

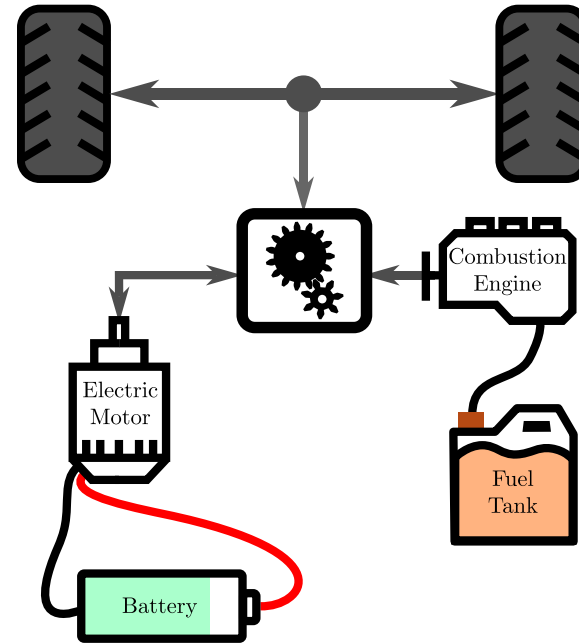
# Scalable FPGA Implementation of Dynamic Programming for Optimal Control of Hybrid Electrical Vehicles

Frans Skarman, Oscar Gustafsson

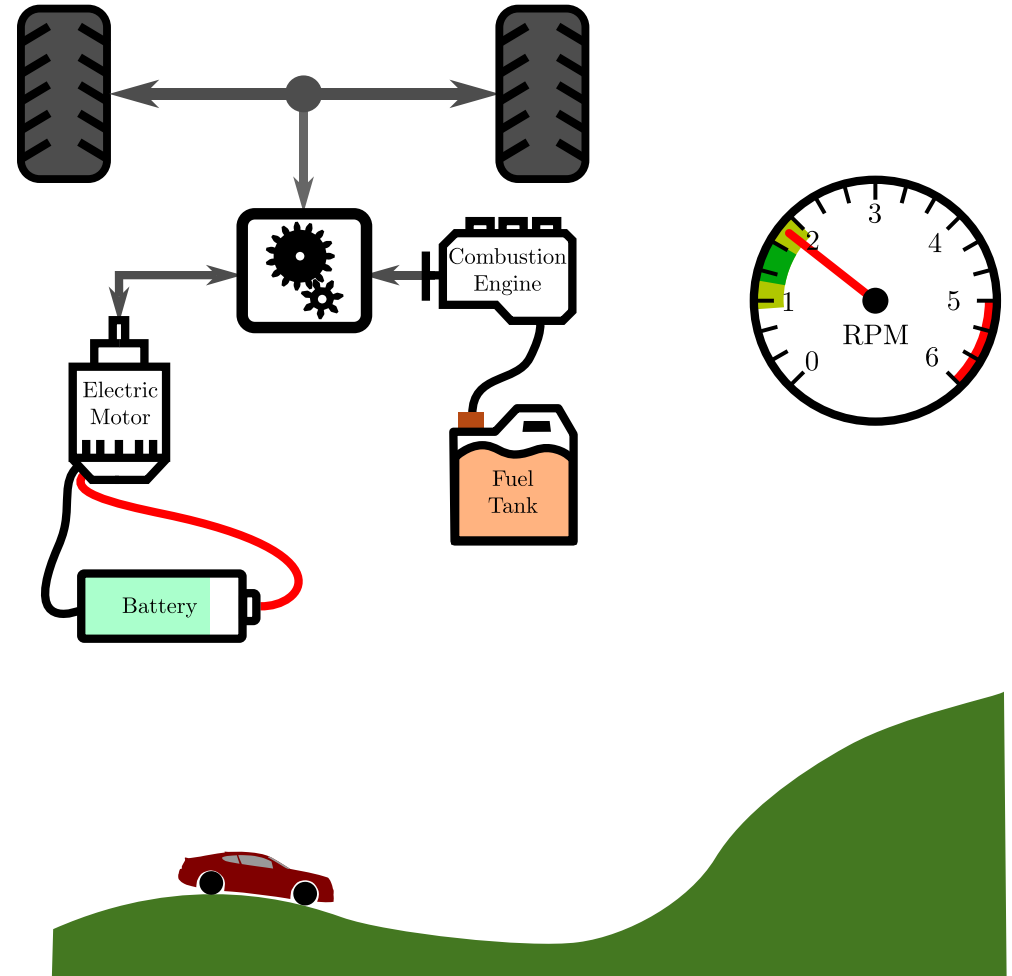
Linköping University



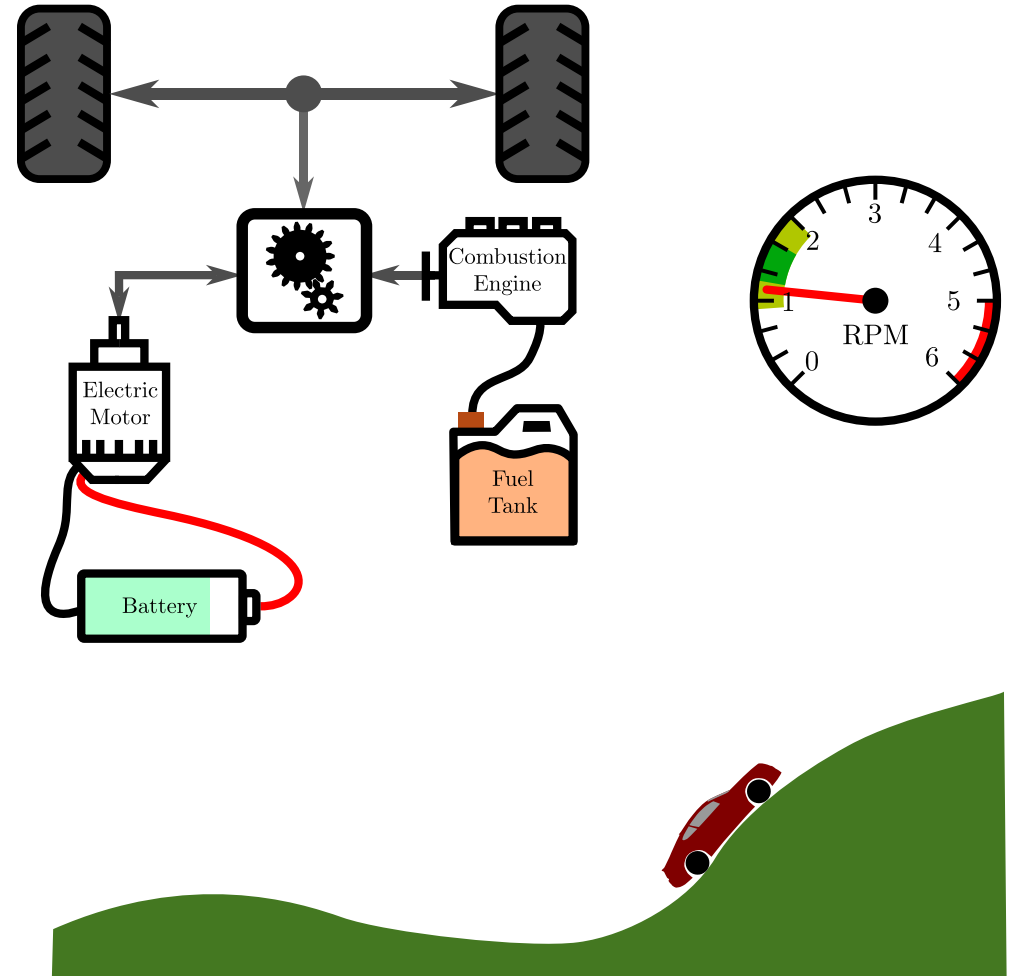
Determine optimal split  
between electric motor and  
combustion engine



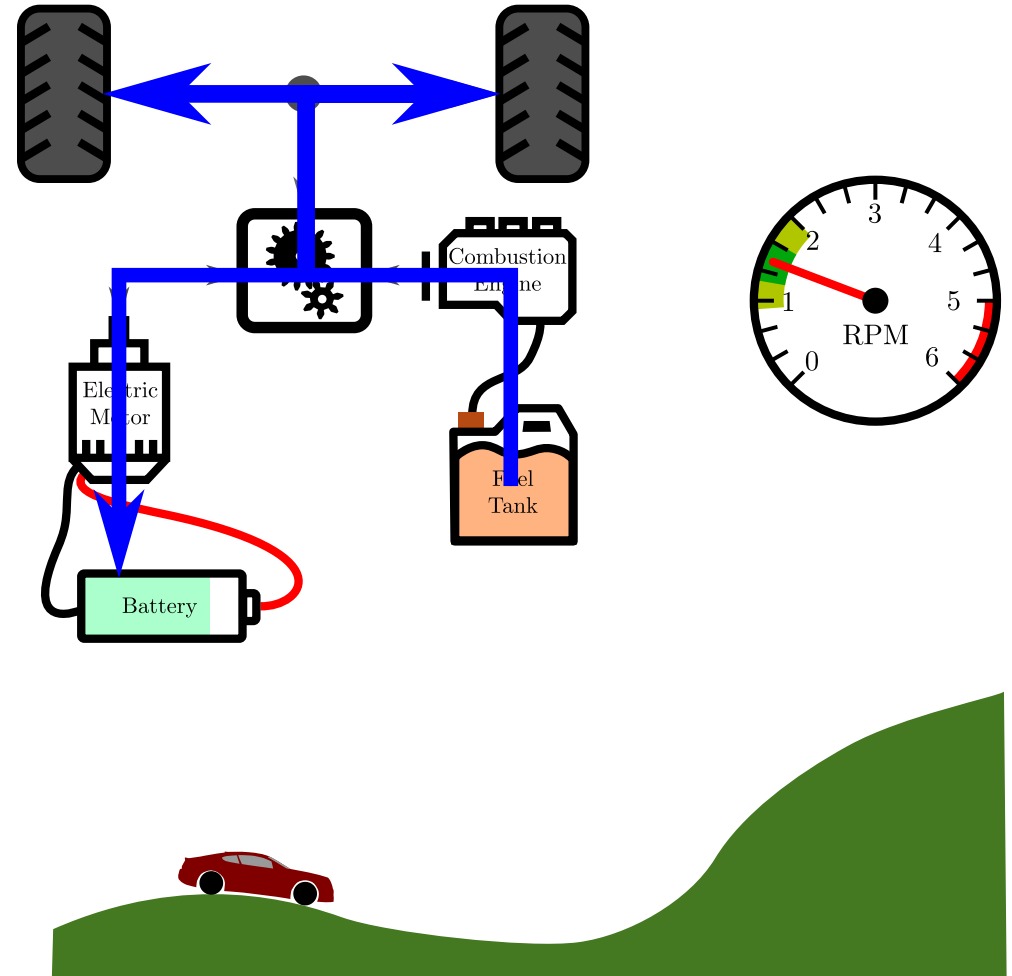
Determine optimal split  
between electric motor and  
combustion engine



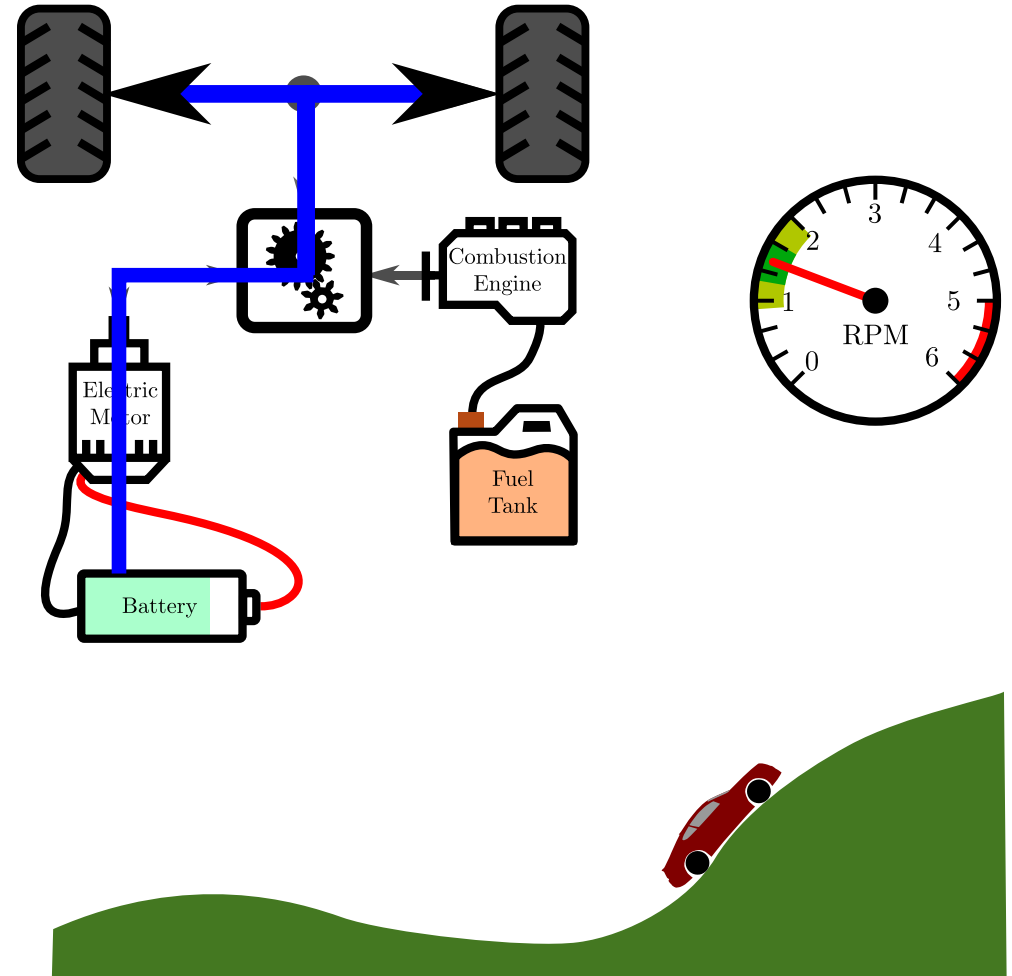
Determine optimal split  
between electric motor and  
combustion engine



Determine optimal split  
between electric motor and  
combustion engine

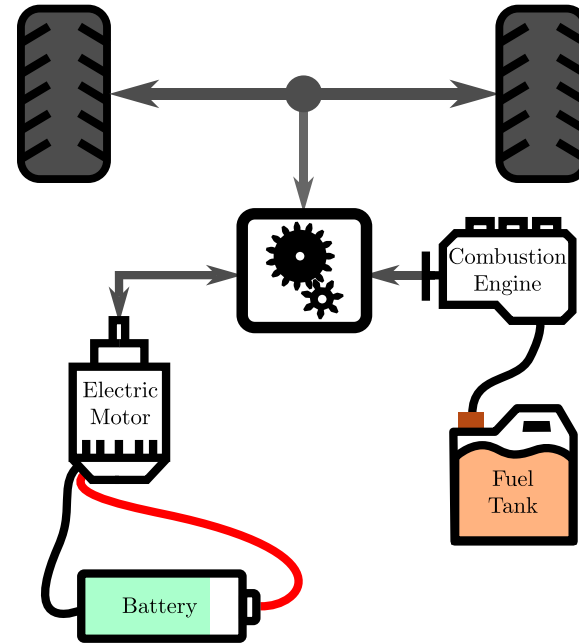


Determine optimal split  
between electric motor and  
combustion engine



Determine optimal split  
between electric motor and  
combustion engine

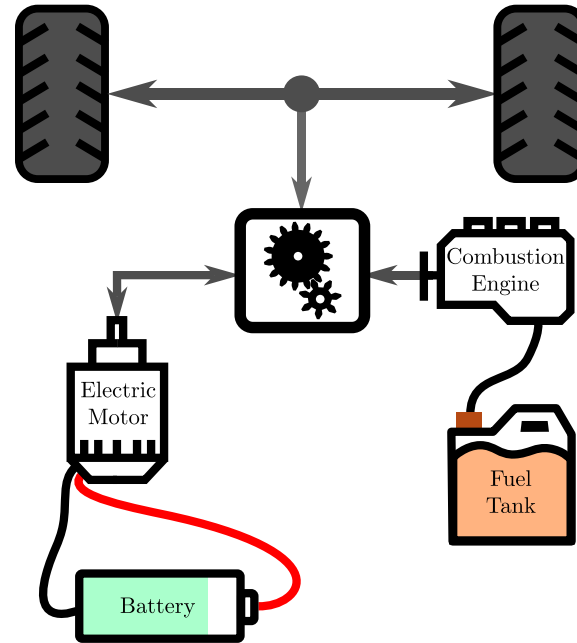
Dynamic programming is not  
fast enough on CPUs



Determine optimal split  
between electric motor and  
combustion engine

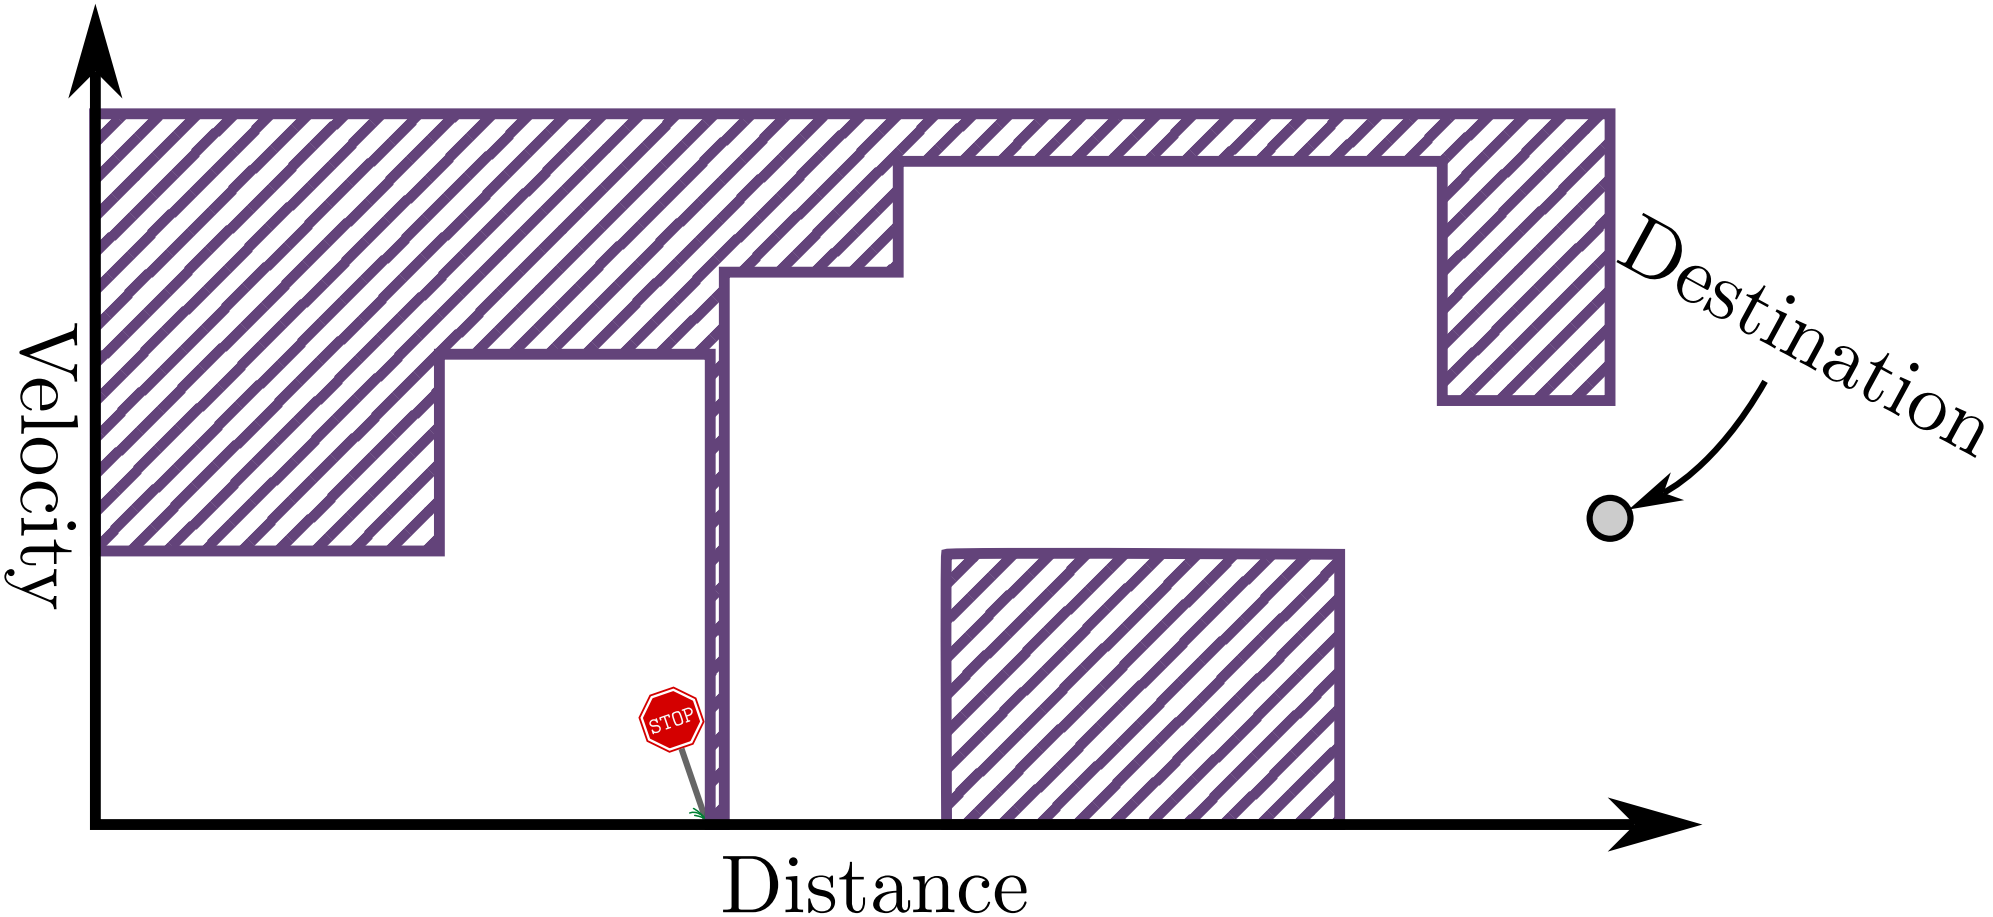
Dynamic programming is not  
fast enough on CPUs

**A scalable FPGA architecture  
that makes real-time use  
possible**

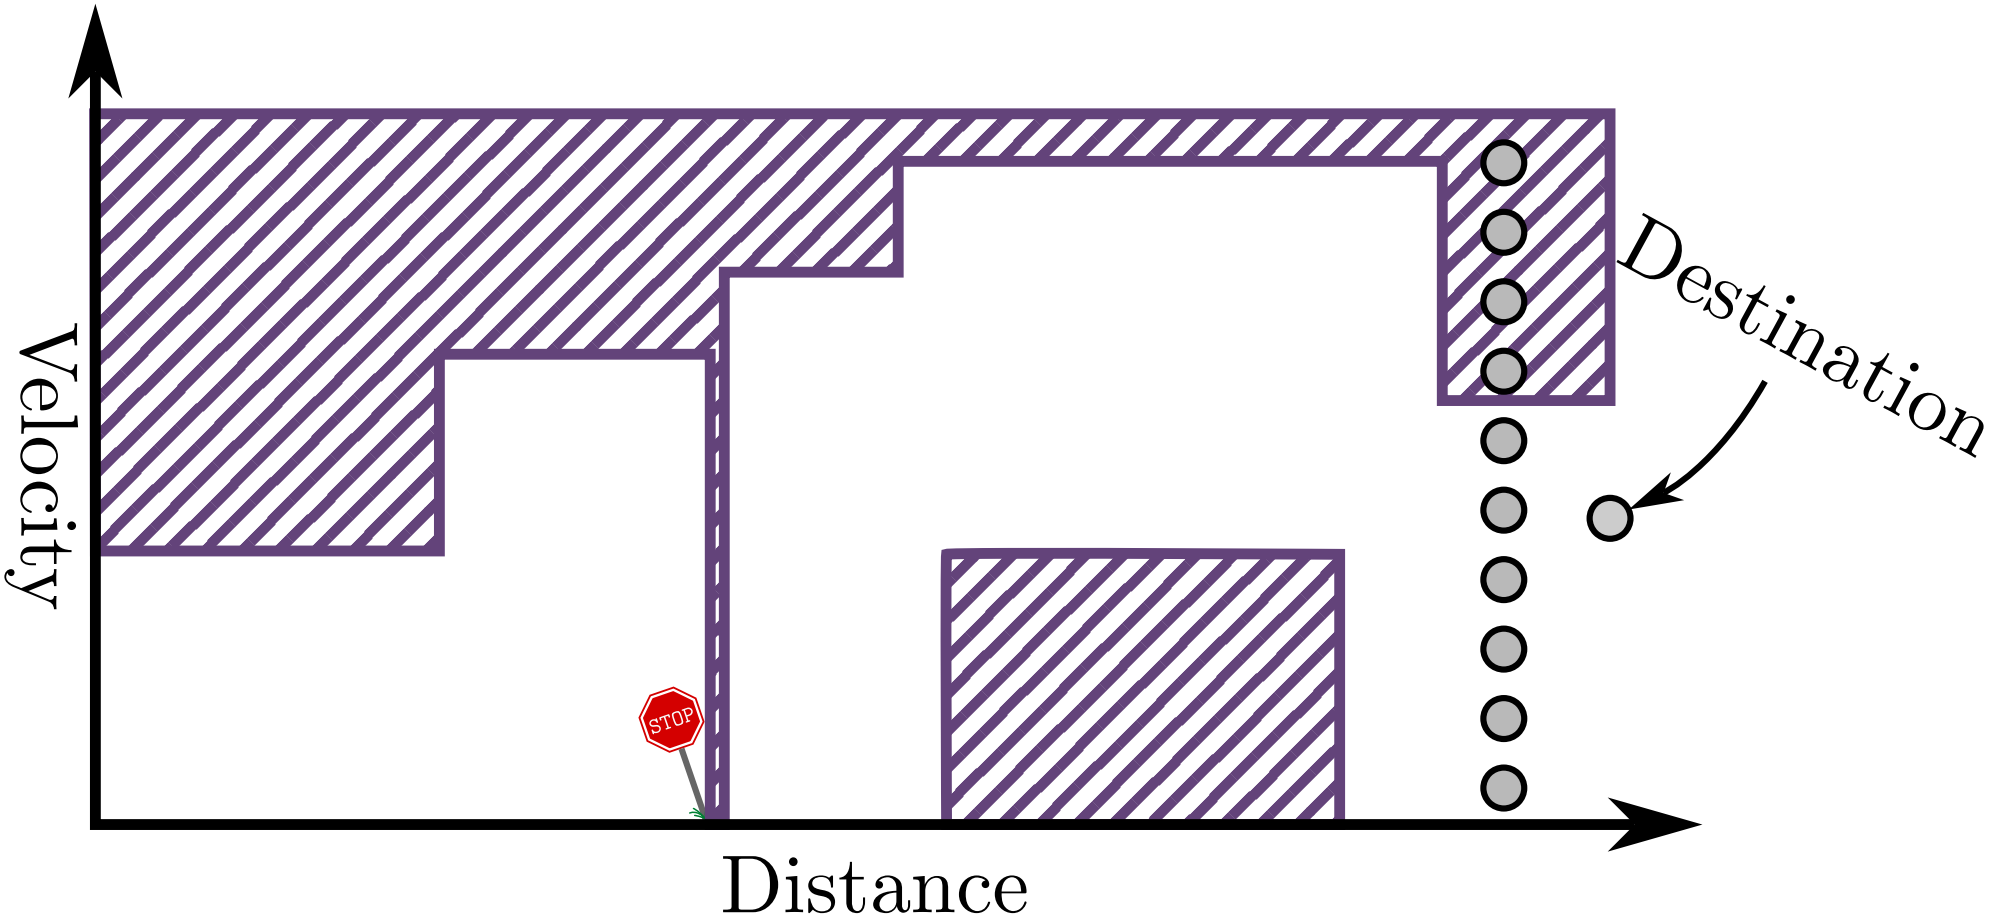




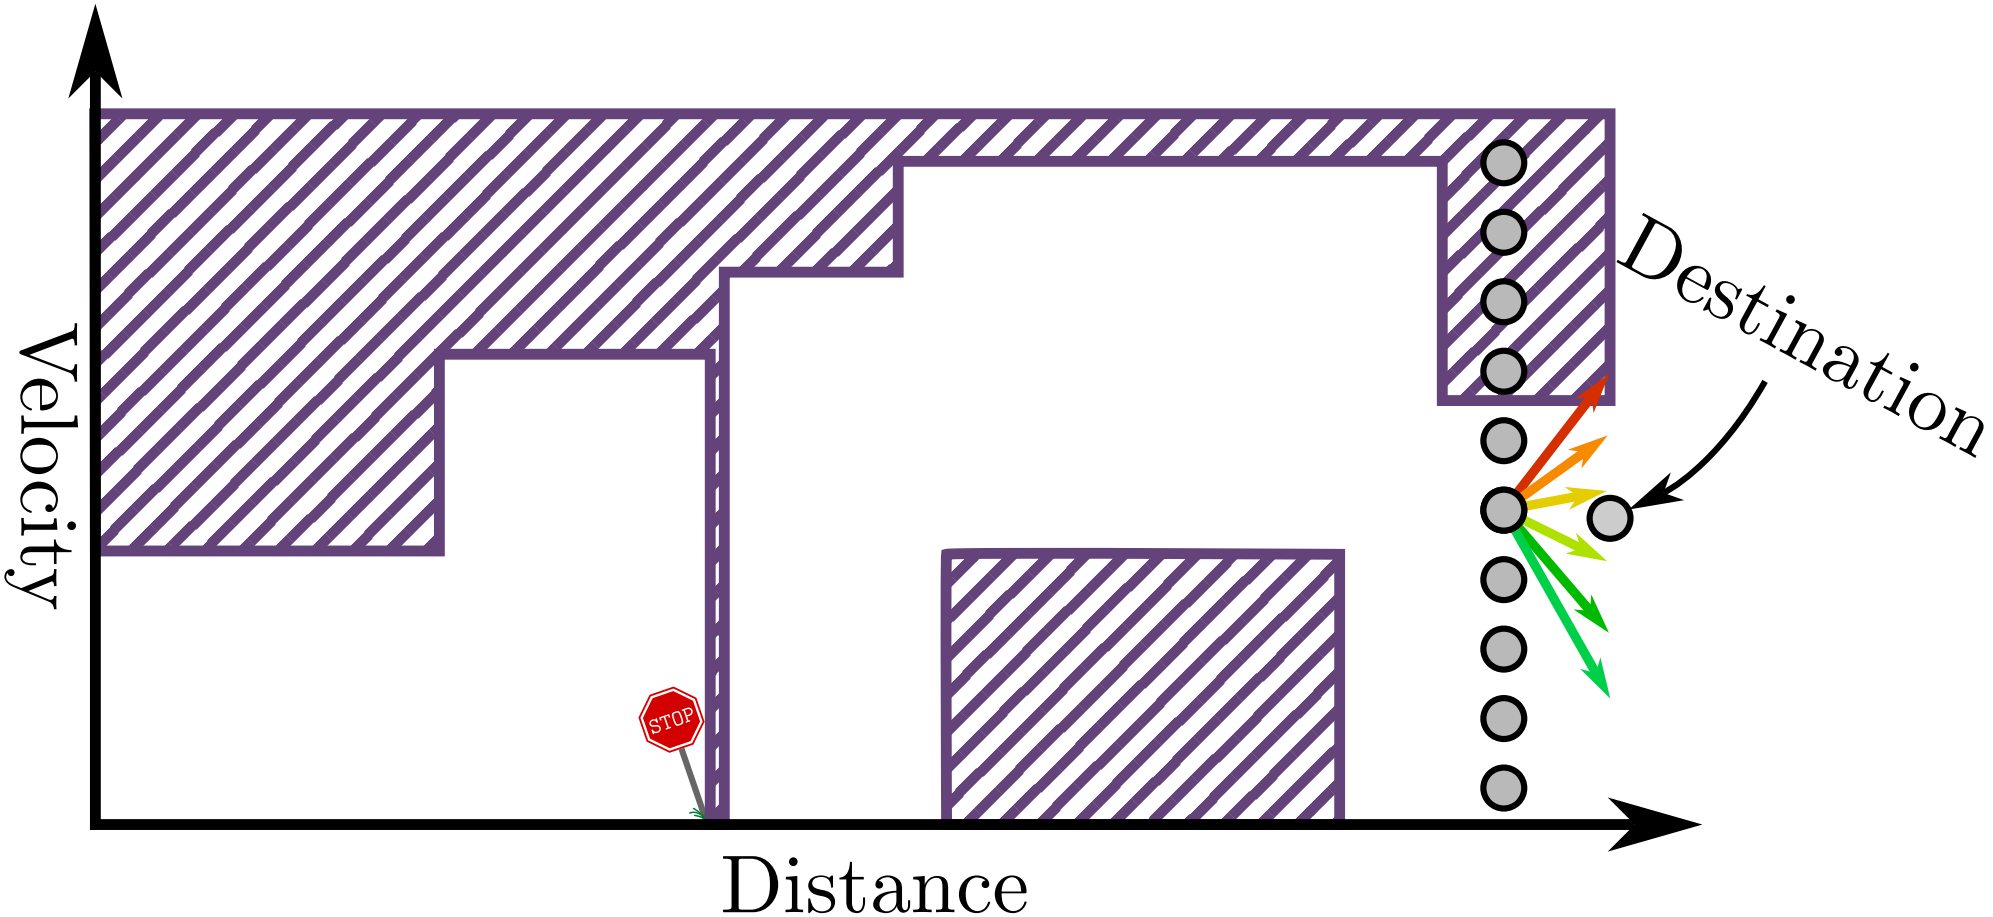
# The dynamic programming algorithm



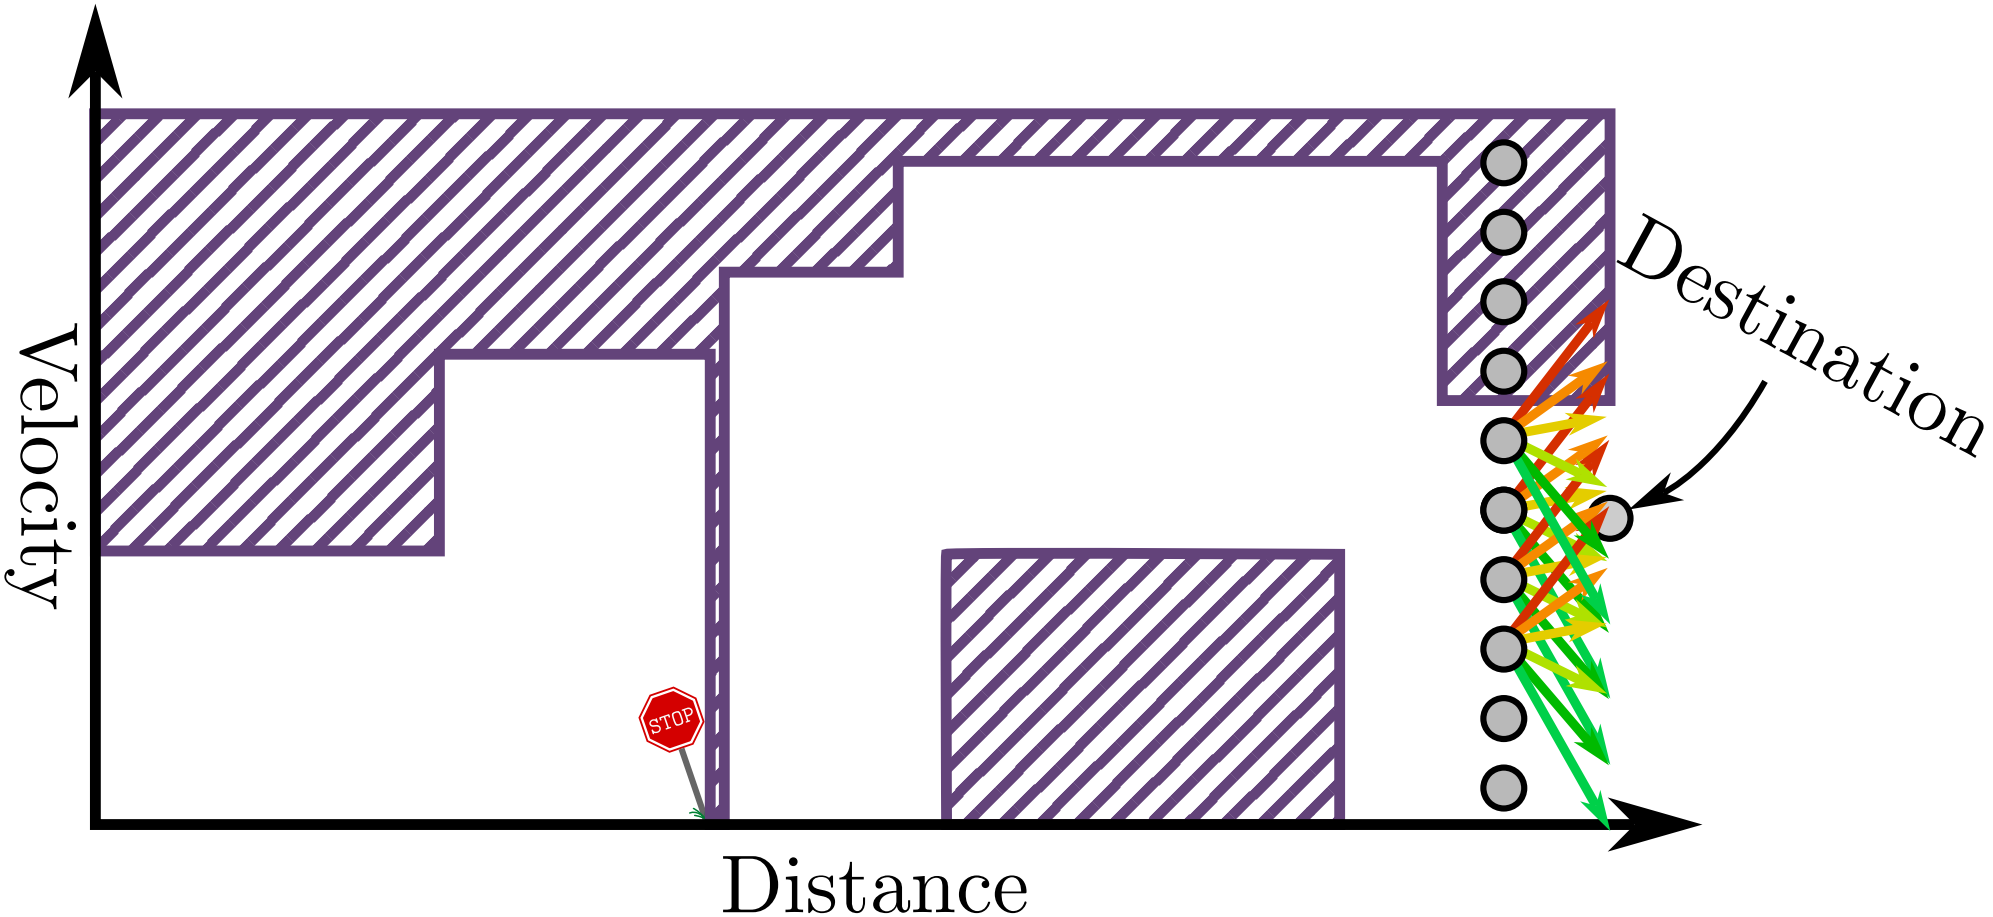
# The dynamic programming algorithm



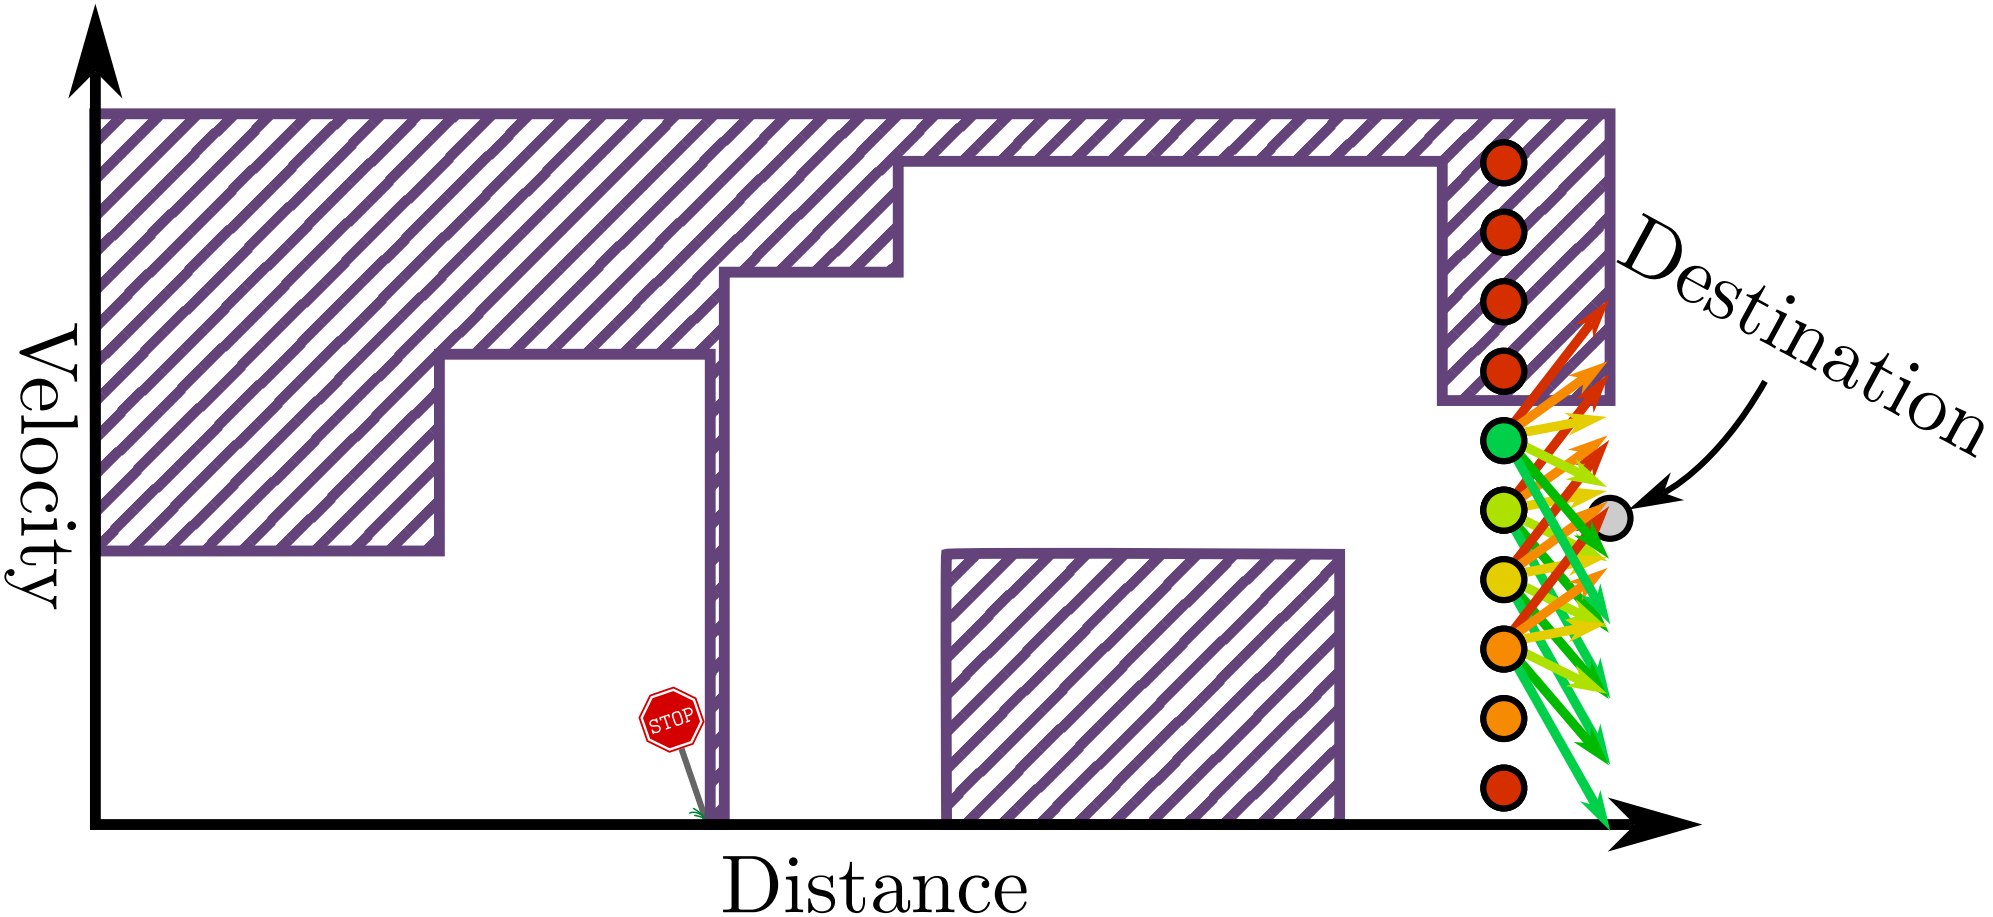
# The dynamic programming algorithm



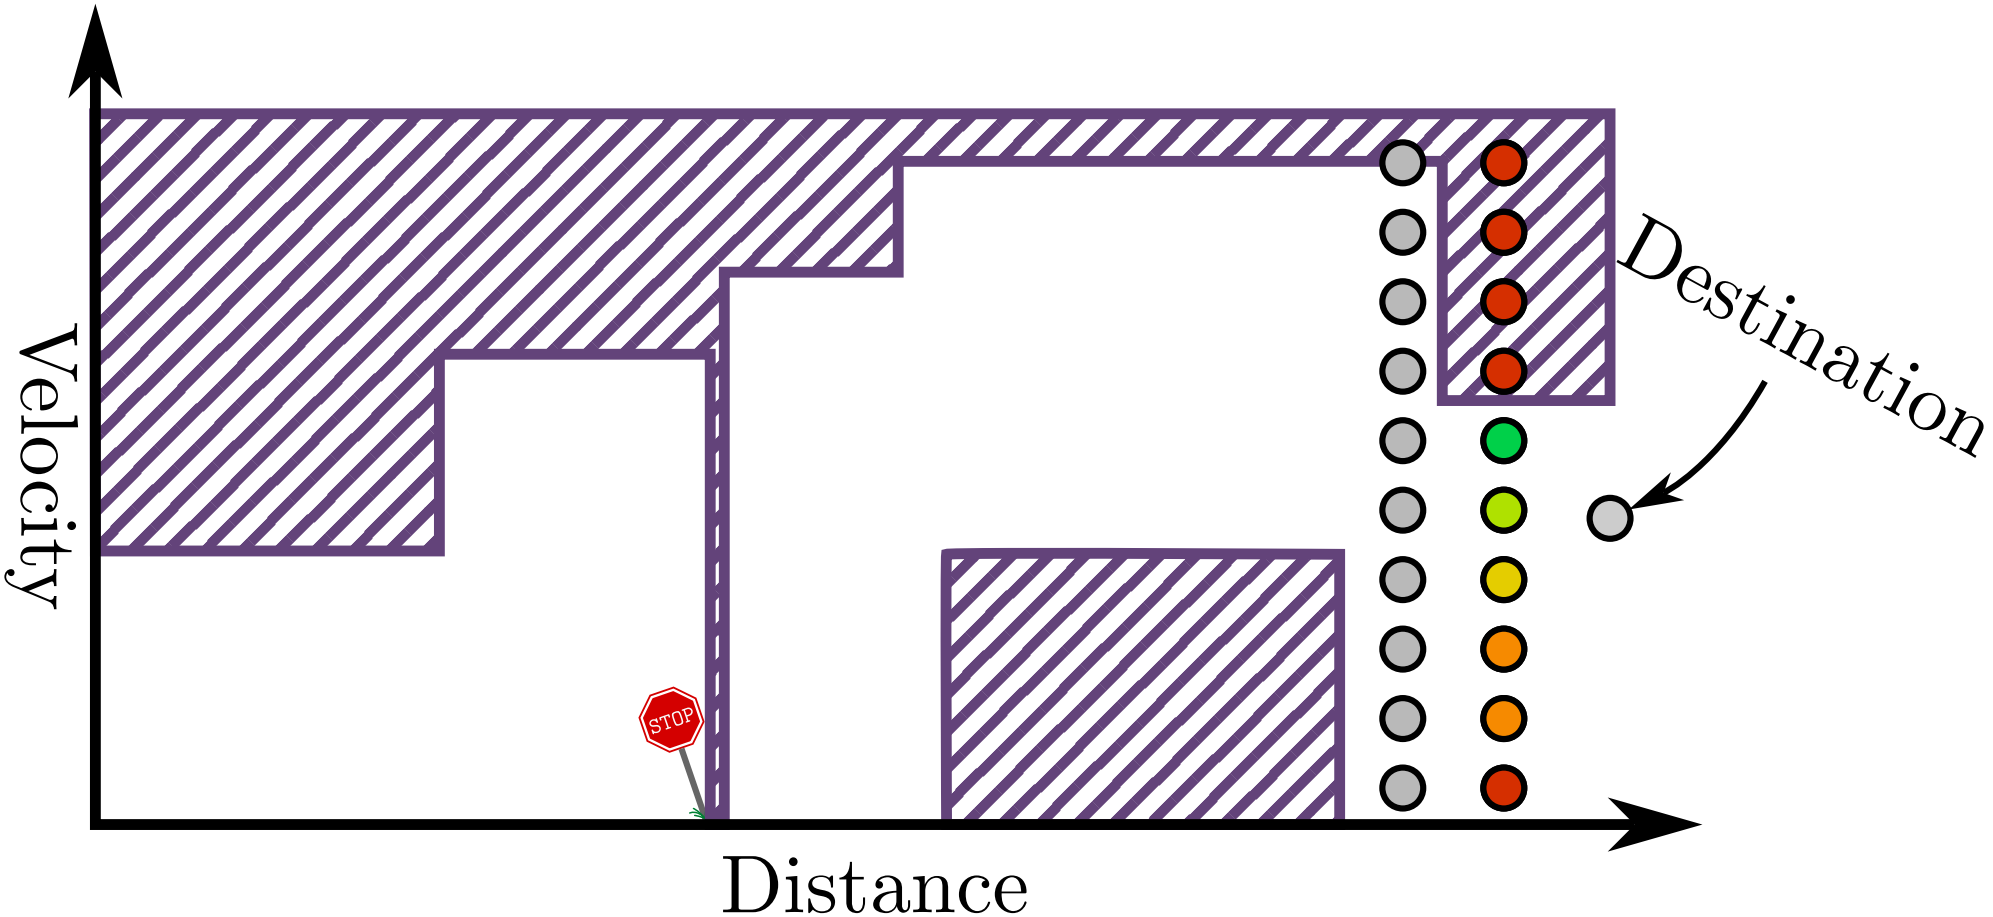
# The dynamic programming algorithm



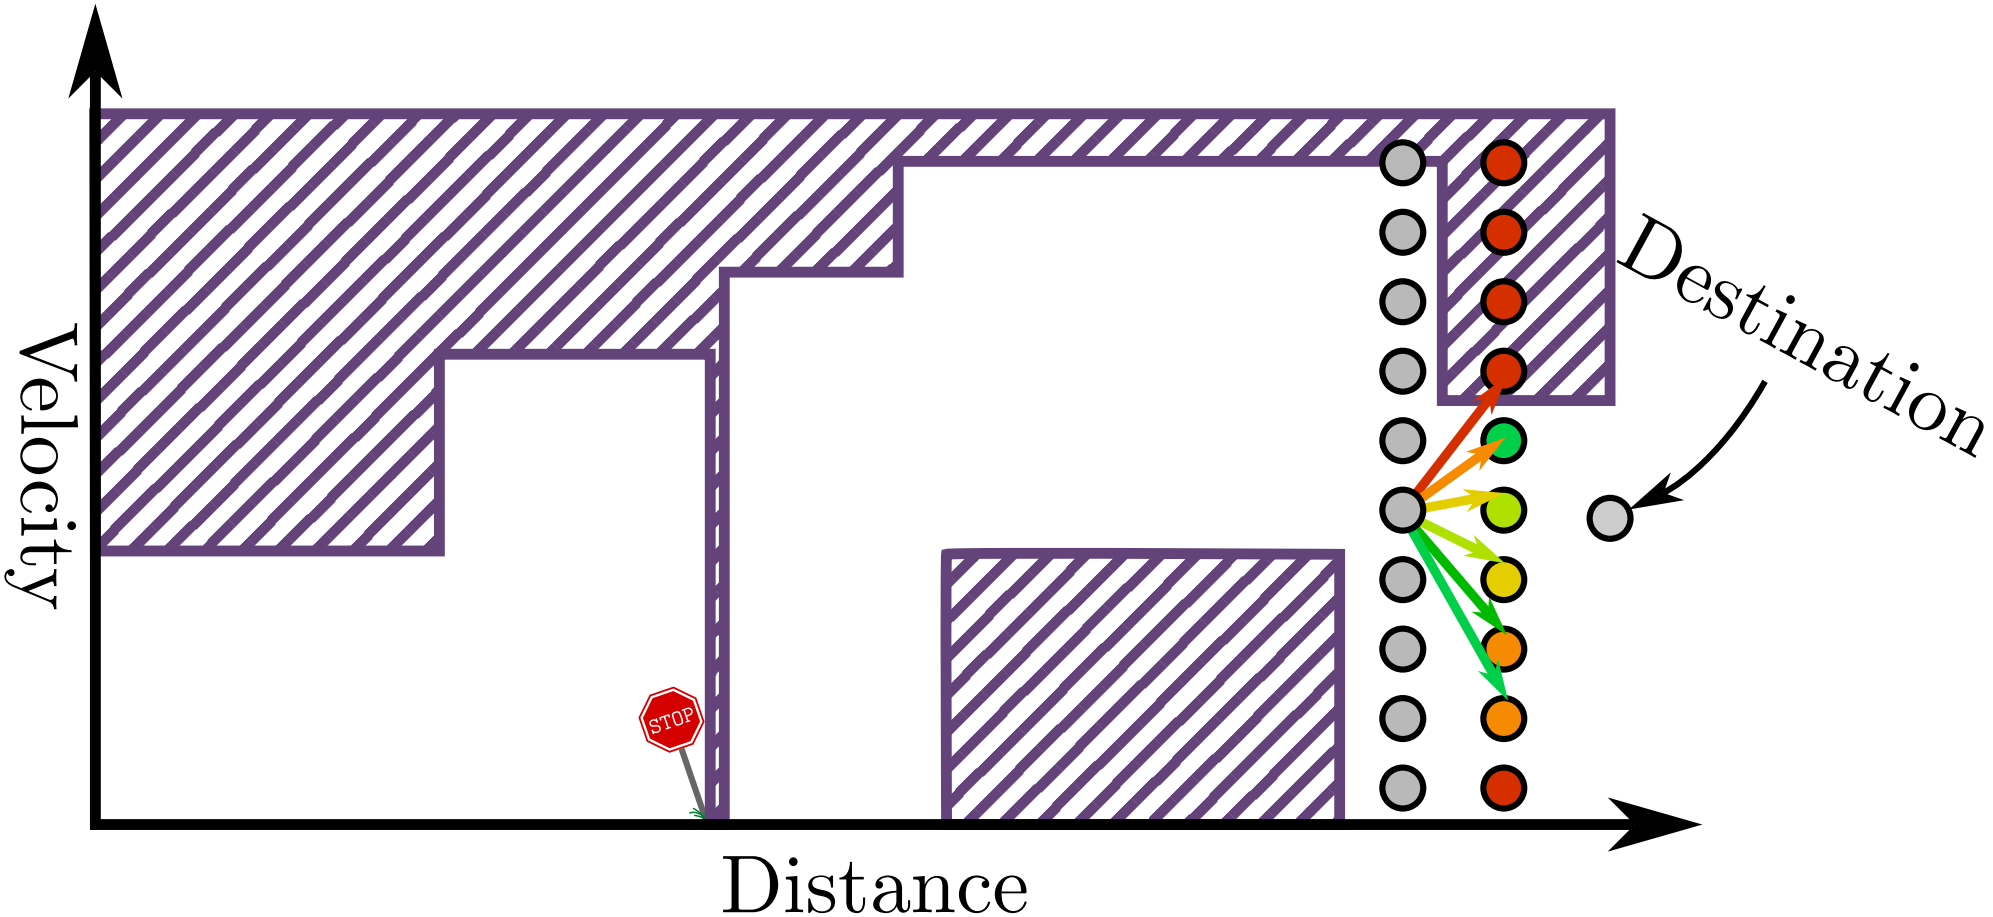
# The dynamic programming algorithm



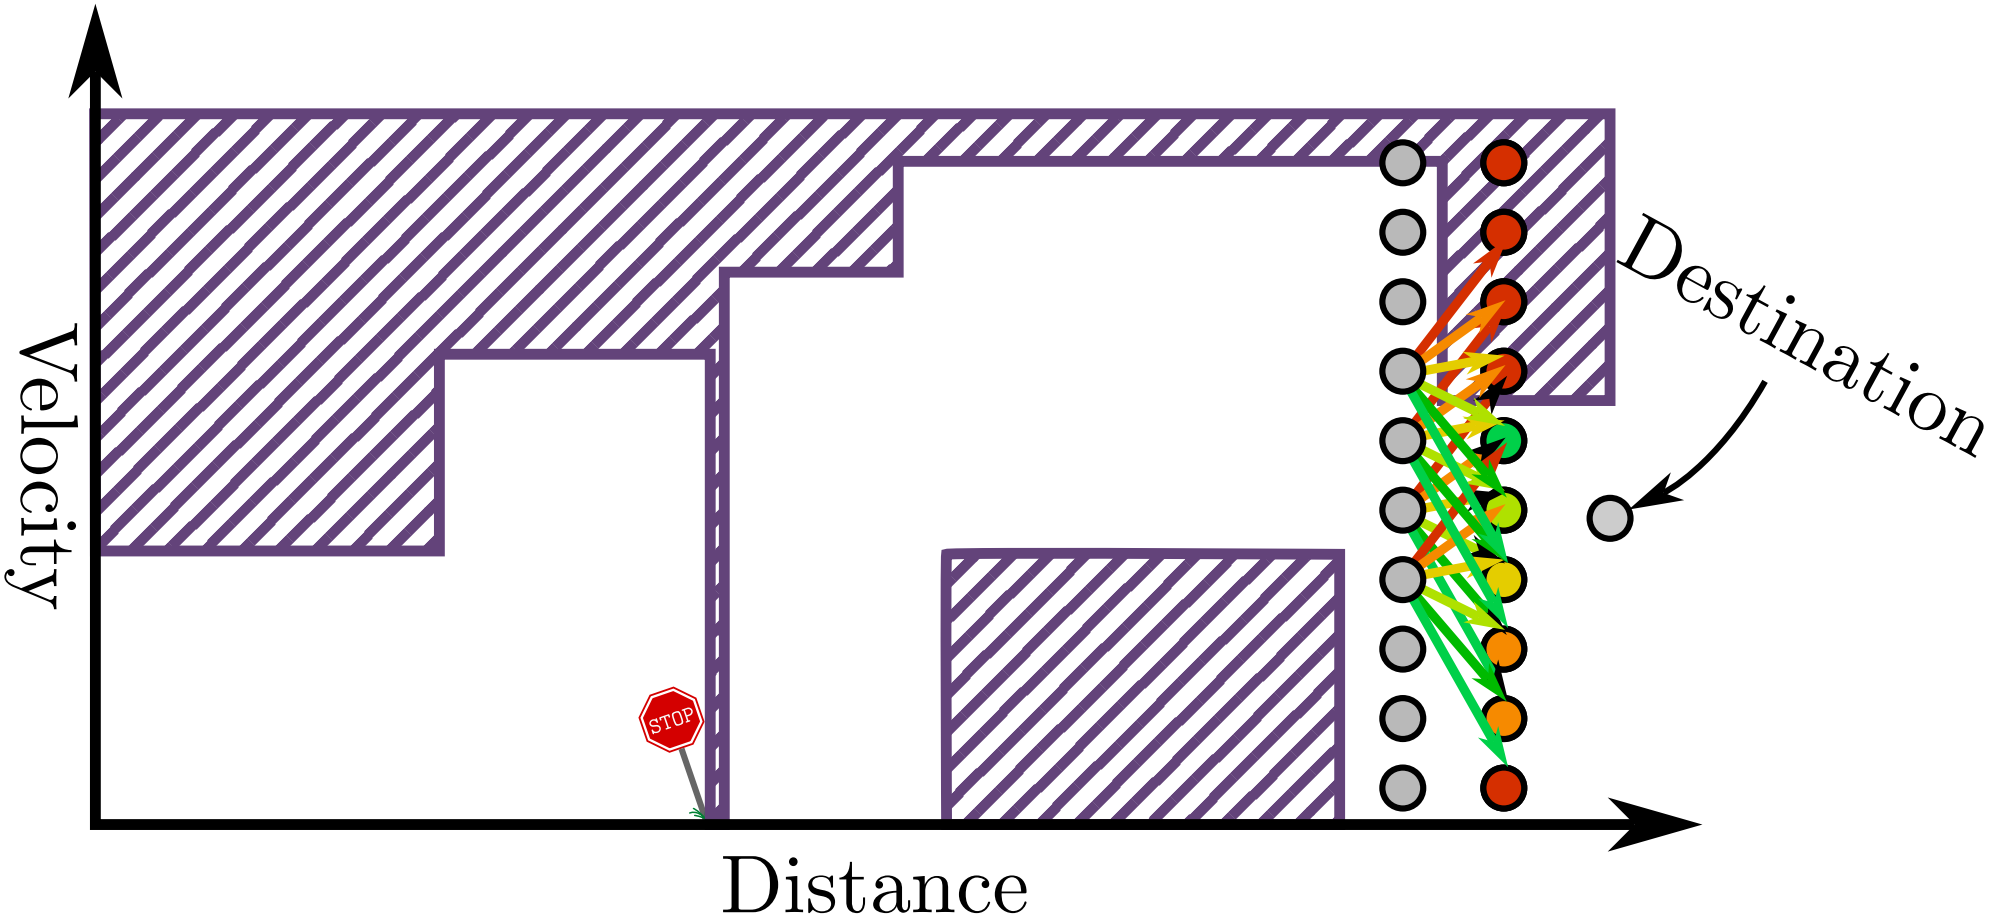
# The dynamic programming algorithm



# The dynamic programming algorithm

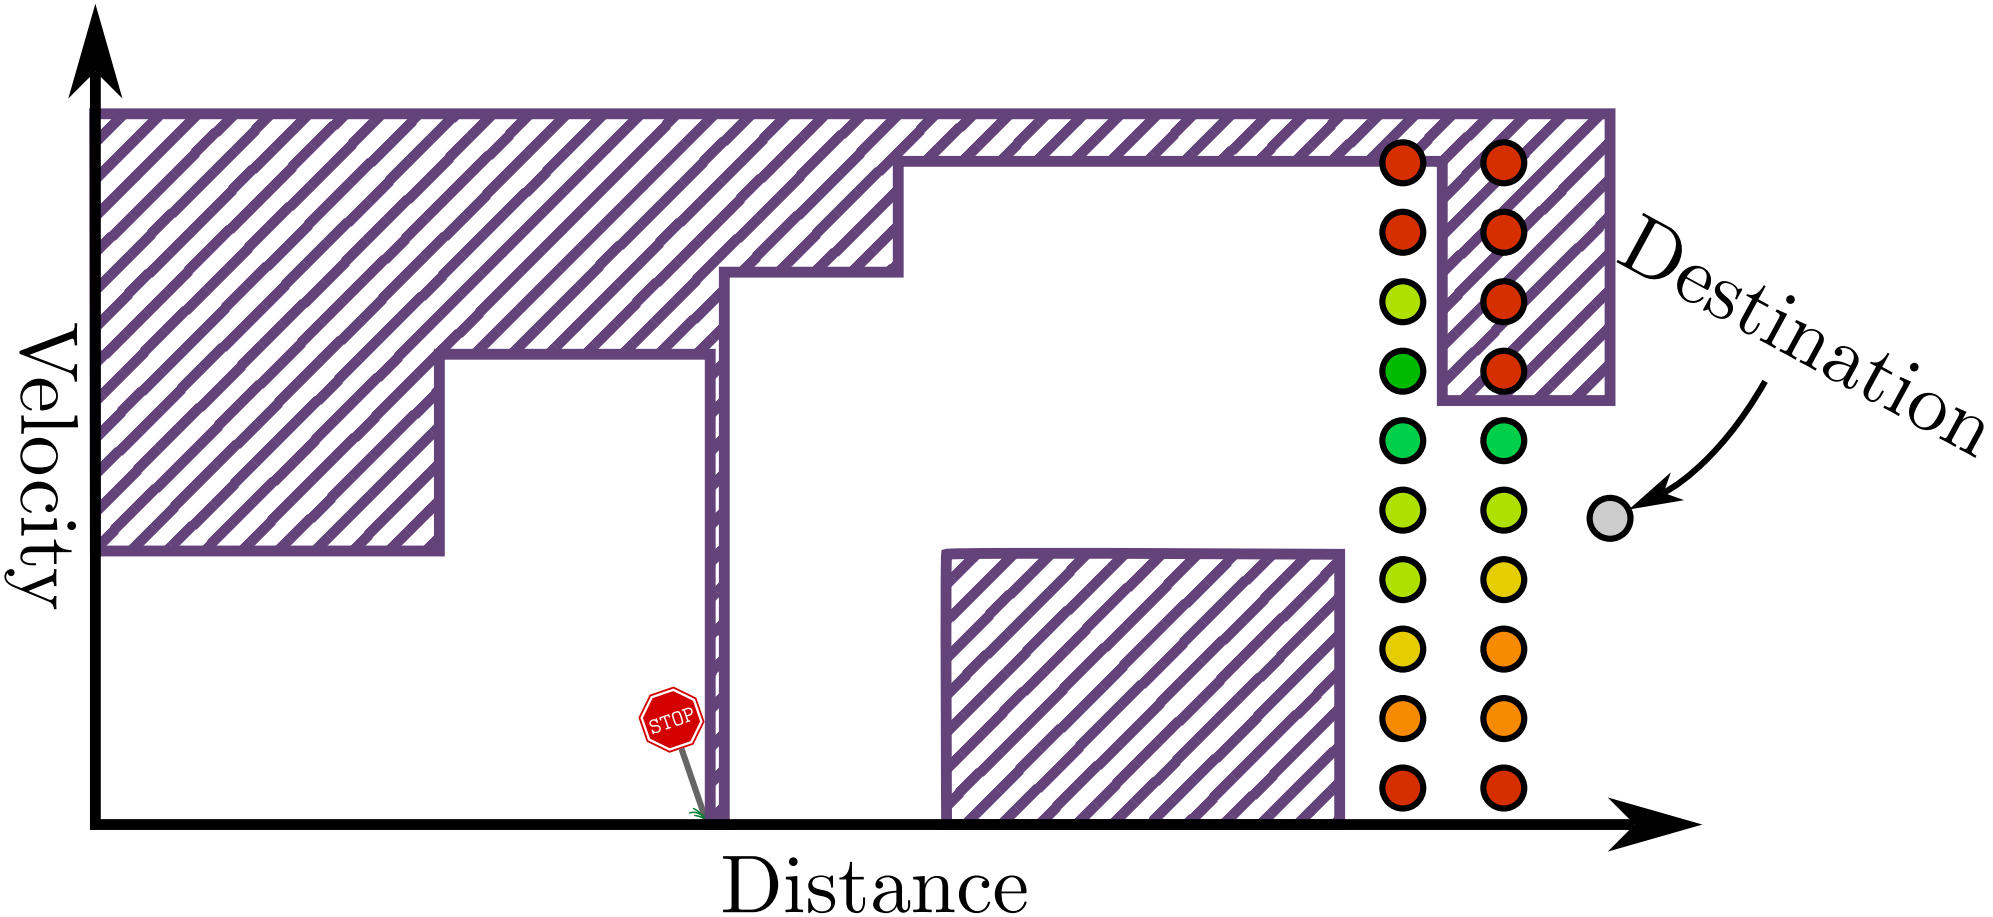


# The dynamic programming algorithm

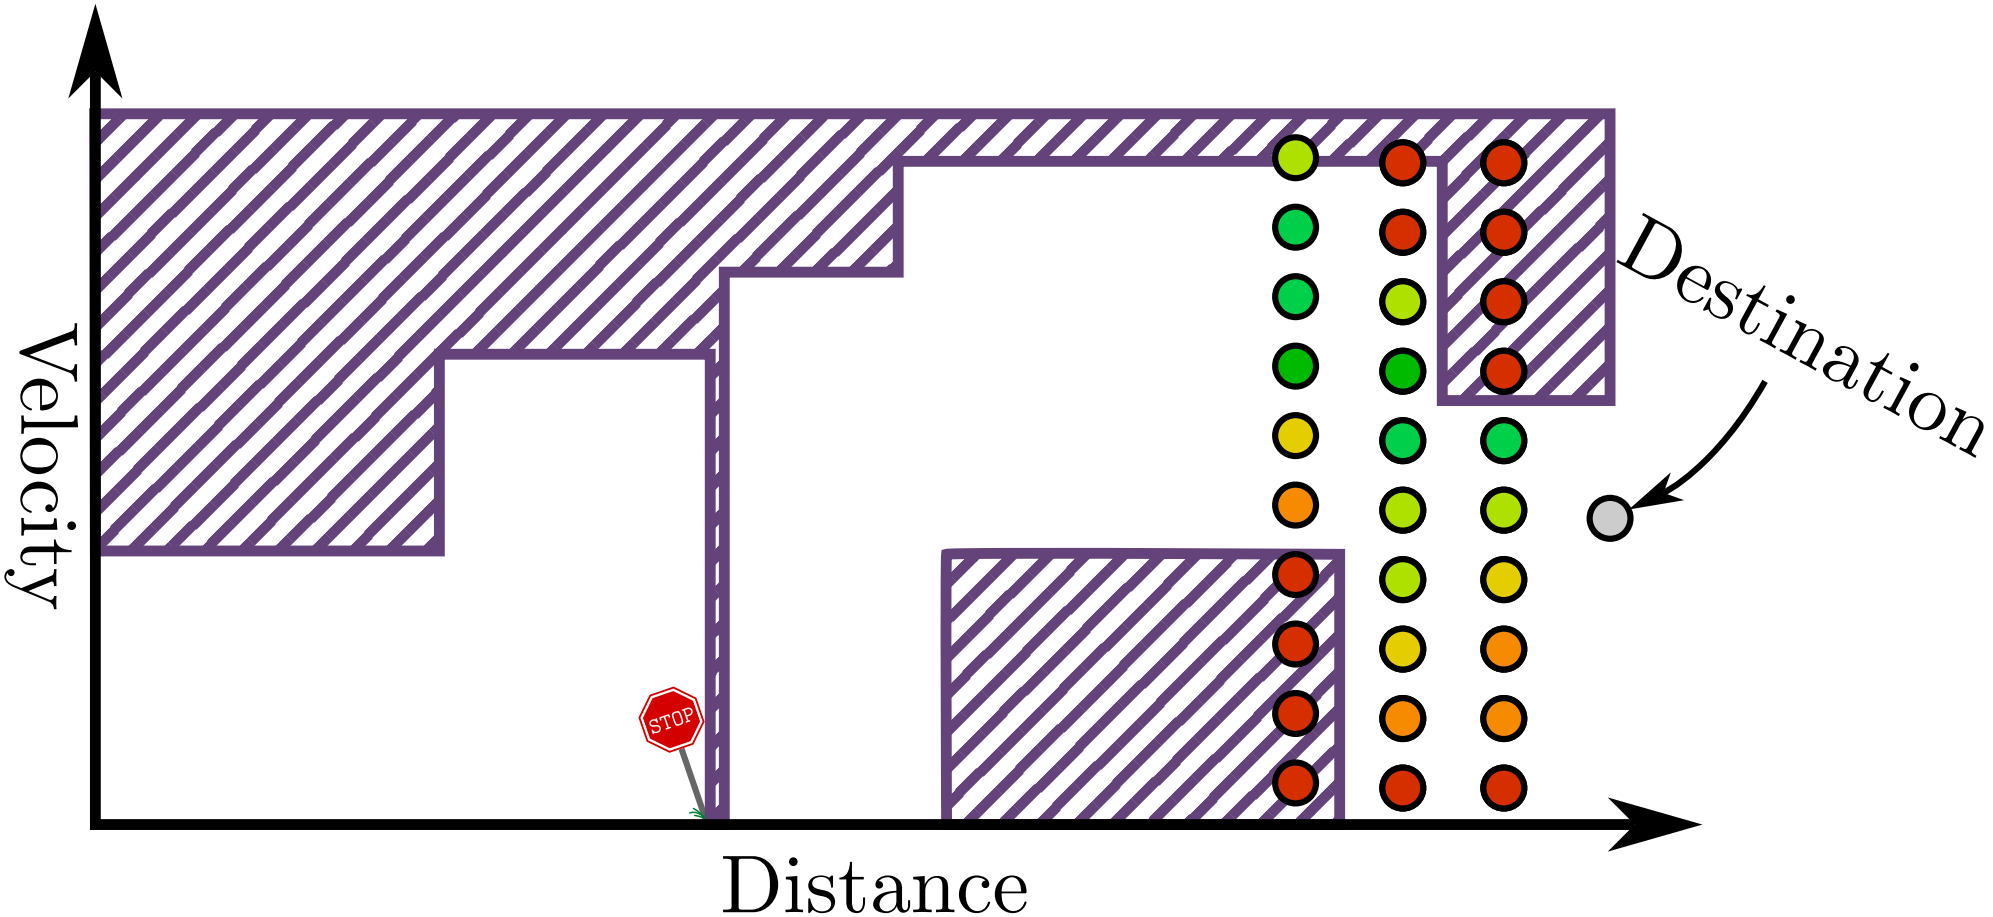




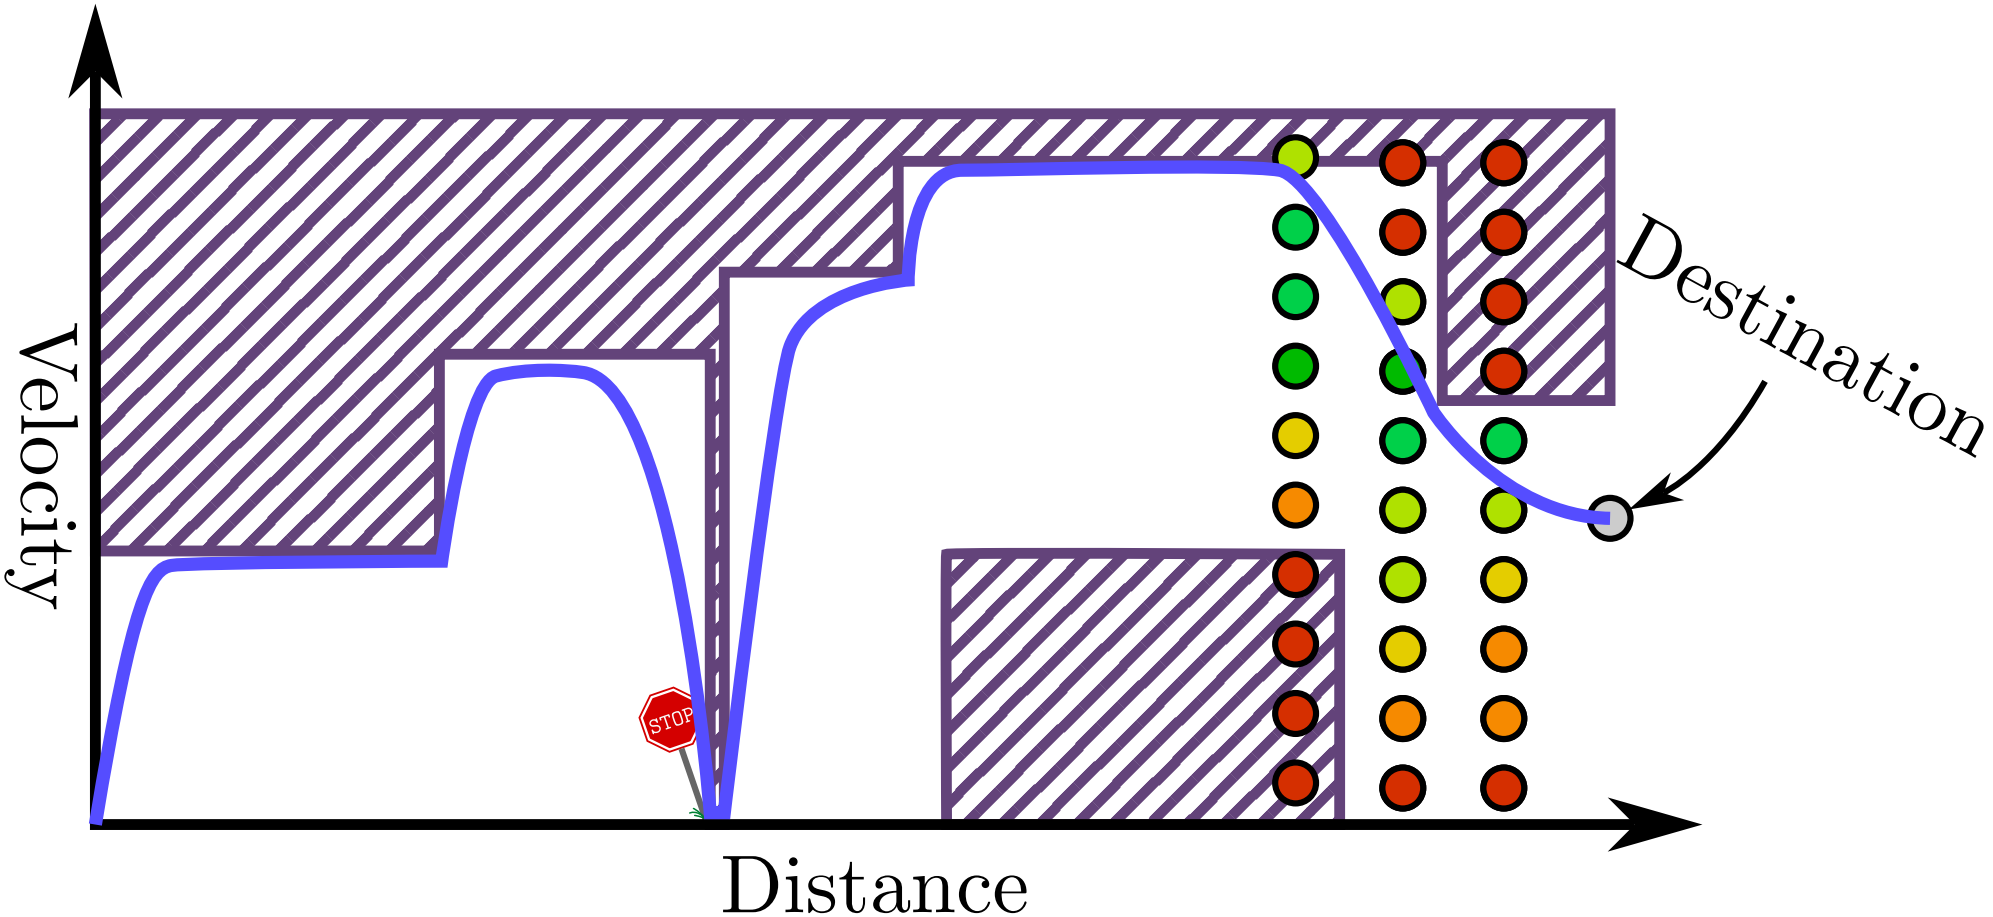
# The dynamic programming algorithm



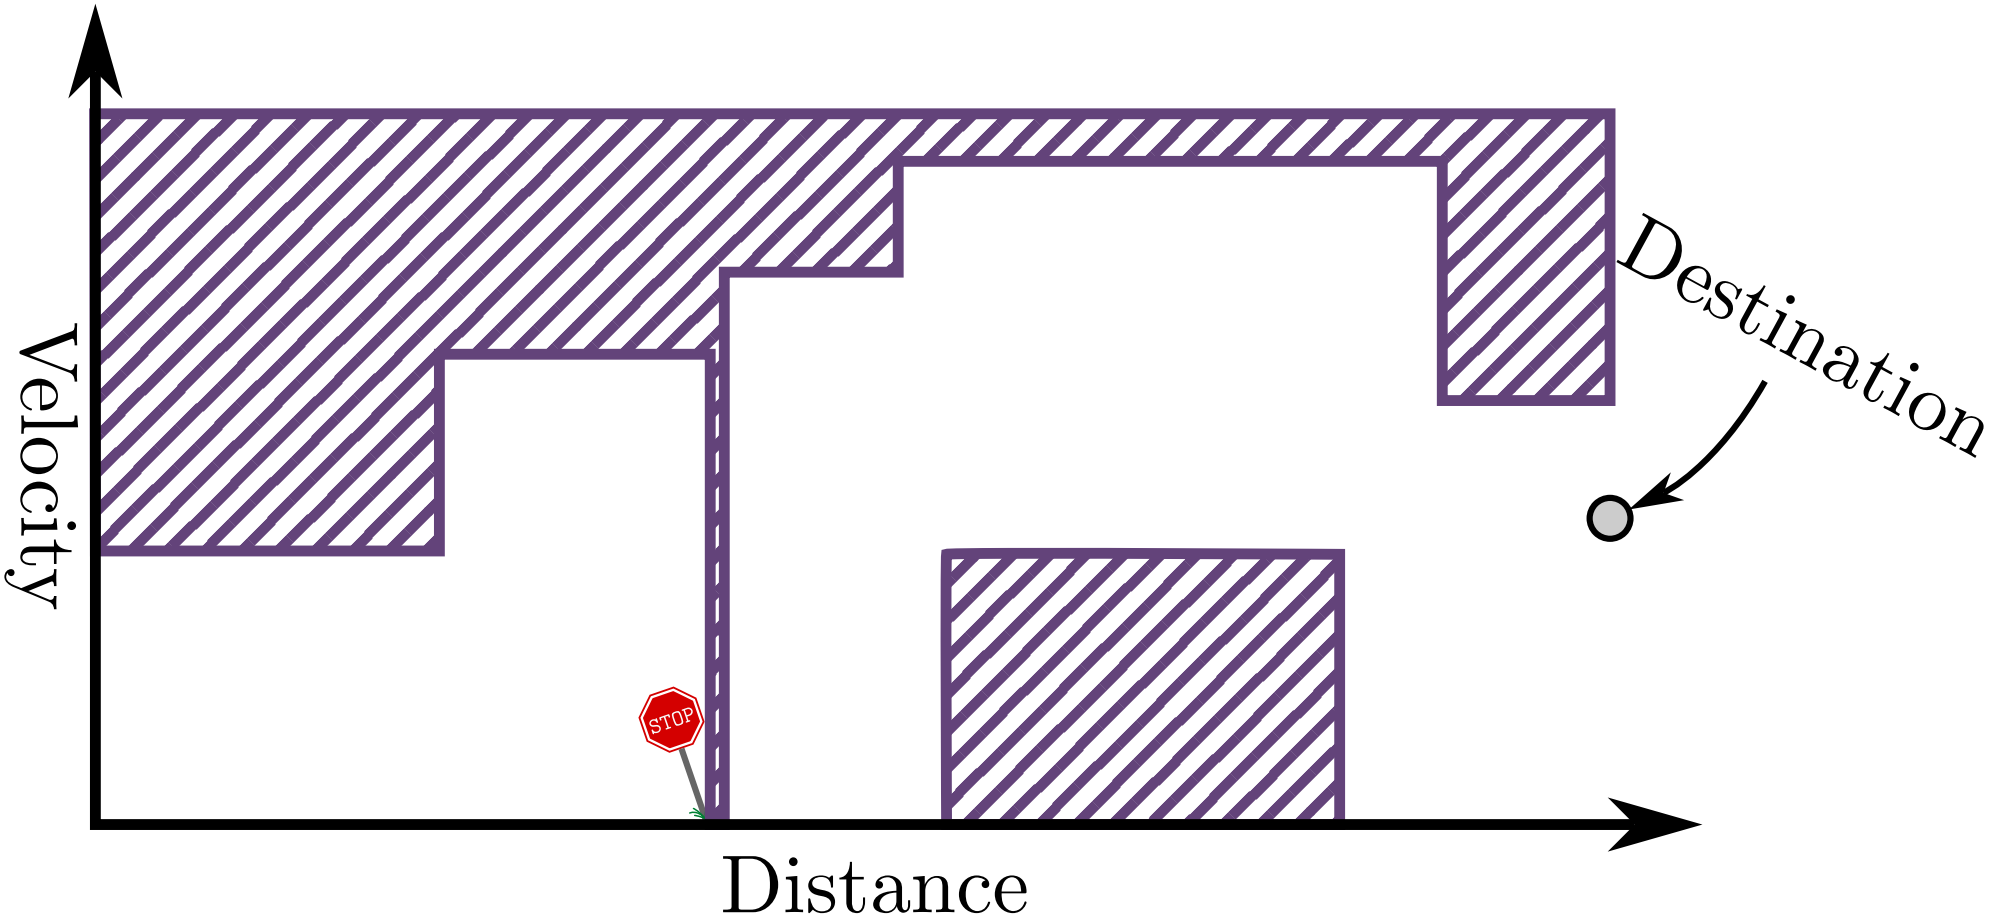
# The dynamic programming algorithm



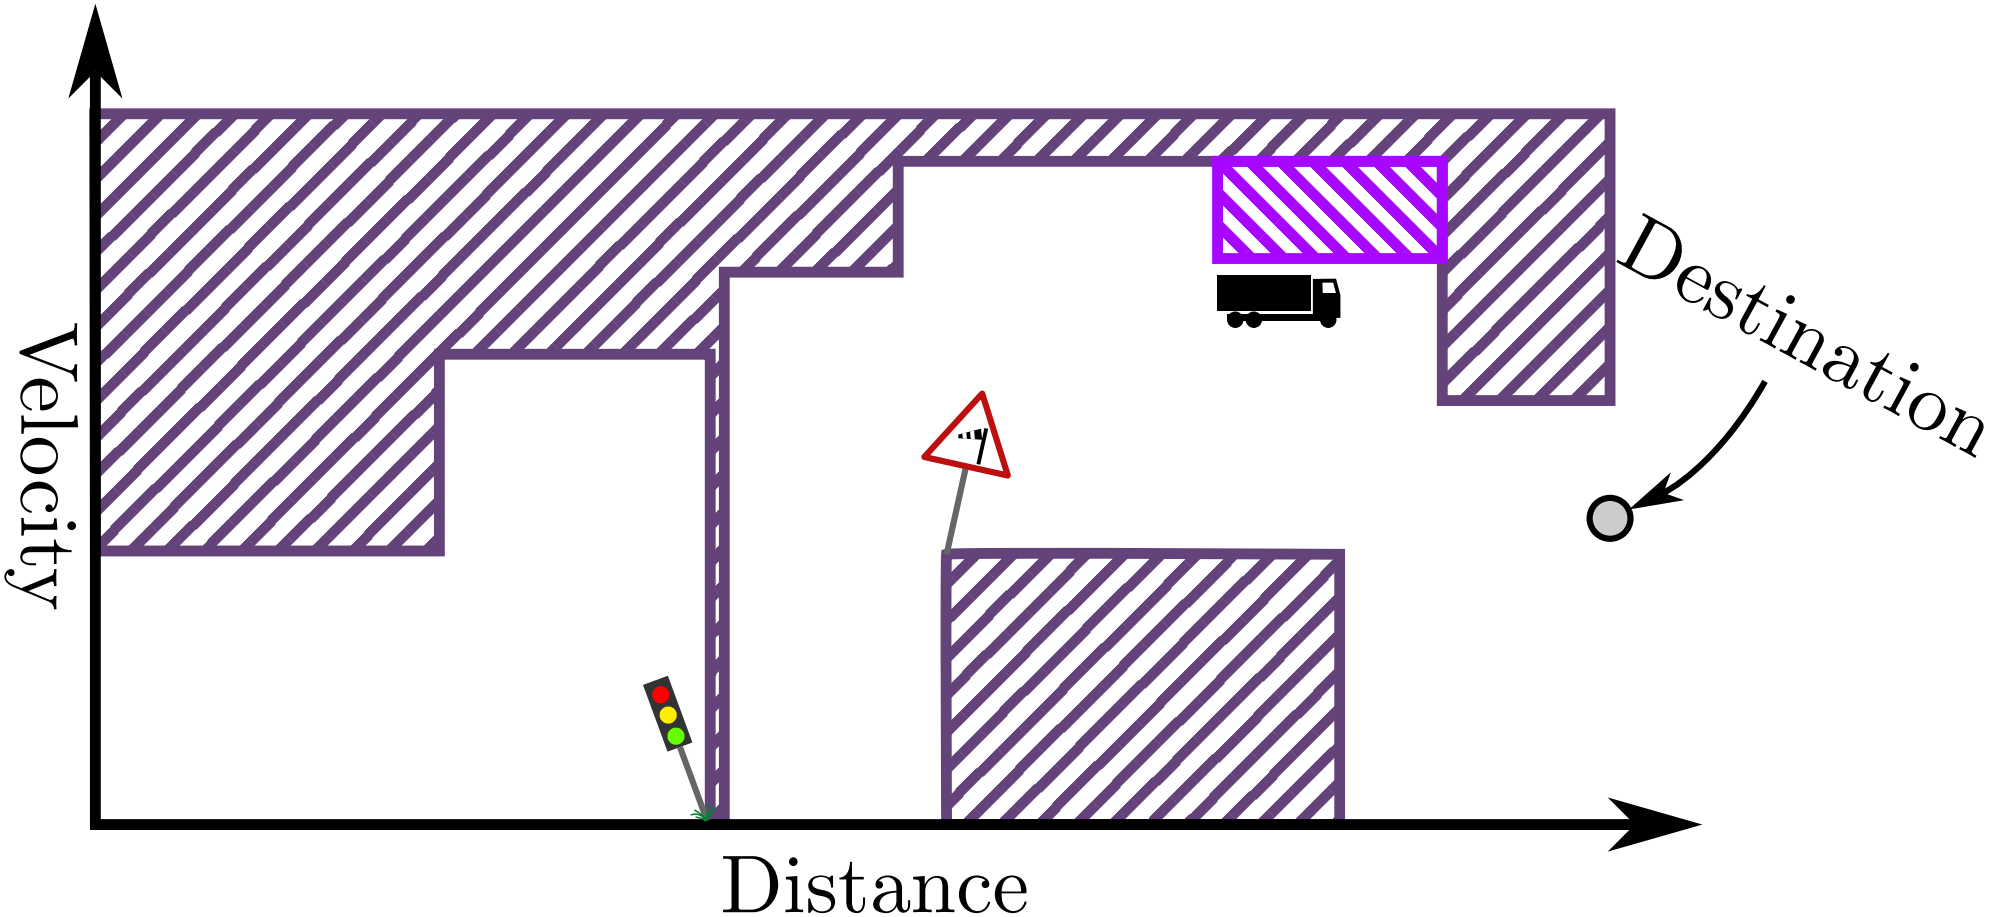
# The dynamic programming algorithm



# The dynamic programming algorithm



# The dynamic programming algorithm



# Compute Requirements

- 5.12 km Search horizon with 10 m steps

# Compute Requirements

- 5.12 km Search horizon with 10 m steps

## States

- 30 velocity steps
- 30 State of Charge (SOC) steps

# Compute Requirements

- 5.12 km Search horizon with 10 m steps

## States

- 30 velocity steps
- 30 State of Charge (SOC) steps

## Inputs

- 30 steps of electric torque
- 30 steps of combustion torque
- 6 gears



# Compute Requirements

The diagram illustrates the calculation of compute requirements. It features a mathematical expression with four main components highlighted in colored boxes: a green box containing '512', an orange box containing '30 x 30', a purple box containing '30 x 30 x 6', and a pink box containing '2.4 billion'. Colored arrows point from labels to these components: a green arrow labeled 'Distance Steps' points to '512'; an orange arrow labeled 'States' points to the '30 x 30' in the orange box; a purple arrow labeled 'Inputs' points to the '30 x 30 x 6' in the purple box; and a pink arrow labeled 'Model evaluations' points to '2.4 billion'. The entire expression is:  $512 \times 30 \times 30 \times 30 \times 30 \times 6 \approx 2.4 \text{ billion}$ .

# Compute Requirements

Distance Steps

Inputs

States

Model evaluations

$$512 \times 30 \times 30 \times 30 \times 30 \times 6 \approx 2.4 \text{ billion}$$

≈ 2 seconds for real time

# Compute Requirements

Distance Steps

Inputs

States

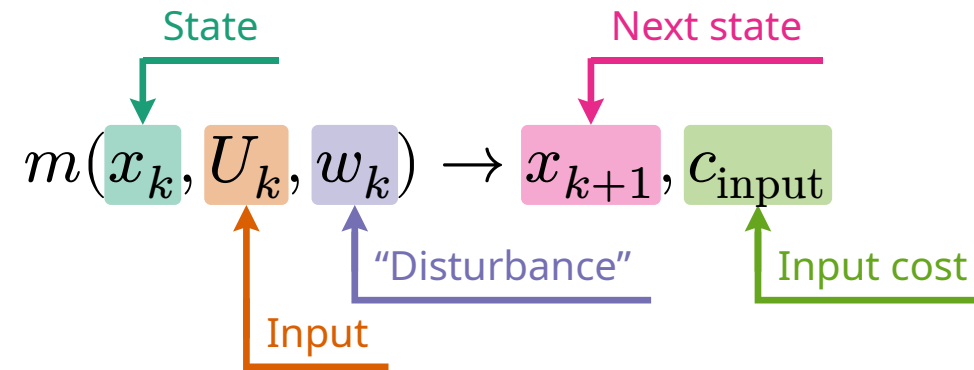
Model evaluations

$$512 \times 30 \times 30 \times 30 \times 30 \times 6 \approx 2.4 \text{ billion}$$

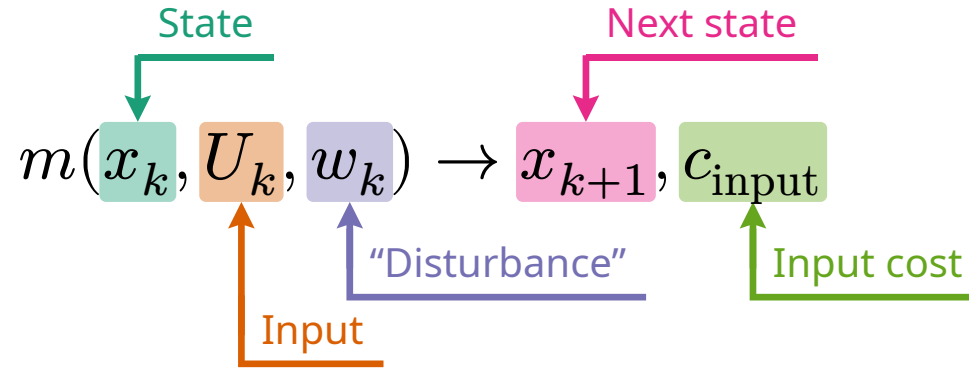
≈ 2 seconds for real time

> 1 model execution every clock cycle

# Vehicle Model



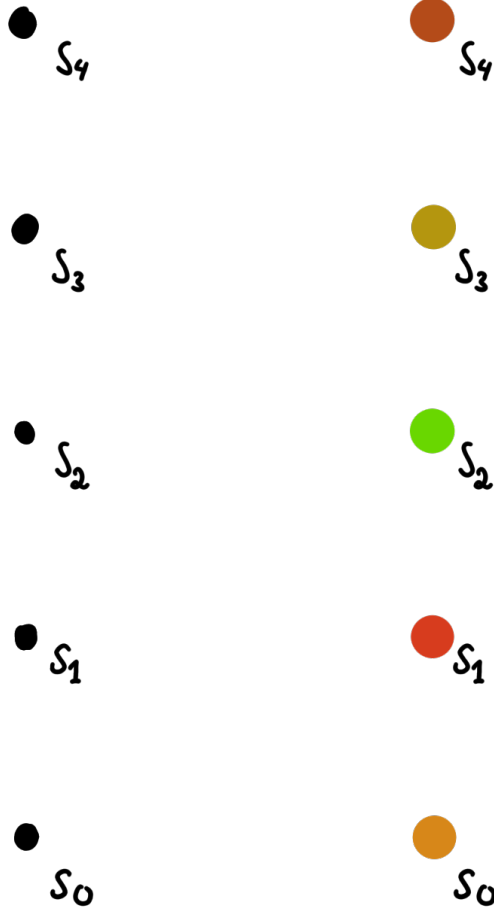
# Vehicle Model



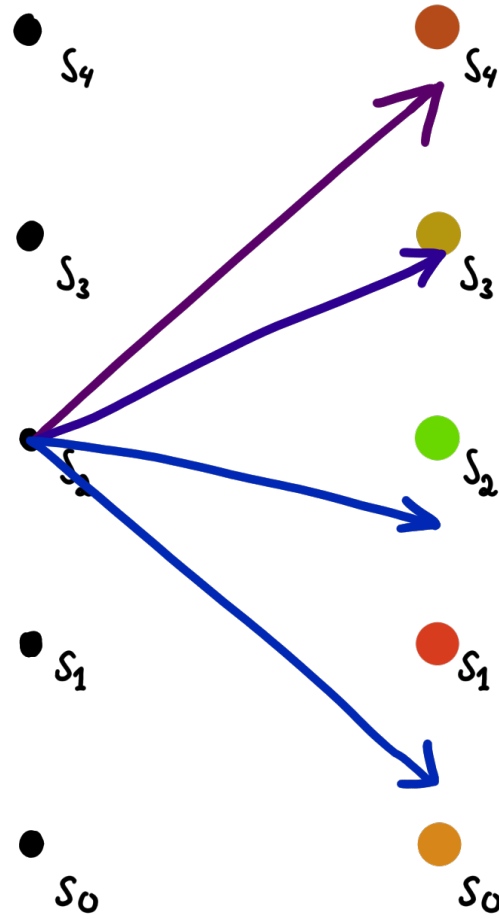
C++ model converted with HLS

Pipelined with Initiation Interval = 1

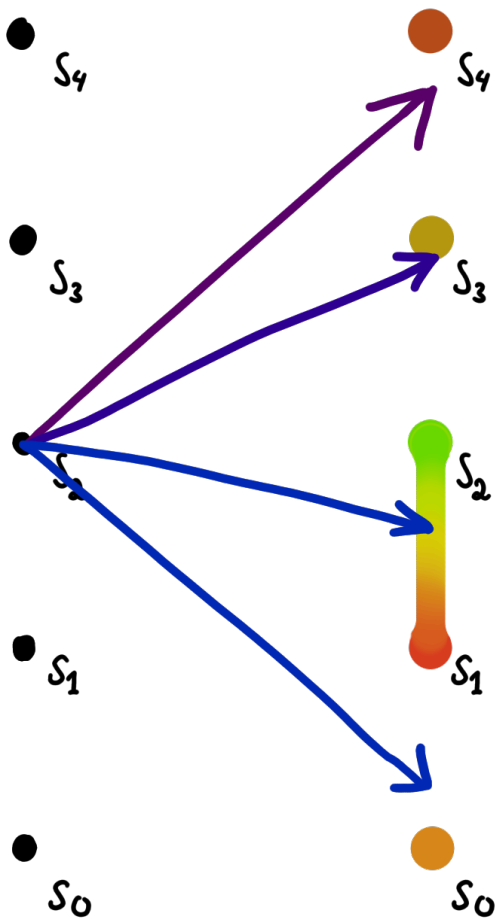
# Interpolation



# Interpolation

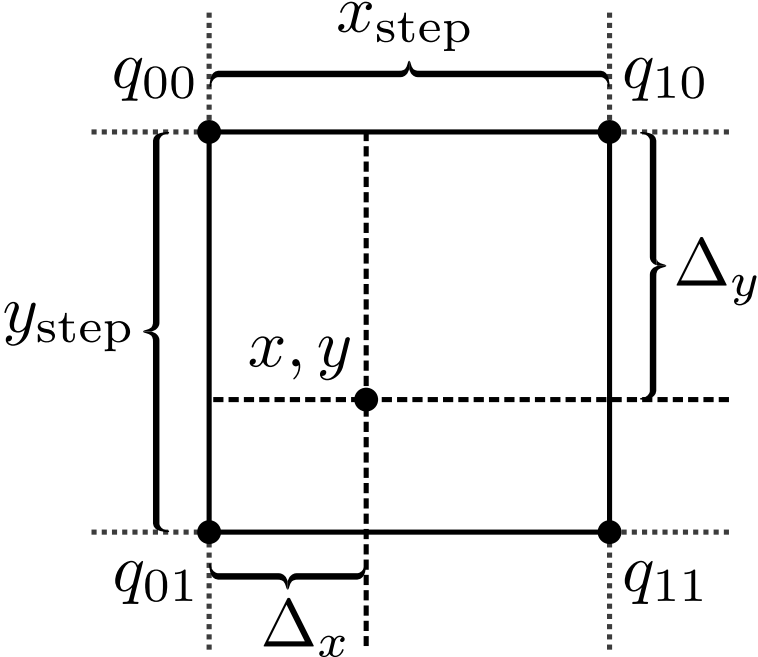


# Interpolation

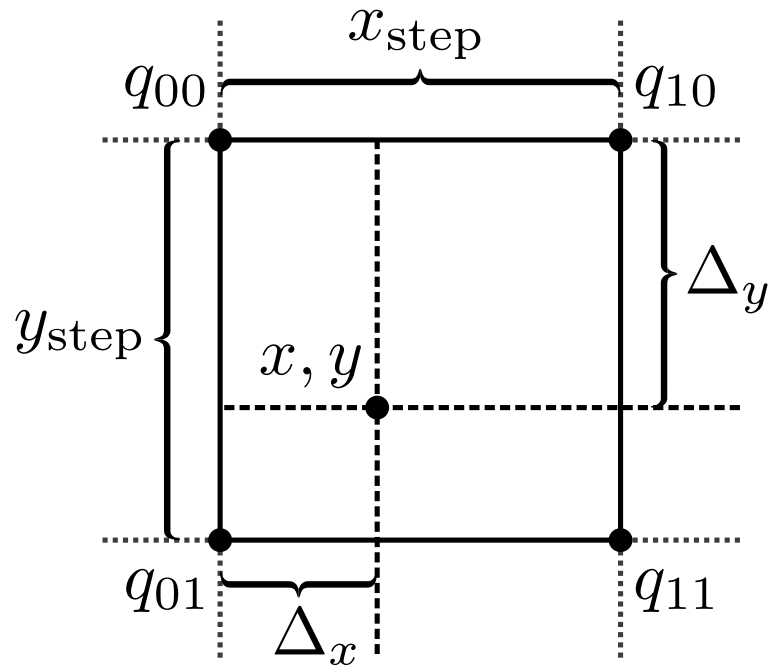




# Interpolation

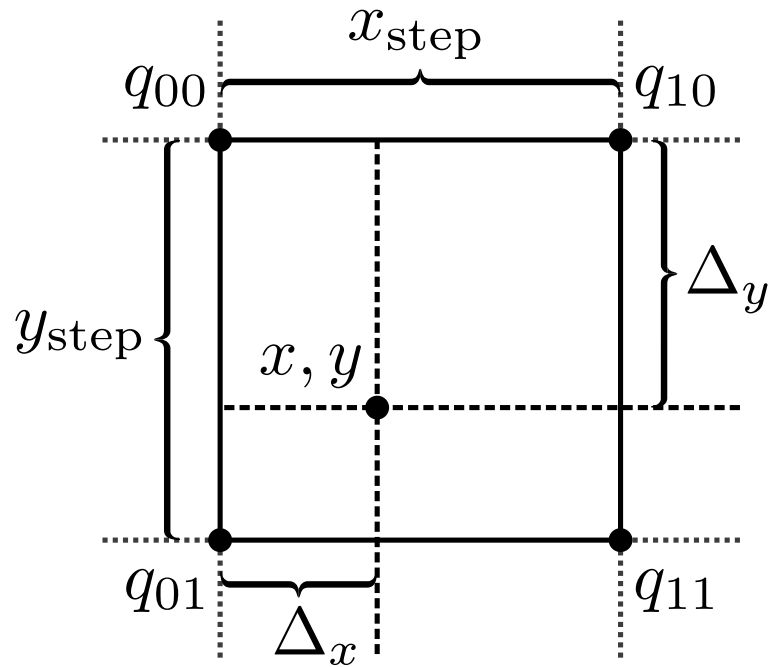


# Interpolation



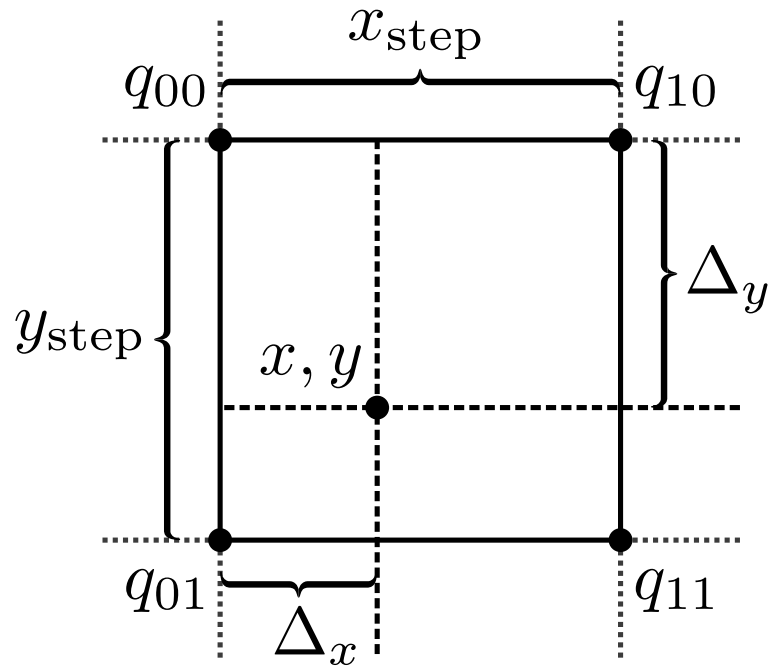
- 2D linear interpolation

# Interpolation



- 2D linear interpolation
- Requires 4 memory accesses per value

# Interpolation



- 2D linear interpolation
- Requires 4 memory accesses per value
- Simultaneous writeback is required

# Memory Read Partitioning

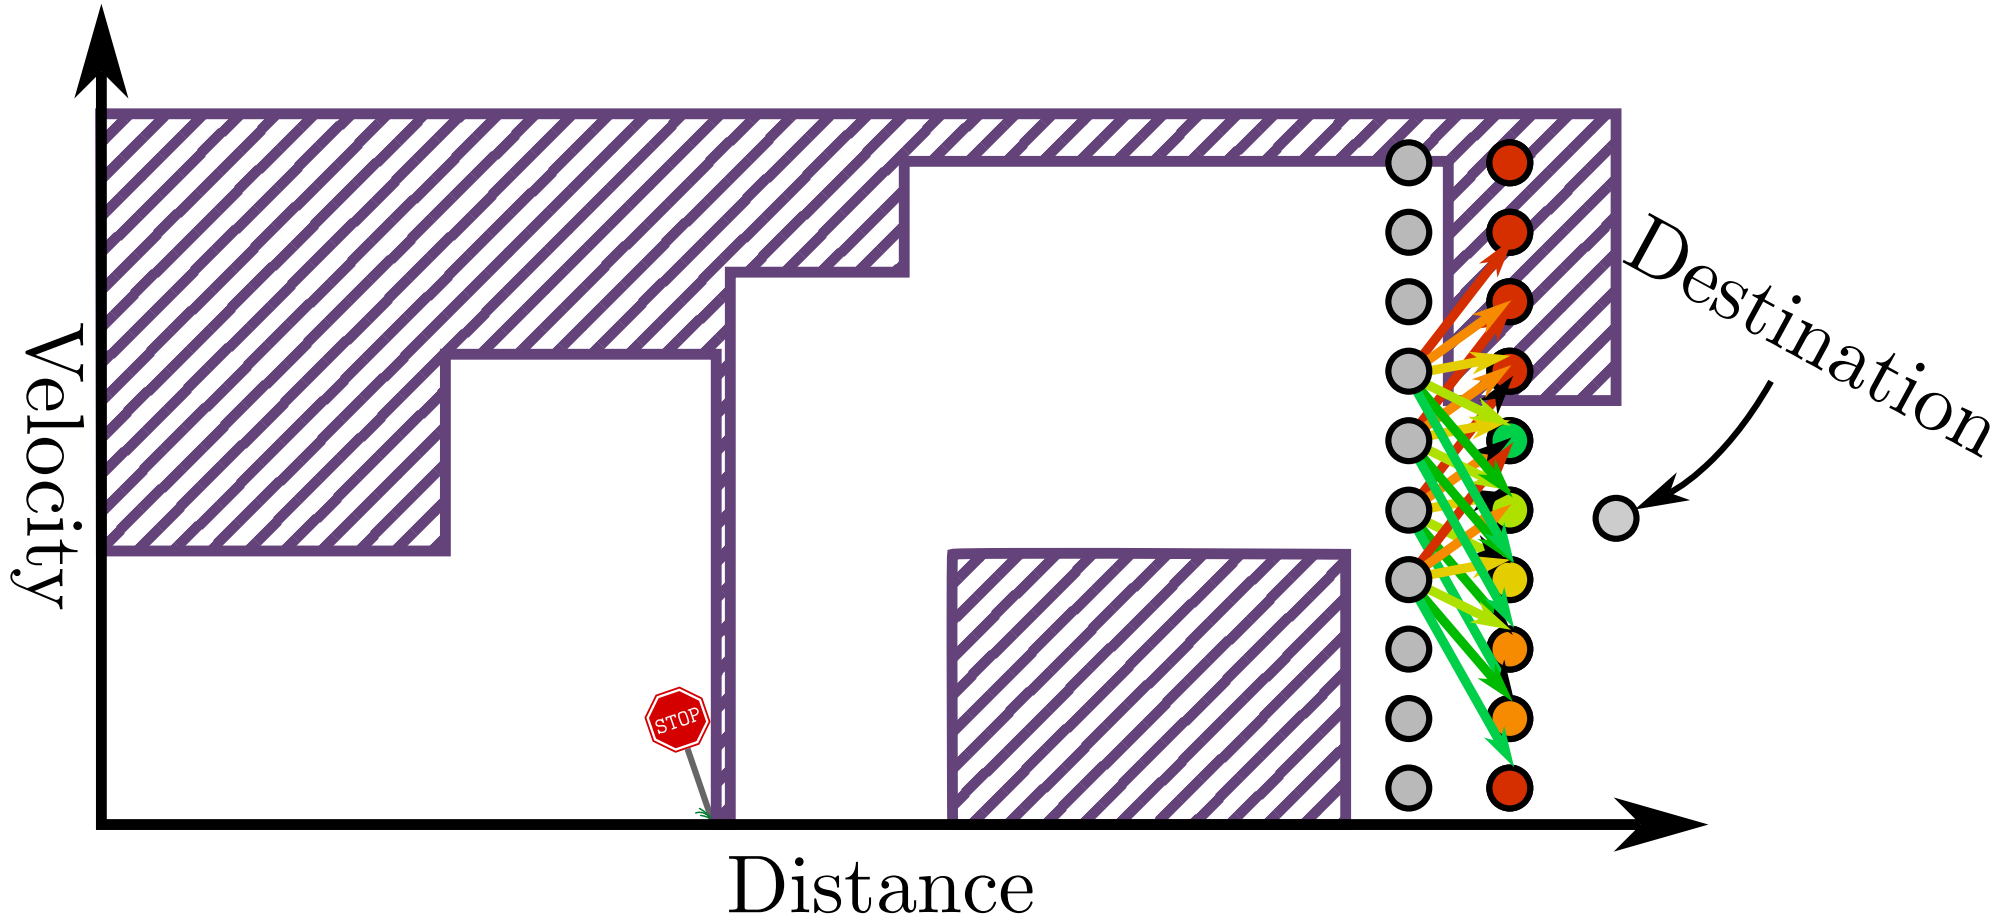
| $y \backslash x$ | 0     | 1     | 2     | 3     | ..... |
|------------------|-------|-------|-------|-------|-------|
| 0                | 0     | 1     | 0     | 1     | ..... |
| 1                | 2     | 3     | 2     | 3     | ..... |
| 2                | 0     | 1     | 0     | 1     | ..... |
| 3                | 2     | 3     | 2     | 3     | ..... |
|                  | ..... | ..... | ..... | ..... | ..... |

# Memory Read Partitioning

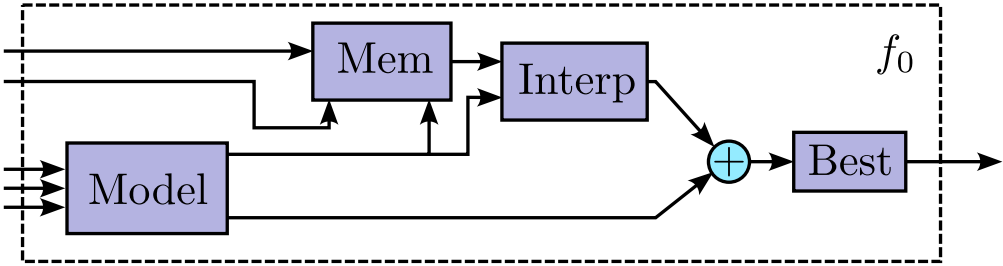
| $y \backslash x$ | 0     | 1     | 2     | 3     | ..... |
|------------------|-------|-------|-------|-------|-------|
| 0                | 0     | 1     | 0     | 1     | ..... |
| 1                | 2     | 3     | 2     | 3     | ..... |
| 2                | 0     | 1     | 0     | 1     | ..... |
| 3                | 2     | 3     | 2     | 3     | ..... |
|                  | ..... | ..... | ..... | ..... | ..... |

- 4 separate memories
- $x, y$  evenness determines indexing

# Memory Write Partitioning

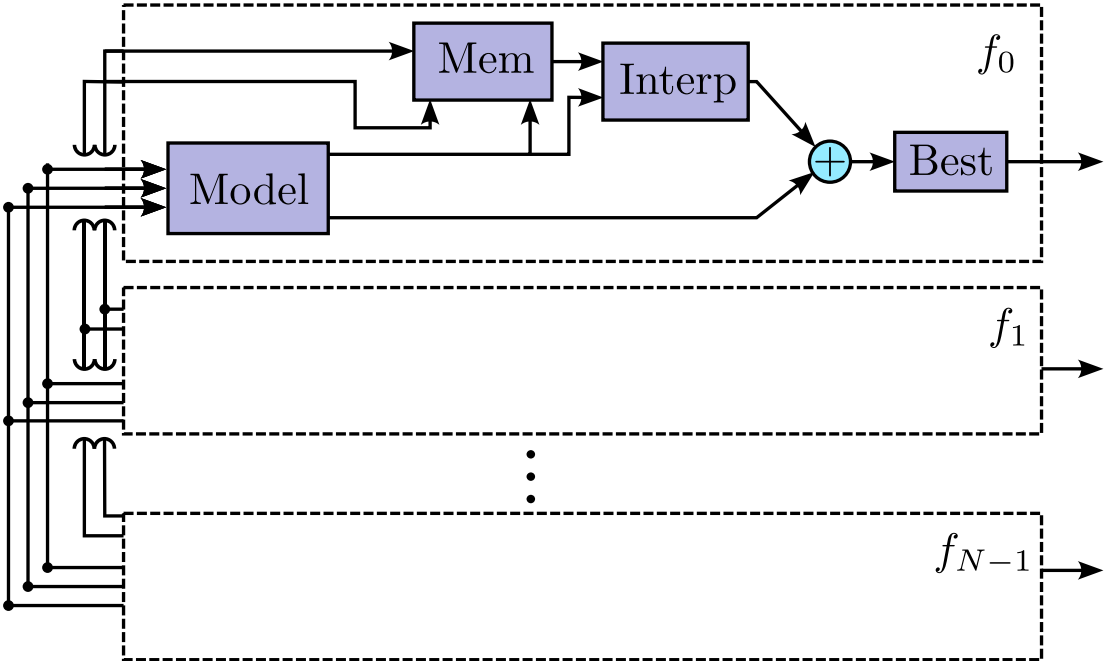


# Architecture Overview

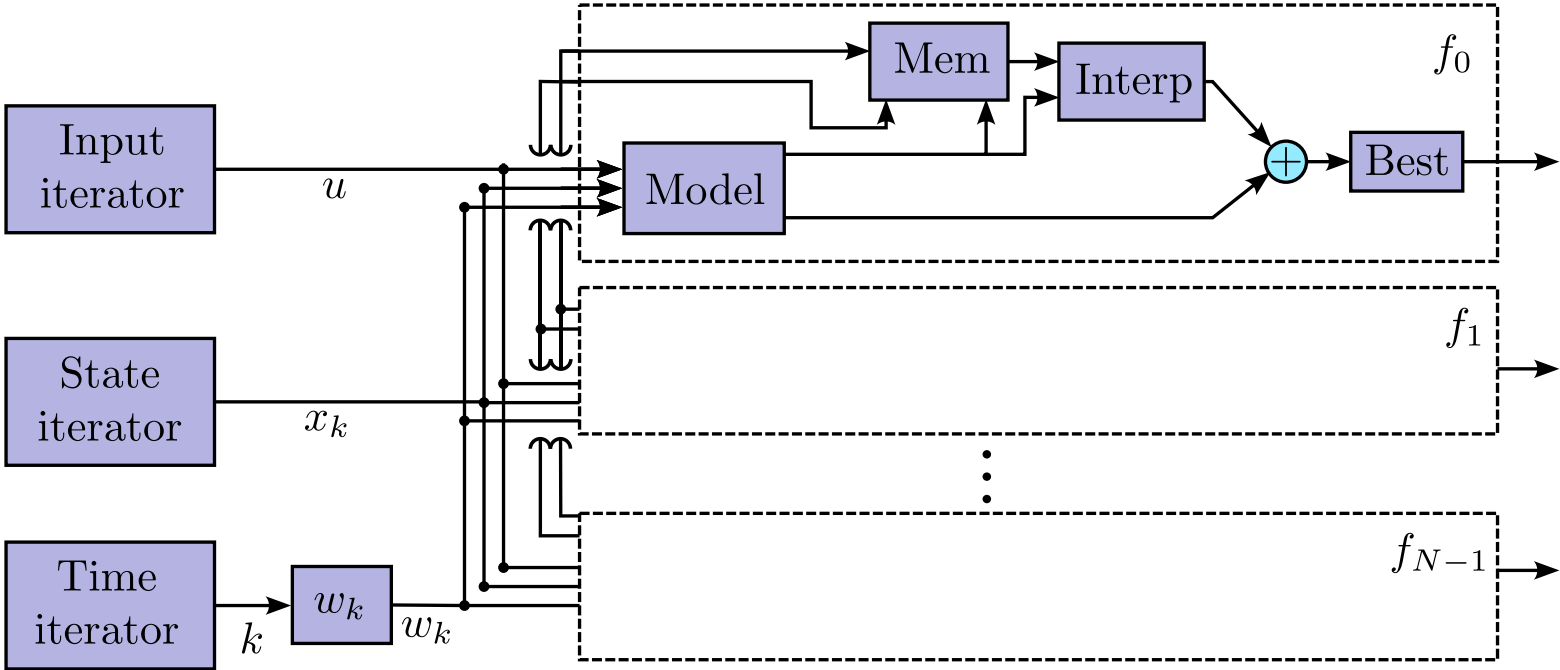




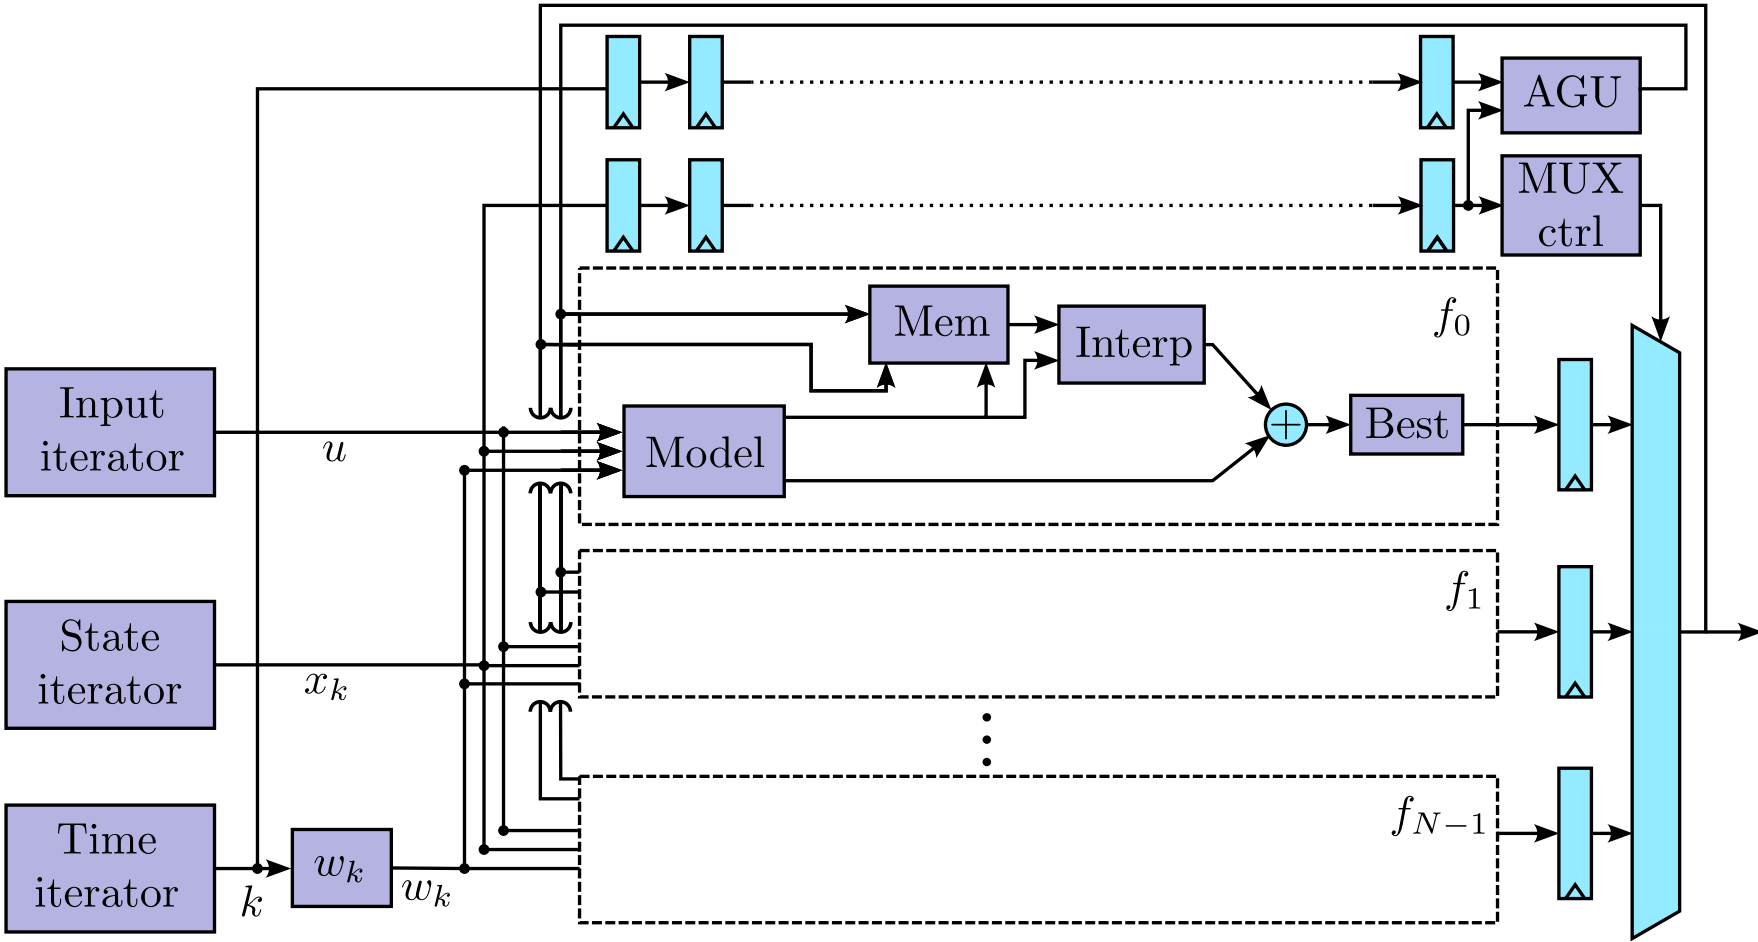
# Architecture Overview



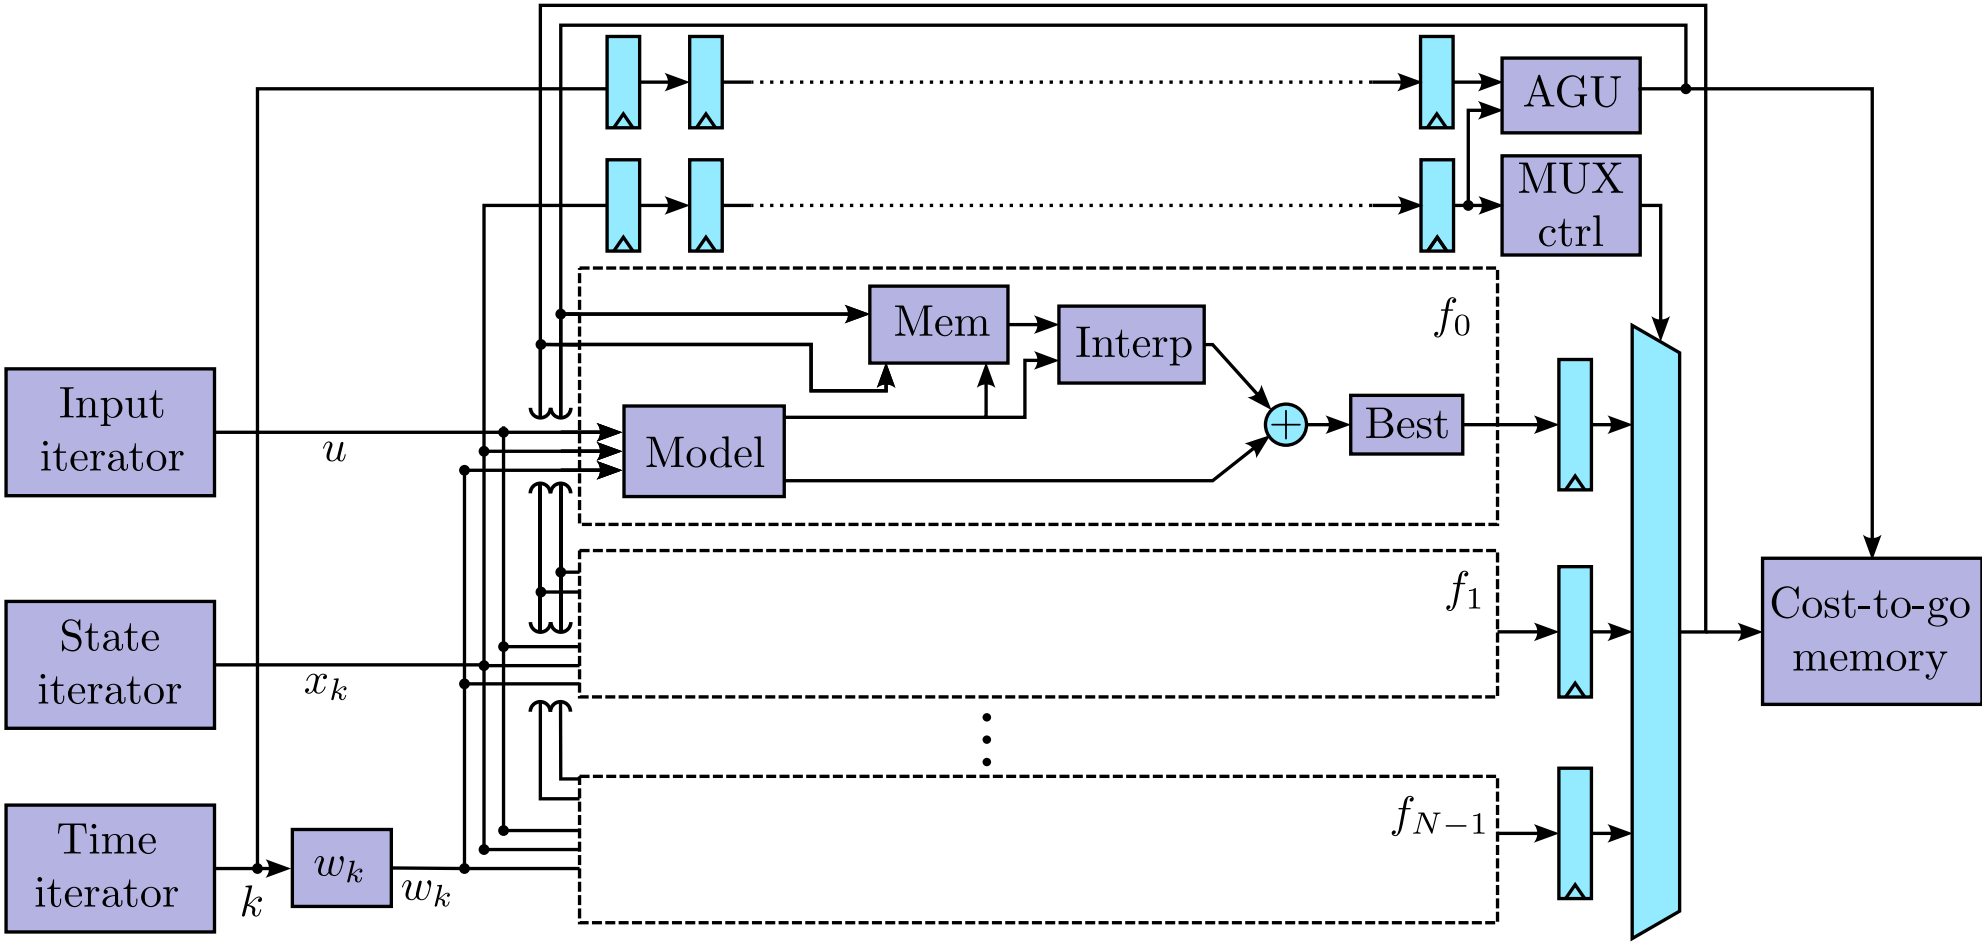
# Architecture Overview



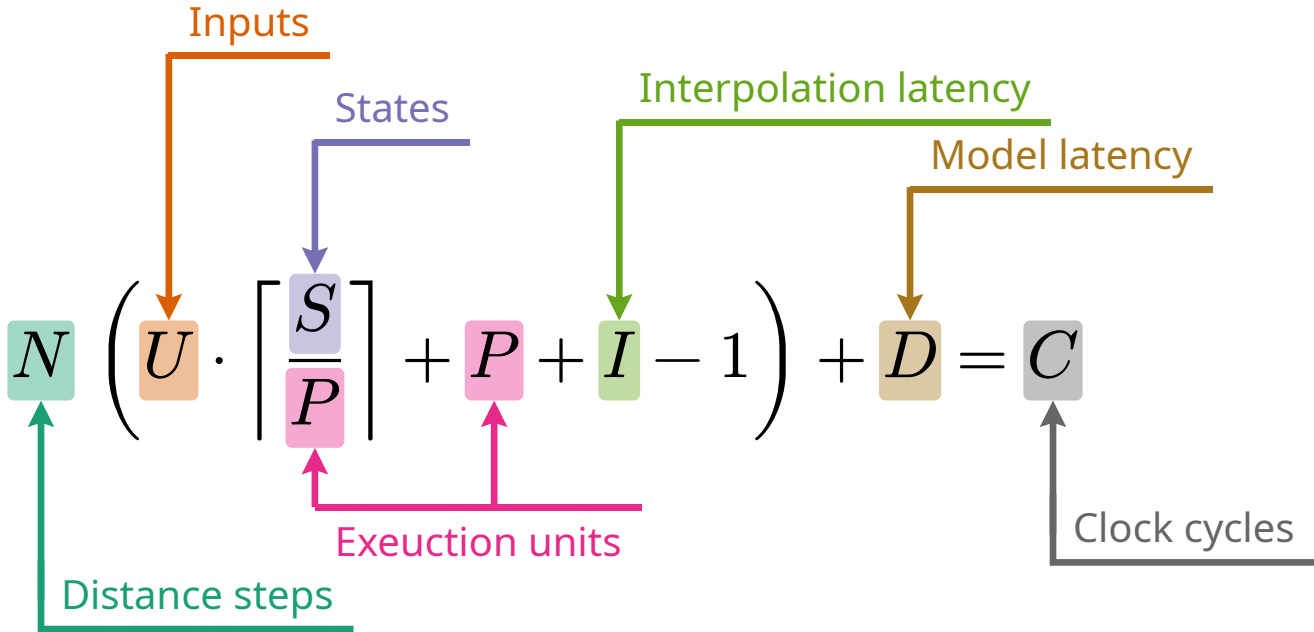
# Architecture Overview



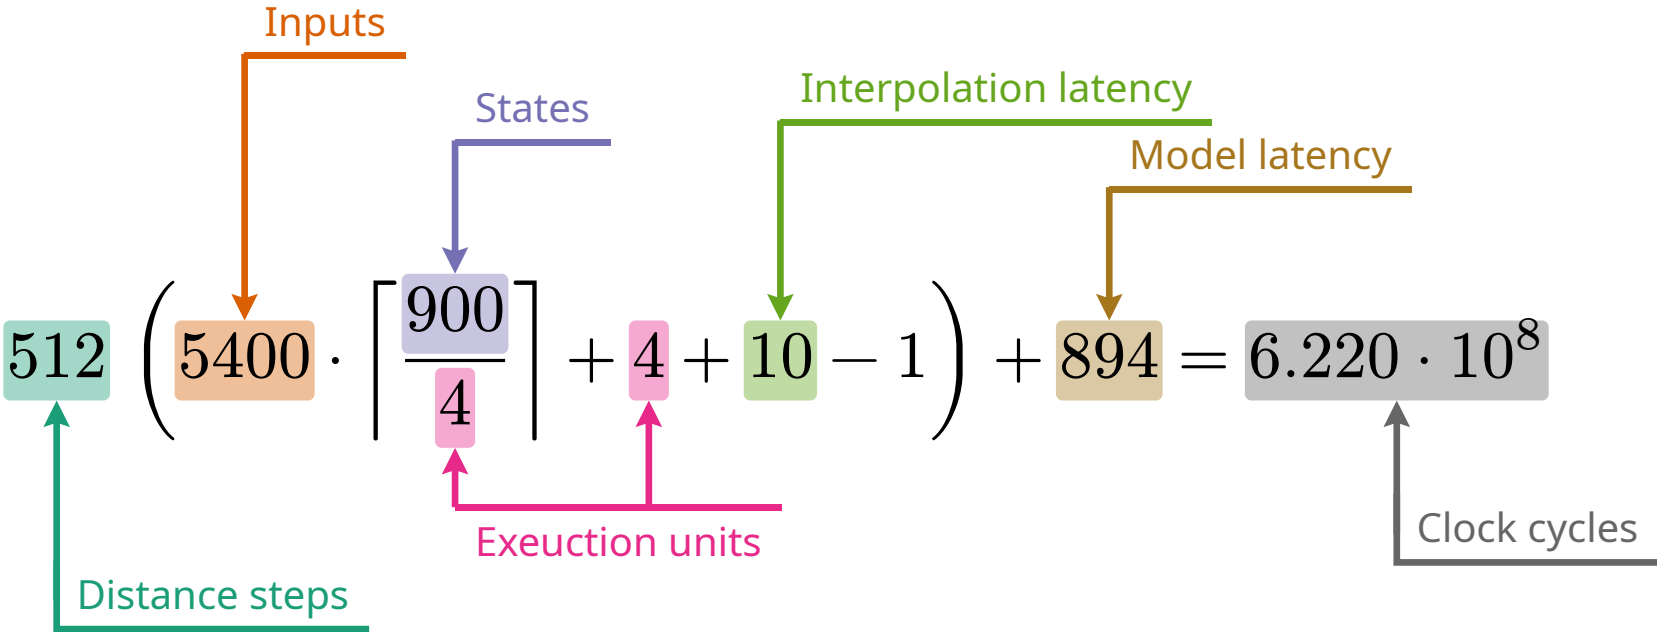
# Architecture Overview




# Schedule and Performance



# Schedule and Performance



# Resource Usage and Runtime

- Implemented in Spade HDL 
- Tools:
  - Vitis HLS 2022.1
  - Vivado 2022.1
  - AMD Virtex UltraScale+ xcvu13pfhga2104-3-e

# Resource Usage and Runtime

| <b>EUs</b> | $f_{\max}$ (MHz) | <b>CCs required</b> | <b>Run time (s)</b> | <b>Speedup</b> | <b>CLB</b> | <b>DSP</b> | <b>BRAM</b> | <b>URAM</b> |
|------------|------------------|---------------------|---------------------|----------------|------------|------------|-------------|-------------|
| 1          | 370              | $2.488 \cdot 10^9$  | 6.73                | 18 ×           | 10196      | 515        | 238         | 154         |
| 2          | 336              | $1.244 \cdot 10^9$  | 3.70                | 34 ×           | 21384      | 1030       | 476         | 154         |
| 4          | 335              | $6.220 \cdot 10^8$  | 1.85                | 68 ×           | 41654      | 2060       | 952         | 154         |
| 6          | 288              | $4.147 \cdot 10^8$  | 1.44                | 87 ×           | 62627      | 3090       | 1428        | 154         |
| 9          | 272              | $2.764 \cdot 10^8$  | 1.02                | 123 ×          | 94874      | 4635       | 2142        | 154         |



# Resource Usage and Runtime

| <b>EUs</b> | $f_{\max}$ (MHz) | <b>CCs required</b> | <b>Run time (s)</b> | <b>Speedup</b> | <b>CLB</b> | <b>DSP</b> | <b>BRAM</b> | <b>URAM</b> |
|------------|------------------|---------------------|---------------------|----------------|------------|------------|-------------|-------------|
| 1          | 370              | $2.488 \cdot 10^9$  | 6.73                | 18 ×           | 10196      | 515        | 238         | 154         |
| 2          | 336              | $1.244 \cdot 10^9$  | 3.70                | 34 ×           | 21384      | 1030       | 476         | 154         |
| 4          | 335              | $6.220 \cdot 10^8$  | 1.85                | 68 ×           | 41654      | 2060       | 952         | 154         |
| 6          | 288              | $4.147 \cdot 10^8$  | 1.44                | 87 ×           | 62627      | 3090       | 1428        | 154         |
| 9          | 272              | $2.764 \cdot 10^8$  | 1.02                | 123 ×          | 94874      | 4635       | 2142        | 154         |

# Resource Usage and Runtime

| <b>EUs</b> | $f_{\max}$ (MHz) | <b>CCs required</b> | <b>Run time (s)</b> | <b>Speedup</b> | <b>CLB</b> | <b>DSP</b> | <b>BRAM</b> | <b>URAM</b> |
|------------|------------------|---------------------|---------------------|----------------|------------|------------|-------------|-------------|
| 1          | 370              | $2.488 \cdot 10^9$  | <b>6.73</b>         | 18 ×           | 10196      | 515        | 238         | 154         |
| 2          | 336              | $1.244 \cdot 10^9$  | <b>3.70</b>         | 34 ×           | 21384      | 1030       | 476         | 154         |
| 4          | 335              | $6.220 \cdot 10^8$  | <b>1.85</b>         | 68 ×           | 41654      | 2060       | 952         | 154         |
| 6          | 288              | $4.147 \cdot 10^8$  | <b>1.44</b>         | 87 ×           | 62627      | 3090       | 1428        | 154         |
| 9          | 272              | $2.764 \cdot 10^8$  | <b>1.02</b>         | 123 ×          | 94874      | 4635       | 2142        | 154         |

# Resource Usage and Runtime

| <b>EUs</b> | $f_{\max}$ (MHz) | <b>CCs required</b> | <b>Run time (s)</b> | <b>Speedup</b> | <b>CLB</b> | <b>DSP</b> | <b>BRAM</b> | <b>URAM</b> |
|------------|------------------|---------------------|---------------------|----------------|------------|------------|-------------|-------------|
| 1          | 370              | $2.488 \cdot 10^9$  | <b>6.73</b>         | 18 ×           | 10196      | 515        | 238         | 154         |
| 2          | 336              | $1.244 \cdot 10^9$  | <b>3.70</b>         | 34 ×           | 21384      | 1030       | 476         | 154         |
| 4          | 335              | $6.220 \cdot 10^8$  | <b>1.85</b>         | 68 ×           | 41654      | 2060       | 952         | 154         |
| 6          | 288              | $4.147 \cdot 10^8$  | <b>1.44</b>         | 87 ×           | 62627      | 3090       | 1428        | 154         |
| 9          | 272              | $2.764 \cdot 10^8$  | <b>1.02</b>         | 123 ×          | 94874      | 4635       | 2142        | 154         |

Single threaded Xeon W-1250: **126 seconds**

# Resource Usage and Runtime

| <b>EUs</b> | $f_{\max}$ (MHz) | <b>CCs required</b> | <b>Run time (s)</b> | <b>Speedup</b> | <b>CLB</b> | <b>DSP</b> | <b>BRAM</b> | <b>URAM</b> |
|------------|------------------|---------------------|---------------------|----------------|------------|------------|-------------|-------------|
| 1          | 370              | $2.488 \cdot 10^9$  | <b>6.73</b>         | 18 ×           | 10196      | 515        | 238         | 154         |
| 2          | 336              | $1.244 \cdot 10^9$  | <b>3.70</b>         | 34 ×           | 21384      | 1030       | 476         | 154         |
| 4          | 335              | $6.220 \cdot 10^8$  | <b>1.85</b>         | 68 ×           | 41654      | 2060       | 952         | 154         |
| 6          | 288              | $4.147 \cdot 10^8$  | <b>1.44</b>         | 87 ×           | 62627      | 3090       | 1428        | 154         |
| 9          | 272              | $2.764 \cdot 10^8$  | <b>1.02</b>         | 123 ×          | 94874      | 4635       | 2142        | 154         |

Single threaded Xeon W-1250: **126 seconds**

# Resource Usage and Runtime

| <b>EUs</b> | $f_{\max}$ (MHz) | <b>CCs required</b> | <b>Run time (s)</b> | <b>Speedup</b> | <b>CLB</b> | <b>DSP</b> | <b>BRAM</b> | <b>URAM</b> |
|------------|------------------|---------------------|---------------------|----------------|------------|------------|-------------|-------------|
| 1          | 370              | $2.488 \cdot 10^9$  | 6.73                | 18 ×           | 10196      | 515        | 238         | 154         |
| 2          | 336              | $1.244 \cdot 10^9$  | 3.70                | 34 ×           | 21384      | 1030       | 476         | 154         |
| 4          | 335              | $6.220 \cdot 10^8$  | 1.85                | 68 ×           | 41654      | 2060       | 952         | 154         |
| 6          | 288              | $4.147 \cdot 10^8$  | 1.44                | 87 ×           | 62627      | 3090       | 1428        | 154         |
| 9          | 272              | $2.764 \cdot 10^8$  | 1.02                | 123 ×          | 94874      | 4635       | 2142        | 154         |

- CLB, DSP, BRAM increase linearly
- URAM only stores final cost-to-go map
- Vehicle model dominates resource usage and  $f_{\max}$

# Conclusions & Contributions

- **Scalable** architecture for hybrid electric vehicle optimization
- Enabling **real-time** use
- $> 100 \times$  **speedup** over CPU implementation

frans.skarman@liu.se