Adventures in Analog Computing

Michael Schapira

האוניברסיטה העברית בירושלים
THE HEBREW UNIVERSITY OF JERUSALEM

Microsoft

Microsoft's Analog-Optical Computer (**AOC**) Project:

A *new* kind of (analog) optical computer

potentially *100x more efficient* than GPUs

built using *consumer-sector technologies*

to *accelerate optimization and AI inference*

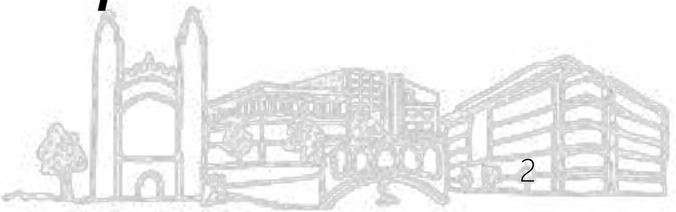

Two questions

- What is the **expressiveness** of AOC?
- What are the **killer apps** for AOC?

Answers?

- A hierarchy of analog computing problem classes
- **Real-time control optimization**

The AOC Team

Disciplines: Optics, Electronics, Mechanics, ML, Optimization, Algorithms, Maths, Systems

Quadratic Unconstrained Mixed Optimization (QUMO)

$$\min_{x} \; x^T Q x + q^T x + c$$

Where:

Q – matrix

q – vector

x – variables (binary and continuous)

Linear constraints on x can be introduced
by adding penalty terms to the QUMO objective

AOC to the rescue!

(gen-2 machine)

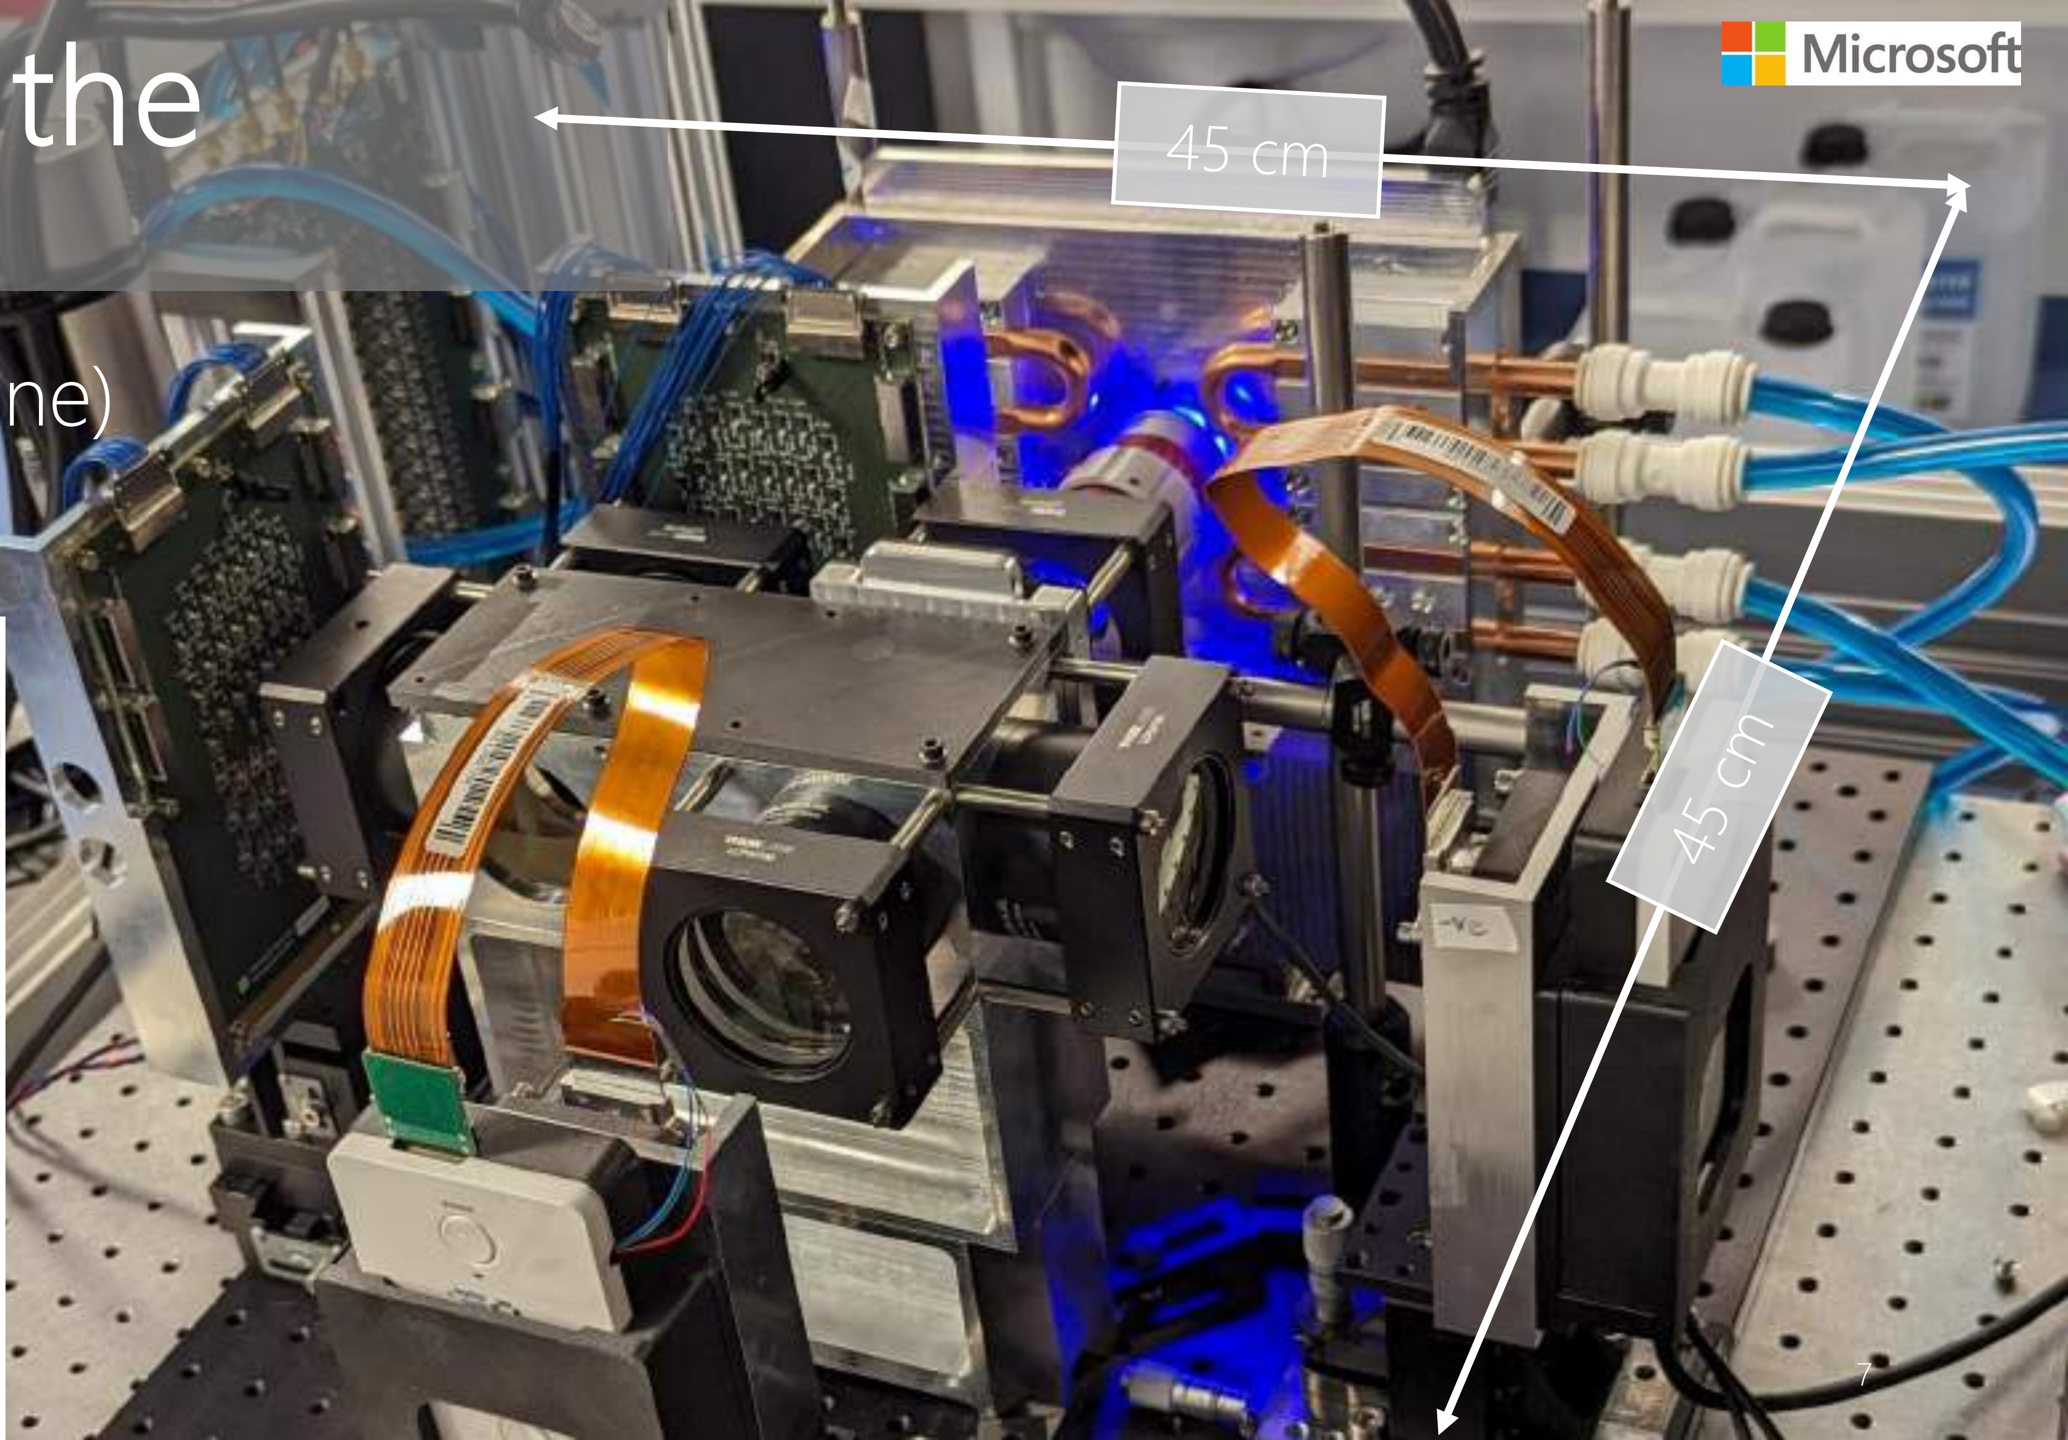

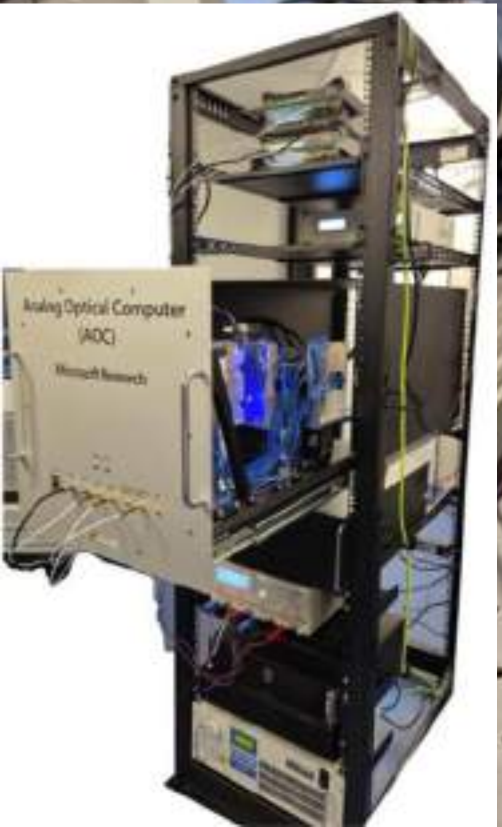

Massively-parallel (enhanced) analog gradient descent

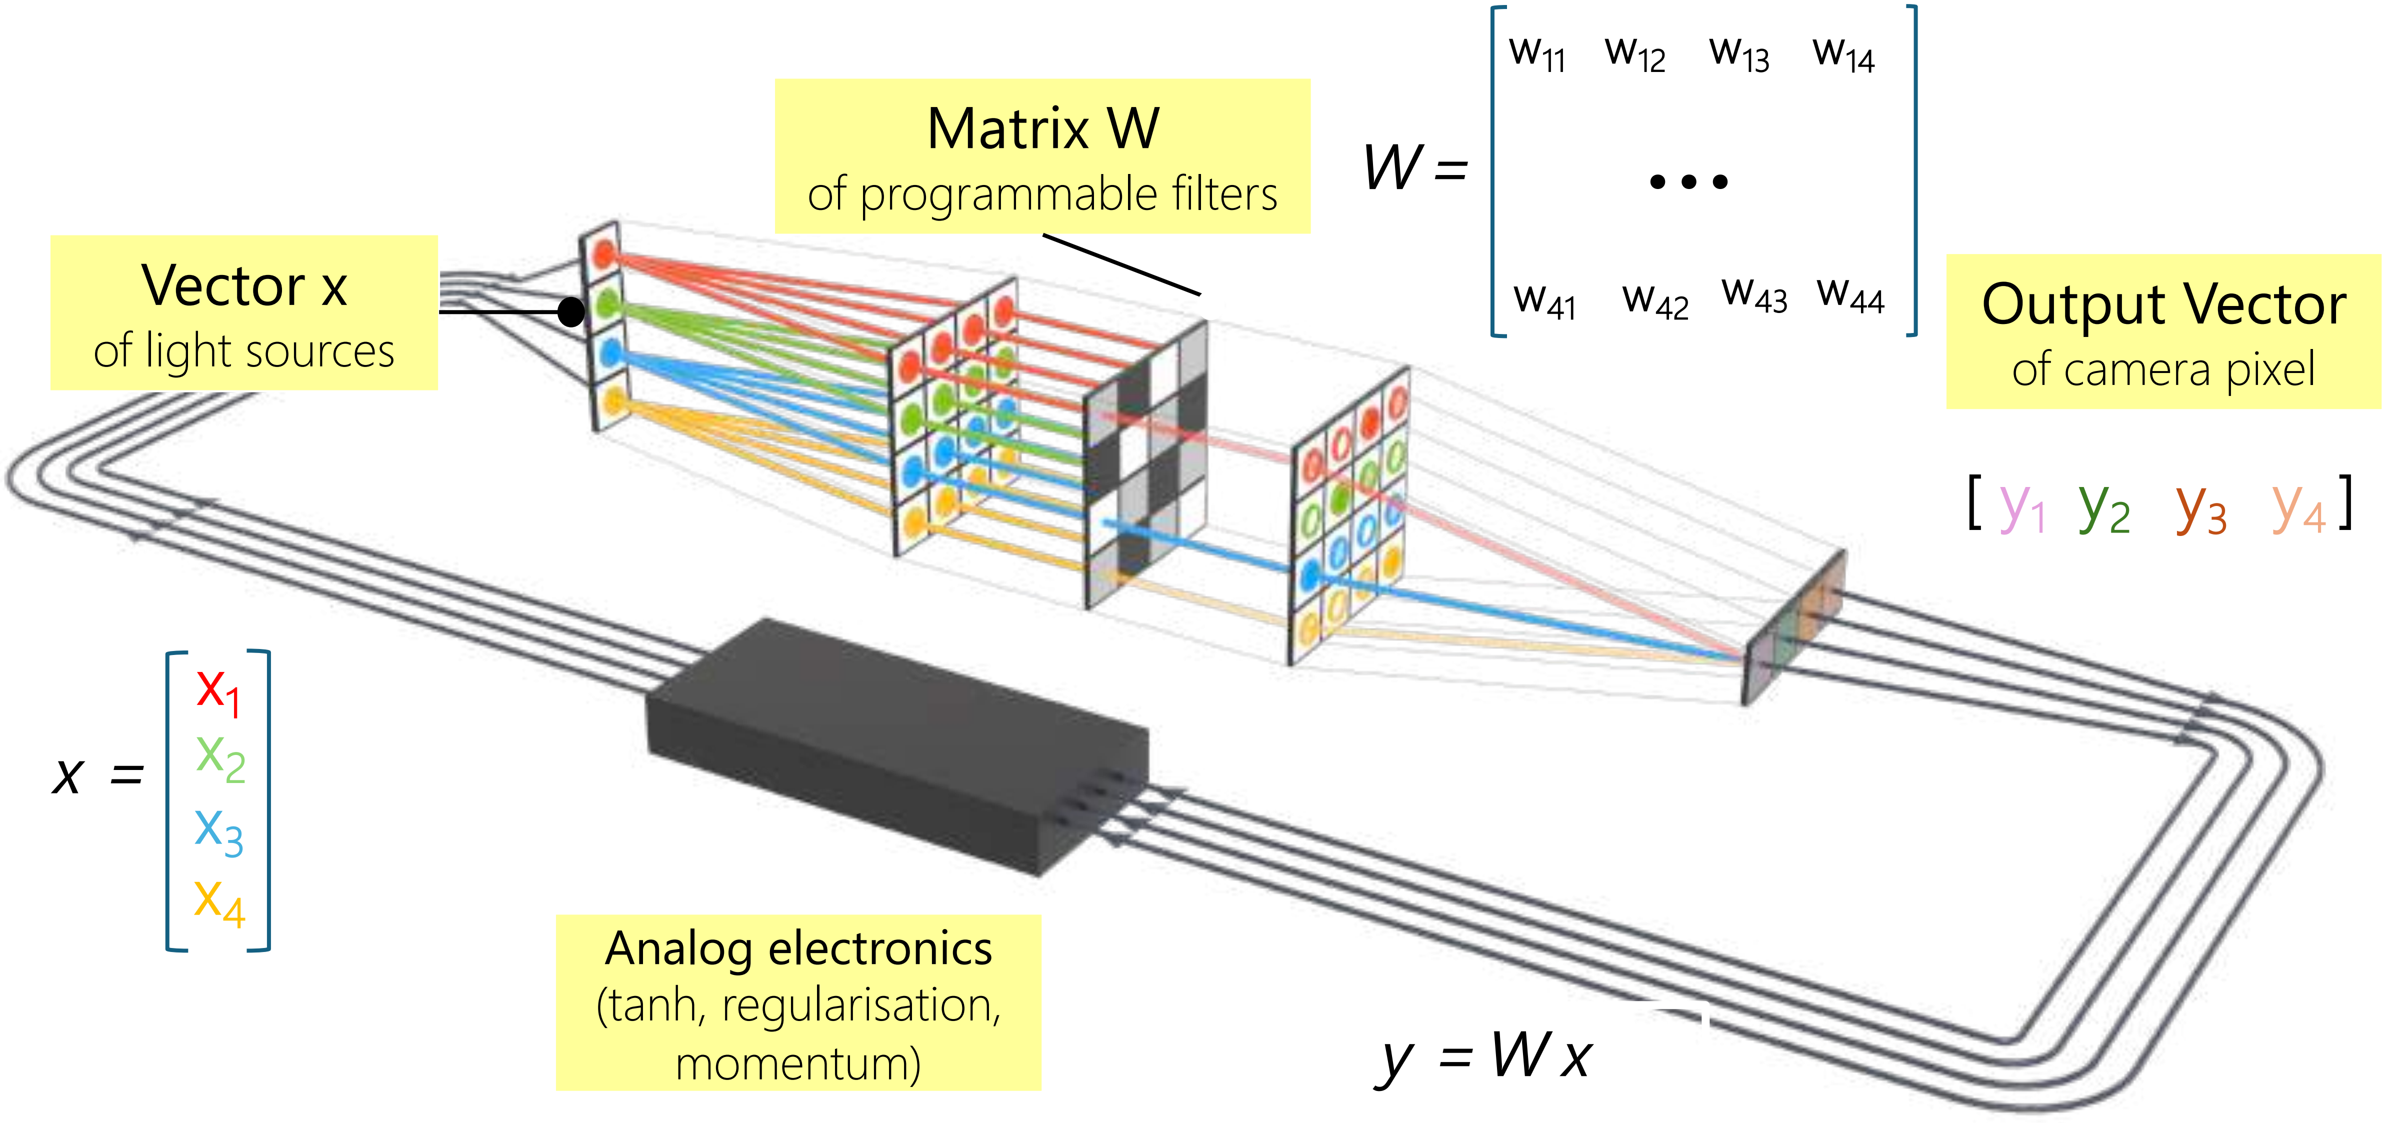

Iterative Update Rule: $y_{t+1} = y_t + \Delta t[W * (f_{nonlinear}(y_t))]$

Model Predictive Control (MPC)

<u>Repeat</u>:

- Observe environment state
- **Optimize action trajectory wrt to a model of the environment**
- Apply next action

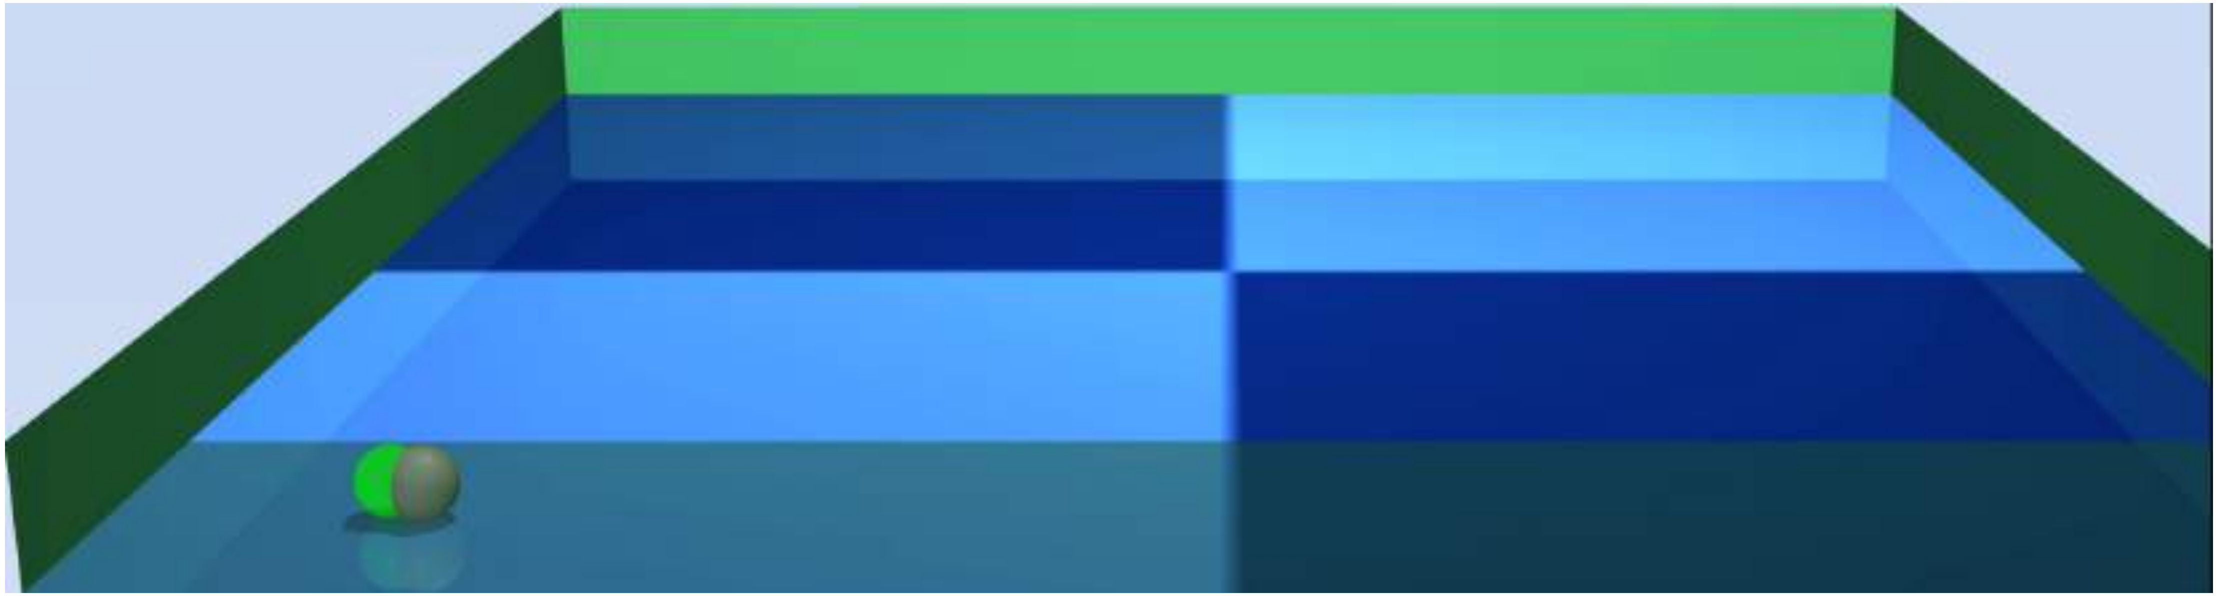

Microsoft

Our Simple Example: Quadratic Objective with Linear Constraints

Inputs: current state x_0, target states $\{x_t^*\}_{t=1}^{T}$

$$\underset{x,u}{\text{minimize}} \sum_{t=1}^{T} w_1 \cdot \|x_t - x_t^*\|_2^2 + w_2 \cdot \|u_t\|_2^2$$

Quadratic objective function

$$s.t.$$

$$u_{min} \le u_t \le u_{max}$$

Linear constraints on controls

Linear approximation of the system dynamics

$$x_{t+1} = f(\hat{x}_t, \hat{u}_t) + J_f(\hat{x}_t, \hat{u}_t) \cdot \begin{bmatrix} x_t - \hat{x}_t \\ u_t - \hat{u}_t \end{bmatrix}$$

Model Predictive Control

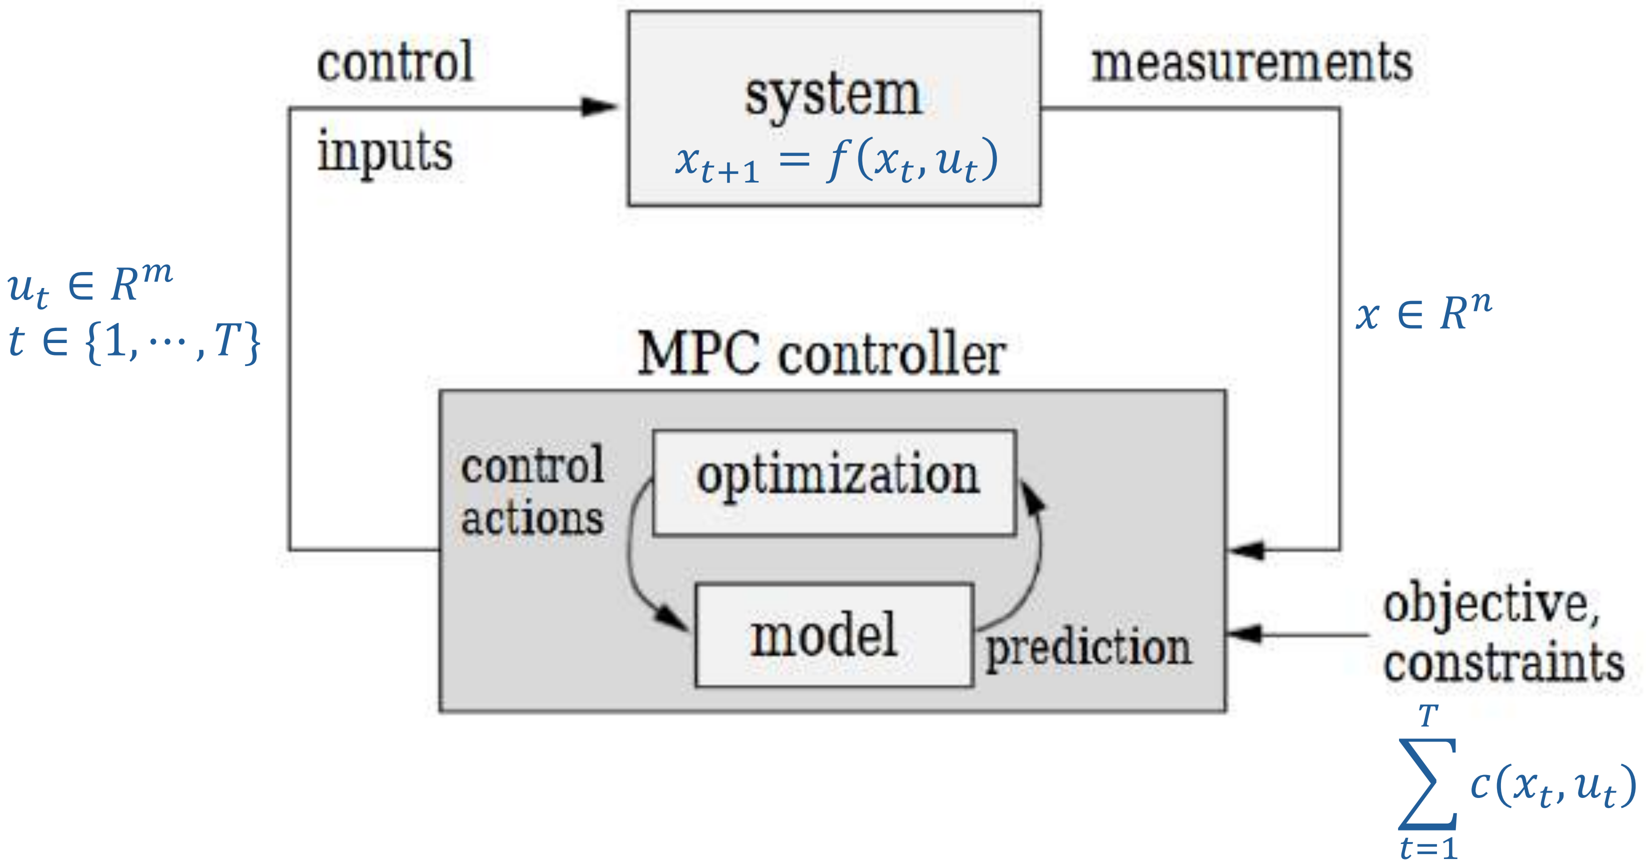

Model Predictive Control

- Fundamental to control theory and practice

- Applications in industrial control , robotics, autonomous vehicles, ...

- **Natural candidate for (analog) hardware acceleration!**
  - Frequent and intensive **online computation**
  - Continuous and **discrete** variables
  - **Non-convex** objectives
  - **QUMO**-compatible
  - Manageable **scale**

Research Agenda

- Devise and implement a **prototype for AOC-powered MPC**
  - Building on MuJoCo MPC in C/C++, AOC simulator in Julia
- **Identify MPC tasks** for which AOC is well-suited
  - Control with discrete variables
  - Non-convex optimization
  - Example: "Track and Avoid" tasks
- **Evaluate AOC-based MPC** and contrast with standard baselines

Model Predictive Control

- Fundamental to control theory and practice
- Applications in industrial control , robotics, autonomous vehicles, ...
- **Natural candidate for (analog) hardware acceleration!**
  - Intensive **online computation**
  - Continuous and **discrete** variables
  - **Non-convex** objectives
  - **QUMO**-compatible
  - Manageable **scale**

Use Case I - **Discrete Actions**

AOC Planner

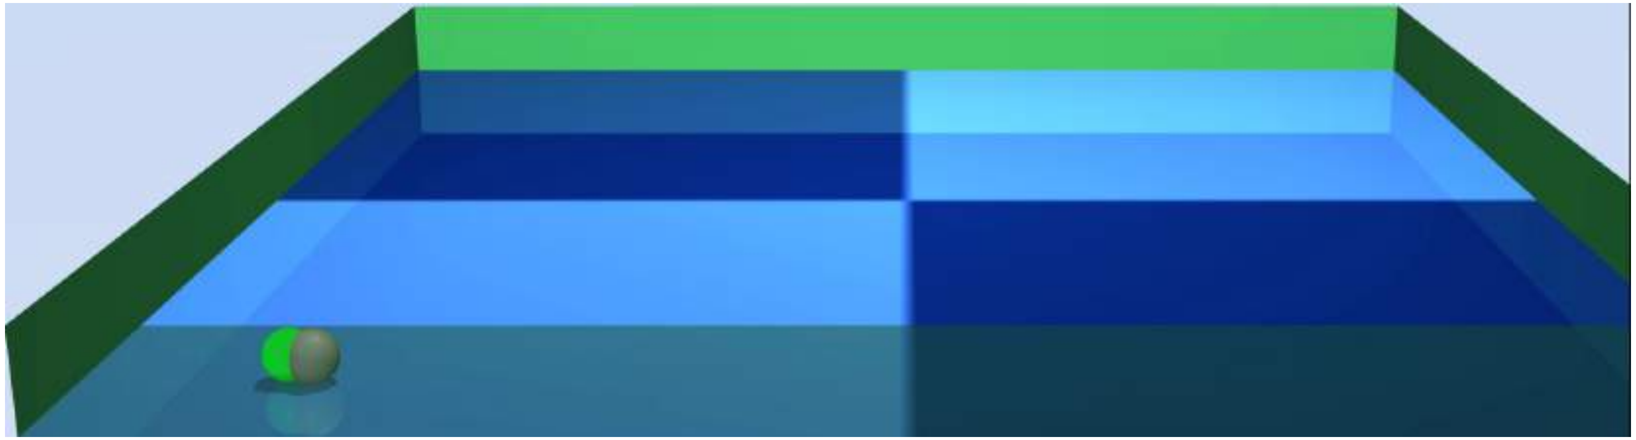

Gradient Descent Planner

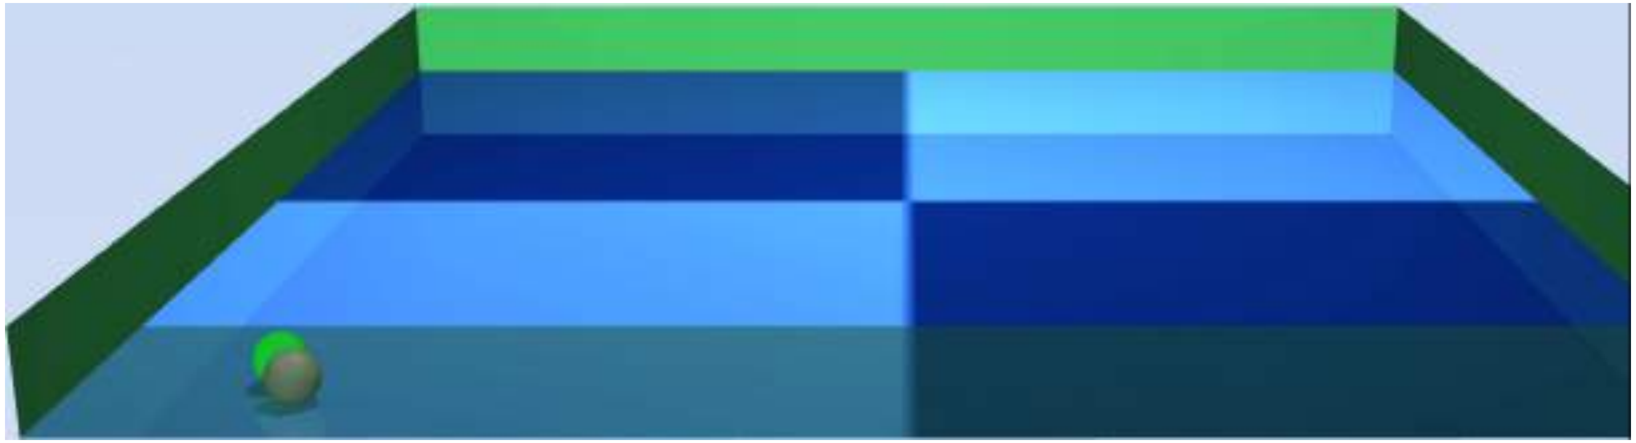

Use Case II - **Non-convex Objective Functions**: "Track and Avoid" (1D)

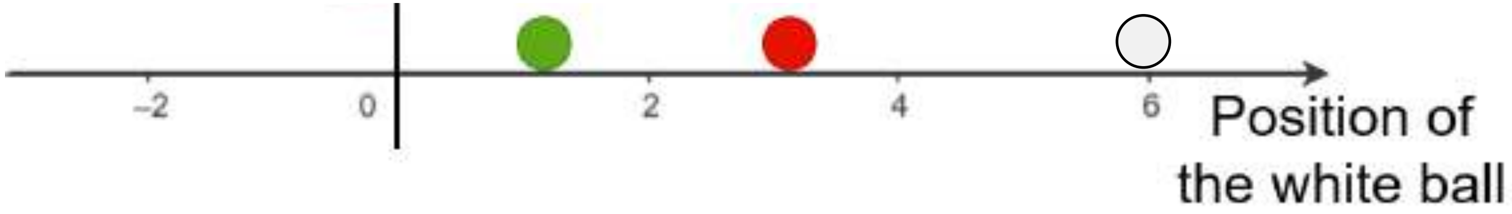

Use Case II - **Non-convex Objective Functions:** "Track and Avoid" (1D)

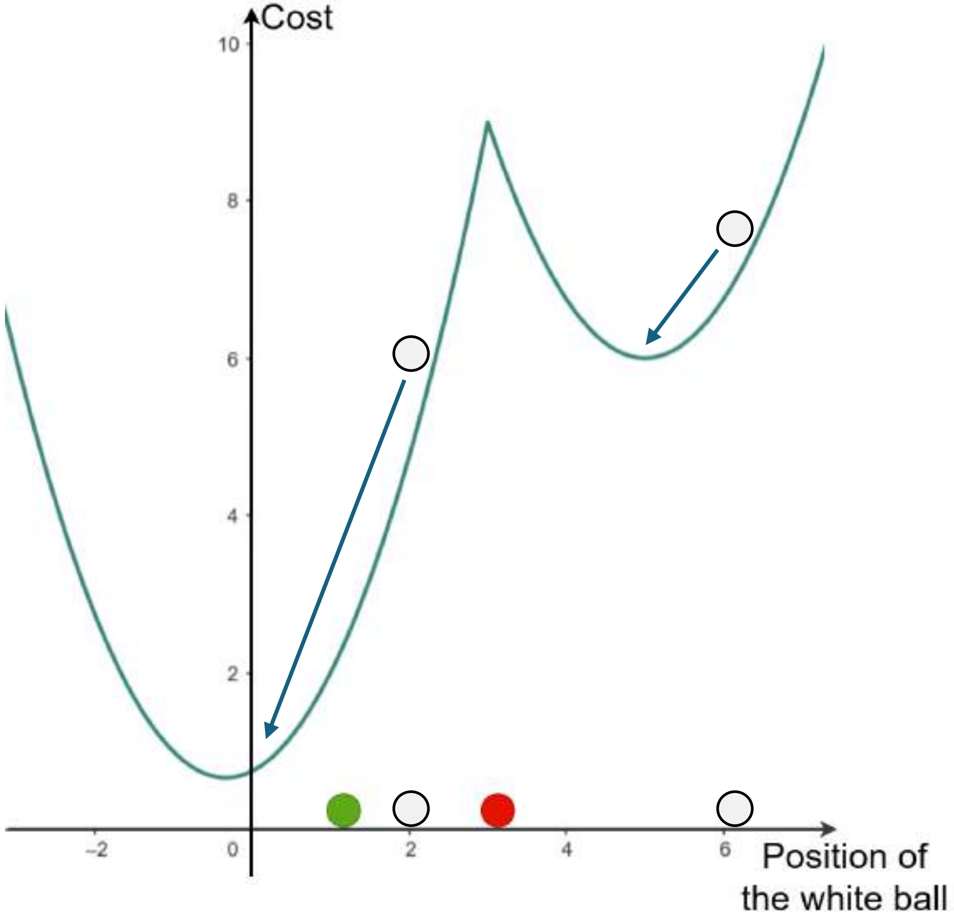

Cost formulation leverages QUMO's expressiveness (binary variables, linear constraints)

DEMO – "Track and Avoid" in 2D

AOC Planner

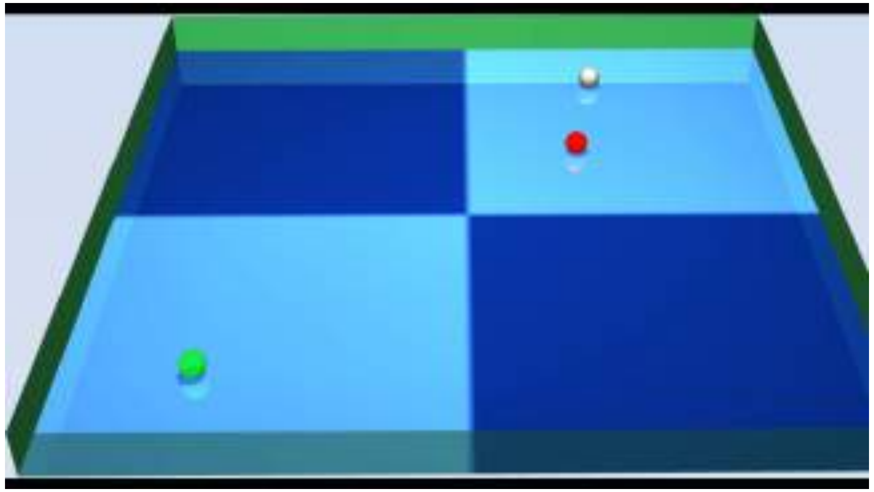

ILQG Planner

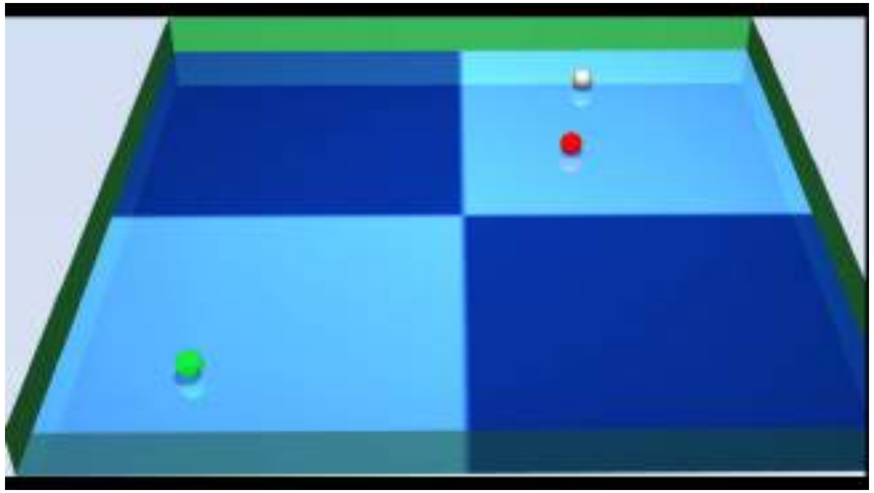

Future Directions

- **Better models** of the environment
  - Extend linear model by linearize around multiple points
  - A learned world model

- **More efficient** implementations
  - "Recycle" hyperparameter tuning across MPC iterations
  - More efficient linearization

- Extensions to **more complex tasks**
  - e.g., 3D track and avoid
  - collaboration with the datacenter robotics team

- **Extensions to MPC logic**
  - Minimax MPC
  - Multi-agent MPC

Thank you!