South African Council for Automation and Control

Control Systems Day 2026: Programme

Friday, 13 November 2026

Time	Title	Speaker	Venue	Supervisors	Institution
08h30	Registration & Coffee/Tea				
09h00	Welcome Address	Tobi Louw	UCT	–	Stellenbosch University
09h10	Invited Speaker: Advanced Process Control in Practice	Albertus Richter	WITS	–	Stone Three
10h00	**Session 1 — Process Modelling & Digital Twins**				
10h00	Dynamic modelling and validation of an industrial steam system for model-based control	Andisiwe Mbadamana	WITS	Andries Wiid Derik le Roux Ian Craig	University of Pretoria
10h20	Digital Twin Technology for Control of Distillation Column Feed Variability	Carolyne Hugo	WITS	Antony Higginson Kevin Brooks	University of the Witwatersrand
10h40	Closed-Cycle Distillation Platform for Multivariable Control Education	Malhar Apte	WITS	Paul Sonnendecker Lehlogonolo Tabana	University of Pretoria
11h00	Dynamic modelling of green ammonia production under variable renewable energy supply	Abulele Nogaya	UCT	Tobi Louw Prathieka Naidoo	Stellenbosch University
11h20	Coffee/Tea Break				
11h40	**Session 2 — Data-Driven Modelling & Controller Tuning**				
11h40	How Much Physics Does a Hybrid Model Need? Neural ODEs for Nonlinear MPC, from Quiet Plant Data to Simulation-Trained Priors	Greg Laycock	WITS	Derik le Roux	University of Pretoria
12h00	SafeOpt for the Automated Tuning of MIMO PID Controllers	Emmanuel Aiyede	WITS	Derik le Roux Ian Craig	University of Pretoria
12h20	Time Series Foundation Models for Chemical Process Modelling	Imker Hoogenhout	UCT	Tobi Louw	Stellenbosch University
12h40	Lunch Break				
13h40	**Session 3 — Guidance, Navigation & Control of Aerial Vehicles**				
13h40	Observer-Based Disturbance Rejection for Quadcopter Attitude Control in Windy Environments	Mpilonhle Nxasana	UCT	Arnold Pretorius	University of Cape Town
14h00	Cascaded Linear Model Predictive Control for Agile Quadcopter Trajectory Tracking	Tony Mwangi	UCT	Arnold Pretorius Ed Boje	University of Cape Town
14h20	Development of an Autonomous Parafoil-Based Recovery System for Suborbital Rockets and High-Altitude Payloads	Joshua Spies	UCT	Willem Jordaan Jean Pitot	Stellenbosch University
14h40	Actuated Tether Proprioception for GNSS-Denied Localisation of a Marsupial UAV-UGV System	Guy Hasewinkel	UCT	Arnold Pretorius	University of Cape Town
15h00	Vote for Best Presentations	Antony Higginson	WITS	–	University of the Witwatersrand
15h20	Closing Address	Tobi Louw	UCT	–	Stellenbosch University

Venues		
UCT	EM 4 & EM 5, Electrical & Mechanical Engineering Building, Library Road, Rondebosch, Cape Town	View map
WITS	Richard Ward Building, Room 333, 1 Jan Smuts Avenue, Braamfontein, Johannesburg	View map

Please note: online attendance will not be an option this year.

Invited Speaker

Advanced Process Control in Practice

Presenter: Albertus Richter (Stone Three)

APC engineers are required to be all-rounders, needing working knowledge of PLCs, instrumentation, mechanical components, computer networks, OPC servers and, finally, the advanced algorithms. In this talk I will walk through what an APC engineer does on a project: analysing plant data to see what each site requires, integrating with the plant control system, commissioning on site, and maintaining the solution together with site personnel.

Short bio: Albertus Richter is an APC engineer at Stone Three, with 5 years of experience. He graduated in 2021 from the University of Pretoria with a bachelor's degree in electronic engineering then went on to specialise in control systems obtaining his master's degree, also at Pretoria. His research focussed on automatic tuning of level controllers in a flotation bank using Bayesian Optimisation. Since then, he has worked as both a Process Control Engineer and Software Engineer in the Mineral Processing field developing, installing and maintaining Advanced Process Control systems.

Session 1: Process Modelling & Digital Twins

Dynamic modelling and validation of an industrial steam system for model-based control

Presenter: Andisiwe Mbadamana (University of Pretoria; Sasol Projects and Engineering)
Supervisors: Andries Wiid, Derik le Roux, Ian Craig

This study presents an integrated nonlinear dynamic model of an operating industrial steam system comprising boilers coupled to a steam header network. Established boiler and pipeline modelling concepts are adapted to the plant configuration, with the distributed pipeline equations solved separately using a Galerkin spectral element method and an implicit finite-difference method. Validation against 50 h of plant data shows that both formulations reproduce the dominant pressure and steam flow dynamics.

Neither formulation demonstrates a consistent advantage in validation accuracy, while the implicit finite-difference formulation is computationally more efficient for the system and numerical settings considered. The validated framework provides a basis for future studies of constrained multivariable control of boiler loading and header pressure, as well as economic optimization of the steam system.

Digital Twin Technology for Control of Distillation Column Feed Variability

Presenter: Carolyne Hugo (University of the Witwatersrand)
Supervisors: Antony Higginson, Kevin Brooks

The Department of Transport reported by the end of 2023 that there were over 13 million registered vehicles in South Africa. Each of these vehicles showcases a range of petrochemical products derived from crude oil. As the global industrial environment moves to reduce environmental impact and maximise resource use, improving the efficiency of distillation processes becomes vital.

Distillation presents several problems, including fouling, product quality drift, feed variability and high energy consumption. Feed variability challenges the safe and efficient operation of distillation columns. This variation directly influences the product quality, yield, sustainability and overall economics of the column.

Digital twins are dynamic, data-driven virtual models of physical assets or processes. Digital twins simulate real-world operations, integrating historical data, real-time sensor input, and machine learning algorithms to mirror actual system behaviour. By simulating operational conditions and predicting process outcomes, digital twins can enable precise control over energy use, yield ratios, and emissions. The integration of digital twin technology into petrochemical processes potentially supports more efficient and environmentally responsible production.

This virtual mirroring enables proactive decision-making. Operators can experiment with process changes virtually before implementing them in physical infrastructure, reducing downtime and operational risks. Digital twins can also flag inefficiencies and suggest improvements by continuously learning from incoming data. The purpose of this research investigation is to further our knowledge of the advantages of digital twins and the justification of their implementation. A comparison will be made between various machine learning techniques and simulations. The benefit of implementing a digital twin specifically to accurately simulate feed changes before they are made will be investigated. The intent of the presentation is to outline the scope of research and progress thus far.

Closed-Cycle Distillation Platform for Multivariable Control Education

Presenter: Malhar Apte (University of Pretoria)
Supervisors: Paul Sonnendecker, Lehlogonolo Tabana

While distillation columns are standard textbook examples of multivariable control, physical units are rarely implemented in instructional laboratories due to solvent handling hazards and inventory costs. Closed-

cycle distillation eliminates these barriers by combining distillate and bottoms product streams in a buffer vessel and continuously recycling the mixture to the column as feed. Closing this inventory cycle fixes total solvent mass, reducing steady-state operation to a single degree of freedom dictated by light-key holdup, while preserving multivariable transient dynamics. Thermally conditioning the feed and reflux streams fixes inlet enthalpies. However, at small laboratory scales, the column is highly sensitive: throughput is strictly bounded by column diameter, flowmeter deadbands limit low-flow precision, and uncompensated inventory shifts cause rapid inventory exhaustion or overflow. Unmanaged operation easily enters a pseudo-steady state that eventually drains vessel holdups, preventing the column from reaching a true, repeatable steady state. To manage this constrained operating window, this work formulates the supervisory control strategy required to automate startup, enforce inventory bounds, and restore the platform to a repeatable base state following student experimentation. By shifting constraint enforcement to the background, this control strategy permits students to directly manipulate regulatory loops, including flow, temperature, and cascade loops while logging real-time data for dynamic modelling.

Component mass balances cannot be closed using volumetric flowmeters alone because liquid density varies non-linearly with temperature and composition. To resolve this, an inline density measuring instrument was developed to soft-sense feed composition in real time, which, when coupled with column temperature profiles, provides full composition observability across the column and buffer vessel. This mass balance closure enables a model-based state observer to track internal inventory distributions, providing the state feedback required to calculate low-level setpoint adjustments, accelerate startups, and recover from perturbations. This work establishes a safe, repeatable, and inventory-efficient platform for process control education.

Dynamic modelling of green ammonia production under variable renewable energy supply

Presenter: Abulele Nogaya (Stellenbosch University)
Supervisors: Tobi Louw, Prathieka Naidoo

Ammonia production accounts for approximately 1–2% of global carbon dioxide emissions, owing to its strong reliance on fossil-fuel-derived hydrogen and energy-intensive processes. Renewable-driven green ammonia production via Power-to-Ammonia systems has emerged as a promising pathway to reduce emissions while enabling renewable energy storage and transport. However, integrating intermittent renewable energy into the Haber-Bosch synthesis loop poses significant operational challenges, as ammonia reactors were historically designed for continuous, steady-state operation under stable feed and energy conditions.

Existing literature has either focused on system-level studies that simplify the synthesis loop using black-box representations, or on high-fidelity dynamic reactor studies that utilise idealised disturbances rather than realistic renewable variability. Consequently, the interaction between variable renewable intermittency and physically resolved synthesis-loop dynamics remains insufficiently characterised. Furthermore, many dynamic Haber-Bosch models are implemented in proprietary simulation environments or in unreleased research code, limiting reproducibility and accessibility.

This study develops a dynamic first-principles model of a renewable-driven Haber-Bosch synthesis loop subject to stochastic hydrogen-feed variability representative of intermittent renewable operation. The framework incorporates dynamic reactor mass and energy balances, recycle-loop dynamics, simplified upstream hydrogen and nitrogen production models, hydrogen buffering, and regulatory control. The modelling framework is being implemented in Python as open-source code.

The model formulation has been developed, and its Python implementation is ongoing; simulation results are therefore not yet available. Once implementation is complete, it will be used to assess reactor temperature stability, recycle-loop behaviour, constraint violations, and the extent to which hydrogen buffering mitigates feed disturbances. These metrics will help characterise the loop's flexible operating range and inform the evaluation of its regulatory control requirements.

The anticipated contribution is a reproducible, open-source framework to support the investigation of con-

trol strategies for flexible, renewable-driven ammonia production. This presentation outlines the modelling framework, disturbance representation, and planned assessment of synthesis-loop dynamics and regulatory control under stochastic feed variability.

Session 2: Data-Driven Modelling & Controller Tuning

How Much Physics Does a Hybrid Model Need? Neural ODEs for Nonlinear MPC, from Quiet Plant Data to Simulation-Trained Priors

Presenter: Greg Laycock (University of Pretoria)
Supervisor: Derik le Roux

Nonlinear model predictive control is only as good as its model when a plant moves to a new operating point. Data-driven models, including the Gaussian-process and Bayesian approaches now widely studied, predict well near recent data but cannot extrapolate to conditions they have not seen. First-principles models extrapolate when correct, but mismatch from simplifying assumptions, unmodelled phenomena and equipment ageing cannot be removed by refitting constants. Hybrid neural ODEs keep the mass and energy balances as written physics and let neural networks represent the uncertain terms, trained on whole trajectories. We present a training method for stiff systems (an implicit integrator, exact adjoint gradients and a second-order optimiser) that fits everything from a purely mechanistic model to a pure neural ODE.

On a simulated exothermic CSTR we walk a ladder of structures with progressively more of the reaction rate learned. Trained only on normal operation, models predict an unseen thermal runaway 146 K beyond their data only where the Arrhenius form is built into the structure and the learned correction is kept from taking it over. A learned rate otherwise fails above the data, whatever the activation function, and no selection criterion that does not know the truth identifies the structures that extrapolate.

Plant data rarely contains upsets, but a simulator usually exists. We therefore train hybrid models on the simulator's full operating envelope and fine-tune them on poorly excited data from plants that differ in progressively harder ways: changed constants, new effects visible in the data, and effects that are not. Each is compared with the first-principles model with refitted constants on the plant's unseen runaway and checked for violations of conservation laws. We close by outlining how the information content of plant data indicates how much physics a model must carry.

SafeOpt for the Automated Tuning of MIMO PID Controllers

Presenter: Emmanuel Aiyede (University of Pretoria)
Supervisors: Derik le Roux, Ian Craig

The automated tuning of multiple-input multiple-output (MIMO) systems is an emerging area of research that has attracted growing interest. Among the various automation techniques, Bayesian optimisation (BO) has been identified as an ideal method for this purpose. However, real-world applications must account for safety systems, especially when evaluating certain parameters that can lead to unstable behaviours or catastrophic damage. This proposal explores automated tuning methods for process control systems by evaluating BO under safety constraints. It explores BO as a powerful algorithm for optimal tuning and describes its optimisation technique, including recent advancements and current limitations. To address the identified limitations in control systems, the proposal introduces the Safe Bayesian optimisation (SafeOpt) framework, which integrates safety constraints into the optimisation process, along with its promising variants for multivariable controllers. This work outlines a comparison with Bayesian optimisation to determine the performance of the SafeOpt framework. The paper compares the SafeOpt framework with existing work on BO for tuning MIMO controllers in process control, particularly for a bulk tailings treatment plant and grinding mill circuit.

Time Series Foundation Models for Chemical Process Modelling

Presenter: Imker Hoogenhout (Stellenbosch University)
Supervisor: Tobi Louw

Digital process models for the chemical industry have contributed greatly to process understanding and optimisation, and are used for a variety of sophisticated tasks, such as state estimation, model predictive control, process monitoring, operator advisory functions, and process forecasting. The most dominant type of process model is the first-principles model (FPM), typically derived from material and energy conservation laws and coupled with auxiliary equations that describe mass and heat transfer, chemical reactions, and thermodynamics of the system. Such models require a detailed mechanistic understanding of the process and a substantial effort to develop, yet they remain inadequate for some complex systems. Time-series foundation models (TSFMs) offer a potential alternative. The main application of TSFMs is zero-shot forecasting, where future values of a time series are predicted based on previous values given to the model as context. Since the second half of 2025, TSFMs have not only improved in their capabilities, but also in their functionality and versatility. Most major TSFM releases over the last year have added support for multivariate forecasting and past-and-future known covariates (inputs). This has allowed applications to go beyond just time-series forecasting; it has created a pre-trained input-output model that can learn dynamic relationships between target variables (outputs) and covariates (inputs) in context. These new capabilities have recently gained attention in applications such as MPC, fault detection, and knowledge distillation. This work aims to implement TSFMs as chemical process models, enabling what-if analysis, and providing operator assistance, whilst only requiring process data. TimesFM 3.0, Chronos-2, and a LoRA fine-tuned Chronos-2 were implemented on the Tennessee Eastman process (TEP) as a benchmark case study. The predicted responses of each model to set-point changes and disturbances were validated against those of the first-principles TEP model. All models showed promising results and proved simple to implement; however, limitations such as insufficient excitation in the context data must be addressed before direct industrial application, and potential ways to address these are presented.

Session 3: Guidance, Navigation & Control of Aerial Vehicles

Observer-Based Disturbance Rejection for Quadcopter Attitude Control in Windy Environments

Presenter: Mpilonhle Nxasana (University of Cape Town)
Supervisor: Arnold Pretorius

Wind acts on the attitude dynamics of a quadcopter as an unknown, time-varying torque. The linear quadratic regulator (LQR) handles the coupled multivariable dynamics but offers no explicit disturbance rejection, and added integral action removes only constant disturbances, and only once the error has accumulated. Since real wind fluctuates, this work develops an attitude controller that also rejects time-varying disturbances.

The nominal LQR is augmented with a linear extended state observer (ESO), a second-order observer on each axis, which lumps model mismatch and external torques into a single disturbance on the rate dynamics, estimates it from the measured body rates and cancels it at the control input. This restores the nominal plant seen by the LQR without changing its gains. With all observer poles placed at a common bandwidth, that bandwidth becomes the only additional design parameter, raised until the expected disturbances are rejected but limited by the amplification of gyroscope noise.

The LQR, an LQR with integral action (LQI) and the LQR+ESO approach were run onboard a 70 g micro-sized quadcopter at 142 Hz. The quadcopter was mounted on a three-axis attitude rig, where the gravitational torque about the pivot acts as a repeatable proxy for wind: constant while a tilt is held and sinusoidal under a conical reference. In sequential 15° step tests, integral action reduced the root-mean-square tracking error by 66–81% relative to LQR, and the ESO reduced it by a further 50–57%, an overall reduction of 83–

92%. Under the conical reference, the LQI improved on the LQR by only 17%, whereas the ESO reduced the mean attitude error from 10.0° with LQI to 1.6°, without raising the feedback gains, which already sat at the onset of instability. Future work will test the scheme under fan-generated wind in free flight before extending it to a full-sized tiltrotor quadcopter.

Cascaded Linear Model Predictive Control for Agile Quadcopter Trajectory Tracking

Presenter: Tony Mwangi (University of Cape Town)
Supervisors: Arnold Pretorius, Ed Boje

Agile quadcopter trajectory tracking requires a controller that accounts for the vehicle's actuation limits, anticipates the reference trajectory, and remains valid far from hover, where the small-angle approximations underpinning conventional linear designs are difficult to justify. Nonlinear model predictive control (MPC) meets these demands at a computational cost that challenges onboard deployment, while linear MPC is typically built on singularity-prone Euler-angle models. This work presents a cascaded linear MPC architecture on a quaternion-based quadcopter plant. An outer position loop, running at 20 Hz, commands a world-frame force vector, which an analytic force-to-attitude mapping converts into a thrust command and a desired quaternion for an inner attitude loop running at 100 Hz. This mapping decouples the translational and rotational dynamics, so that the outer loop controls a linear model while the inner loop regulates a singularity-free exponential-coordinate attitude error along the shortest rotational path. Tilt constraints derived from the vehicle's thrust-to-weight ratio are imposed on the commanded force to keep aggressive manoeuvres feasible. Three controller configurations are evaluated in MATLAB/Simulink on an identified micro-UAV model: a hover-linearised controller without reference trajectory preview on the position loop, the same controller with preview, and an adaptive controller that relinearises the attitude dynamics at the operating point at each step. The controller configurations are tested on a lemniscate speed sweep and two time-optimal trajectories, the latter subjected to external disturbances, sensor noise, actuation lag, and parameter uncertainty. The results show that reference preview is the dominant factor in tracking accuracy, reducing the position error by up to two orders of magnitude on the lemniscate and by 41–59% on the time-optimal trajectories. Under non-ideal conditions, adaptive MPC proves the most robust, more than halving the remaining tracking error on both aggressive trajectories by capturing the body-rate-dependent gyroscopic coupling term in the rotational dynamics. The attitude controller's performance has been validated on the physical quadcopter platform through constrained bench tests, and experimental flight testing is ongoing.

Development of an Autonomous Parafoil-Based Recovery System for Suborbital Rockets and High-Altitude Payloads

Presenter: Joshua Spies (Stellenbosch University)
Supervisors: Willem Jordaan, Jean Pitot (University of KwaZulu-Natal)

Successful recovery of suborbital rockets and other high-altitude payloads remains complex and challenging; autonomous parafoils are a potential solution. Parafoils offer a low-cost means of delivering payloads from high altitude, but landing them accurately is difficult. The vehicle has no thrust, is steered only by asymmetric brake deflection, is strongly affected by wind, and can spend altitude but never regain it. This work presents a complete guidance, navigation and control (GNC) system for an autonomous parafoil, developed in simulation and structured for deployment.

The vehicle is modelled with Jann's four-degree-of-freedom model, in which asymmetric brake banks the canopy and bank produces the turn. Guidance is a three-phase state machine. A Dubins transit path leads directly to the target, giving a known line of approach, and then enters a loiter circle sized from the current altitude and a glide ratio measured over the preceding descent. It leaves the loiter on the last lap that still provides enough altitude for the final approach, whose length is solved numerically so that the remaining path matches the available glide distance, with a margin so it lands slightly long rather than short. The

path is tracked by Park's nonlinear L1 guidance law, feeding a feedforward and proportional feedback that commands the brakes.

Navigation uses a 15-state error-state Kalman filter driven by inertial measurements and corrected by GPS, barometer and magnetometer, requiring no model of the vehicle dynamics. Running live in the closed-loop simulation with realistic sensor models, it estimates position to 0.78 m RMS from 1 m GPS noise, altitude to 0.12 m, roll to 0.32° and turn rate to 0.31°/s, and guidance operates entirely on these estimates. Simulations show promising results, robustness against wind and the path towards flight testing.

Actuated Tether Proprioception for GNSS-Denied Localisation of a Marsupial UAV-UGV System

Presenter: Guy Hasewinkel (University of Cape Town)
Supervisor: Arnold Pretorius

Autonomous inventory auditing in a high-bay warehouse requires long-duration missions and highly mobile robotic platform(s) to inspect goods at high vantage points on shelves up to 12 m. The densely populated warehouse environment denies global navigation satellite system (GNSS) signals. This presents a complex and challenging task for a single robot to achieve (long-duration autonomous missions with non-trivial inspection points); however, if an unmanned aerial vehicle (UAV) and an unmanned ground vehicle (UGV) are coupled together to form a marsupial system, the two robots can work in tandem to achieve the task and mitigate each other's drawbacks. The UGV is constrained in vantage point, and is dynamically slow compared to the UAV. Conversely, UAVs are naturally constrained by payload. However, by coupling the UAV to the UGV, the battery and compute can be offloaded, thus increasing the flight time and allowing the UAV to perform the inspection routine. To provide the required energy and communication transfer between the UAV and UGV, the physical couple is a tether cable, which in turn can be used as a controllable input to the UAV. If the tether hangs freely (not impeded by any objects, forming a catenary shape), the physics are simple and can be exploited as a means to localise the UAV relative to the UGV. To this end, this work develops an actuated proprioceptive tether management system (TMS) and sensor fusion scheme to localise the UAV and UGV. In particular, the designed TMS actuates and senses the azimuth, elevation and length (in turn the tension) of the tether. The sensor fusion scheme, an extended Kalman filter (EKF), exploits the known tether physics and uses pose estimates from an onboard camera on the UAV to reliably estimate the UAV's position. This localisation methodology improves the UAV localisation accuracy when compared against existing localisation strategies.