Variational Methods & Optimal Control

lecture 26

Matthew Roughan
<matthew.roughan@adelaide.edu.au>

Discipline of Applied Mathematics
School of Mathematical Sciences
University of Adelaide

April 14, 2016

Pontryagin Maximum Principle

Modern optimal control theory often starts from the PMP. It is a simple, concise condition for an optimal control.

General control problem

Minimize functional

$$F = \int_{t_0}^{t_1} f_0\left(t, \mathbf{x}, \mathbf{u}\right) dt$$

subject to constraints $\dot{\mathbf{x}} = \mathbf{f}(t, \mathbf{x}, \mathbf{u})$, or more fully,

$$\dot{x}_i = f_i(t, \mathbf{x}, \mathbf{u})$$

- notice no dependence on $\dot{\mathbf{x}}$ in f_0
  - this differs from many CoV problems
- no dependence on $\dot{\mathbf{x}}$ in f_i because we rearrange the equations so that derivatives are on the LHS

Pontryagin Maximum Principle (PMP)

Let $\mathbf{u}(t)$ be an admissible control vector that transfers $(t_0, \mathbf{x}_0)$ to a target $(t_1, \mathbf{x}(t_1))$. Let $\mathbf{x}(t)$ be the trajectory corresponding to $\mathbf{u}(t)$. In order that $\mathbf{u}(t)$ be optimal, it is necessary that there exists $\mathbf{p}(t) = (p_1(t), p_2(t), \ldots, p_n(t))$ and a constant scalar p_0 such that

- $\mathbf{p}$ and $\mathbf{x}$ are the solution to the canonical system

$$\dot{\mathbf{x}} = \frac{\partial H}{\partial \mathbf{p}} \quad \text{and} \quad \dot{\mathbf{p}} = -\frac{\partial H}{\partial \mathbf{x}}$$

- where the Hamiltonian is $H = \sum_{i=0}^{n} p_i f_i$ with $p_0 = -1$
- $H(\mathbf{x}, \mathbf{u}, \mathbf{p}, t) \geq H(\mathbf{x}, \hat{\mathbf{u}}, \mathbf{p}, t)$ for all alternate controls $\hat{\mathbf{u}}$
- all boundary conditions are satisfied

PMP proof sketch

Consider the general problem: minimize functional

$$F\{\mathbf{x},\mathbf{u}\} = \int_{t_0}^{t_1} f_0\,(t,\mathbf{x},\mathbf{u})\,dt$$

subject to constraints

$$\dot{x}_i = f_i(t,\mathbf{x},\mathbf{u})$$

We can incorporate the constraints into the functional using the Lagrange multipliers λ_i, e.g.

$$\tilde{F} = \int_{t_0}^{t_1} L(t,\mathbf{x},\dot{\mathbf{x}},\mathbf{u})\,dt = \int_{t_0}^{t_1} f_0\,(t,\mathbf{x},\mathbf{u}) + \sum_{i=1}^{n} \lambda_i(t)\left[\dot{x}_i - f_i(t,\mathbf{x},\mathbf{u})\right]\,dt$$

PMP proof sketch

Given such a function we get (by definition)

$$p_i = \frac{\partial L}{\partial \dot{x}_i} = \lambda_i$$

So we can identify the Lagrange multipliers λ_i with the **generalized momentum** terms p_i

- the p_i are known in economics literature as **marginal valuation** of x_i or the **shadow prices**
- shows how much a unit increment in x at time t contributes to the optimal objective functional $\tilde{F}$
- the p_i are known in control as **co-state variables** (sometimes written as z_i)

PMP proof sketch

By definition (in previous lectures) the Hamiltonian is

$$\begin{aligned} H(t,\mathbf{x},\mathbf{p},\mathbf{u}) &= \sum_{i=1}^{n} p_i \dot{x}_i - L(t,\mathbf{x},\dot{\mathbf{x}},\mathbf{u}) \\ &= \sum_{i=1}^{n} p_i \dot{x}_i - f_0\left(t,\mathbf{x},\mathbf{u}\right) - \sum_{i=1}^{n} \lambda_i(t)\left[\dot{x}_i - f_i(t,\mathbf{x},\mathbf{u})\right] \\ &= -f_0\left(t,\mathbf{x},\mathbf{u}\right) + \sum_{i=1}^{n} p_i f_i(t,\mathbf{x},\mathbf{u}) \end{aligned}$$

because $\lambda_i = p_i$, so the $\dot{x}_i$ terms cancel. The final result is just the Hamiltonian as defined in the PMP.

PMP proof sketch

From previous slide the Hamiltonian can be written

$$H(t,\mathbf{x},\mathbf{p},\mathbf{u}) \quad = \quad -f_0(t,\mathbf{x},\mathbf{u}) + \sum_{i=1}^{n} p_i f_i(t,\mathbf{x},\mathbf{u})$$

which is the Hamiltonian defined in the PMP. Then the Canonical E-L equations (Hamilton's equations) are

$$\frac{\partial H}{\partial p_i} = \frac{dx_i}{dt} \quad \text{and} \quad \frac{\partial H}{\partial x_i} = -\frac{dp_i}{dt}$$

Note that the equations $\frac{\partial H}{\partial p_i} = \frac{dx_i}{dt}$ just revert to

$$f_i(t,\mathbf{x},\mathbf{u}) = \dot{x}_i$$

which are just the system equations.

PMP proof sketch

Finally, note that Hamilton's equations above only relate x_i and its conjugate momentum p_i. What about equations for u_i? Take the conjugate variable to be z_i, and we get (by definition) that

$$z_i = \frac{\partial L}{\partial \dot{u}_i} = 0$$

and the second of Hamilton's equations is therefore

$$\frac{\partial H}{\partial u_i} = -\frac{dz_i}{dt} = 0$$

which suggests a stationary point of H WRT u_i. In fact we look for a maximum (and note this may happen on the bounds of u_i)

PMP Example: plant growth

Plant growth problem:

- market gardener wants to plants to grow to a fixed height 2 within a fixed window of time $[0, 1]$
- can supplement natural growth with lights (at night)
- growth rate dictates

$$\dot{x} = 1 + u$$

- cost of lights

$$F\{u\} = \int_0^1 \frac{1}{2} u^2 \, dt$$

PMP Example: plant growth

Minimize

$$F\{u\} = \int_0^1 \frac{1}{2} u^2 \, dt$$

Subject to $x(0) = 0$, and $x(1) = 2$ and

$$\dot{x} = f_1(t, x, u) = 1 + u$$

Hamiltonian is

$$\begin{aligned} H &= -f_0(t, x, u) + p f_1(t, x, u) \\ &= -\frac{1}{2} u^2 + p(1 + u) \end{aligned}$$

PMP Example: plant growth

Hamiltonian is

$$H = -\frac{1}{2}u^2 + p(1+u)$$

Canonical equations

$$\begin{array}{rclcrcl} \dfrac{\partial H}{\partial p} & = & \dfrac{dx}{dt} & \quad\text{and}\quad & \dfrac{\partial H}{\partial x} & = & -\dfrac{dp}{dt} \\ & \Downarrow & & & & \Downarrow & \\ 1+u & = & \dot{x} & & 0 & = & -\dot{p} \end{array}$$

LHS =¿ system DE

RHS =¿ $\dot{p} = 0$ means that $p = c_1$ where c_1 is a constant.

PMP Example: plant growth

Maximum principle requires H be a maximum, for which

$$\frac{\partial H}{\partial u} = -u + p = 0$$

So $u = p$, and $\dot{x} = 1 + u$ so

$$x = (1 + c_1)t + c_2$$

The solution which satisfies $x(0) = 0$ and $x(1) = 2$ is

$$x = 2t$$

So $u = c_1 = 1$, and the optimal cost is $1/2$.

PMP and Transversal conditions

Typically we fix t_0 and $\mathbf{x}(t_0)$, but often the right-hand boundary condition is not fixed, so we need transversal, or natural boundary conditions. Here, they differ from traditional CoV problems in two respects:

- The terminal cost ϕ
- The function f_0 is not explicitly dependent on $\dot{x}$

The resulting transversal conditions are

$$\sum_i \left(\frac{\partial \phi}{\partial x_i} + p_i \right) \delta x_i \bigg|_{t=t_1} + \left(\frac{\partial \phi}{\partial t} - H \right) \delta t \bigg|_{t=t_1} = 0$$

for all allowed δx_i and δt.

PMP and Transversal conditions

The resulting transversal condition is

$$\sum_i \left(\frac{\partial \phi}{\partial x_i} + p_i \right) \delta x_i \bigg|_{t=t_1} + \left(\frac{\partial \phi}{\partial t} - H \right) \delta t \bigg|_{t=t_1} = 0$$

Special cases

- when t_1 is fixed and $\mathbf{x}(t_1)$ is completely free we get

$$\left(\frac{\partial \phi}{\partial x_i} + p_i \right) \bigg|_{t=t_1} = 0, \quad \forall i$$

- when $\mathbf{x}(t_1)$ is fixed, $\delta x_i = 0$, and we get

$$\left(\frac{\partial \phi}{\partial t} - H \right) \bigg|_{t=t_1} = 0$$

Example: stimulated plant growth

Plant growth problem:

- market gardener wants to plants to grow as much as possible within a fixed window of time $[0,1]$
- supplement natural growth with lights as before
- growth rate dictates $\dot{x} = 1 + u$
- cost of lights

$$F\{u\} = \int_0^1 \frac{1}{2} u(t)^2 \, dt$$

- value of crop is proportional to the height

$$\phi(t_1, \mathbf{x}(t_1)) = x(t_1)$$

Plant growth problem statement

Write as a minimization problem

$$F\{u,x\} = -x(t_1) + \int_0^1 \frac{1}{2} u^2 \, dt$$

Subject to $x(0) = 0$,

$$\dot{x} = 1 + u$$

- the terminal cost doesn't affect the shape of the solution
- but we need a natural end-point condition for t_1

Plant growth: natural boundary cond.

The problem is solved as before, but we write the natural boundary condition at $x = t_1$ as

$$\left(\frac{\partial \phi}{\partial x_i} + p_i \right) \bigg|_{t=t_1} = 0, \quad \forall i$$

which reduces to

$$-1 + p\big|_{t=t_1} = 0$$

Given p is constant, this sets $p(t) = 1$, and hence the control $u = 1$ (as before).

Autonomous problems

Autonomous problems have no explicit dependence on t.

- time invariance symmetry
- hence H is constant along the optimal trajectory
- if the end-time is free (and the terminal cost is zero) then the transversality conditions ensure $H = 0$ along the optimal trajectory.

PMP Example: Gout

Optimal Treatment of Gout:

- disease characterized by excess of uric acid in blood
  - define level of uric acid to be $x(t)$
  - in absence of any control, tends to 1 according to

$$\dot{x} = 1 - x$$

- drugs are available to control disease (control u)

$$\dot{x} = 1 - x - u$$

  - aim to reduce x to zero as quickly as possible
  - drug is expensive, and unsafe (side effects)

PMP Example: Gout

Formulation: Minimize

$$F\{u\} = \int_0^{t_1} \frac{1}{2}(k^2 + u^2)\, dt$$

given constant k that measures the relative importance of the drugs cost vs the terminal time. End-conditions are $x(0) = 1$, and we wish $x(t_1) = 0$, with t_1 free. The constraint equation is

$$\dot{x} = 1 - x - u$$

Hamiltonian

$$H = -\frac{1}{2}(k^2 + u^2) + p(1 - x - u)$$

PMP Example: Gout

Canonical equations

$$\begin{array}{rclcrcl} \dfrac{\partial H}{\partial p} & = & \dfrac{dx}{dt} & \quad \text{and} \quad & \dfrac{\partial H}{\partial x} & = & -\dfrac{dp}{dt} \\ & \Downarrow & & & & \Downarrow & \\ 1-x-u & = & \dot{x} & & -p & = & -\dot{p} \end{array}$$

LHS =¿ system DE

RHS =¿ $\dot{p} = p$ has solution $p = c_1 e^t$

Now maximize H wrt to u, i.e., find stationary point

$$\frac{\partial H}{\partial u} = -u - p = 0$$

So $u = -p = -c_1 e^t$

PMP Example: Gout

Note

- this is an autonomous problem so $H = const$
- this is a free end-time problem so $H = 0$

Substitute values of p and u into H for $t = 0$ (i.e. $p = c_1 = -u$, and $x(0) = 1$), and we get

$$\begin{aligned} H &= -\frac{1}{2}(k^2 + u^2) + p(1 - x - u) \\ &= -\frac{k^2}{2} - \frac{c_1^2}{2} - c_1^2 \\ &= 0 \end{aligned}$$

and so $c_1 = \pm k$

PMP Example: Gout

Finally solve $\dot{x} = 1 - x - u$ where $u = -ke^t$ to get

$$x = 1 - \frac{k}{2}e^t + \frac{k}{2}e^{-t} = 1 - k\sinh t$$

The terminal condition is $x(t_1) = 0$, and so

$$t_1 = \sinh^{-1}(1/k)$$

- when k is small the prime consideration is to use a small amount of the drug, and as $k \to 0$ then $t_1 \to \infty$
  - no optimal for $k = 0$
- when k is large, we want to get to a safe level as fast as possible, so as $k \to \infty$ we get $t_1 \sim 1/k$

PMP Example: Lunar lander

Atari game, 1979

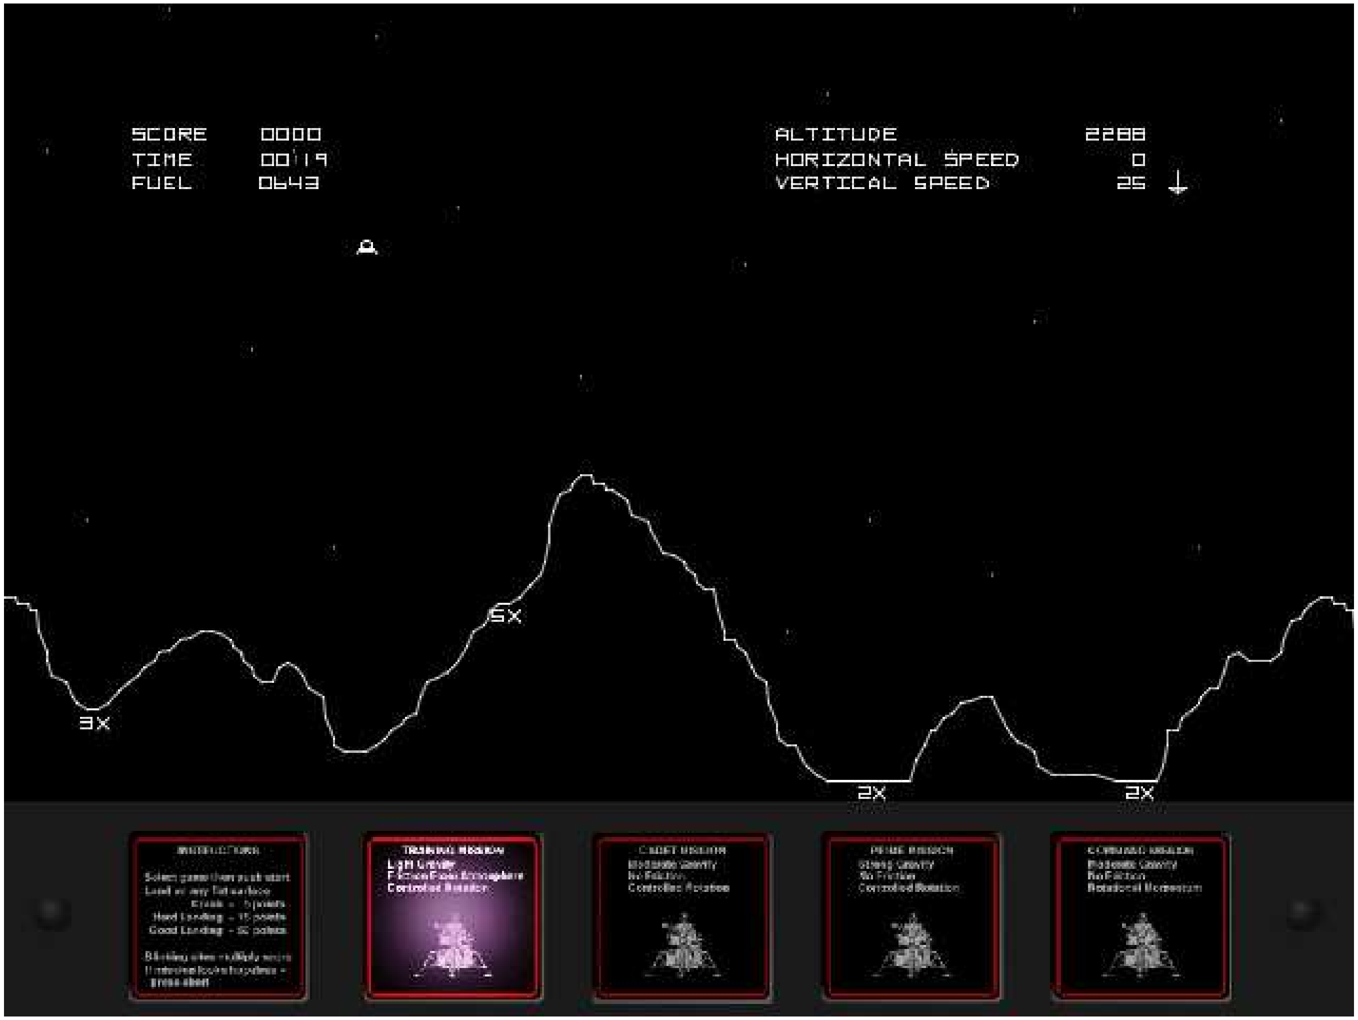

http://www.klov.com/game_detail.php?letter=L&game_id=8465

PMP Example: Lunar lander

- need to land surface-module on the moon
  - Module mass M (ignore fuel load), uniform gravitational acceleration g (might not be $9.8m/s^2$)
  - initial height $y(0) = h$
  - initial velocity $\dot{y}(0) = v$
- controlled descent so landing is "soft"
  - height of module, and downward velocity brought to zero simultaneously
- thrust f either up or down
  - thrust is bounded, so $|f| \leq f_{\max}$
  - want to minimize fuel cost $|f|$ over time

PMP Example: Lunar lander

System defined (at any time t) by

- position y
- velocity $\dot{y}$

State equations (mass $\times$ acceleration = force)

$$M\ddot{y} = -Mg + f$$

Initial state

$$y(0) = h, \quad \text{and} \quad \dot{y}(0) = v$$

Desired final state (t_1 is free)

$$y(t_1) = 0 \quad \text{and} \quad \dot{y}(t_1) = 0$$

and we wish to minimize

$$F\{f\} = \int_0^{t_1} |f|\, dt$$

PMP Example: Lunar lander

Convert the problem to standard form by taking

$$\begin{aligned} x_1 &= y \\ x_2 &= \dot{y} \\ u &= f/M \end{aligned}$$

So the state equation becomes

$$\begin{aligned} \dot{x}_1 &= x_2 \\ \dot{x}_2 &= -g+u \end{aligned}$$

And the initial and final conditions are

$$\begin{aligned} x_1(0) = h \quad &\text{and} \quad x_2(0) = v \\ x_1(t_1) = 0 \quad &\text{and} \quad x_2(t_1) = 0 \end{aligned}$$

PMP Example: Lunar lander

Hamiltonian

$$H = -|u| + p_1 x_2 + p_2(u - g)$$

Canonical equations

$$\frac{\partial H}{\partial p_i} = \frac{dx_i}{dt} \quad \text{and} \quad \frac{\partial H}{\partial x_i} = -\frac{dp_i}{dt}$$

Give the constraints $\dot{x}_1 = x_2$ and $\dot{x}_2 = -g + u$ and

$$\begin{aligned} \frac{\partial H}{\partial x_1} &= 0 = -\dot{p}_1 \\ \frac{\partial H}{\partial x_2} &= p_1 = -\dot{p}_2 \end{aligned}$$

Solution $p_1 = c_1$ and $p_2 = -c_1 t + c_2$.

PMP Example: Lunar lander

Now we have to choose u to maximize H

- $|u|$ is bounded by $f_{\max}/M$

Ignore the terms in H that are constant WRT to u and we have to maximize $-|u| + p_2 u$.

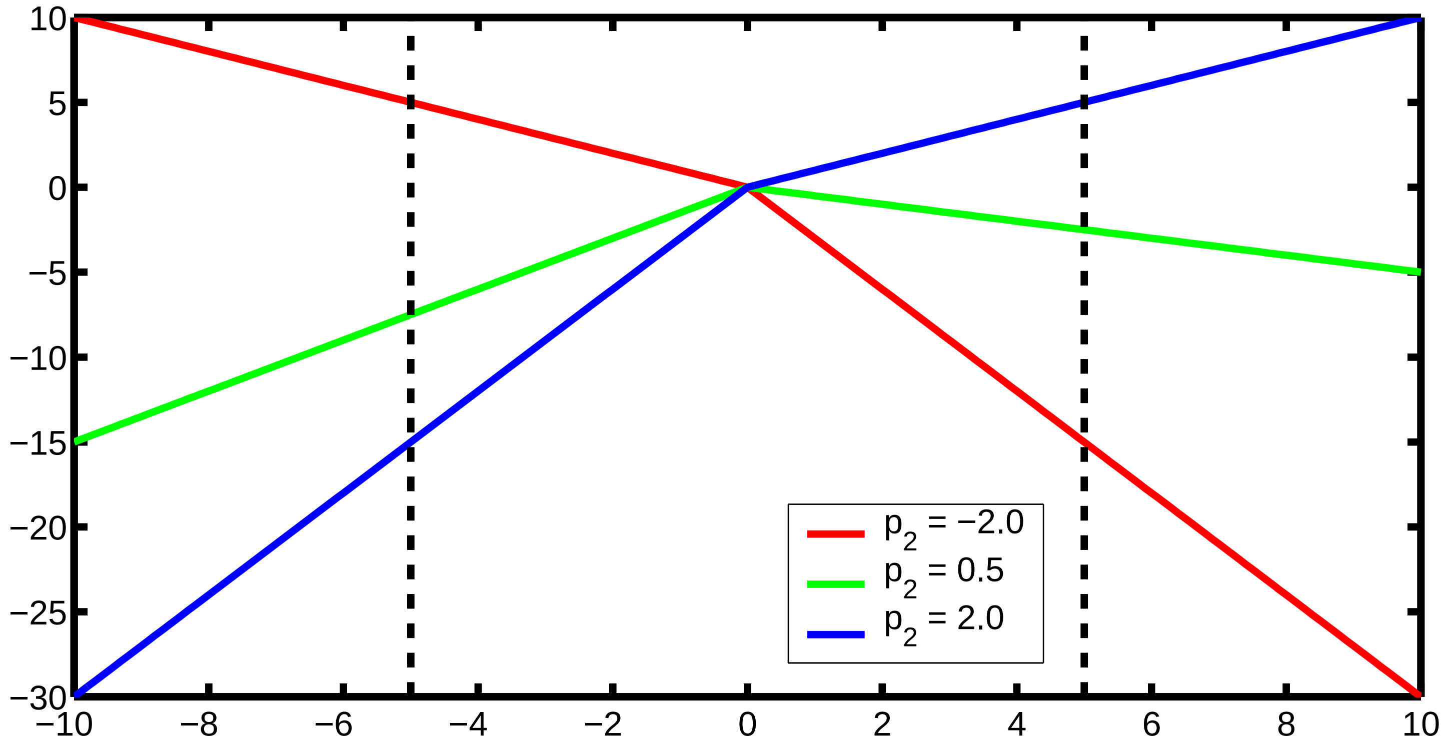

PMP Example: Lunar lander

Maximize $f(u) = -|u| + p_2 u$, with $|u| \leq 1$

- three possible locations for a maximum
  - left or right boundary, or $u = 0$
- The three values (in order from left to right) are

$$f(u) = -1 - p_2, \quad 0, \quad -1 + p_2$$

- Three cases $p_2 < -1$, $-1 < p_2 < 1$ or $p_2 > 1$
- maximum occurs at

$$u = \begin{cases} +1, & \text{if } p_2 > 1 \\ 0, & \text{if } -1 < p_2 < 1 \\ -1, & \text{if } p_2 < -1 \end{cases}$$

- If bounds are $|u| \leq f_{\max}/M$, then the solution scales.

PMP Example: Lunar lander

Call p_2 a switching function, and note that we have

$$p_2 = -c_1 t + c_2$$

- during the final descent, $x_2 < 0$
  - we must be going down just before we land
- but $x_2(t_1) = 0$, so $\dot{x}_2 > 0$ near t_1
  - we must be decelerating, so that we stop at t_1
  - hence we must have positive thrust
  - optimal thrust must be at max, e.g. $u = f_{\max}/M$
- so the equations for motion during final descent are

$$\begin{aligned} \dot{x}_1 &= x_2 \\ \dot{x}_2 &= -g + f_{\max}/M = k > 0 \end{aligned}$$

PMP Example: Lunar lander

Given final conditions the solution near landing is

$$x_1 = \frac{1}{2}k(t - t_1)^2 \quad \text{and} \quad x_2 = k(t - t_1)$$

- note $k > 0$ in final stages of landing
- note $u = f_{\max}/M$ in final stages of landing
- given $p_2 = -c_1 t + c_2$ we must have $c_1 < 0$
- hence prior stages of control include
  - a stage when $u = 0$ (free fall)
  - a stage when $u = -f_{\max}/M$ (accelerating down)
- in each stage we get an equation as above, but with different constant k, for $u = 0$ and $u = -f_{\max}/M$ the constant $k < 0$

PMP Example: Lunar lander

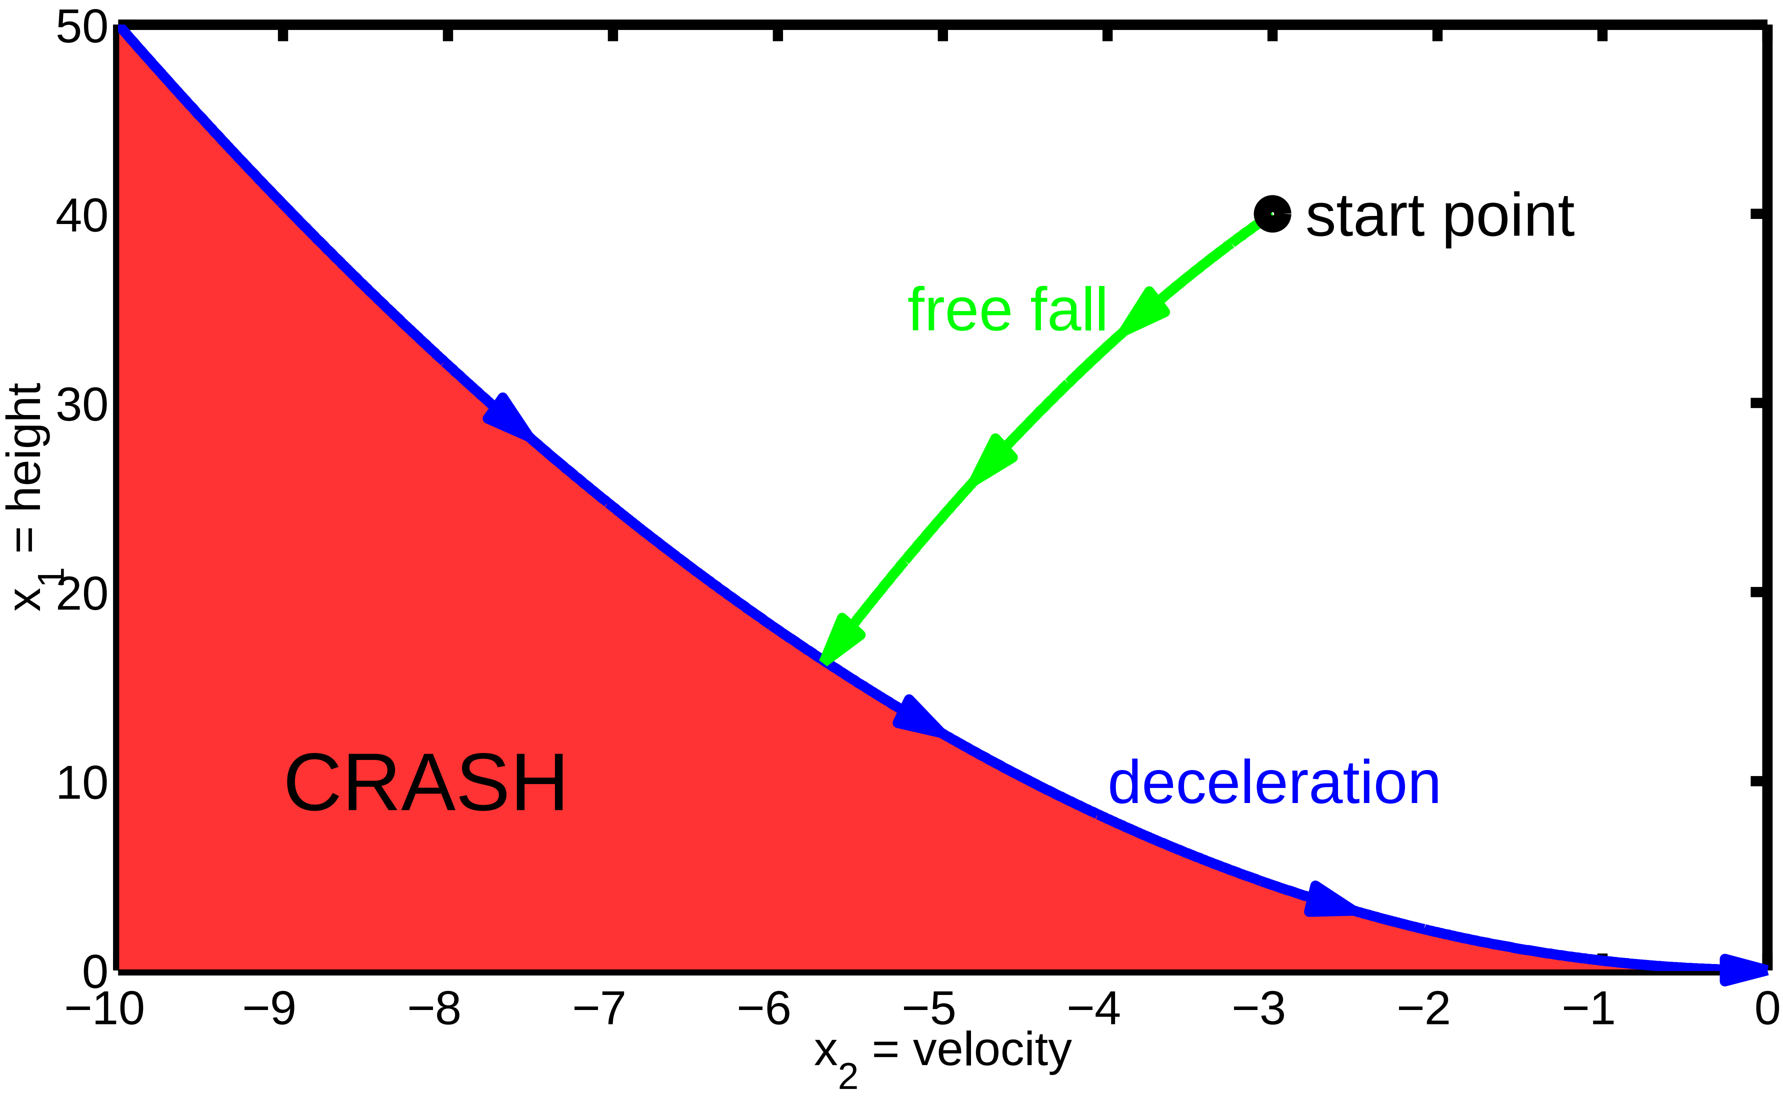

PMP Example: Lunar lander

Solution:

- if start above, or on the critical curve
  - if travelling upwards, max thrust down to cancel upwards velocity
  - then free-fall, until on the critical curve

$$x_1 = \frac{1}{2}k(t-t_1)^2 \quad \text{and} \quad x_2 = k(t-t_1)$$

  - max thrust up until stop on the surface
- if lie below the critical curve
  - you crash

Remember: this is rocket science!

PMP Example: Lunar lander

- What's the point of this example
  - previously, we couldn't easily deal with and objective like $|u|$
  - the function isn't "smooth"
  - PMP can work for such examples
  - it doesn't require smoothness, you just need to be able to find a maximum