

Reconfigurable pipelined control systems

R. Medina Sanchez¹, S. Tabatabaei Nikkhah², D. Goswami², W.P.M.H. Heemels², S. Stuijk², T. Basten^{2,3}

¹TNO, Powertrains, The Netherlands

²Eindhoven University of Technology, The Netherlands

³ESI, TNO, The Netherlands

Abstract – Pipelining is a promising technique to deal with long delays in control loops, for instance due to intensive data processing. Reconfigurable pipelined control (RPC) is of particular interest, as it has an improved performance compared to a static configuration. We present design and implementation techniques for RPC systems.

Keywords— Data intensive control systems, Pipelined control, reconfiguration.

I. INTRODUCTION

Intensive sensor-data processing in a control loop occurs in many application domains including ADAS (Advanced Driver Assistance Systems) [7] and vision-guided control [8]. Such sensing technologies allow to extract a rich set of information from the environment and is instrumental to develop reliable functionality in various modern systems. An example is image processing in the loop where a camera and an image processing algorithm are used for sensing [2]. The key challenge in designing such systems is to deal with the long *sensing delay* resulting from the heavy computation demand for sensor-data processing.

Strategies for coping with a long sensing delay are reported both in the embedded systems [1], [9] and in the control theory literature [5], [8]. In [5], sensing delay is compensated by advanced estimation techniques and in [8] multi-rate actuation strategies are used. These strategies heavily rely on the system model and are vulnerable to modelling errors. Embedded systems approaches aim to reduce the data-processing latency

by creating parallel implementations of the algorithms in specialized hardware such as GPUs [1] or FPGAs [9]. These strategies require a great effort to speed up the data-processing algorithms and the final latency may still not be acceptable for safety-critical applications.

A control loop typically requires to perform sensing, control computation and actuation operations periodically in order to regulate the behavior of a dynamic system. In a classical implementation (Fig. 1), these operations are performed sequentially in which long sensor-data processing would lead to a long execution time. A sampling period h is chosen at least equal to the total execution time τ of the loop. A camera, for example, can capture image frames with a fixed period f_h . Typically, the camera frame rate can be much higher than the rate at which the frames are processed, i.e., $f_h \ll h$. For instance, in Fig. 1, $h = 6f_h$; only every sixth frame is processed due to the long data-processing time. Such a long sampling period can lead to a degraded control performance or even instability and is often unacceptable in safety-critical domains.

A potential solution may be found in cross-layer design taking into account both controller design and its embedded implementation. The key idea is to increase the rate of sensor-data processing by processing frames in parallel on multiple cores. In this way, the sensor processing rate can be increased while the sensing delay remains unaffected. To realize such a pipelined implementation, we need scheduling and reconfiguration mechanisms to match model-level assumptions made in the control design. Further, we need control algorithms taking into account implementation-level artifacts to deal with the long delay and to translate a higher sensing rate to a better performance. Such cross-layer design problems require novel techniques to handle challenges across multiple layers that are usually dealt with in isolation.

Why pipelined control? Pipelined control systems (illustrated in Fig. 3(a)) can outperform a sequential implementation as was shown in [10]. In a pipelined loop, the time between two consecutive completions of sensor-data processing is reduced thereby essentially shortening the sampling period, while the delay is unaltered and thus, remains long. This is illustrated in Fig. 3, which shows an example with three control cores. The sampling period is reduced to $2f_h$ compared to the sequential case with sampling period $6f_h$. The method is particularly suitable when consecutive data samples (images, or any other type of data) can be processed independently. In general, there is a trade-off between the possible overhead caused by




Fig. 1: Data-intensive feedback control systems. The execution time of sensor-data processing may vary. Control computation and actuation are time triggered, taking into account the worst-case sensor-data processing time, to ensure a constant sampling period. Color code: green for sensing, blue for computing and yellow for actuating.

independence between data samples and the (performance) gains of pipelining. Pipelined control is enabled by current trends in computer architecture integrating multi/many-cores, GPUs, and FPGAs in embedded systems.

Why reconfigurable pipelined control? Parallel processing platforms may be *shared* among multiple applications for cost reasons. Availability of cores may change at runtime depending on the activation of (e.g., sporadic) tasks. Statically allocating cores to the pipelined loops limits interference between applications, however, would lead to a small number of available cores for the control loop. An alternative is to *dynamically* allocate cores to the sensor-data processing depending on their availability. Fig. 5 shows an example where the number of available cores temporarily changes from three to two (and the system has to reconfigure accordingly). If the pipelined system can run most of the time with three cores while it may only occasionally need to run with two cores, a static allocation would lead to using only two control cores. In contrast, reconfigurable pipelined control (RPC) could use most of the time three cores and thereby has the potential to outperform static pipelined control in a long run, if designed and implemented properly. The RPC design and analysis aspects are studied in [11]. The current work focuses on how the proposed method can be implemented on embedded platforms. Our implementation is based on a proof-of-concept implementation architecture for pipelined control in [2].

In an RPC system, we need to take into account the implementation-layer timing artifacts in the control model, design, and optimization. Moreover, we need to schedule the control cores to realize the temporal behavior assumed at model-level and in control design. Therefore, an RPC system requires a reconfiguration mechanism that respects timing assumptions made at the model-level under all configurations as well as a control design technique that guarantees stability and an improved performance under possible configuration switching. In this work, we emphasize the implementation aspects of RPC systems and illustrate these on an experimental setup.

II. SYSTEM SETUP

This section introduces the control problem at hand.

A. Motivating application: xCPS

The Explore Cyber-Physical Systems (xCPS) platform (Fig. 2) is an industrial assembly line testbed [3]. The machine assembles pieces that come in two shapes: lower and upper parts. Conveyor belts move the assembly pieces through the machine actuators: a turner, multiple stoppers, and a pick-and-place unit, among others. Regulating the speed and positions of the assembly blocks is crucial for guaranteeing the productivity of the assembly process. To do so, a camera running at 60 frames per second (fps) and an image-processing algorithm based on the Hough transform for circles are used to measure the speed and position of the assembly blocks. The control operations and sensor processing are performed on a 4-core

embedded board. This is a representative system capturing a number of common challenges in many real-life applications that use vision-based processing in the feedback loop [7].

B. Feedback control systems

A wide class of dynamic systems, including the motion system related to the assembly blocks mentioned above, can be described by the linear state-space model

$$\begin{aligned}\dot{x}(t) &= A_c x(t) + B_c u(t), \\ y(t) &= C_c x(t),\end{aligned}\quad (1)$$

where $A_c \in \mathbb{R}^{n \times n}$, $B_c \in \mathbb{R}^{n \times m}$, and $C_c \in \mathbb{R}^{p \times n}$ are system, input and output matrices, respectively and $x(t) \in \mathbb{R}^n$ is the state vector, $u(t) \in \mathbb{R}^m$ is the control input used for actuation, and $y(t) \in \mathbb{R}^p$ is the output to be regulated, at time $t \in \mathbb{R}_{\geq 0}$. With sensor-to-actuator delay τ and zero-order-hold mechanism with sampling period $h \in \mathbb{R}_{\geq 0}$, $u(t)$ becomes a piecewise constant signal, being constant in the intervals $[kh + \tau, (k+1)h + \tau)$ for $k \in \mathbb{Z}_{\geq 0}$. A pipelined control system is relevant when τ is long as motivated in the introduction. We assume that the full state is available for feedback although controller-output based design could be considered as well.

Example 1. The model of an xCPS conveyor belt has two states: $x = [x_1 \ x_2]^T$ (position and velocity), and one input u (motor voltage), one output $y = x_1$ (position of the belt). The system matrices are given by

$$A_c = \begin{bmatrix} 0 & 1.7 \\ -9 & -2.5 \end{bmatrix}, B_c = \begin{bmatrix} 0 \\ 10 \end{bmatrix}, C_c = [1 \ 0].$$

C. Sensor-to-actuation loop execution

A control or sensor-to-actuator loop is executed by periodically performing sensing, control computation and actuation operations. For the system in Eq. 1, the sensing, reading and processing of system state $x(t)$, starts periodically at $t = kh$, $k \in \mathbb{Z}_{\geq 0}$. The sampling period h is the interval between two consecutive activations of the sensing operations that require image frames for processing. A camera captures the images at




Fig. 2: xCPS platform [3].

discrete instants $t_f \in \{if_h | i \in \mathbb{Z}_{\geq 0}\}$, where f_h is the frame arrival period. In xCPS

$$f_h = 1/60 \text{ s} = 16.67 \text{ ms}.$$

We align the start of a sensing operation with t_f (i.e., frame arrival times) as shown in Fig. 3; hence, h is an integer multiple of f_h . Sensor processing is followed by control computation and actuation operations, which generally take short and nearly constant time for execution. Sensing takes much longer time due to the computationally heavy processing, i.e.,

$$\tau_s \gg \tau_c + \tau_a,$$

where τ_s , τ_c and τ_a are the worst-case execution times of sensor-data processing, control computation and actuation operations, respectively. The total (worst-case) execution time τ_t of a sensor-to-actuator loop is thus given by

$$\tau_t = \tau_s + \tau_c + \tau_a. \quad (2)$$

Due to alignment of the start of a sensor operation with t_f , the effective sensor-to-actuator execution time τ is given by

$$\tau = \left\lceil \frac{\tau_t}{f_h} \right\rceil f_h. \quad (3)$$

Clearly, $\tau \geq \tau_t$. While the sensor processing has variable execution time (see, e.g., the blue arrows in Fig. 1), time-triggered activations result in a constant input-to-output delay τ (an integer multiple of f_h).

Example 2. A camera and an image-processing algorithm are used as a sensor for measuring the states x_1 and x_2 in Example 1. The total worst-case execution time τ_t equals 84 ms with $f_h = 16.67$ ms, $\tau = 6f_h = 100$ ms following (3). That is, every sixth frame is processed by a core.

III. PIPELINED CONTROL SYSTEMS

Pipelined control is realized by executing sensor-to-actuator loops in parallel in γ globally synchronous cores (see Fig. 3(a) for $\gamma = 3$) and a camera with a fixed frame-rate. The choice of γ is made based on the requirement on sampling period h (which depends on the system dynamics (1)) and estimation of τ on a given platform. A proof-of-concept implementation architecture in [2] uses time-triggered activations. The cores are denoted by **Core 1** to **Core γ** , where **Core 1** is referred to as the *base* core. We denote the start of the sensor-to-actuator loop for the k^{th} sample in **Core j** as $t_j(k)$. The pipelining is initiated by the base core at $t_1(0) = 0$. The triggering offset between **Core $j+1$** and **Core j** is related as

$$t_{j+1}(k+1) = t_j(k) + h, \quad (4)$$

where h is the sampling period given by

$$h = \frac{\tau}{\gamma}. \quad (5)$$

Further, we trigger the loop on **Core j** as

$$t_j(k+\gamma) = t_j(k) + \tau. \quad (6)$$

In essence, the duration τ is divided into γ samples of length h . Clearly, a larger γ gives a shorter sampling period h . To align the control operations with the availability of the frames

at discrete intervals, $\left\lceil \frac{\tau_t}{f_h} \right\rceil$ is chosen to be integer divisible by γ . Hence, a set of finite configurations are considered.

Example 3. We know from Example 2 that $\tau = 6f_h$ and we have a quad-core platform. With $\gamma = 2$, the sampling period $h = 3f_h$. Similarly, $\gamma = 3$ gives $h = 2f_h$. The choice $\gamma = 4$ causes miss alignment with the frame arrival and, therefore, is not used here. Hence, $\gamma \in \{1, 2, 3\}$ is admissible in the example.

A. Discretized pipelined model

We discretize the continuous-time system (1) with sampling period h and sensor-to-actuator delay τ :

$$x((k+1)h) = A_d x(kh) + B_d u(kh - \gamma h), \quad (7)$$

where A_d, B_d are the discretized state and input matrices [11]. The control input $u(t)$ at $t = kh$ uses γh time units (or γ frames) old sensing information in any sampling interval kh to $(k+1)h$ due to the sensor-to-actuator delay τ . This is reflected in (7) as the delayed input $u(kh - \gamma h)$.

For designing the delayed control input $u(kh - \gamma h)$, one design option is to transform system (7) into standard non-delayed form and apply standard control design techniques. For this purpose, we define a new system state vector including γ old input signals as

$$z(kh) = [x^T(kh) \quad u(kh - \gamma h) \quad \cdots \quad u(kh - h)]^T. \quad (8)$$

Using the state definition in (8), the system model (7) is transformed into the following non-delayed form

$$\begin{aligned} z(kh + h) &= \Phi_d z(kh) + \Gamma_d u(kh), \\ y(kh) &= C_d z(kh) \end{aligned} \quad (9)$$

with Φ_d, Γ_d and C_d are the corresponding augmented discretized matrices and z is the new augmented state vector. See [11] for details.

Example 4. From Example 2, we know that $\tau = 100$ ms and $n = 2$ (two states). Consider $\gamma = 2$ and the resulting sampling period $h = 3f_h$. The augmented state vector of Eq. 8 has $(n + \gamma)$ states and is given by $z(kh) = [x_1(kh) \quad x_2(kh) \quad u(kh - 2h) \quad u(kh - h)]^T$ with measured states $x(kh)$ at $t = kh$ and two older input signals.

B. Pipelined control law and performance

System (9) is in standard discrete-time form for which standard discrete-time control design techniques such as LQR [2] can be used in the form

$$u(kh) = Kz(kh), \quad (10)$$

where K is the state feedback gain to be designed. We apply $u(t) = 0$ for $t < \tau$ since we need to wait for the first sample to be processed before the first feedback-based actuation can take place. As a performance metric, we consider the settling time, which is defined here as the time that the system output takes to reach and stay within a bound of 2% around the reference or zero. Other metrics can also be considered depending on the system requirements.

Example 5. Let us continue with $\tau = 100 \text{ ms} = 6f_h$, $\gamma = 2$ and the resulting sampling period $h = 3f_h$. Control law (10) uses the system state vector $z(kh)$ defined in (8), which has both measured state $x(kh)$ and γ old (stored) input signals. From (7), old state vector $z(kh - \tau)$ is used by (10) in the interval $[kh, (k+1)h)$, i.e., the “sensing delay” is modelled as the “actuator delay”.

$$u(kh - \gamma h) = K \begin{bmatrix} x_1(kh - 6f_h) \\ x_2(kh - 6f_h) \\ u(kh - 12f_h) \\ u(kh - 9f_h) \end{bmatrix}.$$

Clearly, the measured signal x is six frames old and there are other older input signals in the control law. K is designed using the technique presented in [2] leading to

$$K = \begin{bmatrix} -30.9 & -7.3 & -2.7 & -1.9 \end{bmatrix},$$

which results in a settling time of 123.5 ms.

The settling time improves with a higher number of pipelined control cores [2]. A settling time below 120 ms is achieved with $\gamma = 3$, whereas a sequential implementation with $\gamma = 1$ gives a settling time larger than 220 ms.

C. Implementation architecture for pipelined control

Our implementation architecture for the pipelined control is shown in Fig. 3(b), based on [2]. A sensor-to-actuator loop is triggered as per (4) and (6) and it has an (effective) constant execution time of τ , which implies that the processed states are $x(kh - \tau)$ or τ time old. The sensor-data processing is performed on a single image to obtain $x(kh - \tau)$ and stored in local memory (no data dependency between the sensing pipes). The computing task uses $x(kh - \tau)$ and the feedback gain K from the local memory since they are not dependent on other cores. It uses γ old input signals from shared (among the γ control cores) memory. Access to the shared memory is realized with a predictable memory controller (by non-overlapping allocation) reported in [4]. Overall, the control tasks have nearly constant execution times under this predictable time-triggered architecture.

IV. RECONFIGURABLE PIPELINED CONTROL SYSTEMS

Fig. 4 illustrates the idea of reconfigurable pipelined control (RPC) systems. The system may run either in the *maximal configuration* or in a *reduced configuration* depending on the availability of the cores. A maximal configuration (MC) uses γ_{MC} control pipes and the feedback gain K_{MC} in the control law (10). The sampling period h_{MC} in the MC is given by

$$h_{MC} = \frac{\tau}{\gamma_{MC}}. \quad (11)$$

Stability and performance of the MC depend on the closed-loop system matrix $(\Phi_{MC} + \Gamma_{MC}K_{MC})$. Similarly, a reduced configuration (RC) uses $\gamma_{RC,i}$ control pipes (which is smaller than γ_{MC}) and the feedback gain $K_{RC,i}$ in the control law (10). The sampling period $h_{RC,i}$ in a RC is given by

$$h_{RC,i} = \frac{\tau}{\gamma_{RC,i}}. \quad (12)$$




Fig. 3: Pipelined control systems (a) Timing (b) Implementation architecture.

The closed-loop matrix $(\Phi_{RC,i} + \Gamma_{RC,i}K_{RC,i})$ dictates the stability and performance in a RC. There may be multiple reduced configurations.

Example 6. For the system introduced in Section II-A, the possible configurations are given by $\gamma \in \{1, 2, 3\}$. Here, $\gamma_{MC} = 3$ and there are two reduced configurations, e.g., $\gamma_{RC,1} = 1$ and $\gamma_{RC,2} = 2$.

A. Reconfiguration controller design

We design the feedback gains K_{MC} and $K_{RC,i}$ for overall system stability and performance optimization following [11]. As per the definition of the augmented state in (8), the state vector differs in dimension in various configurations. Stability analysis of the switched systems can be performed among the systems of identical dimension. We match the dimension in the configurations by augmenting exactly γ_{MC} old input signals as shown in (8).

Example 7. Recall Example 2. $\tau = 100 \text{ ms}$. The maximal configuration $\gamma_{MC} = 3$. The resulting sampling period $h_{MC} = 2f_h$. The augmented state vector has $(n + \gamma_{MC}) = 5$ states:




Fig. 4: Reconfiguration pipelined control systems.

$$z(kh_{MC}) = \begin{bmatrix} x_1(kh_{MC}) \\ x_2(kh_{MC}) \\ u((k-3)h_{MC}) \\ u((k-2)h_{MC}) \\ u((k-1)h_{MC}) \end{bmatrix} = \begin{bmatrix} x_1(kh_{MC}) \\ x_2(kh_{MC}) \\ u(kh_{MC} - 6f_h) \\ u(kh_{MC} - 4f_h) \\ u(kh_{MC} - 2f_h) \end{bmatrix}.$$

Next, we consider the reduced configuration with $\gamma_{RC,i} = 2$. The resulting sampling period $h_{RC,i} = 3f_h$. The augmented state vector has $(n + \gamma_{MC}) = 5$ states and is given by

$$z(kh_{RC,i}) = \begin{bmatrix} x_1(kh_{RC,i}) \\ x_2(kh_{RC,i}) \\ u((k-3)h_{RC,i}) \\ u((k-2)h_{RC,i}) \\ u((k-1)h_{RC,i}) \end{bmatrix} = \begin{bmatrix} x_1(kh_{RC,i}) \\ x_2(kh_{RC,i}) \\ u(kh_{RC,i} - 9f_h) \\ u(kh_{RC,i} - 6f_h) \\ u(kh_{RC,i} - 3f_h) \end{bmatrix}.$$

It can be seen that a reduced configuration uses older input signals compared to the maximal configuration.

B. Reconfiguration mechanism

We consider $\gamma \in \{\gamma_{MC}, \gamma_{RC,i}\}$ globally synchronous cores and a camera with a fixed frame-rate. **Core 1** is the base core, which initiates the pipelined sensing at $t_1(0) = 0$ and continues at least till all the γ cores complete their first sensor-to-actuation loop as per (4) and (6) before any configuration switching occurs. The reconfiguration process is initiated by the base core at $t_{init} = t_1(k)$ (at a triggering instant of a sensor-to-actuator loop at the base core), while actual switching of the feedback gains (i.e., K_{MC} , $K_{RC,i}$) and sampling periods (i.e., h_{MC} and $h_{RC,i}$) occurs at

$$t_{sw} = t_{init} + \tau.$$

The switching between sampling periods needs offset adjustments ($t_j(k)$ in **Core j**) for which the reconfiguration process must be initiated at least τ time earlier at t_{init} . The offsets are computed as

$$t_j(k+j-1) = t_{init} + \Delta_j. \quad (13)$$

Δ_j (offset adjustment in **Core j**) differs in different configuration switching.

Maximal to reduced configuration switching: Core 1 to Core $\gamma_{RC,i}$ continue in the RC while the other $(\gamma_{MC} - \gamma_{RC,i})$

cores are freed up after completion of an ongoing sensor-to-actuation loop. With $t_{init} = t_1(k)$, the offset adjustments are computed as

$$\Delta_j = (j-1)h_{RC,i}, j \in \{1.. \gamma_{RC,i}\}. \quad (14)$$

Fig.5(a) shows an example reconfiguration from $\gamma_{MC} = 3$ to $\gamma_{RC,i} = 2$; $\Delta_1 = 0$ and $\Delta_2 = 3f_h$.

Reduced to maximal configuration switching: $(\gamma_{MC} - \gamma_{RC,i})$ new cores are activated. With $t_{init} = t_1(k)$, the offset adjustments are computed as

$$\begin{aligned} \Delta_1 &= 0, \\ \Delta_j &= (\gamma_{MC} - \gamma_{RC,i})h_{MC} + (j-1)h_{MC}, j \in \{2.. \gamma_{RC,i}\}, \\ \Delta_j &= (j - \gamma_{RC,i})h_{MC}, j \in \{\gamma_{RC,i} + 1.. \gamma_{MC}\}. \end{aligned} \quad (15)$$

Fig.5(a) shows an example reconfiguration from $\gamma_{RC,i} = 2$ to $\gamma_{MC} = 3$; $\Delta_1 = 0$, $\Delta_2 = 4f_h$ and $\Delta_3 = 2f_h$.

C. Implementation architecture for RPC

The basic architecture for RPC is the one used for pipelined control (Fig. 3(b)). The stored data needs RPC-specific updates: (i) The feedback gains and $t_j(k)$ updated with the offset adjustments Δ_j should be stored in the local memory. (ii) The shared memory needs to store γ_{MC} old input signals. The other important ingredient is the reconfiguration mechanism, which realizes switching between the configurations at runtime using techniques proposed in [6]. In this method, control codes for the sensor-to-actuation loop are already loaded on γ_{MC} cores. Activating/freeing up a core boils down to activating/deactivating a preloaded code instance since the sensor-data processing is performed independently in each core. This also implies that this mechanism has no context switch overhead and has a negligible reconfiguration overhead [6].

V. DISCUSSION AND CONCLUSIONS

A pipelined control (PC) implementation can outperform traditional sequential implementation by executing sensor data-processing in parallel in multiple cores. For example, in the xCPS system, a PC with two (and three) cores achieves a settling time of 123.5 ms (and around 115 ms) as opposed to a sequential implementation which achieves a settling time 220 ms. A RPC switches between available configurations based on available cores can achieve a performance close to the one in the maximal configuration in the long run. In the xCPS system, for example, the RPC design achieves a settling time close to 115 ms. We presented how such a RPC scheme can be implemented on an embedded architecture.

Acknowledgment: This work is partially supported by the H2020 projects FitOptiVis (GA no.783162) and I-MECH (GA no. 737453).




Fig. 5: Reconfigurable pipelined control (a) Timing (b) Implementation architecture.

REFERENCES

- [1] R. Agrawal et al. A GPU based Real-Time CUDA implementation for obtaining Visual Saliency. In *ICVGIP*, 2014.
- [2] R. Medina Sanchez et al. Designing a controller with image-based pipelined sensing and additive uncertainties. *ACM Transactions on Cyber-Physical Systems*, 3(3):1–26, 2019.
- [3] S. Adyanthaya et al. xCPS: a tool to explore cyber physical systems. *SIGBED Review*, 14(1):81–95, 2016.
- [4] Sven Goossens et al. A reconfigurable real-time sdr controller for mixed time-criticality systems. In *CODES+ISSS*, 2013.
- [5] H. Fujimoto. Visual Servoing of 6 DOF Manipulator by Multirate Control with Depth Identification. In *CDC*, 2003.
- [6] L. Gantel, S. Layouni, M. E. A. Benkhelifa, F. Verdier, and S. Chauvet. Multiprocessor task migration implementation in a reconfigurable platform. In *2009 International Conference on Reconfigurable Computing and FPGAs*, pages 362–367, 2009.
- [7] K. Jo et al. Development of autonomous car—part I: Distributed system architecture and development process. *Industrial Electronics, IEEE Trans. on*, 61(12):7131–7140, 2014.
- [8] A. Kawamura et al. Robust visual servoing for object manipulation with large time-delays of visual information. In *IROS*, 2012.
- [9] S. Kestur et al. Emulating mammalian vision on reconfigurable hardware. In *FCCM*, 2012.
- [10] P. Krautgartner et al. Performance evaluation of vision-based control tasks. In *ICRA*, 1998.

- [11] R. Medina Sanchez et al. Reconfigurable pipelined sensing for image-based control. In *SIES*, 2016.

Rbinson Medina completed his MSc in Industrial Control Engineering from Ghent University (Belgium) and University of Ibagu in 2012, as part of a dual exchange program. He currently works as a Scientist Innovator at TNO-automotive. His research interests are in the fields of applied control, automotive systems, and embedded systems.

Contact: Eindhoven University of Technology, Electrical Engineering, PO Box 513, 5600 MB Eindhoven; robinson.medina@tno.nl.

Shayan Tabatabaei Nikkhah received the M.Sc. degree in Electrical Engineering from the University of Tehran in 2018. He is currently pursuing the Ph.D. degree in Electrical Engineering with the Eindhoven University of Technology. His research interests include real-time systems, quality and resource management, and approximate computing.

Contact: Eindhoven University of Technology, Electrical Engineering, PO Box 513, 5600 MB Eindhoven; s.tabatabaei.nikkhah@tue.nl.

Dip Goswami received his PhD in 2009 from National University of Singapore. He is an assistant professor in the Department of Electrical Engineering at Eindhoven University of Technology. His research focuses on design and analysis for embedded control systems. Dr. Goswami is a member of IEEE.

Contact: Eindhoven University of Technology, Electrical Engineering, PO Box 513, 5600 MB Eindhoven; d.goswami@tue.nl.

Maurice Heemels is a Professor at Eindhoven University of Technology and Fellow of IEEE. Dr. Heemels served/s on the editorial boards of *Automatica*, *NAHS*, *Annual Reviews in Control*, and *IEEE Trans. Automatic Control*. He received the 2019 IEEE L-CSS Outstanding Paper and will serve on IEEE-CSS Board of Governors (2021-2023).

Contact: Eindhoven University of Technology, Mechanical Engineering, PO Box 513, 5600 MB Eindhoven; W.P.M.H.Heemels@tue.nl.

Sander Stuijk received his PhD in 2007 from the Eindhoven University of Technology. He is an associate professor in the Department of Electrical Engineering at Eindhoven University of Technology. His research focuses on design trajectories for embedded signal processing application.

Contact: Eindhoven University of Technology, Electrical Engineering, PO Box 513, 5600 MB Eindhoven; S.Stuijk@tue.nl.

Twan Basten is a professor in the Electronic Systems group at Eindhoven University of Technology. His research focuses on model-driven design of embedded and cyber-physical systems. He is a senior member of IEEE and a life member of ACM.

Contact: Eindhoven University of Technology, Electrical Engineering, PO Box 513, 5600 MB Eindhoven; A.A.Basten@tue.nl.