

# Model-Based Optimization for Intelligent Robot Control

Tom Erez

Motion Control Lab  
Computer Science  
University of Washington

# Goal:

## optimal control of humanoid robots

- Quick, autonomous adaptation to novel circumstances.
- Unified framework for different tasks (e.g. grasping, walking).
- User guidance should be easy and intuitive.

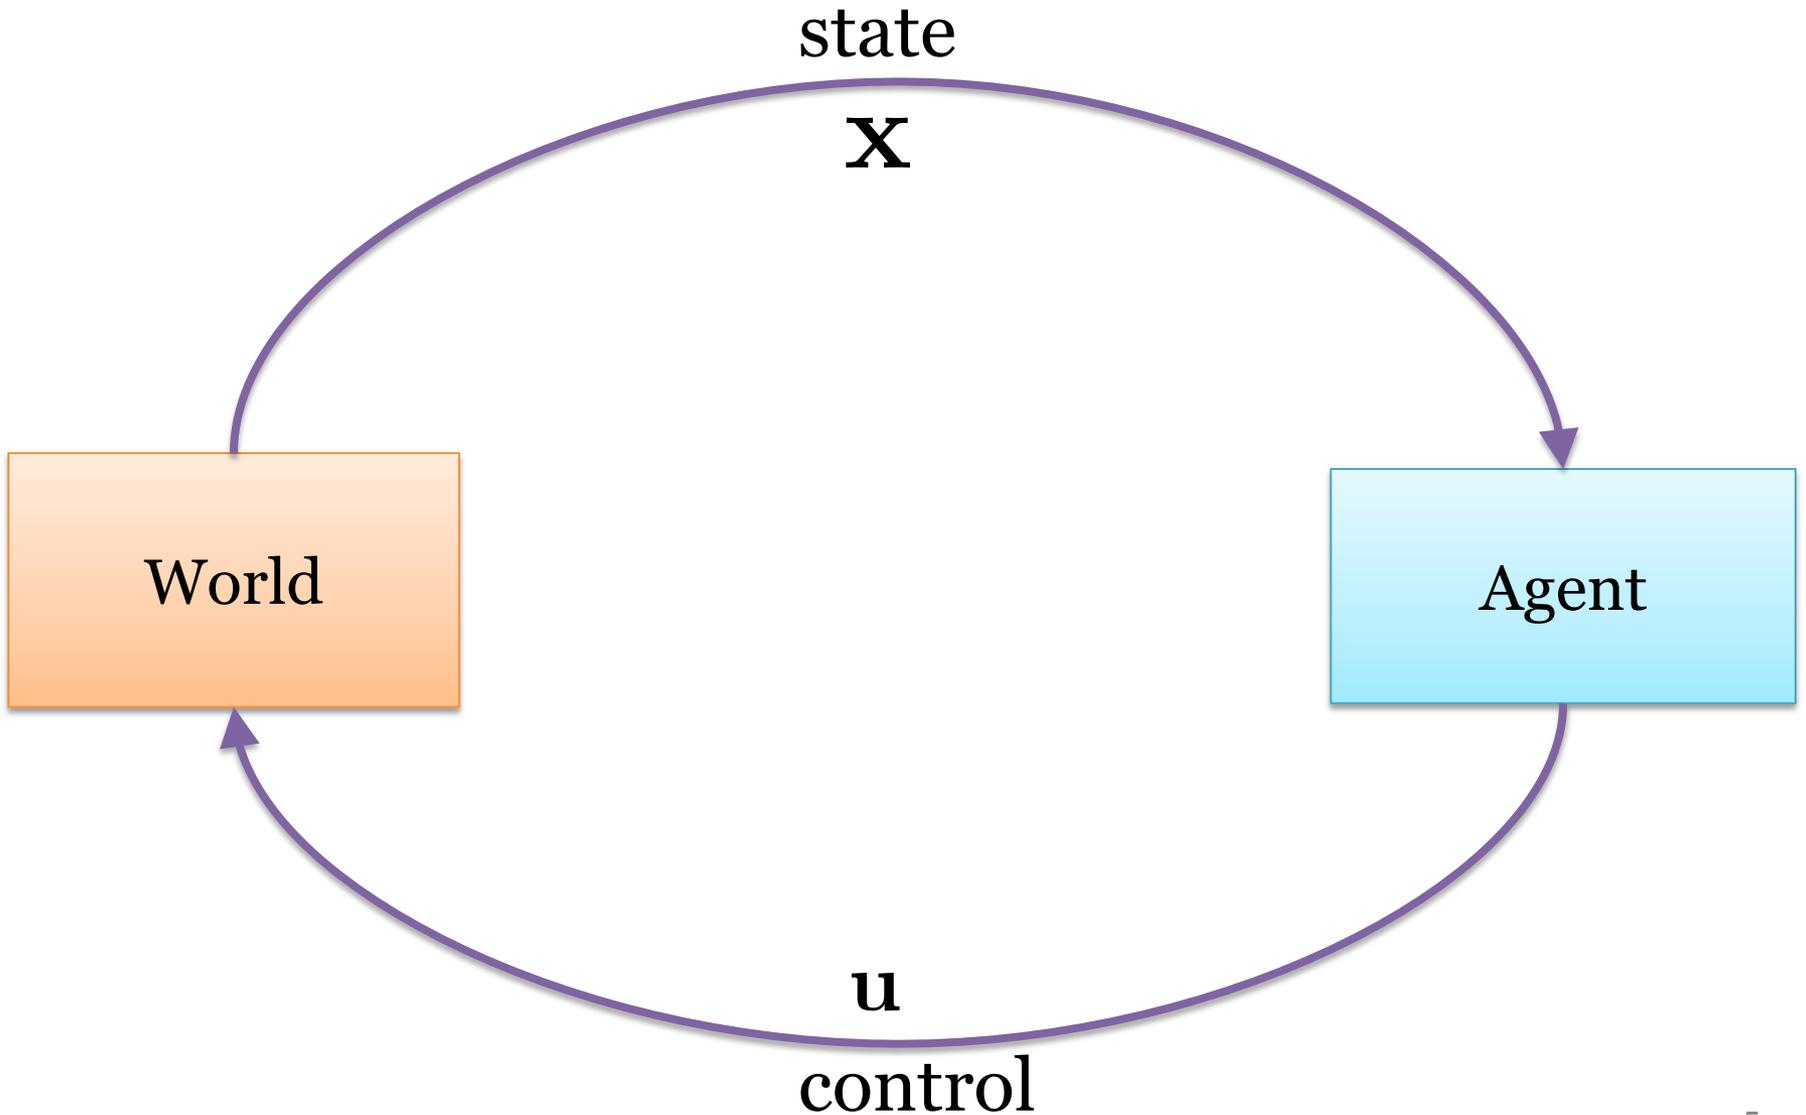
# Challenges

- **Simulation:** torques, friction, contact.
- **Optimization:** nonlinear, nonconvex, high-DoF.
- **Systems:** estimation, integration, feedback control.
- **Human operator:** flexibility, visualization, interfaces.

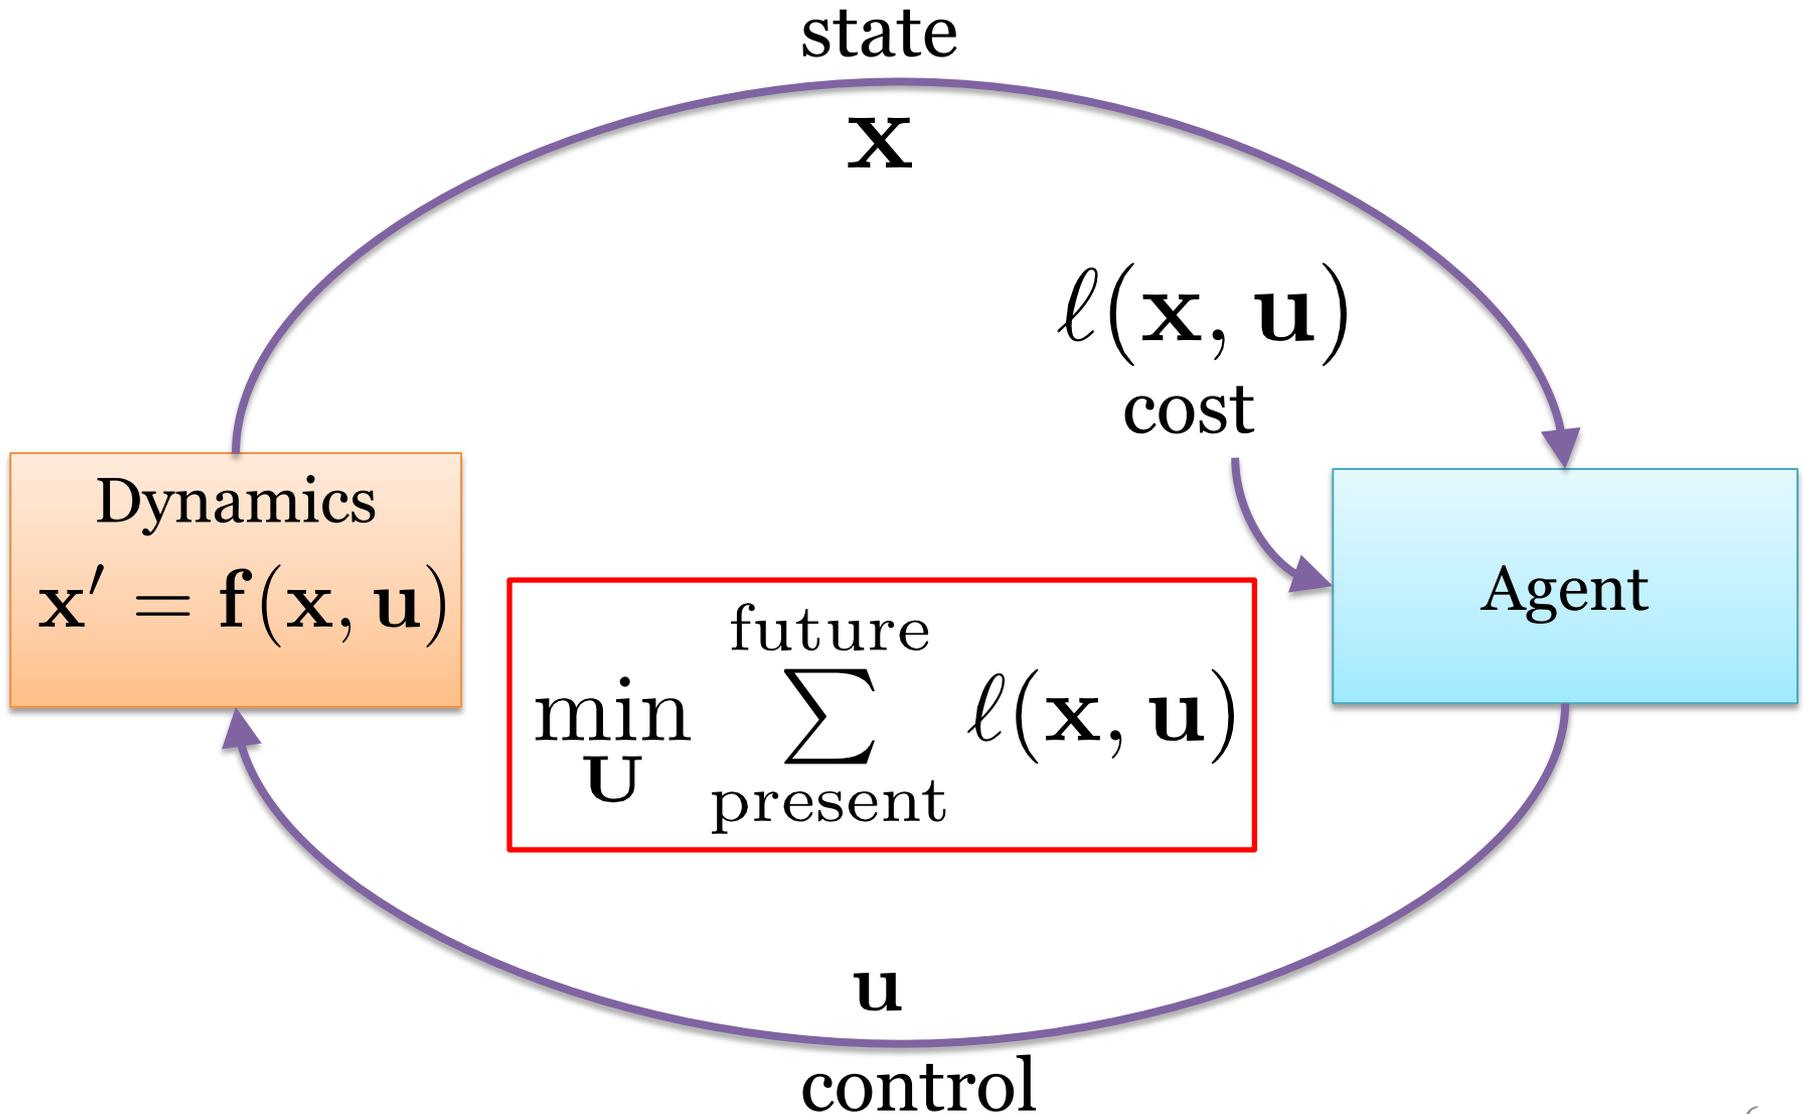
# MPC in MuJoCo

- A new physics engine: MuJoCo (Multi-Joint simulation with Contacts).
- An optimization framework around MuJoCo that can operate in real-time.
- An integrated system for controlling articulated robots.

# Optimal Control



# Optimal Control



# Model-Based Optimization

Physics

$$f(\mathbf{x}, \mathbf{u})$$

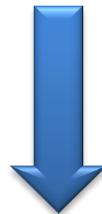
Task

$$\ell(\mathbf{x}, \mathbf{u})$$



Optimization

$$\min_{\mathbf{U}} \sum_{\text{present}}^{\text{future}} \ell(\mathbf{x}, \mathbf{u})$$



Behavior

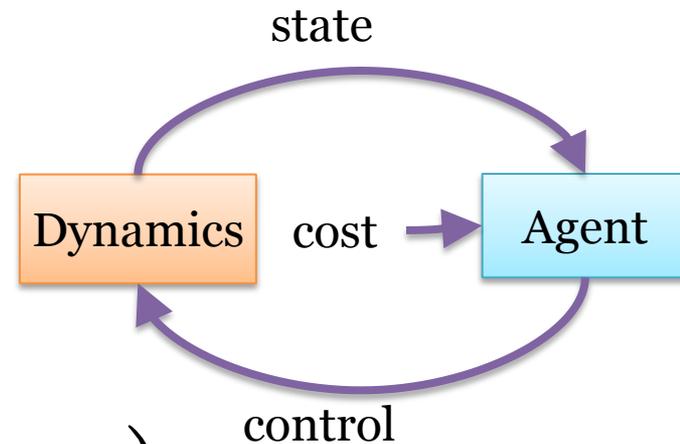
# Trajectory Optimization

# Trajectory Optimization

State  $\mathbf{x}$ , control  $\mathbf{u}$

Dynamics:  $\mathbf{x}' = \mathbf{f}(\mathbf{x}, \mathbf{u})$

Cost:  $\ell(\mathbf{x}, \mathbf{u}), \Omega(\mathbf{x}_{\text{final}})$



# Trajectory Optimization

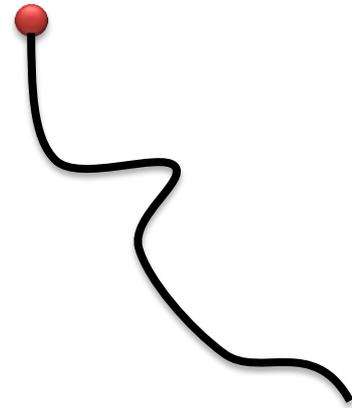
State  $\mathbf{x}$ , control  $\mathbf{u}$

Dynamics:  $\mathbf{x}' = \mathbf{f}(\mathbf{x}, \mathbf{u})$

Cost:  $\ell(\mathbf{x}, \mathbf{u}), \Omega(\mathbf{x}_{\text{final}})$

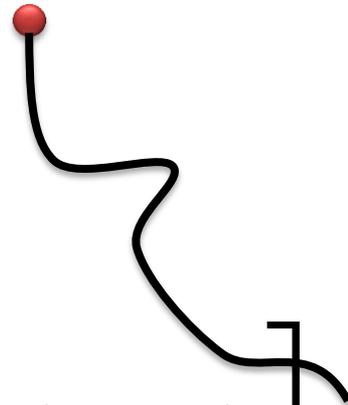
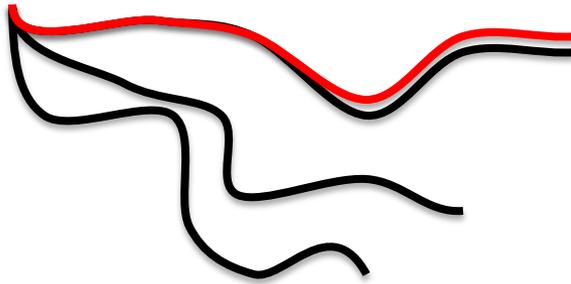
Initial state:  $\mathbf{x}_0$

Control  
sequence:  $\mathbf{u}_1, \mathbf{u}_2 \dots \mathbf{u}_N$



(Single shooting / indirect optimization)

# Trajectory Optimization

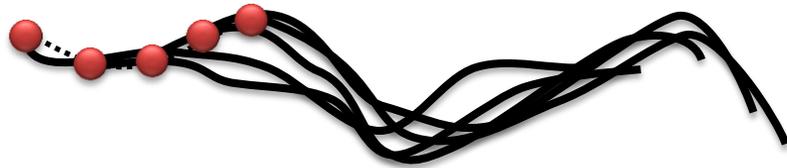


$$\min_{\mathbf{u}_1 \dots \mathbf{u}_N} \left[ \sum_{k=1}^N \ell(\mathbf{x}_{k-1}, \mathbf{u}_k) + \Omega(\mathbf{x}_N) \right]$$

# One Plan is not Enough

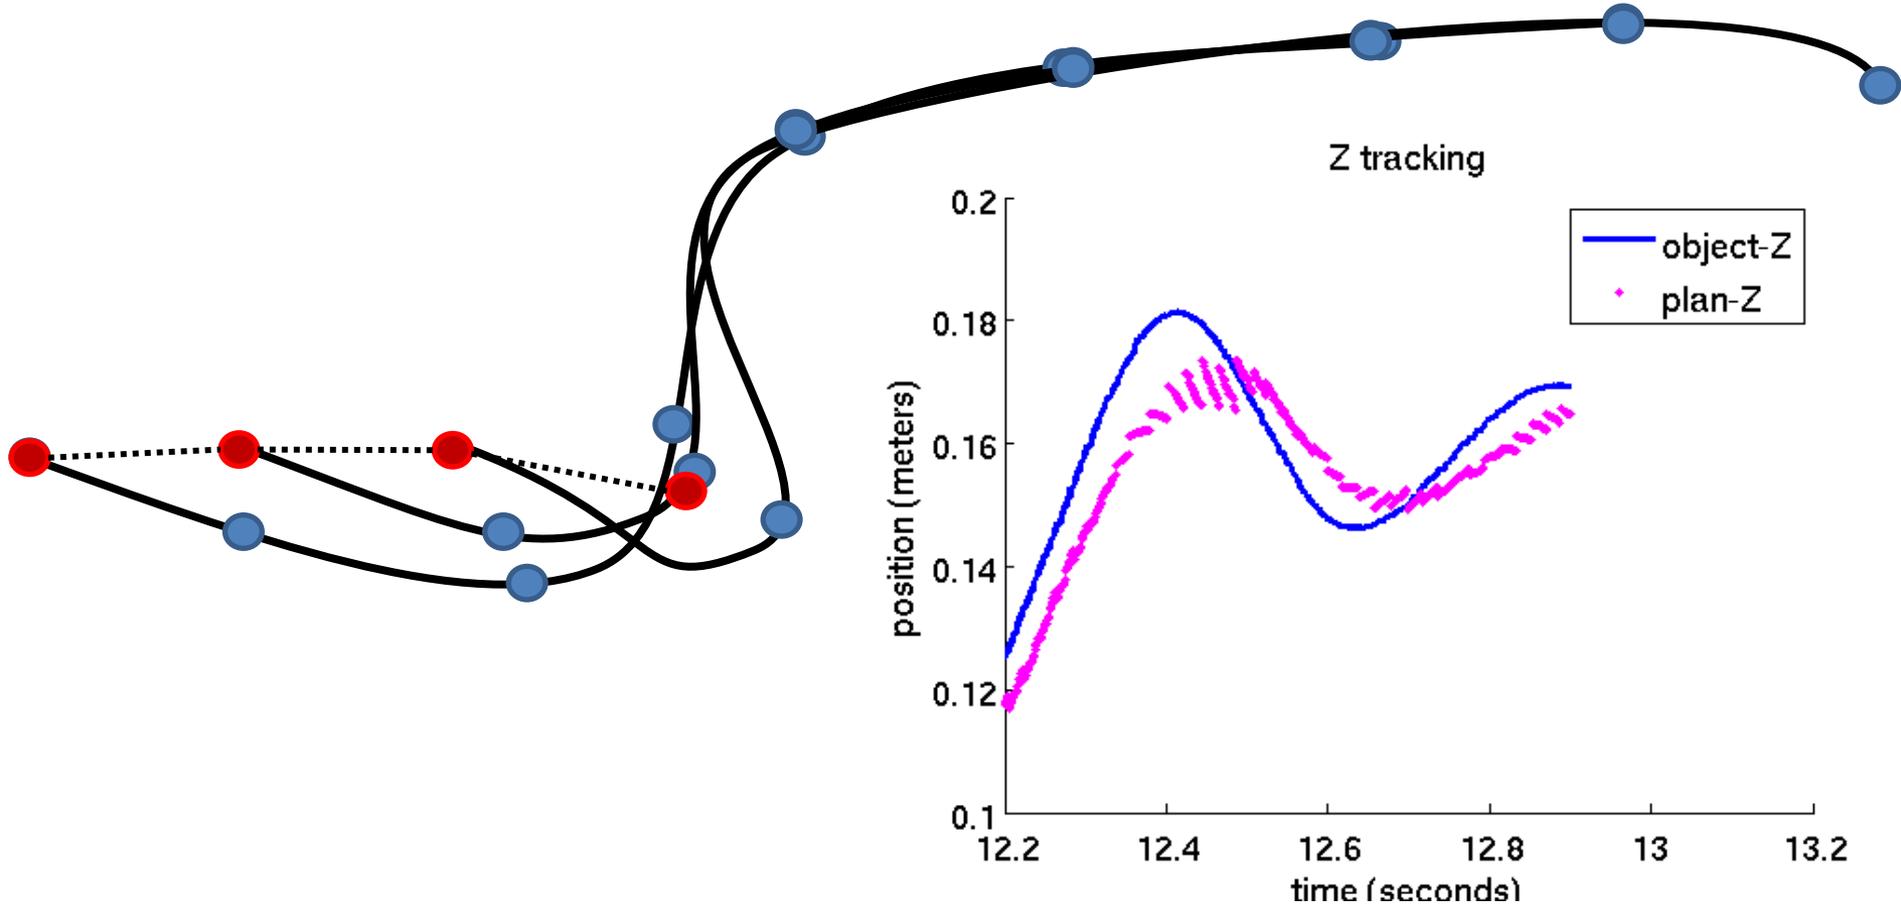
- Modeling errors
- Estimation errors
- Dynamic environment

# Model-Predictive Control

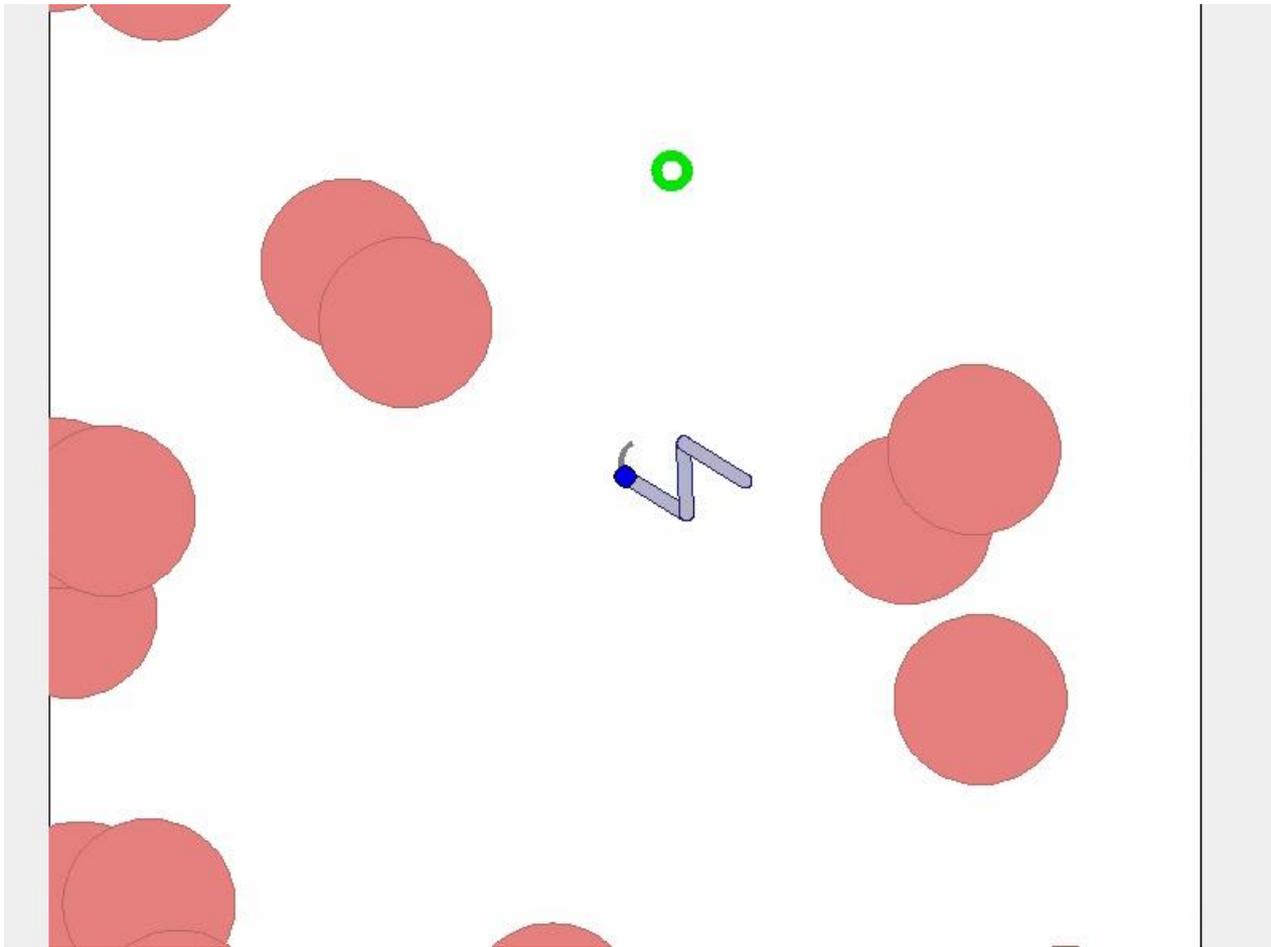


- There is always a plan.
- The plan is constantly updated.
- Only the initial portion is ever executed.

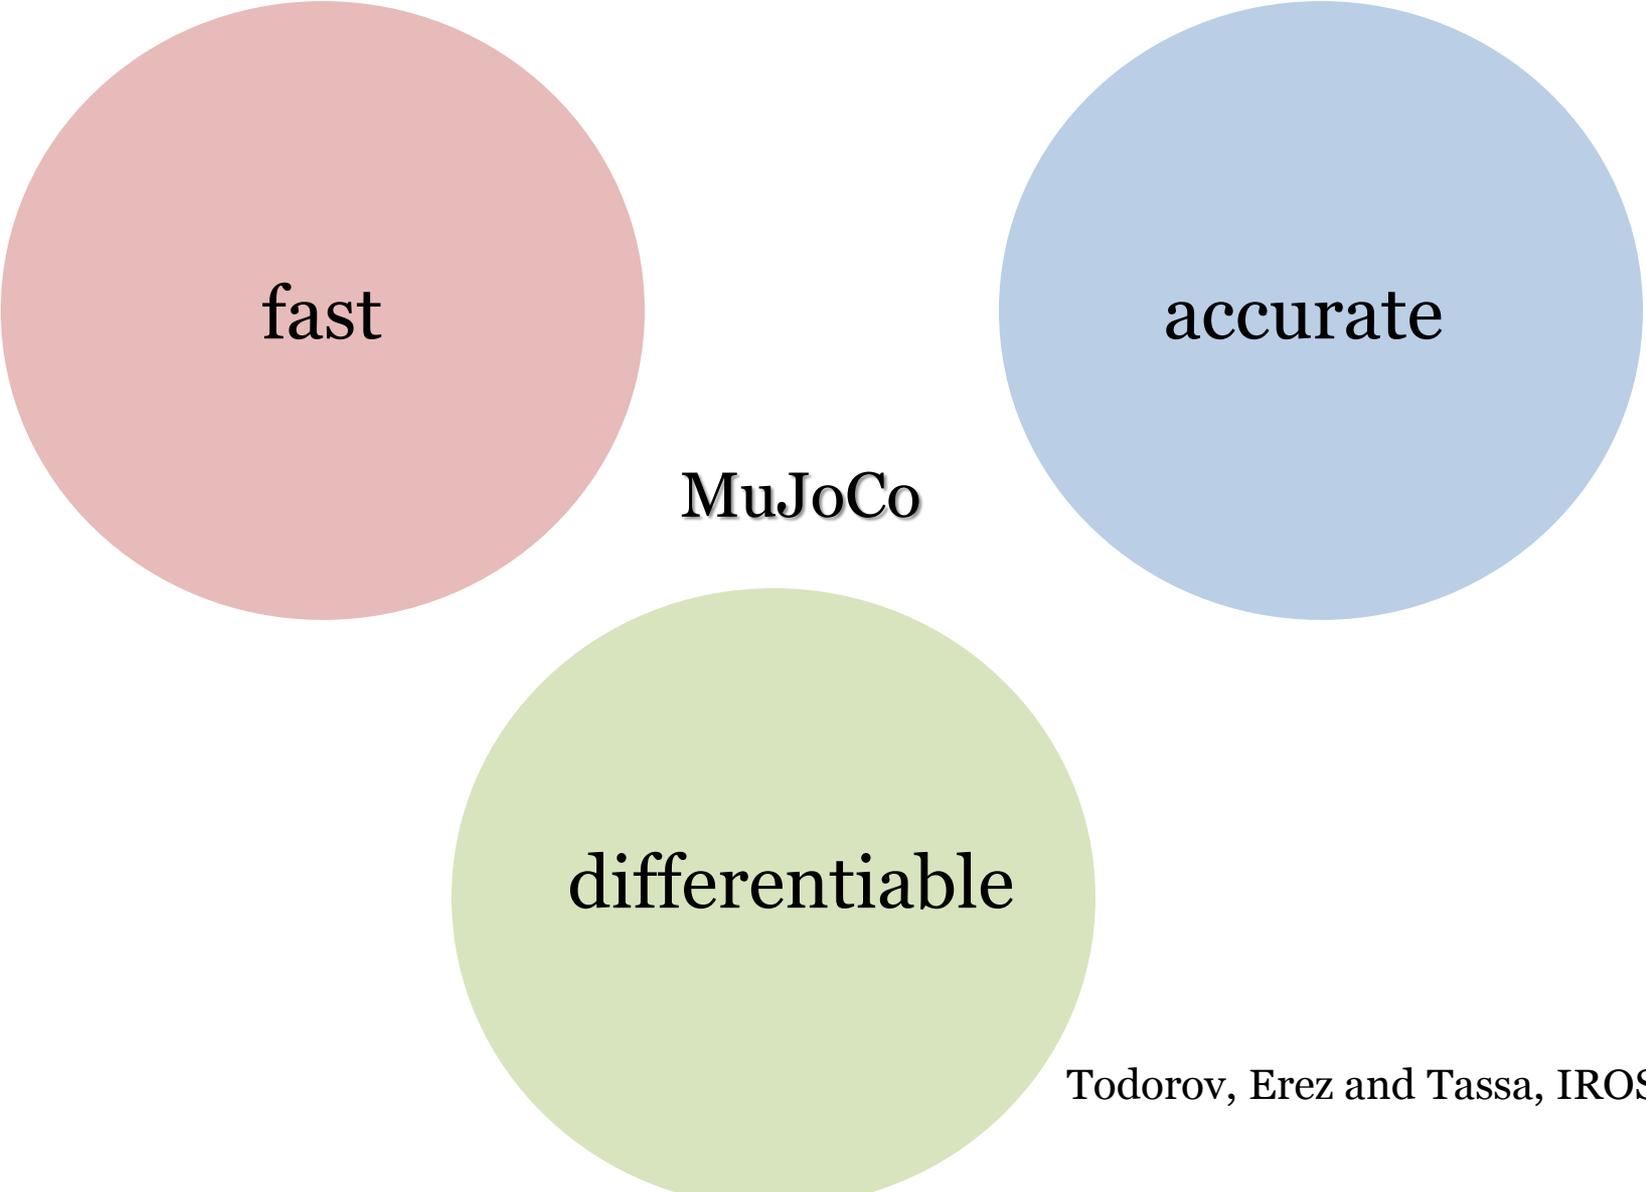
# Modeling Errors



# The Swimmer



# Physics Simulation



fast

accurate

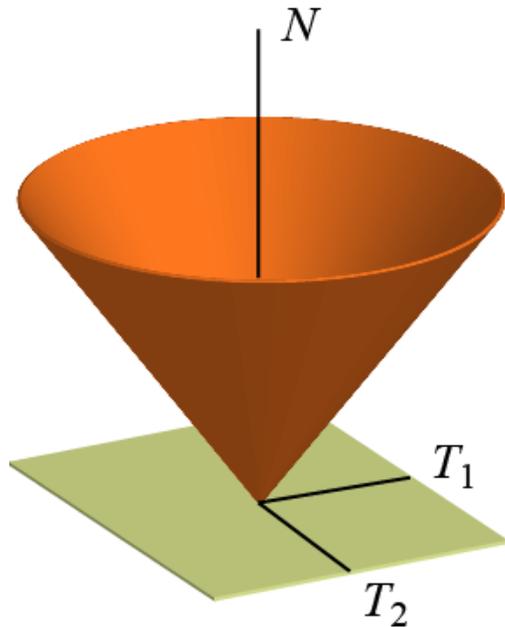
**MuJoCo**

differentiable

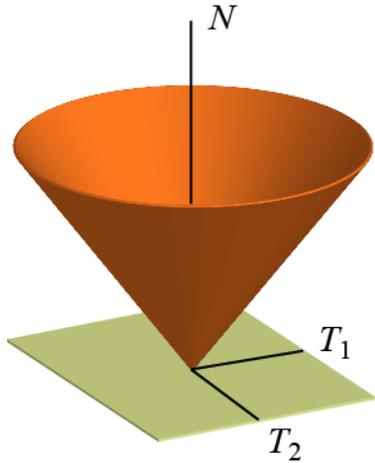
Todorov, Erez and Tassa, IROS 2012

# Contact Modeling

- Non-convex
- Non-smooth
- NP-hard



# Contact dynamics



Coulomb friction (NP hard)

$$v = Af + b$$

$$v_N \perp f_N$$

$$\|v_T\| \perp \mu f_N - \|f_T\|$$

$$v_T = \alpha f_T, \alpha \leq 0$$

$x \perp y$	complementarity: $x \geq 0, y \geq 0, xy = 0$
$A$	contact inverse inertia: $A = JM^{-1}J^T$
$b$	contact velocity before impulse
$f$	<b>contact impulse</b>
$v$	<b>contact velocity after impulse</b>

Define the contact impulse by minimizing contact-space kinetic energy  $\frac{1}{2} v^T A^{-1} v$ , subject to  $v_N \geq 0, f_N \geq 0, \mu f_N \geq \|f_T\|$  for each contact.

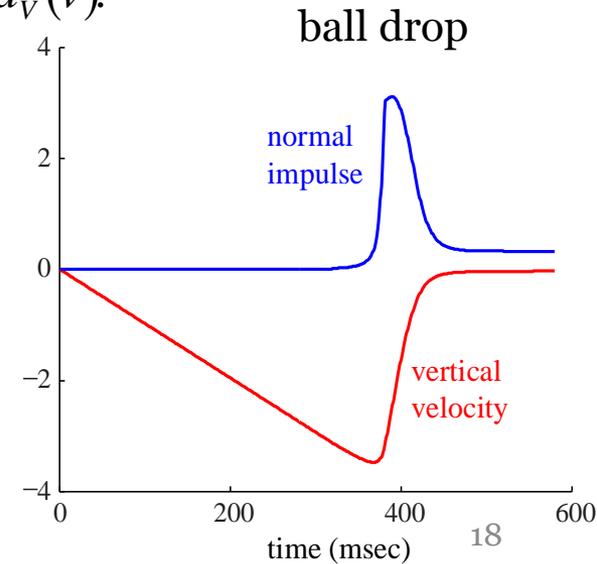
Replace the penetration constraint with a penalty function  $d_v(v)$ .

Forward contact dynamics:  $(A, b) \rightarrow (f, v = Af + b)$

$$f^* = \arg \min_f \frac{1}{2} f^T A f + f^T b + d_v(Af + b)$$

Inverse contact dynamics:  $(A, v) \rightarrow (f, b = v - Af)$

$$f^* = \arg \min_f f^T (v + A \nabla d_v(v))$$



# Timing (Single Core)

One dynamics step  
(35 DOF, 8 contacts)

0.01 ms = 100,000/sec

# Harnessing Parallelization

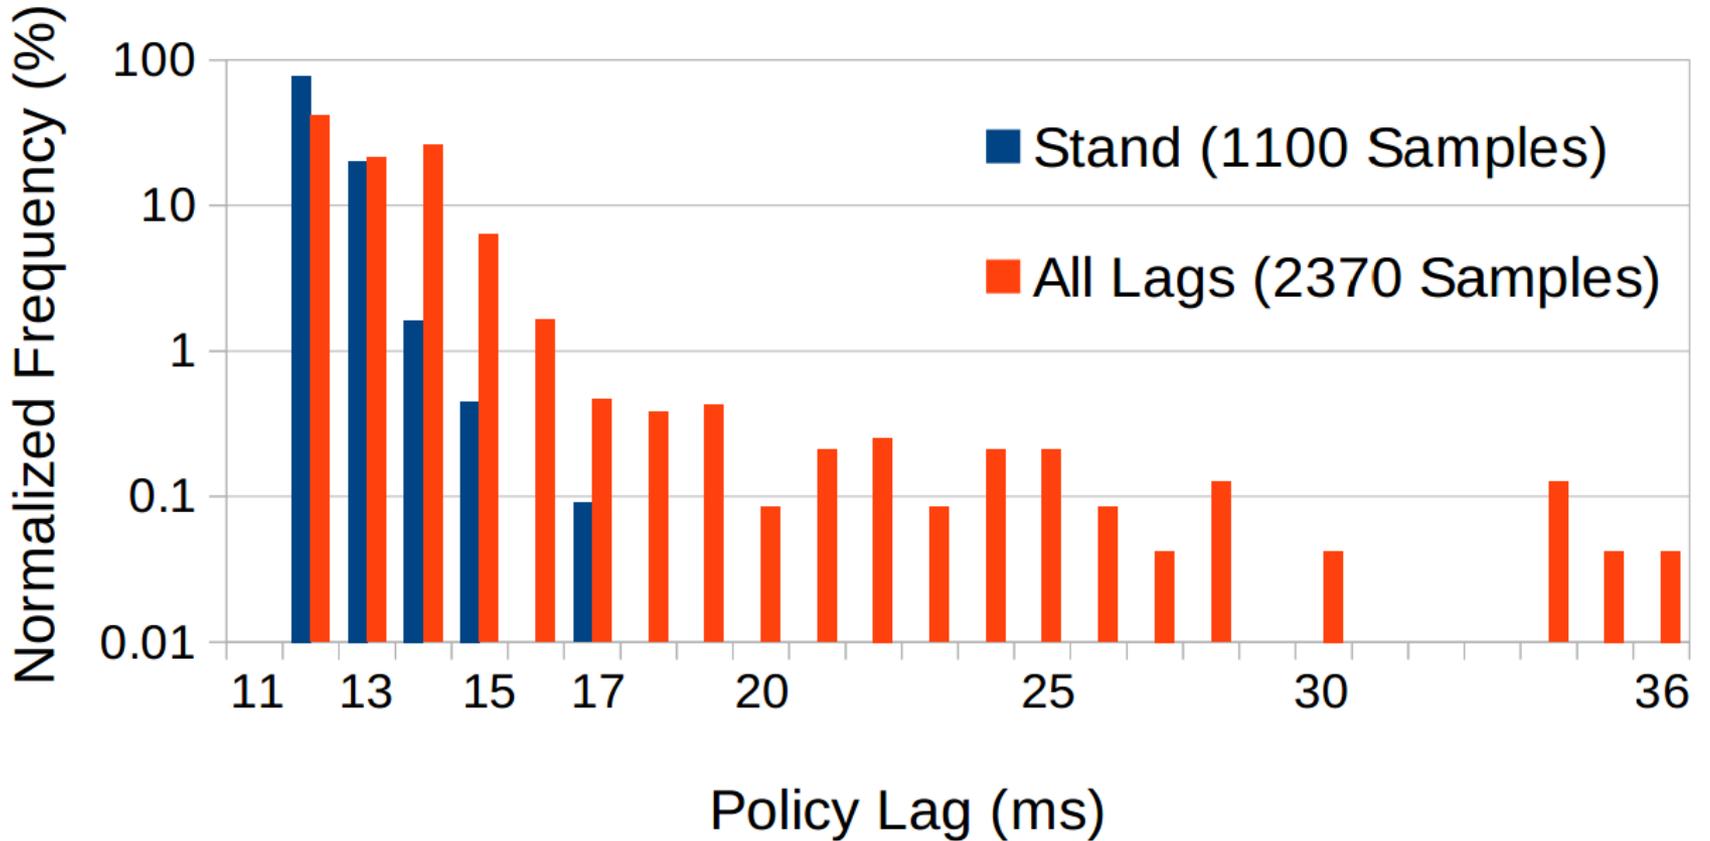
One optimization step

Single i7 core: 9 updates/second

Quad-core i7: 33 updates/second

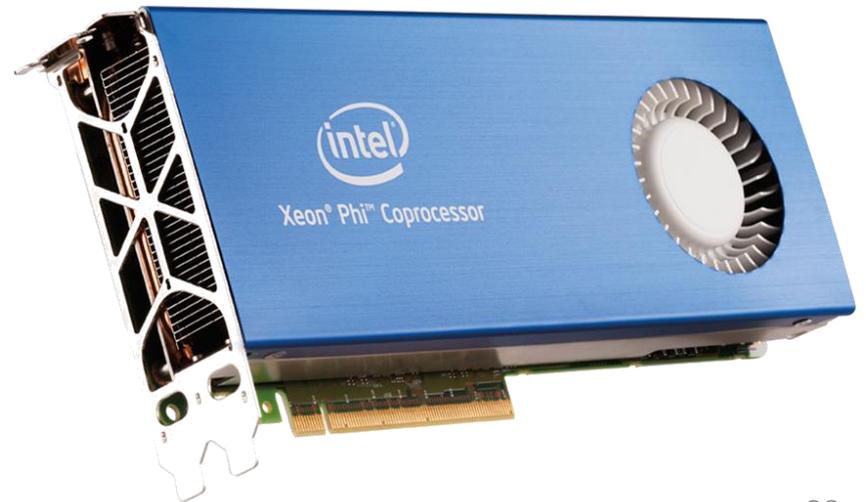
8-core Xeon  
(Amazon EC2): 50 updates/second

## Minimum Policy Lags



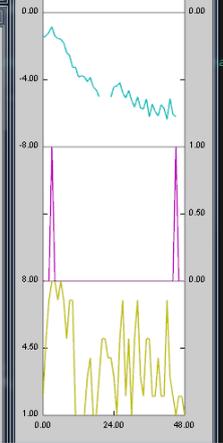
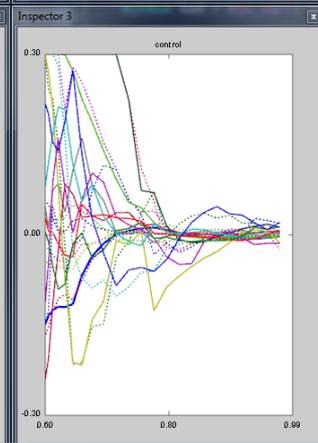
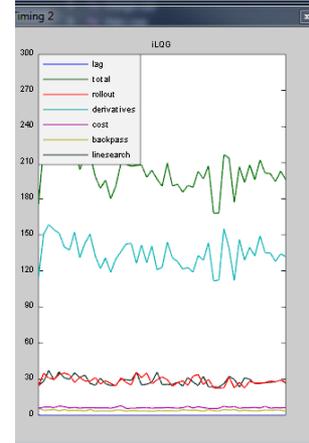
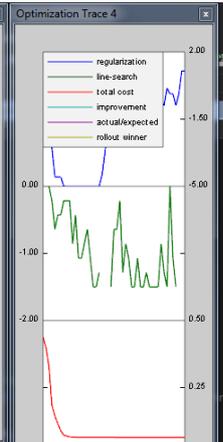
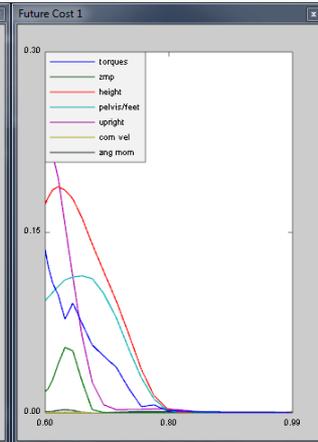
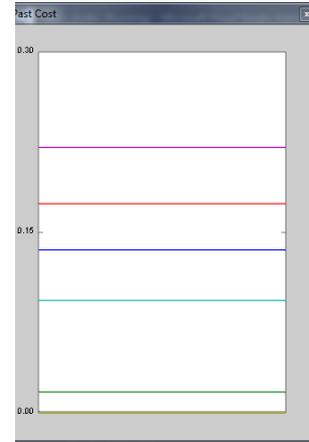
# Harnessing Parallelization

- GPGPU
- Intel's Xeon Phi co-processor



# System integration

- Specification language for models and cost functions, transitions and alterations.
- Visualization for monitoring the system's state – timing, estimation, prediction errors, optimization status.
- GUI – switching between tasks, adjusting weights.



Option window with Physics, Mode, and Style tabs. The Physics tab is active, showing settings for Integrator (Implicit), Collision (Pair), Algorithm (GS), and A-type (Full). The Setting section includes Timestep (0.005), Gravity (0 0 -9.81), Viscosity (0), Wind (0 0 0), EqErReduce (0.01), ImpErReduce (0), EqSoftWeight (100), Soft Clamp (0), Sci Softness (1.1), Exp Dist (3), and Max Iter (5). The Disable section includes checkboxes for Constraint, Impulse, Limit, Contact, Sparse Jac, Passive, Gravity, Bias, Jnt Friction, Clamp Vel, Clamp Frc, Perturbation, Energy, Filter Parent, Cb Control, and Cb Passive.

User window showing MPC, Costs, Running, Final, and Physics tabs. The Running tab is active, showing a table of user-defined parameters.

MPC	Costs	Running	Final	Physics
torques	0.1			
zmp	1			
height	1			
pelvis/feet	1			
upright	1			
com vel	0.01			
ang mom	0.0001			
		not in use		
		not in use		

# Future work

- Estimation with contact
- Combining direct optimization to minimize cost tweaking
- Learning from experience
- “Instincts” – low-level overriding controllers
- Stiff, geared robots
- Pneumatic systems



Emo Todorov

Thank you!



Yuval Tassa