WG 1.4 (theoretical astrophysics) report

Committee members: Mark Krumholz (ANU, Chair), Krzysztof Bolejko (UTas), Tamara Davis (UQ), Christoph Federrath (ANU), Lilia Ferrario (ANU), Alexander Heger (Monash), Jarrod Hurley (Swinburne), Amanda Karakas (Monash), Aaron Ludlow (ICRAR/UWA), Bernhard Mueller (Monash), Danail Obreschkow (ICRAR/UWA), David Pontin (Newcastle), Chris Power (ICRAR/UWA), Daniel Price (Monash), Amit Seta (ANU), Stas Shabala (UTas), Aditi Vijayan (ANU)

Executive summary

Theoretical astrophysics in Australia is an incredibly broad category that embraces research on scales ranging from planetary to cosmological, and using techniques from pencil-and-paper to calculations on the world’s largest computers. Theory represents approximately 25% of the astronomy research effort in Australia, roughly the same as reported in the previous decadal survey; a majority of this effort is in computational theory and numerics, with smaller efforts in analytic theory and in AI and machine learning. The community has had a number of successes over the past decade, including world-class efforts in semi-analytic galaxy formation and simulations of MHD turbulence.

While the theory community in Australia is vibrant, it suffers from the challenge of being very thinly spread. Whereas access to small- and medium-scale computational resources is good, there is poor access to both software support and large-scale computational resources; the latter means that most of the highest-profile Australian computational theory work must be carried out on facilities outside Australia. The theory community is also younger than the average of the astronomical community, resulting in theory being underrepresented in senior decision-making roles relative to its share of the community.

Based on the current state of theoretical astrophysics research in Australia, we recommend the following:

- Theoretical astrophysics research in the coming decade will become increasingly dominated by large-scale “flagship” simulations run with highly-tuned and efficient codes that require significant resources to develop and run. To be competitive internationally, Australian theorists will require access not just to the small-scale resources currently available – access to which should be maintained, particularly for student-led projects – but to larger-scale resources equivalent to roughly one-third of a top-100 level machine. This level of access was also a recommendation of the last decadal plan, but the goal has not been met: actual resources are only about 25% of the target. To be competitive internationally Australian theorists will also need access to increased software development support. The current ADACS Merit Allocation Program for software support is valuable, but is not sufficiently resourced to meet this need.

- The Australian observational landscape in the next decade will be increasingly dominated by large surveys, most prominently by the SKA, but also international surveys where Australia plays a part. For these projects to be successful they will require theoretical support and interpretation, but at present no mechanism beyond standard ARC grants exists to fund this work or the personnel who carry it out. Large surveys should consider pathways to supporting theoretical work that enables them, in much the same way that support is provided for instrumental work.

- The distributed nature and relatively junior status of the theory community creates difficulties in achieving critical mass and influencing the direction of the community. To remedy this, the community should consider ways to develop a theory-focused institute or department similar to the Canadian Institute for Theoretical Astrophysics, possibly with additional fellowships at other institutions to support local theory groups similar to CITA national fellowships.

Full report

WG scope and approach

Working group 1.4 on theoretical astrophysics is distinct from many of the other decadal plan working groups in two important ways. First, it had no precursor in previous decadal plans because the current one is the first to identify theory as a distinct area covered by its own working group. Second, theoretical research is an approach rather than a topic and thus crosses into all areas of astronomy and astrophysics, including those covered by other, topical working groups. Given this situation, this WG report will identify forefront problems from various areas of astronomy and astrophysics focusing specifically on theoretical challenges, and will then review the state of infrastructure for theory and of the theoretical workforce. This report ends by identifying trends and unifying themes for priorities and challenges that cross disciplines and apply to theoretical astrophysics in Australia as a whole. The appendices provide a timeline of the working group's consultations and links to the reports compiled by the various subgroups.

Progress and problems by field

Cosmology and astroparticle physics

Cosmology in the past decade has seen precise measurements of many cosmological parameters through a combination of cosmic microwave background (CMB) and galaxy survey measurements, with large-volume simulations playing a vital role in interpreting the latter. Thus far these measurements have not provided definitive evidence for physics beyond the standard ΛCDM cosmological model, but tensions in the model are beginning to emerge, most prominently over the Hubble parameter, H_0, and the power spectrum normalisation, σ_8. The last decade has also seen the emergence of two major techniques – gravity waves (GWs) and fast radio bursts (FRBs) – that promise to allow ever more precise future measurements. In the coming decade, measurements from these new techniques will join measurements of the epoch of reionization, improved wide-area galaxy surveys, next-generation CMB experiments, X-ray and high-energy gamma-ray observatories, and GW and neutrino observatories. All of these can constrain both cosmological parameters and beyond-standard-model particle physics. Australia plays a leading role in one of these experiments (SKA) and important supporting roles in others.

Theoretical work will be required to extract maximum science from all these measurements. Whereas theory and simulations for large-scale structure formation in the standard ΛCDM model are reasonably advanced, a forefront problem for the next decade that the Australian community is well-positioned to tackle is to extend this level of understanding to alternative dark matter models (e.g., axions, ultra-light dark matter, self-interacting dark matter) to enable their assessment against the data on equal footing with standard CDM. This will require close collaboration between the cosmological simulation, astroparticle physics, observational, and statistical inference communities. A second area where theoretical work will be needed is in exploiting SKA and similar wide-area, low frequency radio facilities. Maximising science return from these will require mock galaxy catalogues from cosmological simulations to be used in carrying out inferences on the data, which in the coming decade will be increasingly extended and refined using machine learning methods to generate realistic predictions without the expense of a full simulation for every possible model or set of parameters. This will require collaborations between galaxy formation theorists and simulators, survey designers, and machine learning specialists. A final area in need of significant theoretical work is FRBs, the source population for which is presently unknown, and will require theoretical work to understand. FRBs also allow exquisite measurement of the state of the gas in the intergalactic and circumgalactic medium, but translating these measurements to constraints on physics will require comparison against models.

Galaxies, star formation, and the interstellar medium

The past decade has seen significant advances in theoretical efforts to model the assembly of galaxies and the stars within them. At larger scales, both cosmological simulations and semi-analytic models (SAMs) now include far more physics than was the standard a decade ago, with newly-introduced processes including AGN and jet feedback, radiation feedback, magnetic fields, metal production and transport, and cosmic rays. Zoom-in simulations are now beginning first attempts to capture the complexity of the multiphase interstellar medium (ISM) starting from cosmological initial conditions. As a result of these improvements, simulations and SAMs can now pass first-order tests such as reproducing the stellar mass to halo mass relation for moderate to massive galaxies. On smaller scales, simulations of the ISM and star formation have also grown in physical realism, and now routinely include sophisticated treatments of radiative transfer, magnetic fields (including non-ideal and dynamo effects), non-equilibrium chemistry and dust physics, and multiple forms of stellar feedback, and are beginning to produce quantitative answers to questions about the origin of the stellar mass distribution, the formation of bound star clusters, and regulation of the star formation rate. We now have multiple small-scale simulations that pass first-order tests such as reproducing the stellar initial mass function and yielding star formation rates and efficiencies that are consistent with observationally-determined values. Finally, intermediate-scale simulations, which target isolated galaxies (often modelled directly on the Milky Way) with a level of physical detail intermediate between small- and large-scale simulations, are beginning to reveal some of the physics of galactic-scale phenomena such as bars and stellar migration.

In the coming decade a forefront challenge will be to build bridges between the large and small scales: whereas large-volume cosmological simulations remain unlikely to be able to resolve the structure of the ISM, zoom-in ones will increasingly be able to do so. For large-scale simulations and SAMs, we will require sub-grid models for unresolvable processes (e.g., pressure from cosmic rays and dynamo-amplified magnetic fields, and energy and momentum injection by supermassive black holes) that are calibrated against better-resolved but smaller-volume simulations. These efforts will be required for interpretation of data from the coming generation of Australian-led HI surveys, which are sensitive to some phases of the ISM but not others. Conversely, questions about the structure and evolution of the ISM on small scales, for example, the origin of turbulence and magnetic fields, and the distribution of metals, cannot be meaningfully answered outside the context of whole galaxies or even cosmological galaxy formation. This will require expanding current small-scale simulations that focus on individual molecular clouds or star-forming regions to larger and more realistic contexts, likely using adaptive techniques to keep the cost manageable. Intermediate-scale simulations may well form a bridge between these two extremes. On all scales these simulations must be tightly coupled to simulated observations, for example HI light cones, an effort that will require extensive collaboration with survey designers.

Stars, planets, and supernovae

The last decade has seen several transformational advances in the understanding of the formation, evolution, and impact of stars, planets, and supernovae. JWST enabled observations of very high-redshift galaxies, revealing early enrichment of the intergalactic medium and high nitrogen abundance that requires an explanation in terms of stellar nucleosynthesis. Simulations have modelled Population III stars and the formation of intermediate-mass black holes in detail. The field of planet formation has been transformed since 2015 by resolved observations of substructures and embedded planets in protoplanetary disks. Both large stellar surveys with chemical tagging and the 2017 discovery of neutron star mergers as the source of r-process elements changed our view of nucleosynthesis, and drove advances in stellar and Galactic chemical evolution models, particularly relating to the s- and i-process. Gravitational wave observations also helped constrain the neutron star equation of state. These have driven qualitative advances in our understanding and modelling of multiple star and planetary systems. Finally, detailed models of core collapse supernovae, tidal disruption events, and black hole formation have been developed.

Significant open problems remain, which are ripe for progress in the coming decade. These include finding and characterising habitable exoplanets and their atmospheres, studying the nature and demographics of the first stars, identifying second-generation stars, and finding *any* sign of pair instability supernovae; understanding the internal mixing and angular momentum transport processes in stars and planets; understanding how the magnetic field of the Sun and stars is generated and how it provides a conduit for energy transport through stellar systems; and revealing the origins and properties of globular clusters and their connection to reionization. We aim to explain the origin of the alpha-Fe bimodality in the Milky Way, and to understand the chemical and structural evolution of galaxies in general, discerning specific contributions from different nucleosynthetic processes.Theoretical studies will benefit from the interaction of Gaia kinematic data with simulations of star and planet formation and galaxy evolution. We expect to make progress on the formation mechanism and early evolution of planetary systems, including ours, aided by long-wavelength follow up of ALMA observations of protoplanetary disks with the SKA and to better understand the dynamics and evolution of stellar and planetary systems such as observed by PLATO. LSST, Roman, and high-resolution spectroscopic surveys will enhance many of these areas and constrain our theoretical models for the progenitors and relative statistics of Type Ia supernovae. Finally, through gravitational waves, we will develop a deeper understanding of the initial mass functions of black holes and neutron stars, and the impact of binary evolution on nucleosynthesis beyond explosive events. In addition to this astrophysical work, opportunities exist to build interdisciplinary collaboration with nuclear physics, planetary science, and the space industry, in part driven by advances in machine learning driven by its application to astronomical data.

High-energy astrophysics: jets, compact objects, and cosmic rays

A major theme in these areas in the past decade has been building bridges from understanding the internal mechanisms powering and producing emission from high-energy phenomena – jets, relativistic particles, and accretion onto compact objects – to understanding the demographics of source populations and the way that these processes affect their larger environments. For AGN jets, this has taken the form of significant progress in understanding how both the triggering of jets and their observable emission depend on their larger environment, and in the production of the first realistic synthetic catalogues for upcoming surveys such as the SKA. For compact objects, a similar effort is starting to connect source populations such as pulsars and neutron stars to large surveys of the time domain and gravity waves. For cosmic rays (CRs) there has been a great deal of progress in developing fluid treatments of CR propagation that make it possible to embed CRs in MHD simulations. This in turn has allowed both a greater understanding of CRs as an important dynamical component of galaxies and improved modelling of CR-driven emission. Another notable area of observational progress that has prompted a great deal of theoretical work, if not yet definitive success, is in understanding the origins of the highest energy CRs that have been detected by facilities such as the Auger Observatory, and their corresponding neutrinos as seen by IceCube.

The forefront challenges for the Australian community in these areas in the coming decade will be to solidify and extend the bridges we have spent the last decade building. In the area of jets, an ongoing challenge will be to ensure that subgrid treatments of jet feedback are capturing the correct physics and cross-check our understanding against the flood of data that will emerge from the SKA. The situation is analogous for compact objects, except that rather than the SKA, the flood of data that will challenge and require us to refine our understanding of source populations will emerge from GW observations with LIGO and Advanced LIGO and large transient surveys such as the Legacy Survey of Space and Time (LSST); theoretical work will be required to make sense of event rates and properties from these surveys. For CRs, the next decade will see the Cherenkov Telescope Array (CTA) become the premier international facility for GeV to PeV photons and upgraded IceCube a similar role for neutrinos, and the theoretical challenge will be to understand the population of sources revealed by these new observatories. We will require combinations of analytic and numerical models to understand both the nature of these sources and their relationship to the larger Galactic and extragalactic ecosystem.

Infrastructure for theoretical astrophysics

Whereas theoretical astrophysics does not directly make use of telescopes, it does rely on instruments in the form of computing hardware for simulations. These have become an increasingly prominent part of the theory landscape of the 21st century. The science challenges outlined above will therefore require significant infrastructure in the form of both computing hardware and software development and support for those using that hardware.

Hardware

The current state of Australian computing infrastructure is discussed extensively in the report from WG 2.3, so here we focus only on the aspects particularly relevant to theory, drawing extensively on the data gathered by that report. Australian theory uses high-performance computing (HPC) hardware in a wide range of applications, including simulations of fluid and N-body systems, radiative transfer calculations, and statistical inference and modelling. Jobs range in scale from small HPC applications using single nodes that are essentially powerful desktop machines to large jobs running on hundreds of thousands of nodes on large machines with specialised hardware at dedicated computing centres. At present, Australia has strong hardware infrastructure for small to medium jobs, which are particularly important for student training and student-led projects, which tend to be smaller in scale and where the longer lead-times associated with access to larger facilities can be problematic. This need is met by the OzSTAR and NT facilities, which dedicate much of their time to astronomy, by a small set of individual university clusters, and by the general-purpose national-level facilities hosted by the National Computational Infrastructure and the Pawsey Supercomputing Centre, which have historically awarded roughly 10% of their compute time to theoretical astrophysics projects. These facilities allow competitive small-to-medium simulations and cover roughly half of the theoretical computational work done in Australia.

However, Australia currently lacks the infrastructure to support large HPC projects; both the total amount of computing time and the data storage available are insufficient to run internationally-competitive large-scale simulations. Indeed, some flagship international simulation projects use more computing power than the *entire* amount of open time available in Australia over *all fields* of science and engineering, not just astronomy. The last decadal survey identified a goal that Australian theoretical astrophysics should have access to compute time comparable to one-third of a “top-100” level machine, and as documented in the WG 2.3 report, this goal is far from being met; actual resources are only 25% of this goal.

Consequently, Australian theorists carry out large projects almost exclusively on overseas platforms to which individual researchers gain access by either personal connections (i.e., someone at an overseas institution leads the project and gets the time, which is then provided to an Australian researcher) or through access schemes that do not require residency. All of the large HPC groups in Australia use this approach. This implies that Australian researchers cannot lead large HPC projects and must instead join international groups as junior partners. They cannot direct the science. The situation for HPC infrastructure for theory is analogous to one where Australian optical astronomers had no access to 8m-class optical telescopes except by working with overseas partners who did have access.

Meeting the science challenges identified above will require a significant increase in HPC resources in Australia. This is crucial not just for theory, but for the success of large observational projects that rely on theoretical work for interpretation and analysis. In the absence of a significant investment, Australia risks being in a position to gather much of the data from large projects but having the bulk of the science from it be done overseas. We return to this topic below in our identification of unifying priorities and themes.

Software and support

Successful use of HPC requires software that is both advanced – in the physics it simulates or the statistical techniques it uses – and capable of efficient parallel execution on modern computing hardware. Because of this need, the scientific impact of simulation software internationally over the last decade is hard to overstate: successful codes such as Arepo, PLUTO, Gizmo, FLASH, and Athena have user-bases of hundreds and produce citation impacts comparable to those of major instruments or surveys. For example, the GADGET-2 code paper (Springel 2005) has been cited more than 5000 times, roughly twice the total number of citations to all GALAH survey papers combined. Despite the number of simulation codes developed in Australia that have international usage, for example the PHANTOM code developed at Monash and the SLUG code developed at ANU, no Australian codes have reached the scale or impact of the highest-impact international codes.

Developing simulation software that contains a large number of physical processes and runs in parallel on large modern machines is beyond the capability of any single individual, and instead requires substantial development teams. Traditionally these teams have been largely astronomers, but in recent years have also begun to include software development specialists as well, and this role is only likely to grow in the coming decade as the continued deployment of GPU-based systems increases the complexity of HPC architecture. Overseas this support is often available through competitive application processes; as an example, the Center for Accelerated Application Readiness (CAAR) program run by the Oak Ridge Leadership Computing Facility in the United States offers multiple FTE of professional programming support over time scales of a year to help upgrade existing codes to run efficiently on exascale machines. The primary service providing support of this type to the Australian theory community at present is the ADACS Merit Allocation Program (MAP). While this program has been a success, its scope and resourcing are far smaller – even on a per capita basis – than that of CAAR or similar programs elsewhere in the world. The challenges faced by the theory community require long-term, stable, and strategically-directable access to larger allocations of developer time that the ADACS MAP is not currently capable of providing. As with hardware, we return to the question of software needs for the coming decade below.

Workforce and careers

The 2024 survey of the Australian astronomical community finds that theoretical research constitutes roughly 25% of community research effort; given differences in methodology, this is consistent with no significant change compared to the 30% theory effort reported in the last decadal survey[1]. We therefore conclude that the size of the theory workforce as a fraction of the total astronomy community in Australia is roughly stable. Within theory, roughly 70% of effort is in computational astrophysics or numerical methods, 15% in machine learning, and 15% in analytic theory. There is some level of theory effort at all major astronomical research institutions in Australia. Because it is newer to Australia, this community is younger than the mean of the Australian astronomical community. It is well-connected internationally, and produces high-impact work.

Surveys and outreach to the theory community carried out as part of this working group's consultation identified a number of ongoing challenges. First, because theory remains a minority area and the community is distributed widely, there are concerns about critical mass. Most Australian astronomical institutions include at most a few theorists, and the country as a whole generally includes only a single theoretical group working in any given area – as examples, almost all star formation / ISM theory in Australia is done at ANU, the great majority of cosmology simulation at UWA, and all supernova theory at Monash. Australia also lacks institutions that take theory as their primary focus; there is nothing in Australia analogous to, for example, the Canadian Institute for Theoretical Astrophysics (CITA) or the theory-focused institutes of the German Max Planck Society. This situation makes it difficult to mount the scale of effort required for major projects such as flagship simulations and major code development. It also creates challenges for talent retention, since someone who is trained as a PhD student or a postdoc in theory in Australia either has to change fields or go overseas for their next position, and those who go overseas then have few prospects to return given the very small number of theory jobs.

A second challenge identified by the theory community is a lack of senior leadership and influence within the broader astronomical community. Because it skews young, the theory community is generally underrepresented in decision-making and leadership bodies compared to its fraction of the community as a whole. For example, in the entire history of Australia Astronomy Ltd. (AAL), only three theorists have ever served on the board, and only once has the board included more than one theorist at a time. This is far below the 25% of community research effort in theory. Similarly, far fewer than 25% of heads of school in Australia are

---

[1] To arrive at our 25% figure, and the analogous figures provided later in this paragraph, we treat theoretical research as work in one of the following four research areas listed as options in the survey: machine learning / artificial intelligence, numerical methods / software, theory - analytical, and theory - computational. We then compute the fraction of each individual's effort in these areas, and compute the mean over all individuals weighted by the fraction of time that individual devotes to research. This differs significantly from the method used in the previous decadal survey, which did not attempt to weight by research fraction or fraction of time on theoretical work. We prefer the present methodology because it avoids overcounting people whose work is primarily observational, but who spend a minority of their effort on theory.

theorists. This situation creates challenges for the leadership of the astronomical community in understanding and prioritising the needs of the theoretical community.

Trajectory and recommendations

By combining the science goals and the review of infrastructure and workforce above, we can identify common trajectories over the next decade, and find a few cross-disciplinary priorities and recommendations for the next decade.

Flagship simulation projects and codes

One prominent unifying theme that appears in many of the discussions of individual science topics above is the increasing role of “flagship” simulation codes and projects on the international landscape. Examples include projects such as Illustris and its descendants run using the Arepo code in cosmology and galaxy formation, the Starforge simulation suite built around the Gizmo code in star formation, and simulations of accretion discs built around the Athena code. These projects achieve their outsized impact in part by sheer scale – they typically use hundreds of millions of CPU-hours with a highly-optimised code, thereby reaching combinations of volume and resolution inaccessible to smaller efforts – and in part by data sharing and re-use enabled by good documentation, public-facing archives, and a culture of sharing and collaboration.

The importance of these flagship simulations and codes is only likely to grow in the coming decade as computing systems gain in complexity and specialised hardware for accelerating computations, particularly GPUs, becomes the norm. This trend will both enable larger and faster simulations and demand more complex software that will become increasingly difficult for individual researchers and small groups to maintain. The days of every group writing its own hydro or N-body scheme are truly over.

A central challenge for Australian theoretical astrophysics is that no Australian group has successfully executed a simulation campaign at the international flagship scale. Despite the challenge of data sharing and curation being partly alleviated by the Theoretical Astrophysics Observatory project, overall simulation resources in Australia remain insufficient, and personnel too thinly-spread, to carry out internationally-competitive efforts. To change this situation in the next decade, it is important that Australia overcome these challenges and begin to execute at least some simulation campaigns at the same scale as the large overseas groups. This, in turn, will require meeting the unmet goal from the last decadal plan of having access to computer hardware that is equivalent to roughly one-third of a top-100 level machine.

The problem of running flagship simulations is also tightly coupled with the challenge of simulation software. As discussed above, flagship-level projects require highly-optimised, high-accuracy codes that can make use of increasingly complex modern computer architectures. Australia currently lacks the critical mass and level of professional support required for code development at scales comparable to the largest overseas groups. The ADACS Merit Allocation Program (MAP) is helpful, but currently lacks the resources to match the dedicated and long-term support available to overseas groups. If Australia is to be internationally competitive, it needs an investment in software support to at least partly close this gap. The existing ADACS MAP should be extended with an eye to providing long-term support to flagship code efforts.

Interfacing with large surveys

A second unifying theme that arises from the discussion above is the need for theory to be embedded in large-scale survey projects. Australia is justly well-renowned for its survey science, with flagship projects such as 2dF being among the most productive in the world at their times. Despite Australia having invested heavily in making data collection and curation from surveys a success, there has not always been comparable support for the necessary theory work to interpret and exploit the results.

The coming decade promises a number of new large surveys with Australian leadership, most prominently from the SKA, but also from international facilities in which Australia participates that cover almost all scientific domains, for example CTA and LSST. These surveys will increasingly dominate the Australian astronomical landscape in the next decade, and will drive theoretical as well as observational work. To maximise science return from large surveys, however, there is an urgent need to fund the theoretical work related to them. The cost of such work is relatively modest compared to the cost of carrying out a large survey but increases the science return many-fold. At present, however, there exists no real mechanism to fund either projects or personnel to do theoretical work in support of surveys. There is an urgent need to create such pathways, in much the same way that the community has recognised the need to support instrumentation work. In addition to grant funding, institutions leading such surveys should consider continuing staff positions for theory support for large surveys as they currently support observers and instrumentalists.

Support for a theory-focused institute

Several of the challenges for theory mentioned above – lack of representation at senior leadership levels, lack of career paths, lack of critical mass – point toward the need for a flagship institution for Australian theory. The obvious model here is the Canadian Institute for Theoretical Astrophysics, which provides leadership, critical mass, and career paths in Canada, a community similar in size and scientific impact to the Australian one. Particular aspects of the CITA model would clearly address some of the issues in Australian theory. For example, CITA offers postdoctoral fellowships which provide a career path for theorists, and supports critical mass by distributing roughly half of these fellowships to other institutions around Canada, thereby bolstering theory efforts elsewhere, not just at CITA’s host institution (University of Toronto). Similarly, concentrating more than a few theorists in one place would provide a basis for ongoing support for software development and large-scale simulation efforts.

Appendix A: Working approach and record of consultations

Given the broad scope of this working group, the group decided to first divide into a series of subcommittees, each of which carried out consultations with the community both asynchronously and via town hall meetings. These sub-groups then produced reports (see Appendix B) that flowed to a drafting sub-committee, which merged them together to produce this report. The subgroups and their chairs were as follows:

- Cosmology and particle astrophysics (Chair: Chris Power)
- Galaxies, interstellar medium, and star formation (Chair: Aaron Ludlow)
- Stars, planets, and supernovae (Chairs: Alexander Heger, Daniel Price)
- Cosmic rays, jets, and compact objects (Chairs: Lilia Ferrario, Amit Seta, Stas Shabala)
- High-performance computing hardware and software (Chair: Bernhard Mueller)
- Workforce, employment, and career paths (Chairs: Amanda Karakas, Stas Shabala)
- Drafting (Chairs: Mark Krumholz, Tamara Davis)

The subcommittees held town halls as follows (all by zoom):

- Cosmic rays and jets – Thursday, 21 March, 12:30 - 1:30 pm AEDT
- Cosmology and astroparticle physics (jointly hosted with the extragalactic and cosmology working group) – Monday, 25 March, 1 - 2 pm AEDT
- Stars, planets, and supernovae (jointly hosted with the stars, planets, and Milky Way working group) – Wednesday, 3 April, 1 - 2 pm AEDT
- Workforce and employment – Friday, 19 April, 2 - 3 pm AEST
- Galaxies, ISM, and star formation – Monday, 22 April, 12 - 1 pm AEST
- High performance computing and simulations – Tuesday, 23 April, 2 - 3 pm AEST

Appendix B: Subcommittee reports

Here we provide the reports provided by each of the individual subgroups.

- Cosmology and particle astrophysics (Chair: Chris Power) – Theory Astro Sub-WG Report: Cosmology & Astroparticle Physics
- Galaxies, interstellar medium, and star formation (Chair: Aaron Ludlow) – White Paper: Theory and Galaxies & Cosmology WGs
- Stars, planets, and supernovae (Chair: Alexander Heger, Daniel Price) – Meeting result summary by DP
- High-energy and compact objects (Chair: Lillia Ferrario) – Compact Stars
- Cosmic rays (Chair: Amit Seta) – Cosmic rays subgroup report
- Jets (Chair: Stas Shabala) – Sub-WG on jets report
- High-performance computing hardware and software (Chair: Bernhard Mueller) – WG Theory HPC Subcommitte Draft report
- Workforce, employment, and career paths (Chair: Amanda Karakas, Stas Shabala) – Town Hall - Workforce Subcommittee