Dependable Platforms for Autonomous systems and Control

DPAC

Newsletter
spring 2020

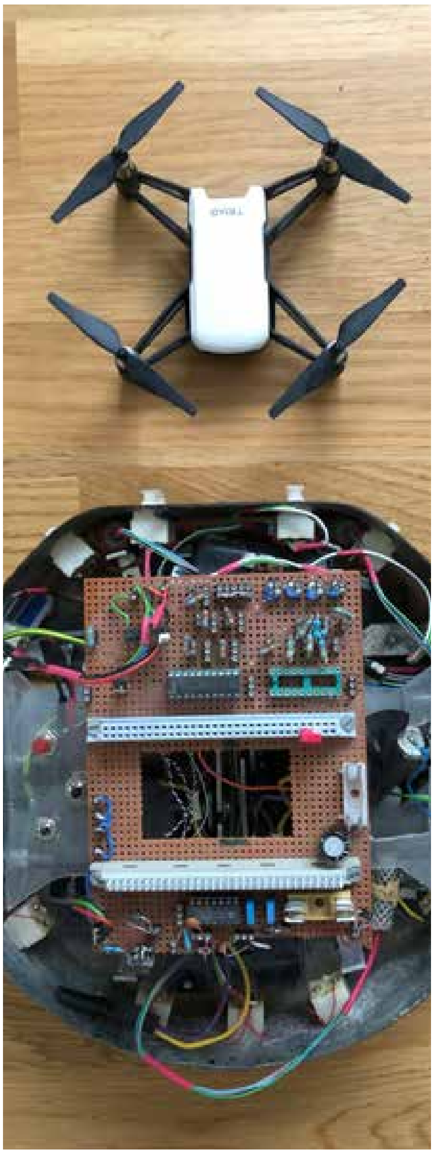

MÄLARDALEN UNIVERSITY SWEDEN

About DPAC

The DPAC profile establishes a leading research profile targeting dependable platforms for autonomous systems and control, hosted at Mälardalen University. DPAC is organized through close collaboration and interaction between several research groups at MDH and a set of participating industrial companies. The profile will leverage our solid track record of close cooperation to conduct excellent research, knowledge transfer, and support commercialization with industrial partners. DPAC shall create synergy effects between the partners and a significant increase in coproduction is to be expected.

The ultimate goal of the DPAC profile is to establish a nationally leading and internationally renowned research centre that facilitate close cooperation between academia and industry to achieve a significant increase in research and available knowhow on advanced dependable platforms for embedded systems. Embedded computer systems are nowadays incorporated in many kinds of products including many mission critical applications such as trains, autonomous utility vehicles, aviation, smart grid power management etc. These systems offer advanced functionality and serve an important role for the competitiveness of companies and the future national and global infrastructure. The scientific and technical results of DPAC will support future innovation by providing dependable platforms that can be used to efficiently realize dependable, reliable and safe electronically controlled products.

CONTENTS

ABB ALTEN Arcticus Systems BOMBARDIER ENEA ERICSSON SAAB Senseair AsahiKASEI UNIBAP VOLVO Volvo Group

MÄLARDALEN UNIVERSITY SWEDEN

Knowledge Foundation

Welcome to the DPAC Newsletter for spring 2020

This spring we opt to share information between DPAC partners using a newsletter instead of the previously planned DPAC spring summit, which was canceled due to Covid-19. The newsletter highlights what is currently ongoing and has been achieved in each of the four main threads of work in DPAC: The unified demonstrator and the three subprojects. In addition to these main activities, the DPAC core-team is preparing for an industrial research school organized around DPAC's main topics and a unifying use-case to which each student's project can contribute.

During 2019 DPAC underwent a half-time evaluation by the KK Foundation. The evaluation was a success and DPAC secured support for the second phase which started during fall of 2019 and will continue until fall 2023. During its first phase, DPAC produced 189 scientific publications and employed 25 researchers. We organized 7 summits, where all DPAC partners from the 12 participating companies and MDH participated. The summits typically run for 25 hours and include both scientific and inspirational lectures, as well as, plenty of time for social interaction and networking. In addition, the DPAC profile have grown to include over 25 additional research projects.

DPAC will continue as planned, with researchers and industrial partners working together towards dependable platforms for autonomous systems and control, and as soon as times allow, we will continue with organizing summits and other joint events.

All DPAC publications can be found via the DPAC homepage: http://www.es.mdh.se/dpac/

In phase 2 DPAC will extend its scope to system of systems, implying (e.g.): the systems will be made up from more than one node (i.e., they will be distributed systems); the systems need to be able to collaborate with each other (and with humans) towards a common goal; and techniques for system verification and validation has to be extended to manage not only one complex system, but also a system of interconnected systems. System of systems are becoming more common and are often associated with variable degrees of autonomy that introduces new research challenges described in the respective work packages below.

The fundamental direction of research performed in phase 1, remains for phase 2 of DPAC and we will continue to organize the profile around the demonstrator and three research areas with one project for each area.

P1: PREDICTABILITY AND DEPENDABILITY IN PARALLEL ARCHITECTURES: This area addresses challenges in designing predictability and dependability in parallel architectures for dependable electronic platforms and challenges in developing dependable and predictable software platforms that execute on such electronic platforms. **Contact:** Saad Mubeen, saad.mubeen@mdh.se

P2: AUTONOMOUS SYSTEMS AND CONTROL: This area addresses challenges with respect to achieving dependability in autonomous systems. We specifically target control intensive systems that should operate in a dependable, reliable, and safe way without a human operator providing detailed control. **Contact:** Mikael Ekström, mikael.ekstrom@mdh.se

P3: DESIGN METHODOLOGIES: This area addresses challenges associated with design methodologies used for developing dependable platforms and systems. It includes methodologies to capture and validate correct sets of requirements and design assurance, and it addresses how we can simplify the complex dependable embedded system models and make analysis more tractable. **Contact:** Håkan Forsberg, hakan.forsberg@mdh.se

DEMONSTRATOR: This area addresses the design and implementation of a predictable autonomous wheel loader, which acts as the unifying use case in DPAC. The wheel loader is based on applying selected results of the research conducted in P1, P2, and P3, with respect to dependable architectures, autonomous intelligent algorithms, and design methodologies. **Contact:** Cristina Seceleanu, cristina.seceleanu@mdh.se

Improving on-board data processing using CPU-GPU heterogeneous architectures for Real-Time Systems

Nandinbaatar Tsog's Licentiate thesis, Mälardalen University, Västerås, December 18, 2019.

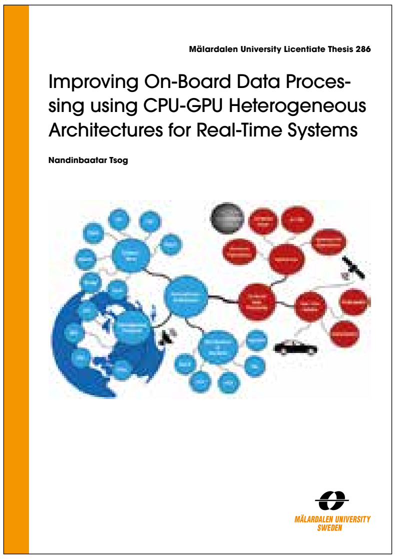

A technique for alternative execution of parallel segments in heterogeneous architectures

The alternative executions of parallel segment eliminates the bottleneck caused by overuse of accelerators such as GPUs. Our preliminary results indicate that up to 90% of improvement in the schedulability of task sets can be achieved as compared to traditional use of parallel segments.

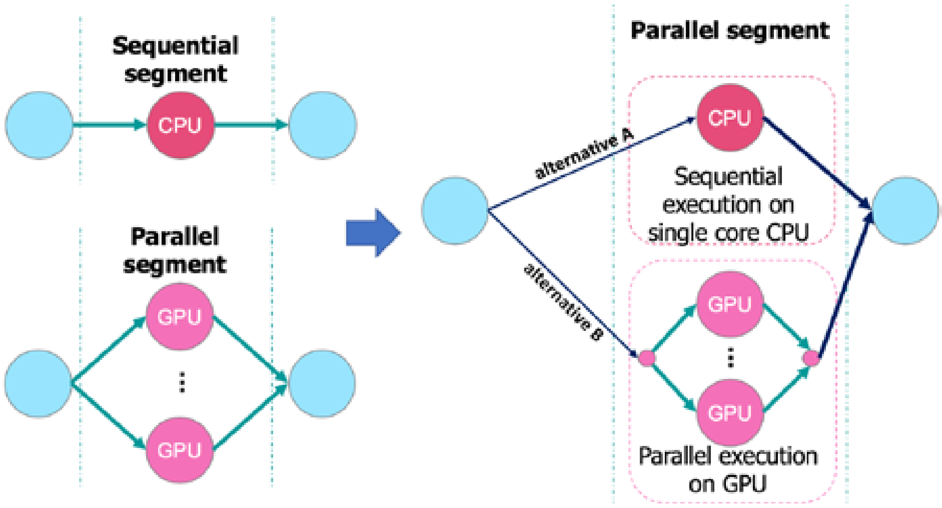

Heterogeneous computing takes the cloud higher and into space

Unibap's collaboration with MDH on heterogeneous computing for space environment is enabling cloud computing in-orbit. Unibap is working with the European Space Agency to standardize SpaceCloud™ services for cloud computing on space missions. This include intelligent data processing and data storage and management similarly to ground based cloud computing.

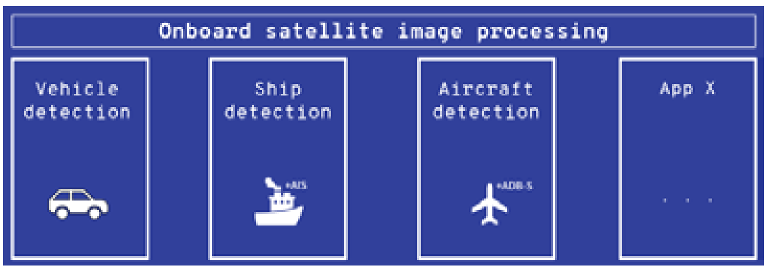

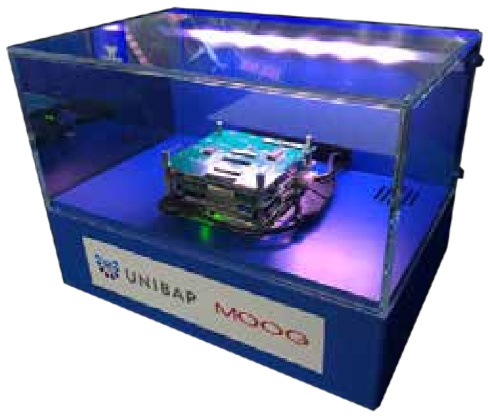

Figures and corresponding material courtesy of Unibap AB.

P1 -PREDICTABILITY & DEPENDABILITY IN DISTRIBUTED PARALLEL ARCHITECTURES

Characterization of Shared Resource Contention in Multi-core Systems

Methods for finding performance bottlenecks of algorithms due to shared resource contention

We investigate parallelization bottlenecks of parallel fork-join versions of feature detection algorithms. In collaboration with Ericsson, we developed methods utilizing performance counters that pinpoints performance bottlenecks that happens as a consequence of shared resource contention.

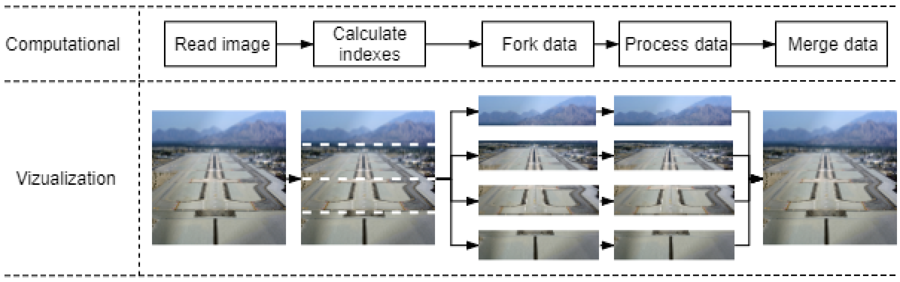

Methods for partitioning systems to mitigate shared resource contention

In collaboration with ENEA, we investigated the jailhouse hypervisor – a virtualization technique used isolate local resources such as CPU and local caches. We furthermore developed a last-level cache partitioning controller using the PALLOC framework in collaboration with Ericsson to optimize cache partition assignments.

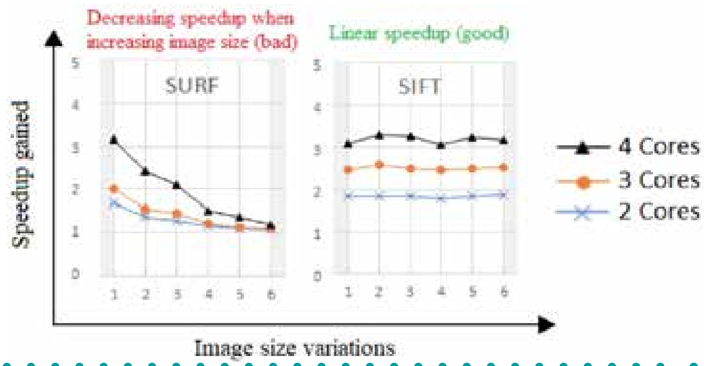

DenseDisp: A Multi-Objective Search Method for Improving the Accuracy of Disparity Estimation on Heterogeneous Platforms

Stereo cameras are multi-modal vision sensors which can extract depth information (disparity map). Neural networks provide the most accurate disparity map results. However, execution of neural vision algorithms needs huge computing capacity which is extremely challenging for real-time constraints and resource-limited heterogeneous hardware. DenseDisp proposes a multi-objective and fast Neural Architecture Search (NAS) method that discovers hardware-friendly neural networks by considering network accuracy and network floating-point operations as the search objective. DenseDisp also deploys the optimized neural network on a wide range of heterogeneous multi/many-core platforms such as Intel® NC2, Nvidia Jetson TX2, Google TPU, and FPGA. The figures illustrate the results of DenseDisp for predicting the distance of objects for autonomous vehicles. Our results indicate that DenseDisp provides up to 9.4x faster execution time while losing only 5% accuracy compared to the state-of-the-art results on heterogeneous many-core platforms.

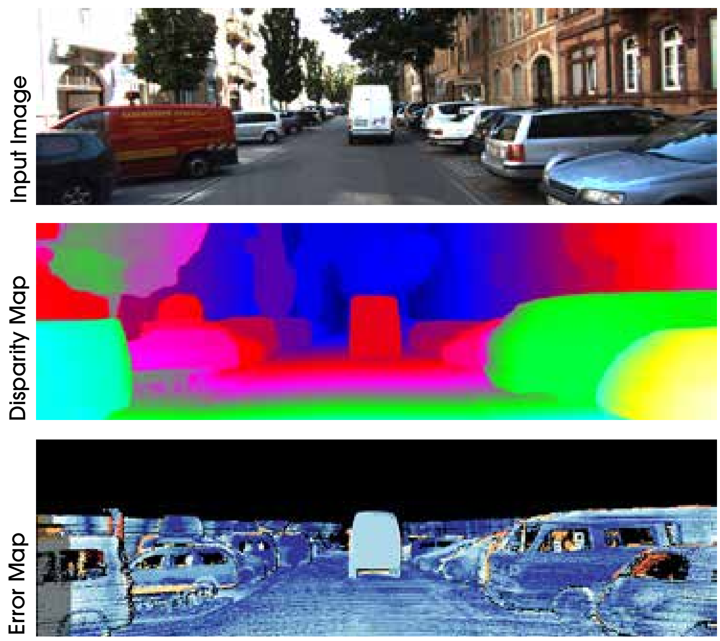

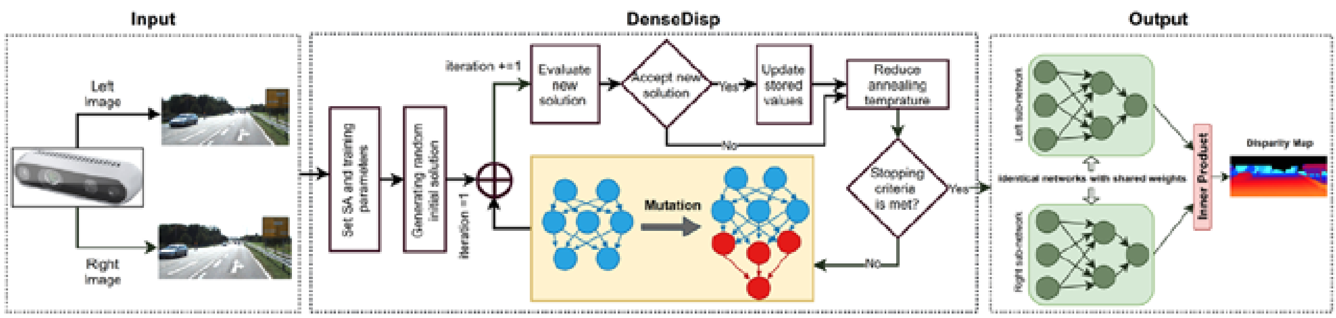

P2 -AUTONOMOUS SYSTEMS AND CONTROL

Dependable multi-path planning with obstacle avoidance for multiple robotic agents

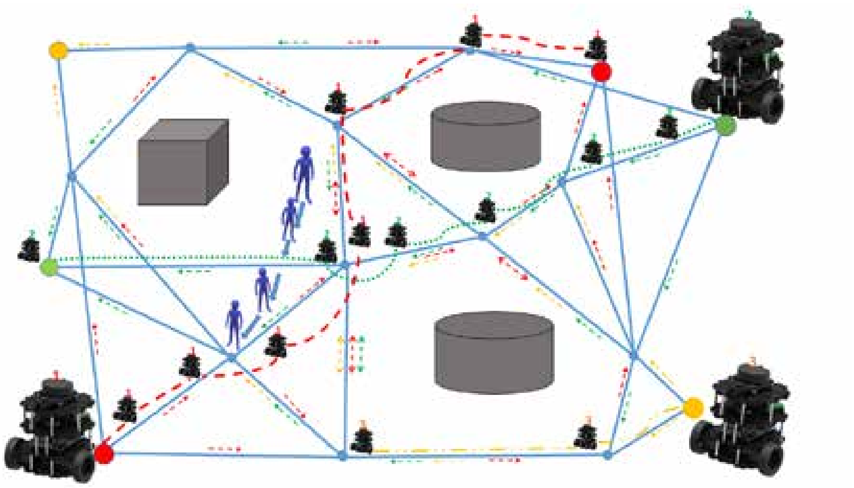

Navigation and path planning is one of the grand challenges in robotics. It has implication for robotic applications in numerous cases in the professional and civil settings including autonomous driving, autonomous control, automated warehouse systems, etc. The main aim here is to build a dependable, i.e. safe, reliable and effective, path planning algorithm for a group of fully autonomous robots that share their working space with humans. In this project, path planning for multiple co-existing robots is combined with moving obstacle avoidance to develop congestion control. The objective is to implement fault tolerance and to ensure safe navigation of robots to avoid collisions with operators, other robots and moving objects in working environment. So far, successful implementation of the dipole flow field for obstacle avoidance and Petri Net for fault tolerance analysis and congestion control algorithms has been implemented and demonstrated. Dependable attributes of proposed algorithms have been evaluated in the Gazebo simulator. In addition, the algorithms have been implemented and tested with the robotic operating system (ROS) and real robots (Husqvarna, TurtleBot3, etc.) as well.

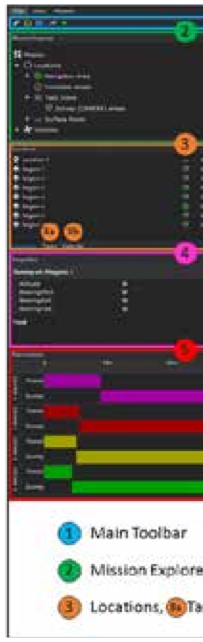

Autonomous agent architecture for testing the multi-path planning algorithm

The core modules and their functions are Map generation for 2-D binary maps with static objects and obstacles, Path initialization for the path from start to destination, Static flow field configuration for driving the agent back to the designed path, Collision avoidance for agents and Velocity planning for adapting the agent's velocity to the environment. There are also external modules: Sensor data collection gathers information from the environment, the internal model based on the changes in the environment are managed and applied to the control commands in Update, The Object classification module receives the data from the Sensor data collection module to determine which objects in the environment that are static objects and which ones are moving objects

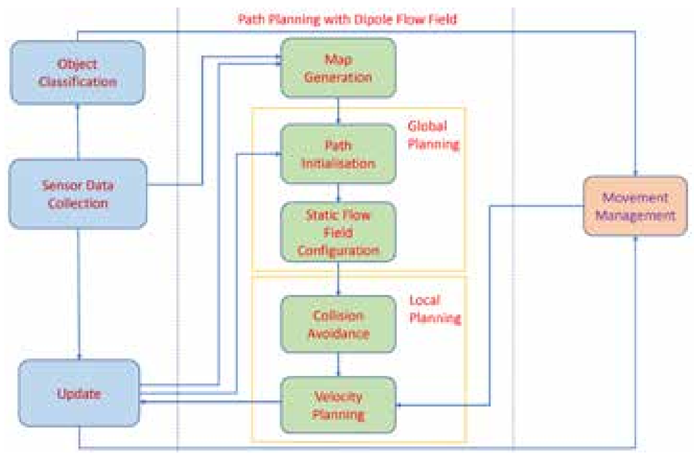

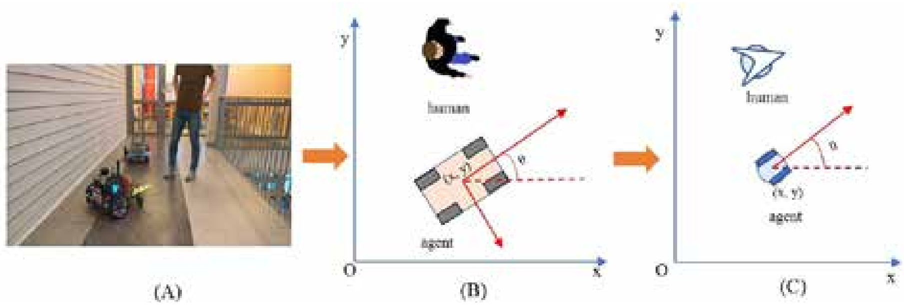

Visualisation of an agent with kinematic parameters and human from (A) a real world space in (B) a 2D mapping space, and (C) a simplified visualization used in the proposed work.

P2 -AUTONOMOUS SYSTEMS AND CONTROL

Path simulating, multi-robotic agent orchestration, and mission monitoring and supervision

The Mission management tool (MMT) is a unifying visualisation and interaction tool for critical multi-robot missions. It allows an operators to manage complex mission through vehicle configuration, multi-agent plan generation, during the mission monitoring as well as intervention. The MMT is thus critical for improving the situational awareness of the operators. In DPAC, the MMT also functions as the interface to the software solutions for simulations and generation of dependable paths for multi-robots in complex environments. The MMT allows plug-in solutions for example map integration, and communication with mobile solutions. It is also integrated with generic unmanned aerial vehicles. For demonstration purposes a DJI Mavic Air 2 is used.

MMT - Mission view

Main functionalities are path simulation and path generation (Project 3), task-level path planning. orchestration, monitoring

Main Toolbar (1): Provides access to general tools and functions.

Mission Explorer (2): This panel contains all the mission data and relationship between them. The operator will define mission goals by dragging and dropping mission items into this panel.

Locations, Tasks and Vehicles (3): Contains a list of all user-defined locations, tasks and vehicles.

Properties (4): Provides a list of all user-definable properties that a selected object (location, vehicle, task) might have. This can be used by the operator to for example change a specific task parameter, or the colour used to represent a region on the map.

Plan Outline (5): A Gantt chart that presents all the missions, both the ones that are running in real time as well as the ones planned for the future

Map & Map Tools (6): Provides a map of the area and tools to mark points or regions of interest. The mission plan, vehicle locations, and visualizations of some mission results can be overlaid on this panel.

MMT - Asset view

The asset management view helps to visualise how the assets/vehicles have moved during the course of a mission or a working day. Information on speed, fuel, etc. can be seen to contribute to the operator's situation awareness.

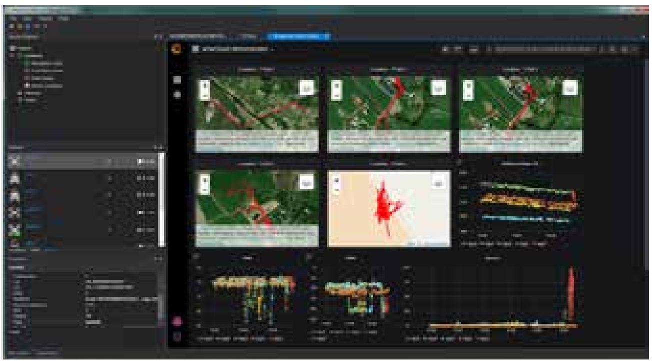

MMT - Data visualisation view

This view allows visualisation and analysis of data collected from indoor and outdoor sensors pre-, during, and post-mission. It also allows visualisation of fused data, or data from other resources.

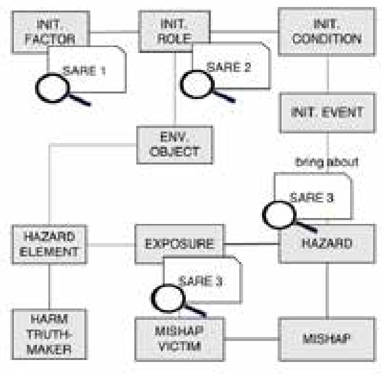

Hazard Analysis, Requirements Elicitation and Validation

In the first phase of DPAC we developed an ontology (called Hazard Ontology), that can be used to perform a structured hazard analysis, along with a method to elicit safety requirements (called Safety Requirements Elicitation approach, in short SARE) based on the ontology. Even though the ontology is general, in the sense that it can be applied to various types of systems, it was however developed primarily for traditional single systems. As a result, SARE is used to discover the safety requirements that mitigate the hazards identified for these systems. In the second phase of DPAC we have applied the ontology on the quarry use-case to evaluate if the ontology can be used on System of Systems (SoS). Preliminary results indicate that it can be used on SoS, that is, it supports the common characteristics of those systems and that we are able to identify emergent hazards. We have also proposed a method to abstract system descriptions to further facilitate the use of the ontology on SoS, and to capture the safety requirements that mitigate hazards in this kind of systems. The safety requirements, elicited through SARE, are specified in a way so that they can easily be translated into formal verification methods. In addition, with the help of our industrial partners we are continuing our work on tool support that will facilitate the use of the ontology.

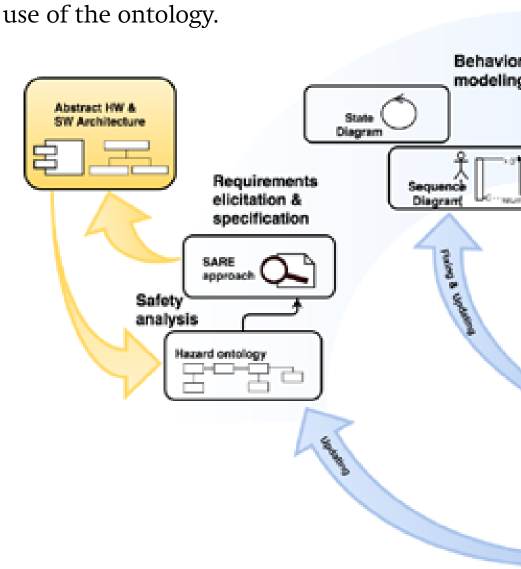

Fault-Tolerant Deep Neural Networks for Autonomous Systems

The use of Deep Neural Networks (DNNs) in safety-critical applications requires a reliable and secure platform but also a reduction of systematic faults, e.g., design faults in both hardware and software. In previous research, we have suggested the use of Overarching Properties (intent, correctness and acceptability) together with assurance cases to argument that assurance objectives can be met for future computing platforms, including those based on machine learning.

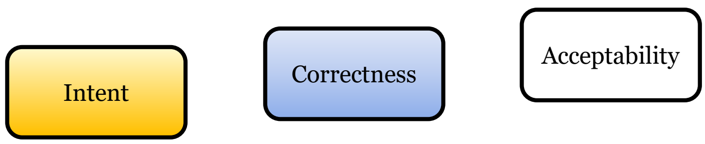

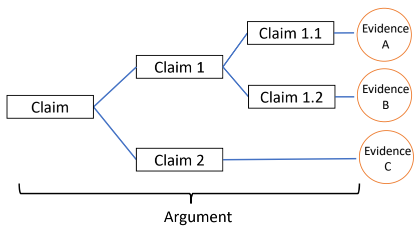

A graphical presentation of an assurance case. The top-level claim (leftmost) is decomposed until each sub-claim can be substantiated by evidence. The argument part consists of strategies used to decompose claims and sub-claims.

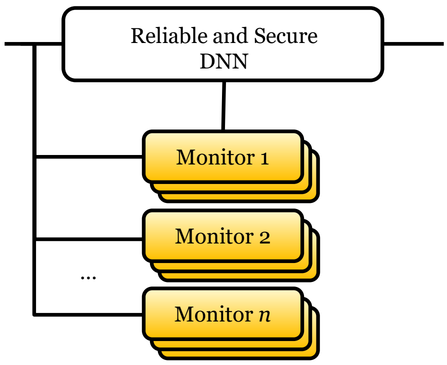

A reliable and secure DNN backed up with diverse fault tolerant architectures, each specified for a task such as detecting transients in the DNN, adversarial attacks, data input distortions, untrained input data, reduced false negatives or false positives

In the case of image classification using DNNs, correct classification is only one concern for safety. In some scenarios it is equally important to minimize incorrectly classified images.

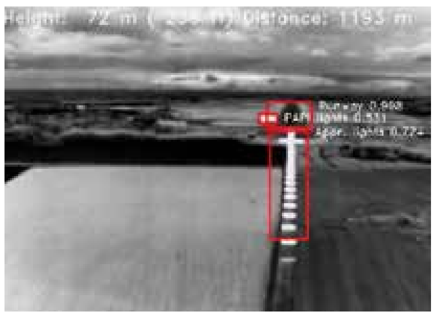

Guided approach – an example of multiple different object detection scenarios

Consider for instance the scenario of an aircraft to perform machine vision-guided approach. In this case, it is of high importance not to mix a highway for a runway. Once the aircraft is approaching the runway it is more important to detect obstacles even if they are incorrectly classified rather than not detecting them at all (it is of little importance whether we detect a motorcycle or a truck on the landing strip). We believe diverse redundant systems are needed to cope with the scenarios above. These systems may or may not include the time domain (i.e. history of classified objects and moving targets) and may consist of deterministic or statistical monitors. Additional redundant architectures may be necessary for symmetric faults.

Verification-Driven Iterative Development of Cyber-Physical System

We used an iterative and incremental approach to build formally verified models. The actor-based textual modeling language, Rebeca, with model checking support is used for formal verification. Starting from structured requirements and system architecture design, the behavioral models, including Rebeca codes, are built.

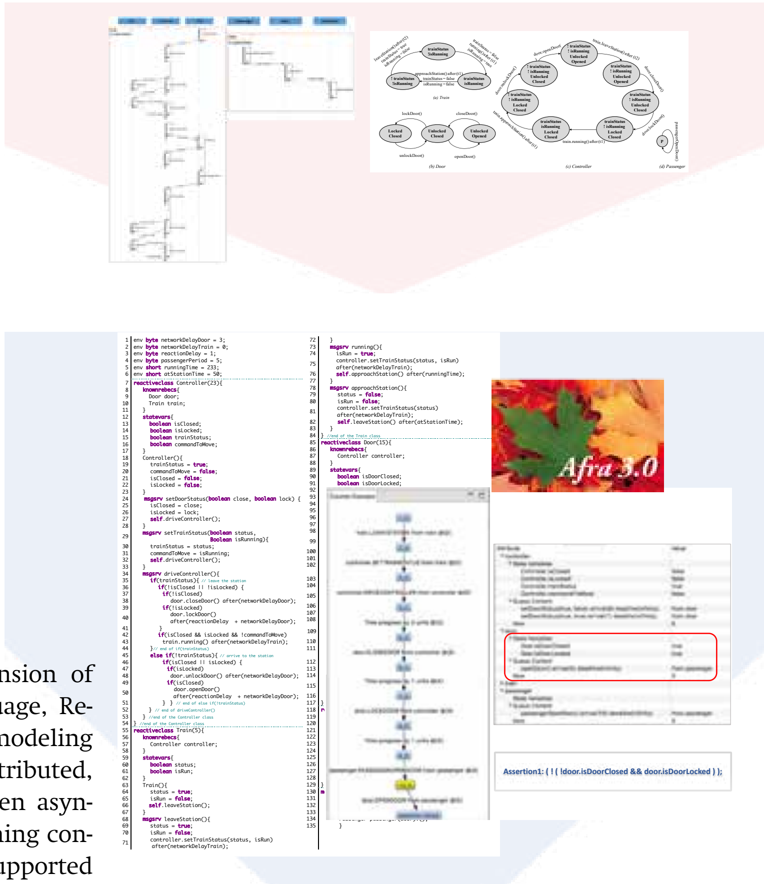

Properties of interest are also derived from the structured requirements, and then model checking is used to formally verify the properties.

This process can be performed in iterations until satisfaction of desired properties are ensured, and possible ambiguities and inconsistencies in requirements are resolved. The formally verified models can then be used to develop the executable code.

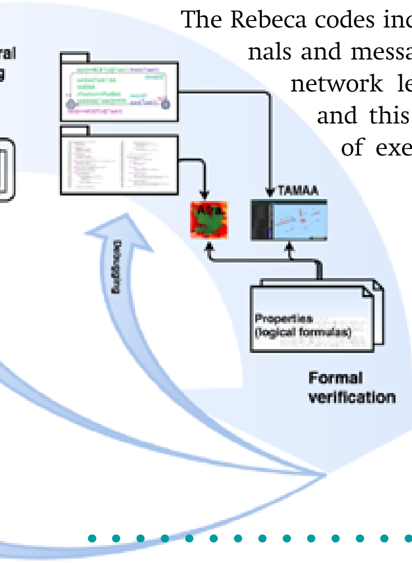

The Rebeca codes include the details of the signals and messages that are passed at the network level including the timing and this facilitates the generation of executable code.

Timed Rebeca is an extension of the Reactive Object Language, Rebeca, and is designed for modeling and verification of distributed, concurrent and event-driven asynchronous systems with timing constraints. Timed Rebeca is supported by a model checking tool.

TAMAA: Timed Automata Based Mission Planner for Multiple Autonomous Agents

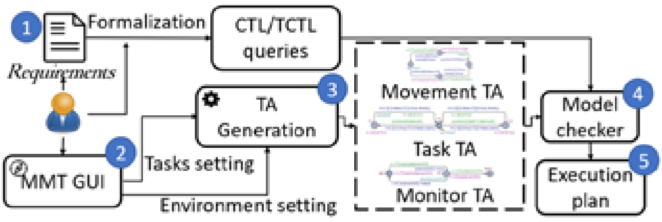

Mission planning for autonomous agents (e.g., vehicles, mobile robots, etc.) involves path planning and task scheduling. While path planning is supported by some algorithms, e.g., Dijkstra, A*, Theta* algorithms, its correctness needs to be guaranteed in the context of autonomous agents. Task scheduling is a well-known NP-hard problem that remains outstanding for decades especially when facing multi-agent systems.

We design a timed-automata-based mission planner for multiple agents to solve this problem, namely TAMAA. This approach is based on a state-of-the-art model checker for real-time systems, called UPPAAL, and is associated with a GUI called Mission Management Tool (MMT) that is developed at Mälardalen University.

The figure above depicts the process of TAMAA method. First, users specify the requirement of autonomous vehicles and configure the environment and tasks for the vehicles in MMT (steps 1 and 2). Next, TA models that represent the movement and task execution of agents, as well as the monitors for special events are automatically generated by TAMAA, based on the information of the environment and tasks (step 3).

These TA models are then verified in the UPPAAL model checker against the requirements that are formalized as CTL/TCTL ((Timed) Computation Tree Logic) queries (step 4). The model checker will generate witnesses of execution that satisfy or violate the requirements, which are used to synthesize mission plans (step 5).

To enable TAMAA to solve multi-agent systems that contain a large numbers of agents, we adopt reinforcement learning and manage to synthesize comprehensive mission plans that consider all possible execution and movement time of the agents. In short, a Q-learning algorithm processes the traces obtained by simulation in UPPAAL, and populates a Q-table, which is then used to form a new restricted model.

We show that the method is applicable to complex scenarios like autonomous quarries, overcoming the scalability problem that is not solved otherwise.

DEMONSTRATOR
-AUTONOMOUS WHEEL LOADER FROM VOLVO CE

AWL Scenario and Challenges

Scenario: The use case focuses on autonomous wheel loaders (AWLs) that are driver-less vehicles that are utilized to transport materials in a quarry site. An AWL digs a given stone pile and carries an amount of stones to a primary crusher that crushes the stones at given fractions, after which the crusher unloads the stones onto the conveyor belt. Next, the AWL moves to the other end of the conveyor belt and loads the crushed stones. It then continues moving to the secondary crusher to unload the stones, which finishes one round in the job cycle. The AWL must avoid static obstacles (e.g., holes and rocks on the ground) as well as possible dynamic obstacles (e.g., other mobile machines or humans), in an optimal manner in order to keep productivity high.

Design and implementation challenges: (i) dependable and accurate algorithms for obstacle detection, path planning, and collision avoidance (P2), (ii) design-time verification of requirements and analysis of behavior assuming unpredictable occurrence of dynamic obstacles (P3), (iii) predictable platform for AWL to run the mentioned algorithms (P1), and (iv) implementation and performance evaluation of the proposed algorithms on the proposed platform (P1, P2, P3).

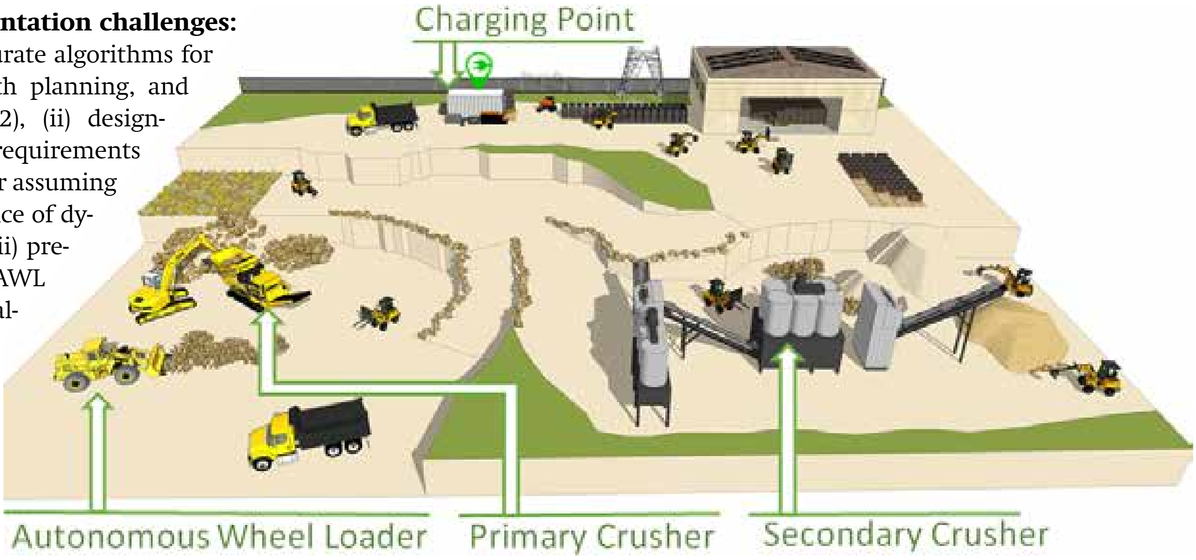

Intelligent Path Planning and Collision Avoidance via Dipole Flow Field Technique

Intelligent Path Planning: We employ the A*/Theta* algorithm to generate an initial shortest path for AWL, from the starting position to the destination. The A* or Theta* algorithm generates smooth paths with few turns, by exploring the map and calculating the cost of nodes by the function f(n) = g(n) + h(n), where n is the current node being explored, g(n) is the Euclidean distance from the starting node to n, and h(n) is the estimated cheapest cost from n to the destination. In each search iteration, the node with the lowest cost among the nodes that have been explored is selected, and its reachable neighbors are also explored by calculating their costs. The iteration is eventually ended if the destination is found or all reachable nodes have been explored.

Dipole Flow Field (DFF) for Collision Avoidance: We model every object as a source of magnetic dipole field, in which the magnetic moment is aligned with the moving direction, and the magnitude of the magnetic moment is proportional to the velocity. DFF has two components: (i) The static flow field force that attracts the vehicle to its goal, ensuring that the vehicle avoids the static obstacles on the map, and (ii) the dynamic dipole field that generates forces that push the vehicle away from the dynamic obstacle, once encountered, based on the latter's respective moving direction and velocity. The obstacle detection relies on data collected from AWL sensors and on executing recognition algorithms based on deep learning, to determine the presence of obstacles.

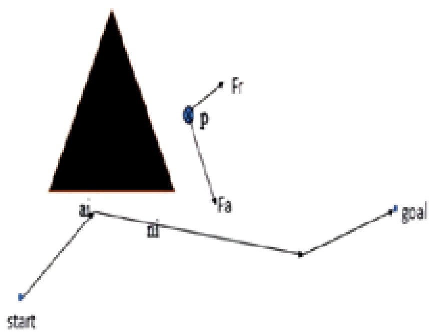

Figure a.

- Assume that the shortest path is the connection of several line segments (Figure a).
- The static flow field is created within the neighbourhood of the line segment (Figure b).
- The flow field force at a point p is given by

$$\boldsymbol{F}_{flow}(\boldsymbol{p}) = F_a(\boldsymbol{p}) + \boldsymbol{F}_r(\boldsymbol{p})$$

$F_a(\boldsymbol{p})$: Attractive forces to the initial path
$\boldsymbol{F}_r(\boldsymbol{p})$: Repulsive forces from static obstacles

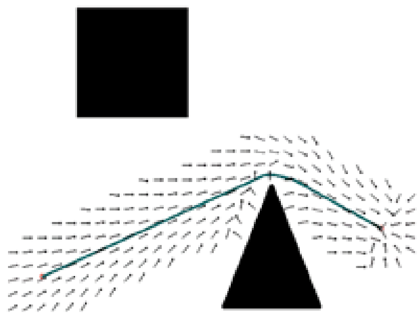

Figure b.

DEMONSTRATOR

-AUTONOMOUS WHEEL LOADER FROM VOLVO CE

Formal Modeling, Verification and Performance Evaluation of AWL Behavior

The AWL has a large set of functional and extra-functional requirements, including end-to-end deadlines for a mission, and energy consumption constraints. In order to ensure the dependability of the AWL at design time, we perform the following:

- Employ our framework TAMAA (Timed-Automata-based planner for Multiple Autonomous Agents) to automatically generate mission plans (including path planning and operation scheduling) for the AWL, by automating the generation of formal models as networks of timed automata (TA) or stochastic timed automata (STA), and integrating the model-checking tool (UPPAAL) and the Mission Management Tool (MMT) in one platform (Figure c).
- Model and verify the AWL's behavior (path computation, obstacle avoidance, end-to-end deadlines) with respect to the synthesized mission plan, and the A*/Theta* and dipole flow field algorithms encoded in UPPAAL (Figure d).
- Formally verify the safety and performance of the AWL's behavior under certain adaptive policies and dynamic changes encoded in Rebeca, by employing the integrated development environment Afra.

TAMAA-based Synthesized Mission Plans for AWL

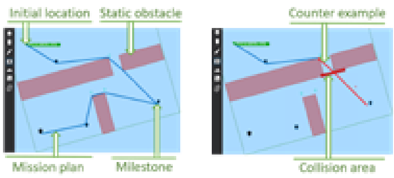

Figure c.

UPPAAL-based AWL Modeling and Formal Verification

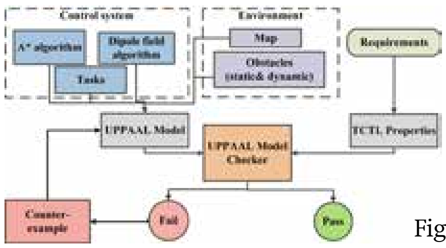

Figure d.

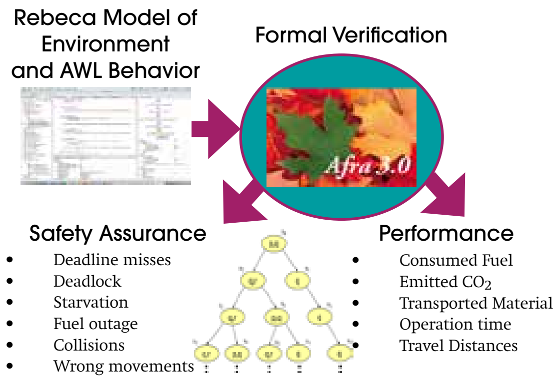

Figure e.

Implementation of the Path Planning and Dipole Flow Field Algorithms on a Heterogeneous System Architecture (HSA)-compliant Reference Platform

To assess the predictability and performance of our solutions, we implement the path planning and dipole flow field algorithms on an HSA architecture (Figure f), which uses 4 CPU cores, 6/8 GPU computing units, all on the same chip. We use an open source run-time API called Heterogeneous-compute Interface for Portability (HIP), with a C++ kernel language. The code is portable on AMD and NVIDIA GPU hardware.

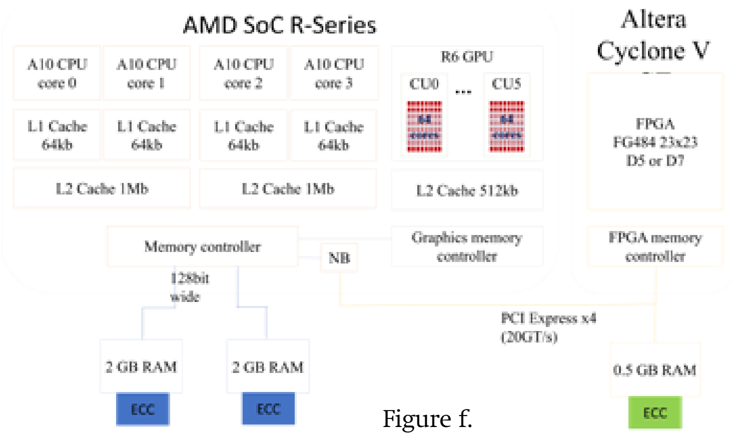

Figure f.

```
__global__
void kernel_Theta2FlowForce(Point2f *s, Point2f *end,  float *segment_len,
        float *x, float *y, float *nx, float *ny,
        int num_of_segment, Point2f *return_force_)
{
  // Calculate the coordinate in which the kernel function runs
  unsigned i = hipThreadIdx_x + hipBlockIdx_x*hipBlockDim_x;

  if(i < num_of_segment) {
    dist = sqrt((s->x - x[i])*(s->x - x[i]) + (s->y - y[i])*(s->y - y[i])) + 1e-12;
    if (dist < min_dist) {  ...  }
      // Wait until all threads finish computing the minimum distance
      // and updating dipole force components
      __syncthreads();
      if(i == 0) {
        // Determine the thread with minimum distance ,
        // obtained from the parallel computation
          min_dist_shared = min_dist_arr[j];
    }}
  // Wait until the minimum distance is determined by the first thread
  __syncthreads();
    if(i == min_index){
        // Calculate attractive force
      // Only the unmasked kernel thread does this
    }}}
```

Figure g.

Solution: Initialize an attractive force (return_force in Figure g) from the start to the goal points of the path in the map grid. Compute the minimum distance, min_dist, in the CPU, and on each generated segment calculate the distance to the goal. This is implemented as parallel GPU threads.

Each GPU thread computes the attractive force only if the distance is lower than min_dist.

The threads synchronize, and only the GPU thread that returns the minimum distance yields the final result to the CPU, the rest are masked.

MÄLARDALEN UNIVERSITY SWEDEN

Knowledge Foundation

About MDH

MDH is one of Sweden's largest HEIs, with 16 000 students reading courses and programmes in Business, Health, Engineering and Education. At MDH, research is conducted within all areas of education to address the challenges of society, and of this the research in future energy and embedded systems is internationally outstanding. MDH's close cooperation with the private and public sectors enables us to help people feel better and the earth to last longer. MDH is located on both sides of Lake Mälaren, with campuses in Eskilstuna and Västerås.

About KKS

The Knowledge Foundation funds research and competence development at Sweden's university colleges and new universities with the purpose of strengthening Sweden's competitiveness. We provide funding when activities are conducted in collaboration between academic staff and business sector partners. The aim is to build internationally competitive, integrated research and education environments. Our mission is to strengthen Sweden's competitiveness, and we know that collaborative projects between academia and industry create great benefits for both parties. The Foundation was established in 1994 with a founding capital of 3.6 billion SEK, and has now invested some 9.3 billion SEK in over 2 500 projects.

DPAC

Profile leader: Kristina Lundqvist
kristina.lundqvist@mdh.se
+46(0)73 960 74 40
www.es.mdh.se/dpac

MÄLARDALEN UNIVERSITY SWEDEN

Address: P.O. Box 883, SE-721 23 Västerås. Sweden
Address: P.O. Box 325, SE-631 05 Eskilstuna. Sweden
E-mail: info@mdh.se **Web:** www.mdh.se