UTS UNIVERSITY OF TECHNOLOGY SYDNEY

Building a Climate Responsive Neighbourhood

A report on low-cost environmental monitoring at Melrose Park

TULIP

A project of the Technology for Urban Liveability Program (TULIP)

This report was written and produced solely by the University of Technology Sydney and opinions and positions discussed are not the official positions of project partners.

Published in November 2020

Project Partners

City of Parramatta

PAYCE

mProjects

[ui!] the urban institute

ESRI Australia

Lead author

Andrew Tovey, Senior Research Consultant, Institute for Sustainable Futures, University of Technology Sydney

Key contributors

Frank Zeichner, Director, Business Development (Smart Cities), Faculty of Engineering and IT, University of Technology Sydney

Kerryn Wilmot, Research Principal, Institute for Sustainable Futures, University of Technology Sydney

BUILDING A CLIMATE RESPONSIVE NEIGHBOURHOOD

A REPORT ON LOW-COST ENVIRONMENTAL MONITORING AT MELROSE PARK

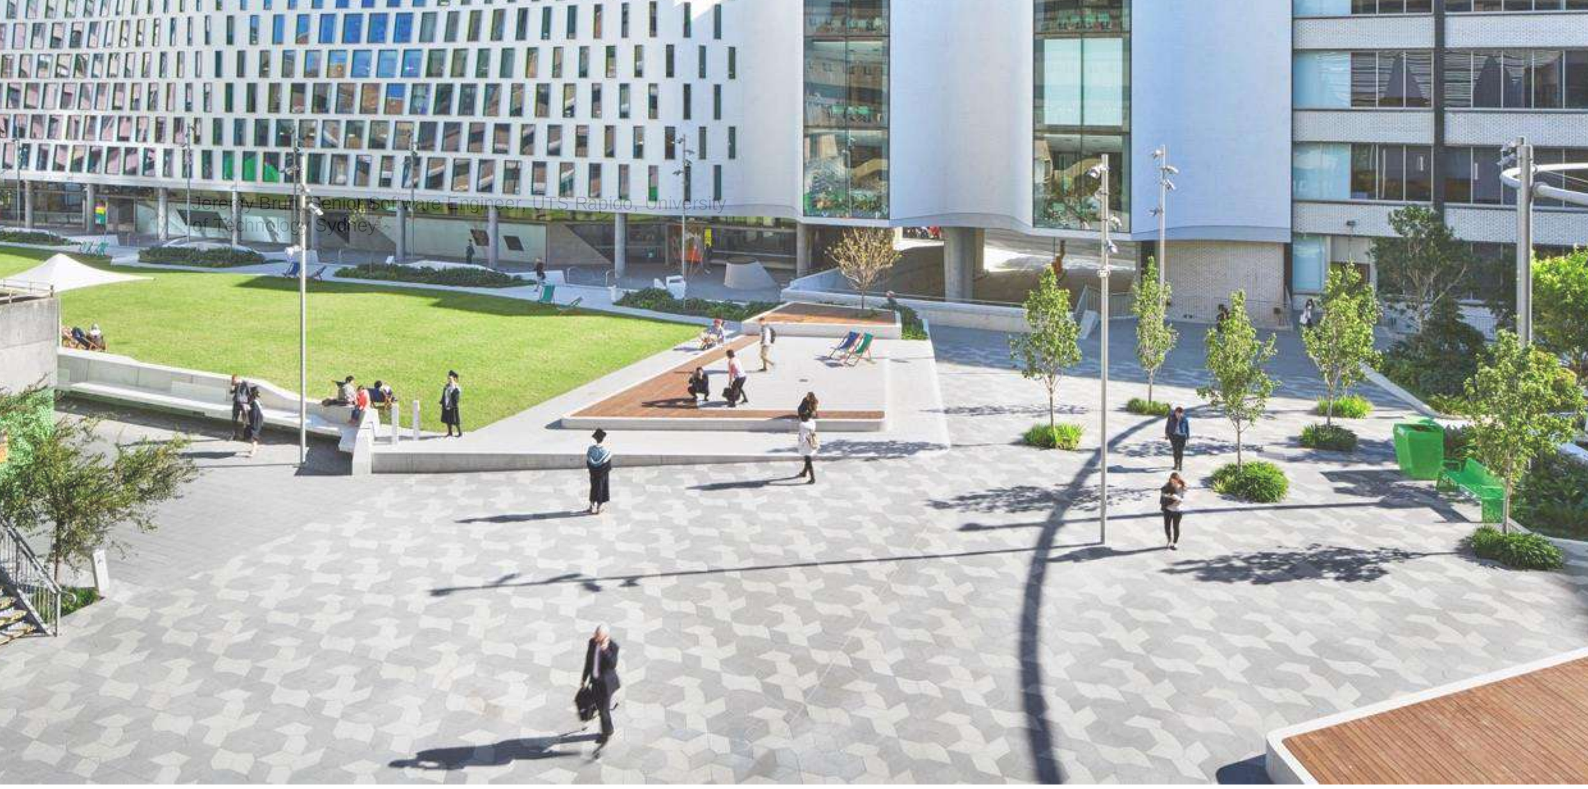

UTS Research Team

- Andrew Tovey
- Kerryn Wilmot
- Frank Zeichner
- Jeremy Brun
- Yohan Kim (urban heat)
- Quang Ha (air quality)
- Xiaojun Qui (noise)
- James Ball (stormwater)

Citation

Tovey, A., Wilmot, K., Zeichner, F. and Brun, J., 2020. *Building A Climate Responsive Neighbourhood: A Report On Low-Cost Environmental Monitoring At Melrose Park*. Sydney: UTS Institute for Sustainable Futures.

Institute for Sustainable Futures
University of Technology Sydney
PO Box 123 Broadway, NSW, 2007
www.isf.uts.edu.au

Contents

Introduction 5
- About the project 5
- About this report 6
- Low-cost environmental sensing: an emerging smart city paradigm 7
- Four use cases 9

Use case 1: Urban Heat 10
- Understanding urban heat at Melrose Park 11
- Urban heat monitoring study design 12
- Constraints of the urban heat monitoring study 14
- What data did we collect and how can we use it? 16
- Insights from urban heat monitoring at Melrose Park 17
- Future directions for urban heat monitoring at Melrose Park 21

Use case 2: Air Quality 24
- Understanding particulate pollution at Melrose Park 25
- Particulate monitoring study design 26
- Constraints of the particulate monitoring study 28
- What data did we collect and how can we use it? 30
- Insights from particulate monitoring at Melrose Park 31
- Future directions for particulate monitoring at Melrose Park 33

Use case 3: Noise 36
- Understanding noise at Melrose Park 37
- Noise monitoring study design 38
- Constraints of the noise monitoring study 39
- Insights from noise monitoring at Melrose Park 41
- Future directions for noise monitoring at Melrose Park 44

Use case 4: Stormwater 46
- Understanding stormwater at Melrose Park 47
- Stormwater monitoring study design 49
- Constraints of the stormwater monitoring study 56
- What data did we collect and how can we use it? 58
- Insights from stormwater monitoring at Melrose Park 61
- Future directions for stormwater monitoring at Melrose Park 67

Practical lessons for designing and delivering low-cost environmental sensor networks 73
- Collaboration and successful project delivery 74
- Device deployment choices and the iterative nature of smart city experimentation 79
- Technical design 81

References 82

Introduction

About the project

Melrose Park: Smart planning for Climate Responsive Neighbourhoods was a twenty month project led by the City of Parramatta Council, in partnership with the University of Technology Sydney, PAYCE, mProjects, the urban institute [ui!], and ESRI Australia, with significant grant support from the Australian Government, through the second round of the *Smart Cities and Suburbs program*. The project commenced in November 2018 and was completed in June 2020.

The project explored the application of real-time low-cost environmental sensing technology to four critical use cases at Melrose Park, on the eastern edge of the City of Parramatta. The area is a mix of residential and light industrial zones on the northern side of the Parramatta River and is dominated by a thirty hectare brownfield site, formally a light industrial park, now under development by PAYCE. The new Melrose Park development is being constructed over a ten year period and is expected to include up to ten thousand new residences as well as mixed commercial and retail space. The development features a significant number of new buildings, extensive parkland and tree canopy cover, as well as new transport links. Combined with the significant growth in local population resulting from new residences, the liveability of the area is expected to change substantially as the site develops and matures. By monitoring the local microclimate at Melrose Park, the project sought to establish a baseline of environmental data, against which future changes can be compared. The provision of real-time data, aggregated from multiple sources, also supported a range of experimental operational outcomes for the management and future planning of the development. The project worked closely with the local community, exploring citizen's concerns about environmental liveability and delivering novel and accessible new technologies, such as an open community LoRaWAN network and a public data discovery dashboard.

Image: An urban heat sensor deployed on the development site (credit: Andrew Tovey)

About this report

This report details the methodology and research findings relating to low-cost environmental sensing at Melrose Park, NSW. The focus is on environmental science and the use of emerging data gathering technologies to support new insights for urban liveability. The work sits within a broader exploration of smart city technology.

The *Melrose Park: Smart planning for Climate Responsive Neighbourhoods* project was about more than the four environmental use cases presented in this report. Project partners also sought to explore the operational and technical challenges of smart city development.

For a more comprehensive overview of all aspects of the project, including details of project governance and co-design approach, smart city technologies, data architecture, storage and flows, community engagement programs and social research findings, please refer to:

> City of Parramatta, 2020. *Melrose Park: Blueprint For Climate Responsive Neighbourhoods - A Report On The Framework, Learnings And Findings From The Smart Cities And Suburbs Initiative*. Parramatta: City of Parramatta Council.

Low-cost environmental sensing: an emerging smart city paradigm

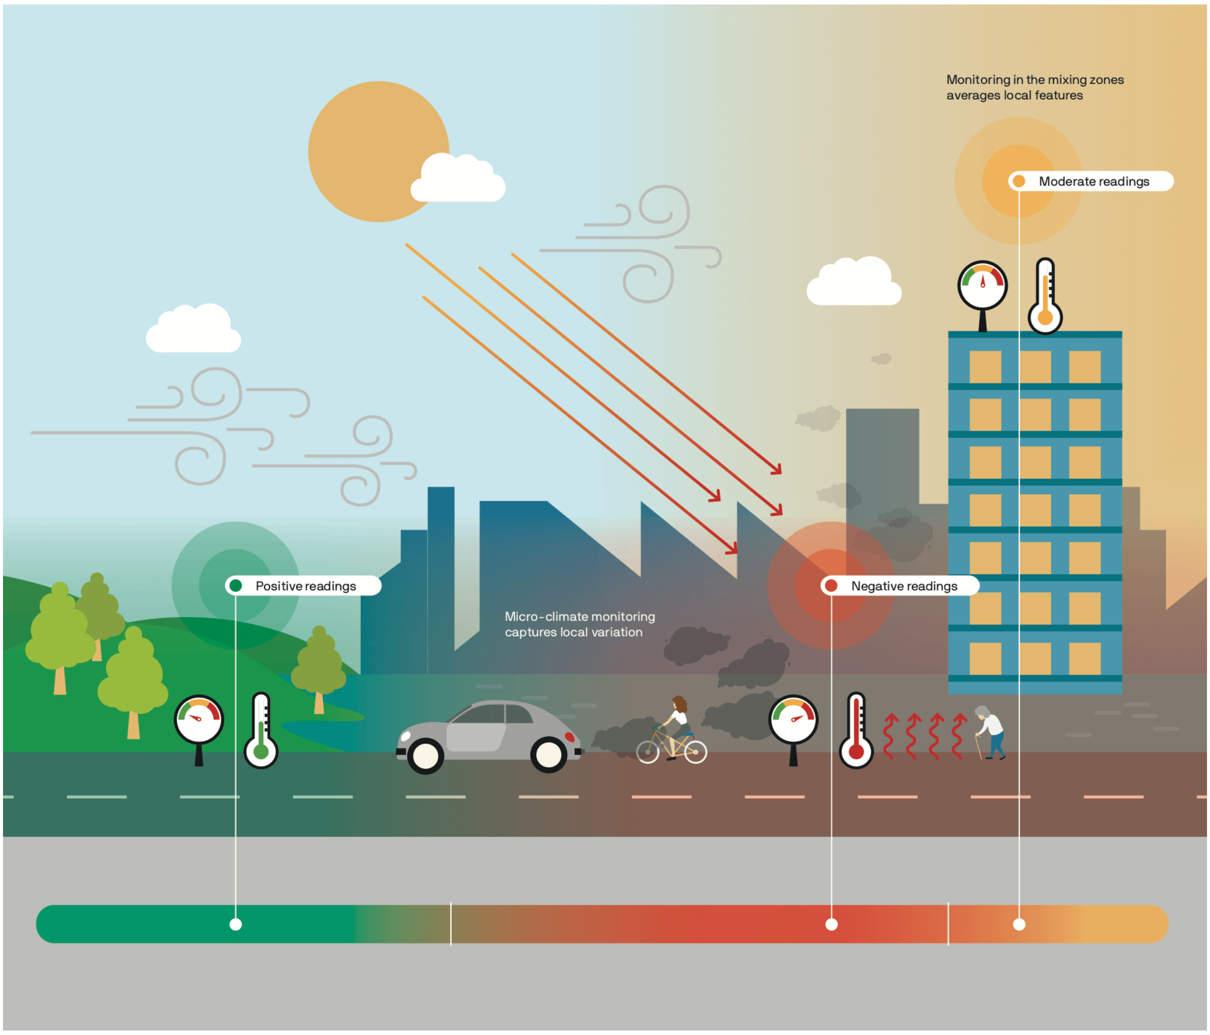

Over the past decade, the concept of the Smart City has risen to prominence in Australia and around the world. While initially characterised by the application of rapidly emerging communications, computing, internet of things and big data management technologies to the urban environment, the Smart City has transcended a purely technological definition and is now generally understood to occupy the intersection of technology, societal, environmental and governance considerations of city development. The aim is to produce tangible positive outcomes for citizens in the context of significant intersecting challenges.

Cities worldwide face enormous challenges and these are expected to increase dramatically in coming decades. Cities are under strain from rapid population increase and densification, with over two thirds of the global population expected to live in cities by mid-century (United Nations, 2015[1]). This places increasing strain on existing infrastructure and services, as well as on the health of the urban environment, resulting in increasing pressure upon the lives and wellbeing of city dwellers[2]. The impacts of climate change by 2050, while still dependent upon the degree of mitigation achieved in coming years, are projected to significantly intensify the challenges faced by cities, even under best case scenarios. In Australia and around the world, extreme heat, bushfires and smoke are a growing concern for public health, infrastructure and city economies. Flooding, intense storms and coastal

---

[1] United Nations (2015). World urbanization prospects. The 2014 revision. New York: Department of Economic and Social Affairs, http://esa.un.org/unpd/wup/Publications/Files/WUP2014-Report.pdf (accessed 22.1.2017).

[2] Desa, U.N., 2016. Transforming our world: The 2030 agenda for sustainable development.

erosion have also increased in recent decades, with increasing future impacts that are directly attributable to climate change[3],[4].

The challenges faced by cities underscore an urgent need for dramatic changes to urban thinking, with a more holistic, systemic and long-term perspective gaining increased prominence. In the past two decades, the role of ICT has risen as a powerful new tool in the urbanist playbook, with trends of widespread interconnectivity and data-driven outcomes forming the technological backbone of a rising smart city agenda. It is argued that the complexity of the urban system and its interrelated co-dependent parts, demands the complex information processing and operational outcomes promised by a rapidly evolving suite of smart technologies[5]. These technologies promise to increase efficiency, reduce redundancy, connect disparate parts of a complex urban system and as such, address foundational sustainability challenges faced by cities.

A *Climate Responsive Neighbourhood* is a term that describes a location where strategies are available to enable the monitoring of and adaption to changes over time. In its broadest sense these can include various social, technological, environmental and political factors that can impact a neighbourhood and the wellbeing of its residents. At the heart of this concept is the practical application of low-cost smart sensor technologies to the urban environment. Advances in sensors, micro-processing, communications and big data management technologies are heralding a new paradigm that is enabling interested parties to capture detailed real-time hyper-local environmental data that was previously unavailable. The ways in which this data might be used to improve urban liveability, promises to redefine almost every aspect of how we design and manage our cities over the coming years.

The project at Melrose Park was a Smart City pilot project, working at the leading edge of experimentation and innovation in a rapidly emerging space. The aim was to explore real-world operational use cases for real-time hyper-local environmental data. It is hoped that insights emerging from this work will support a variety of future applications of smart sensor technology to brownfield development projects and to smart city agendas more broadly.

[3] Climate Council, 2017. *Intense Rainfall And Flooding: The Influence Of Climate Change*. [online] Climate Council. Available at: <https://www.climatecouncil.org.au/uploads/5dafe61d7b3f68d156abd97603d67075.pdf> [Accessed 27 August 2020].

[4] Steffen, W., Hunter, J. and Hughes, L. (2014). *Counting the costs: Climate change and coastal flooding*. [online] Climate Council. Available at: https://www.climatecouncil.org.au/uploads/coastalflooding.pdf [Accessed 7 Mar. 2020].

[5] Bibri, S.E. and Krogstie, J., 2017. Smart sustainable cities of the future: An extensive interdisciplinary literature review. *Sustainable cities and society*, *31*, pp.183-212.

Four use cases

The project explored four distinct use cases for hyper-local environmental data.

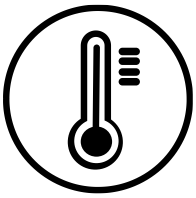

Use case 1: Urban Heat

Urban heat is a growing concern in Australian cities during the summer months, with the frequency and severity of extremely hot days increasing with climate change. Heat has a major impact on human health and wellbeing, however the design and management of our built environment can play a large role in mitigating that impact. By collecting detailed continuous data about urban heat across varied microclimates at Melrose Park the project sought to understand how differences in the urban environment can create localised variation in urban heat. The aim was to develop new insights that might inform more effective heat mitigation in the area.

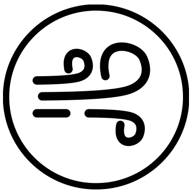

Use case 2: Air quality

Urban air quality is a major and rising health concern globally. Pollution sources include vehicle emissions, industry, wildfire smoke, and dust from drought, desertification, mining and construction. The focus of this project was dust generated by construction activity on the Melrose Park development site. Low-cost particulate sensors were used to gather continuous data at fifteen locations, allowing analysis of where, when and under what conditions dust is generated on site and how it might impact the surrounding residential community.

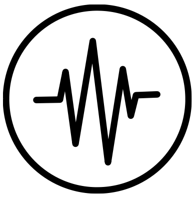

Use case 3: Noise

Noise pollution can have a significant impact on human health and wellbeing and has been linked to heightened stress and cardiovascular disease. Sources of noise in the urban environment include motor vehicles, planes, trains, civil works (such as roadbuilding), and construction. The Melrose Park development is notably large and has a ten year construction period, with the potential for significant construction noise. By deploying smart ambient noise monitors around Melrose Park, the project sought to understand where and when noise occurs on site, what sorts of activities it is connected to, and how much it impacts the surrounding community.

Use case 4: Stormwater

Urban stormwater management is a complex area of concern, with implications for flood and erosion mitigation as well as the health of urban waterways. Construction sites disturb the ground for prolonged periods and if not properly managed, heavy rainfall can cause the release of sediment and soil contaminants into rivers and the ocean. A major change in land-use also alters the stormwater profile of an area, which may either improve or worsen flood risk. By using smart sensors to monitor the flow of stormwater at Melrose Park and water quality in the Parramatta River, the project sought to establish a baseline stormwater profile for the area as well as explore new approaches to flood and erosion mitigation on the construction site.

Use case 1

Urban Heat

Use case 1: Urban Heat

- **Understand baseline urban heat at Melrose Park on a microclimate scale, including:**
  - Temporal variations (daily and seasonal)
  - Extreme heat events
  - Varied land-use, ground cover, and influence of trees and vegetation
  - Relationship of urban heat to weather conditions
- **Use baseline data and insights into microclimate variance to:**
  - Support long-term urban heat management and mitigation in the Melrose Park area
  - Inform broader urban heat strategy in the LGA (Council) and for future development sites (PAYCE)

Understanding urban heat at Melrose Park

The study of urban heat is concerned with the impacts of temperature and humidity on our cities and communities. Extreme heat carries significant implications for human health and wellbeing and has been linked to a range of social issues, from increased crime and violence to higher incidence of accidents. Heat is a major concern for workplace health and safety, particularly for people working outdoors on a construction site. As Melrose Park develops, the community will grow and with hotter and longer summers forecast as a result of climate change, the wellbeing and comfort of that community will relate strongly to urban heat and its mitigation.

Urban heat can vary over very short distances, in what we refer to as urban microclimates. A number of factors contribute to this localised variation, notably: the radiative potential of nearby surfaces and bodies; shading; aspect; the degree of enclosure or exposure to wind and rain; nearby vegetation (type, quantity and height); and to some degree, elevation.

Radiative potential relates to a number of factors, including the thermal mass of nearby surfaces (e.g. a building, a road or a body of water), the colour and reflectivity of those surfaces, and the sun exposure of those surfaces. It is this radiative potential that generates the so-called ‘urban heat island effect’, which is the effect of large amounts of concrete and asphalt absorbing the sun’s energy and radiating it back as heat, often well into the evening. This effect causes urban areas to have an ambient temperature much higher than the surrounding countryside. It also means that at certain times of day a parking lot or highway can have localised temperatures that are significantly warmer than say, a city park.

Shading from trees and buildings is another major variable and is tied to aspect. A location with tall trees or buildings to its north is likely to experience a lot of shade and cooler temperatures in the middle of the day. This will reduce the amount of solar energy that surfaces in the location absorb, in turn reducing their radiative potential and resulting in cooler ambient air temperatures in the late afternoon and evening as well. If shade comes from the east then it mitigates early morning heat and this can help to reduce midday temperatures. Aspect also relates to so-called ‘sun traps’, which tend to be north facing locations sheltered from wind and rain, with thermal mass (generally a building) to the south, east or west.

Vegetation has the greatest impact in terms of tree shade, however ground cover is also a major factor, even in unshaded locations, because it prevents the sun from directly heating up the ground and becoming a radiative thermal mass. Plants also transpire (they evaporate water through their leaves) and this can create a significant localised cooling effect. Some tree species are particularly effective for evapotranspirative cooling and combined with the shade they provide, they can create localised microclimates that are as much as ten degrees cooler than the surrounding area.

Image: An urban heat sensor deployed on the development site (credit: Andrew Tovey).

Urban heat monitoring study design

Urban heat is measured as a combination of temperature and humidity. Smart low-cost sensors that measure these parameters are amongst the lower cost options for environmental monitoring, in the range of $200-$1000 per device, depending which model is chosen. For this project, we chose to work with the Netvox R712 temperature and humidity sensor node ($250 per unit), which comes with a weather-proof radiation shield that helps to minimise the effects of direct sun and infrared thermal interference on sensor readings, while maintaining adequate air flow.

In addition to thirty seven Netvox nodes, the Melrose Park urban heat monitoring network includes fifteen TULIP Environmental Monitoring Systems (EMS) that monitor temperature and humidity alongside air quality and noise. This provided a total of 52 urban heat monitoring locations across the study area.

The number of devices used was in part limited by cost, but was also higher than other types of node due to a commitment to a deployment of a minimum number of devices through the project. With temperature and humidity nodes being the lower cost device option it was possible to deploy more of them for the same budget allowance, helping us to meet our deployment quota.

From a research perspective, using a larger number of urban heat nodes makes sense given that temperature and humidity are known to vary significantly over very short distances, creating complex and variable microclimates. The number of devices chosen allowed a relatively dense on-site deployment density of roughly one device every 100-120m and an off-site deployment of roughly one device every 150-200m, with fairly uniform coverage across an area of approximately one square kilometre. All devices were deployed at three metres off the ground, to standardise readings and avoid too much direct thermal radiation from the ground. Large thermal masses such as thick metal poles or north-facing walls were avoided, to reduce thermal interference.

The choice of specific sensor deployment locations was a balance between achieving a minimum desired density and coverage of the target area, and capturing as full a range as possible of the various microclimate types present in the study area. It was hypothesised that on a hot day, an exposed location covered with concrete (such as a large parking lot) will be significantly hotter than a location by the Parramatta River, shaded by trees. Furthermore, a large number of location types with various characteristics presented themselves, such as: exposed grass, exposed concrete, large shade trees, high building shadow from north, riverside parkland, residential streets with mixed ground cover, and a whole range of finer variables and intermediate location types that were expected to produce subtly different microclimates. Following extensive field observations, sensors were deployed to explore these variations in location type.

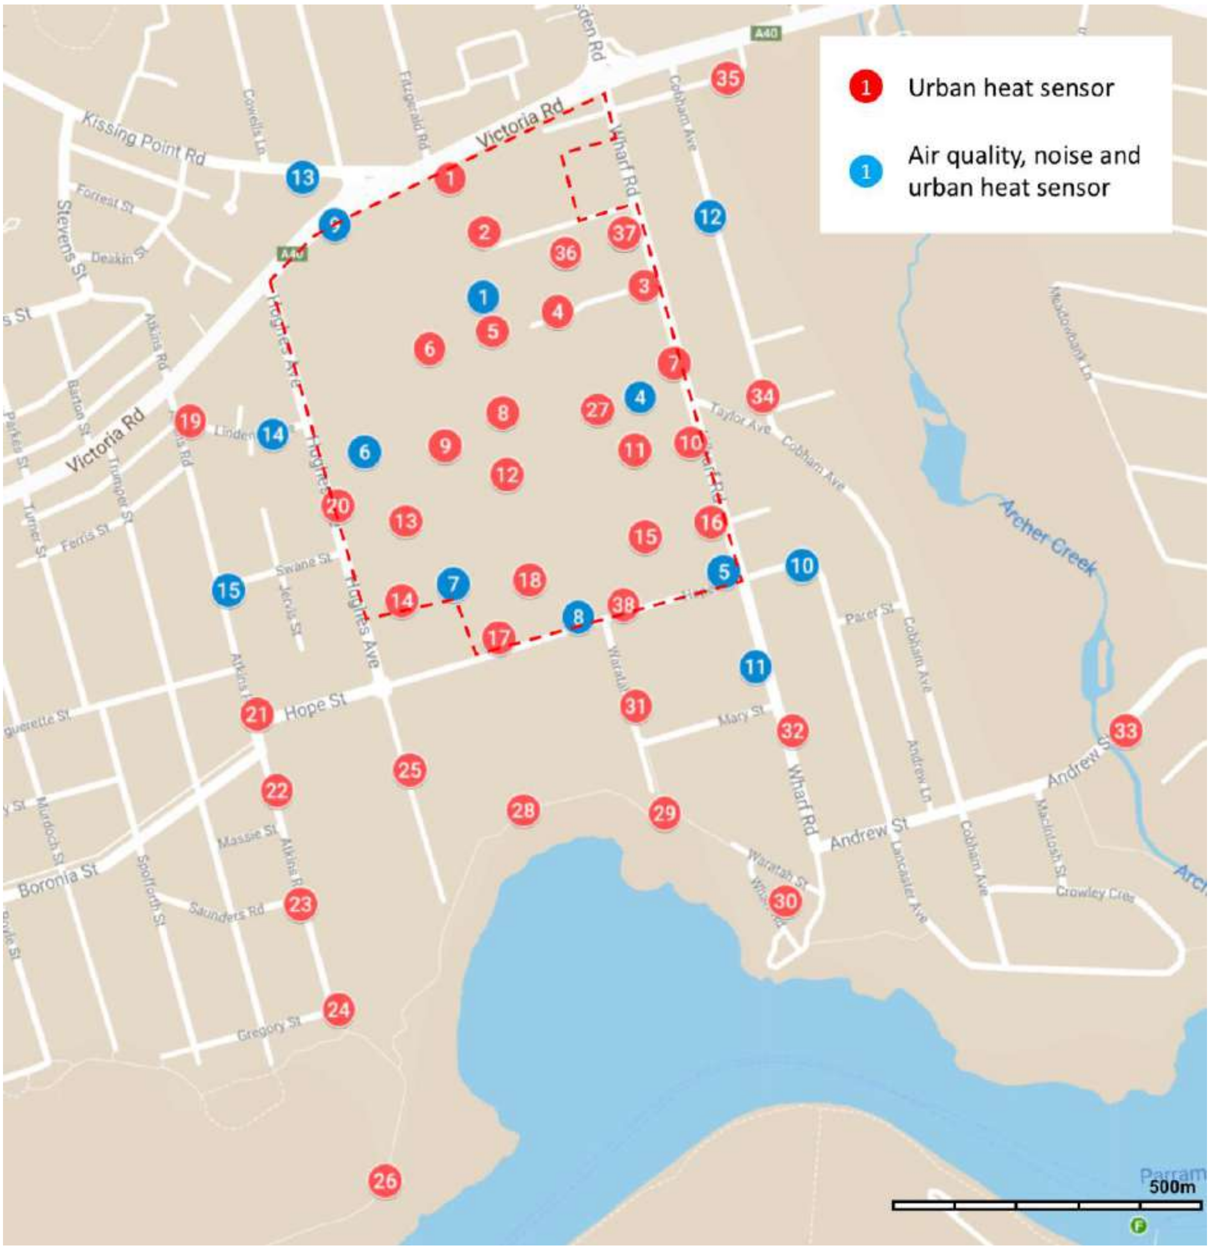

Image: The distribution of urban heat monitoring nodes deployed for the study. Red dots are the Netvox R712; blue dots are the TULIP EMS

Constraints of the urban heat monitoring study

Capturing microclimate deployment context

Urban heat varies by microclimate, with significant variation in temperature and humidity possible over very short distances. Ground cover, shading, nearby thermal mass, and exposure to wind are all factors that influence this microclimate variation. Standardised meteorological monitoring stipulates where and how a sensor can be deployed in order to minimise the influence of very localised conditions. This allows readings from different locations to be compared. One constraint of standardised monitoring is that the locations that are appropriate for monitoring in a given area tend to be somewhat limited. The challenge emerges when we specifically aim to undertake microclimate monitoring, because to do that we must deploy sensors in many different locations, the majority of which do not comply with meteorological standards. Indeed, rather than seeking to minimise our measurement of very localised conditions, we are seeking to capture and better understand them.

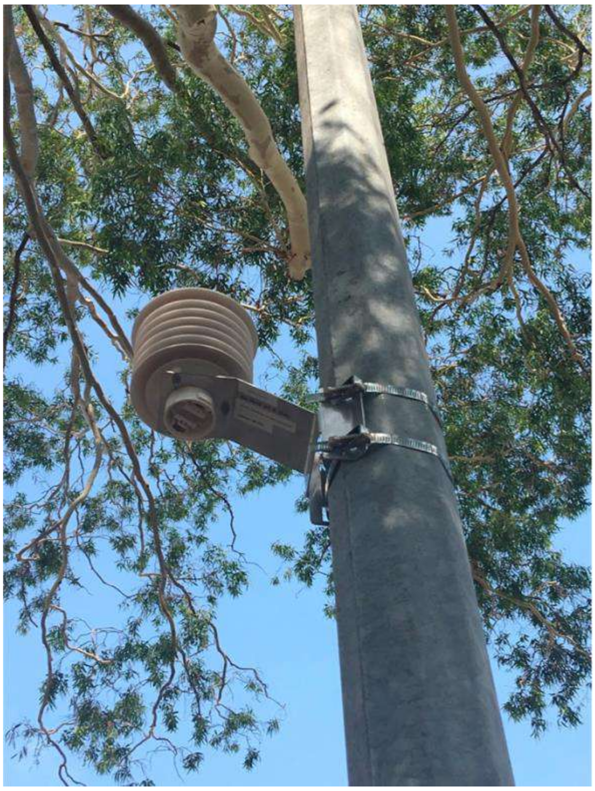

Images: Examples of the differing deployment contexts of our sensors (credit: Andrew Tovey).

Microclimate data can start to tell us a great deal about urban heat at a human scale but to make sense of it we must understand the context in which it was collected. To do this we need standardised ways of measuring and recording that context: the ground cover, shading, nearby thermal mass, and exposure to wind. This project lacked these standards and indeed, we are not aware of existing standards that might be applied due to the still emerging nature of low-cost sensing and microclimate monitoring. While key aspects of each deployment were recorded, they do not capture all of the contextual variables known to influence urban heat, for the most part due to the shear complexity involved. We hypothesise that distributed low-cost urban heat monitoring will grow into a significant emerging field of research in the coming years and we may therefore expect to see new methodologies and standards begin to emerge. It is likely that this will include an intersection of low-cost sensor networks with GIS data modelling and correlation of sensor data with remote infrared imaging.

Heat maps are hard to do right

The current inability to accurately record the specific microclimate deployment context of our urban heat sensors has implications for spatial interpolation of our data. Spatial interpolation refers to our ability to infer what the temperature might be at a location between two points with known measurements. It is the basis upon which we are able to construct an accurate heat map.

A heat map is arguably the most user-friendly visualisation of urban heat data. Users are familiar with heat maps as remote infrared imaging snapshots. These provide a comprehensive spatial map of heat but they only capture a single point in time. Fixed sensors provide an understanding of urban heat as a dynamic variable over time, allowing conditions to be explored in relation to other dynamic variables such as wind, sun and cloud cover. By adding a temporal dimension to our urban heat data we can also explore heat in terms of how it impacts people over time.

The trade-off for this view of temporal variation achieved using fixed sensors is a loss of spatial resolution. Where an IR snapshot tells us exactly what temperature any point on the map is, we are left to interpolate it when all we have is two sensor readings either side of that point. The difficulty lies in the fact that the temperature or humidity at our interpolated location is not a function of distance from the sensors to either side. These parameters are a product of the specific location, not the conditions that exist even tens of metres away.

Consider two sensors spaced fifty metres apart, each in an open concrete parking lot. Now imagine a location halfway between those two sensors that is a pocket of urban bushland; what is the temperature there? A standard approach to heat mapping would interpolate the average from the temperature readings of the two sensors and conclude that, on a hot day, the temperature in the bushland is roughly equal to the high temperature in the middle of the parking lots. However, we know that the temperature in that bushland might be as much as ten degrees cooler. We have no way of interpolating that localised difference from just using the sensor data that we collect. This is why it is hard to do heat maps right when we are dealing with distributed urban heat monitoring.

The only way we can start to get heat maps right is by introducing *additional* data. That additional data must capture aspects of the local geography. Remote imaging can do this, providing information about thermal radiation, albedo, vegetation height and species type. Topographic and building models can add information that allows modelling of air movements and shadowing. This sort of data may be combined with sensor data in GIS models and such an approach, that starts to work with many complex and varied data sets is likely the only pathway towards accurate heat mapping with a temporal dimension.

What data did we collect and how can we use it?

The project deployed a network of 52 urban heat sensors that measure temperature and humidity. Due to time constraints, initial data analysis was limited to temperature only and did not explore humidity. A more complete and in-depth exploration of urban heat data might explore both these variables as well as 'Apparent Temperature, which is derived from temperature, humidity and wind data.

Urban heat data collected at Melrose Park was reviewed for a summer and autumn period, namely: Summer period, from 2019-12-20 to 2020-01-17 (28 days), and Autumn period, from 2020-03-16 to 2020-04-30 (45 days). Due to staggered deployment of sensors, the summer period refers only to EMS devices 002-008 (7 in total). The autumn period refers only to Netvox devices (35 in total). Due to a systemic variation in observed temperature between the two device types, it was decided to avoid analytics where data from both was directly compared. It is believed that the readings from the EMS devices, which were on average around 1 degree warmer than Netvox readings, result from differences in the design of the sensor housing.

The Australian Bureau of Meteorology (BOM) has a monitoring station roughly 2.5km from Melrose Park, at Sydney Olympic Park. This station complies with international meteorological standards and the temperature data that it provides was used as a comparison for our own sensor data. When measuring urban microclimates, we are undertaking a fundamentally different exercise to what the BOM does with their weather stations. Where the BOM take measures to avoid localised variation and interference from thermal mass, shade and other factors, we explicitly seek to capture these variations. As a result, our data is 'messier' than the BOM's and it does not comply with meteorological standards, which are not appropriate for microclimate monitoring. Our sensors produce data that is highly contextual to their deployment location and mounting solution, and localised variation may result from a number of complex overlapping factors. As a result, many of the insights that have emerged are rough interpretations that support general hypotheses, and in most cases simply highlight areas for more focused future study.

The greatest limitation of the work to date is the limited data sets at Melrose Park. It was not until mid March 2020 that a majority of urban heat monitoring devices were deployed and actively accruing data, meaning that summer data was somewhat limited. In order to establish a reliable baseline for urban heat microclimates at Melrose Park it is necessary to gather at least 12 months of data, across all four seasons, and ideally at least 24 months to account for annual variation and anomalies.

A number of project sensors produced data sets with gaps and in a couple of more significant cases (e.g. NVX024), sensors were excluded from analysis because of this. As a rule, data gaps can be expected for sensors in locations where the LoRaWAN signal might be weak or marginal, perhaps due to topology or buildings blocking signal. During poor weather or high bushfire smoke events, such marginal signals can fail altogether, resulting in gaps. Sometimes, sensors closer to the LoRaWAN gateway (that would otherwise be expected to have a strong signal) can experience problems (e.g. EMS008). There can be oddities of the built environment and topology that create radio black spots, even in locations quite close to a gateway. Overall, issues with sensor connectivity and data gaps were a relatively minor concern for this project, though not insignificant. Certain locations proved to be unviable for sensor deployment and this did place limits on the study design. The larger issue was the creation of time-bound data gaps associated with poor weather and bushfire smoke. Longer term, the solution to this issue would be to install additional LoRaWAN gateways to provide greater coverage. This would be the single most effective step towards data gap reduction.

Insights from urban heat monitoring at Melrose Park

Urban heat data collected at Melrose Park was reviewed for a summer and autumn period. A summer period of 28 days, from mid December to mid January, refers only to EMS devices (6 in total) and coincided with drought, extreme heat and a severe bushfire season. An autumn period of 45 days, from mid March to the end of April, refers only to Netvox devices (35 in total). Due to a systemic variation in observed temperature between the two device types, it was decided to avoid analytics where data from both was directly compared. It is believed that the readings from the EMS devices, which were on average around 1 degree warmer than Netvox readings, result from differences in the design of the sensor housing.

1. We detected an urban heat island effect, resulting from localised elevation of microclimate temperatures

We found a tendency for our sensors to observe warmer average temperatures than those measured by the BoM station at Sydney Olympic Park. The average variation was up to 0.85°C warmer, but local events showed much greater difference. Further investigation into this systematic difference is required, but a likely cause is that our sensors are detecting radiated heat from surrounding objects, at dawn and through the early to mid-afternoon, whereas the BoM station is more shielded from such effects due to its tighter compliance with meteorological standards.

Image: Installing a Netvox R712 temperature and humidity sensor on site (credit: Andrew Tovey).

2. The urban heat island effect is intensified during 'hot' and 'very hot' days

The number of 'hot' (32°C+) and 'very hot' (41°C+) days recorded by one or more of our sensors in the summer study period exceeded those recorded by the BOM at Sydney Olympic Park (two to three times as many 'hot' days; two times as many 'very hot' days). Jan 4th was found to be the hottest day, with all of the locations hotter than the BOM reading. Peaks exceeded 50°C by 4pm, when all of the locations monitored at Melrose Park were hotter than the BOM reading by 5°C to 7.5°C.

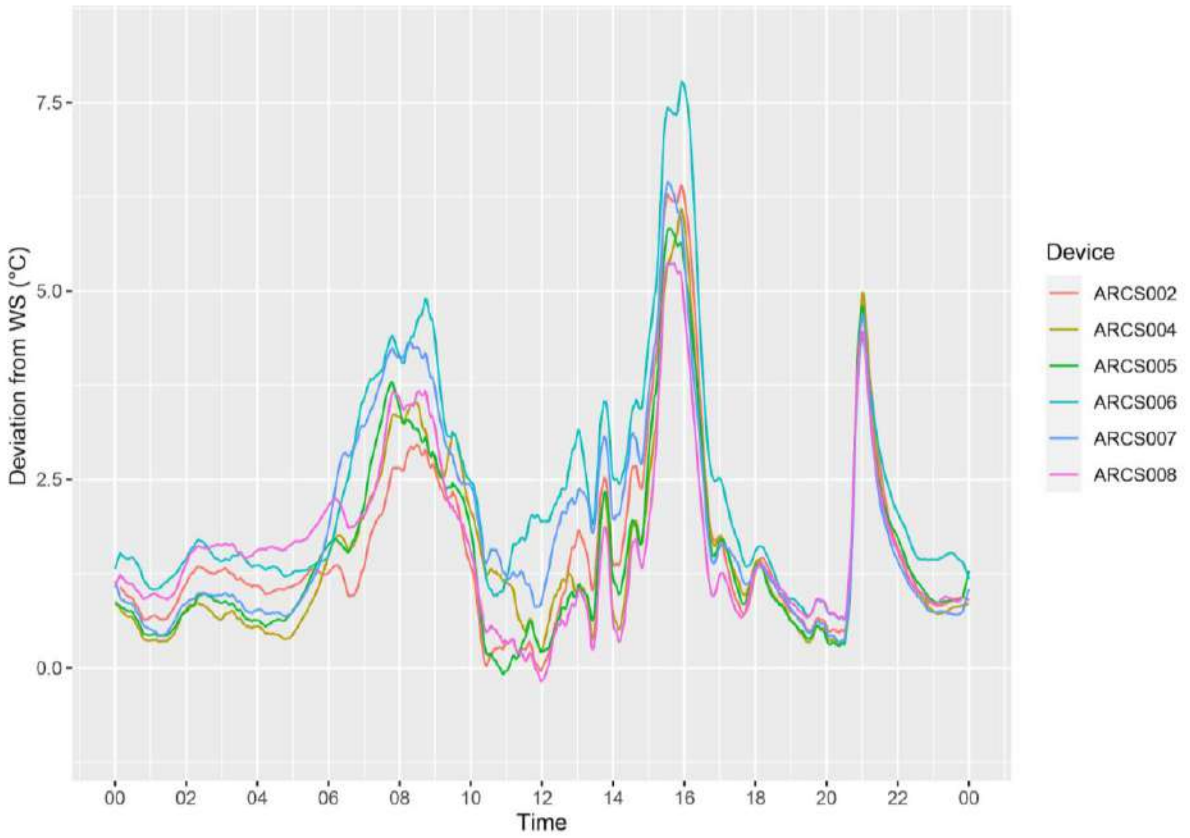

Image: A graph showing how each EMS sensor varies against the BOM on a "very hot" day (Jan 4th, 2020). There are clear peaks in the early morning (8am) and late afternoon (4pm) that can be interpreted as localised microclimate temperature elevations on the Melrose Park site

These insights suggest that we are measuring temperatures at Melrose Park that relate to highly localised microclimates. The BOM weather station is positioned to avoid any elevations in temperature that might result from the localised effects of the built environment, topography or vegetation. What our data shows is that the ambient temperature of the urban environment on hot and very hot days often far exceeds the 'official' temperature for the area, due to the localised interferences discussed.

3. The urban heat island effect is most evident after sundown

The urban heat island effect is strongest when thermal inertia comes into play at the end of the day. That is, solar heat which builds up in asphalt and concrete during the day is radiated as infrared in the hours after sunset. As a result, we found that ambient air temperatures in more built up locations (much of the development site and residential streets) remained warmer than locations characterised by high tree canopy or parkland, such as alongside the Parramatta River. It was found that on average, non-vegetated locations remain warmer than vegetated locations throughout the day, but that the difference between them is strongest just after sunset. This effect was most pronounced in the evenings of hot and very hot days (32°C+), when localised temperature variation makes built up areas on average 0.93°C to 1.14°C warmer than vegetated areas, with the temperature difference being up to 2.5 times greater after sunset than in the afternoon.

These figures relate to averages and indicate general trends, however specific paired examples show even more pronounced variation. The most shaded and vegetated location studied was Netvox 028, on the Ermington Bay cycle track. Early evening data from this location was compared to a number of nearby locations (<250m distant) with high amounts of concrete and low amounts of vegetation. The temperature difference between NVX028 and its neighbours on these hot evenings ranged from +1.6 to +2.1°C, indicating pronounced variation over short distances. The greatest variation observed between two sensors at a single point in time (2.4°C) was found between NV028 and NVX001 (the display suite car park), 900m to the north, marking the greatest difference detected between any pair of sensors at a single point in time. It is not clear if distance from the river played a role in this variation.

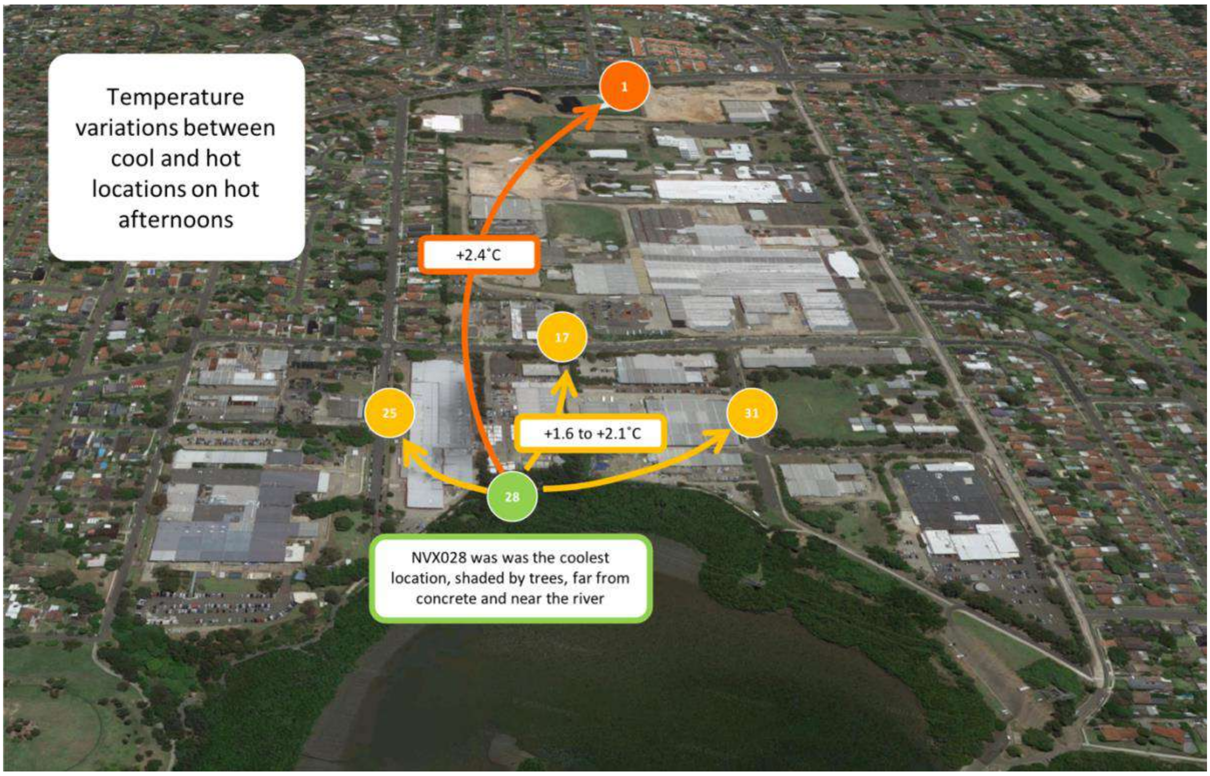

Image: Illustration showing the maximum variance in temperature between the coolest location (Netvox 028) and locations characterised by larger amounts of concrete and direct sun. The latter locations absorb more heat on a hot day and remain warmer in the early evening.

It should be noted that the greatest difference in temperature recorded between two of our sensors (2.4°C) is considerably lower than the 7.5°C difference recorded between one of our sensors and the BOM station on the 4th of January. The main reason for this is most likely due to the fact that the 7.5°C variation relates to a mid summer event, with peak temperatures above 50°C, whereas the 2.4°C variation is from a cooler autumn event. The effects of urban heat island are most pronounced during extreme heat. We only had six working sensors during the height of summer so we had fewer non-BOM locations to compare, hence our focus on autumn, when we had 32 devices to study. Additional mid-summer data would almost certainly have revealed a larger peak variation between two of our sensors than we were able to observe in autumn.

4. Shading noticeably reduces temperature during extreme heat events

Netvox 028, deployed along the Ermington Bay cycle track, is well shaded in the morning and throughout most of the day. As a result, it's location is cool in the early morning, warms up slowly over the day, and stays below temperatures recorded by the BOM weather station (by up to 3°C). Another sensor (NVX029) is 200m away. It is a comparable distance from the river and is surrounded by grass, with low thermal radiation interference, however it's immediate vicinity receives full sun for most of the day. The ambient temperature at the sunny location is closer to that recorded by the BOM station and notably warmer (average +1 °C) than the well-shaded location. This difference between the two locations peaks to over 2.5°C during a hot day. The greatest variation is in the early morning as the sun rises, and again at 4pm, when the afternoon sun has warmed the air and surrounding objects over several hours.

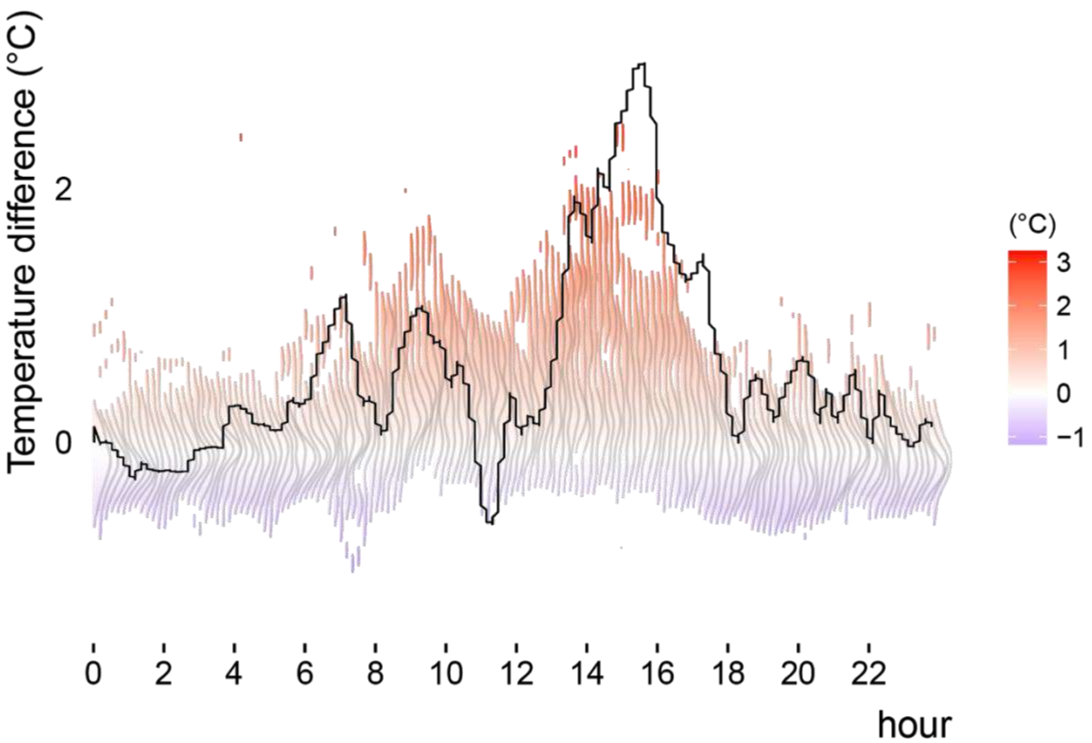

Image: Hourly temperature difference between NVX028 (full shade) and NVX029 (full sun). Both devices are in similar settings, the only significant difference being shading.

Future directions for urban heat monitoring at Melrose Park

The relatively dense deployment of fifty two urban heat sensors in a single square kilometre at Melrose Park is a unique opportunity for urban heat data collection in Australia, presenting a focused microcosm in a high density area due for major development and a significant rise in population over the coming decade. The network is building a valuable baseline of data that may prove invaluable as the area is transformed.

Extend our data sets for all devices across 12 months or more

This is critical to further in-depth study of urban heat at Melrose Park. Analysis to date has been limited to data sets for summer and autumn 2020. Temperature and humidity vary seasonally, meaning that 12 months is a minimum period needed in order to establish a baseline for annual heat variation. The Australian climate is also highly variable from year to year, meaning that data from just 12 months is also somewhat limited. 2019 to 2020 was notable for an extended period of drought in NSW. Ideally, two or more years' worth of data are needed to build up a more representative picture of how heat works at Melrose Park.

Co-locate our sensors in order to understand relative performance characteristics

The project used two device types for monitoring urban heat, the Netvox R712 and the TULIP EMS. It would be worthwhile co-locating Netvox R712 devices at three locations where an EMS is already deployed. Comparison of data from both devices showed a difference of about 1.5°C between the recorded temperatures of the two device types, most notably at higher temperatures. It seems likely that the EMS records higher temperatures due to thermal interference relating to the sensor housing. The Netvox R712 has a large Stevenson shield that we believe effectively minimises such interference, giving a more accurate reading.

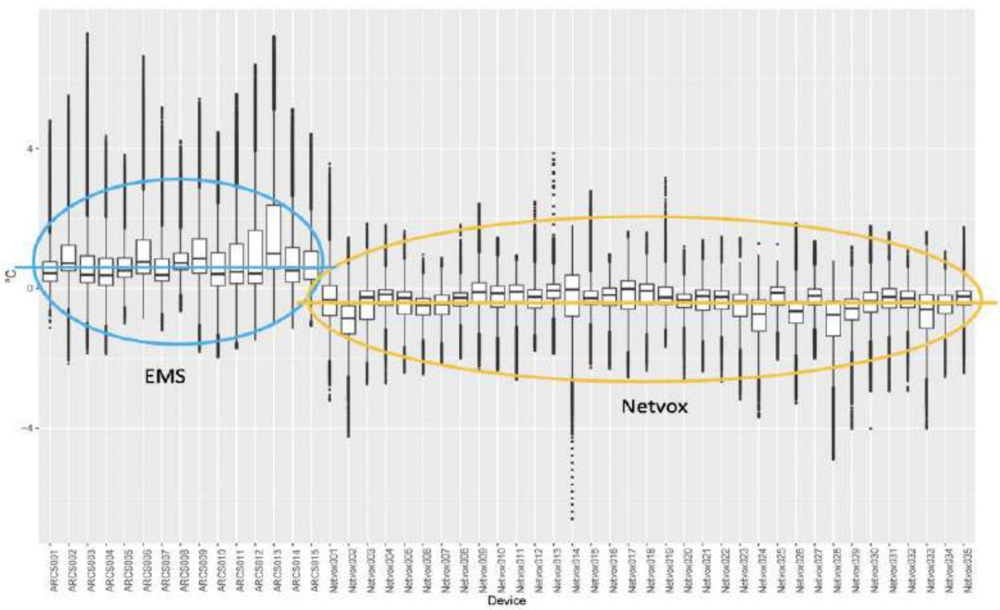

Image: Graph showing the average difference between each device and the average of all devices in the network. It is clear that the EMS consistently produces readings that are, on average, around 1.5°C higher than the Netvox

With no co-locations during the project we were not able to directly compare both devices in a single deployment context. If a few different co-locations were set up then it would allow comparison of the devices in multiple different contexts, for example, in full sun, partial shade and full shade. The reason why this would be useful is that it would establish a more accurate understanding of how each device performs relative to the other and in a variety of deployment contexts. Findings could then support a correction factor, allowing data sets from a network of mixed device types to be directly compared. This ability, to work with multiple device types that may exhibit small variations in performance, is likely critical to the scalability of distributed environmental monitoring. In the longer term, the most cost effective way of expanding sensor coverage in cities may be the ability to draw upon data from many different sources, which in term enables a crowd-sourced approach to data, or the 'Uberfication' of environmental monitoring. If we consider the position of a Council and a developer we might see a scenario where both entities have deployed sensors and share data with each other, both seeing mutual benefit. Citizens, schools and local businesses may also be sensor owners and may decide to share their data with a broader community of users in exchange for access to that same broader community of data sources. In such a scenario we might expect a variety of devices to used by different stakeholders and it would be necessary to understand how data of varying quality can be usefully collated and synthesised into a single coherent view of urban heat. While such a scenario is some way off, early exploration of how different device types perform relative to each other is one major area of research that will help to bring the scenario into reality.

Explore the impacts of pole type and pole orientation on temperature readings

Temperature sensors were deployed almost exclusively on poles, three metres above the ground. The type of pole differed, ranging from thin metal poles of around 6cm diameter, to thicker metal street poles of 15-20cm diameter, to very large wooden street poles of 30cm diameter or more. Pole type matters because a large metal pole can heat up in the sun and radiate that heat back out to warm the surrounding air. Thin poles do not create this effect and neither do wooden poles. Orientation of a sensor on a pole is also a consideration if the pole is thicker. A thick wooden pole creates shade that may reduce the temperature reading, meaning that if two identical sensors were placed on the northern and southern sides of such a pole, we might expect to see a variation between their readings. For the most part, all Netvox sensors were positioned on the southern side of their pole and all EMS sensors were placed on the northern side of their pole. This is because the Netvox devices do not require solar energy, so can sit in shade, which is likely to support a more accurate reading, whereas the EMS devices needed to have their solar panels pointed north for power optimisation.

We do not fully understand these types of subtlety in sensor deployment. However, if such differences can lead to reading variations of, say, half a degree, then that is significant to our discussion about the relative impacts of the built environment and factors like vegetation and tree canopy cover. We need to understand and rule out temperature variation that results from how we deploy sensors, so that we are left with only the variation resulting from where we deploy sensors. To examine this, we might set up two sensors on opposite sides of each type of pole, in locations that experience full sun throughout the day. Analysis of the results, with attention paid to very hot days, would allow us to understand this issue better and perhaps support measures for greater standardisation of device deployment methods.

Explore the impacts of other significant variables that are known to impact thermal microclimates, such as shading

Initial data analysis has shown a clear urban heat island effect evident at Melrose Park, with the strongest localised variation in temperature on hot days being linked to the presence of vegetation and trees, versus concrete and asphalt. However, comparison of various pairs of devices that are deployed in locations which might reasonably be characterised as vegetated and non-vegetated produced mixed results. Some pairs show a clear variation in temperature that is undoubtedly due to the amount of vegetation of concrete. With others, variation is far less evident, or only evident at certain times, with no clear pattern.

Thermal microclimates vary according to many factors other than vegetation and thermal mass. Our analysis to date has not explored these factors, in part because they are quite difficult to explore. Shading is likely to be the most significant and may even create localised effects that run counter to assumptions about the relevance of nearby vegetation or concrete. For example, location with lots of concrete which has buildings to the east, north or west, may have sufficient shading that it remains cool throughout the day and is comparable to a nearby tree-covered location. Time of year is also a factor as the angle of the sun creates seasonal variation in shading. To explore the impacts of shading, the hours of shade for a subset of sensors would need to be determined, to support a systematic study of the impacts. Other possible factors to explore would be the effects of aspect, recent rainfall and wind.

Explore the interplay with humidity and 'Apparent Temperature' and incorporate feedback from local residents about their experience of extreme heat in different locations around Melrose Park

Analysis so far has not explored humidity. Humidity and temperature work together to influence the human experience of heat. When we include wind speed, we can calculate 'Apparent Temperature' or what is referred to as the 'feels like' temperature. High humidity at 35°C feels a lot more uncomfortable than low humidity at 35°C. We know that humidity can vary over short distances due to microclimate effects, just like temperature. The combination of both factors can make the experience of localised hot and cool spots feel more pronounced. The way that people experience heat can be explored through social research and it may be interesting to combine sensor data with people's perceptions of heat in different locations during extreme heat events.

Develop an accurate heat map

One clear future direction for urban heat monitoring at Melrose Park is to explore the challenge of accurate spatial interpolation and heat mapping by way of GIS modelling. The aim is to develop an accurate, temporally dynamic, real-time heat map of the area by combining live sensor data with other existing data sets. Such work would constitute a significant standalone study. The benefits may be quite valuable for understanding the complexities of urban heat microclimates, with specific application to the future design and construction of major developments like Melrose Park. With increased pressures from climate change combined with growing urban density, such insights and new tools may become critical contributors to more liveable Australian cities.

Use case 2

Air Quality

Use case 2: Air Quality

- Understand dust creation on site, how and where it is generated, and how it may impact surrounding areas, particularly residents. Insights should support practical mitigation measures.
- Build baseline information about air quality in the area to understand what constitutes elevated levels and to compare future data sets to.
- Understand the impact of Victoria Road as a pollution source, allowing us to distinguish between traffic pollution and dust from site.

Understanding particulate pollution at Melrose Park

Particulate pollution has a significant impact on human health and wellbeing and is a growing problem in our cities. Vehicle emissions are a major and much talked about source, however construction activities are increasingly recognised as a significant contributor too. On a 30 hectare brownfield site like Melrose Park, dust may arise from demolition activity, earth moving, heavy vehicle movements on unsealed site roads, and cutting of concrete. Creation of dust, the degree to which it becomes airborne, and its dispersal to adjacent areas, is dependent upon a combination of on-site activities, mitigation measures and weather conditions. Dry weather and windier conditions are most associated with the generation and dispersal of dust. Likewise, rain has a cleansing effect on air quality, washing dust from the atmosphere and reducing new dust formation on the ground.

The primary means of mitigation for construction site dust is informed planning of activities that are known to generate dust so as to avoid them coinciding with weather conditions most associated with dust generation and dispersal. It can also include the spraying of unsealed roads with water to reduce the amount of dust becoming airborne. Management of such mitigation measures may be improved with the addition of real-time particulate pollution data from the site and the surrounding area.

Image: An air quality monitoring device deployed on the development site (credit: Andrew Tovey).

Particulate monitoring study design

The focus for air quality monitoring at Melrose Park was on particulates of 2.5 and 10 microns in size. This includes fine dust particles that might be associated with on-site construction activity as well as particulate pollution from vehicle exhaust, and smoke from bushfires. Smart low-cost air quality monitoring devices for air quality monitoring range in cost between $2500 and $10,000, or roughly 'Tier III' of the US EPA's categorisation system[6].

For the project at Melrose Park a device called the TULIP Environmental Monitoring System (EMS) was used. The EMS was jointly developed by UTS in collaboration with industry Partner The ARCS Group through a collaboration supported by an earlier round of the Australian Government's Smart Cities and Suburbs Program. The EMS monitors PM2.5 and PM10 as well as a range of noxious gases. Solar power and low power optimisation, together with a compact design, allows deployment of the EMS in a variety of locations with no need for mains power. LoRaWAN communications provide a 15 minute reporting interval without draining limited battery reserves.

Fifteen EMS devices were deployed at Melrose Park, with eight on site and seven in the surrounding streets. All devices were deployed at 3 metres on free-standing poles. The average distance between two EMS devices within the study area was 175m, with a non-uniform clustering around locations of higher interest. Onsite deployments were clustered to capture areas of known construction activity and construction vehicle access in the north and south of the site. Devices were also spaced around the perimeter of the site to capture air quality at the boundary between the site and the surrounding area. Off-site deployments aimed to do two things. Two devices on either side of Victoria Road were deployed to gain an understanding of the profile of traffic pollution from that major road and how it might impact the wider air quality of the Melrose Park area. The other devices were deployed in residential streets around the east, south and west of the development site, at varying distances from active construction. These were deployed to provide an understanding of the baseline air quality in these locations as well as to test if detectable spikes in particulates at off-site locations might be correlated to on-site spikes and known dust creation activities, or to other known events, including rush hour traffic and bushfire smoke. Detection of the latter relationships would be heavily reliant upon wind direction and speed.

[6]Watkins, T., 2013. Draft roadmap for next generation air monitoring. *Environmental Protection Agency*, *2*. [online] US EPA. Available at: https://www.epa.gov/sites/production/files/2014-09/documents/roadmap-20130308.pdf

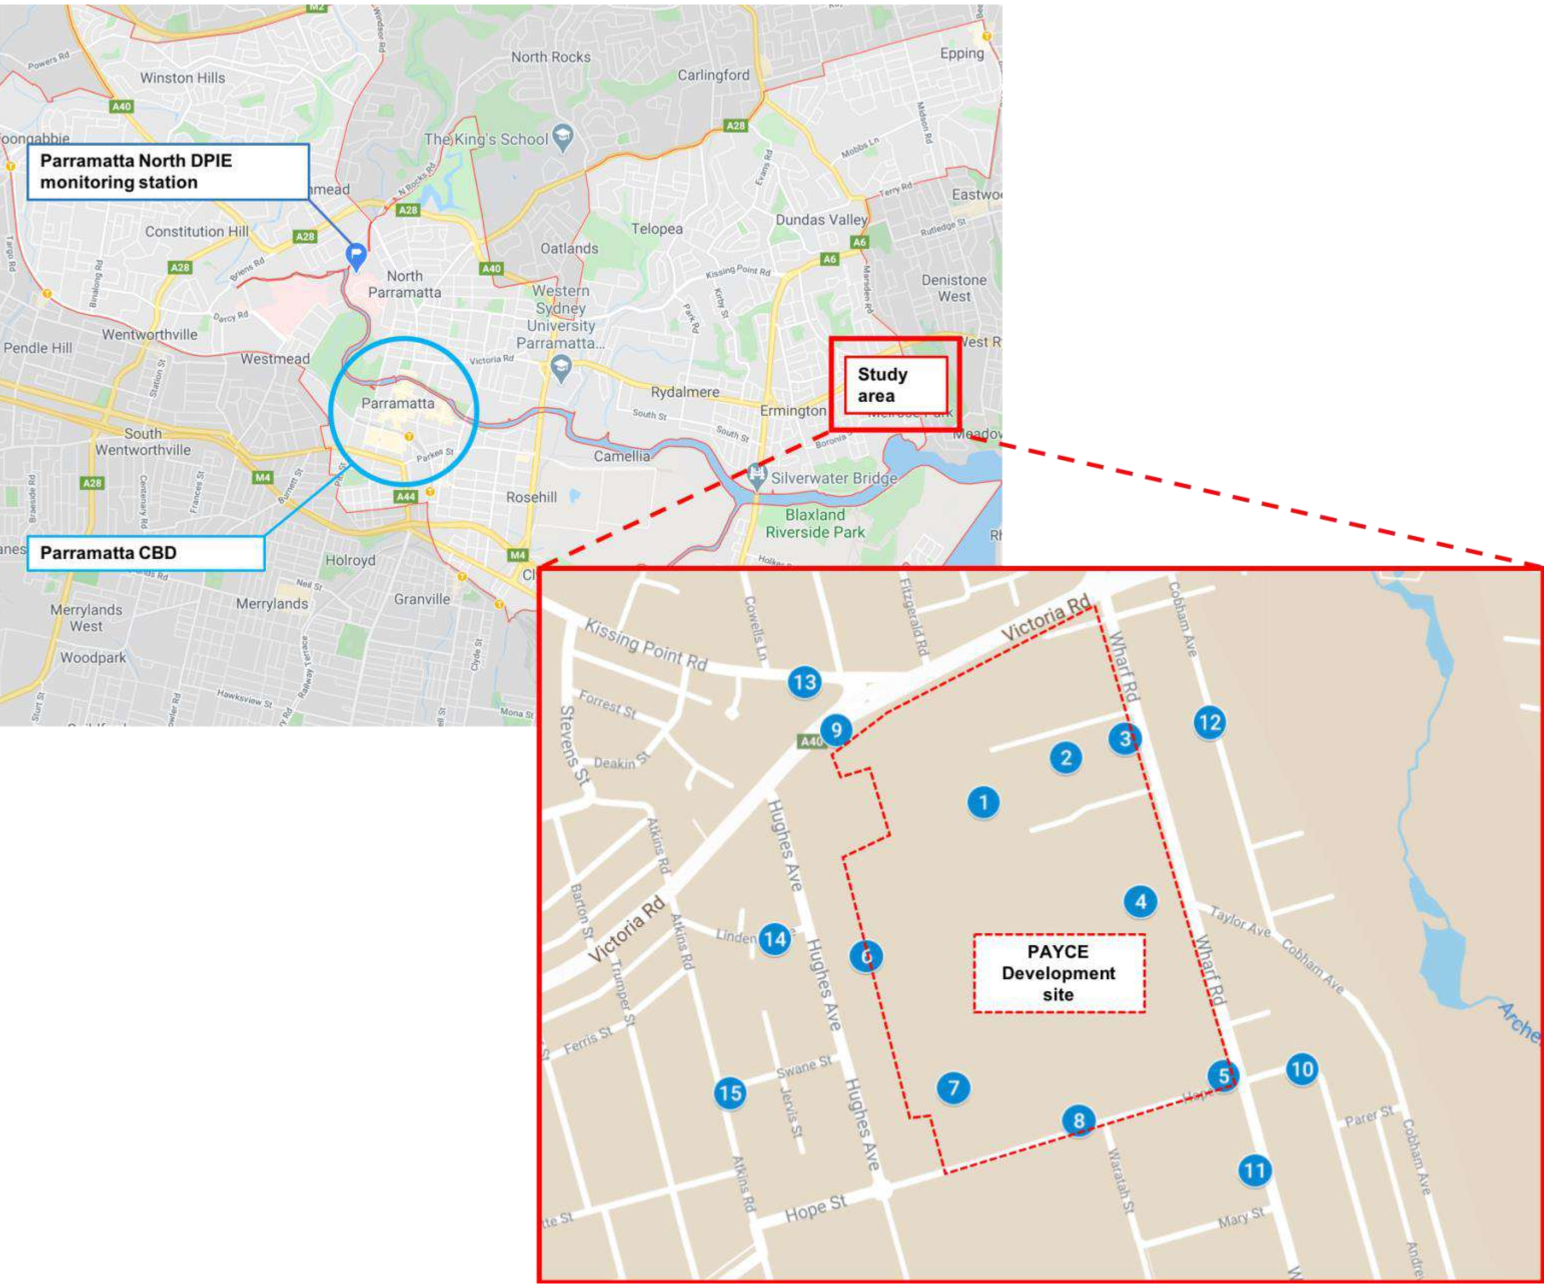

Image: *A map showing EMS deployments relative to the City of Parramatta LGA. Note that the same deployment of 15 EMS devices was also used for noise and urban heat monitoring.*

Constraints of the particulate monitoring study

Spatial resolution

The spatial resolution of the air quality sensor network was perhaps the greatest constraint. Devices were positioned, on average, 175m apart. In designing the device deployment strategy, a balance was sought between clustering of devices around areas of known activity (e.g. Hope St, the northern construction zone), and comprehensive coverage of the site perimeter and surrounding streets. It is clear that if devices were placed closer together it would be possible to find stronger correlation between their data. This would allow us to far more clearly discern the presence and nature of localised dust creation events. Any attempts to increase the future density of device deployment would yield clearer understanding of dust creation and dispersal on and around the site.

Take for example the northern construction zone. Three EMS devices captured the southern boundary of this area, however their data was often inconclusive. Direct observation of the site over many months has confirmed that dust creation and dispersal can be highly localised, and visible clouds of dust can pass clearly between sensor locations, never being detected. An additional four or five devices in this area would have increased our ability to detect such localised ‘micro-events’.

Image: Northern construction zone, June 2019. A cloud of localised dust is visible in the sunlight, a result of earth moving. Such micro-events are difficult to capture with a low sensor distribution (credit: Andrew Tovey).

The number and density of EMS deployments in this study was a result of:

a) Limited budget (and in turn, the current cost of devices)

b) A need to initially deploy for broad coverage of the site and surrounding area in order to establish a general noise baseline and to identify locations of interest for further study.

c) Practical limits on the locations where it was possible to place a device. For example, it was initially the aim for a device to be located at the Wharf Road heavy vehicle access gate to capture the noise of passing construction traffic. This was not possible during the project period because that gate was not operational and there was no clarity about suitable mounting infrastructure at that location until after the new roadway was completed. On a construction site it is often difficult to identify fixed semi-permanent infrastructure, which is needed in order to mount a sensor for any useful period of time. Furthermore, very active areas were more difficult for researchers to safely access on a regular basis. Thus, many locations that may have proven strategically useful to deploy a sensor were not used.

Temporal resolution

The temporal resolution of the air quality sensor network was another constraint. This refers to the fifteen minute reporting interval of the TULIP EMS. Generally speaking, air quality monitoring involves taking regular samples of air to provide an indication of background pollution. In most urban settings, where air pollution relates principally to traffic fumes, industry and bushfire smoke, concern relates to ambient air quality over hours. Changes in ambient air quality occur over timeframes where 15 minute intervals capture major trends. Very short duration events of elevated pollution are not captured and are less of a concern. If we consider dust creation on a construction site, it may occur in highly localised non-uniform clouds relating to specific disturbances. These ‘micro-events’ may pass a sensor over seconds or minutes and thus may not be detected by a fifteen minute sample period, even if the dust cloud completely engulfs the device.

One way to address this challenge is to increase the sampling interval of the sensor. This would likely involve near continuous running of the nephelometer (particulate sensor), with a corresponding increase in power demand. The sensor uses a heated element and on a fifteen minute cycle this is allowed to cool and then has to be reheated. By shortening the cycle to, say, under a minute, it is likely that the sensor would need to run continuously. The only practical way to accommodate this would be to run the device using mains power. While this would be possible, it would significantly reduce the range of deployment locations available, thus significantly impacting spatial resolution. For this reason there is almost always a trade-off between increases in spatial and temporal resolution.

The other way to approach the challenge of limited temporal (and indeed spatial) resolution, is to focus more on ambient dust levels around the site, rather than on trying to capture these ‘micro-events’. Under dry conditions with high site activity, the ambient dust levels in the air around the site will be elevated. This particulate load may not be visible in the way that ‘micro-events’ are but it will be fairly evenly spread, and trends are likely to fluctuate over longer periods that can be more readily detected by sensors operating on fifteen minute cycles. It is also worth noting that if a higher deployment density of devices is attained in a relatively small area then this effectively increases the temporal resolution for that area because any two devices do not synchronise their cycles. If we are merely concerned with the ambient dust levels in the vicinity then having five devices in that vicinity cuts the sampling rate from fifteen minutes to more like three minutes. Obviously spacing and wind direction would come into play but the key point is that the more devices you have in a distributed network, the greater the overall resolution, not just spatially, but temporally too.

What data did we collect and how can we use it?

The project produced an initial assessment of highly localised particulate matter impacts on air quality associated with construction activity and vehicle emissions at Melrose Park. Data from a network of fifteen 'low-cost' monitoring devices was available for a period from 18th Nov 2019 to 15th Jan 2020. Eight devices were deployed on site and another seven were deployed in surrounding streets.

High levels of site activity during a dry period may contribute to a more evenly distributed elevation of dust levels across the site, however the majority of dust 'events' can be expected to be very localised, associated with a specific vehicle movement or other activity. Our ability to detect such highly localised events is therefore linked to the density of our sensor network. With only 8 on-site sensors, this density was somewhat limited and despite some amount of clustering of devices around areas of active construction, we believe that much of the highly localised dust being generated during the study period was not detected by our sensors. We recognise that the project was a pilot and that future development of construction dust monitoring would likely demand investment in a higher density of sensors.

Construction site activities, and indeed the layout of the site itself, change on a near weekly basis. This means that the location, timing and significance of dust generating activities changes all the time, with few fixed trends evident beyond a day/night week/weekend variance. This made it difficult to clearly link data to known activities. Ideally you would have a period of months to establish the link between elevated dust levels and one known activity, in one place, at a recurring time. Instead, we found we were only able to capture one-off events before the situation changed. Furthermore, in the absence of reliable activity records, it did not prove possible to identify known activities and link them to dust data, beyond a high level knowledge that construction was under way in a particular site quadrant in a given month. In reality, much of our dust data needed to be interpreted in terms of one-off events that implied localised activity but which could not be reliably ascribed to it. A great deal of contextual interpretation was often needed, with consideration of multiple overlapping factors and a working knowledge of the site for any given event.

Data gaps were minimal for the study period, though data sets were limited due to late deployment of around half of the EMs devices. It is worth noting that since the data analytics was undertaken, many of the EMS devices have experienced intermittent connectivity, resulting in substantial data gaps. This is believed to be a result of low levels of sunlight in winter, combined with overcast weather, leading to power outages.

Insights from particulate monitoring at Melrose Park

1. Our sensor data tended to agree with data from nearby state-run (DPIE) monitoring stations

On average, the measurement of particle pollution with low-cost sensors at Melrose Park fits relatively well with the data collected by the regulatory instruments of the NSW DPIE monitoring station at Parramatta North, supporting the use of the low-cost EMS sensing system for air quality monitoring in municipal suburbs over time.

2. Our chosen device was appropriate for the job

The EMS provided relatively accurate measures across a wide range of readings, indicating that it is suitable for accurately monitoring very good to very poor air quality conditions. This assessment was supported by the occurrence of heavy bushfire smoke events during the study period and our ability to closely fit EMS data with DPIE data across large ranges. This supports the idea of low-cost distributed sensing networks as a viable and affordable approach for monitoring urban air quality.

3. Construction activity produced detectable dust (but we are not clear on specific causes)

There is evidence that construction activity did contribute to emissions of particulate matter (PM2.5 and PM10), especially in the contribution of fine PM2.5 particles. Data showed consistently varying particulate levels across the development site, implying localised dust generation. However, further study is needed to rule out the influence of vehicle traffic emissions, and the correlation of particulate levels with actual observed construction events over a longer period.

On days where particulate pollution may be attributable to construction activity on site, levels appear to fall within national standards. However, this assessment is based on average levels for the day and does not account for highly localised short-lived peaks in particulates. Further study is needed into highly localised short-lived peaks in our data sets to determine whether breaches of standards are occurring due to possible site activity.

4. Air quality was worse during the day

There is a diurnal variation in air quality consistent with poorer air quality during the day and more activities in the area and on site. It is not clear whether this relates to local or regional variations or whether it relates to construction activity, vehicle traffic or a combination of overlapping factors. Further study and larger data sets are needed to explore this trend in more detail

5. Distributed low-cost sensor networks support internal cross-validation of data

For the purpose of using low-cost air quality sensors for local and regional scale monitoring, we have shown that deployment of multiple devices in close proximity, as part of a dense network, can allow for cross-validation of data between adjacent devices or between any one device and the wider local network. We have not explored whether similar approaches might be used to cross-validate microclimate variations associated with data from a single device, or how to validate such highly localised variation in the context of local and regional scale validation. This would be an area for future research.

6. Rain improves air quality but its role in mitigating dust creation requires further study

The effect of rain on air quality was demonstrated during the bushfire period when particulate levels dropped dramatically, coincident with precipitation.

Analysis has not yet explored the link between rain and construction dust. It is hypothesised that there is likely to be a threshold for accumulated rainfall across the previous week that would dramatically reduce dust creation if exceeded, thus it should be possible to position rainfall as a critical factor for calculating the risk of onsite dust creation. It should also be possible to understand how rainfall impacts elevated dust levels on and around the site, particularly the speed and strength of mitigation afforded by rainfall of different intensities. This represents an area of recommended further study.

7. The role of wind in the generation and dispersal of construction dust is inconclusive, and represents a clear area for further enquiry

Analysis of wind speed data from a two day period in mid December indicates that there may be a rough correlation between wind speed and macro regional trends in air quality that are common to average readings at Melrose Park and the North Parramatta DPIE monitoring station. Higher wind speeds during that period match higher particulate levels. Given that the period occurred during the height of the 2019/2020 bushfire season and biomass smoke was causing widespread air quality issues across Sydney, it seems likely that windy conditions were blowing smoke in from the extensive fire grounds to the north west, resulting in the elevated particulate levels.

No analysis has yet been undertaken to explore the relationship between localised events of elevated particulates at Melrose and local wind conditions. It is hypothesised that very low wind speed combined with dry weather may result in highly localised dust elevation on the construction site that may be more broadly detectable by our sensor network due to the fact that particulate laden air is likely to remain in the general vicinity. It is also thought that low to moderate wind speed (e.g. a light breeze) combined with dry weather may generate airborne dust in areas of the site where the ground is disturbed, such as exposed earth moving or on unsealed site roads. It is not known whether still air or light breezes might result in greater dust generation. Furthermore, it is hypothesised that more moderate to high wind speeds are associated with increased dispersion of site dust to off-site locations.

The role of wind in the formation and dispersal of airborne dust from the construction site is an area that is recommended for future study. One aim might be to identify if there are critical thresholds for wind speed and direction where risk of dust creation and dispersal increase or decrease. If such thresholds are evident they will be dependent upon additional factors such as amount of rain in recent days, temperature and soil moisture.

Future directions for particulate monitoring at Melrose Park

Increase the deployment density of devices

Increase the density of deployed devices around known areas of probable dust creation. These include current and future heavy vehicle access gates and current and future areas of heavy earth works. A greater density of device deployments allows us to capture more highly localised dust creation events that might otherwise pass between existing sensors. It also supports the cross-validation of data between adjacent devices which is critical for building confidence in micro-climate data.

Explore the role of wind in dust creation and dispersal

Initial analysis of Melrose Park particulate data did not include a deeper exploration of on-site wind data. Wind is the principal dispersion factor for dust, in terms of how far away dust impacts off-site locations and in what direction. Wind also plays a significant role in dust creation; even a light breeze on a dry day can pick up significant dust from unsealed site roads. Existing EMS deployments continue to gather data and a follow-up study of data sets from all devices should be undertaken with a specific focus on wind relationships. Analysis should explore data from outlying EMS devices in residential streets around the site to understand how wind speed and direction impact measurable dust levels.

Integration of video or image capture of dust sources for on-site validation

If it were possible to integrate video or photographic surveillance technology into the network in order to capture a visual record of dust creation events on the construction site (for example, active earth moving or heavy vehicle movement), this may significantly augment the data captured by the EMS devices and provide a strong validation for patterns detected in the air quality data. This is only relevant to dust creation analysis, hence cameras would only need to be deployed at onsite locations, not in residential streets.

Focused study of dust in the northern construction zone

Focus attention on the northern construction zone with the aim of building a deeper understanding of how weather conditions and ground conditions, plus measurable activity, relate to ambient dust levels. Insights from a focused location can then be used to start developing a dust risk model for the site, which might be applied to other construction zones as they open up. Initially, this work might make use of the existing EMS devices in the northern zone (EMS001, 002 and 003), as well as weather station data. The relatively underutilised nCounter people counting technologies deployed through the project should also be explored in greater detail in this context, as indicators of onsite activity. An expansion of deployed devices should include additional EMS devices in the northern construction zone, specifically along the northern zone perimeter as well as on the new heavy vehicle access on Wharf Road. An upcoming new site access off Victoria Road (late 2020) should also be covered.

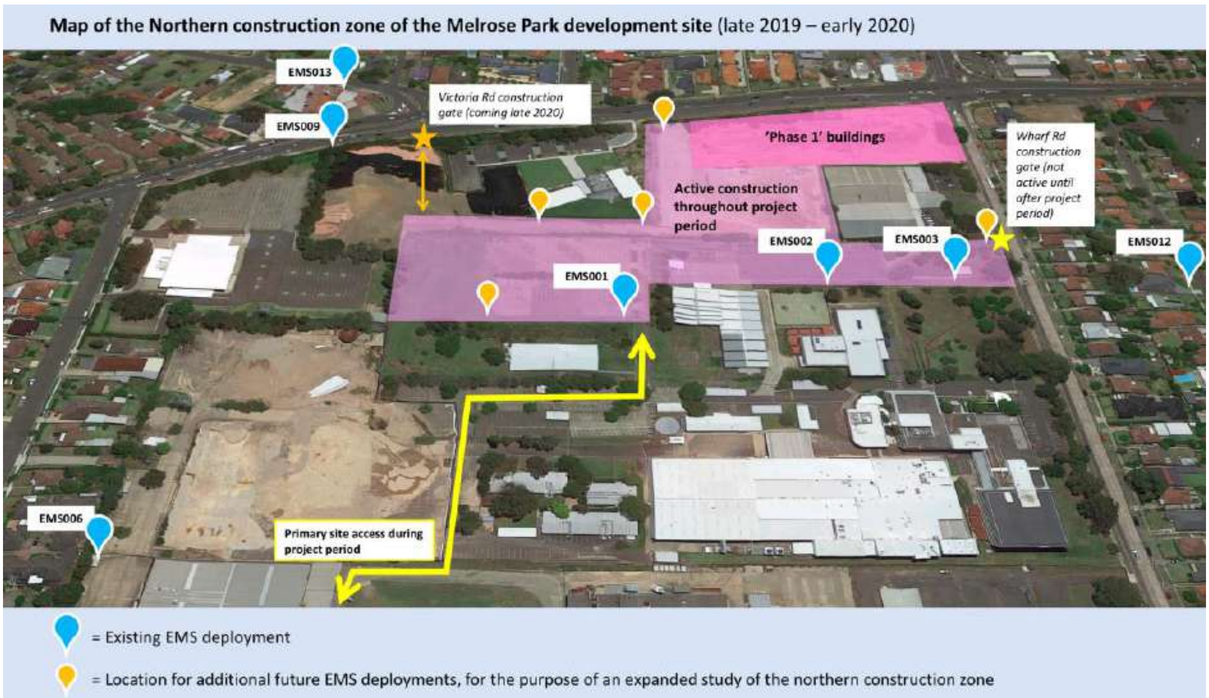

Image: Map of the northern construction zone showing existing EMS deployments, plus potential future deployments for an expanded study.

Dispersion modelling through use of transects

Deploy denser linear transects of devices spreading outwards from the site. Wind direction data from the Lufft weather station shows a general day/night trend of easterly wind during the day, switching to westerly wind at night. Daytime winds coincide with drier conditions and onsite activity and thus we might expect, as a general rule of thumb, that areas to the west of site have a higher chance of being impacted by dust than areas to the east. Therefore, transects designed to explore dust dispersion away from site should consider EMS deployments on the western edge of the RB Remediation site, Hughes Avenue, northern Atkins Road, north Jervis Street, and potentially Trumper Street. An additional five or six devices spread across these locations would support a significant extended study of off-site dust dispersion. Any such study should be preceded by a general wind/dust correlation analysis, outlined above.

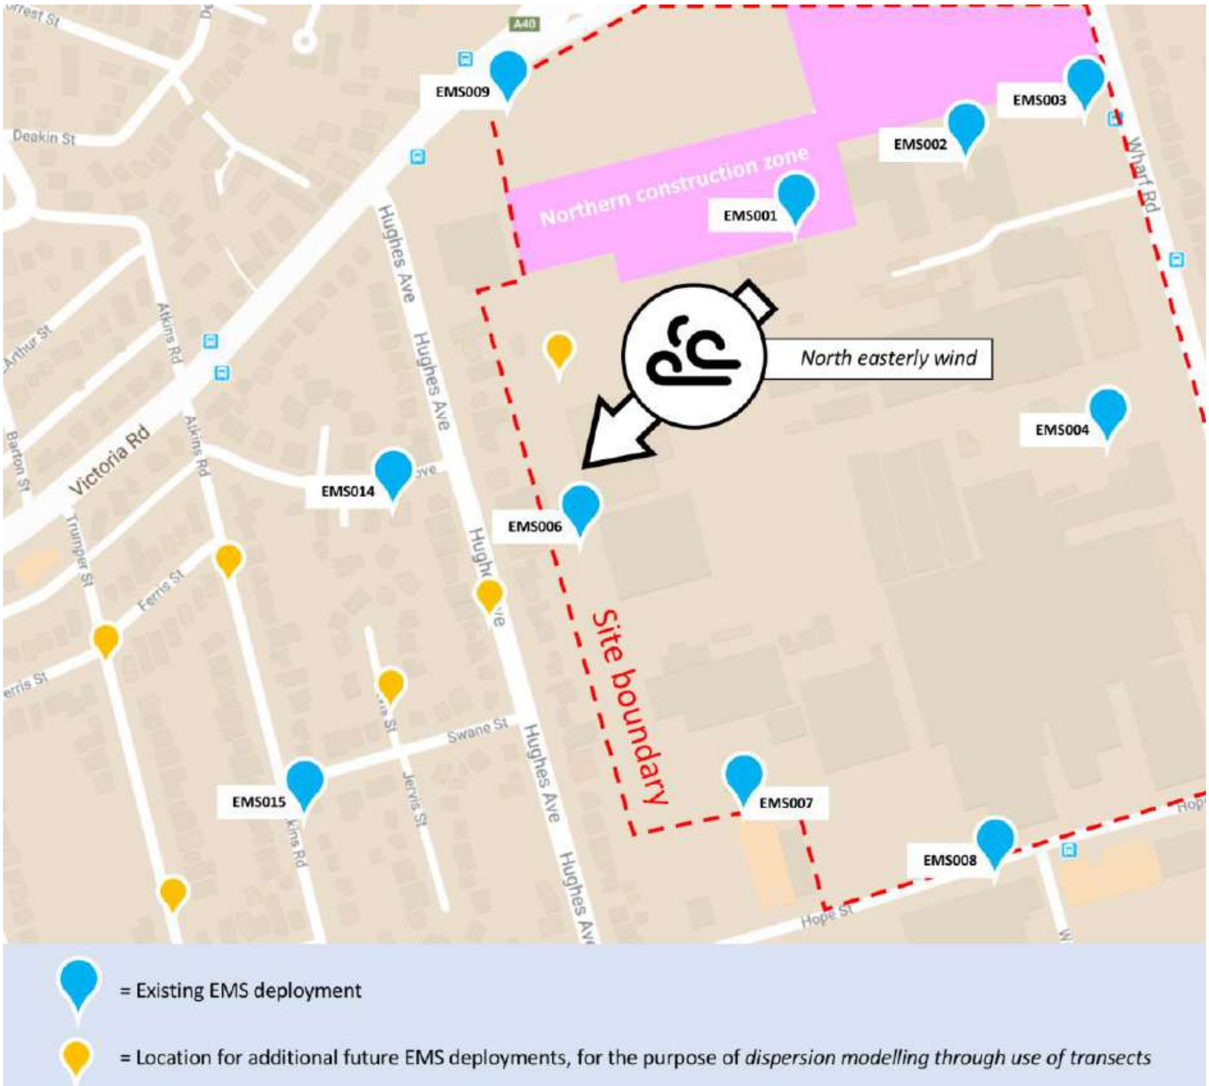

Image: Map showing possible future EMS deployments to establish dispersion modelling transect

Use case

3

Noise

Use case 3: Noise

- Understand noise creation on site and how this may impact surrounding areas, particularly residents.
- Build baseline information about noise in the area in order to understand what constitutes elevated levels and to compare future data sets to.
- Understand the impact of Victoria Road as a noise source, allowing us to distinguish between traffic noise and noise from site.

Understanding noise at Melrose Park

Noise levels in cities are a concern primarily due to their impact upon human amenity and wellbeing. Consistent high background noise has been extensively linked to elevated stress levels and in turn, a range of health concerns, including increased risk of cardiovascular disease. At the simplest level, quiet places are just nicer places to live. A major construction site like Melrose Park will see active work for up to a decade. From demolition to earth moving and heavy construction vehicles, development activities can generate noise well above what might be considered acceptable background levels. Such noise tends to be intermittent, and is often quite localised. With a construction site, concern rests around how much noise is created (amplitude), the frequency of noise events, and where they occur relative to people who might be impacted. The aim is for construction noise to be kept within certain hours, and for attention to be paid to how and where noise is created so as to minimise its impact on residents. The sensor network deployed at Melrose Park sought to identify when, where and how loud construction noise is at Melrose Park.

It was hypothesised that the two most likely sources of noise associated with early stage construction would be the movement of heavy vehicles to and from site (including along adjacent public roads such as Wharf Road), and earth moving activity at specific locations within the construction site (e.g. the northern construction zone). Sensors were deployed to record noise at these locations, with others positioned in what were likely to be quieter locations. Other sensors were also placed in residential streets to explore if noise associated with the site was travelling into the community at significant levels.

Noise in a city environment comes from a number of places, some of which are artificial and some of which are natural. Traffic from large roads can be a major source of noise and this was expected at Melrose Park. The greatest natural contributor to noise is wind, as it passes trees and objects in the built environment. Other natural sources of noise include rain and thunder. When we monitor noise at Melrose Park we have to distinguish highly localised noise events that relate to site activity from broader effects of weather, and from traffic.

Image: Heavy construction vehicles are a major source of noise on the site (credit: Andrew Tovey).

Noise monitoring study design

Noise was measured as both a peak and an average value across fifteen minute intervals, in a-weighted decibels (dBa). This was done by the TULIP Environmental Monitoring System (EMS), which was also deployed in the project to monitor air quality. The EMS measures sound by sampling noise every two seconds. Peak noise values capture the loudest single noise that occurred within the fifteen minute reporting period. The average noise value is the average of all the ~450 samples (two seconds apart) taken over the 15 minute sample period. This provides a sense of the baseline noise level across that period.

Noise was monitored by the same fifteen EMS devices that were deployed for air quality monitoring. We have already discussed the limitations and trade-offs involved with deployment locations of the EMS and it is worth noting that another constraint was that we used a single device type to service two different use cases (air quality monitoring and noise monitoring). Fortunately noise and air pollution tend to originate from some of the same sources, notably vehicles, be they construction vehicles on site or off-site traffic, meaning that it made sense to monitor both from the same locations. There was some degree of divergence however. The noise data collected has indicated that detectable noise that is most likely associated with the site relates to (likely heavy) construction vehicles passing along access roads. With this hypothesised when deployments were planned, some EMS units were placed along Wharf Road and Hope Street to capture noise, with possible traffic pollution as a secondary point of interest. A number of EMS devices were located in residential streets away from the site, to establish a baseline noise profile in these locations with the aim of building an understanding of how noise from a variety of sources might impact residents.

Image: Map showing distribution of noise sensors on and around the development site

Constraints of the noise monitoring study

Spatial resolution

As with air quality, the spatial resolution of the noise sensor network was perhaps the greatest constraint, with both noise and particulates measured using the same fifteen EMS devices positioned on average, 175m apart. With noise and dust creation sources expected to be in the same locations, the same broad deployment strategy was sought for both use cases, with the same trade-off between broad coverage and focused clustering. It is clear that if devices were placed more closely together it would be possible to find stronger correlation between their data. This would allow us to far more clearly discern the presence and nature of localised noise events, such as the passing of heavy vehicles on a specific route.

Furthermore, where on-site noise is occurring at a specific location, a denser network of on-site sensors would also enable us to roughly triangulate a likely zone of occurrence. The current density of devices does not allow this to be done. Any attempts to increase the future density of device deployment would yield clearer understanding of noise on and around the site.

The number and density of EMS deployments in this study was a result of:

a) Limited budget (and in turn, the current cost of devices)

b) A need to initially deploy for broad coverage of the site and surrounding area in order to establish a general noise baseline and to identify locations of interest for further study.

c) Practical limits on the locations where it was possible to place a device. For example, it was initially the aim for a device to be located at the Wharf Road heavy vehicle access gate to capture the noise of passing construction traffic. This was not possible during the project period because that gate was not operational and there was no clarity about suitable mounting infrastructure at that location until after the new roadway was completed. On a construction site it is often difficult to identify fixed semi-permanent infrastructure, which is needed in order to mount a sensor for any useful period of time. Furthermore, very active areas were more difficult for researchers to safely access on a regular basis. Thus, many locations that may have proven strategically useful to deploy a sensor were not used.

Temporal resolution

The temporal resolution of the noise sensor network was another constraint. This refers to the sampling rate of noise by the device, which in the EMS is once every couple of seconds. If a noise occurs between two noise samples captured by the EMS then it won't register, either as a peak or as part of an average noise calculation. When monitoring passing vehicles, a couple of seconds also makes a big difference to the detected amplitude of noise by the sensor. The sensor is likely to miss capturing noise when the vehicle is closest and therefore loudest. By reducing the sampling rate we would more accurately capture peak noise of passing vehicles, providing more accurate peak noise readings as well as a larger set of samples to draw average noise readings from. The limitation on sampling rate is the processing power of the device. The EMS has been optimised for low power consumption which places limits on its processing power and hence on its maximum possible sampling rate. We believe that for most of the project needs, the existing sampling rate suffices, however a discussion of noise sampling rate and how it might be increased was one of the outcomes of data analysis for the project. It may be worth engaging with The ARCS Group (manufacturer of the EMS) about ways in which the device's sampling rate might be slightly increased without compromising its low-power optimisation.

What data did we collect and how can we use it?

The project deployed fifteen noise sensors to measure peak and average noise levels around known areas of activity, notably the main site gate and approaches, and the primary areas of construction activity. Others captured traffic noise on Victoria Road and noise in residential streets at various distances from the development site.

Study of the data collected so far has shown that the EMS is an effective device for monitoring average and peak noise levels in the urban environment. There are some limitations to its performance that result from the rate at which it is able to sample noise, in turn a product of its core processing capacity. However, it has enabled a number of useful conclusions to be drawn about noise levels at Melrose Park

The density of deployed sensors is high by usual standards, in that any deployment of fifteen sensors within a single square kilometre might be considered to be dense with respect to standard noise monitoring approaches. However, in terms of understanding and pinpointing highly localised noise sources, our deployment can be understood as fairly low density, with sensors positioned, on average, 175m apart. Despite this spacing, we were able to learn a fair bit about localised noise generation, because noise propagates outwards from a source in all directions, providing there is no physical barrier to block it. As such, it acts quite differently to something like air pollution, which is far more dependent upon wind and air movement and tends to disperse in just one direction from a source. We did still run into sensor deployment density issues however, and this was one of the findings of the data analytics we undertook. Our existing network was enough to identify probable hotspots (e.g. around the Hope Street site gate), however we lacked the density of deployed devices needed to confirm precisely what was occurring.

Unlike other environmental parameters measured for the project, there is no external reference for noise against which we might validate our data. A degree of cross-validation between multiple devices was attempted but this was hindered by the low deployment density of devices.

As with other use cases, the short period of the noise data set was a significant limitation, though perhaps less so than with a seasonal variable such as urban heat. Larger data sets would certainly help deepen our insights, though only to a point. As with dust creation, on-site noise is associated with temporary activity that tends to alter its patterns on a near-weekly basis, meaning that continual fixed patterns were hard to identify. A location might be noisy one week because trucks would be driving past, but quiet the next, once that particular piece of site work was completed. As such it was difficult to validate patterns of noise over longer periods. An increase in deployment density of noise sensors would likely generate greater insights over a shorter period than longer term monitoring using only the existing deployed devices.

Data gaps were minimal for the study period, though data sets were limited due to late deployment of around half of the EMs devices. It is worth noting that since the data analytics was undertaken, many of the EMS devices have experienced intermittent connectivity, resulting in substantial data gaps. This is believed to be a result of low levels of sunlight in winter, combined with overcast weather, leading to power outages.

Insights from noise monitoring at Melrose Park

The study has shown that the TULIP EMS is an effective device for monitoring average and peak noise levels in the urban environment. Despite the constraints outlined, it has enabled a number of useful conclusions to be drawn about noise levels and trends at Melrose Park.

1. Day/night pattern

A clear day/night pattern in noise that roughly corresponds with work hours.

2. Weekly pattern

A clear weekly pattern in noise that shows greater day-time noise levels on weekdays and quieter daytime noise levels on weekends.

3. Localised impacts

A clear variation in average and peak noise levels recorded by the 12 devices studied which would seem to be a result of their positions in different locations, subjected to differing noise sources.

4. The detection of noisy zones by multiple sensors

A clear correlation in noise levels detected by certain clusters of devices, indicating a relationship that suggests they are detecting the same noise sources. The strongest example of this is a correlation during weekday daylight hours were between EMS005, EMS008 and EMS007, which are located respectively on the corner of Wharf Road and Hope Street, Hope Street and Waratah Street, and just north of the Hope Street site access gate on the main internal site access road. This relationship may be explained by site traffic accessing the Hope Street gate via Wharf Road.

5. Noise in specific areas tends to occur over shorter periods, coinciding with shifting site activity

Over shorter periods of time, patterns of noise evident amongst clusters of devices are more pronounced. Over longer periods, the patterns are less evident, with less correlation evident between any two devices. One explanation for this may be that patterns of noise-creating activity on the Melrose Park site change on a near-weekly basis. Over a period of a week, a single pattern may be in place, leaving a clear imprint in the data. Over a longer period, multiple different patterns may appear and disappear, overlapping each other. When looking at a longer data set, the patterns from these competing signals overlap and create noise, resulting in the weakening of direct correlations between any two devices. Over the longest periods of time some of the only consistent patterns of noise creation are likely to be traffic from major roads. We see this in the data from EMS009, which is deployed on the side of Victoria Road.

6. Loud and quiet

The loudest consistently recorded peak and average noise is from the side of Victoria Road, on the northern boundary of the site. The quietest locations are internal to the site, furthest from roads.

7. We probably detected construction noise (but it's inconclusive)

While there are some noise patterns that may be interpreted as construction noise from the northern end of the site, this is not conclusive. Additional data from EMS001 (the weather station, which was offline during this study period) is expected to assist in more clearly identifying construction noise.

8. Wind creates noise and is a major factor to be considered and understood

Extreme weather events are known to have occurred at the same time as significant spikes in noise for all devices, indicating that high winds can have a significant impact on noise levels detected. It also appears that the direction of wind may result in different patterns due to the localised interplay of the wind and the environment immediately surrounding each sensor. This study has not included a deep exploration of wind data, however it is reasonable to expect that there will be some strong degree of correlation with noise.

9. If there is a link between noise and dust we have not yet found it

Particulate pollution shows a very slight correlation with noise, however this is most likely connected to the day/night cycle and possibly, to a weak extent, traffic exhaust. While certain site based activities may produce dust and noise at the same time, it has thus far not been possible to clearly identify any such events.

10. There is a strong correlation between temperature and noise and humidity and noise

This is almost certainly due to the day/night cycle, as we know that noise levels rise in the early morning, and drop significantly over night, in line with changes in temperature and humidity. There is no reason to believe that the noise/temperature/humidity relationship is causative.

Image: An EMS (008) deployed on the site boundary with Hope Street. Noise events detected here closely matched events detected further up the road, and just inside the site gate. This indicates that all three sensors are detecting heavy construction vehicles as they access the site via the Hope Street gate (credit: Andrew Tovey).

Future directions for noise monitoring at Melrose Park

Increase the deployment density of devices

The spatial resolution of the sensor network was perhaps the greatest constraint on our noise data and the understandings we could gain from it. It is clear that if devices were placed more closely together it would be possible to find stronger correlation between their data. This would allow us to far more clearly discern the presence and nature of localised noise events, such as the passing of heavy vehicles on a specific route. Where on-site noise is occurring at a specific location, a denser network of on-site sensors would also enable us to roughly triangulate a likely zone of occurrence. The current density of devices does not allow this to be done. Any attempts to increase the future density of device deployment would yield clearer understanding of noise on and around the site.

Explore the role of wind in noise creation and dispersal

Initial analysis of Melrose Park noise data did not include a deeper exploration of on-site wind data and how it relates to the creation and dispersal of noise on and around the site. There are two areas of inquiry worth investigating.

High wind is known to generate noise, which has been observed in elevated noise levels recorded across all EMS devices during extreme weather events in early February 2020. A correlation of on-site wind data with average and peak noise data should be run to explore this relationship. Analysis should seek to identify whether there is a linear or non-linear relationship between wind speed and noise, and whether there is a critical threshold of wind speed that causes a significant elevation in recorded noise, effectively muting our ability to discern more localised noise events during high winds.

Wind direction strongly influences dispersal of noise and is critical to our understanding of how construction noise impacts residents. The proposed dispersal modelling transects for future particulate monitoring described above could also be used to study noise dispersal. Analysis could explore how variation in wind speed influences dispersal of noise of different amplitudes.

Focused study of noise in the northern construction zone

Focus attention on the northern construction zone with the aim of building a deeper understanding of ambient noise levels on site: when and where it is generated; how it disperses, with consideration of wind and topography; and how it relates to known site activity. Insights from a focused location can then be used to start developing a noise risk model for the site, which might be applied to other construction zones as they open up. Initially, this work might make use of the existing EMS devices in the northern zone (EMS001, 002 and 003), as well as weather station data. The relatively underutilised nCounter people counting technologies deployed through the project should also be explored in greater detail in this context, as indicators of onsite activity. An expansion of deployed devices should include additional EMS devices in the northern construction zone, specifically along the northern zone perimeter as well as on the new heavy vehicle access on Wharf Road. An upcoming new site access off Victoria Road (late 2020) should also be covered.

Focused study of noise along Hope Street and the southern site access road

Data analysis has shown strong correlation between data from devices located along Hope Street and onto the southern site access road (EMS005, 008 and 007). This correlation is thought to correspond with the passage of heavy construction vehicles travelling between the northern construction zone and Victoria Road, via the southern site entrance on Hope Street. Longer period data sets from existing devices can be analysed to see if stronger long term trends can be established, such as those relating to daily and weekly cycles and to known construction activity on site. Additional EMS devices might be deployed along Hope Street, Wharf Road and Hughes Avenue to increase the spatial resolution of the data and increase our ability to analyse precisely where and when noise is generated.

Integration of video or image capture of noise sources, for on-site validation

If it were possible to integrate video or photographic surveillance technology into the network in order to capture a visual record of noise creation events (for example, passing vehicles), this may significantly augment the data captured by the EMS devices and provide a strong validation for patterns detected in the noise data. This may also conceivably make up for any shortfall in sampling rate.

Development of the TULIP EMS to increase the sampling rate

Any future efforts to increase the sampling rate of the EMS, or to deploy devices with higher sampling rates, would help to generate more detailed data and clearer insights.

Use case

4

Stormwater

Use case 4: Stormwater

- Understand the current stormwater profile of the site during rainfall events of varying size
- Seek insights about the relationship between rainfall and stormwater outflow to inform an understanding of flood and erosion risk on site
- Build a data record of water quality in Ermington Bay, to establish a baseline at the start of the development and to explore possible links between river water quality and run-off from site.
- Establish a system to gather stormwater and river water quality data over longer seasonal periods.

Understanding stormwater at Melrose Park

Understanding stormwater at the scale of a suburb required us to explore a number of complex interrelated processes including weather, land use, topography and stormwater infrastructure. Unlike other environmental variables explored for the project, stormwater is a linear system of flows that relates to a large number of complexly intersecting factors that must be approached in an entirely different way to something like urban heat or air quality. The primary concerns of the use case relate to flooding and erosion risk on the construction site, as well as river water quality in the Parramatta River (Ermington Bay).

Flooding occurs when rainfall intensity exceeds the capacity of the land or stormwater system to move the water away, resulting in a build-up. Given that stormwater infrastructure is designed to handle extreme events, flooding tends to be a localised occurrence that is commonly associated with faults and blockages and is not necessarily contingent upon extreme rainfall. On a construction site, existing stormwater infrastructure may be interrupted or retrofitted. Sediment and debris can also accumulate at drainage pits causing blockages. Furthermore, large-scale earthworks require the installation of temporary pumped drainage systems which can alter the dynamics of expected stormwater outflow. Site managers are keenly aware of how stormwater is managed on a site and with site profile changing on a near-weekly basis, the flows and accumulations of stormwater tend to vary quite dynamically. Flooding is to be planned for and mitigated as it can lead to delayed work schedules, not to mention physical damage. The aim of the project was to explore low-cost methods of gathering stormwater data with the hope that insights might emerge regarding flood risk and mitigation.

Erosion and sediment loss on a construction site are tied closely to flooding, though with disturbed ground even moderate rainfall can generate stormwater with high sediment load, which can wash off-site and into aquatic ecosystems, causing ecological damage. Developers are required to implement sediment control measures on construction sites to minimise erosion and outwash of disturbed sediment. At Melrose Park, this is principally through the use of an on-site stormwater retention pit and flocculation system. Increased flood risk is likely to increase erosion risk, though a more significant factor is the physical status of the site at any given time. That is, how much disturbed ground there is, how much heavy vehicle movement is occurring, and what flood and erosion mitigation measures are in place. While it may be possible to gather data about rainfall and stormwater, it is much harder to accurately track elements of site status. Furthermore, there are no low-cost sensors that enable automated monitoring of sediment load in stormwater. As such, a concern with erosion control guided our efforts to explore stormwater at Melrose Park, however it was recognised that direct utilisation of stormwater data for erosion mitigation is likely out of reach for the time being.

Stormwater from Melrose Park flows into the Parramatta River (Ermington Bay). A southern catchment empties into Ermington Bay while a northern catchment flows east, through residential streets, across a golf course, and into the river about a kilometre downstream. River water quality is a notable concern for riverside councils like the City of Parramatta, as well as for local communities. The amount and cleanliness of stormwater entering the river is the greatest factor for determining water quality. Thus, an understanding of how river water quality varies over time, adjacent to a major construction, provides a valuable insight into the possible impacts of the construction process on the river.

Image: Stormwater collects on site after heavy rain and must be managed by the developer. In areas with disturbed soil, the water becomes loaded with sediment. Various measures are taken to control the flow and retention of stormwater on site and to mitigate erosion and sediment loss (credit: Andrew Tovey).

Stormwater monitoring study design

Stormwater monitoring was the most complex and experimental of the four use cases explored for the project. It involved measurements at five points in the 'flow' of water through the system, illustrated in the diagram below.

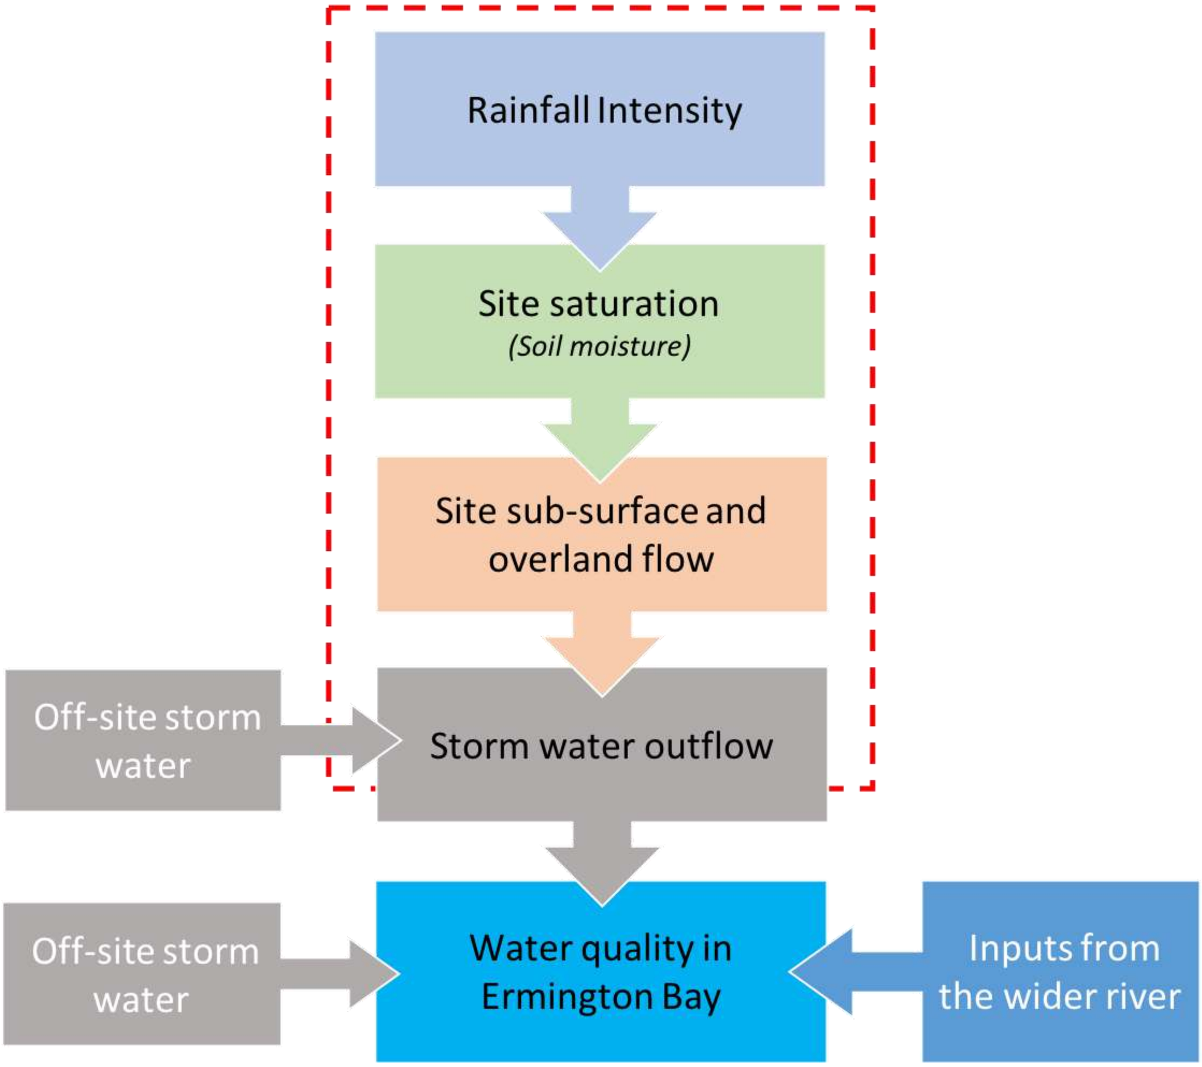

Image: Diagram showing the five points in the stormwater system where measurements were recorded.

Rainfall

Rainfall was captured by an onsite weather station and a tipping bucket rain gauge (TBRG). The weather station was an extension of the TULIP Environmental Monitoring System (EMS) developed by The ARCS Group and UTS, to include a solid state multi-parameter meteorological sensor, the WS10 by German manufacturer Lufft (complete system from The ARCS Group, $6,500[7]). Solid state weather stations are increasingly popular in smart city applications as they are very compact, have no moving parts and are less prone to fouling. Such practical benefits may well outweigh any small constraints on data quality. Solid state rainfall sensors can lose accuracy during high rainfall events and as such are not compliant with meteorological standards used by the Australian Bureau of Meteorology. The use of a solid state weather station was experimental in its own right, enabling us to assess its performance and appropriateness in the context of operational use case development. In order to provide a validation reference, a YSI BOM-compliant tipping bucket rain gauge was installed on site, 25m from the solid state weather station. This device has a real-time 3G data uplink with 5 second temporal resolution, providing a highly accurate reference for rainfall on site and validation for the Lufft WS10 data. Localised rainfall data collected on site can be understood to be more accurate than BOM data collected at Sydney Olympic Park, roughly 2km away, as rainfall patterns can vary over very short distances.

Image: YSI brand BOM-compliant Tipping Bucket Rain Gauge (TBRG) installed on site (credit: Andrew Tovey).

[7] Price differs from cost to project, which was $7,750. The system cost has since dropped.

Site saturation

Site saturation is measured in terms of soil moisture and is one factor that influences the ‘buffering capacity’ of the site during rain – that is, how much water the site is able to absorb rather than lose to run-off. In rural locations this is critical. For an urban site like Melrose Park, with large areas of impermeable ground (concrete, compacted earth), it was not clear to what extent soil moisture actually relates to buffering capacity. Part of the research was to explore this question. Soil moisture was measured in two locations on site using an Enviropro EP100GL multi-depth probe with LoRaWAN communications, supplied by ICT International (via Meshed, $1710 per system). The locations chosen differed slightly in order to provide an indication of on-site variation, one being an exposed and well-drained raised grassy area, the other being a more sheltered and slightly less-well-drained location near trees.

Image: An Enviropro multi-depth soil moisture sensor installed on site (credit: Andrew Tovey).

Sub-surface and overland flow

Sub-surface and overland flow is water that is actively travelling across or just below the ground and entering the stormwater system. This happens during moderate to high rainfall events and may continue after rain if soils continue to drain. We can expect such flow to be higher if soils are saturated or close to saturated prior to the rainfall event.

The presence of sub-surface and overland flow was detected using a smaller and simpler soil moisture probe called the Decagon 10HS ($434 per device). Two devices were deployed at locations where overland flow had been directly observed during or immediately after rainfall. This was a somewhat experimental approach that does appear to have seen some success, with peaks in the data corresponding with rainfall events. Due to the proxy nature of the measurement and the somewhat experimental methodology, the data is interpreted in terms of a low, medium, high or very high probability of overland or sub-surface flow occurring. By identifying when sub-surface and overland flow is likely to be occurring, we have a potentially important data point for localised flood risk. Such flow must have somewhere to drain into and if there is a blockage in onsite drainage infrastructure then water will pool.

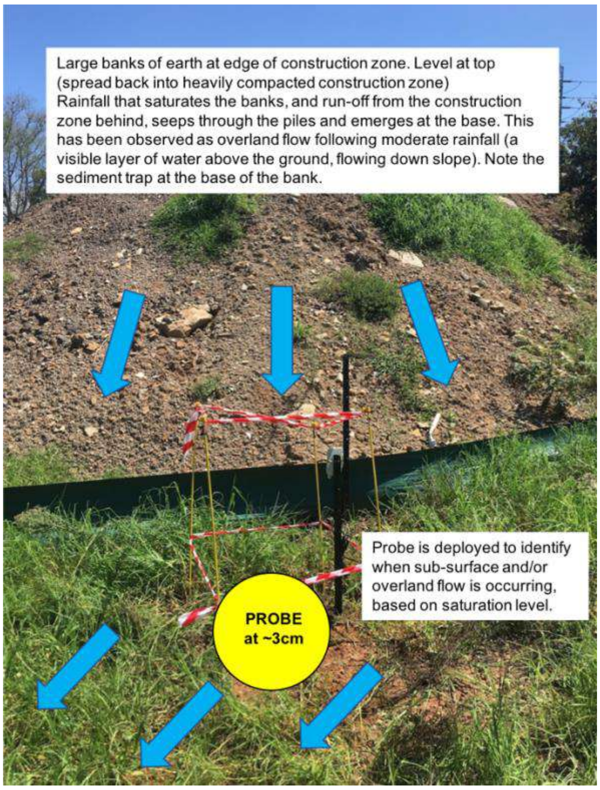

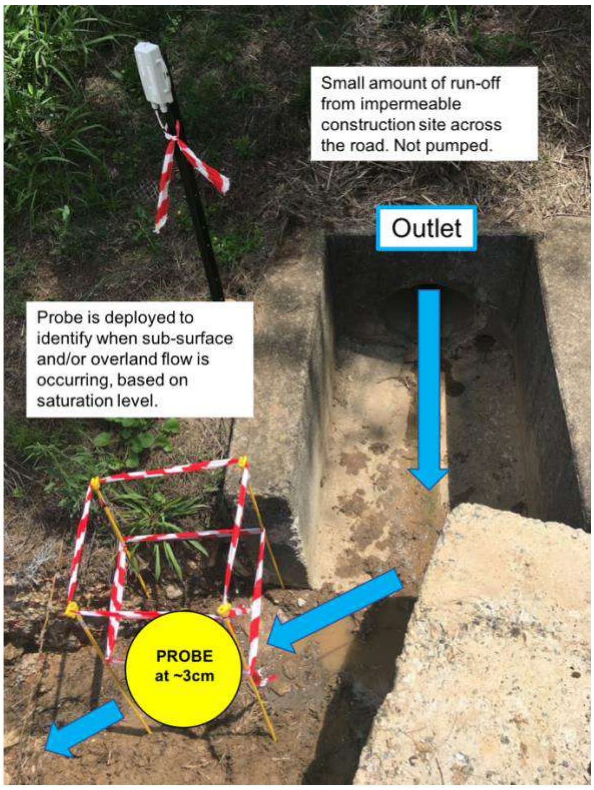

Image: Decagon 10HS soil moisture sensors installed to detect saturation levels concurrent with sub-surface and overland flow (credit: Andrew Tovey).

Stormwater outflow

Stormwater outflow refers to water that has entered the stormwater system of pipes and channels that are designed to drain an area of rain water and transport it into the river. Most but not all stormwater originates as rain - on a construction site, some potable water is used for cleaning of equipment and concrete and this adds to stormwater outflow. Furthermore, the Melrose Park construction site uses a stormwater retention pit and automated flocculation system to remove suspended sediment and improve the water quality of stormwater leaving the site. Discharge from the pit thus originates as rainfall though its release does not correspond with natural run-off from rainfall events.

The most common method for remote sensing of stormwater flow is to use a fully immersed probe that sits inside the pipe or channel. Apart from significant practical constraints relating to access and real-time communications, such devices are relatively costly (a quote acquired by the project for such a device was over $20,000 for one system). The project set out to explore a lower cost option for stormwater monitoring that is accessible for scaled deployments by Council and developers like PAYCE. The solution arrived at was low-cost ultrasonic sensors used to measure the height of flow through a stormwater pipe inflow or outflow. LoRaWAN ultrasonic sensors are an established commercially available product. The sensor emits a burst of high frequency sound, records the time elapsed until an echo is detected, and converts this into a distance, accurate to one millimetre. Their traditional commercial use is for tank and silo fullness monitoring. At Melrose Park we deployed six ultrasonic Maxbotix sensors, connected to Elsys LoRaWAN nodes ($653 per device). Sensors were mounted above an outflow or inflow, pointing straight down and providing a 15 minute update on height of flow.

Ultrasonic sensors were deployed at two on-site locations where stormwater enters pipes, including the point where pumped discharge enters the system from the on-site retention pit. Remaining sensors were deployed at major stormwater outlets off site, two of which served the site and two of which served surrounding residential streets. The aim was to develop a record of stormwater flow trends through the different pipes and look for correlation with recorded rainfall, as well as possible hydraulic relationships with soil moisture, sub-surface and overland flow data.

Image: An Elsys/Maxbotix ultrasonic sensor positioned above a stormwater outflow. The device measures the distance to the water, which we convert into 'height of flow' (credit: Andrew Tovey).

Water quality monitoring

Water quality in Ermington Bay was monitored using the EXO2 sonde from Xylem, which carries an array of six YSI-brand water quality sensors. The Sonde is mounted on a buoy, with solar power and 3G communications for a live data feed. A full spectrum of water quality parameters were measured, including temperature, salinity, pH, ORP, turbidity, dissolved organic matter, dissolved oxygen, chlorophyll-a and blue green algae. The complete system, including visibility on a dedicated dashboard, cost $29,381, and constituted the single most expensive device deployed for the project. Marine environments are notoriously harsh on sensitive equipment and accurate water quality monitoring is itself a highly precise science. Despite some lower cost options in the space, we decided to use the Xylem/YSI solution, not least because it was recommended by CSIRO as the instrument of choice for remote water quality monitoring. The buoy was deployed in the mouth of Ermington Bay, immediately south of Melrose Park, in around 1.7m of water (low tide depth).

The aim was to understand a baseline for water quality in Ermington Bay at the start of the project, and beyond this, to look for possible causative relationships between rainfall, stormwater monitoring at Melrose Park, and water quality in the bay. The interplay between the various parameters in an estuarine environment is highly complex and the statistical noise generated by the flow of the river and tidal flux makes it difficult to link observations in the bay to specific local processes, as opposed to wider changes in the river system. Thus the buoy was deployed with a strong interest in the river water quality for direct environmental reasons, with the hope (but no certainty) that relationships with the construction site may also emerge. It should also be noted that most construction during the project period was in the northern site catchment, which does not flow to Ermington Bay. Only stormwater from the southern site catchment flows to the bay, however this area is yet to undergo any major demolition or construction work. Therefore, any correlation between water quality data and on-site monitoring can be understood as a record of baseline processes, with the true value of such insights only emerging as the site is developed further in coming years.

Image: The Xylem EXO2[s] water quality monitoring buoy, deployed in Ermington Bay (Credit: City of Parramatta Council)

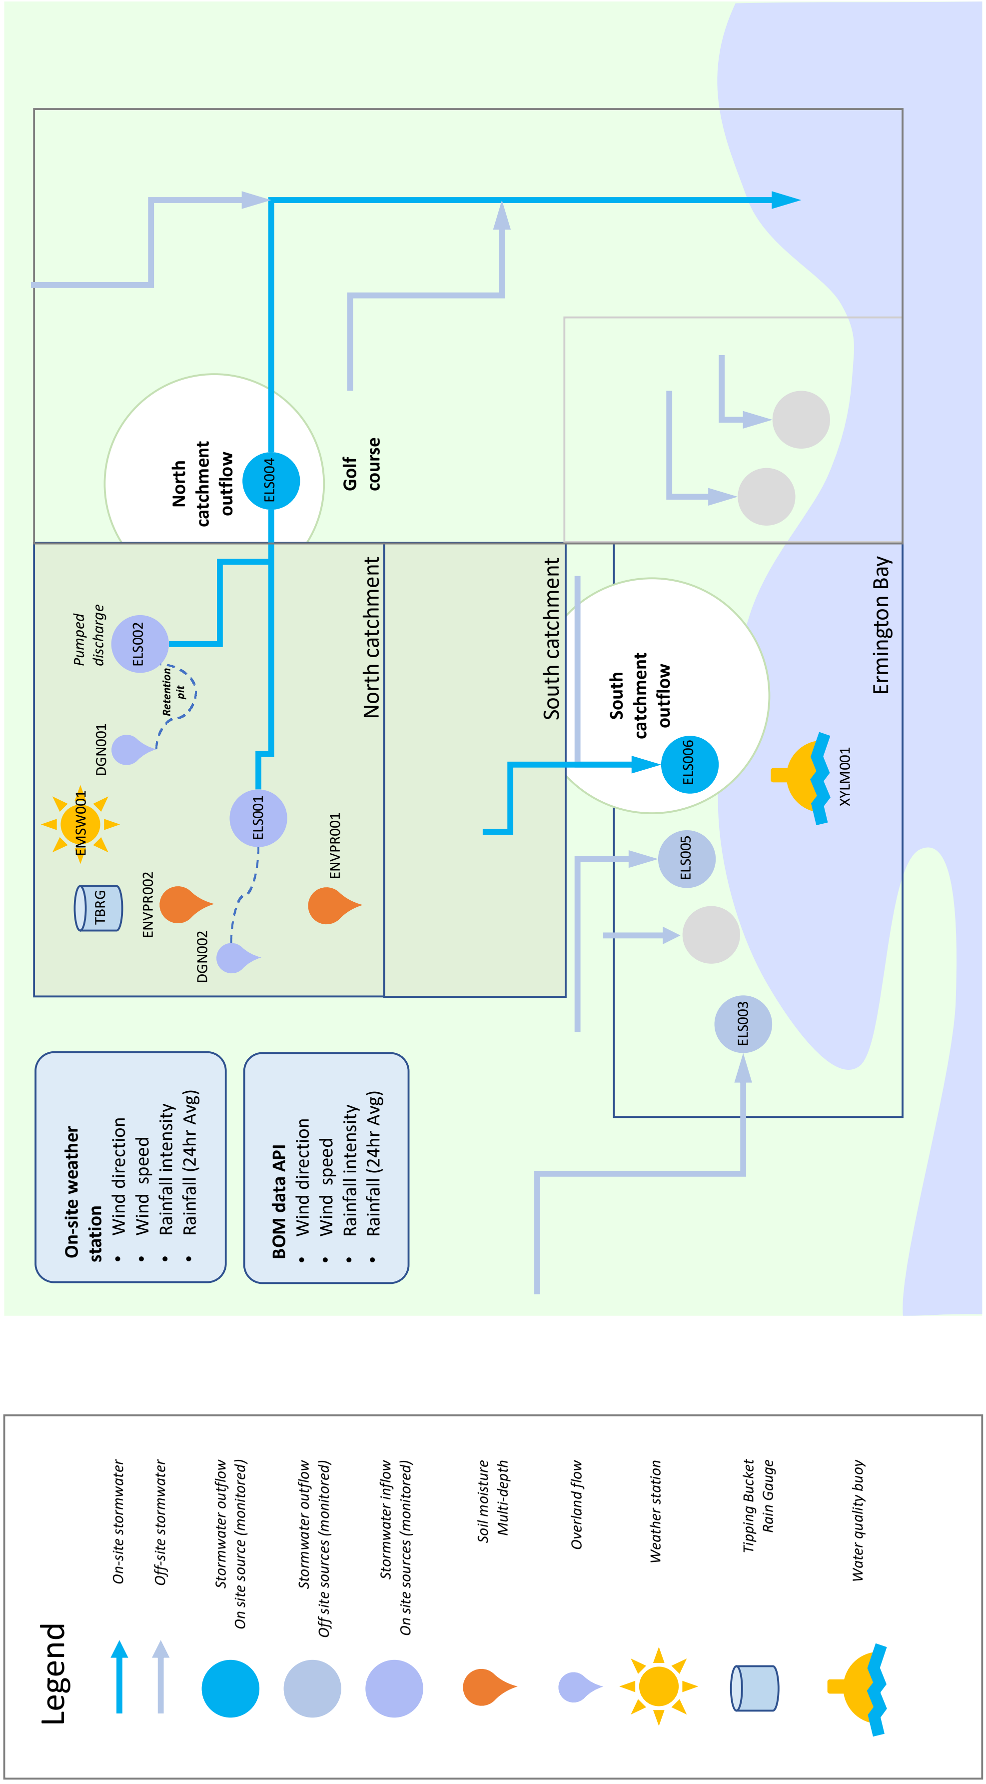

Image: Schematic map of the stormwater monitoring system deployed at Melrose Park

Constraints of the stormwater monitoring study

Limited budget and an experimental approach

As noted, many of the monitoring methods explored for this project were low-cost and experimental in nature. High cost solutions to stormwater monitoring do exist and are commercially available. However, if we were to take stormwater flow monitoring as an example, the higher cost solution (fully immersed sonde probes) would cost in the region of $130,000 for six devices, as opposed to around $4000 that was spent on low-cost non-immersed ultrasonic sensors. The higher cost approach is over thirty times more costly. Our aim was to explore what was possible if we applied existing low-cost sensors in novel new ways. This is driven by a belief that stormwater monitoring is likely to have operational value for a construction site, but that the value may require solutions that cost tens of thousands of dollars, rather than hundreds of thousands. This position was further emphasised by the limited budget available within the project for stormwater monitoring. The experimental nature of the stormwater monitoring undertaken for this project mean that deep insights about stormwater trends for the site were only partially realised. Furthermore, some of the most valuable outcomes of the work were methodological, being insights into how we might use low-cost sensors to deliver relatively high value information.

Limited time period

Another major constraint of this study was the limited timeframe available for collection of stormwater data prior to analysis. Design and implementation of the stormwater monitoring system deployed for this project took longer than the other use cases. This was due to the complexity of the solution, involving multiple device types, all of which were new to the project team, with significant work needed to commission devices, calibrate them, and usefully interpret data. As a result, a complete data set incorporating all device types was only available for around one month, prior to analysis. This period did include some moderate rainfall events. The EXO_2 buoy, the weather station and some of the ultrasonic sensors were commissioned earlier, giving some extended data sets for certain parameters, including coverage of a couple of heavier rainfall events. However, overall the total available data was somewhat limited. Stormwater data relates most strongly to rainfall, which can be highly seasonal. In order to gain a stronger understanding of stormwater on the site it is recommended that a full twelve months of data from all devices is explored as a future study.

Calibration of devices and development of deployment-specific models and thresholds

All parameters measured for the project had a series of thresholds associated with them tied to a ‘low’, ‘medium’, ‘high’ and ‘very high’ status. For urban heat, air quality, and noise, these thresholds were linked to established external reference scales, as well as to early accumulated data sets from Melrose Park. Thresholds applied to all devices that measured the same parameter and were what we referred to as ‘universal’. For stormwater monitoring, things were different because these thresholds were relative only to specific devices. Firstly, due to the experimental nature of the monitoring methodologies, there were often no external references against which data could be compared. Secondly, the data from an individual ultrasonic sensor or soil moisture probe turned out to be highly relative to the deployment context but not directly comparable even to other devices in the study. For example, the two deep soil moisture probes showed general agreement in their trends, however one location was significantly more drained and drier than the other, meaning that an assessment of soil moisture for the site could not be associated with one value in both locations. If we ask what a ‘highly saturated’ site looks like we need to look for quite different values at each sensor. No doubt a third sensor, with different drainage and soil type, would require its own relative threshold value. Likewise, each ultrasonic sensor was deployed above a different sized stormwater pipe and each pipe was connected to a different network of pipes and drains. Therefore, the height of flow and indeed the ‘fullness of pipe’ (a derived metric used to compare flow through pipes of differing size) varied significantly by location. The only way to determine what constitutes ‘low’, ‘medium’, ‘high’ or ‘very high’ flow through each pipe is to build up enough data from each device in its deployment context that a long- term pattern can emerge that is specific to that precise pipe and deployment location.

The constraint then, was the lack of time available to build up a representative data set at each location. This constraint actually applies to all stormwater parameters measured, including water quality. While water quality parameters do have external references, these tend to be highly variable by the type of aquatic environment, latitude, soil and vegetation type. The specific context of water quality in the upper Parramatta River is therefore unique to acidic soils of central coastal NSW in a highly tidal inner estuarine bay. The urban nature of the catchment adds a further dynamic. This all serves to make our particular deployment of the buoy somewhat unique, meaning that while external references for water quality parameters are a useful and relevant touch point, they cannot define fixed thresholds for our particular use case. In fact, the literature relating to water quality monitoring is quite clear that the only way to build a reliable reference framework for most data sets is to build up an initial body of site-specific data, ideally over at least twelve months.

Limited access to hydraulic plans showing stormwater systems on private property

City of Parramatta Council was able to provide plans for stormwater drainage infrastructure on public land and was able to obtain the equivalent documents from the City of Ryde Council for areas east of Wharf Road. However, public plans do not contain details of stormwater infrastructure on private property. Each private property has its own set of plans and these are not accessible through any central database, public or otherwise. The Melrose Park development site is made up of a number of formerly separate private properties. The project team sought access to these plans but was only successful in obtaining plans for the former Pfizer property, which covers a substantial area in the north and central-eastern part of the site. Luckily these plans related to the main area of focus for onsite stormwater monitoring. However, stormwater plans for the rest of the site, including the entire southern catchment, were not available. This limited our understanding of how stormwater behaves on the site.

Much of the land to the south of the development site is light industrial and the drainage plans for these properties were also not able to be obtained. This was not a concern for our study of the northern site catchment, however it does relate to the southern catchment and to water quality monitoring in Ermington Bay. It is likely that stormwater impacts on the Bay will strongly relate to stormwater management and potential contamination from that light industrial zone. As the southern part of the PAYCE site begins to undergo active development, stormwater and sediment from the site will start to flow through the southern catchment system and into Ermington Bay. It will then be important to understand the relative impact of the light industrial zone.

What data did we collect and how can we use it?

Real-time data about localised rainfall, soil saturation, and the flow of stormwater off site provides a continuous record of how the site behaves during storm events. The project explored the use of low-cost sensors to gather this data, recognising that approaches taken were somewhat experimental. It was recognised that an urban hydrological system and its connection with human activity on the one hand, and river health on the other, amounts to a significantly complex system. With so much complexity and so many variables to consider, modelling the system is an enormous challenge. More so than other use cases developed for the project, stormwater monitoring took on the attributes of a so-called *smart living lab*, whereby the collection of data was approached in a somewhat open-ended fashion and the relative usefulness of that data was explored iteratively as it accumulated and as we refined our methods. To date, actual data sets that capture significant rainfall events are very limited. Much longer periods of data collection are needed in order to start drawing conclusions about the relative usefulness of the data. Furthermore, users such as site managers and Council need to explore data over more time, particularly when large rainfall events occur within the context of real-world operations. Novel applications for data may emerge, from on-site rainfall to stormwater outflow on the golf course. It is this real-world engagement that must guide the development of future enquiry.

Rainfall

Two parallel sets of rainfall data were gathered, from a solid state weather station on-site and from a Tipping Bucket Rain Gauge (TBRG) also deployed on-site. The TBRG can be taken as the more reliable sensor as it complies with BOM standards and also has a more frequent reporting rate. Longer term analysis of data from both devices to determine the relative performance of the Lufft WS10 would be a useful further study. Only one month of TBRG data was available, with just one relatively insignificant rainfall event, meaning that conclusions about on-site rainfall monitoring are very limited for the time being.

It is worth noting that both devices provide a high temporal definition compared to available BOM data, which is only available as a daily accumulation. Our data tells us when rain occurred during the day, to the nearest minute. It also tells us the rainfall intensity, which is not available from the BOM. This level of detail is vital for understanding how the site behaves during stormburst events.

Soil moisture

Two Enviropro multi-depth soil moisture probes were deployed on the site, in locations with slightly different characteristics of soil, slope, elevation and exposure. The differences between locations, in terms of their different peak saturation points and the rising and falling trends of moisture, are not yet well understood. The variation indicates that no one location for soil moisture measurement can act as a site-wide proxy, though a combination of locations may provide a more reliable picture. For now, attention should be on the timing of trends, rather than any absolute VWC figures.

Soil moisture data may be useful for a number of future inquiries. We can use it to build an understanding of how the site becomes saturated in relation to long-term weather patterns. This has implications for understanding 'muddiness' and localised flood risk (which may impact construction activity), and the likelihood of dust creation (which may impact air quality). It also relates to the relative impact of a given rainfall event, thus if we forecast rain and the ground is very wet then we can anticipate a higher flood risk than if the same rain is forecast with dry ground.

Overland flow

Two Decagon 10HS soil moisture probes were deployed on-site to measure soil moisture at approximately 3cm depth. Both sensors were deployed in locations where surface run-off has been observed during moderate rainfall events. The purpose of these deployments was to observe when surface run-off is likely to have occurred. As we are measuring soil moisture rather than flow directly, data acts as an imperfect proxy indicator for flow.

Additional complexity arises from the detection of flow events in the drainage channel (DGN001) that do not correspond to rainfall. The most likely explanation is that mains water is being used as part of active construction activities, possibly for cleaning earthmoving equipment, leading to low-level run-off through the channel.

Going forward, the data from these two locations may be useful for understanding how much rain needs to be falling on the site for overland flow to occur. Such flow may be an indicator of localised flood risk during lower volume rainfall events in instances where there are blockages to stormwater infrastructure. There may also be scope for exploring more contextual questions, for example, investigating the effectiveness of sediment traps used up-slope from DGN002.

Stormwater outflow

Four ultrasonic sensors were deployed above stormwater outflow pipes: one on the golf course, which carries all run-off from the northern site catchment, as well as adjacent suburbs; one that carries water from the southern catchment into Ermington Bay; and two others that enter Ermington Bay but which do not carry water from the site. All four sensors have produced near continuous readings that can be interpreted as 'height of flow' in millimetres. These readings may have a +/- accuracy of up to 20mm, despite a 1mm precision, so the data cannot be reliably modelled using traditional analytics techniques in order to indicate outflow volume. However, these sensors do provide a record of general flow trends that match rainfall events, indicating lag time and the rising and falling profile of the flow, even if we can't identify actual volumes. Over time, we can build up a relative model for each location, where we can understand how height of flow varies by rainfall characteristics and other critical factors. Initial focus is on the two locations that connect with site run-off. The additional Ermington Bay locations may be relevant to longer term analysis of water quality data from the Bay.

A fifth device was deployed over a stormwater inflow on the development site, capturing natural flow through a swale. This data may compliment the outflow data as well as the overland flow data discussed above.

It should be noted that the use of these ultrasonic devices is highly experimental. There are no existing references that we are aware of. Data sets must be allowed to accumulate to establish range, and thresholds relative to rainfall. All data sets contain extreme erroneous negative values. We are not clear what causes these but given that the issue appears worse with large pipes we believe that something may be occurring where the ultrasonic signal bounces around inside the pipe opening. To add to the complexity, more turbulent water, associated with higher flow, may increase the likelihood of error, meaning that high flow events may be more likely to give false readings. As such, all data must be checked and cleaned before analysis and a great deal of further work is needed to understand the limitations of this monitoring approach.

There are a number of possible uses for this data. It provides a clear record of relative outflow in specific locations, which may be useful for PAYCE's engagements with local stakeholders such as the nearby golf course. We may also be able to build models of how stormwater outflow relates to rainfall events with different characteristics, or explore the ratio of site to non-site stormwater entering the golf course by comparing to data relating to onsite stormwater inflows. In the longer term, this data may be incorporated into a machine learning model of the catchment that ultimately may be able predict outflow based on weather forecasts, soil moisture and other critical variables.

Pumped discharge

A sixth ultrasonic sensor was positioned above the outflow from a flocculation system, which removes sediment from water held in an on-site retention pit and pumps it into the stormwater system. This water then flows out through the golf course to the river. Pumping is automated and occurs on a timer. It also results in stormwater outflow, detectable downstream, that does not correspond with rainfall events. The data from this sensor provides a 'height of flow' metric, that can be interpreted to show if pumping was or was not occurring. It also indicates when the system is being used at a higher capacity, aligned with high rainfall.

Pumped discharge data may be useful to PAYCE as an independent record of when the flocculation system operated and at what capacity. This may assist with management of the on-site retention pit. The data also helps us to interpret height of flow data recorded on the golf course as it is likely to be the only major input during dry periods.

River water quality

A multi-parameter water quality monitoring probe was deployed on a buoy in Ermington Bay in December 2019, providing a continuous record of critical water quality parameters. These parameters are physical and biological attributes - it is not possible to automatically monitor chemical contaminants with a low-cost sensor. Data relates to the trends of water quality in Ermington Bay. Stormwater from the smaller southern site catchment enters the Bay, alongside stormwater from surrounding industrial and residential areas. The long term goal is to understand how water quality trends relate, if at all, to stormwater influx to the Bay. This may prove to be quite difficult given the much larger movements of water resulting from the tidal flow of the river. During the project, no significant construction occurred on the southern site catchment, meaning that existing stormwater infrastructure remains unaltered and there are no activities that might generate sediment loss into Ermington Bay. Therefore, we can assume that the impact of the site on Ermington Bay has not changed from pre-development conditions. Activity on-site is in the northern catchment, which does not connect to Ermington Bay. Therefore, the data collected by the buoy to date should be viewed as the accumulation of a baseline. As the southern catchment begins to see active construction, the possible impacts on water quality in the Bay can be explored.

Detailed discussion of early data sets from the various stormwater sensors can be found in two detailed project documents (available upon request):

1. *Melrose Park: Stormwater data interpretation - Regarding data produced by soil moisture and ultrasonic sensors*
2. *Melrose Park: Water quality data interpretation - Regarding data produced by the Xylem EXO2s smart buoy in Ermington Bay*

Insights from stormwater monitoring at Melrose Park

Stormwater monitoring data and analysis has been very limited due to a delay in device deployment and a small number of rainfall events captured at the time of initial data analysis. The strongest insights at the end of the project relate to the appropriateness of the various methods of stormwater monitoring. Deeper insights into environmental processes will require much more data. However, we can make some high level observations, as follows:

Rainfall

1. **Rainfall is highly localised**

   We see a general agreement in trends between the record of daily rainfall recorded by on-site sensors and the record of daily rainfall recorded by the BoM at Sydney Olympic Park, 2.6km away. Our TBRG can be assumed to produce data of a comparable quality and accuracy to the BoM. We see significant variation in total accumulated daily rainfall between the two locations. For example, on the 22nd of May 2020, rainfall at Sydney Olympic Park was 41mm, nearly double the 21mm recorded at Melrose Park. This shows that rainfall can be highly localised, with heavy rain in one location and moderate rain just a couple of kilometres away.

2. **Rainfall varies significantly by season**

   We have also captured limited rainfall events to date. It is therefore useful to collect at least 12 months of rainfall data in order to build up an understanding of normal range that is specific to the site.

Soil moisture

An initial focus was made on readings at 10cm, over a seven week period from March to May 2020. With too few rainfall events on record at the time of initial analysis, it was not possible to draw strong conclusions about how soil moisture relates to other factors. However, we can say the following:

3. **Soil moisture spikes fast after rain, then reduces slowly**

   Both locations show similar variations in line with rainfall events, with rapid increases in soil moisture as rain occurs, followed by slow reductions over many days, associated with dryer weather.

4. **The saturation peak associated with rainfall depends upon the characteristics of the rainfall event *and* the existing saturation of the soil**

   A high level review of data shows a marked increase in soil moisture at both measurement locations at the start of winter, coinciding with significant rainfall on the 22nd of May, and the 8th of June. We can also see initial evidence that the saturation peak recorded by the sensor is dependent upon a combination of rainfall *and* existing saturation at the time when the rainfall began. There was nearly 50% more rain on the 8th of June than on the 15th of June, however the device registered a significantly higher soil saturation on the 15th. This is likely due to a ten day dry period prior to the 8th, presenting well-drained soil with a high capacity to absorb rain. On the 15th, despite less rain, the ground was already very wet as a result of intermittent rain since the 8th, causing a higher saturation peak than on the 8th.

   This example illustrates how soil moisture information can help us to predict the relative impact of rainfall on a site, and notably, the likelihood that standing water will pool on the surface - something that has been directly observed at ENPR001 following heavy rain. Standing water is its own form of localised flooding and is a management concern in active construction zones.

Overland flow

5. The devices work as intended, but larger data sets and further direct visual observation is needed to calibrate thresholds and determine baselines

Early data analysis of a seven week period from March to May 2020 supports the idea that the Decagon soil moisture sensors are able to capture overland flow events that coincide with rainfall of a certain duration and intensity. However more analysis is needed to identify the critical thresholds that lead to these measurements. Early data has allowed provisional thresholds to be set for each device, with ‘high’ values likely corresponding to active flow events. Each location is unique, making data highly contextual. It can only be interpreted clearly once larger data sets have accumulated and ongoing direct observation of conditions, ideally during rainfall events, is most likely needed to calibrate our data analysis over the longer term.

The device deployed below the earth mounds, behind the site office (DG002) clearly captured all major rainfall events. We see very sharp increases in soil moisture immediately after the start of a rain event, and supported by direct observation of surface water at the location during rain, we can be confident that the reading stands as a reliable proxy for overland flow. Very small rainfall events are evident as an increase in soil moisture, however this is not interpreted as a proxy for the presence of overland flow.

The critical question remains: what is the threshold for soil moisture at this location which corresponds to the start of overland flow? To answer this, detailed visual records would need to be made before, during and after rainfall events, allowing us to build a more accurate proxy model. For now, we must remain conservative and assume that only VWC readings greater than 0.34 are likely to stand as a reliable proxy.

Stormwater outflow

6. The devices work as intended and peak outflow events match peak rainfall events

Outflow to the golf course (ELS004) for May and June 2020 shows a series of peak outflow events that correspond in size and time to peak rainfall events. In total, ten events of clearly elevated flow are evident for the two month period, all matching to rainfall data.

The other ultrasonic sensors have performed similarly, supporting the experimental method. As discussed, we do not know if we can trust flow heights as absolute figures so it is prudent to focus on general trends and timings. A 349mm reading is 'very high', but so is a 300mm reading ,and for now at least, two such readings should be treated as broadly the same until we build more confidence in the method. For this reason, the data dashboard emphasises 'low', 'medium', 'high', and 'very high' thresholds as operational references, reducing the focus on specific values. As data accumulates the initial thresholds outlined for each ultrasonic device should be reviewed.

7. Sensors detect flow associated with very low rainfall and moderately high rainfall

Very low rainfall of 1mm is visible as a 20-30mm rise in height of flow (e.g. 2nd June).

The largest rainfall event of that period was 35mm, which fell over twelve hours on the night of the 21st of May, corresponding with a peak outflow at ELS004 of 349mm a couple of hours after the start of the storm and a return to base flow around four to five hours after the end of the storm.

Image: Stormwater outlet on the golf course (ELS004) to the east of the site (credit: Andrew Tovey)

River water quality in Ermington Bay

Nine water quality variables are being measured in Ermington Bay. A detailed review of all of these can be found in project document: '*Melrose Park: Water quality data interpretation - Regarding data produced by the Xylem EXO2s smart buoy in Ermington Bay*'.

Certain variables are known to be more closely tied to sediment load in stormwater and might be considered to be primary areas of focus with respect to the Melrose Park development. These are Blue Green Algae, pH, and Turbidity and are discussed below.

i) Blue Green Algae (BGA)

Elevated BGA can indicate what we call a harmful algal bloom (HAB). These have been linked to increased temperature, lowered salinity, raised turbidity, dissolved oxygen (growth in high DO, leading to DO depletion), pH changes, chlorophyll increase and DOM increase – most of the other parameters that we measure. A rule of thumb is that a bloom is a symptom of disruption to normal ecological balance. If critical parameters change away from levels that sustain healthy biological communities, certain communities die, others grow, and often algal blooms come out on top. The scenarios that lead to a harmful bloom are numerous and there are many complex interrelating factors. However, as with Chlorophyll a, one common driver is high rainfall events. These can wash pollutants into the water, or nutrients, or change the pH, or the temperature or the salinity. A combination of these changes following high rainfall can sometimes trigger a bloom. There is generally a lag of 10-14 days associated with the growth profile of the BGA community in response to the new nutrients and changed conditions.

The highest level of BGA recorded over four months from December 2019 to April 2020 was 383.07ug/L. The lowest level recorded for the same period was 2.38ug/L. There are no official guidelines for 'safe' or 'normal' levels that can be applied to our data, however these observations can be considered somewhat typical for an estuarine system in NSW.

Image: Ermington Bay at sunset (credit: Andrew Tovey)

Between mid December and late April, two significant 'bloom' events are evident. The first (3 day duration) peak occurred 15 days after a moderate rainfall event, which itself occurred after a prolonged period of extreme drought throughout December and early January. It is likely that the rain washed accumulated nutrients into the Bay and an algal bloom appeared two weeks later. A similar pattern can be seen with the second larger peak (10 day duration), which began 7 days after an extreme rainfall event and peaked 14 days after. It is highly unlikely that these blooms are connected to development activity at Melrose Park because active work is occurring in the northern site catchment meaning that outflow does not enter the bay. Therefore, these kinds of bloom event can be understood to be part of the normal seasonal variation at Ermington Bay.

Outside of bloom events, tidal fluctuation is strongly evident and there appears to be a baseline fluctuation between around 15ug/L and 35ug/L. This may be seasonal and we might expect lower levels in winter. It is difficult to assess what a normal baseline is without larger quantities of data across at least twelve months.

ii) pH

pH is a figure expressing the acidity or alkalinity of a solution on a logarithmic scale on which 7 is neutral, lower values are more acid and higher values more alkaline. Ocean water is very slightly alkaline, with a pH of 8.1. Estuarine water can be more variable because fresh water pH varies as a result of interaction with the land. River water will tend to be more acidic than ocean water. Impacts on freshwater and estuarine pH generally originate from three sources: rock/soil; man-made environment and chemical pollutants; vegetation.

Construction sites are a well-recognised source of stormwater that can alter the pH of aquatic environments. Pollutants can cause changes that are either acid or alkaline. Concrete is strongly alkaline, however many soils (and notably those around the Upper Parramatta River) are relatively acidic. Most authorities dictate that stormwater leaving a construction site must be within a pH range of 6.5 to 8.5 and the most effective way to ensure this is to mitigate sediment loss from site. ANZEC guidelines for a NSW estuary state lower and upper pH values as 7.0 and 8.5.

Our data shows that significant changes to pH are more likely to coincide with algal growth than they are to indicate the direct impact of stormwater runoff. This makes it harder to attribute spikes in pH to land-based activity such as on the construction site. However, algal growth does respond to run-off, there is just a lag-time of approximately 14 days. Sure enough, the 22nd February pH peak relates to a 21st February algal peak, which is 11-14 days after a severe and prolonged rainfall event between the 7th and 10th of February. It is likely that the rain washed land-based nutrients into Ermington Bay, kicking off an algal bloom formation that peaked two weeks later, in turn resulting in a pH peak. It is clear that while pH levels in the Bay may be connected to future construction activity, they also exist as part of a complex balance with multiple other variables and it is likely to be very difficult to pinpoint unequivocal cause and effect.

It is vital that longer term trends are observed, over at least twelve months, to build up a clearer baseline for pH variation in the Bay.

iii) Turbidity

Turbidity is the indirect measurement of the concentration of suspended solids in water. Sources include silt, clay, sand, algae and organic matter. Sediment from construction sites is recognised as a major source of solids, making turbidity a water quality parameter of high interest to the project.

Turbidity can fluctuate with stormwater ingress. It also tends to be less impacted by rainfall after a longer rainy period because all the particles are washed into the waterways early on and subsequent rain is not washing much more away. Thus the most pronounced impact of rainfall on turbidity comes after a longer dry spell. Tidal fluctuations are evident, though they are subtle.

Based on the first four months of data, a 'normal' range for turbidity at Ermington Bay is between 5FNU and 10 FNU, with occasional spikes to 15 or 20FNU, perhaps associated with small rainfall events, though potentially related to a number of other occurrences. For example, rainfall events wash nutrients into waterways and this can cause a localised phytoplankton bloom a couple of weeks later, which shows up as elevated turbidity. One example of this may be the concurrent FNU/BGA spike around the 2nd of February, two weeks after a moderate rainfall event on the 17th of January, which occurred after a period of prolonged drought and may well have washed a high amount of dust and nutrients into water (which acted as food for the algae).

Extreme weather can also impact turbidity over shorter periods. Sunday 9th February was an extreme rainfall event in Sydney. The event created a significant signature in the turbidity data from Ermington Bay, with a spike in FNU from 30 to 348 over a five hour period. This can only be attributed to the extreme weather and may be a combination of stormwater influx and choppy surface conditions associated with gale-force winds, which would then have disturbed sediment on the river bed.

ANZECC guidelines for turbidity in an estuarine environment are not useful as they are conflated with a purely marine environment and also differ markedly from the figures given for lowland rivers. Thus the figures are far too low to be useful (0.5 FNU is considered ‘low’ and 10 FNU is considered ‘high’). As such, we must build an assessment of what constitutes a ‘normal’ baseline for Ermington Bay using data that we accumulate, over at least 12 months.

As with other parameters, the relationship of turbidity with the ecological and hydrological processes of the bay is highly complex and it is difficult to make clear causative links to stormwater and sediment load. In the future, as the Southern end of the construction site becomes active and we see the potential for sediment loss into the bay, it is certainly possible that we might detect associated elevations in turbidity. However, to establish clear links we will need to build longer term baselines of data for multiple variables and increase our understanding of some of the key interrelationships through ongoing study of the location.

Future directions for stormwater monitoring at Melrose Park

Extend data sets for a more complete understanding

From rainfall to soil moisture, and from ultrasonic data to water quality data, our greatest constraint is limited data sets. Stormwater variation is highly seasonal, meaning that a bare minimum of twelve months of data is needed to develop clear baselines for the location and provide a reference for assessing future change. Once this data has been gathered, we can:

a) Review and update device-specific thresholds for all stormwater devices, including water quality parameters. Most stormwater thresholds are calculated based upon early data sets, which lack a full range of normal events or seasonal variation.

b) Explore ‘high’ and ‘very high’ rainfall events at all five levels of system monitoring. We have not yet captured extreme weather events with all devices. A series of such events did occur in early February 2020 and these were captured by the Lufft weather station and by the EXO2 buoy, however the soil moisture, overland flow and stormwater outflow sensors were not operational at that time.

Run a 12 month co-location study on the two rainfall sensors

Compare rainfall data from the Lufft WS10 solid state weather station, the YSI Tipping Bucket Rain Gauge, and available BOM data from Sydney Olympic Park over a twelve month period.

There are some concerns about the relative accuracy of the solid state weather station, which uses a doppler radar technology to monitor falling rain, converting this into a proxy value for accumulated rain in millimetres. While the manufacturer claims high accuracy and sampling rate, there would appear to be some discrepancy between Lufft data and TBRG data in our early data sets.

The two devices are positioned 25m apart, both with clear sky, meaning that they should be measuring the same rainfall. The TBRG can be assumed to be accurate as it measures actual (rather than proxy) accumulation and it is BOM-compliant. However, there is value in a solid state weather station. It has no moving parts and no areas that can be fouled, meaning that maintenance is very low. It is also many sensors in one compact device, so while there may be some compromise in accuracy for rainfall monitoring, it also measures wind, sun, pressure and temperature. A traditional weather station, which might include a TBRG as one of its components, would be a large collection of parallel devices. Due to the compact nature of solid state systems, they can be deployed easily and they have a discreet profile. This may be critical for deployment of weather stations in public spaces - for example, if PAYCE wishes to maintain a system at Melrose Park once it is open to the public.

Given the benefits of working with solid state weather stations it is necessary to develop a better understanding of the WS10’s accuracy and limitations, to inform the possible future use of other solid state systems. The performance of both systems may also be validated to some extent against BOM data, despite the fact that rainfall will vary between the two locations.

Explore additional methods for acquiring stormwater data using low-cost sensors

Initial monitoring explored through the project has been a useful first foray into low-cost stormwater monitoring on a construction site but it was by no means comprehensive. The following things might be explored in an effort to expand effective data capture methods in the area:

i. Stormwater retention pit level monitoring

A stormwater retention pit on the northern end of the site could be monitored for its depth by using an ultrasonic sensor of the same type deployed over stormwater outflows. This would provide a record of the rise and fall of the pit level. Such data is useful as it can conceivably assist with management of the pit. Pit discharge should occur prior to major rainfall events to ensure that it is not inundated, so knowing its current buffering capacity is useful if we also know the likelihood of upcoming heavy rainfall.

ii. Stormwater drainage pit monitoring

This involves the installation of a device in a drainage pit to measure the depth of water in the pit, and potentially the build-up of sediment (a concern as it may block the pit). UTS is aware of some experimental work using multi-depth soil moisture probes such as the Enviropro EP100GL to determine both water level and sediment build-up, however results are currently inconclusive and while this was considered for the project it was decided that it was not yet a sufficiently proven method to be worth attention. This may well change however. There are also a number of other remote sensing options for pit monitoring, including the use of ultrasonic sensors. A review of options and basic trial of one or more promising approaches may add a great deal to our toolbox of low-cost stormwater monitoring. This may of particular relevance to the use case, as blocked or backed-up drainage pits tend to be the primary cause of localised flooding in a stormwater system.

iii. Trunk monitoring with additional ultrasonic sensors

Stormwater exists the development site at specific locations, via large drainage trunks that gather run-off from a network of smaller drains. These trunks have access shafts[8] where an ultrasonic sensor might be deployed to measure height of flow. While such a deployment was considered for this project, it was not pursued due to anticipated technical difficulties. These difficulties are likely to include:

- Poor communications from beneath a steel manhole cover. The likely solution would require a LoRaWAN node mounted above ground with a sensor mounted in the shaft beneath, on a length of cable. The Elsys/Maxbotix device as sold does not support this configuration. It is likely that other communications technologies (e.g. 3G) would run into similar difficulties.

- Erroneous data resulting from high echo. This is an issue already experienced with Elsys/Maxbotix deployments over the larger stormwater outflow pipes. The ultrasonic sensor can sometimes return a measured distance that is far in excess of the possible distance that can be measured. In such instances, the reading is ignored as an anomaly. The explanation for such a reading is that the ultrasonic pulse echoes around the inside of the pipe before being detected. It is possible that such an effect might be amplified in a closed shaft and this should be considered when the device is

[8] E.g. at coordinates -33.810916, 151.071254

chosen for this sort of application. A tighter cone of detection and larger sampling rate may be solutions to this issue but would likely need to be actively explored.

- Trunk monitoring at the point where water exits the site would be a very useful addition to the overall stormwater monitoring effort as it could be compared with downstream outflow and used to interpret the relative proportion run-off from site compared with surrounding residential and industrial areas. For example, the major trunk exiting the northern site catchment to Wharf Road flows across to the golf course, where an Elsys/Maxbotix sensor is located. By comparing flow at the golf course with flow exiting site, we might expect to gain a better understanding of how much impact the site has on the golf course, as opposed to run-off from the surrounding residential streets.

Focused study of stormwater retention and run-off in the southern site catchment

Water quality data collected to date can be understood as a baseline of normal seasonal variation, typical for Ermington Bay. The interplay between pollutant influx and ecological and hydrological processes is influenced by the broader mix of suburban and light industrial land use but we can essentially rule out impacts from the Melrose Park development at this time. During the project, no active construction was occurring in this part of the site (with the exception of some small amount of earth moving in the south west). Stormwater drained from this area enters Ermington Bay through a large stormwater outlet monitored by ultrasonic sensor ELS006, which for now is capturing baseline outflow data relating to the existing system and land use. As the southern end of the site moves into an active construction phase, this will change, because stormwater leaving the site and entering the bay will be influenced by construction activity for the first time. Future study might focus on the southern site catchment as demolition and construction begins on the Melrose Park 'town centre. 'We should be able to compare changes observed in water quality to the 2020 baseline. As development expands we can expect two things to change.

Firstly, the drainage system serving the site (north and south) is likely to change. This will alter not only the amount of water retained on site and the speed at which stormwater exits the site, but may also shift the north-south catchment divide, potentially altering overall flows out of both catchments. We can expect such changes to create noticeable signatures in downstream data, from stormwater outflow sensors and from the EXO2 buoy.

Secondly, demolition and construction activity in the southern catchment will create erosion risk and possible sediment loss. Apart from the impacts of rainfall on disturbed ground, demolition and construction activities can involve the active use of water for cleaning equipment and concrete and for dust control, all of which generates potential for contaminated run-off.

With such changes in mind, it is recommended that a future study explores data from the Ermington Bay outflow sensors (003, 005 and 006) and the EXO2 buoy following the commencement of active development work in the southern site catchment. Data available since April 2020 will provide a baseline, against which changes can be analysed. In order to undertake this work most effectively it will be necessary to obtain hydraulic plans for as many of the private light industrial properties that lie between the site and the river as possible. The area currently represents something of a blank spot in our understanding. It may also be worth engaging actively with tenants and land owners in the light industrial zone and ideally, identifying if there are any accessible locations where additional ultrasonic sensors might be deployed.

Focused study of water quality in Ermington Bay

Undertake a deeper exploration of water quality in Ermington Bay and the possible impacts of the construction site. As noted above, the profile of a specific estuarine location is generally quite unique and what constitutes ‘normal’ or ‘problematic’ will be quite locally contextual. It is necessary to explore larger data sets to establish a more reliable baseline for the Bay. A more detailed analysis of stormwater outflow into the Bay and correlation with buoy data and rainfall data may yield deeper insights into the relative impact of stormwater from site and from the surrounding suburbs and industrial areas on water quality. This work should ideally go hand in hand with increased focus on the southern site catchment described above.

Beyond this, it is recommended that lab analysis of water samples is undertaken and compared with water quality data at the buoy. Samples should be collected at stormwater outflows to the Bay and ideally at trunk access points where stormwater exits the development site. The traditional way of collecting samples is by hand, however this provides data for only a limited snapshot in time and will tend to miss large rainfall and stormwater outflow events. Therefore it would be necessary to deploy a temporary automated sample collection system. This can be set up to collect a series of samples triggered by high rainfall and store them in a refrigerated compartment until they are collected for analysis. Such equipment is used widely by water authorities and should be available to loan. Alternatively, a third party may be contracted to undertake the work.

Image: A researcher checking an automatic water sampler. These compact devices can be battery powered and are able to detect high-flow events in order to capture samples at optimal times (credit: Endress+Hauser)

Analysis of water samples would allow the following angles of inquiry:

Sediment load can be compared with turbidity data to establish the relationship between inputs to the Bay from Melrose Park and general water quality in the Bay. This is helpful because it would allow us to more clearly determine the influence of upstream and downstream factors on the Bay and distinguish the role of local inputs, such as the development site. It would also enable us to explore the impact of stormwater from site versus non-site locations, given that some stormwater outflows (e.g. ELS003) do not serve the site at all, but rather the surrounding residential or industrial areas.

The upper Parramatta River, including Ermington Bay, is known to have sediments contaminated with relatively high levels of heavy metals and toxic persistent organic compounds[9]. This contamination originates on land and is the product of industrial activity in the area throughout the 20th Century. Contaminated soils are a widespread problem in the area and any disturbance of these soils can result in contaminants entering stormwater and polluting the Parramatta River. Naturally, the disturbance associated with a thirty hectare construction site should be of concern and is one reason why stormwater and erosion control measures are so important. Parts of the site with known contamination have been remediated, however this does not rule out lower level contamination across the wider site. There are no smart sensors that allow real-time in-situ monitoring of chemical pollutants in water and there are certainly none onboard the EXO2. The parameters that we do measure at the buoy are physical attributes of water quality, detectable using methods like light scatter and absorbance. To detect chemical signatures, a chemical reaction needs to occur and this can only really be done in a lab. Following lab analysis of water samples associated with high stormwater outflow we might seek to establish if there are any direct relationships between physical attributes of the water that can be measured by the buoy (e.g. turbidity) and levels of chemical pollutants. If such a relationship does exist it may be possible to use the physical parameter as a rough proxy for the chemical contaminant. This may then allow EXO2 data to be interpreted as a low accuracy proxy for chemical pollution.

Build an understanding of how different data sources relate and explore machine learning to model stormwater flows and interdependencies at Melrose Park

Preliminary exploration of stormwater data has tended to review the different types of data separately. There is a great deal of scope for building an understanding of how the data recorded at different levels of the stormwater system interacts. Ultimately, we should look at ways to build working models that are adapted to the specific context of Melrose Park Such models might be used to estimate risk, plan mitigation efforts or even forecast future events.

Part of the challenge with developing such models is that much of the data we are gathering is a-typical for hydrological study. By experimenting with low-cost sensors we are exploring what is possible with limited resources and information, working with constraints and generating novel new types of data that require new approaches for interpretation and analysis. Much of our data is highly relational, meaning that it can only really be made sense of as a trend line against its own long-term baseline, and against the trends from neighbouring sensors. Models based upon such data may need to be quite innovative, departing somewhat from more standard methods of hydrological research.

One starting point for exploring the relationships between different data sets is to undertake a series of correlation analyses. Water flows through a system and the effects of that flow should be detectable at each level. The following comparisons might be explored:

- Compare all six ultrasonic sensors to each other, to rainfall intensity and to rainfall accumulation data. This can help to tell us the relationship between flow measured at any two points in the system, and the relationship between rainfall and flow at each outflow. The aim might be to develop outflow models that are calibrated with twelve month data sets, allowing us to understand how height of flow tends to vary against various rainfall characteristics. This relationship is complicated by the fact that factors other than rainfall impact outflow.

- Compare the two Decagon 10HS soil moisture probes being used to observe overland flow with each other and with rainfall intensity and rainfall accumulation data. This can help to

[9] Parramatta City Council (2016), Strategic Analysis of Water Quality in the Parramatta River Technical Analysis Report

confirm that we are detecting rainfall events and explore the critical duration and intensity thresholds associated with detectable flow in each location.

- Compare soil moisture data from the two multi-depth soil moisture probes with rainfall data and with ultrasonic data for the primary catchment outflows (ELS004 and ELS006). This can help to tell us the degree to which soil saturation has a significant effect on run-off.

Ultimately, given the complexity of the system and the many relationships between data sets, it may be necessary to explore the options afforded by machine learning. One approach would be to use deep learning to discover relationships between all of our data sets (and we may look at including urban heat and air quality). The other would be to provide more of an initial structure by building a virtual hydrological model for Melrose Park using known parameters, then exploring data sets with AI within that structure.

Engage with site managers and explore options for identifying and logging flooding and erosion issues as they occur, as well as the nature and location of flood and erosion mitigation efforts

We are currently somewhat blind to actual day to day realities of stormwater on site, from an operational perspective. The data we collect quantifies the presence of water in various locations and we can infer things from this. However, if we can build a record of actual events such as localised flooding or pooling of surface water then we can explore sensor data in more targeted and informed ways and ensure that we refine our enquiry so that we address real operational challenges. To do this, we may consider asking site managers to log where and when stormwater issues occur.

Image: The stormwater retention pit on site (credit: Andrew Tovey).

5 Practical lessons for designing and delivering low-cost environmental sensor networks

Practical lessons for designing and delivering low-cost environmental sensor networks

- Collaboration and successful project delivery
- Device deployment choices and the iterative nature of smart city experimentation
- Technical design

Collaboration and successful project delivery

1. The importance of engaged and enthusiastic project partners

A collaborative project demands continuous active engagement by all stakeholders. If this is achieved then collaboration tends to succeed and produce outcomes that serve the needs of those stakeholders. Central to such active engagement is the importance of enthusiasm for the project and its success, grounded in a shared vision for what might be achieved. Perceived value may vary by stakeholder but the means to achieve it must be held in common consensus.

The three major partners for this project were the City of Parramatta Council, PAYCE and UTS, all of whom were actively and engaged and enthusiastic throughout:

- City of Parramatta Council showed early commitment to a broader smart city agenda, as well as to environmental liveability, with multiple past and present projects exploring urban heat, air quality and river health. From the outset, Council approached the Melrose Park project with an openness and appetite for experimentation and innovation, through deep and genuine collaboration. Council's aims are ultimately for improved strategy, reporting and governance, leading to improved operational capability and improved customer experience of the City.

- PAYCE were committed to developing their engagement with smart city technology and to recognise the potential such technology could hold for transformation of their business and the industrial sector more generally. They wished to develop expertise and leadership in smart cities and are comfortable with the idea of open and iterative experimentation, realising that this is the reality of work at the leading edge of new technology adoption.

- UTS is building a collective smart city research practice leveraging success with other Round 1 and Round 2 Smart City grant projects. A number of research clusters across multiple faculties were actively engaged with critical elements of the practice, from the harder engineering considerations, through to environmental science, urban design, social wellbeing, public health and policy development. A project like this was an opportunity to build an emerging smart city research community through collaborative engagement with real-world problems, end users and practical realities. UTS realises that smart city research cannot occur in an academic vacuum and that it is inherently transdisciplinary and collaborative in nature.

Executive buy-in is a critical foundation for this sort of collaborative engagement. The original Council project bid was led by the Head of Future City, aligning the project with Council's broader smart city strategy and ensuring top-down support. Furthermore, it laid the foundations for discussing how the legacy systems of the project will be taken forward into the future.

2. The importance of internal champions to the success of pilots

Council, UTS and M Projects (on behalf of PAYCE) assigned dedicated staff to the project, forming a core collaborative team. All of these individuals showed a strong personal engagement with the project and support for what it sought to achieve and may be thought of as internal champions. The dedication of these internal champions to the ideas and aims of the project was likely critical to the success of this project.

3. The importance of a Business Requirements Document (BRD)

The first task undertaken by the project team following the launch of the project was to develop a consensus about the business requirements of the project and record this in a single shared reference document. The BRD defined the aims of the project and the problems that it sought to address. It confirmed the project scope, the opportunities and benefits, as well as the challenges and risks. It identified agreed measures of success against clear milestones and it established a clear governance structure. The BRD was vital for exploring the assumptions, expectations and agendas of project partners from the get go. As the project developed, the design and nature of deliverables adapted agilely as needed, while remaining focused on core requirements of the BRD. The project was completed on time, within budget and successfully met the expectations of all partners – an outcome that was almost certainly contingent upon the early work undertaken on the BRD.

4. The importance of creating time and space for people to learn about new technologies, concepts and vocabularies

It takes time for people to grasp new technologies and applications. This time is a vital part of the smart city journey that ensures partners take full ownership of solutions. The design and implementation of the technical solutions delivered by the project required familiarity with concepts and vocabularies that were new to many of the project team. Over the course of the project, individuals built their knowledge and understanding of the new technologies. However, this meant that progress, particularly during the earlier design stages, was slower than originally expected. Ultimately, if staff across the various partner organisation leave the project with a firm working knowledge of new systems then they are best placed to ensure the continuation and optimisation of those systems. Too many projects fizzle out after the official delivery period. If the goal is to create a lasting endeavour, that grows to produce a range of future benefits, then it is vital to pursue an inclusive project delivery model.

5. The importance (and challenge) of iterative experimentation

The project has been a highly experimental foray into the use of new technologies and how they might be applied within the context of major residential developments like Melrose Park. Design and delivery has been somewhat iterative and many activities (such as stormwater monitoring) have been undertaken without full understanding of the outcomes. Technical delays and hurdles appeared where they were not anticipated. Many of the specific technical milestones provided to the grant application were based on early technical assumptions that became obsolete by the time the system was implemented and deployed. All partners learnt a great deal as a result of the project and the open experimental approach has been critical to this learning. The iterative approach is now being recognised as valuable and is increasingly being incorporated into smart city strategies around the world.

6. The importance of procuring devices early

A major constraint on project outcomes was the limited amount of data available for analysis. This was a result of many devices only being fully commissioned quite late in the project (some as late as April 2020). To avoid this occurring in future projects, efforts should be made to procure devices as early as possible.

"

Do not underestimate the timeframes required for design and delivery of experimental sensor networks and aim to procure new technologies as early as possible in order to allow adequate time for data collection.

Experience of the design and delivery process for this project has highlighted quite how long each stage can take. The process of developing the BRD took several months; time well spent on a vital outcome, but certainly a more complex process than was anticipated. Development of detailed plans for the four use cases followed, culminating in device procurement decisions and orders by September 2019. It was found that the lead time on device shipping extended well beyond what was anticipated. Notably, this was prior to the Covid-19 crisis, which has since proven how unexpected events can significantly delay procurement schedules.

Once received, the process of commissioning a device to a point where data is being reliably and accurately captured can take a significant amount of time, as experienced in this project. Devices can have technical difficulties, sometimes requiring months of work and returns to the supplier for recalibration before they function as required. Even once a device is functioning correctly, the challenge of interpreting its data remains, particularly in cases where a device is being used in a novel new way. For example, the interpretation of data from ultrasonic sensors used to monitor the flow of stormwater relied upon contextual information and custom calculations that were unique to each specific device deployment. This required custom coding in the data management platform and until that coding was developed, tested and fine-tuned, no data could be captured. A final hurdle is storage of data. You can have devices deployed, with data received and correctly interpreted, but if the data is not being stored in a reliable and accessible format then this counts for nothing.

Experience of these processes, and the often unavoidable delays and complications that accompany experimentation with new technologies and applications, has emphasised what we believe to be an important learning for future projects. Do not underestimate the timeframes required for design and delivery of experimental sensor networks and aim to procure new technologies as early as possible in order to allow adequate time for data collection.

7. The importance of ‘boots on the ground’

From detailed study design to successful deployment of devices and informed interpretation of initial data, having one or more project team members with ‘boots on the ground’ has proven to be vital to project success. In this project, this role was undertaken by one of the project leads at UTS, who worked closely with all partners on a daily basis throughout the project and spent time on site on a roughly bi-weekly basis throughout the project period.

Detailed use case design requires a responsiveness to place that can only be achieved through direct physical exploration. Locations for particulate and noise sensors were chosen based upon direct observation of construction activity on multiple occasions, combined with a working knowledge of how the site was scheduled to be developed over the timeframe of the project. Stormwater monitoring was planned and fine-tuned based upon observations of the site during heavy rainfall, a necessarily agile and opportunistic design approach.

Once detailed plans are made, they often need to be adapted in response to changing circumstances or unforeseen practical constraints. A proposed mounting location might become inaccessible due to changed construction schedules. Another location might turn out to have unviable communications coverage. Someone needs to make informed decisions when these issues arise and a person with close active working knowledge of the site is best placed to make them. With a large and complex project comprising over 60 device deployments, nine device types and four separate use cases, such agile decision-making was vital to the timely and successful delivery of the project.

Once devices are deployed, the need for a hands-on role continues. Devices can fail due to water ingress or faulty batteries, requiring devices or components to be retrieved from the field, returned to the manufacturer and replacements installed. Even functional devices tend not to work ‘out the box’ as needed, requiring calibration to their specific context. For example, each ultrasonic sensor was installed above a stormwater outflow pipe. To calibrate, it was necessary to determine the precise distance between the sensor and the base of the pipe, which equalled zero outflow. Determining this distance required a comparison of data output to direct observations on multiple occasions. For some devices deployed for the project, if we combine visits for deployment planning, initial deployment, troubleshooting and calibration, five or six visits were needed before they were functioning correctly.

A final critical benefit of having had ‘boots on the ground’ throughout the project is the ability to interpret data based on accumulated site-specific knowledge and long-term observations of site activity. Here is a brief example:

Speed humps create a noise trap

Notably high peak noise readings were recorded by EMS device 007, located on the heavy vehicle access road on the southern end of the site. Very large earth-moving trucks pass by the sensor on a regular basis and are likely to be responsible for peak noise events. These trucks also pass two other sensors nearby, but noise levels at EMS007 are highest. Why might that be?

Either side of this sensor there is a speed hump for traffic calming, causing trucks to slow (applying compression brakes) and then accelerate (revving their engine), immediately adjacent to the sensor. Only through a detailed knowledge of the location, grounded in multiple site visits, is it possible to interpret what is occurring in this instance. The inherent messiness of the urban landscape means that accurate interpretation of environmental phenomena commonly requires this sort of highly contextual understanding of place.

The importance of direct relationships with land managers

The land manager for the Melrose Park site is M Projects, who are contracted by PAYCE and were an official project partner. M Projects acted as a proxy stand-in for PAYCE interests throughout much of the project, particularly once high-level design was captured in the BRD. M Projects dedicated a staff member to an actively engaged role on the project team. When agile decisions needed to be made, M Projects were able to assess, advise and approve a swift course of action. Site-based work by UTS also benefited greatly from this direct relationship. The UTS project lead for on-site work built personal rapport with a number of key M Projects personnel, through regular presence on site and face-to-face engagement, further supporting agile decision-making.

> **From detailed study design to successful deployment of devices and informed interpretation of initial data, having one or more project team members with 'boots on the ground' has proven to be vital to project success.**

Image: Installing a stormwater sensor (credit: Andrew Tovey)

Device deployment choices and the iterative nature of smart city experimentation

This project can be viewed as a complex series of interconnected investigations using experimental new technology, on a modest budget. The work was inherently iterative. Decisions were made about device deployments and study design made on best available information at the time. An ethos of 'just getting on with it' and seeing how things go gave us the freedom to rapidly test ideas, discarding what didn't work and evolving what did. By the end of the project it was clear that some things should be done differently, yet the value in what we did do remains clear. An area that perhaps best illustrates this approach is the way in which we chose to deploy a limited number of air quality and noise sensors to support competing priorities.

The greatest limitation of the study design for air quality and noise was the number of available EMS devices and the constraint this placed upon deployment density. Working with just fifteen devices made it difficult to provide both widespread coverage of the study area, as well as focused coverage of specific areas of interest. The trade-off between these two goals may have resulted in a reduced capacity to test operational hypotheses in focused locations and so we might reasonably ask if we could have made different choices. However, if we consider our alternatives, things are not so simple. The cost of an extra four EMS devices at locations of greatest interest would amount to more than the entire budget for thirty seven urban heat monitors, so a reassignment of some device budget from, say, urban heat, would not have been a practical alternative. What might have been done is to reduce the number of EMS devices deployed in certain other locations, however this is not a clearly better alternative and would certainly be a trade-off; devices in locations like Atkins Road and Cobham Avenue provide valuable baseline data about the air quality and noise in residential streets set back from the site, allowing Council and PAYCE to understand the levels and trends that residents experience. Reallocating some of these devices would reduce this understanding.

Ultimately, the EMS deployment choices were made when it was not clear where locations of high interest might emerge. A more generalist deployment pattern was chosen to 'cover all bases' and develop a general overview of air quality and noise trends in the area. Only once data began to arrive did it become clear where the 'hot spots' were, and this in turn indicated locations where denser sensor deployment would be justified.

An additional factor in this process was direct observation of the site by researchers over more than a year of project delivery. A large construction site is complex and ever-changing and it takes a good grasp of operations and on-the-ground observation of site infrastructure and physical events like large vehicle movements and excavations, in multiple locations and under various weather conditions, to build a strong working knowledge of how noise and dust are formed and travel. The working knowledge accrued by the end of the project period was significantly greater than was held at the time when device deployment decisions were first made and in some instances, with hindsight, some devices may have been more effectively located in different locations.

Our evolving understanding of the site and how to most effectively use the technologies at our disposal to answer critical questions emphasises the inescapably iterative and experimental nature of the project. In order to confirm certain areas of interest, we first needed to take a general approach. Only through trial and error did patterns and insights emerge. These point towards more focused and expanded research in possible future phases of this work. Experimental smart city projects, in keeping with the living lab model and action research methodologies, are inherently open-ended and always point towards new areas of inquiry. Much was learned over the course of the project and some of this was how *not* to do things, or where *not* to deploy sensors. Such insights define a journey of ongoing exploration where the location, the use cases and the technologies form a complex and highly contextual system.

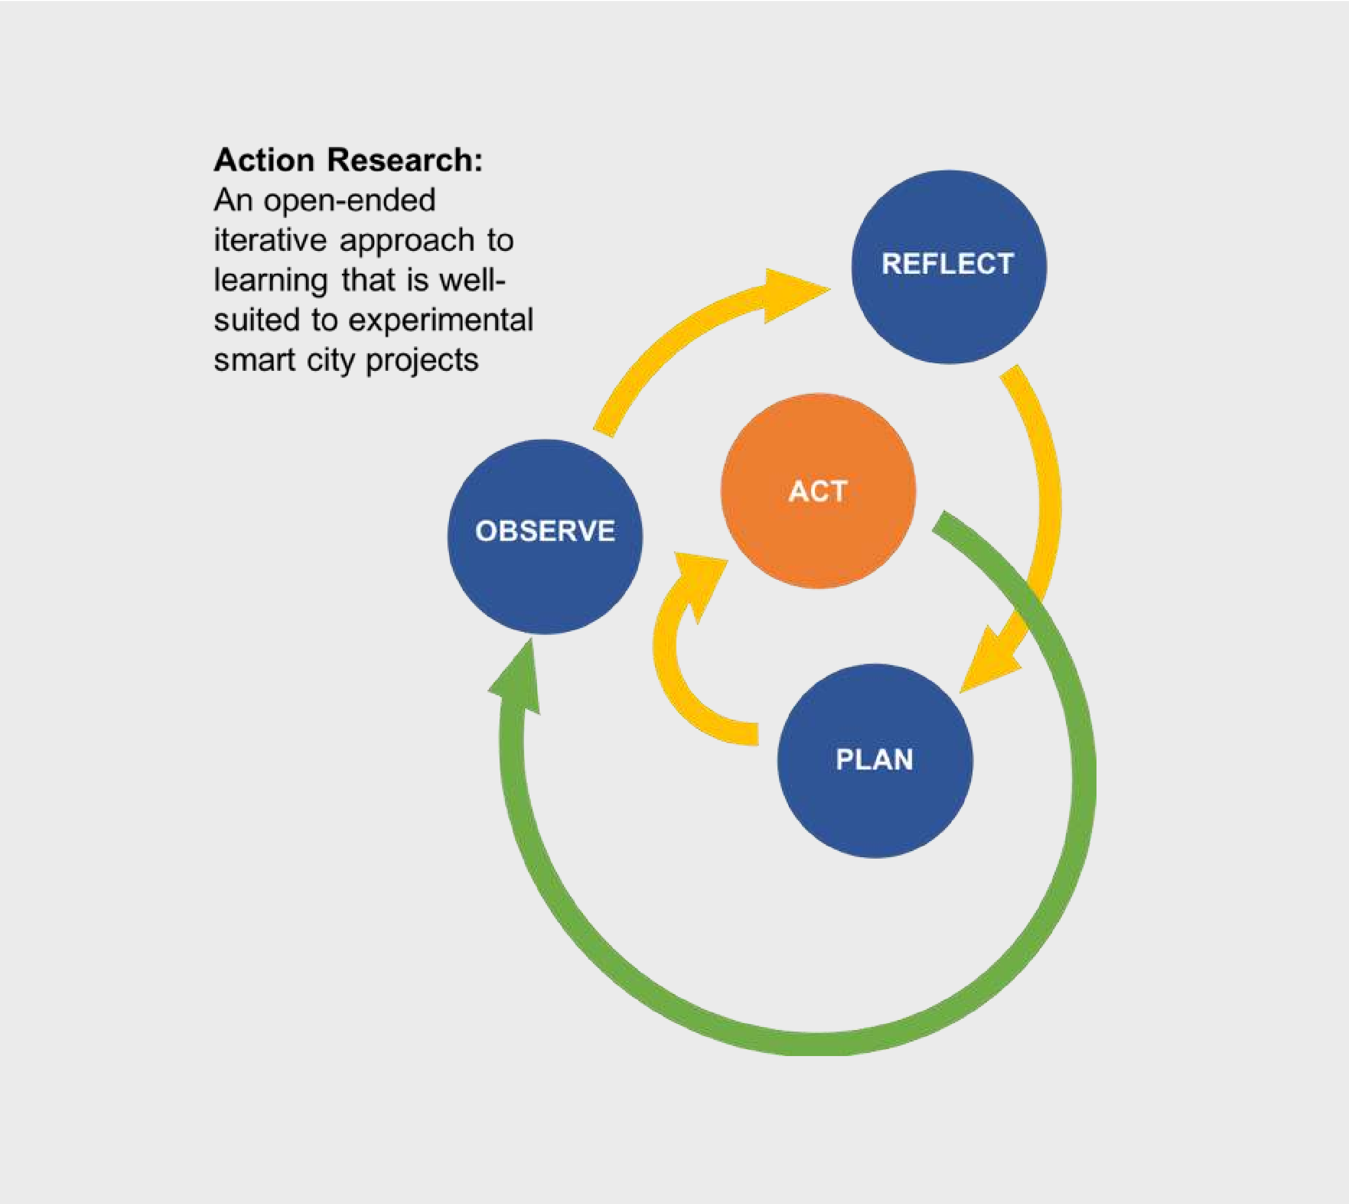
Action Research:
An open-ended iterative approach to learning that is well-suited to experimental smart city projects
REFLECT
ACT
OBSERVE
PLAN

Technical design

The importance of strong in-house technical capacity

UTS was the technical and research lead for the project. While study design and production was managed by the Institute for Sustainable Futures (ISF), the technical needs of the project were met by UTS Rapido, an advanced technology development unit within the Faculty of Engineering and IT (FEIT). Rapido brought considerable technical expertise and direct delivery capacity to the project. Their involvement was critical for integration of multiple data sources into a single centralised data aggregation and management system (the TULIP platform). Rapido were able to work directly with devices at a technical level and were able to design and implement code for novel data interpretation. As part of the project, Rapido developed a 'data lake' for storage of data from all sources, including our own deployed devices as well as data from third party sources like the Bureau of Meteorology (BOM) and the NSW Department of Planning Infrastructure and Energy (DPIE). They also developed a centralised metadata management database for all devices, providing a single place to store and edit metadata, accessible by multiple end users (notably the [ui!] and ESRI dashboards). This supported the integration of new sources of data in days instead of weeks.

The importance of flexible and adaptable data architecture and an agile engineering process

The UTS technical team applied the architectural principles of flexibility and adaptability to any data source or sink, combined with an agile work process. This allowed the engineers to adapt to unforeseen changes, various technical issues and other "surprises" inherent to experimental work.

In particular, the decision to forego a defined end-to-end data model in favour of a scoped API-specific format enabled the project to:

- Make sure that a necessary change in part of the system would not propagate to the rest of the project.
- Enable the asynchronous integration of the multitude of data sources and sinks in a relatively independent manner.
- Pursue a very short code-test-deploy loop, saving valuable integration time.

The importance of direct relationships with external technology providers

As technical lead for the project, UTS has direct working relationships with most of the commercial technology providers for the project. This proved vital for a pilot project of this nature, where inherently experimental work requires non-standard deployments, configurations and integrations of technology. Notable relationships include The ARCS Group and Meshed, both of which are Australian SMEs who worked on a previous *Smart Cities and Suburbs* project with UTS. The ARCS Group co-developed the TULIP EMS device. UTS was able to continue a working relationship with The ARCS Group throughout the Melrose Park project to refine the functionality of the EMS and the associated Lufft weather station. Meshed provided the LoRaWAN communication for the project but beyond this they were also the supplier for over 80% of devices procured. A close working relationship with Meshed allowed UTS to plan and execute experimental device applications, with strong external technical support. Other hardware providers for the project were Australian companies, notably Xylem and ICT International. UTS was able to develop direct connections with technical staff at both companies and this proved to be invaluable for configuration and troubleshooting. Working with Australian companies also turned out to be critical in the final months of the project as the Covid-19 crisis took hold, as an emphasis on local suppliers meant that there were no disruptions to shipping.

References

Bibri, S.E. and Krogstie, J., 2017. Smart sustainable cities of the future: An extensive interdisciplinary literature review. *Sustainable cities and society*, *31*, pp.183-212.

Climate Council, 2017. *Intense Rainfall And Flooding: The Influence Of Climate Change*. [online] Climate Council. Available at: <https://www.climatecouncil.org.au/uploads/5dafe61d7b3f68d156abd97603d67075.pdf> [Accessed 27 August 2020].

Desa, U.N., 2016. Transforming our world: The 2030 agenda for sustainable development

Parramatta City Council (2016), Strategic Analysis of Water Quality in the Parramatta River Technical Analysis Report

Steffen, W., Hunter, J. and Hughes, L. (2014). *Counting the costs: Climate change and coastal flooding*. [online] Climate Council. Available at: https://www.climatecouncil.org.au/uploads/coastalflooding.pdf [Accessed 7 Mar. 2020].

United Nations (2015). World urbanization prospects. The 2014 revision. New York: Department of Economic and Social Affairs, http://esa.un.org/unpd/wup/Publications/Files/WUP2014-Report.pdf (accessed 22.1.2017).

Watkins, T., 2013. Draft roadmap for next generation air monitoring. *Environmental Protection Agency*, *2*. [online] US EPA. Available at: https://www.epa.gov/sites/production/files/2014-09/documents/roadmap-20130308.pdf

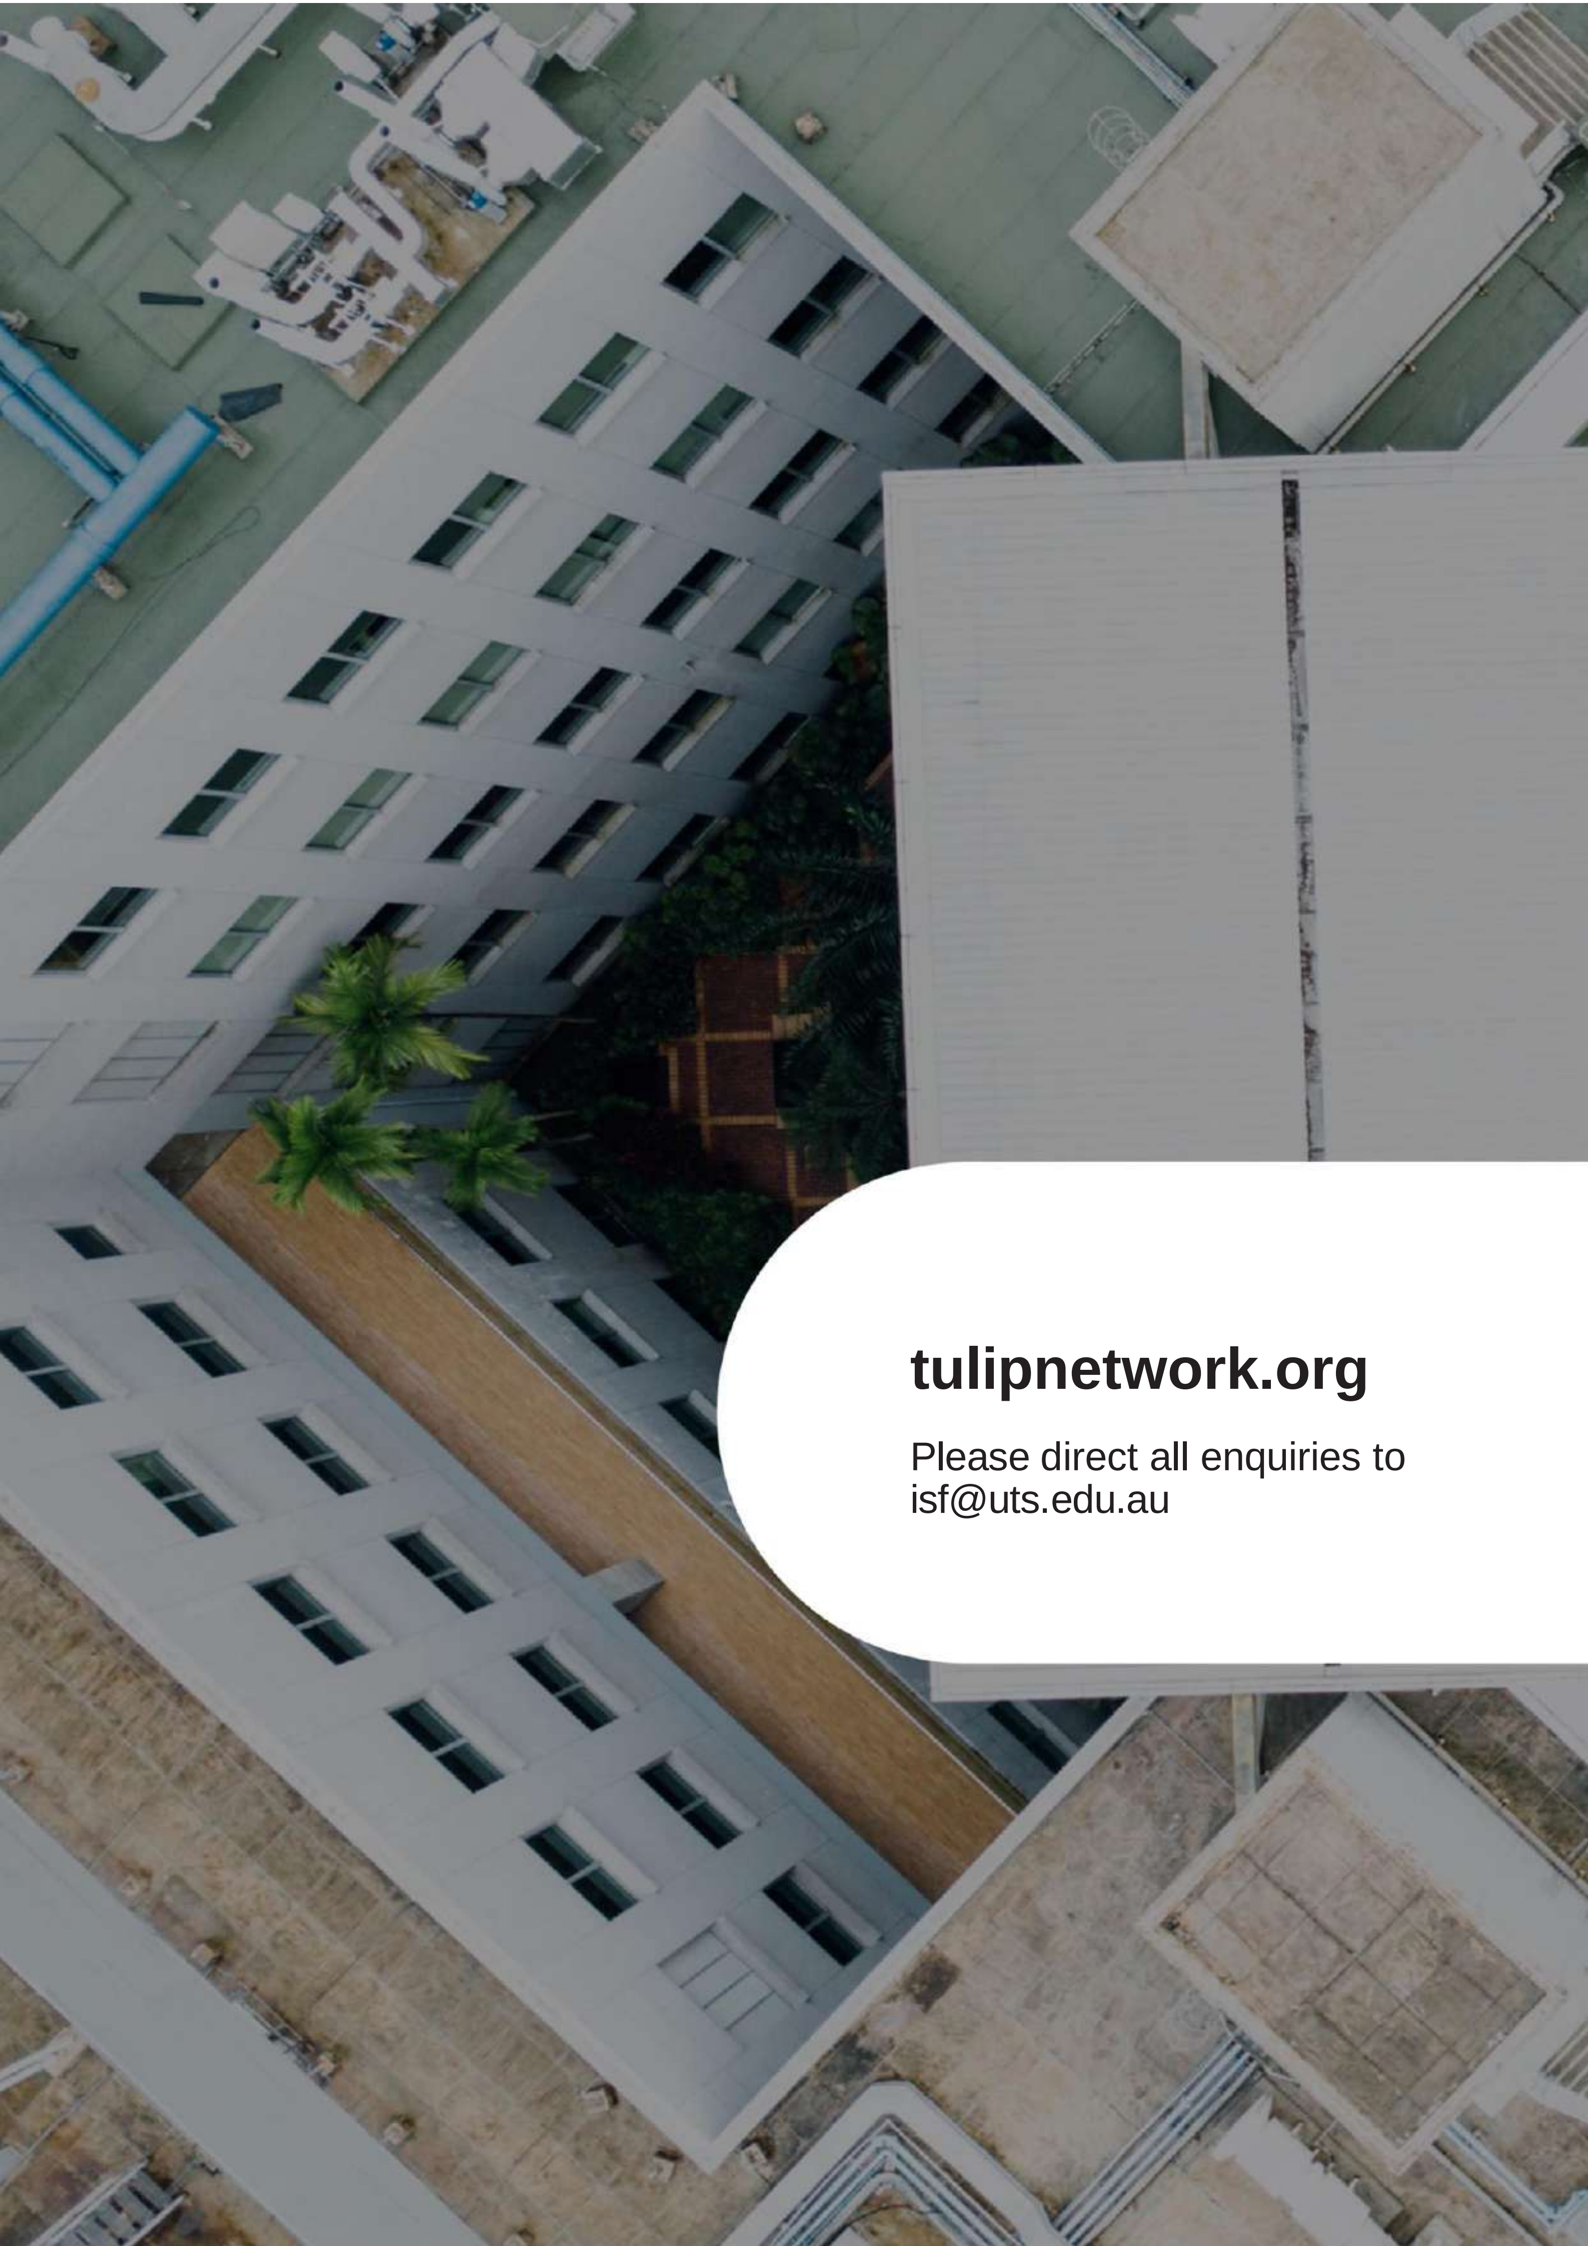

tulipnetwork.org

Please direct all enquiries to
isf@uts.edu.au