RC3: Resilient Computing and Cybersecurity Center

جامعة الملك عبدالله للعلوم والتقنية
King Abdullah University of Science and Technology

Towards Cyber Resilient Autonomous Vehicle Ecosystems: a historical perspective

https://rc3.kaust.edu.sa

Paulo Esteves-Veríssimo, **Professor, Director**
King Abdullah University of Science and Technology,
Resilient Computing and Cybersecurity Center – RC3
paulo.verissimo@kaust.edu.sa
https://cemse.kaust.edu.sa/people/person/paulo-verissimo

Informal Workshop on Reliable Communication in Autonomous Vehicles' Ecosystems, KAUST, Feb'22

High-level Strategic Objectives and Mission

جامعة الملك عبدالله للعلوم والتقنية
King Abdullah University of Science and Technology

- Relevance and value creation for KAUST and its units
- Relevance and value creation for faculty/researchers
- Relevance and value creation for the KSA

- Become recognised as leaders in the cyber resilience scientific arena
- Promote career progression of members (international and/or local)
- Induce technological and societal transformation in the region
- In the long run, promote digital sovereignty of the KSA.

جامعة الملك عبدالله للعلوم والتقنية
King Abdullah University of Science and Technology

Resilient Computing and Cybersecurity Center

مركز أبحاث الأمن السيبراني والحوسبة الصامدة

PRINCIPAL INVESTIGATORS

Paulo Esteves-Veríssimo

Professor/Director

Cyber Resilience Research Group

Architectures, middleware and algorithms for resilient modular and distributed computing. Security and dependability, autonomous vehicles from earth to space, digital health and genomics, SDN-based infrastructures, blockchain and cryptocurrencies.

Marc Dacier

Professor/Associate Director

SeRBER Lab

Intrusion detection, intrusion tolerance, network security, cybersecurity, threat intelligence, fraud detection.

Marco Canini

Associate Professor

SANDS Lab

Cloud computing, distributed systems and networking. Recent interest is in designing better systems support for AI/ML and providing practical Implementations deployable in the real-world.

Elmootazbellah (Mootaz) Elnozahy

Professor

Fault-tolerance, trusted environments, power management, distributed systems, operating systems,
high-performance computing, computer architecture, simulation tools and recently, cryptography, recently, Cryptography and AI/ML implementations.

Shehab Ahmed Elsayed

Professor

MERGE

Power Conversion and Distribution, Subsurface Mechatronics, Renewable Energy Systems

Suhaib Fahmy

Associate Professor

ACCL

Hardware acceleration, hardware virtualisation, embedded systems and networks security, FPGA cloud computing

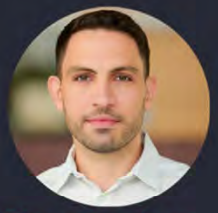

Charalambos (Harrys) Konstantinou

Assistant Professor

SENTRY Lab

Secure, trustworthy, and resilient cyber-physical and embedded IoT systems. Critical infrastructures security and resilience with special focus on smart grid technologies, renewable energy integration, and real-time simulation.

Basem Shihada

Associate Professor

NETLAB

Broadband wired and wireless communication networks, including multi-hop, sensors, cognitive networks, fiber-wireless network integration, optical networks, and green communication protocol. Also, resilience of cyber-physical system infrastructures and control, internet and cloud infrastructure resilience

CybeResil - Cyber Resilience Research Group

Prof. Paulo Esteves-Veríssimo, KAUST, RC3

https://cemse.kaust.edu.sa/cyberesil

The world has become an immense interconnected infrastructure!

... promising great functionalities ...

... under equally great threats ...

HOW TO ADDRESS THIS DILEMMA?

Resilience is the Cybersecurity of the XXI century!

- built-in baseline defences
- incremental and automatic response to threats
- any quality of threat from attacks to errors
- adaptive, self-recoverable and sustainable operation

This research is solving problems in:

- autonomous vehicles from earth to space
- distributed control systems
- smart health and genomics
- e-gov, fintech, blockchain and cryptocurrencies

Early projects on connected and/or autonomous vehicles

MAFTIA - Malicious and Accidental Fault Tolerance for Internet Applications

2000-2003

http://www.research.ec.org/maftia

Computer systems fail for many reasons

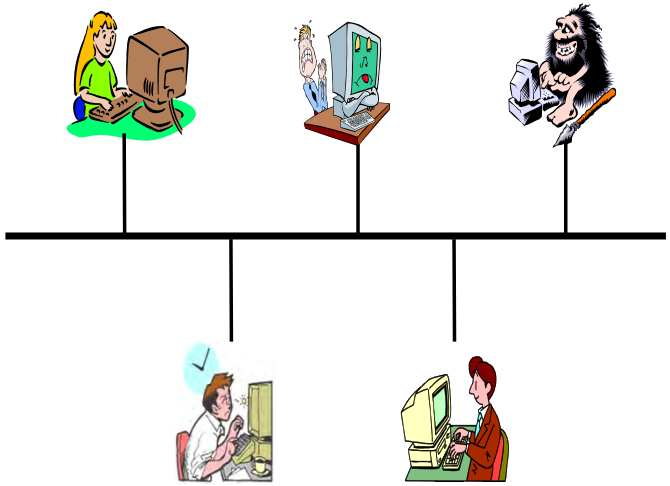

MAFTIA is investigating ways of making computer systems more dependable in the presence of both accidental and malicious faults

- **Architectural framework and conceptual model**
- **Mechanisms and protocols:**
- **Validation and assessment techniques**
- **Partners**
  - **DERA/Qinetiq, Malvern (UK) – T McCutcheon / S Creese**
  - **IBM, Zurich (CH) – Marc Dacier / M. Waidner**
  - **LAAS-CNRS, Toulouse (F) – Y. Deswarte / D. Powell**
  - **Newcastle University (UK)(Coord.) R. Stroud / Brian Randell**
  - **Universität des Saarlandes (D) – Michael Steiner**
  - **Universidade de Lisboa (P) - P Veríssimo / N. Neves**
- **EU coordinator – Andrea Servida**

Paulo Verissimo, Nuno Ferreira Neves, Miguel Correia, "Intrusion-Tolerant Architectures: Concepts and Design", in Architecting Dependable Systems, ser. LNCS. Springer-Verlag, Jun. 2003, vol. 2677, pp. 3–36.

2001-2004

Project Info

INFORMATION SOCIETY TECHNOLOGIES
(IST) PROGRAMME

Project acronym: ***CORTEX***
Project full title:
CO-operating Real-time senTient objects:
architecture and EXperimental evaluation

- **Members**:
  - ***Fac. de Ciências da U. de Lisboa (PT) (proj. coord.),***
  - ***Trinity College Dublin (IR)***
  - ***U. of Lancaster (UK)***
  - ***U. of Ulm (DE)***
- Duration:
  - 3 years, started April 2001
- Budget:
  - 2 MEURO

- **Challenges**:
  - ***Communication, Co-ordination, Control***
  - ***Heterogeneity, Hierarchy and Scope***
  - ***Predictability and Adaptability***

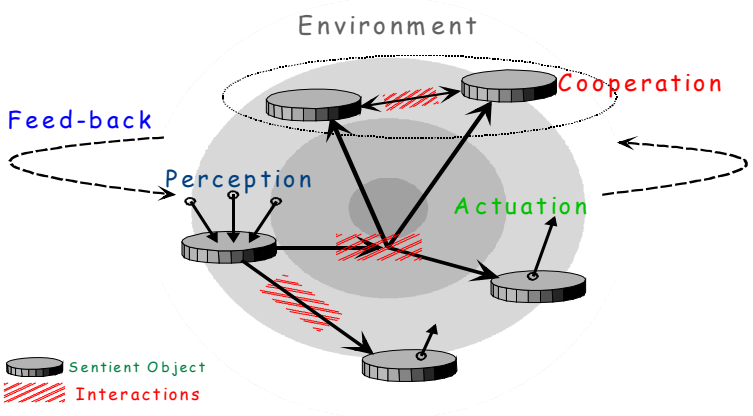

Sentient objects interaction model emulator

- Emulates physical environments in real-time

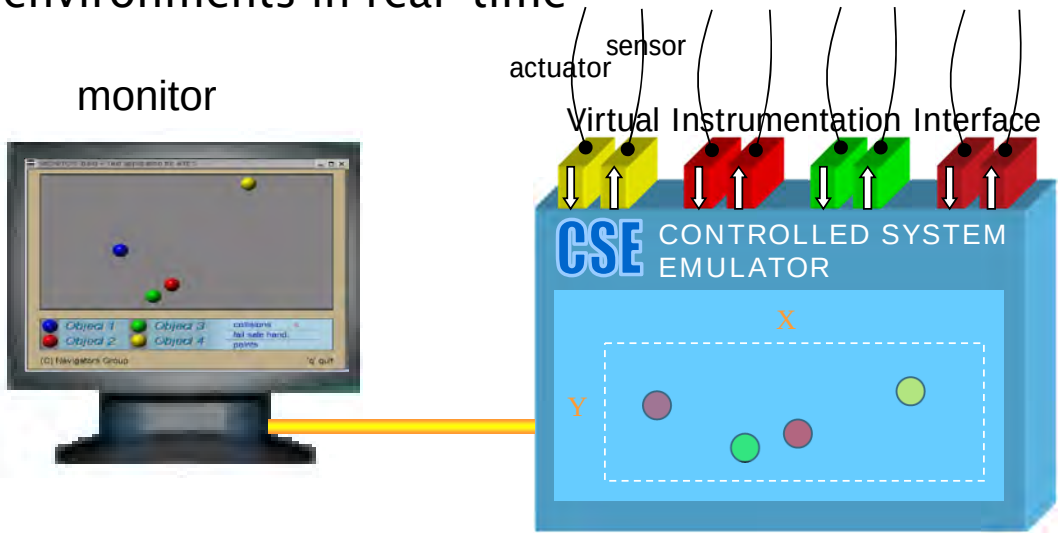

- Emulated environment: four entities shaped as colored balls move in a space with a certain speed and direction
- The Virtual Instrumentation Interface allows to:
  - acquire ball positions, directions and speeds;
  - change ball movement (speed and direction)

Froodvlrq#dyrlgdqfh#lq#uhh#p ryhp hqw#edoov

Multimedia paper, MAIN FEATURE of May 2005 issue of IEEE Distributed Systems On-Line Journal.

A New Programming Model for Dependable Adaptive Real-Time Applications. Pedro Martins et al., IEEE Distributed Systems Online, vol. 6, no. 5, 2005.

Uncontrolled balls

Fail-safe behaviour

Surwrw|sh#p schp hqwdwlrq

2004

- Windows CE / iPAQ pocket PC
- Payload: Windows CE + Regular 802.11b ad-hoc channel
- Windows CE RT tasks + Dedicated 802.11b ad-hoc channel
  - Mockup of a RT protocol (e.g. TBMAC, 802.11e)

HIDENETS

ghly DEpendable IP-based NETworks and Services

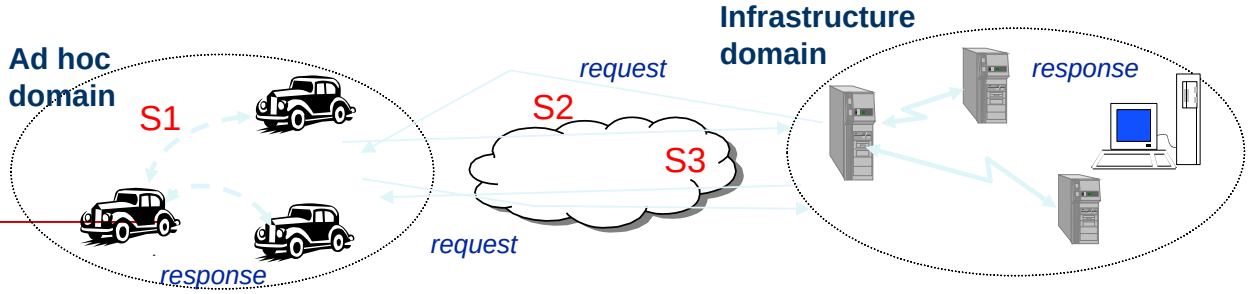

2006-2009

- *Partners from eight countries:*
  - Industry:
    Carmeq (GER), FSC (GER), WMC (NL), Telenor (NO)
  - Academia/research institutions:
    AAU (DK), BME (HU), LAAS (FR), Uni-Fi (IT), FCUL (PT)
- *Develop and analyze end-to-end resilience solutions for AV*
  - for scalable distributed applications and mobility aware services
  - in ubiquitous communication scenarios
  - assuming highly dynamic, unreliable communication infrastructures

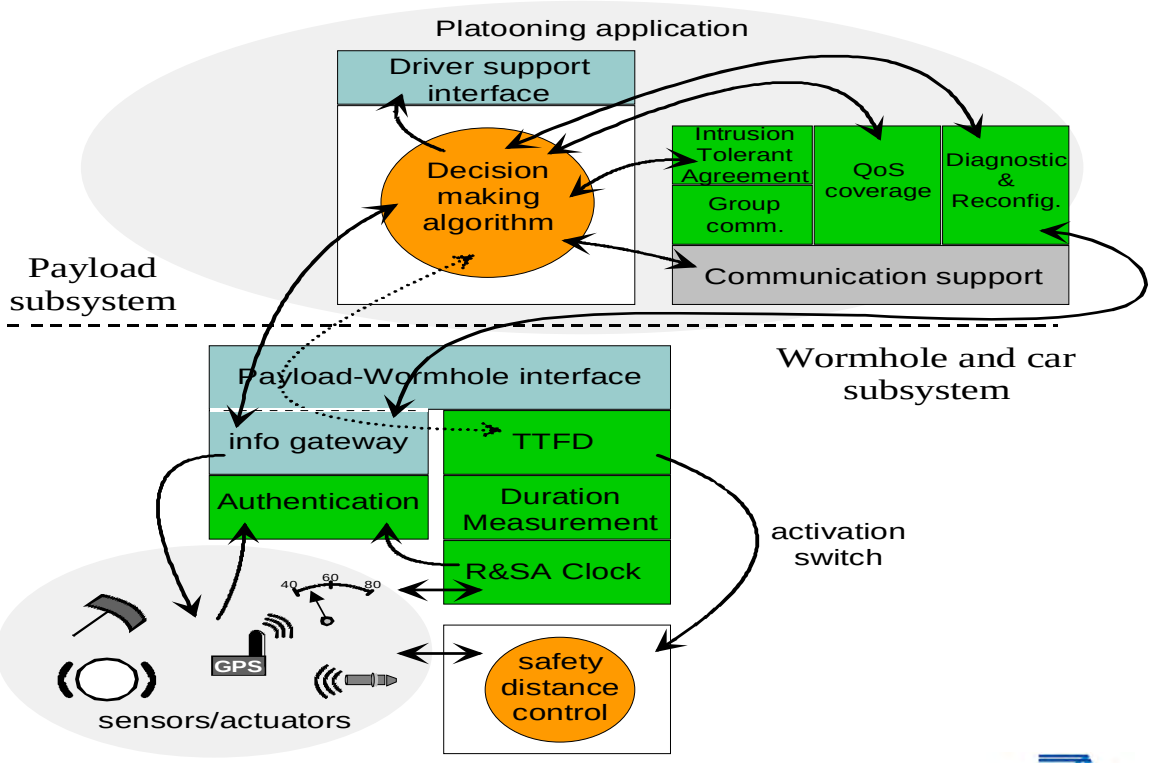

Ghp rqvwudwlrq
V|vwhp #Sodwrrqlqj,#dufklwhfwxuh

KARYON PROJECT : Kernel Based Architecture for safety critical cONtrol

2011-2014

KARYON

SEVENTH FRAMEWORK PROGRAMME

Academia & Research Institutes
SMEs and Industry

Proof-of-concept prototypes
Simulations

FACULDADE DE CIÊNCIAS UNIVERSIDADE DE LISBOA

gmv INNOVATING SOLUTIONS

EMBRAER

OTTO VON GUERICKE UNIVERSITÄT MAGDEBURG

SP Science Partner

CHALMERS UNIVERSITY OF TECHNOLOGY

4S Group Technology for Sustainability

Avionics

UAS/Aircraft flight mission

Automotive

Adaptive cruise control
Coordinated lane change
Coordinated intersection crossing

- Provide system solutions for predictable and safe coordination of smart vehicles that autonomously cooperate and interact in an open and inherently uncertain environment

Vfhqdulrv#vwxglhg

- **Automotive Scenarios:**
  - **Adaptive Cruise Control Systems**
    - More efficient platooning capabilities should improve fuel consumption.
  - **Crossing road intersections**
    - Improved safety measures should help avoid collisions
  - **Coordinated lane change**
    - One of the key collision reasons is the changing of lanes with other vehicles in the driver “blind spot”.
- **Avionics Scenarios:**
  - **Common trajectory traffic in the same direction**
    - Increased usage of air corridors.
  - **Levelled crossing trajectories**
    - Improved safety measures should help avoid collisions.
  - **Coordinated flight level change**
    - Improved safety measures should help avoid collisions.

Intel Collaborative Research Institute for

Collaborative Autonomous & Resilient Systems *(CARS)*

securityandtrust.lu

CritiX

https://www.icri-cars.org/

2017-2020

uni.lu UNIVERSITÉ DU LUXEMBOURG

KEY MESSAGES

Cyberspace today vs.
Research and practice in classic cybersecurity
paradigms and techniques are today insufficient!

- distributed infrastructure:
  - –
- highly exposed to threats:
  - –
- steadily increasing software vulnerabilities:
  - –
- degradation of the threat plane:
  - –

Find the
ECOSYSTEM THREAT SURFACE
in the autonomous vehicles area
... *(land, air, space)*

Autonomous Vehicle Ecosystem

جامعة الملك عبدالله للعلوم والتقنية
King Abdullah University of Science and Technology

2016

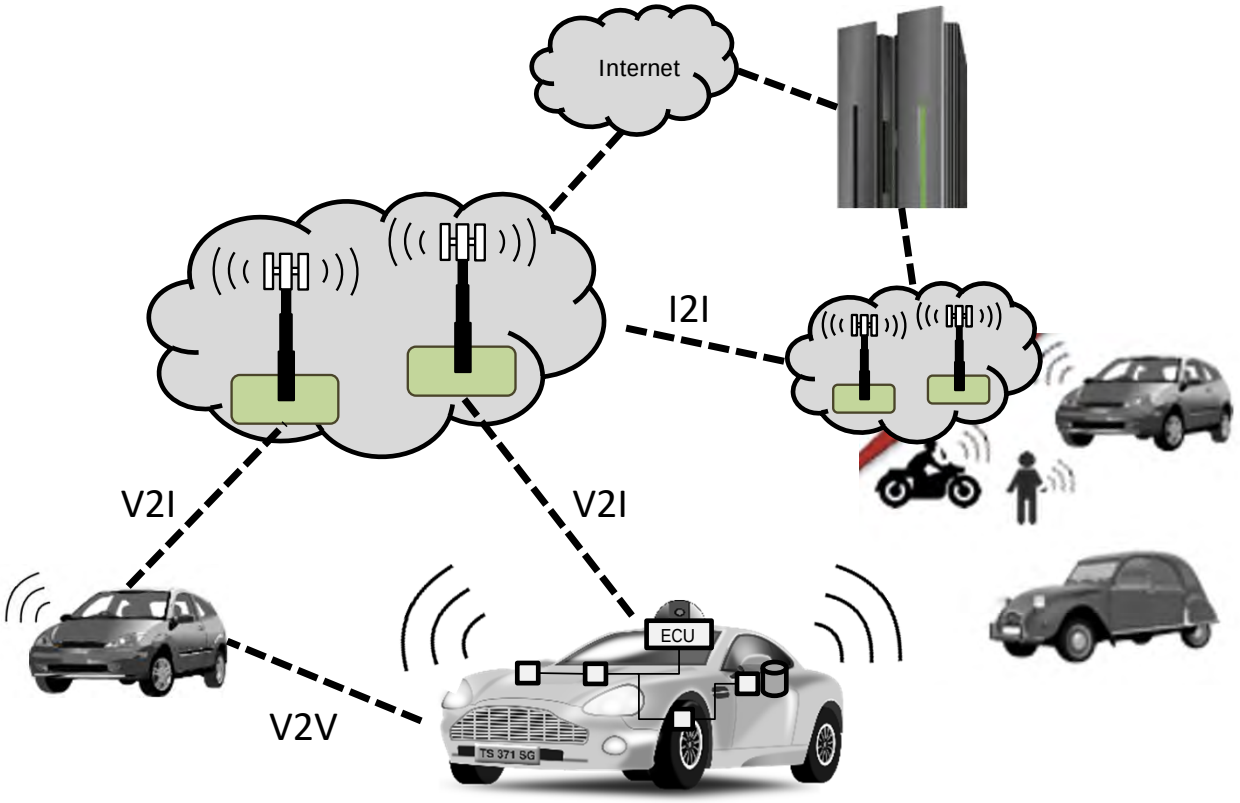

FIRST COMPREHENSIVE STUDY OF THE THREAT PLANE AND SAFETY-SECURITY GAP OF AUTONOMOUS AND COOPERATIVE VEHICLE ECOSYSTEMS

Towards Safe and Secure Autonomous and Cooperative Vehicle Ecosystems. Lima, A; Rocha, F; Volp, M; Verissimo, P. in Proc's 2nd ACM Workshop on Cyber-Physical Systems Security and Privacy (2016, October) @CCS, Vienna-Austria

Autonomous vehicle ecosystem threat surface perhaps wider than many think

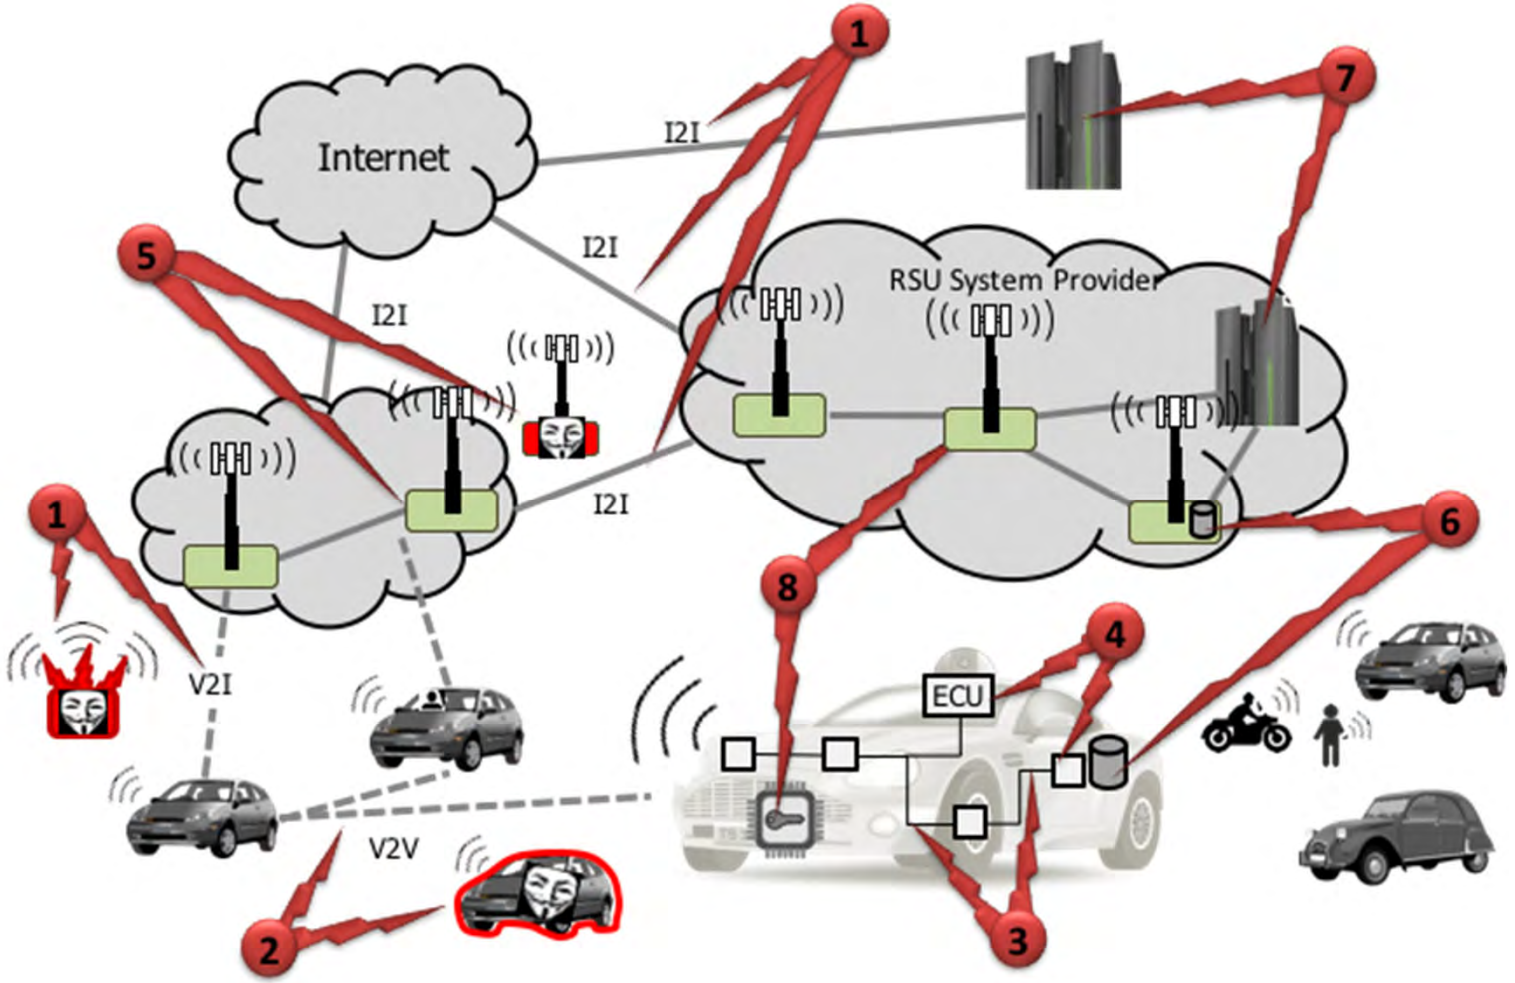

Autonomous vehicle ecosystem threat plane

perhaps wider than many think

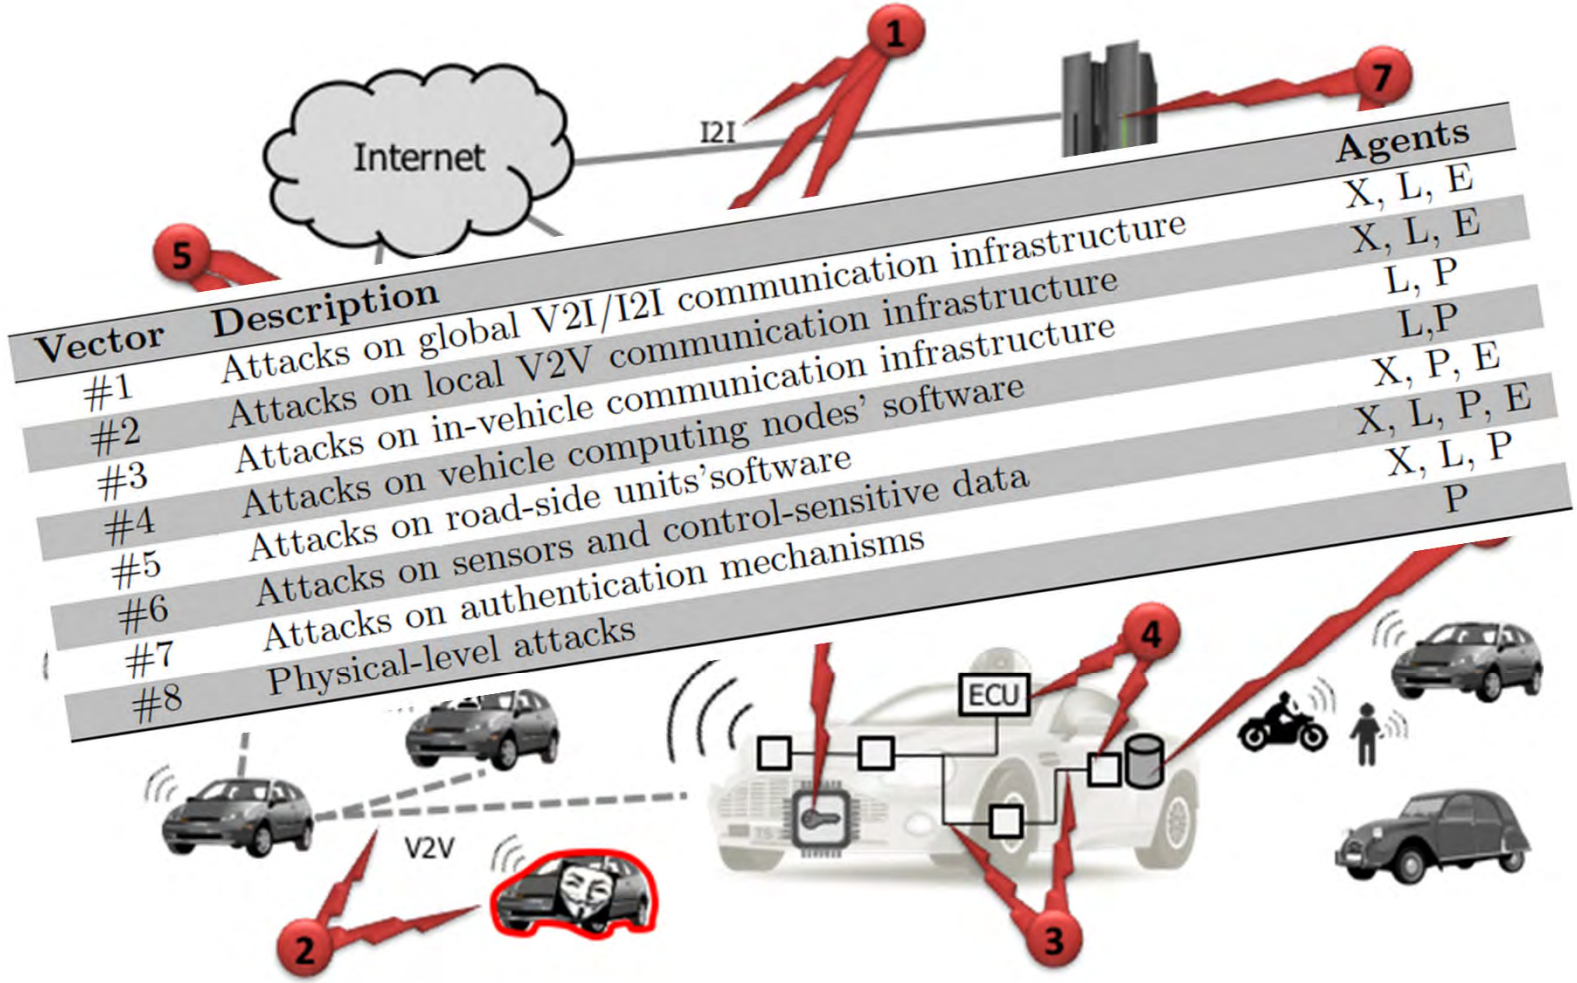

Vector	Description	Agents
#1	Attacks on global V2I/I2I communication infrastructure	X, L, E
#2	Attacks on local V2V communication infrastructure	X, L, E
#3	Attacks on in-vehicle communication infrastructure	L, P
#4	Attacks on vehicle computing nodes' software	L,P
#5	Attacks on road-side units'software	X, P, E
#6	Attacks on sensors and control-sensitive data	X, L, P, E
#7	Attacks on authentication mechanisms	X, L, P
#8	Physical-level attacks	P

The *SAFETY-SECURITY GAP* in the autonomous vehicles area ... ... *(land, air, space)*


Toyota "Unintended Acceleration" Has Killed 89
A 2005 Toyota Prius, which was in an accident, is seen at a police station in Harrison, New York, Wednesday, March 10, 2010. The driver of the Toyota Prius told police that the car accelerated on its own, then lurched down a driveway, across a road and into a stone wall. (AP Photo/Seth Wenig) AP PHOTO/SETH WENIG
Unintended acceleration in Toyota vehicles may have been
in the deaths
of 89 people over the past decade, upgrading the
possibly linked
to the massive recalls, the government


TESLA'S AUTOPILOT HAS HAD ITS FIRST DEADLY CRASH
Woman dead after being struck by self-driving Uber
Pedestrian killed in accident involving self-driving Uber


TESLA FSD 10.4 tries to turn into incoming traffic, AGAIN!


naked security by SOPHOS
Award-winning computer security new
The Jeep hackers retur
ditch a car going 60 mp
03 AUG 2016 5
Security threats, Vulnerability
hacked

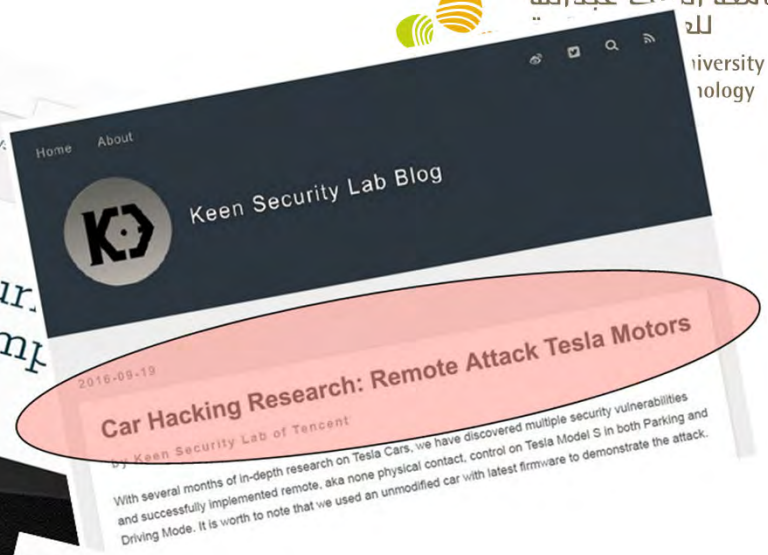
Home About
Keen Security Lab Blog
2016-09-19
Car Hacking Research: Remote Attack Tesla Motors
By Keen Security Lab of Tencent
With several months of in-depth research on Tesla Cars, we have discovered multiple security vulnerabilities and successfully implemented remote, aka none physical contact, control on Tesla Model S in both Parking and Driving Mode. It is worth to note that we used an unmodified car with latest firmware to demonstrate the attack.


naked security by SOPHOS
Award-winning computer security news
SOPHOS.COM
Hackers take over Tesla Model S while car is moving
21 SEP 2016 2
Security threats

Close the Safety-Security gap
in vehicle ecosystems design, verif. & certific.

SnT
securityandtrust.lu
CritiX

Faults in a well designed car lead to an **infinitesimal and acceptable** probability of catastrophic failure;

Vulnerabilities in a car **will** lead, rather sooner than later, to catastrophic failures;

Towards Safe and Secure Autonomous and Cooperative Vehicle Ecosystems. Lima, A; Rocha, F; Volp, M; Verissimo, P. in Proc's 2nd ACM Workshop on Cyber-Physical Systems Security and Privacy (2016, October) @CCS, Vienna-Austria

uni.lu
UNIVERSITÉ DU LUXEMBOURG

Ecosystem approach: Cooperation is key!

Individualistic cars worsen safety!

securityandtrust.lu

TECHNOLOGY NEWS | Mon Feb 29, 2016 | 6:31pm EST

Google says it bears 'sc
after self-driving car hit

pilot crash under federal

Cooperation is key!

TESLA'S AUTOPILOT HAS HAD ITS FIRST DEADLY CRASH

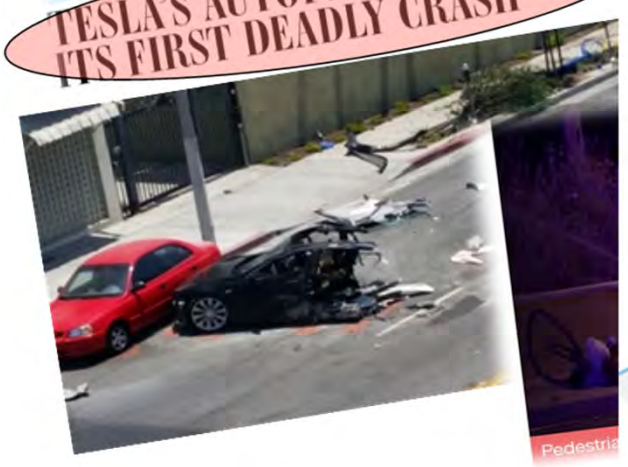

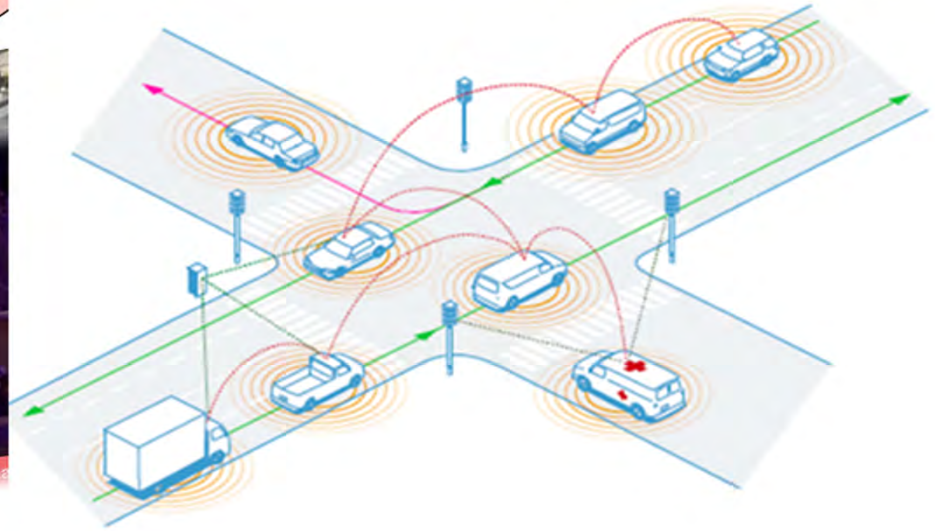

tries to
ing traffic,

Towards Safe and Secure Autonomous and Cooperative Vehicle Ecosystems. Lima, A; Rocha, F; Volp, M; Verissimo, P. in Proc's 2nd ACM Workshop on Cyber-Physical Systems Security and Privacy (2016, October) @CCS, Vienna-Austria

UNIVERSITÉ DU LUXEMBOURG

«Resilience is the Cybersecurity of the XXI century»

«Resilience is the Cybersecurity of the XXI century»

- Classic cybersecurity paradigms insufficient today, *given the challenging cyberspace threat landscape*
- *Resilient Computing*, marrying
  - *security/privacy (against attacks),*
  - *dependability/safety (against accidents),*
  - *and AI/ML (against uncertainty)*

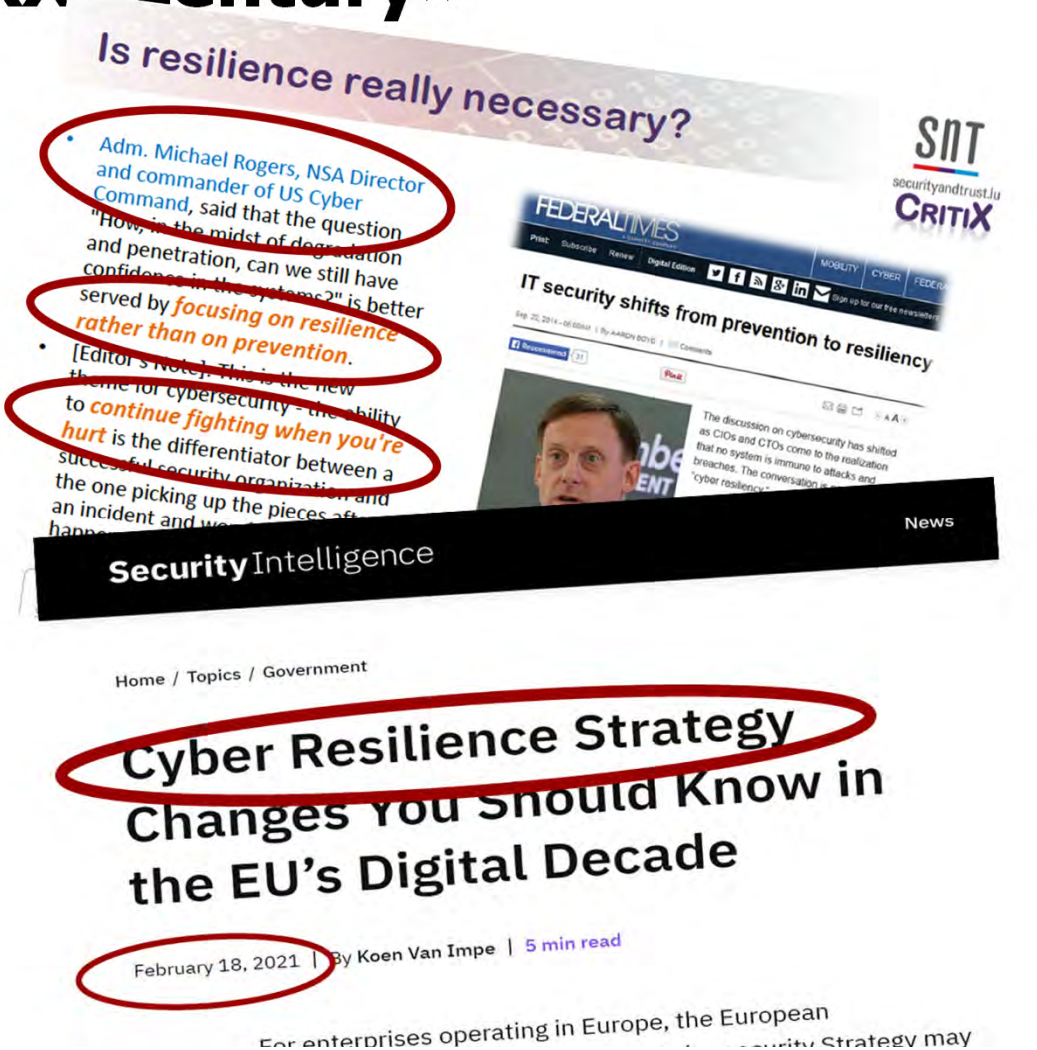

Architecting and designing for resilience

- have built-in baseline defences
  - from first principles and first hour, for whatever mission;
- cope with virtually any quality of threat
  - accidental faults, design errors, cyber-attacks, or unexpected operating conditions;
- protect in an incremental way
  - since not all threats are extreme, or all systems critical;
- automatically respond and adapt to threats
  - in a dynamic range of their severity;
- provide unattended and sustainable operation

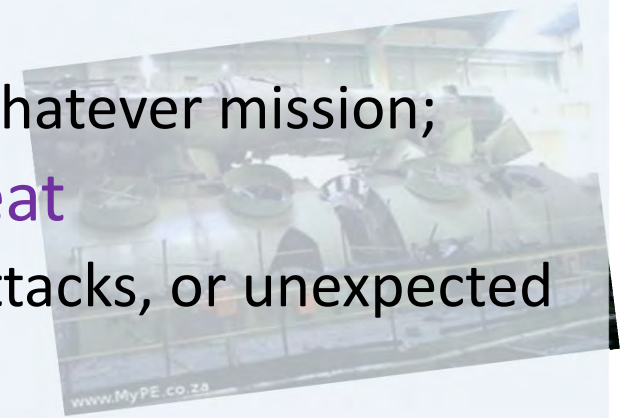

Resilience as general definition -- ***"The property of resuming original shape or position after being bent, stretched, or compressed; elasticity and plasticity."***

Reality behind the AI/ML hype

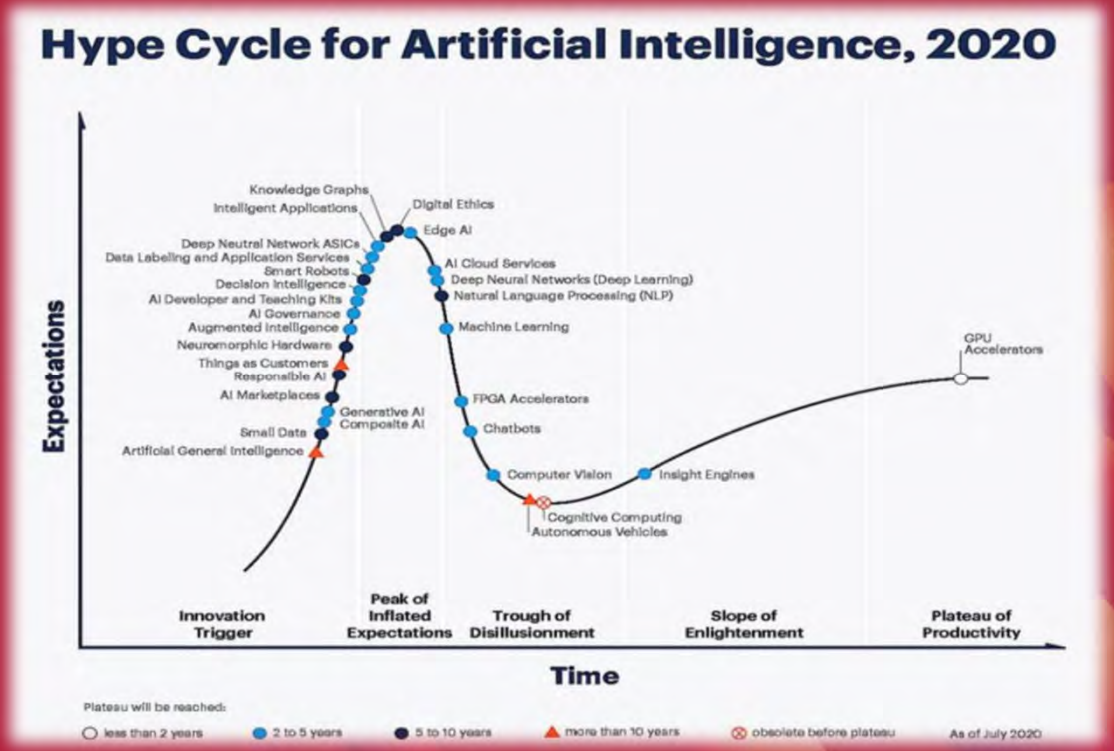

AI vs. Security vs. Safety

AI/ML for Cybersecurity and Safety

Cybersecure and Safe AI/ML

Tesla radar did not recognize a camel, causing an accident in the UAE

جامعة الملك عبدالله للعلوم والتقنية
King Abdullah University of Science and Technology

AI / ML and Resilient Computing ...

- Where Resilient Computing can help AI/ML :
  - *Strike a balance between model-based and statistics-based reasoning*, possibly by leveraging ***architectural hybridisation***
  - *Enforcing deterministic guarantees whilst leveraging the efficiency of deep learning*, possibly by leveraging ***architectural hybridisation***
  - (see PoC Karyon Project slides ahead)

KEY POINTS of ACTION for Cyber Resilient Autonomous Vehicle Ecosystems

Finding Design Solutions!

COMPONENT-BASED, INDIVIDUALIZED

ATTACK PREVENTION, ACCESS CONTROL, FWALLS, ETC.

VULNERABILITY PREVENTION AND REMOVAL

HUMAN-STEERED AD-HOC MITIGATION

Real-Time Intrusion Tolerance & Recovery

Intrusion Detection OK for a start, and for forensics but...

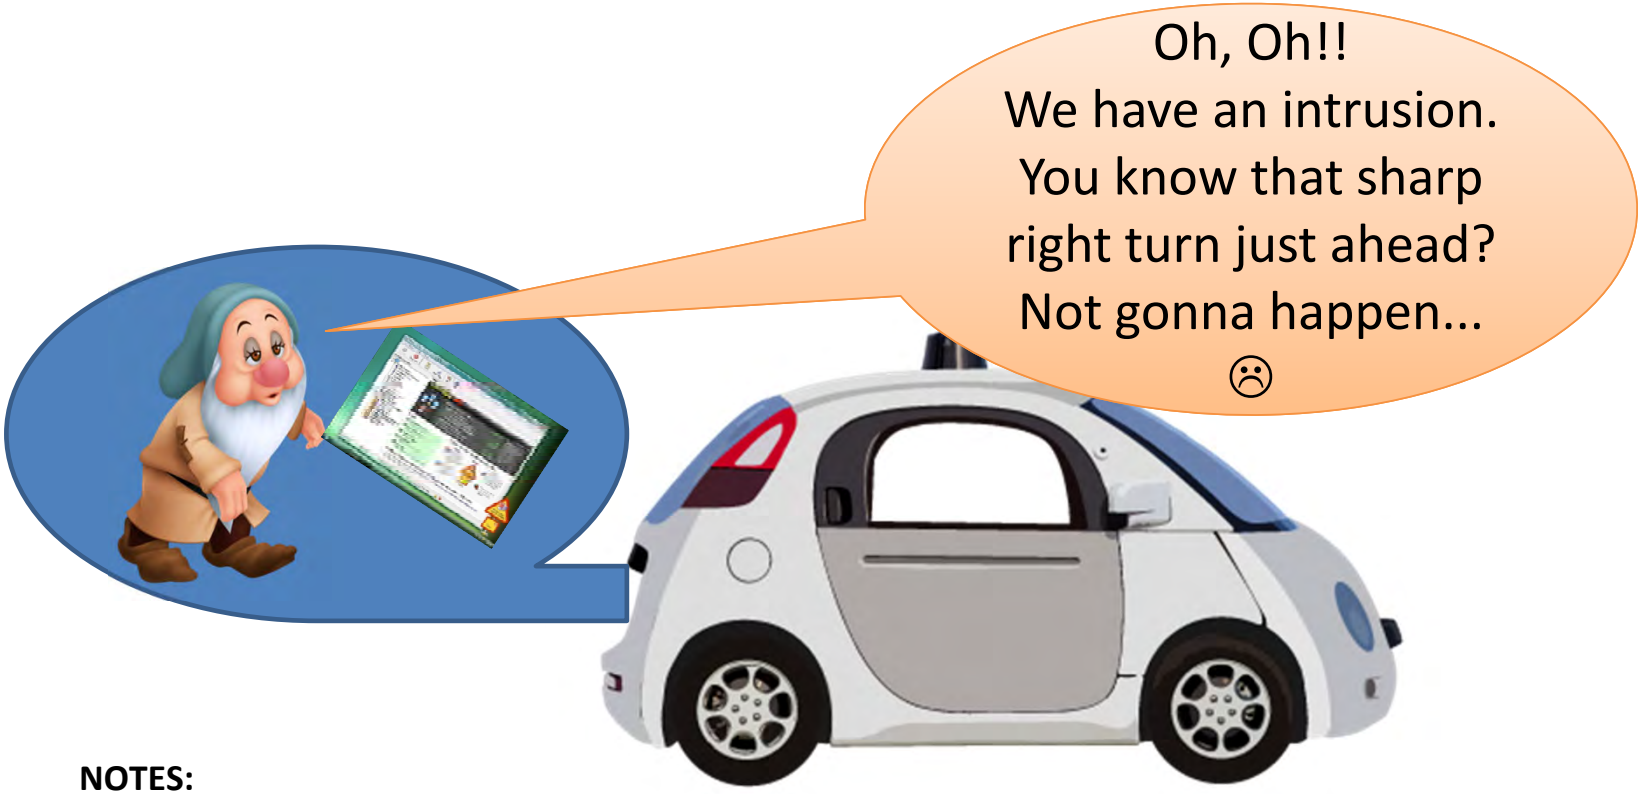

NOTES:

(i) after intrusion, the system is in the path to failure, so incompleteness or slowliness of intrusion detection and/or processing/mitigation, is the path to catastrophe.

(ii) IDS are at best useful as off-line fault diagnosis

Towards sustainable security and safety

(inspired by precursor project ICRI CARS, with INTEL)

securityandtrust.lu

CritiX

Collaboration among autonomous vehicles (V2V, V2I)

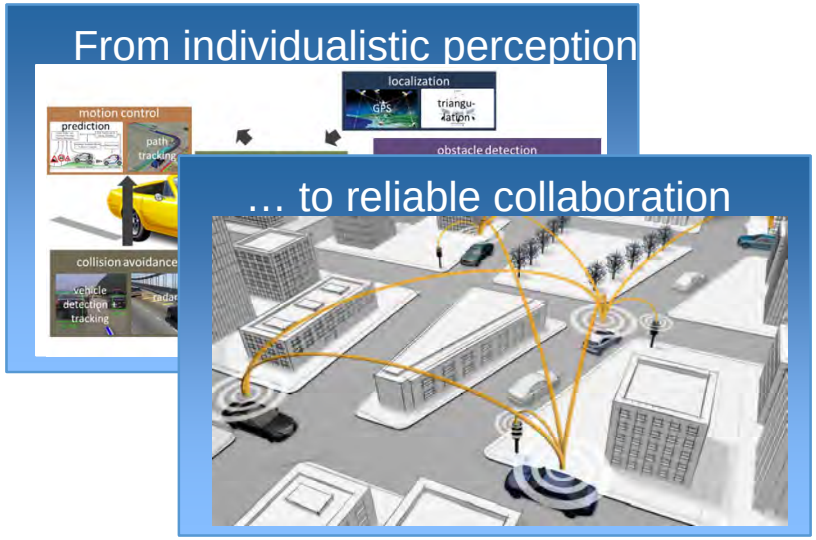

Fault and intrusion tolerant control in-vehicle by eliminating SPOFs, in particular at operating-system level

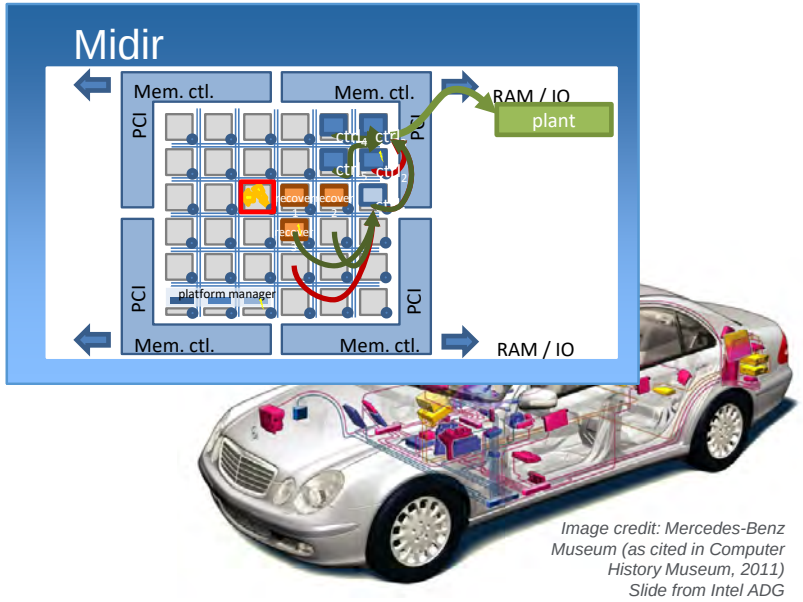

Image credit: Mercedes-Benz Museum (as cited in Computer History Museum, 2011)
Slide from Intel ADG

uni.lu UNIVERSITÉ DU LUXEMBOURG

Research strategy enablers
for autonomous and collaborative vehicles

- **Powerful architectures**
  - (e.g. manycores), capable of: high-power computing, enabling security/safety defenses
- **Automatic in-car resilience**
  - mechanisms for safety and security (gateway, ECU, trusted components/enclaves)
- **Secure and dependable *real-time* communication**,
  - V2V and V2I, despite accidents and attacks

Research strategy enablers
for autonomous and collaborative vehicles

- **Powerful architectures**
  - (e.g. manycores), capable of: high-power computing, enabling security/safety defenses
- **Automatic in-car resilience**
  - mechanisms for safety and security (gateway, ECU, trusted components/enclaves)
- **Secure and dependable *real-time* communication**,
  - V2V and V2I, despite accidents and attacks

Recent Automated Driving Solutions:
seem to be aligned with this philosophy

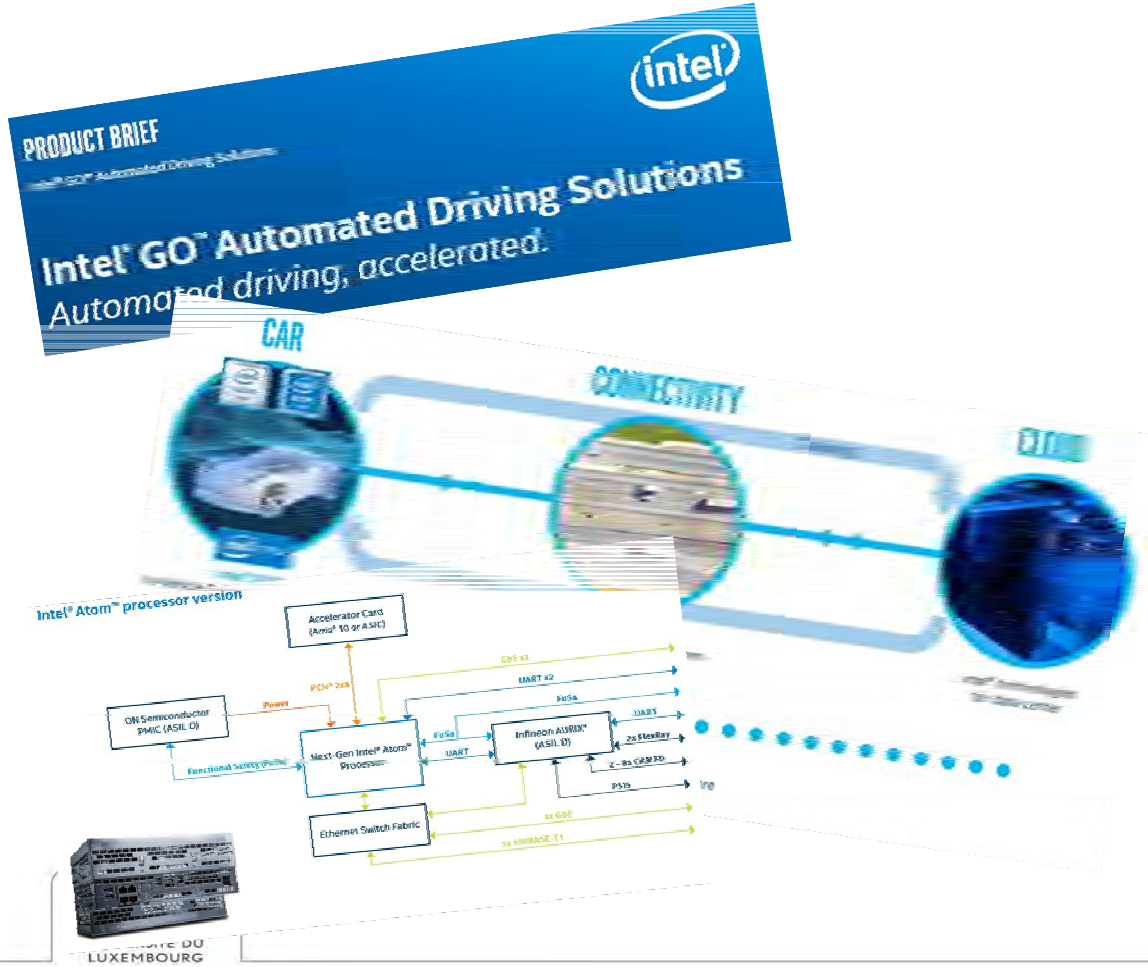

Tesla Hardware 3 (Full Self-Driving Computer)

Research strategy enablers
for autonomous and collaborative vehicles

- **Powerful architectures**
  - (e.g. manycores), capable of: high-power computing, enabling security/safety defenses
- **Automatic in-car resilience**
  - mechanisms for safety and security (gateway, ECU, trusted components/enclaves)
- **Secure and dependable *real-time* communication**,
  - V2V and V2I, despite accidents and attacks

Integrating security into safety-critical real-time distributed control systems

securityandtrust.lu

CritiX

- novel approaches which **treat deadline misses as timing faults** masked through novel timing fault and intrusion tolerant SMR algorithms
- **Collaboration, consensus and control with real-time constraints, under accidental and malicious threats**
- Byzantine-Resilient Real-time SMR control

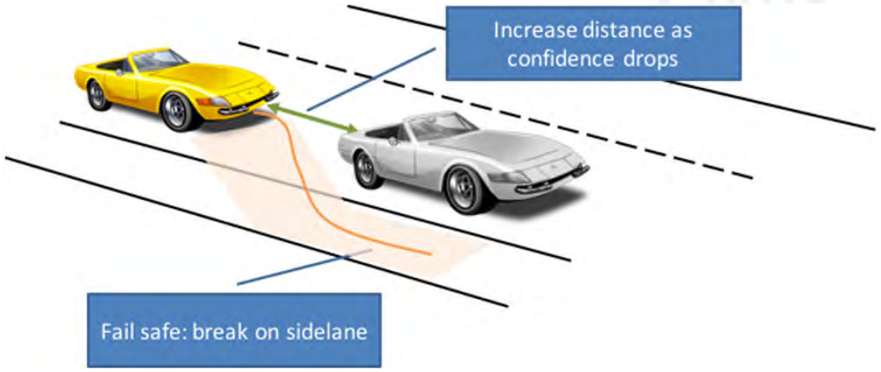

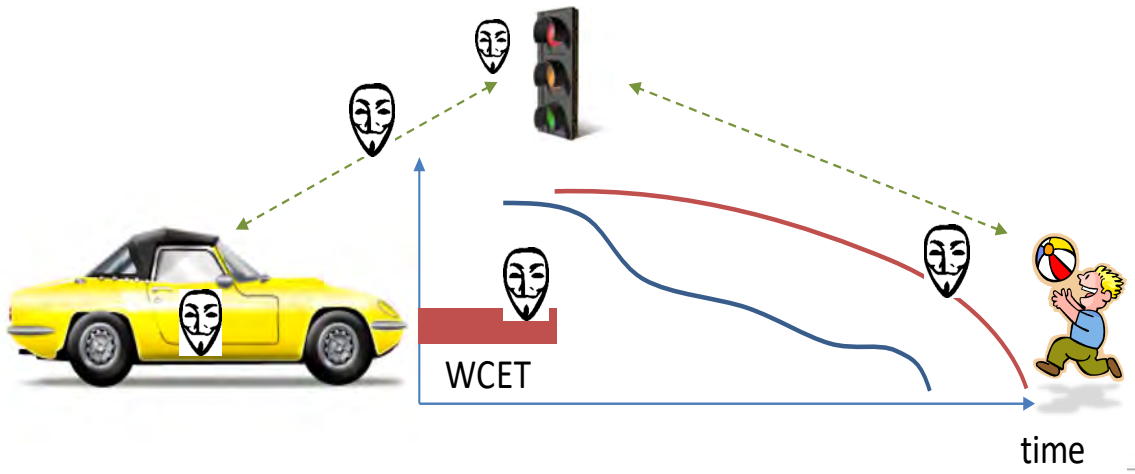

Ultra-resilient minimal roots-of-trust and enclaves

Ultra-reliable minimal roots-of-trust and enclaves

securityandtrust.lu

CRITIX

=> Patent applications

MIDIR

- Threats have been permeating all levels of architecture.
- And we are always one step “late”:
  - we rely on high-level protection (Paxos, BFT,…)
  - threats haunt below (hyp, ME, hw)
  - lost battle: general 0-defect infeasible
- Leverage properties of manycore systems:
  - create **distributed systems-on-a-chip (DisSoC)**
  - reinstantiate protection techniques at low enough level (detection, self-check, tolerance)

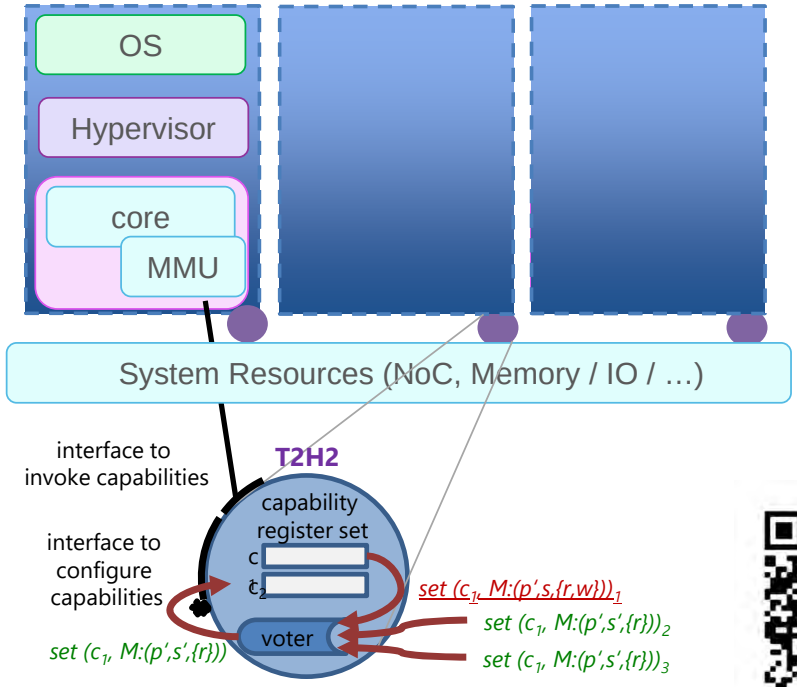

Behind the Last Line of Defense -- Surviving SoC Faults and Intrusions. Pinto Gouveia, Ines; Volp, Marcus; Esteves-Verissimo, Paulo. arXiv preprint arXiv:2005.04096 (2020)

uni.lu UNIVERSITÉ DU LUXEMBOURG

Midir: a chip level evolution path towards resilience

securityandtrust.lu

CRITIX

Hybridisation-aware algorithms: models, architect., and control

Midir:

- ***Fault containment*** through tile/sector level privilege enforcement
- ***Fault tolerance*** through consensual execution of all critical operations, in particular privilege change
- ***Resilience*** through rejuvenation and relocation

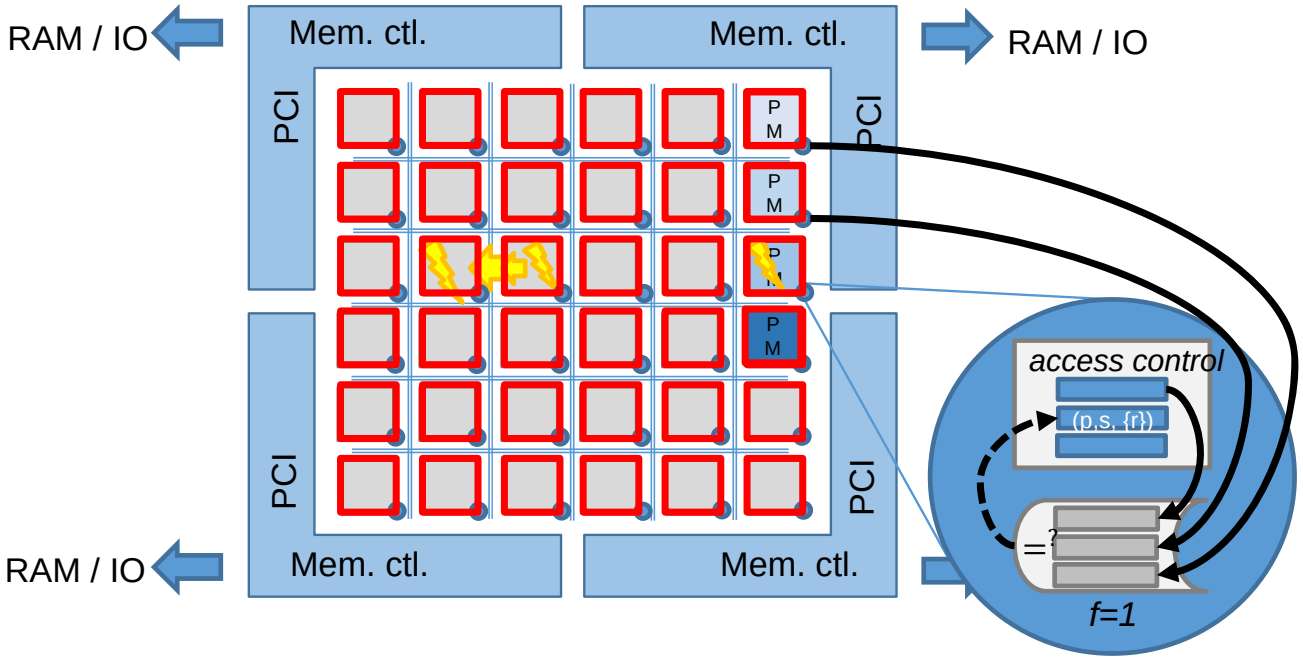

Research strategy enablers
for autonomous and collaborative vehicles

- **Powerful architectures**
  - (e.g. manycores), capable of: high-power computing, enabling security/safety defenses
- **Automatic in-car resilience**
  - mechanisms for safety and security (gateway, ECU, trusted components/enclaves)
- **Secure and dependable *real-time* communication,**
  - V2V and V2I, despite accidents and attacks

Application- and context-aware protocols for security- and safety-critical R/T ops.

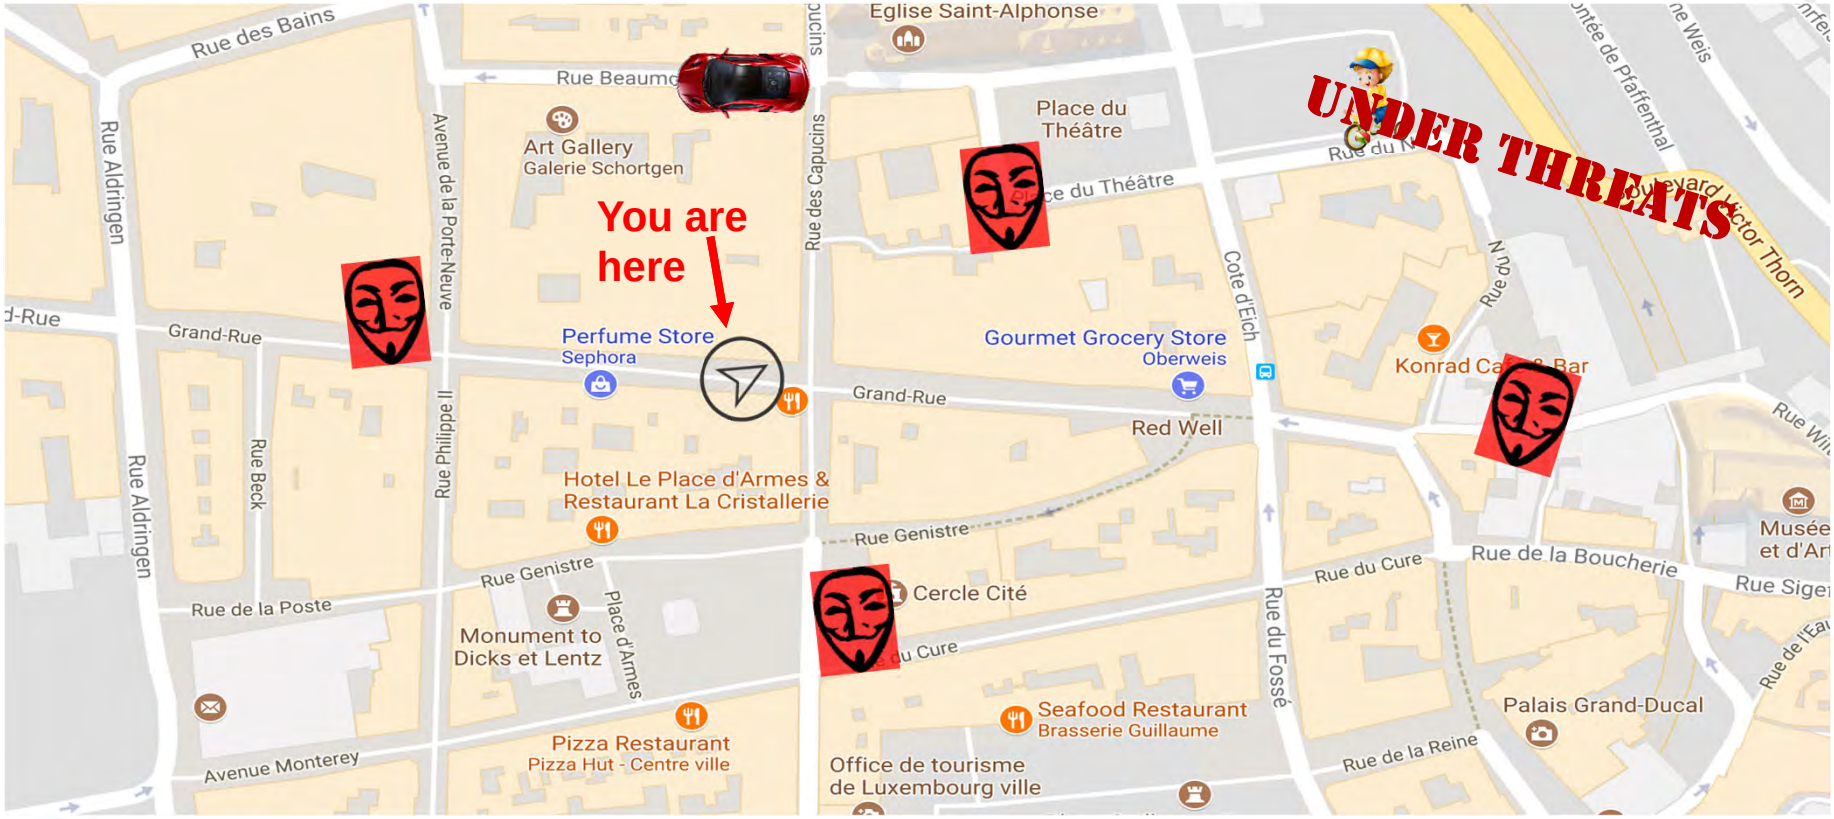

Autonomous Vehicles – ADAS, V2X

Beyond Real-Time Data Dissemination

securityandtrust.lu

CritiX

Modularly build

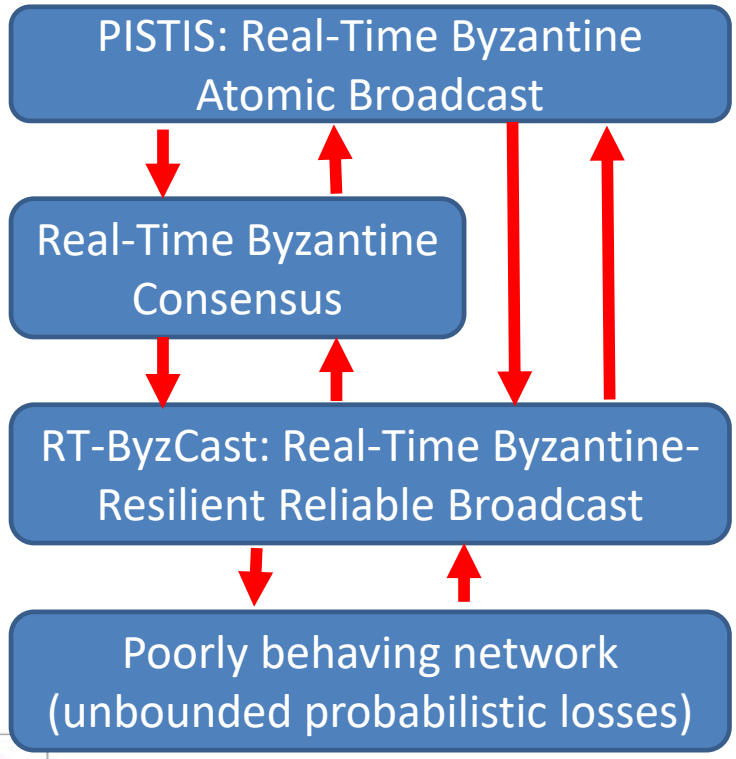

intel Research Institute for Collaborative Autonomous and Resilient Systems

Accurate Real-Time Digital Maps for Autonomous Driving

*D. Kozhaya, J. Decouchant and P. Esteves-Veríssimo, "RT-ByzCast: Real-Time Byzantine-Resilient Reliable Broadcast" , IEEE Transactions on Computers 2019, Core A**

*Kozhaya, D., Decouchant, J., Rahli, V., & Esteves-Verissimo, P. (2021). PISTIS: An Event-Triggered Real-time Byzantine Resilient Protocol Suite. IEEE TPDS. doi:10.1109/tpds.2021.3056718, Core A**

Find the right
ARCHITECTURAL SOLUTION
for autonomous in-vehicle control

Hybridisation-aware distributed algorithms, models, and architectures

Hybridisation-aware algorithms: models, architect., and control

جامعة الملك عبدالله للعلوم والتقنية
King Abdullah University of Science and Technology

Leveraging trusted-trustworthy components and TEE, with the right set of simple functions (failure detectors, monotonic counters, reliable timers or clocks, PRG, signatures, indelible logs, binary consensus)

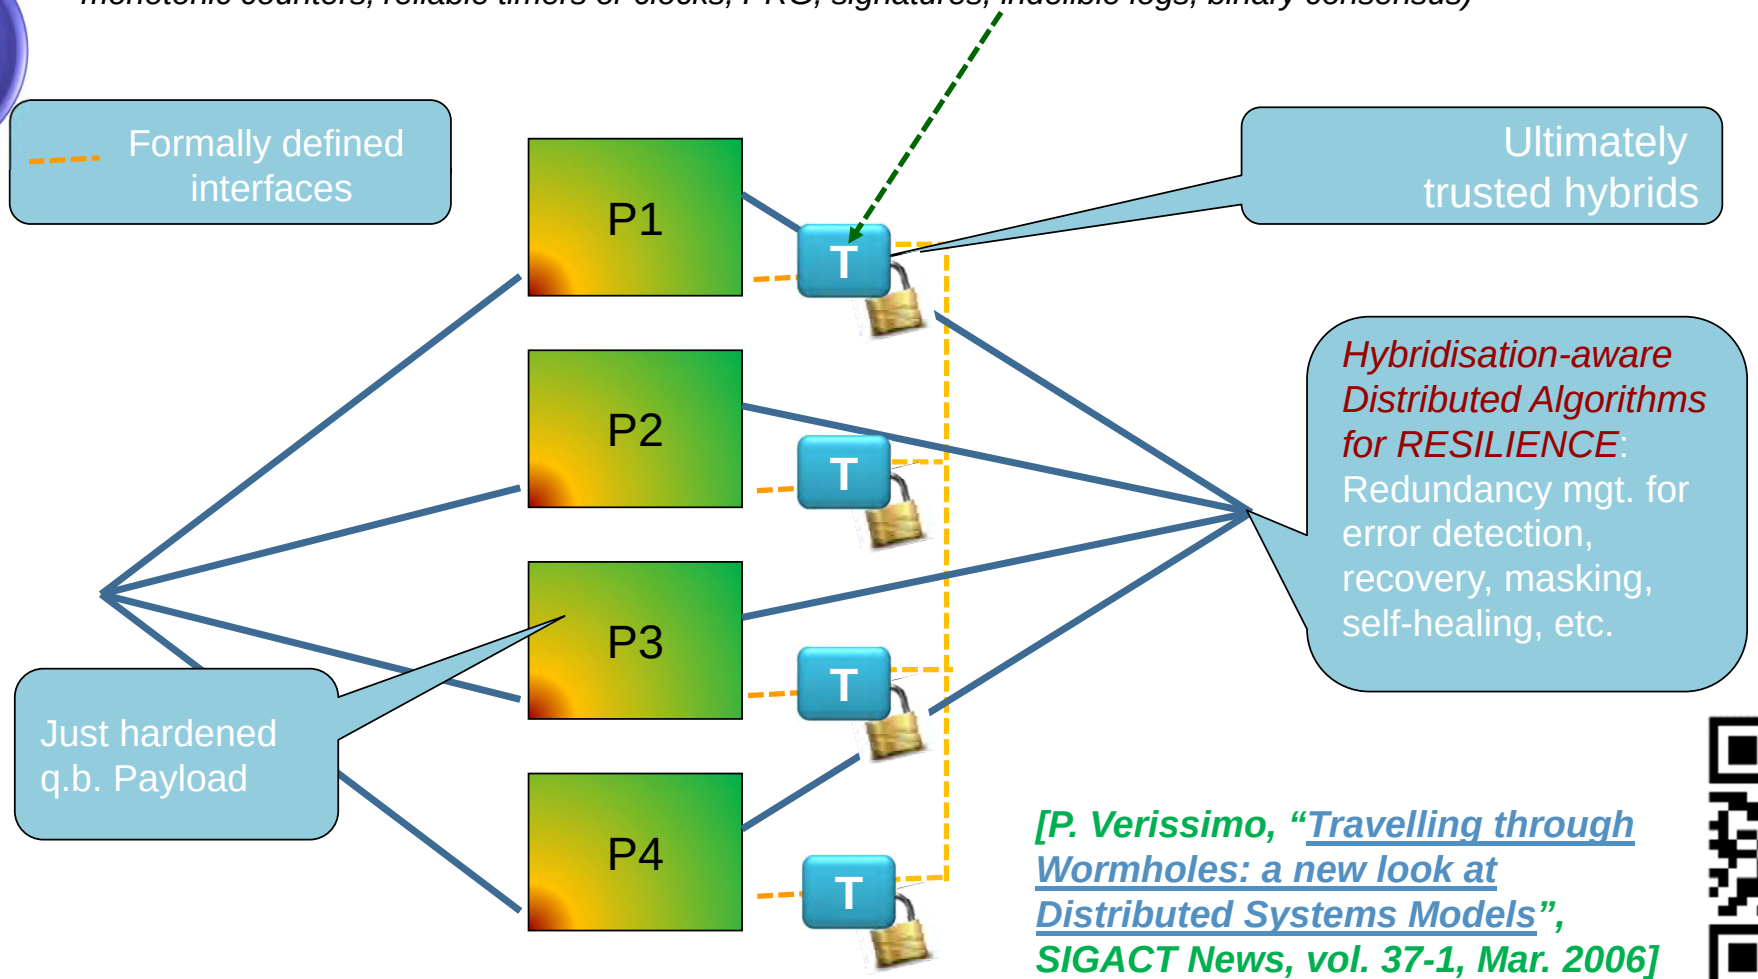

[P. Verissimo, "Travelling through Wormholes: a new look at Distributed Systems Models", SIGACT News, vol. 37-1, Mar. 2006]

Divide-and-conquer I: **Hybrid models and architectures**
Giving substance to assumptions

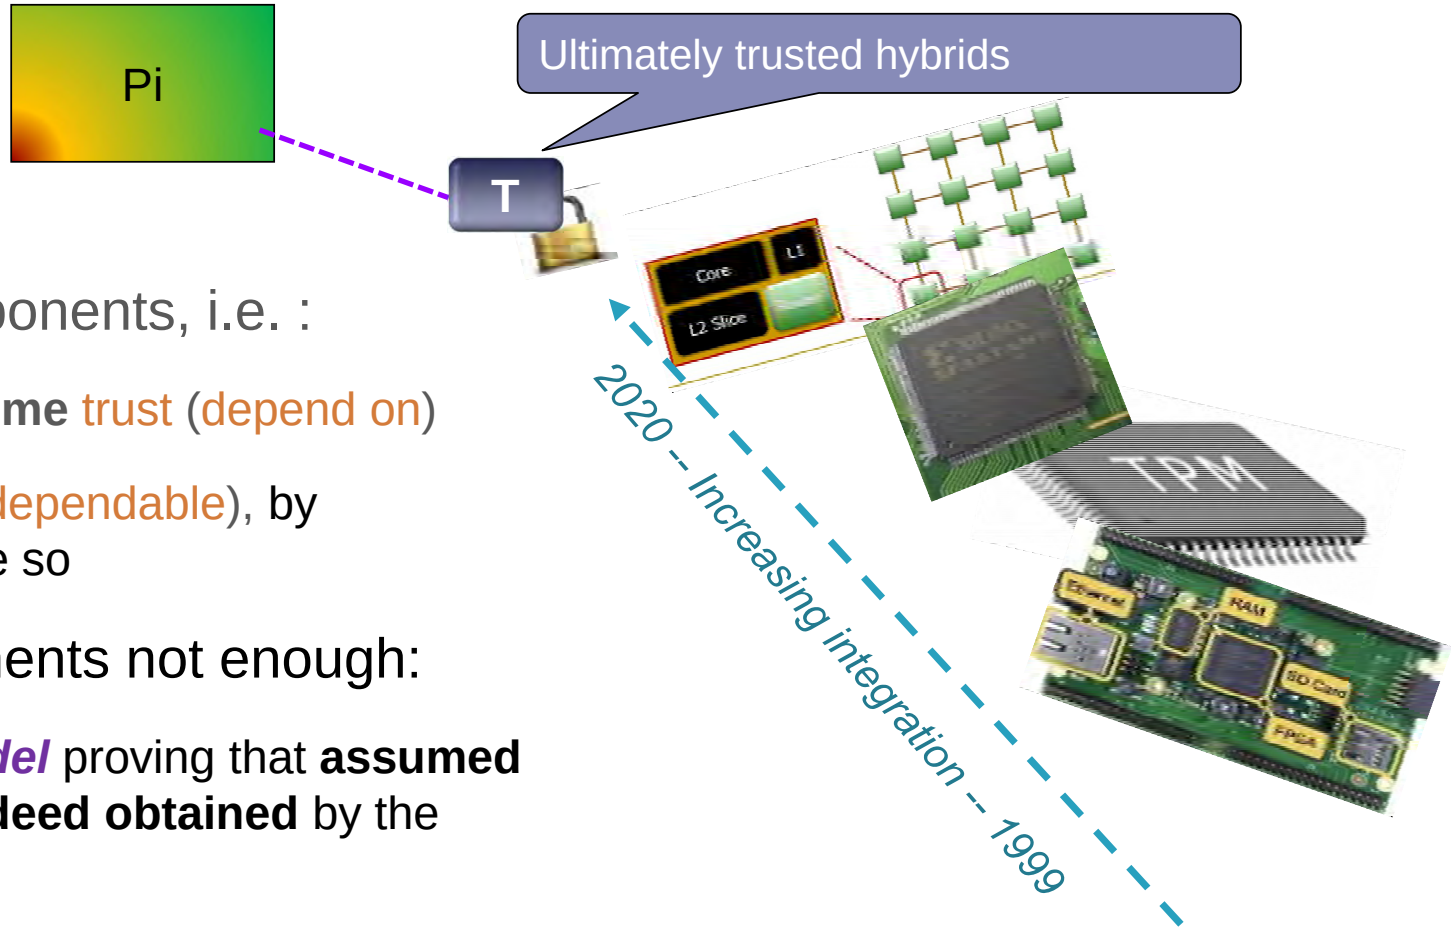

- ***Trusted-Trustworthy*** components, i.e. :
  - components on which you **assume** trust (depend on)
  - because they are trustworthy (dependable), by **construction**, and verified to be so
- Trusted-trustworthy components not enough:
  - We need a ***computational model*** proving that **assumed** trusted comp. properties **are indeed obtained** by the payload

Hybrid models and architectures for VLSI FIT over tile-based many-core SoC

Leveraging power at right place right time

securityandtrust.lu

CRITIX

Replication for fault and intrusion tolerance

Incoming requests

CORE HV1 TILE

CORE HV2 TILE

CORE HV3 TILE

RE TILE

Hybridisation-aware distributed algorithms

Results consolidation

Ultimately trusted hybrids

Just hardened q.b. payload

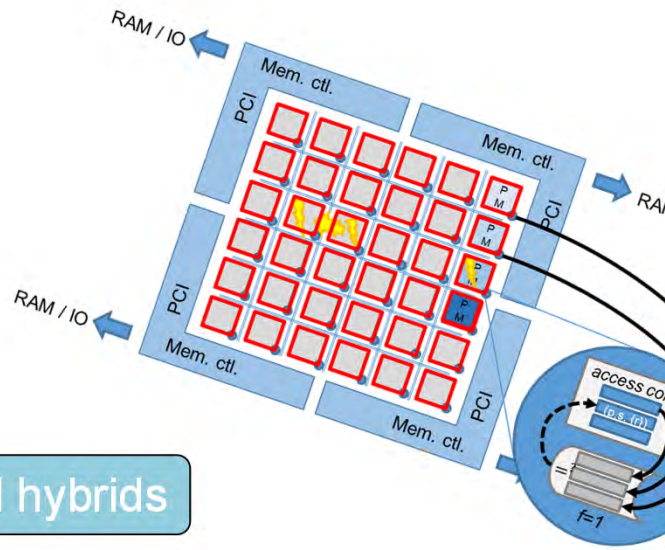

uni.lu UNIVERSITÉ DU LUXEMBOURG

3f+1 replicas (f=1 in this example)

KARYON architecture: proof of concept of hybridisation for safety

- Main Concepts:
  - Level of Service
  - Data validity
  - Safety kernel

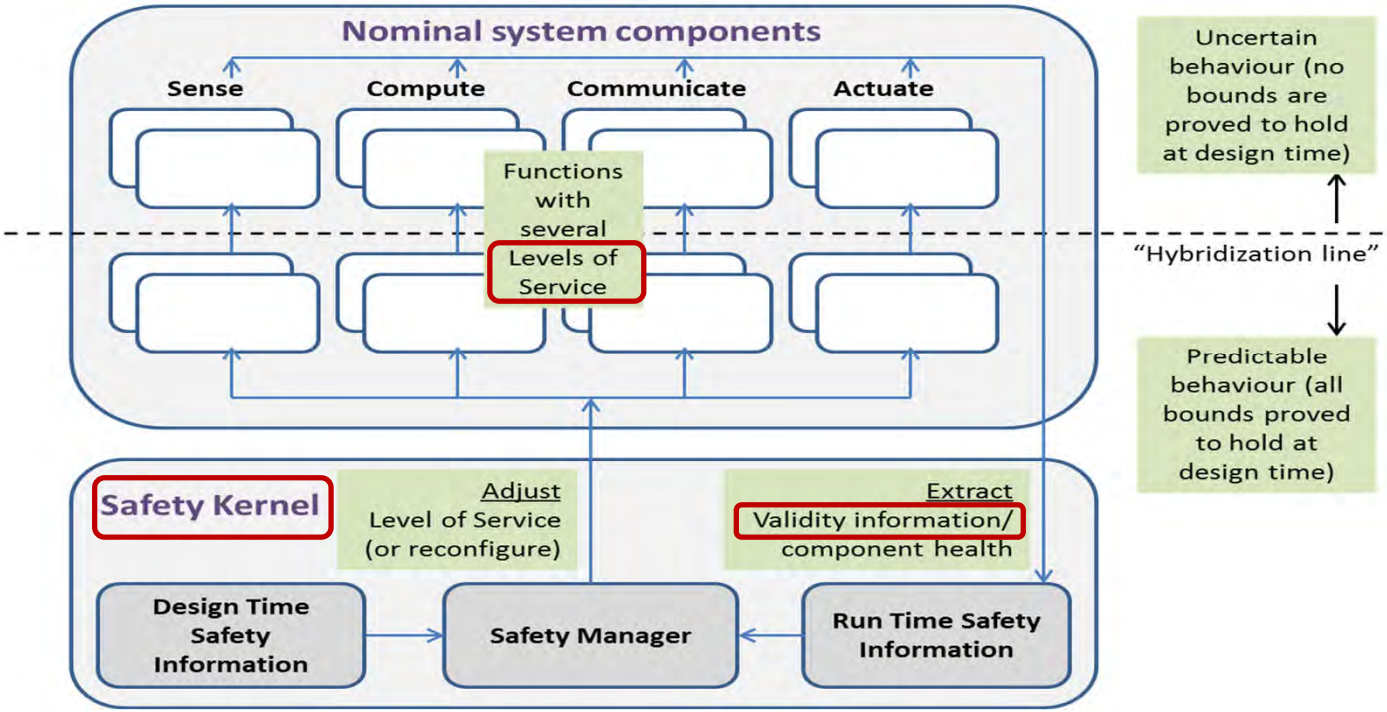

Karyon Architecture

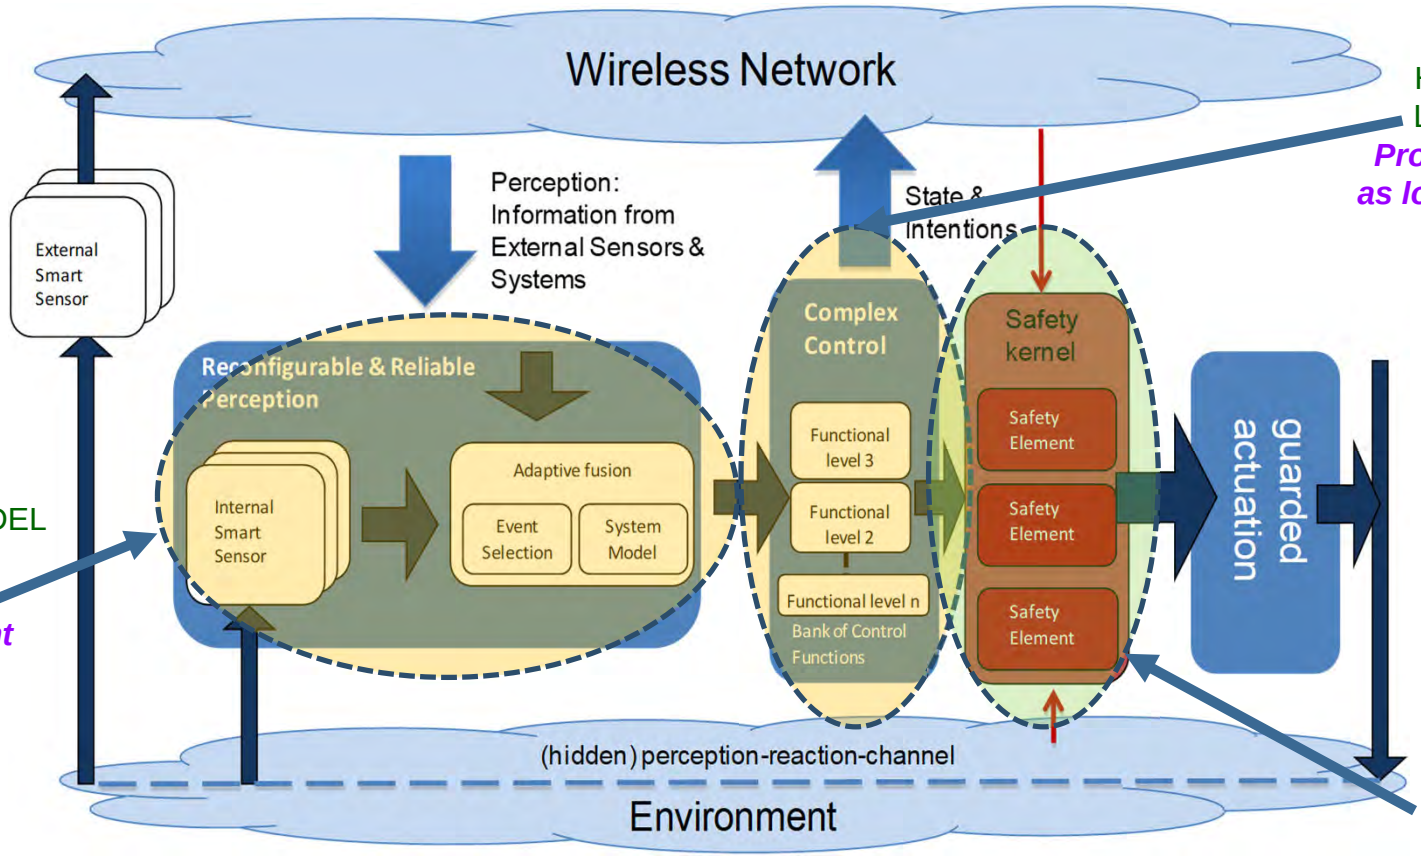

HYBRIDISATION BY LEVEL OF SERVICE
Proven functional safety, as long as assumptions for level are satisfied

ABSTRACT SENSOR MODEL
integrity attributes are evaluated against a set of different safety requirements

SAFETY KERNEL HYBRID
takes over, toward fail-operational or fail-safe termination

KARYON prototypes
Proof of concept -- Cars demonstration

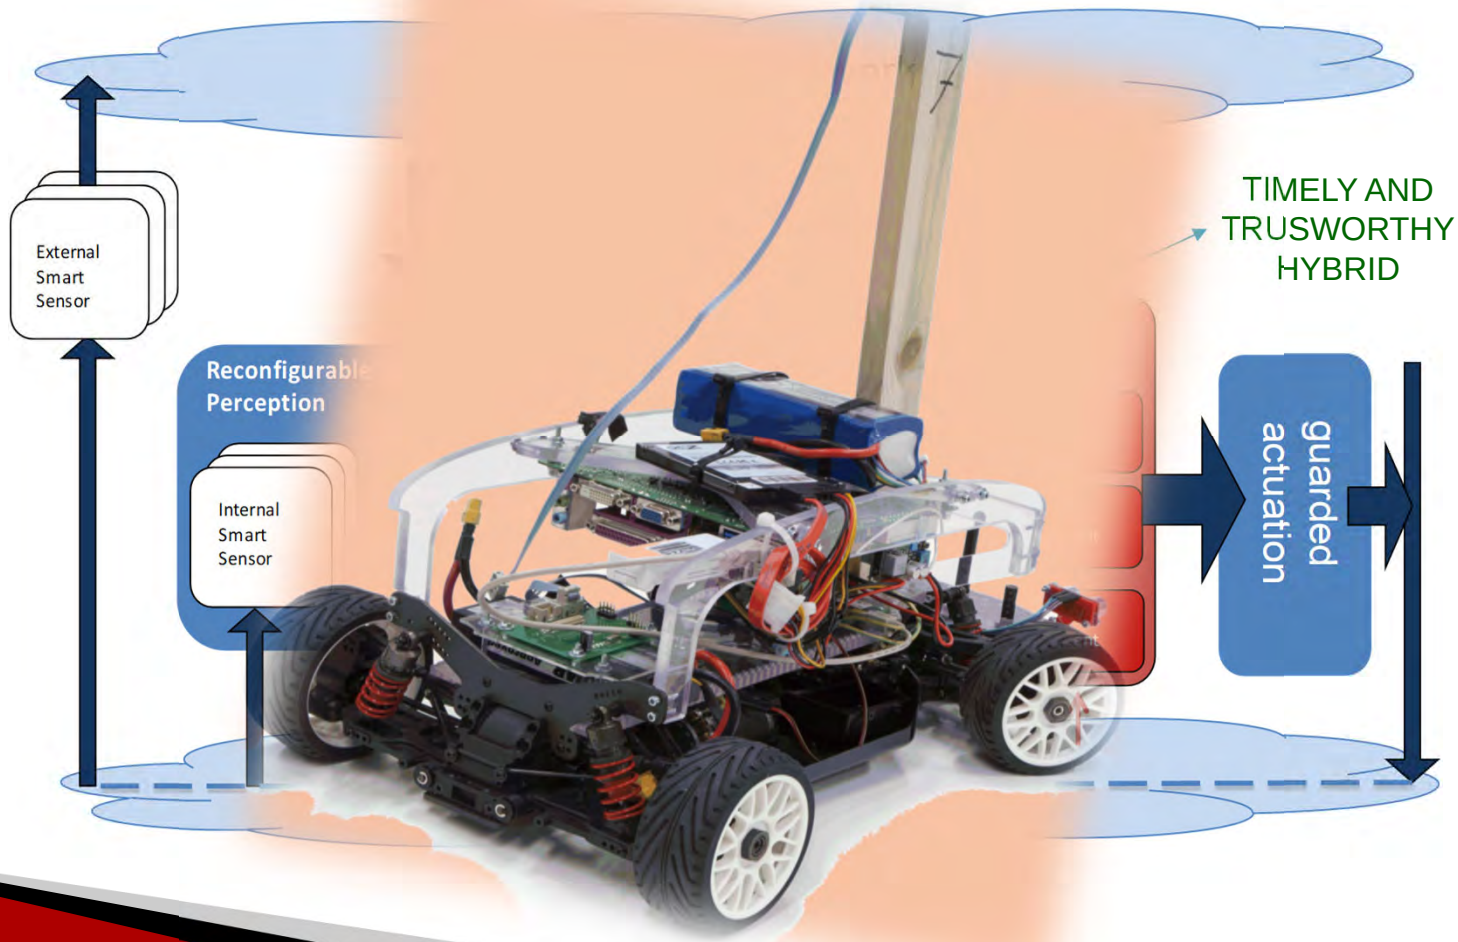

KARYON prototypes

Proof of c onstration

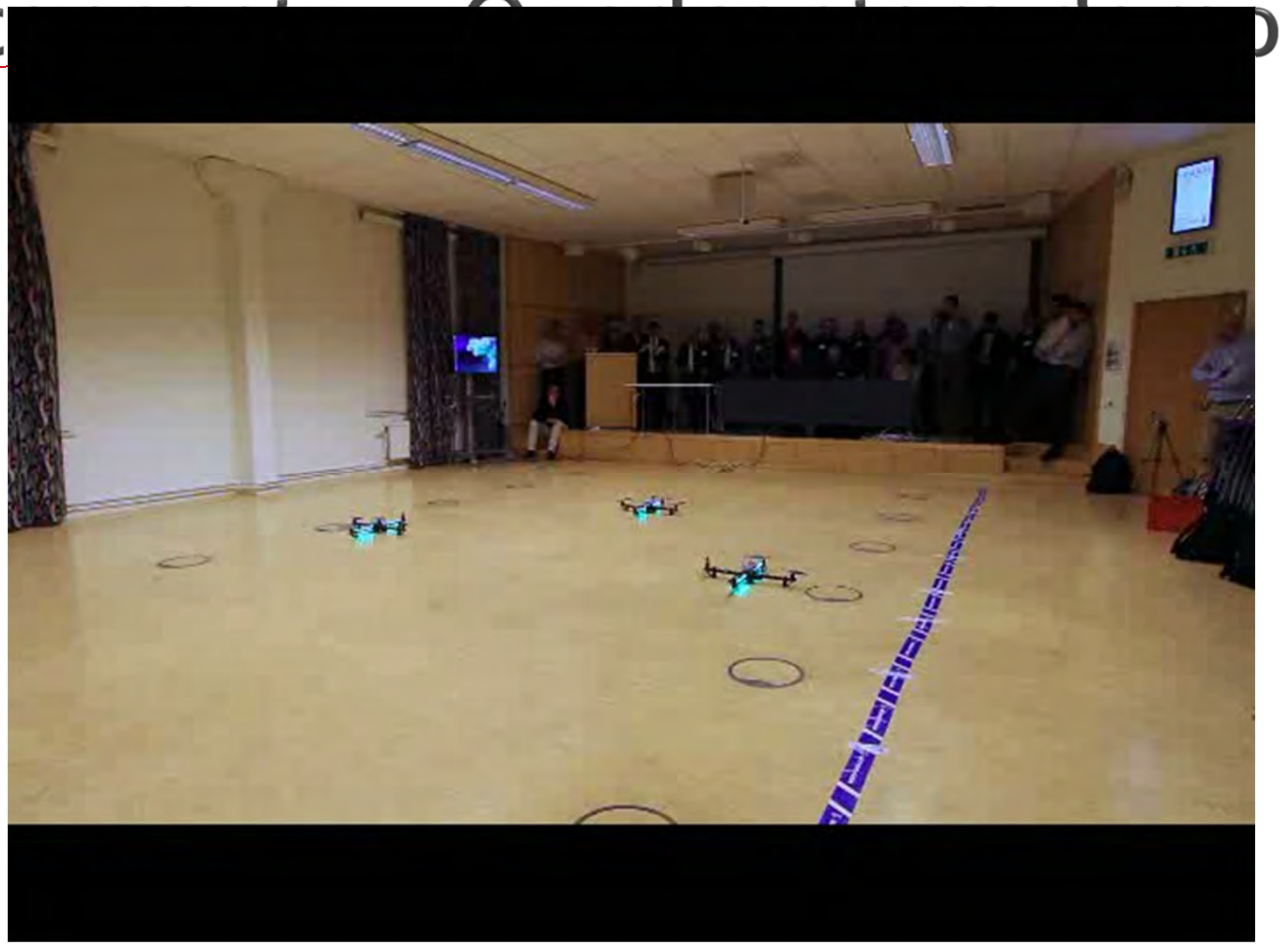

The end.

There is no trusted Artificial Intelligence (the data) on a non-trustworthy Artificial Brain (the system)

Thank you!

Paulo Esteves-Veríssimo

PhD STUDENTS
POSTDOCTORAL &
RESEARCH SCIENTISTS

Opportunities in Resilient Computing

at Resilient Computing and Cybersecurity Center (RC3),
King Abdullah University of Science and Technology (KAUST)

We seek candidates who are prepared to participate in the shaping of such a challenging and innovative research theme and contribute to the ambitious research agenda of the RC3 center, supported by a very substantial funding landscape.

rc3.kaust.edu.sa