A Multi-Pronged Approach to Computational Challenges in HJ Reachability

Mo Chen | mochen@cs.sfu.ca

Multi-Agent Robotic Systems Lab | sfumars.com

School of Computing Science, Simon Fraser University

Introduction

- Assistant Professor, School of Computing Science, Simon Fraser University
 - SFU Multi-Agent Robotics Lab
 - SFU Visual Computing Group
- I love multidisciplinary research and collaborations
 - Postdoc in Stanford AA with Marco Pavone
 - PhD in Berkeley EECS with Claire Tomlin
 - Bachelors in UBC Engineering Physics
- Dimensionality reduction
 - System decomposition



- Optimized software
 - Memory locality and parallelism

Outline

- The challenge: curse of dimensionality
 - Finite difference: a numerically convergent approach
- Dimensionality reduction
 - "Chained" systems and projections
- Optimizing software
 - Maximizing brute-force computation capabilities
- Novel applications of HJ PDE solutions in machine learning

Outline

- The challenge: curse of dimensionality
 - Finite difference: a numerically convergent approach
- Dimensionality reduction
 - "Chained" systems and projections
- Optimizing software
 - Maximizing brute-force computation capabilities
- Novel applications of HJ PDE solutions in machine learning

Challenges in Safety-Critical Systems

- Account for all possible system behaviors
 - Formal verification is needed
- Complex environment
 - Weather conditions
 - Unpredictable
 - Adversarial agents
- Complex dynamics
- High-dimensional system dynamics



The Hamilton-Jacobi PDE

- One version of the equation $\frac{\partial V}{\partial t} + \max_{u} \min_{d} \left[\left(\frac{\partial V}{\partial x} \right)^{\mathsf{T}} f(x, u, d) \right] = 0, \quad t \le 0, \quad V(0, x) = l(x)$
- Solved on a grid
- Numerically convergent
 - As state space and time discretization go to zero, we approach the exact solution

Advantages

- Globally optimal solution
- Flexibility: variants of value/cost functions

$$V(t, x(t)) = \min_{\Gamma[u](\cdot)} \max_{u(\cdot)} l(x(0)) \qquad \Rightarrow \frac{\partial V}{\partial t} + \max_{u} \min_{d} \left[\left(\frac{\partial V}{\partial x} \right)^{\mathsf{T}} f(x, u, d) \right] = 0$$
$$x(t)) = \min_{u} \max_{d} \min_{u} \left[\left(s, x(s) \right) \right] \qquad \Rightarrow \min_{u} \left\{ \frac{\partial V}{\partial t} + \max_{u} \min_{d} \left[\left(\frac{\partial V}{\partial t} \right)^{\mathsf{T}} f(x, u, d) \right] \right] \left[l(t, x) - V(t, x) \right] = 0$$

- $V(t,x(t)) = \min_{\Gamma[u](\cdot)} \max_{u(\cdot)} \min_{s \in [t,0]} l(s,x(s)) \qquad \Rightarrow \min\left\{\frac{\partial V}{\partial t} + \max_{u} \min_{d} \left[\left(\frac{\partial V}{\partial x}\right) f(x,u,d)\right], l(t,x) V(t,x)\right\} = 0$
 - Several other variants for different reachability problems

Advantages

• Disturbances

$$V(t, x(t)) = \min_{\Gamma[u](\cdot)} \max_{u(\cdot)} l(x(T)) \qquad \Rightarrow \frac{\partial V}{\partial t} + \max_{u} \min_{d} \left[\left(\frac{\partial V}{\partial x} \right)^{\mathsf{T}} f(x, u, d) \right] = 0$$

- Nonlinear dynamics
- Can represent sets of arbitrary shapes





Challenges in Safety-Critical Systems

- Account for all possible system behaviors
 - Formal verification is needed
- Complex environment
 - Weather conditions
 - Unpredictable
 - Adversarial agents
- Continuous-time system dynamics
- High-dimensional system dynamics



Main Challenge: Exponential Computational Complexity

6D: intractable!

5D:

days





Car Experiment

- More complex model (7D)
 - Intractable \rightarrow simplify model
 - Modelling errors resulted in positional errors of up to ~15cm
- Rate of change in turn rate and acceleration is important!



Leung et al., IJRR 2018

Outline

- The challenge: curse of dimensionality
 - Finite difference: a numerically convergent approach
- Dimensionality reduction
 - "Chained" systems and projections
- Optimizing software
 - Maximizing brute-force computation capabilities
- Novel applications of HJ PDE solutions in machine learning

Outline

- The challenge: curse of dimensionality
 - Finite difference: a numerically convergent approach
- Dimensionality reduction
 - "Chained" systems and projections
- Optimizing software
 - Maximizing brute-force computation capabilities
- Novel applications of HJ PDE solutions in machine learning



State Dependency-Based Decomposition



Missing States

- (z_3, z_4) is self-contained
 - Induces a bound for z_4 , given z_3
- Evolution of (z_2, z_3) depends on z_4 , which is unknown
 - Use worst-case z_4 given all available information, in this case value of z_3
 - This induces a bound for z_3 , given z_2
- Evolution of (z_1, z_2) depends on z_3 , which is unknown
 - Use worst-case z_3 given all available information, in this case value of z_2



Result: Over-approximation

- Coloured shapes represent projections of reachable sets onto each subspace
- Over-approximation of true reachable set is obtained by back-projection and intersection
- Safety is guaranteed



Computational Complexity

- Memory:
 - Need to store value functions for each subsystem
 - So maximum dimension of subsystems
 - (Black nodes, 2 in this case)
- Time:
 - Need to compute value function for each subsystem
 - Need to search over missing variables
 - So maximum dimension of subsystems, including missing states
 - (Black + blue nodes, 3 in this case)



Missing States

- (z_2, z_3, z_4) is self-contained
 - Induces a bound for z_4 , given (z_2, z_3)



- Evolution of (z_1, z_2, z_3) system depends on z_4 , which is unknown
 - Use worst-case z_4 given all available information, in this case value of (z_2, z_3)

Approximation vs Groundtruth



Over-approximation \Rightarrow guaranteed to avoid unsafe states

6D Bicycle Model

Commonly used to model autonomous vehicles



6D Bicycle Model: Decomposition



2

10

⁰ Y

-10

-10 0 10 X

0 Y

-10

2

10

-10 0 10 X

Other Examples

System configuration	System dynamics	State Dependency Graph	Decomposed State Dependency Graph	Time and space
5D Car (x, y)-position θ - heading v - speed ω - turn rate u_a - accel. control u_{α} - ang. accel. control	$\begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{\theta} \\ \dot{v} \\ \dot{\omega} \end{bmatrix} = \begin{bmatrix} v \cos \theta \\ v \sin \theta \\ \omega \\ u_a \\ u_\alpha \end{bmatrix}$	$\begin{array}{c} x \\ y \\ \end{array} \\ \theta \\ \end{array} \\ \omega \end{array}$	$ \begin{array}{cccccccccccccccccccccccccccccccccccc$	Ground truth: both $O(k^5)$ Decomposition: $O(k^4)$ and $O(k^3)$

Other Examples

System configuration	System dynamics	State Dependency Graph	Decomposed State Dependency Graph	Time and space
5D Car (x, y)-position θ - heading v - speed ω - turn rate u_a - accel. control u_{α} - ang. accel. control	$\begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{\theta} \\ \dot{v} \\ \dot{\omega} \end{bmatrix} = \begin{bmatrix} v \cos \theta \\ v \sin \theta \\ \omega \\ u_a \\ u_\alpha \end{bmatrix}$	$\begin{array}{c} x \\ y \\ \end{array} \\ \theta \\ \end{array} \\ \omega \end{array}$	$ \begin{array}{cccccccccccccccccccccccccccccccccccc$	Ground truth: both $O(k^5)$ Decomposition: $O(k^4)$ and $O(k^3)$
6D Planar Quadrotor (x, y)-position (v_x, v_z) - velocity θ - pitch ω - pitch rate u_T - thrust control u_{τ} - ang.accel.control	$\begin{bmatrix} \dot{x} \\ \dot{z} \\ \dot{v}_x \\ \dot{v}_z \\ \dot{\theta} \\ \dot{\omega} \end{bmatrix} = \begin{bmatrix} v_x \\ v_z \\ -u_T \sin \theta \\ u_T \cos \theta - g \\ \omega \\ u_\tau \end{bmatrix}$	$\begin{array}{c} x \\ x \\ \end{array} \\ \end{array} \\ \end{array} \\ \end{array} \\ \begin{array}{c} \\ \\ \\ \\ \\ \\ \\ \\ \\ \\ \\ \\ \\ \\ \\ \\ \\ \\ \\$	$\begin{array}{c} x \longrightarrow v_x \longrightarrow \theta \longrightarrow \omega \\ \hline z \longrightarrow v_z \longrightarrow \theta \longrightarrow \omega \\ \hline v_x \longrightarrow \theta \longrightarrow \omega \\ \hline v_z \longrightarrow \theta \longrightarrow \omega \end{array}$	Ground truth: both $O(k^6)$ Decomposition: $O(k^4)$ and $O(k^3)$

Future Work and Limitations

- Leaking corners
 - Reconstructing the full-dimensional reachable set involves taking intersections or union of subsystem reachable sets
 - For collision avoidance, one must take the intersection, not union $\overline{z_1}$
 - Well-known problem in optimal control and differential games
- Sets must exist in other subspaces
 - For example, reachable set cannot be empty in (z_2, z_3)



 Z_2

Future Work and Limitations

• Automatic decomposition based on directed graph



 Z_1

Outline

- The challenge: curse of dimensionality
 - Finite difference: a numerically convergent approach
- Dimensionality reduction
 - "Chained" systems and projections
- Optimizing software
 - Maximizing brute-force computation capabilities
- Novel applications of HJ PDE solutions in machine learning

Outline

- The challenge: curse of dimensionality
 - Finite difference: a numerically convergent approach
- Dimensionality reduction
 - "Chained" systems and projections
- Optimizing software
 - Maximizing brute-force computation capabilities
- Novel applications of HJ PDE solutions in machine learning

Optimizing Numerical Methods for PDEs

- Much of the progress in deep learning depends on being able to utilize a large amount of compute resources
- PDE solvers are most often run locally
- Can we alleviate computational challenges via software optimization?
 - Different hardware such as CPU, GPU, FPGA and cloud compute
 - Enables analysis of higher-dimensional subsystems
 - Enables real-time applications





Time-Dependent HJ PDE

- Consider the time-dependent HJ PDE
 - Similar observations can be made for other variants



Optimization

- Memory locality and loop order
- Parallelism
- HeteroCL
- Tiling
- Dynamics-dependent loop order

Computer Memory Pyramid

Register Memory Cache Increasing order of access time ratio Memory Primary Memory Main Memory **Magnetic Disks** Auxillary Memory **Magnetic Tapes**

Relative access latency comparison:

- Register memory : (1 x)
- Cache : (~ 16 x)
- Main memory: ($\sim 60 \text{ x}$)
- Disk: (~ 100 x 1000 x)

Cache

- keeps a copy of a neighbourhood of accessed memory
 - Flushed and replaced if data is unused after a while 32

Memory Locality and Loop order

- Memory is linear
- When a memory location is accessed, "neighbours" are loaded into the CPU cache to improve access times
- Therefore, in the example it is faster to iterate over rows



Memory Locality and Loop order

- Memory is linear
- When a memory location is accessed, "neighbours" are loaded into the CPU cache to improve access times
- Therefore, in the example it is faster to iterate over rows





Memory Locality and Loop order

- Memory is linear
- When a memory location is accessed, "neighbours" are loaded into the CPU cache to improve access times

1,0

memory

2,0

• • •

• • •

CPU cache

• Therefore, in the example it is faster to iterate over rows



Parallelism

- Memory is linear
- When a memory location is accessed, "neighbours" are loaded into the CPU cache to improve access times
- Therefore, in the example it is faster to iterate over rows



HeteroCL

- HeteroCL is a programming infrastructure composed of a Python-based domain-specific language (DSL) and a compilation flow
 - <u>http://heterocl.csl.cornell.edu/web/</u>



HeteroCL

- Start with a naive version of the code (e.g. simple nested for loops)
 - Transformations of code for optimization purposes automatically done with HeteroCL
 - No need to manually re-implement anything
 - Guarantees algorithmic correctness
- Main algorithm written in Python (user friendly)
 - HeteroCL-specific syntax translates Python code to to high performance C code
- Easy to re-use and optimize code for different hardware platforms in mind (e.g. CPUs, GPUs, FPGAs)

Results: Dubins Car

 $\dot{x} = v \cos \theta$ $\dot{y} = v \sin \theta$ $\dot{\theta} = \omega$

- Memory locality and loop order + parallelism
 - Implemented with HeteroCL
 - Times reported in seconds, per integration time step

Intel i9-9900K CPU at 3.8 GHz (Desktop)

Grid size	80 ³	90 ³	100 ³
MATLAB level set toolbox	0.047	0.0726	0.10

Results: Dubins Car

 $\dot{x} = v \cos \theta$ $\dot{y} = v \sin \theta$ $\dot{\theta} = \omega$

- Memory locality and loop order + parallelism
 - Implemented with HeteroCL
 - Times reported in seconds, per integration time step

Intel i9-9900K CPU at 3.8 GHz (Desktop)

AMD A10-8700P CPU at 3.2 GHz (Laptop)

Grid size	80 ³	90 ³	100³	Grid size	80 ³	90 ³	100 ³
MATLAB level set toolbox	0.047	0.0726	0.10	MATLAB level set toolbox	0.58	0.83	1.05
heteroCL	0.0018	0.0022	0.0025	heteroCL	0.011	0.015	0.018
Speed-up	26 times	30 times	38 times	Speed-up	54 times	57 times	58 times

Results: 6D Underwater Vehicle

$$\begin{aligned} \dot{x}_{\alpha} &= u_{r} + G_{1}(z)\cos(\sigma(x,t)) + d_{x} - b_{x} \\ \dot{z}_{\alpha} &= w_{r} + G_{1}(z)(-\sin(\sigma(x,t))) + d_{z} - b_{z} \\ \dot{u}_{r} &= \frac{1}{m - X_{\dot{u}}}\left(\Phi_{11}(b - X_{\dot{u}})u_{r} + \cdots \right. \\ \dot{w}_{r} &= \frac{1}{m - Z_{\dot{w}}}\left(\Phi_{22}(b - Z_{\dot{w}})w_{r} + \cdots \right. \\ \dot{x} &= u_{r} + G_{1}(z)\cos(\sigma(x,t)) + d_{x} \\ \dot{z} &= w_{r} + G_{1}(z)(-\sin(\sigma(x,t))) + d_{z} \end{aligned}$$

- Tracking error bounds under plane-progressive waves
 - Grid size: 25⁶
 - 5.1 seconds per step
 - Previously intractable

Siriya et al., CDC '20 (submitted)



 $t = 20s, u_r = -0.250000, w_r = 0.350000, x_s = 2.500000, z_s = 6.000000 \qquad t = 20s, u_r = -0.250000, w_r = 0.350000, x_s = 2.500000, z_s = 10.000000, z_s = 10.00000, z_s = 10.0000, z_s = 10.00000, z$



Results: 6D Simplified Humanoid Model

$$\dot{x}_{1} = x_{2}$$

$$\dot{x}_{2} = \frac{g + u_{2}}{x_{3}}(x_{1} + u_{1}) + u_{3}$$

$$\dot{x}_{3} = x_{4}$$

$$\dot{x}_{4} = u_{2}$$

$$\dot{x}_{5} = x_{6}$$

$$\dot{x}_{6} = \frac{x_{3}}{J}u_{3}$$

- Backward reachable set
 - Grid size: 25⁶
 - 2 seconds per step
 - Previously intractable









Future work

- Optimizations
 - Tiling
 - Dynamics-dependent loop order

Tiling

- Load a part of memory into CPU cache
 - Perform all computations that require this data
 - Nested loops that iterate only over this data
- Load another part of memory into CPU cache
 - This replaces old data, which is no longer needed
 - Perform all computations that require this data
- Minimize data movement

0,0	0,1	0,2	•••		
1,0	1,1	1,2	•••		
•••	•	:	*.		

Dynamics-Dependent Loop Order

• Dubins Car dynamics: $\dot{x} = v \cos \theta$

$$\dot{y} = v \sin \theta$$
$$\dot{\theta} = \omega$$

- θ does not depend on x and y
- So it would make sense to iterate over different θ values in the outer loop
 - Values of \dot{x} and \dot{y} would change less often in the inner loops
- Combine this with tiling

Future work

- Optimizations
 - Tiling
 - Dynamics-dependent loop order
- Applications
 - Real-time reachability computations
 - Model fidelity study
- Easy-to-use toolbox



Outline

- The challenge: curse of dimensionality
 - Finite difference: a numerically convergent approach
- Dimensionality reduction
 - "Chained" systems and projections
- Optimizing software
 - Maximizing brute-force computation capabilities
- Novel applications of HJ PDE solutions in machine learning

Outline

- The challenge: curse of dimensionality
 - Finite difference: a numerically convergent approach
- Dimensionality reduction
 - "Chained" systems and projections
- Optimizing software
 - Maximizing brute-force computation capabilities
- Novel applications of HJ PDE solutions in machine learning

Novel Applications in Machine Learning

- Reinforcement Learning
 - Using HJ PDE solutions for reward shaping
- Visual Navigation
 - Using HJ PDE solutions as training data

Model-Free Reinforcement Learning

- Given: simulator of a system (or real system) and reward function
 - Determine (locally) optimal policy without knowing a model of the system
 - Try control policies that involve exploration and exploitation
 - Use control policies that acquire more reward more often
- Advantages
 - Does not require system dynamics
 - Flexible in terms of input data (eg. sensor data)
 - Scalable to large system state spaces when using function approximators
- Disadvantages
 - Poor sample efficiency and generalizability
 - Difficult to incorporate knowledge of system behaviour

Reward Shaping

- Tasks are sometimes specified by a sparse reward
 - Goal: high reward (e.g. +1000)
 - Obstacles: very negative reward (e.g. -5000)
 - Other states: 0 reward
- Very little reward signal!
- A couple of heuristics for reward shaping to achieve better signal during training
 - Distance to goal (closer is better)
 - Inverse reinforcement learning (requires expert demonstration)

TTR-Based Reward Shaping

- TTR: "time-to-reach"
 - Minimum time required to reach the goal from any state *s*, *T*(*s*)
 - Requires system dynamics $\dot{s} = f(s, u)$
 - Intractable to compute for complex systems
- Idea: compute the TTR approximately using a simplified system model $\dot{\hat{s}} = \hat{f}(\hat{s}, \hat{u})$
 - Approximately account for system dynamics
- Approximate TTR: $\hat{T}(\hat{s}) \rightarrow$ reward: $R(s) = -\hat{T}(\hat{s})$
 - "Good" states are those that require less time to reach the target, given system dynamics

52



Planar Quadrotor With LIDAR Example



Quadrotor simulation

Success From Start State: Value-Based RL



Success From Start State: Actor-Critic RL



Visual Navigation







- Navigate through an environment using camera only
 - Cameras are cheap
- End-to-end: pixels \rightarrow actions
- Control-inspired: pixels \rightarrow waypoint
 - Use control theory to plan trajectory to the waypoint
 - No need to learn what we already know
 - Train in simulation (Stanford indoor data set)
 - Results are transferrable to real life



Bansal et al. CoRL 2019 Li et al. L4DC 2020

Visual Navigation

- Challenge:
 - How to generate training data?
- Need to know what is a "good" waypoint
 - Waypoints are chosen based on simple heuristics involving distances to goal and obstacles may not be sufficient





Visual Navigation

- Supervision by HJ PDE solution
 - Time-to-reach and time-tocollision functions can be used as a metric for the quality of a waypoint
- Data generation is timeconsuming





Visual Navigation: Training Data

Distance heuristic





Distance heuristic

Time-to-reach



Time-to-reach

Visual Navigation – Test Performance



Agent	Success (%)	Time taken (s)	Acceleration (m/s ²)	Jerk (m/s ³)
WayPtNav-ReachabilityCost	63.82	21.00 ± 8.00	0.06 ± 0.01	0.94 ±0.13
WayPtNav-HeuristicsCost	52.26	18.82 ±5.66	0.07 ± 0.02	1.06 ± 0.15
WayPtNav-ReachabilityCost-NoDstb	49.24	16.19 ±4.8	0.07 ± 0.01	0.98 ± 0.16
E2E-ReachabilityCost	8.04	19.55 ±4.72	0.07 ± 0.01	2.16 ± 0.30
E2E-HeuristicsCost	31.66	25.56 ±9.85	0.26 ± 0.06	9.06 ± 1.94



Success rate at different difficulty levels

Thank you!

Mo Chen mochen@cs.sfu.ca https://sfumars.com

- The challenge: curse of dimensionality
 - Finite difference: a numerically convergent approach
- Dimensionality reduction
 - "Chained" systems and projections
- Optimizing software
 - Maximizing brute-force computation capabilities
- Novel applications of HJ PDE solutions in machine learning