

Balancing Actuation Energy and Computing Energy in Motion Planning

Soumya Sudhakar, Vivienne Sze, Sertac Karaman

Low Energy Autonomy and Navigation (LEAN) Group

CICS Talk - May 5, 2021



Planning a Path to the Coffeeshop




Planning a Path to the Coffeeshop




flour
bakery + cafe

Don't Think Too Hard



Don't Think Too Hard



flour
bakery + cafe

Don't Think Too Hard




How do we get an energy-constrained robot to decide when it has computed enough?

flour
bakery + cafe

Miniature or Long-Duration Robotics are Power-Constrained



Power-constrained
due to size



Power-constrained
due to duration


Miniature or Long-Duration Robotics Enable New Solutions




Infrastructure inspections



Persistent environmental
monitoring




Noninvasive targeted drug
delivery



Space exploration

Many applications exist for miniature or long-duration robotic platforms that can intelligently navigate

Recent Advances in Low-Power Robotic Platforms




Success in actuating miniature and long-duration robotics at low power in the lab and real-world


Less Attention Paid to Energy Computing Consumes

Avg. Energy per Meter [J/m]


Cassie bipedal robot
[Source: Agility Robotics]
Kashiri et al. 2018



Tesla Model S at 70 mph
[Source: Tesla]
Sherman 2014, *Car and Driver*




KUKA arm
[Source: KUKA]
Grebers et al. 2017



ASIC, FPGA
[Source: Xilinx]
(power dependent on
hardware design)



Cortex-A7



Cortex-A15




Nvidia Jetson TX2
[Source: Nvidia]
GPU

Avg. Energy per Second [J/s]


Less Attention Paid to Energy Computing Consumes

Avg. Energy per Meter [J/m]


Cassie bipedal robot
[Source: Agility Robotics]
Kashiri et al. 2018



Tesla Model S at 70 mph
[Source: Tesla]
Sherman 2014, *Car and Driver*




KUKA arm
[Source: KUKA]
Grebers et al. 2017



ASIC, FPGA
[Source: Xilinx]
(power dependent on
hardware design)



Cortex-A7



Cortex-A15



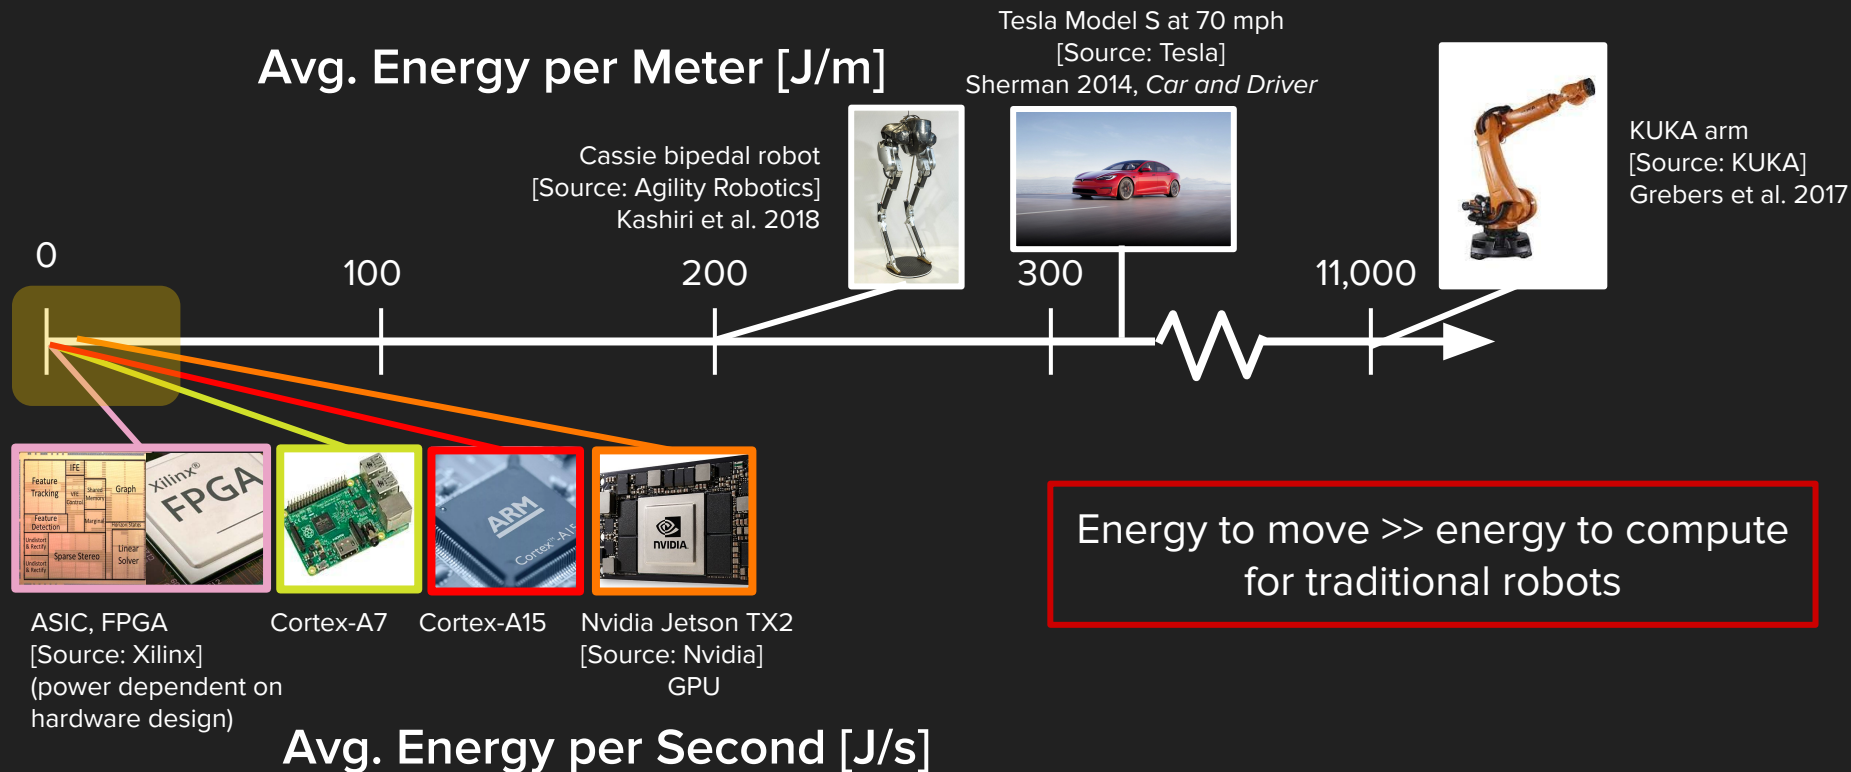
Nvidia Jetson TX2
[Source: Nvidia]
GPU

Energy to move >> energy to compute
for traditional robots

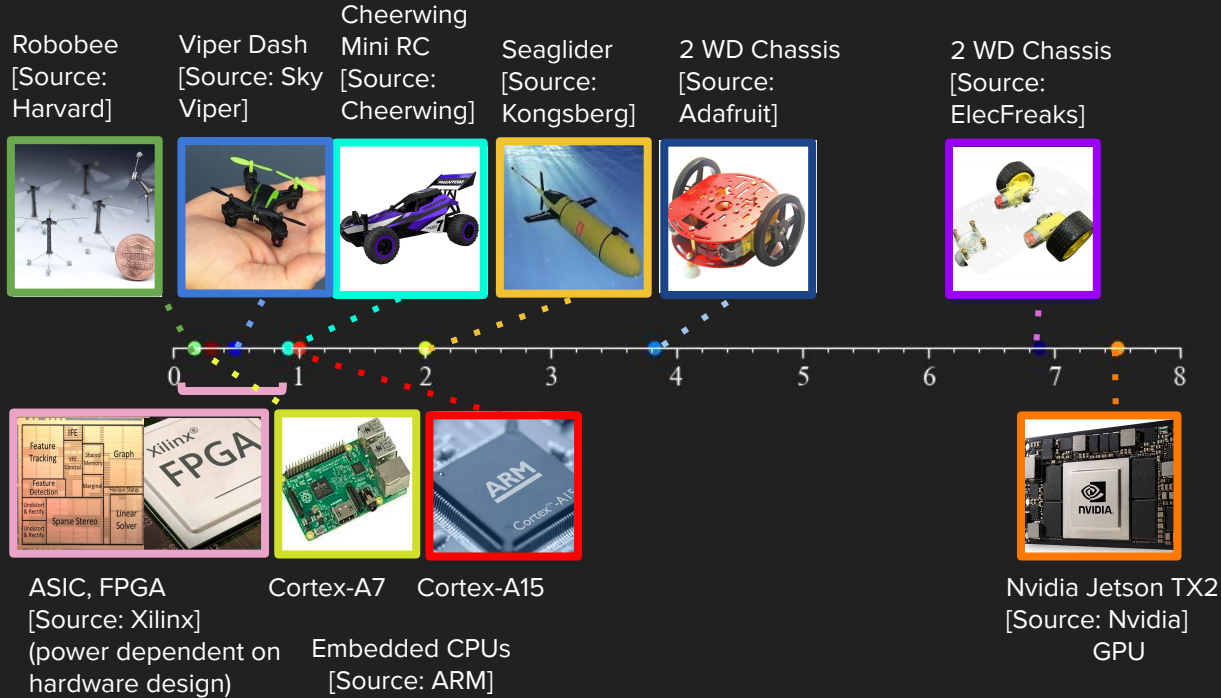
Avg. Energy per Second [J/s]

Less Attention Paid to Energy Computing Consumes

Avg. Energy per Meter [J/m]



Avg. Energy per Meter [J/m]




For low-power robotics, energy to move and energy to compute are on a similar magnitude

Avg. Energy per Second [J/s]

Background: Motion Planning

The motion planning problem


Plan the shortest
path from the
start to the goal
avoiding all
obstacles




Background: Motion Planning

Plan the shortest
path from the
start to the **goal**
avoiding all
obstacles

The motion planning problem



Sampling-based motion planner




● Node
— Edge

Sampling-based motion planner find paths by sampling and connecting nodes in free space


Background: Motion Planning

Plan the shortest
path from the
start to the **goal**
avoiding all
obstacles

The motion planning problem



Sampling-based motion planner




● Node
— Edge

Sampling-based motion planner find paths by sampling and connecting nodes in free space


Background: Motion Planning

Plan the shortest
path from the
start to the **goal**
avoiding all
obstacles

The motion planning problem



Sampling-based motion planner




● Node
— Edge

Sampling-based motion planner find paths by sampling and connecting nodes in free space


Background: Motion Planning

Plan the shortest
path from the
start to the **goal**
avoiding all
obstacles

The motion planning problem



Sampling-based motion planner




● Node
— Edge

Sampling-based motion planner find paths by sampling and connecting nodes in free space


Background: Motion Planning

Plan the shortest
path from the
start to the **goal**
avoiding all
obstacles

The motion planning problem



Sampling-based motion planner




● Node
— Edge

Sampling-based motion planner find paths by sampling and connecting nodes in free space


Background: Motion Planning

Plan the shortest
path from the
start to the **goal**
avoiding all
obstacles

The motion planning problem



Sampling-based motion planner




● Node
— Edge

Sampling-based motion planner find paths by sampling and connecting nodes in free space


Background: Motion Planning

Plan the shortest
path from the
start to the **goal**
avoiding all
obstacles

The motion planning problem



Sampling-based motion planner




● Node
— Edge

Sampling-based motion planner find paths by sampling and connecting nodes in free space


Background: Motion Planning

Plan the shortest
path from the
start to the **goal**
avoiding all
obstacles

The motion planning problem



Sampling-based motion planner




● Node
— Edge

Sampling-based motion planner find paths by sampling and connecting nodes in free space

Background: Motion Planning


Plan the shortest
path from the
start to the **goal**
avoiding all
obstacles

The motion planning problem



Computing more nodes → find shorter paths

Sampling-based motion planner




● Node
— Edge

Sampling-based motion planner find paths by sampling and connecting nodes in free space


Background: Motion Planning

Plan the shortest
path from the
start to the **goal**
avoiding all
obstacles

The motion planning problem



Sampling-based motion planner




● Node
— Edge

Sampling-based motion planner find paths by sampling and connecting nodes in free space


Background: Motion Planning

Plan the shortest
path from the
start to the **goal**
avoiding all
obstacles

The motion planning problem



Sampling-based motion planner




● Node
— Edge

Sampling-based motion planner find paths by sampling and connecting nodes in free space


Background: Motion Planning

Plan the shortest
path from the
start to the **goal**
avoiding all
obstacles

The motion planning problem



Sampling-based motion planner




● Node
— Edge

Sampling-based motion planner find paths by sampling and connecting nodes in free space


Background: Motion Planning

Plan the shortest
path from the
start to the **goal**
avoiding all
obstacles

The motion planning problem



Sampling-based motion planner




● Node
— Edge

Sampling-based motion planner find paths by sampling and connecting nodes in free space


Background: Motion Planning

Plan the shortest
path from the
start to the **goal**
avoiding all
obstacles

The motion planning problem



Sampling-based motion planner



● Node
— Edge

Sampling-based motion planner find paths by sampling and connecting nodes in free space

Background: Motion Planning


E_a

Actuation energy = energy consumed by vehicle's actuators
(e.g., motors) to move along path

E_c


Computing energy = energy consumed by computer
onboard vehicle to compute the path

Total Energy of Actuation and Computing



Simulated vehicle that can travel 1 m/s at 1 Watt,
computing on a Cortex A15 (embedded CPU)

Total Energy of Actuation and Computing




Simulated vehicle that can travel 1 m/s at 1 Watt,
computing on a Cortex A15 (embedded CPU)

Total energy

$$E_t = E_a$$

Actuation energy
(energy to move
along path)

Total Energy of Planning and Moving



Simulated vehicle that can travel 1 m/s at 1 Watt, computing on a Cortex A15 (embedded CPU)

Total energy



$$E_t = E_a + E_c$$



Actuation energy
(energy to move
along path)

Computing energy
(energy to
compute path)

Reducing total energy is now an early stopping problem

The Work of Actuation and Computation

Actuation	Computing
- Path length [m], $l_a(n)$	Num. of nodes [nodes], n Num. of operations [ops], $l_c(n)$

The work of actuation and the work of computing have analogous variables

The Work of Actuation and Computation

Actuation	Computing
- Path length [m], $l_a(n)$ Vehicle speed [m/s], v_a	Num. of nodes [nodes], n Num. of operations [ops], $l_c(n)$ Processing speed [ops/s], v_c

The work of actuation and the work of computing have analogous variables

The Work of Actuation and Computation

Actuation	Computing
- Path length [m], $l_a(n)$ Vehicle speed [m/s], v_a Actuation power [W], $P_a(v_a)$	Num. of nodes [nodes], n Num. of operations [ops], $l_c(n)$ Processing speed [ops/s], v_c Computing power [W], $P_c(v_c)$

The work of actuation and the work of computing have analogous variables

The Work of Actuation and Computation

Actuation	Computing
- Path length [m], $l_a(n)$ Vehicle speed [m/s], v_a Actuation power [W], $P_a(v_a)$ Actuation energy [J], E_a	Num. of nodes [nodes], n Num. of operations [ops], $l_c(n)$ Processing speed [ops/s], v_c Computing power [W], $P_c(v_c)$ Computing Energy [J], E_c

The work of actuation and the work of computing have analogous variables

The Work of Actuation and Computation

Actuation	Computing
-	Num. of nodes [nodes], n
Path length [m], $l_a(n)$	Num. of operations [ops], $l_c(n)$
Vehicle speed [m/s], v_a	Processing speed [ops/s], v_c
Actuation power [W], $P_a(v_a)$	Computing power [W], $P_c(v_c)$
Actuation energy [J], E_a	Computing Energy [J], E_c


The work of actuation and the work of computing have analogous variables

The Work of Actuation and Computation

Actuation	Computing
- Path length [m], $l_a(n)$ Vehicle speed [m/s], v_a Actuation power [W], P_a Actuation energy [J], E_a	Num. of nodes [nodes], n Num. of operations [ops], $l_c(n)$ Processing speed [ops/s], v_c Computing power [W], P_c Computing Energy [J], E_c

The work of actuation and the work of computing have analogous variables

Total Energy of Planning and Moving



Simulated vehicle that can travel 1 m/s at 1 Watt,
computing on a Cortex A15 (embedded CPU)

Actuation	Computing
-	Num. of nodes [nodes], n
Path length [m], $l_a(n)$	Num. of operations [ops], $l_c(n)$
Vehicle speed [m/s], v_a	Processing speed [ops/s], v_c
Actuation power [W], $P_a(v_a)$	Computing power [W], $P_c(v_c)$
Actuation energy [J], E_a	Computing Energy [J], E_c

$$E_t = E_a + E_c$$


$$= \underbrace{\frac{P_a}{v_a} l_a(n)}_{\text{Actuation energy}} + \underbrace{\frac{P_c}{v_c} l_c(n)}_{\text{Computing energy}}$$

Actuation energy (energy to move along path)

Computing energy (energy to compute path)

↑ # of nodes n
 ↓ actuation energy
 ↑ computing energy

Total Energy of Planning and Moving



Simulated vehicle that can travel 1 m/s at 1 Watt,
computing on a Cortex A15 (embedded CPU)

Actuation	Computing
-	Num. of nodes [nodes], n
Path length [m], $l_a(n)$	Num. of operations [ops], $l_c(n)$
Vehicle speed [m/s], v_a	Processing speed [ops/s], v_c
Actuation power [W], $P_a(v_a)$	Computing power [W], $P_c(v_c)$
Actuation energy [J], E_a	Computing Energy [J], E_c

$$E_t = E_a + E_c$$

$$= \underbrace{\frac{P_a}{v_a} l_a(n)}_{\text{Actuation energy}} + \underbrace{\frac{P_c}{v_c} l_c(n)}_{\text{Computing energy}} + \underbrace{\frac{P_c}{v_c} \bar{l}_c(n)}_{\text{Overhead energy}}$$

Actuation energy
(energy to move
along path)

Computing energy
(energy to
compute path)

Overhead energy
(energy to decide
when to stop)

Performance metric includes overhead we introduce

Related Work

2000

2015

2020

Reducing actuation or other energy

Incorporating communication energy
[Yan et al (TCNS 2014)]

Incorporating terrain
[Gaganath et al. (TII 2015)]

Reducing computing energy

Lazy PRM
[Bohlin et al. (ICRA 2000)]

Batch-Informed Trees
[Gammell et al. (ICRA 2015)]

Fast-Marching Trees
[Janson et al. (IJRR 2015)]

FPGA acceleration
[Murray et al. (RSS 2016),
Palossi et al. (IoT 2019)]

Considering actuation energy & computing energy

Hierarchical abstractions
[Larsson et al. (2017)]

Related Work

2000

2015

2020

**Reducing
actuation or
other energy**

Incorporating
communication energy
[Yan et al (TCNS 2014)]

Incorporating terrain
[Gaganath et al. (TII 2015)]

**Reducing
computing
energy**

Lazy PRM
[Bohlin et al. (ICRA 2000)]

Batch-Informed Trees
[Gammell et al. (ICRA 2015)]

Fast-Marching Trees
[Janson et al. (IJRR 2015)]



FPGA acceleration
[Murray et al. (RSS 2016),
Palossi et al. (IoT 2019)]

**Considering
actuation energy &
computing energy**



Hierarchical abstractions
[Larsson et al. (2017)]

CEIMP (this work)
[Sudhakar et al. (2020)]



Technical Gap: Knowing How Much Computing is Enough



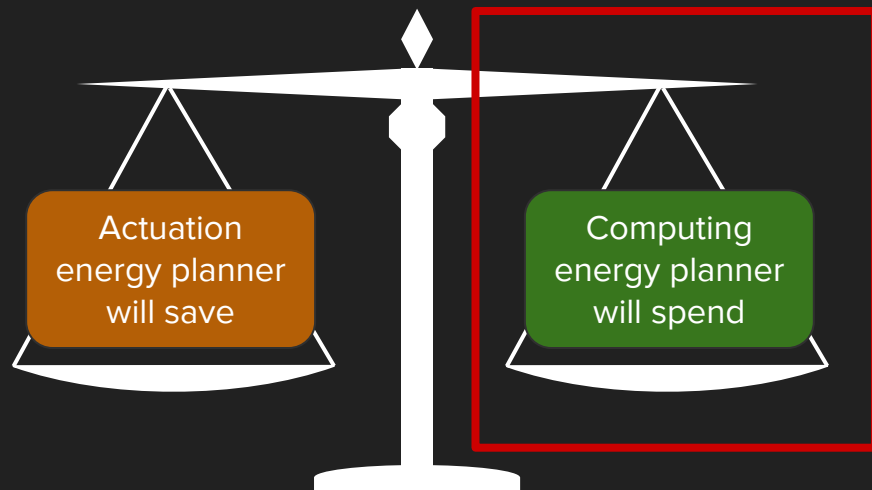
Technical Gap: Knowing How Much Computing is Enough



Technical Gap: Knowing How Much Computing is Enough




Technical Gap: Knowing How Much Computing is Enough





Technical Gap: Knowing How Much Computing is Enough

Improvement in **same**
homotopic class




Improvement from
homotopic class change





Technical Gap: Knowing How Much Computing is Enough

Improvement in **same**
homotopic class




Improvement from
homotopic class change






Technical Gap: Knowing How Much Computing is Enough

Improvement in **same**
homotopic class




Improvement from
homotopic class change





Technical Gap: Knowing How Much Computing is Enough

Improvement in **same**
homotopic class




Improvement from
homotopic class change





Technical Gap: Knowing How Much Computing is Enough

Improvement in **same
homotopic class**



Improvement from
homotopic class change






Technical Gap: Knowing How Much Computing is Enough

Improvement in **same**
homotopic class




Improvement from
homotopic class change






Technical Gap: Knowing How Much Computing is Enough

Improvement in **same homotopic class**




Improvement from **homotopic class change**






Technical Gap: Knowing How Much Computing is Enough

Improvement in **same homotopic class**




Improvement from **homotopic class change**






Technical Gap: Knowing How Much Computing is Enough

Improvement in **same homotopic class**




Improvement from **homotopic class change**






Technical Gap: Knowing How Much Computing is Enough

Improvement in **same homotopic class**




Improvement from **homotopic class change**






Technical Gap: Knowing How Much Computing is Enough

Improvement in **same homotopic class**




Improvement from **homotopic class change**






Technical Gap: Knowing How Much Computing is Enough

Improvement in **same homotopic class**




Improvement from **homotopic class change**






Technical Gap: Knowing How Much Computing is Enough

Improvement in **same homotopic class**




Improvement from **homotopic class change**






Technical Gap: Knowing How Much Computing is Enough

Improvement in **same homotopic class**




Improvement from **homotopic class change**




This work:
Computing Energy Included
Motion Planning (CEIMP)


CEIMP: Underlying Motion Planner




CEIMP: Track edges in collision with the obstacles to “probe” the environment for new homotopic classes




CEIMP: Track edges in collision with the obstacles to “probe” the environment for new homotopic classes



CEIMP: Track edges in collision with the obstacles to “probe” the environment for new homotopic classes

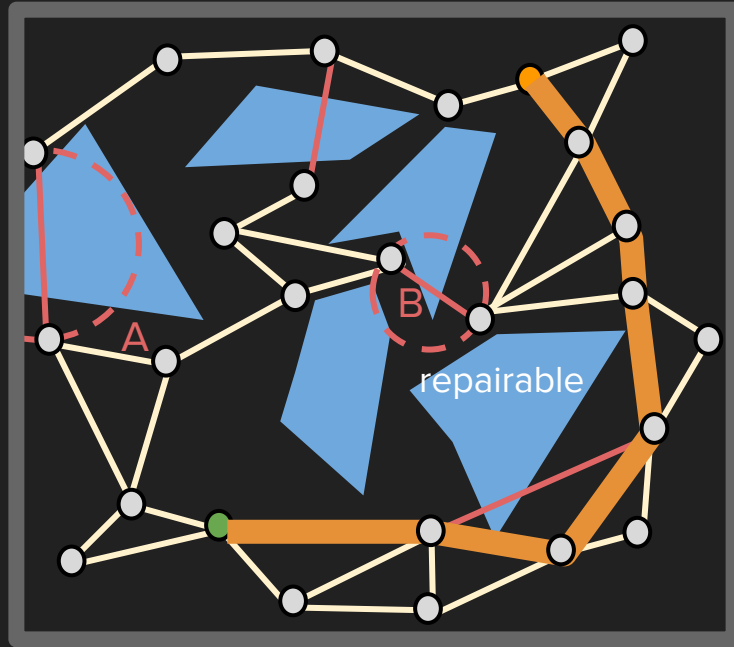


CEIMP: Model the state of each currently edge in collision as 'reparable' or 'unreparable'




CEIMP: Model the state of each currently edge in collision as 'repairable' or 'unrepairable'

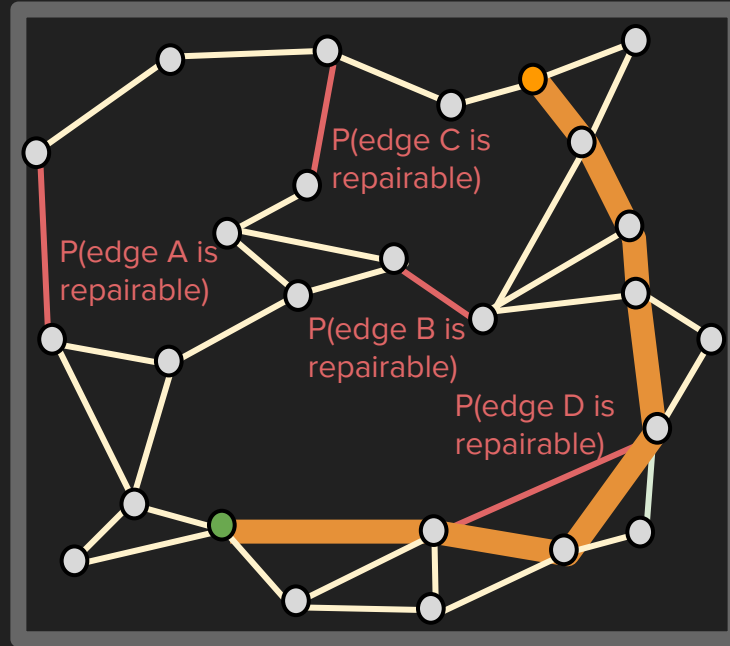
unrepairable



CEIMP: Model the state of each currently edge in collision as 'reparable' or 'unreparable'

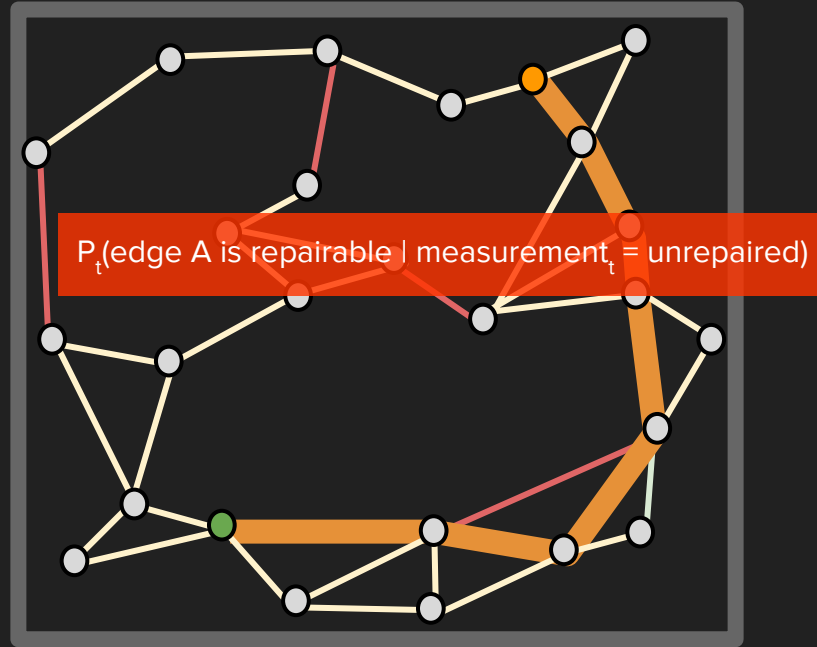


CEIMP: Estimate the probability an edge in collision's state is 'repairable'



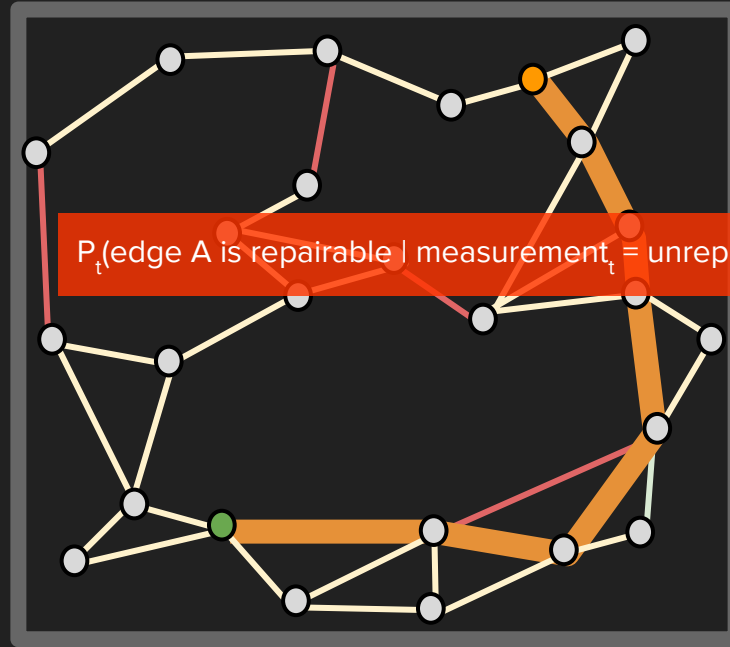
CEIMP: Estimate the probability an edge in collision's state is 'reparable'

Model computing a batch of nodes as a noisy sensor that returns a reading of **repaired** or **unrepaired**



CEIMP: Estimate the probability an edge in collision's state is 'reparable'

Model computing a batch of nodes as a noisy sensor that returns a reading of **repaired** or **unrepaired**



What is the probability an edge in collision is a reparable edge, given we haven't been able to repair it yet?

Computing as a Measurement

edge \in {**repairable**, **unrepairable**}

measurement from sampling a batch of nodes \in {**repaired**, **unrepaired**}

Computing as a Measurement

edge $\in \{\text{repairable}, \text{unrepairable}\}$

measurement from sampling a batch of nodes $\in \{\text{repaired}, \text{unrepaired}\}$

$$P(\text{measurement} = \text{unrepaired}) = \begin{cases} a & \text{if edge is repairable} \\ 1 & \text{if edge is unrepairable} \end{cases}$$

Computing as a Measurement

edge \in {**repairable**, **unrepairable**}


measurement from sampling a batch of nodes \in {**repaired**, **unrepaired**}

$$P(\text{measurement} = \text{unrepaired}) = \begin{cases} a & \text{if edge is } \text{repairable} \\ 1 & \text{if edge is } \text{unrepairable} \end{cases}$$

Binary Bayesian filtering


$$\begin{aligned} P_t(\text{edge A is } \text{repairable} \mid \text{measurement}_t = \text{unrepaired}) \\ = \eta P(\text{measurement} = \text{unrepaired} \mid \text{edge A is } \text{repairable}) P_{t-1}(\text{edge A is } \text{repairable}) \end{aligned}$$

CEIMP: Run a search algorithm on probabilistic graph to return the shortest expected path




Experimental Results

Simulated vehicle that can travel 1 m/s at 1 Watt, computing on a Cortex A15 (embedded CPU)




CEIMP successful at stopping close to true minimum




CEIMP stops too early, misses savings from homotopic class change

Experimental Results


Simulated vehicle that can travel 1 m/s at 1 Watt, computing on a Cortex A15 (embedded CPU)



CEIMP successful at stopping close to true minimum




Example path returned by baseline




Example path returned by CEIMP

Experimental Results


Simulated vehicle that can travel 1 m/s at 1 Watt, computing on a Cortex A15 (embedded CPU)



CEIMP successful at stopping close to true minimum




Example path returned by baseline



Example path returned by CEIMP


On average across 10 MIT floorplans, CEIMP saves **2.1x-8.9x** the energy compared to baseline

As the the energy to compute becomes more expensive relative to the energy to move, CEIMP will increase energy savings




Energy to compute 1 sec relative to the energy to move 1 meter

As the the energy to compute becomes more expensive relative to the energy to move, CEIMP will increase energy savings




Energy to compute 1 sec relative to the energy to move 1 meter

As the the energy to compute becomes more expensive relative to the energy to move, CEIMP will increase energy savings



Energy to compute 1 sec relative to the energy to move 1 meter


As the the energy to compute becomes more expensive relative to the energy to move, CEIMP will increase energy savings



Energy to compute 1 sec relative to the energy to move 1 meter

Key Takeaways


- **Don't think too hard:** A longer path that we have now can be better than a shorter path that we have to compute a long time to find



Sudhakar, Soumya, Sertac Karaman, and Vivienne Sze. "Balancing Actuation and Computing Energy in Motion Planning." *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020.

Key Takeaways

- **Don't think too hard:** A longer path that we have now can be better than a shorter path that we have to compute a long time to find
- **Computing is (noisy) sensing:** Sampling nodes in a motion planner can be modeled as a noisy sensor that returns whether a path is open or closed



Sudhakar, Soumya, Sertac Karaman, and Vivienne Sze. "Balancing Actuation and Computing Energy in Motion Planning." *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020.