*, 18(1-2):67–126, February 2010. doi: 10.1007/s00159-009-0025-1.
- Aaron C Vincent. Dark Matter in Stars. *ArXiv:2009.00663*, 2020. URL <https://github.com/aaronvincent/captngen><http://arxiv.org/abs/2009.00663>.
- Aaron C. Vincent, Pat Scott, and Aldo Serenelli. Possible Indication of Momentum-Dependent Asymmetric Dark Matter in the Sun. *Phys. Rev. Lett.*, 114(8):081302, Feb 2015. doi: 10.1103/PhysRevLett.114.081302.
- Brian Warner. *Cataclysmic variable stars*. Cambridge University Press, New York, 1995. ISBN 9780511586491.
- S. D. M. White and M. J. Rees. Core condensation in heavy halos: a two-stage theory for galaxy formation and clustering. *MNRAS*, 183:341–358, May 1978. doi: 10.1093/mnras/183.3.341.
- R. D. Williams, S. G. Djorgovski, A. J. Drake, M. J. Graham, and A. Mahabal. Skyalert: Real-time Astronomy for You and Your Robots. In D. A. Bohlender, D. Durand, and P. Dowler, editors, *Astronomical Data Analysis Software and Systems XVIII*, volume 411 of *Astronomical Society of the Pacific Conference Series*, page 115, September 2009.
- Norhasliza Yusof, Hasan Abu Kassima, Raphael Hirschi, Paul. A. Crowther, Olivier Schnurr, Richard J. Parkerd, and Simon P. Goodwin. Life and Death of Very Massive Stars. *arXiv e-prints*, art. arXiv:1012.3649, December 2010.
- Andrew R. Zentner. High-energy neutrinos from dark matter particle self-capture within the Sun. *Physical Review D - Particles, Fields, Gravitation and Cosmology*, 80(6), 2009. ISSN 15507998. doi: 10.1103/PhysRevD.80.063501.
- Andrew R. Zentner and Andrew P. Hearin. Asymmetric dark matter may alter the evolution of very low-mass stars and brown dwarfs. *Physical Review D - Particles, Fields, Gravitation and Cosmology*, 84(10):3–8, 2011. ISSN 15507998. doi: 10.1103/PhysRevD.84.101302.
- Monica Zorotovic and Matthias R. Schreiber. Cataclysmic variable evolution and the white dwarf mass problem: A Review. *Advances in Space Research*, 2019. ISSN 18791948. doi: 10.1016/j.asr.2019.08.044.
- Kathryn M. Zurek. Asymmetric Dark Matter: Theories, signatures, and constraints. *Phys. Rep.*, 537(3):91–121, Apr 2014. doi: 10.1016/j.physrep.2013.12.001.