2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) October 25-29, 2020, Las Vegas, NV, USA (Virtual)

Research and Educational Directions

Houssam Abbas EECS, Oregon State University

and the F1/10 team



978-1-7281-6211-9/20/\$31.00 ©2020 IEEE

The role F1/10 can play in preparing students for research

Algorithmic Ethics

Cyber-Physical Security

Fairness in Multi-Agent Control

How should a car behave?

Understanding the vulnerabilities of an autonomous vehicle to attacks Ensuring fair access to the airspace for UAS

F1/10: An Open-Source Autonomous Cyber-Physical Platform

Matthew O'Kelly*, Houssam Abbas**, Jack Harkins, Chris Kao, Yash Vardhan Pant & Rahul Mangharam Department of Electrical and Systems Engineering, University of Pennsylvania, USA {mokelly, harkj, chriskao, yashpant, rahulm}@seas.upenn.edu ** School of Electrical Engineering and Computer Science, Oregon State University, USA houssam.abbas@oregonstate.edu

> Varundev Suresh Babu*, Dipshil Agarwal & Madhur Behl Department of Computer Science, University of Virginia, USA {varundev, dpa4va, madhur.behl}@virginia.edu

Paolo Burgio & Marko Bertogna Dept. of Physics, Informatics & Mathematics, University of Modena & Reggio Emilia, Italy {marko.bertogna, paolo.burgio}@unimore.it

* These authors contributed equally

ABSTRACT

8567v1 [cs.RO] 24 Jan 2019

In 2005 DARPA labeled the realization of viable autonomous vehicles (AVs) a grand challenge; a short time later the idea became a moonshot that could change the automotive industry. Today, the question of safety stands between reality and solved. Given the right platform the CPS community is poised to offer unique insights. However, testing the limits of safety and performance on real vehicles is costly and hazardous. The use of such vehicles is also outside the reach of most researchers and students. In this paper, we present F1/10: an open-source, affordable, and high-performance 1/10 scale autonomous vehicle testbed. The F1/10 testbed carries a full suite of sensors, perception, planning, control, and networking software stacks that are similar to full scale solutions. We demonstrate key examples of the research enabled by the F1/10 testbed, and how the platform can be used to augment research and education in autonomous systems, making autonomy more accessible.

1 INTRODUCTION



Figure 1: It takes only a couple of hours fully to assemble a F1/10 autonomous racecar, using detailed instructions available at http://fitenth.org/

F1/10: An Open-Source Autonomous Cyber-Physical Platform

Matthew O'Kelly*. Houssam Abbas**. Jack Harkins. Chris Kao. Yash Vardhan Pant & Rahul Mangharam Department of Electrical and Systems Engineering, University of Pennsylvania, USA {mokelly, harki, chriskao, vashpant, rahulm}@seas.upenn.edu ** School of Electrical Engineering and Computer Science, Oregon State University, USA houssam.abbas@oregonstate.edu

Dept. of

50

ABSTRACT

In 2005 DARPA labeled the rea hicles (AVs) a grand challenge; a moonshot that could change the question of safety stands bet right platform the CPS communi However, testing the limits of sa cles is costly and hazardous. The the reach of most researchers and F1/10: an open-source, affordable autonomous vehicle testbed. Th of sensors, perception, planning, stacks that are similar to full sc examples of the research enable the platform can be used to au autonomous systems, making a NTRODUCTION

FormulaZero: Distributionally Robust Online Adaptation via Offline Population Synthesis

> Aman Sin Rahul M

amans@stanfo

Balancing performan agent environments. In conservative policies, hi either make simplifying adaptation. This work a realistic, diverse set o

TUNERCAR: A Superoptimization Toolchain for Autonomous Racing

Matthew O'Kelly¹, Hongrui Zheng¹, Joseph Auckley¹, Achin Jain¹, Kim Luong^{1,2}, and Rahul Mangharam¹

Abstract— The objective of this effort is to develop an optimal autonomous racer with safe and reusable core autonomy components that are agnostic to vehicle planning and control software. TUNERCAR is a toolchain that jointly optimizes racing strategy, planning methods, control algorithms and vehicle parameters for an autonomous racecar. In this paper, we detail the target hardware, software, simulators, and systems infrastructure for this toolchain. Our methodology employs a massively parallel implementation of CMA-ES which enables simulations to proceed 6 times faster than real-world rollouts. Besides a massive speed up, we show our approach can reduce the lap times in autonomous racing, given a fixed computational budget. We demonstrate improvements over naive random search of 2.0 seconds per lap, improvements over expert solutions of 0.81 seconds per lap, and beat a human driver by 6.52 seconds. For all tested tracks, our method provides the lowest lap-time, and relative lap-time improvements between 6 and 12 percent. We further compare the performance of our method against hand tuned solutions submitted by over 30 international

This paper introduces the notion that component reuse and adaptation is analogous to creating a new kind of compiler that targets computational, physical, and external environmental details of a robot's operational domain. In general, the goal of a compiler is to validate and then transform a source program in one language to another (usually lower-level *e.g.* assembly) which is suitable for the target domain [2]. Modern optimizing compilers [3] also seek to improve the performance of the transformed program. To concretize the analogy, we define the source program as a parameterized description of the vehicle dynamics, tracking controllers, and local planning method which we wish to transform to perform safely and efficiently in the operational environment represented by a map, physical laws, and sensing capabilities.

We propose a solution to the cyber-physical compilation problem utilizing the concept of superoptimization (c.f. [4],

F1/10: An Open-Source Autonomous Cyber-Physical Platform

Matthew O'Kelly*, Houssam Abbas**, Jack Harkins, Chris Kao, Yash Vardhan Pant & Rahul Mangharam Department of Electrical and Systems Engineering, University of Pennsylvania, USA {mokelly, harki, chriskao, vashpant, rahulm}@seas.upenn.edu

Teaching Autonomous Systems at $1/10^{th}$ **-scale:** Design of the F1/10 Racecar, Simulators and Curriculum

Abhijeet Agnihotri*[†] agnihota@oregonstate.edu Oregon State University Corvallis, Oregon

Rahul Mangharam rahulm@seas.upenn.edu University of Pennsylvania Philadelphia, Pennsylvania

ABSTRACT

Teaching autonomous systems is challenging because it is a rapidly advancing cross-disciplinary field that requires theory to be continually validated on physical platforms. For an autonomous vehicle

> adaptation. This work a realistic, diverse set o

Matthew O'Kelly* mokelly@seas.upenn.edu University of Pennsylvania Philadelphia, Pennsylvania

Houssam Abbas houssam.abbas@oregonstate.edu Oregon State University Corvallis, Oregon



implementation of CMA-ES which enables simulations to proceed 6 times faster than real-world rollouts. Besides a massive speed up, we show our approach can reduce the lap times in autonomous racing, given a fixed computational budget. We demonstrate improvements over naive random search of 2.0 seconds per lap, improvements over expert solutions of 0.81 seconds per lap, and beat a human driver by 6.52 seconds. For all tested tracks, our method provides the lowest lap-time, and relative lap-time improvements between 6 and 12 percent. We further compare the performance of our method against hand tuned solutions submitted by over 30 international

Autonomous Racing

uong^{1,2}, and Rahul Mangharam¹

es the notion that component reuse and to creating a new kind of compiler that physical, and external environmental rational domain. In general, the goal of te and then transform a source program ther (usually lower-level e.g. assembly) e target domain [2]. Modern optimizing

compilers [3] also seek to improve the performance of the transformed program. To concretize the analogy, we define the source program as a parameterized description of the vehicle dynamics, tracking controllers, and local planning method which we wish to transform to perform safely and efficiently in the operational environment represented by a map, physical laws, and sensing capabilities.

We propose a solution to the cyber-physical compilation problem utilizing the concept of superoptimization (c.f. [4],

NTD

Real-time systems research

- Real workloads
- Representative hardware

Real-time systems research

- Real workloads
- Representative hardware

AutoV: An Automotive Testbed for Real-Time Virtualization

Overview

Contributors Demos Publications Related Projects Sitemap

Overview

Automotive systems are becoming increasingly complex. Virtualization is a promising technique to achieve low size, we systems: functionalities on multiple ECUs can be consolidated into multiple virtual machines (VMs) on a commodity must as the <u>RT-Xen project</u>, is a promising technique to achieve the timing isolation for virtualization in automotive systems. with real automotive applications in a non-simulation environment

Real-time reachability



Real-time reachability

Runtime monitoring

Real-time reachability

Runtime monitoring



Reelay is a header-only C++ library and set of tools for system-level verification and testing of real-time systems. Reelay implements state-of-the-art runtime verification techniques to construct runtime monitors that check temporal behaviors of the system against system-level requirements. Hence, Reelay can be used to enhance rigorous systems engineering practices by formalizing and automating the assessment phase.

Main Features

- · Formal specification of temporal properties
- · Provably correct monitor construction from the specification
- · Fast and frugal runtime requirement checking (very low overhead)
- Simple but non-restrictive user interface
- Available for C++ and Python

Real-time reachability

Runtime monitoring

Adversarial traffic

Scalable End-to-End Autonomous Vehicle Testing via Rare-event Simulation

Matthew O'Kelly* University of Pennsylvania mokelly@seas.upenn.edu Aman Sinha* Stanford University amans@stanford.edu Hongseok Namkoong* Stanford University hnamk@stanford.edu

John Duchi Stanford University jduchi@stanford.edu Russ Tedrake Massachusetts Institute of Technology russt@mit.edu

F1/10 Onboard Reachability

Nathan Jewell-Carlton, CS undergrad Autonomous Systems Lab OSU

Problem

- How do we ensure safety?
- Where is the agent going to be?
- How is it going to get there?
- Will it hit an obstacle?
- Find out and react before it happens (speed!)



- Find physical properties of the agent
- Develop dynamical model

- Find physical properties of the agent
- Develop dynamical model
- Minimally over-approximate possible future system states



- Find physical properties of the agent
- Develop dynamical model
- Minimally over-approximate possible future system states
- Ensure there are no failure modes
- Adjust control systems

- Find physical properties of the agent
- Develop dynamical model
- Minimally over-approximate possible future system states
- Ensure there are no failure modes
- Adjust control systems

this is the essence of reachability

Successive Linearization

Nonlinear reachability - slow, accurate

Linearize then run Linear reachability - fast, cheap, less accurate

Successive Linearization

- Multiple linear eq. Combined
- Fast computation algorithms
- Controllable accuracy
- Linearization overhead

Non-Linear Dynamical Model

Variables of state are decided by the systems' dynamics. x: the vehicle's x coordinate in meters y: the vehicle's y coordinate in meters ψ : the vehicle's yaw in Radian

Input variables are decided by the control system. v: the vehicles velocity d: the vehicles change if ψ

$$\dot{x} = v\cos\psi + d \tag{1}$$

$$\dot{y} = v \sin \psi + d \tag{2}$$

$$\dot{\psi} = v \frac{1}{L} \sin(d) \tag{3}$$



Uses HYLAA,

a reachability tool by Stanley Bak



607





GPU Parallelism

- How can performance be increased?
- Use more of our hardware!
- Free cpu cycles
- Improve latency
- Significant refactoring required



Digging into function call timings (pycallgraph output)





Flow diagram for refactorin⁶¹ run on NVIDIA GPU

Demo



RVIZ of car running around demo track - live onboard reachability

Performance

Combined execution time by steps for runtime modes



Performance



Performance

Standard Deviation between trial execution time by steps for runtime mode



Nonlinear MPC and MPCC for F1/10 cars

Niraj Basnet, M.Sc. student Autonomous Systems Lab, OSU

Nonlinear MPC

- Uses nonlinear system model
- Better disturbance rejection capability
- Better step response
- Better tracking performance
- Slower than Linear MPC
- Modern solvers are improving runtimes

$$\min_{\boldsymbol{U},\boldsymbol{X}} \int_{t=0}^{T} \left\| \boldsymbol{x}(t) - \boldsymbol{x}_{ref}(t) \right\|_{\boldsymbol{Q}_{\boldsymbol{x}}}^{2} + \left\| \boldsymbol{u}(t) - \boldsymbol{u}_{ref}(t) \right\|_{\boldsymbol{R}_{\boldsymbol{u}}}^{2} dt + \left\| \boldsymbol{x}(T) - \boldsymbol{x}_{ref}(T) \right\|_{\boldsymbol{P}}^{2} \text{subject to} \quad \dot{\boldsymbol{x}} = \boldsymbol{f}(\boldsymbol{x}, \boldsymbol{u}); \\ \boldsymbol{u}(t) \in \mathbb{U} \\ \boldsymbol{x}(0) = \boldsymbol{x}(t_{0}).$$

Model Predictive Contouring Control (MPCC)

- Spatial reference path sufficient for tracking
- No reference control inputs needed
- Centerline of the racetrack acts as reference contour, parameterized by its arc length θ
- Uses lag and contouring error to generate trajectory with trade-off between tracking precision and speed
- Path planning and tracking combined into single nonlinear optimization problem
- Able to generate locally optimal racing lines



Contouring error(e^c) and Lag error(e^l) with linear approximations \hat{e}^c and \hat{e}^l

Nonlinear MPCC

Problem Formulation

$$\min \sum_{k=1}^{N} \begin{bmatrix} \hat{e}_{k}^{c} \\ \hat{e}_{k}^{l} \end{bmatrix}^{T} \begin{bmatrix} q_{c} & 0 \\ 0 & q_{l} \end{bmatrix} \begin{bmatrix} \hat{e}_{k}^{c} \\ \hat{e}_{k}^{l} \end{bmatrix} - q_{v} v_{\theta,k} + \Delta u_{k}^{T} R_{\Delta} \Delta u_{k}$$
s.t. $x_{0} = x(0)$

$$x_{k+1} = f(x_{k}, u_{k}) \longrightarrow \text{System model}$$

$$\hat{e}^{c}(x_{k}) = -\sin \left(\Phi^{\text{ref}}(\theta_{k})\right) \left(X_{k} - X^{\text{ref}}(\theta_{k})\right) - \cos \left(\Phi^{\text{ref}}(\theta_{k})\right) \left(Y_{k} - Y^{\text{ref}}(\theta_{k})\right) \longrightarrow \text{Contouring error}$$

$$\hat{e}^{l}(x_{k}) = -\cos \left(\Phi^{\text{ref}}(\theta_{k})\right) \left(X_{k} - X^{\text{ref}}(\theta_{k})\right) - \sin \left(\Phi^{\text{ref}}(\theta_{k})\right) \left(Y_{k} - Y^{\text{ref}}(\theta_{k})\right) \longrightarrow \text{Lag error}$$

$$\Delta u_{k} = u_{k} - u_{k-1} \longrightarrow \text{Rate of input change}$$

$$x_{k} \in \mathcal{X}_{\text{Track}} \longrightarrow \text{Path constraints}$$

$$\underline{u} \leq u_{k} \leq \overline{u} \longrightarrow \text{State constraints}$$

$$\underline{u} \leq u_{k} \leq \overline{u} \longrightarrow \text{Actuator constraints}$$

Kinematic vehicle model

- Simplified motion model with easiest parameter identification
- Imposes non-holonomic constraints
- No regards for tyre forces
- Suitable for slow speeds only
- Prone to understeering and oversteering
- Computationally cheap to compute

$$\begin{array}{c} x = [X, Y, \psi, v, \theta] \\ u = [a, \delta, v_{\theta}] \end{array}$$



Rear axle bicycle kinematic model

$$\dot{x} = v\cos(\theta)$$

$$\dot{y} = v\sin(\theta)$$

$$\dot{\theta} = \frac{v}{L}\tan(\delta)$$

 $\dot{v} = a$

Dynamic Vehicle Model

- Considers both lateral and longitudinal dynamics
- Uses linear or nonlinear tyre forces to account for slip at high speeds
- Involves complicated system modeling by an expert
- Handles high speed behaviors
- Defines understeer and oversteer
- Computationally demanding





Hybrid Vehicle Model

- Dynamic model(DM) is ill-defined for slow speeds due to slip angles
- Kinematic model(KM) fails at high speeds
- Combine Dynamic and Kinematic model to get the best of both
- Uses a velocity threshold to switch between two models
- Velocity threshold used : $v_{th} = 0.8$ m/s
- Actuator space is kept same while state space solution is remapped during the transition

$$f_{hyb}(x,u) = \lambda f_{kin}(x,u) + (1-\lambda) f_{dyn}(x,u)$$
$$\lambda = \begin{cases} 1 & \text{if } v_x \leq v_{th} \\ 0 & \text{otherwise} \end{cases}$$

Constraints

bound

Tyre constraints

Slip angle constraints on front wheel and rear wheel



Obstacle constraints

Ellipsoidal approximation of other obstacles(cars)



Solvers for Nonlinear MPCC

- Interior point based solvers like Ipopt , Hpipm, Acados, ForcesPro(Commercial)
- Can be structured into sequence of locally convex QPs to use SQP based solvers
- Casadi framework for Algorithmic differentiation
- Implementation details of Nonlinear MPCC(dynamical model)
 - Sampling time of 50ms
 - Horizon N=40
 - 4th order Runge Kutta discretization
 - Solver: ForcesPro Embotech
 - Controller PC with i7@3.2Ghz
 - Mean Solver time of 10ms (Real time)

MPCC Kinematic (Precision Tracking mode)



MPCC Kinematic (Contouring mode with an obstacle)



MPCC Dynamic(High speed driving)

References

- D. Lam, C. Manzie, and M. Good, "Model predictive contouring control," in Conference on Decision and Control (CDC), pp. 6137–6142, 2010
- Liniger, A., Domahidi, A. and Morari, M.Optimization-based autonomous racing of 1: 43 scale RC cars.Optimal Control Applications and Methods, 2015, 36(5), pp.628-647.
- J. Kong, M. Pfeiffer, G. Schildbach and F. Borrelli, "Kinematic and dynamic vehicle models for autonomous driving control design," 2015 IEEE Intelligent Vehicles Symposium (IV), Seoul, 2015, pp. 1094-1099, doi: 10.1109/IVS.2015.7225830.
- Kirchner, William & Southward, Steve. (2011). An Anthropomimetic Approach to High Performance Traction Control. Paladyn, Journal of Behavioral Robotics. 2. 25-35. 10.2478/s13230-011-0013-9.

The role F1/10 can play in preparing students for research