Questions	Answers
What's the impact of HPC for this study?	We use HPC and cloud resources to run the various models we have. As mentioned during the presentation, we do not necessarily aim at simulating a single scenario as quickly as possible (although this is sometimes needed in order to meet deadlines) but instead simulate several scenarios. In future, with the Digital Urban Climate Twin, there will be many more scenarios that have to be simulated. Using HPC resources is thus very important for us.
How do we verify the simulation results? How large is the problem size for this research?	Model validation is done by comparing the outputs of your simulation with data collected from the real world. We use various data sources for this purpose. For example, this includes sensor data and data provided by various government agencies. The problem size (or domain size) for our research is the whole city (for mesoscale studies) or districts (for microscale studies).
The UHIs in Singapore appear to be strongly spatially correlated to density (of housing, of industry). After all, would the best mitigation strategy mean to dissolve density? Or put it on another scale: Might Singapore's density be a huge disadvantage in future?	'Dissolving density' could be one solution, but not the only one and for sure cannot guarantee the best results. If dissolving density means Singapore will extend the urbanised areas, then the UHI phenomena could improve in some areas but worsen in others. Urban temperature in a district does not only depend on its specific development characteristics, but also its surroundings. Increasing population and extending urbanised areas will increase Singapore's UHI. Implementation of novel strategies could reduce, at least partially, this impact. There is also the possibility to reduce the local UHI while maintaining density. For example, high thermal mass of concrete buildings is an important factor contributing to the UHI. Protecting building façades and roofs from the sun through electricity producing skins will be the most effective in the short and medium-term. In the long-term, every estate must be emission free and produce more energy than it uses.
Would we be able to get a copy of the slides?	Please contact the organisers.
More the question for an opinion: Do you think people will emigrate from Singapore and other tropical regions in future (2050 onwards)?	People often migrate in the hope to seek 'greener pastures' for themselves and their families elsewhere. Whether people migrate or not will probably not so much depend on the local climate but rather on social and economic factors and

	overall liveability in Singapore and the region. Climate change does pose a threat to liveability and may increase the likelihood of emigration if countries and cities (like Singapore) do not adapt in time.
Is there a demand for real-time usage of simulation results?	In a planning environment, if you have an iterative planning process, then it would be helpful to have simulation results available as quickly as possible. Ideally, you change the scenario and you get to see the implications (i.e., simulation results) in real-time. This would make for a very smooth workflow and would allow planners to evaluate different planning/design options in a short period of time. However, at the moment, evaluation of a single scenario can take days and weeks. Full-scale climatic simulations are unlikely to ever run that fast which is why we consider building so-called ‘surrogate models’ for quicker evaluation of scenarios.
Is there a most economic processors number?	In general, with parallel computing, as soon as you increase the number of processors, you get overhead (which results in sublinear speedup). The most economic number of processors would thus be the number of processors at which the overhead is minimal. This may very well mean N=1, i.e., no parallelisation. However, in practice, there are many factors (e.g., estimated run time, cost, ...) and technical constraints (e.g., problem size, hardware constraints, ...) that need to be considered. Without knowing the exact use-case and parallel computing problem at hand, it is difficult to answer this question.
Have any measures which are based on those simulations already been implemented or are we still in the research/setup stage?	We have interesting results for cool roofs, for electric vehicles and district cooling implementation. These are results obtained only recently. We will share and discuss them with the government agencies. It will then be their turn to deliberate and decide which measures and design policies to implement.
Are the outputs mainly meant for policymakers, or are there plans to release some datasets to the public?	Our work aims primarily at urban planners and policy makers - as well as researchers. Data that we collect and own can be released to the public at some point. However, a lot of the data we use for modelling and simulation we do not own. In this case, the decision whether or not to make this data public lies with the respective data owner.

1) What kind of surrogate models do you use in the digital twin? Are they data/machine learning models or only-physics based? 2) Tight/loose coupling refers to model execution or also involves model design?	1) Machine learning models (particularly deep learning models) are being considered. The data for training such models may come from the various simulations we perform using full-scale urban climate models. We have not built this kind of surrogate models yet. 2) Tight/loose coupling refers to model execution.
How long is one simulation run, in general? WRF simulation is the memory-bound workload. The scalability of WRF shall be quite good for the modern system architecture. Do you see any hotspot on the "multi-physics" part?	With our current setup, a single simulation with all the anthropogenic heat load included, using 10 nodes takes 6 days in order to simulate a period of one month. WRF scales well on modern Cray systems. One important aspect - which also contributes to the communication overhead - is the data exchange between the Building Energy Model (BEM) and the atmospheric part of the WRF model. This is also an example of tight coupling which is done internally by WRF.
Since you are running the simulations for the entire Singapore, you are dealing with a huge computational domain. Could you give an idea of typical computational times/resources to run e.g. a CFD simulation for SG? Or are the simulations being run for different districts separately?	In general, running CFD simulations should be considered carefully. What is the purpose? Are the results going to be trustable? What methodology is going to be used to run the CFD model? Different boundary conditions in different areas of the island? Additionally, should the calculations be done for one weather condition or different weather conditions representative of Singapore? Depending on the approach we will have different results that will provide different types of information. We use CFD simulations for district-scale simulations (simulating the entire city is done using mesoscale models such as WRF). Run times and resource utilisation of CFD simulations vary a lot, depending on the scenario and the runtime configuration. For example, with a domain size of 2.4 km x 2.4 km and approx. 17 million cells, we are able to simulate one hour worth of meteorological conditions (steady state configuration) in less than a day.
Currently, the ask and criteria for the “green building” is well developed, but the impact of these “green buildings” on the surroundings – spaces between buildings, public spaces and other buildings is still very limited. Could this research help and how to help with this, both for the people using the city and the better urban governance?	Yes, we will evaluate the impact of vegetation on and around buildings on people’s thermal comfort as well as on its impact on the UHI.
I wonder if the data is available to the public? Similar to 6 Sigma DCX, would there be project files that could be exported to be used by other stimulation models or packages?	See above for publication of data. We generally use standard data formats (e.g., GIS, GeoJSON, ...) that should be compatible with other simulation models and tools.

How much of this contributes to the ongoing urban planning for Singapore?	It will help to assess future master plans and potential what-if scenarios on how the city can be developed in a more climate-responsive way.
How is the current shift in urban life pattern of people ie. work from home lifestyle in response to COVID, impacting the urban heat challenge?	This is a good question. We do not currently have a conclusive answer. One hypothesis is that the tempo-spatial heat emission patterns have changed since many people are now working from home. However, the impact on the local climate (and the UHI in particular) is difficult to measure with the current network of weather stations in place. In addition, there is the issue of temporal dynamics. The circuit breaker, for example, was relatively short in duration (by climate standards). We are now also in a transitional phase in which society moves towards a ‘new normal’. Once such a new normal (e.g., more remote work) has been reached, it should be possible to take long-term measurements which can then be compared against pre-COVID measurements.
Which HPC and data centre operators are providing compute power for the above? Other than NSCC and Google Cloud	So far we have been using primarily in-house computing clusters and, for prototyping, also Google cloud resources. We have also used ETH Zurich HPC resources. In future we plan to work with NSCC and CSCS (Swiss National Supercomputing Centre).