Leapfrog Algorithm for Solving Optimal Control Problems

Workshop on Challenges of
Continuous Optimization
Rhodes, Greece
2-3 July 2004

C. Yalçın Kaya
University of South Australia
&
Universidade Federal do Rio de Janeiro, Brazil

Joint work with

J. Lyle Noakes

University of Western Australia

Outline

1 Optimal control problem and existing solution techniques

2 Leapfrog algorithm

3 Uniqueness of controls and costates

4 Continuity of cost

5 Midpoint maps and total cost

6 Compactness of state-costate space

7 Extreme points and leapfrog splicing

8 Convergence to a critical trajectory

The Problem

$$(P)\begin{cases} \text{minimize} & \displaystyle\int_{t_0}^{t_f} f_0(x(t),u(t))\,dt \\ \text{subject to} & \dot{x}(t) = f(x(t),u(t)), \quad u \in \mathcal{U}\,, \\ & x(t_0) = x_0 \quad \text{and} \quad x(t_f) = x_f\,. \end{cases}$$

- *State* $x(t) \in \mathbb{R}^n$, *control* $u(t) \in \mathbb{R}^m$, and f is C^1 vector field.

- $f_0 : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}$ is C^1 and nonnegative, and *final time* t_f is *free*.

Necessary conditions of optimality

The Hamiltonian:

$$H(x, \psi, u) = f_0(x, u) + \langle \psi, f(x, u) \rangle$$

- $\psi(t) \in \mathbb{R}^n$: costate (or adjoint) variable.

If $x(\cdot)$ and $u(\cdot)$ are optimal, then a nontrivial $\psi(\cdot)$ exists such that, for every $t \in [t_0, t_f]$,

$$\dot{x} = \partial H / \partial \psi = f(x, u) \tag{1}$$

$$\dot{\psi} = -\partial H / \partial x \tag{2}$$

$$H(x, \psi, u) = \min_{v \in \mathcal{U}} H(x, \psi, v) \tag{3}$$

$$H(x, \psi, u) \equiv 0 \tag{4}$$

with $x(t_0) = x_0$ and $x(t_f) = x_f$. We call (1)-(4) the optimality system.

- We call the trajectory pair $(x(\cdot), \psi(\cdot))$ corresponding to some $u(\cdot)$ a critical trajectory if it solves the optimality system.

- Note that a critical trajectory is **not** necessarily optimal.

Existing Solution Techniques

- Direct methods: Discretizes the problem, deals with the resulting (large scale) finite optimization problem

- TPBVPs tackled by a combination of
  - Multiple shooting
  - Collocation methods
  - Continuation (or homotopy) methods

We propose the leapfrog algorithm along the second line of techniques, in particular as an alternative to multiple shooting.

Leapfrog Algorithm

We first partition a given feasible trajectory:

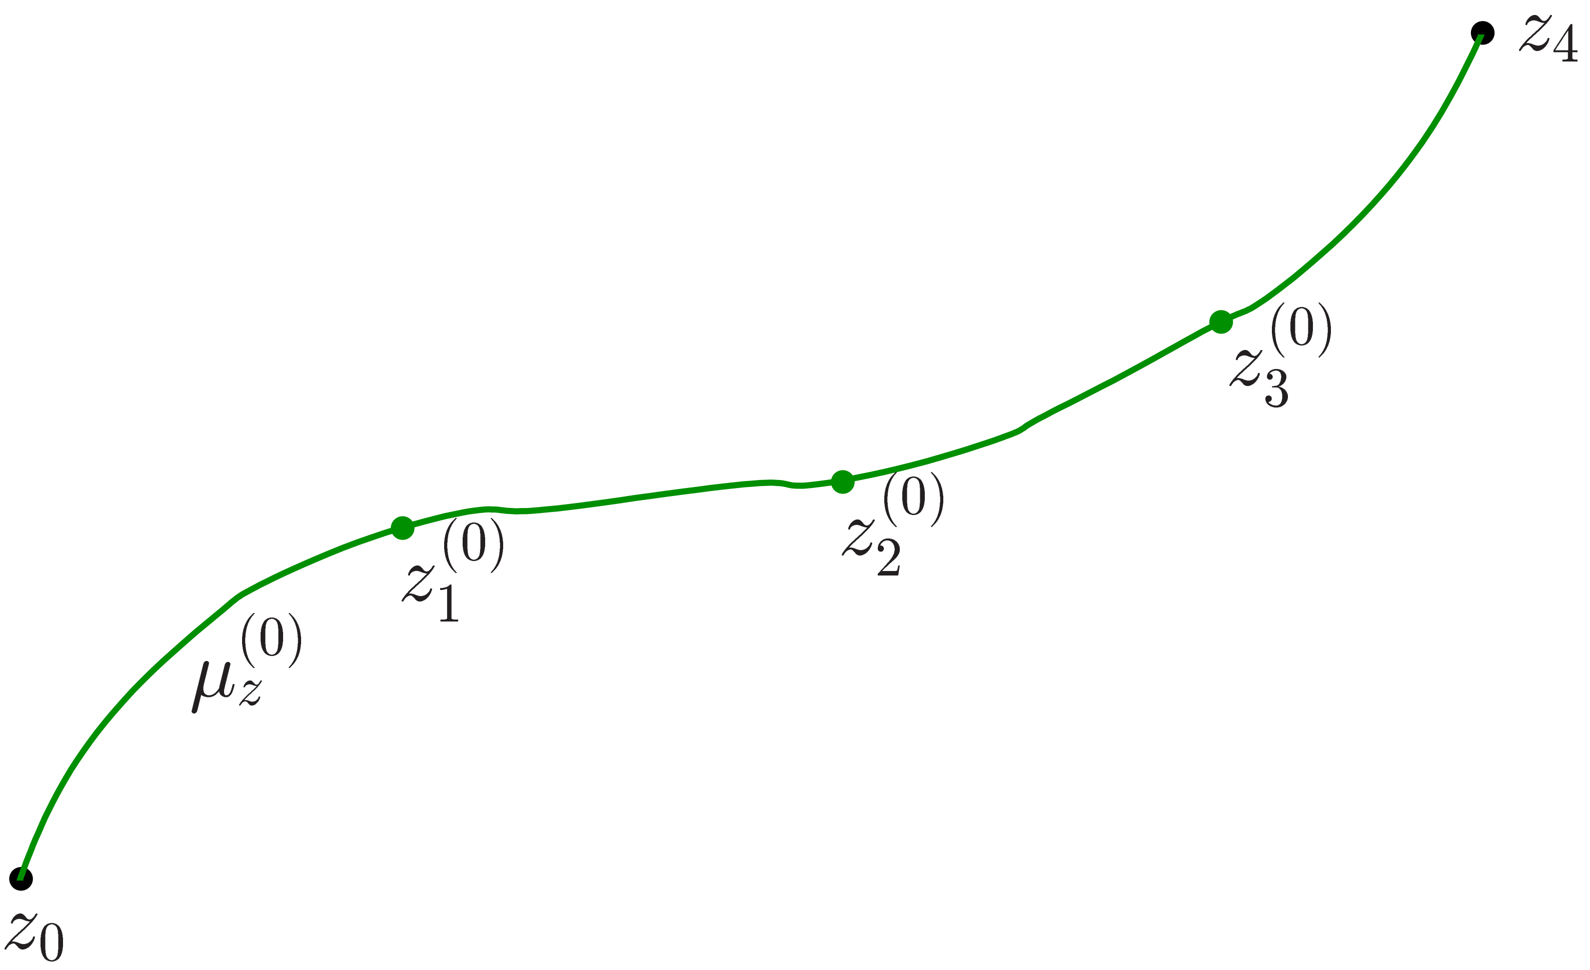

Subproblems

Leapfrog relies on solution of subproblems (P_i) through the given partition:

$$
(\mathrm{P}_i)\quad \left\{
\begin{array}{ll}
\text{minimize} & \displaystyle\int_{t_{i-1}}^{t_{i+1}} f_0(x(t), u(t))\, dt \\
\text{subject to} & \dot{x}(t) = f(x(t), u(t)), \quad u \in \mathcal{U}\ , \\
 & x(t_{i-1}) = z_{i-1} \quad \text{and} \quad x(t_{i+1}) = z_{i+1}\ ,
\end{array}
\right.
$$

where t_{i-1} is fixed by solution of (P_{i-1}) for $i > 1$, and t_{i+1} is free.

Leapfrog Algorithm

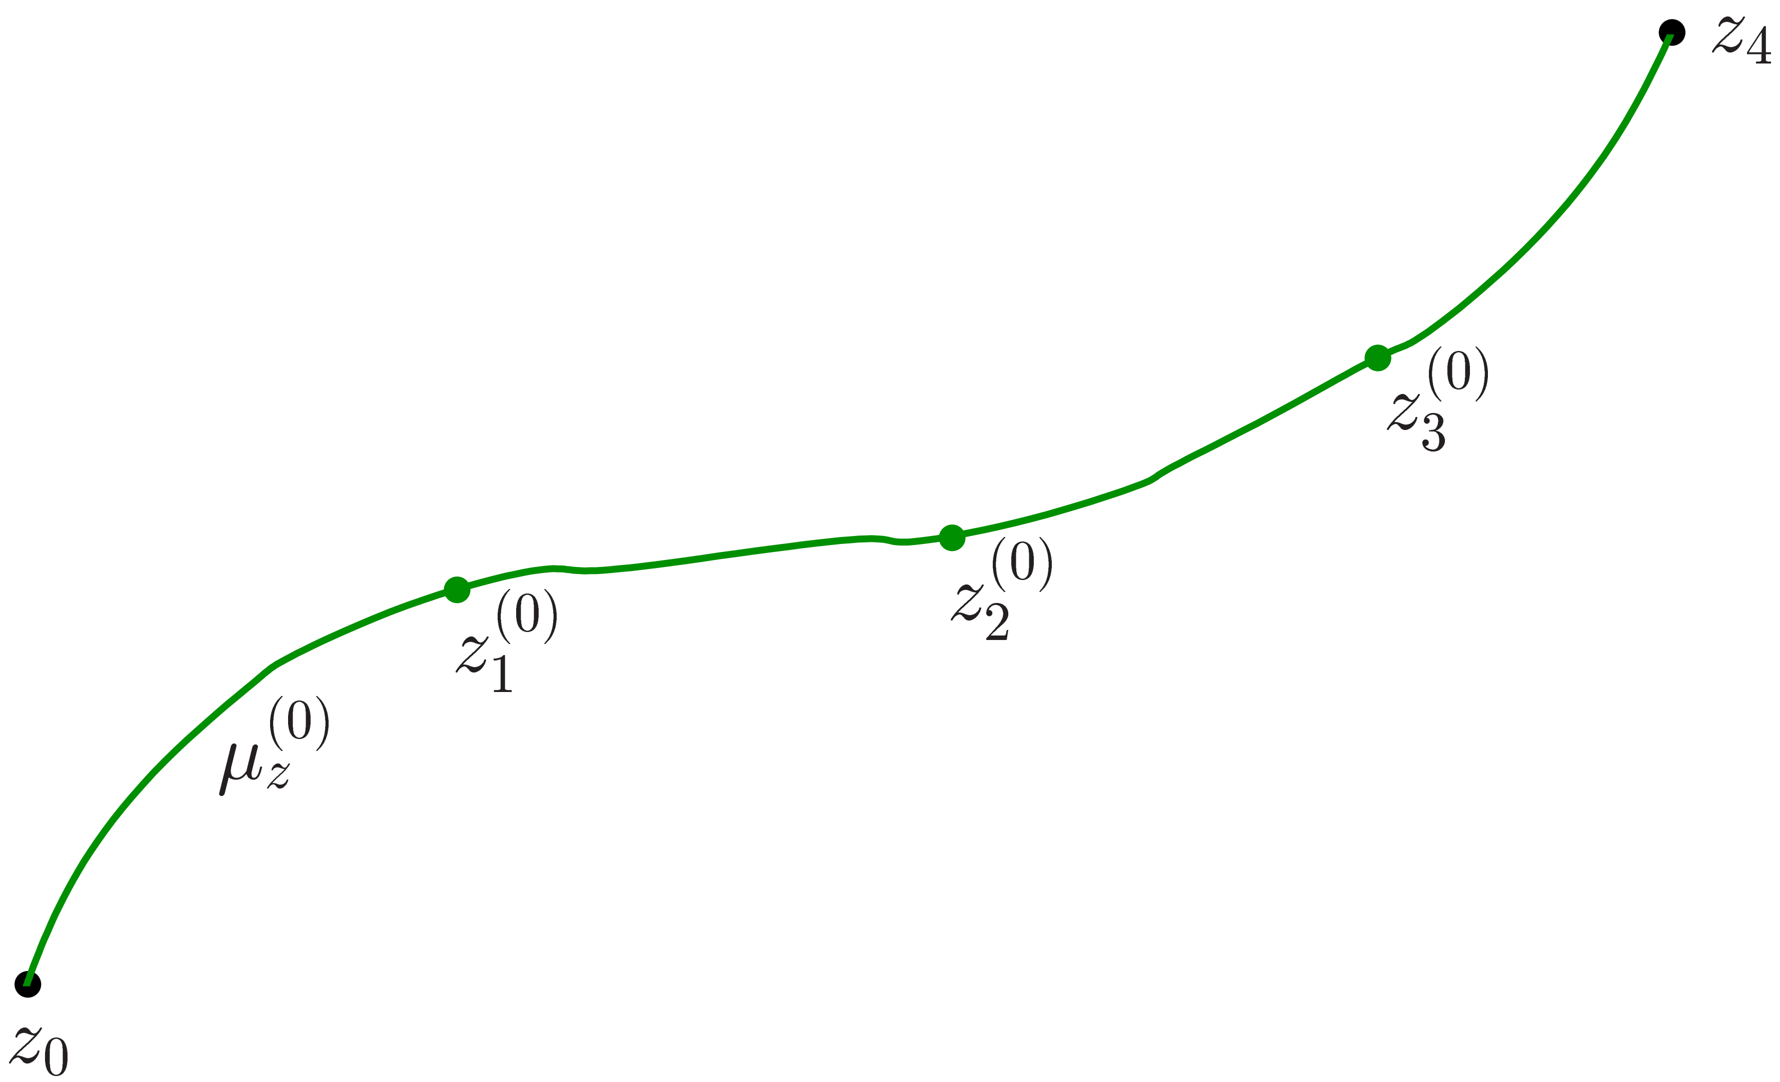

Leapfrog Algorithm

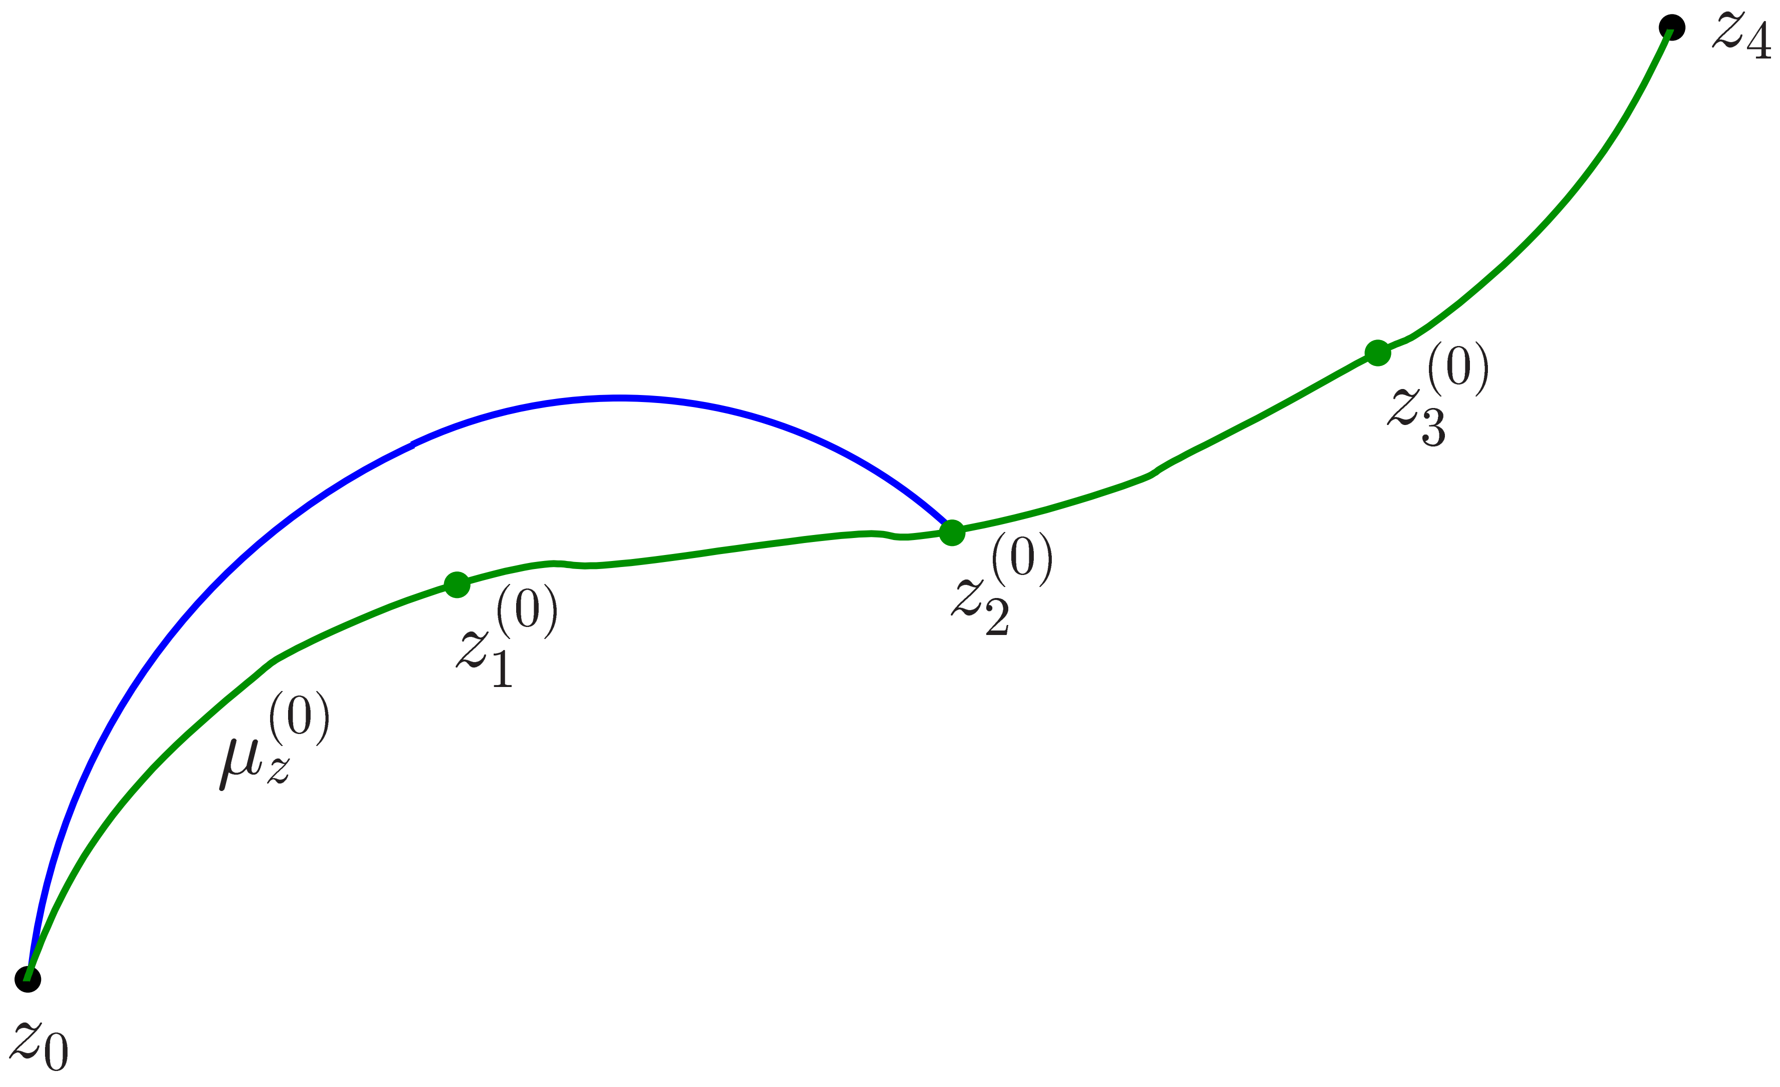

Leapfrog Algorithm

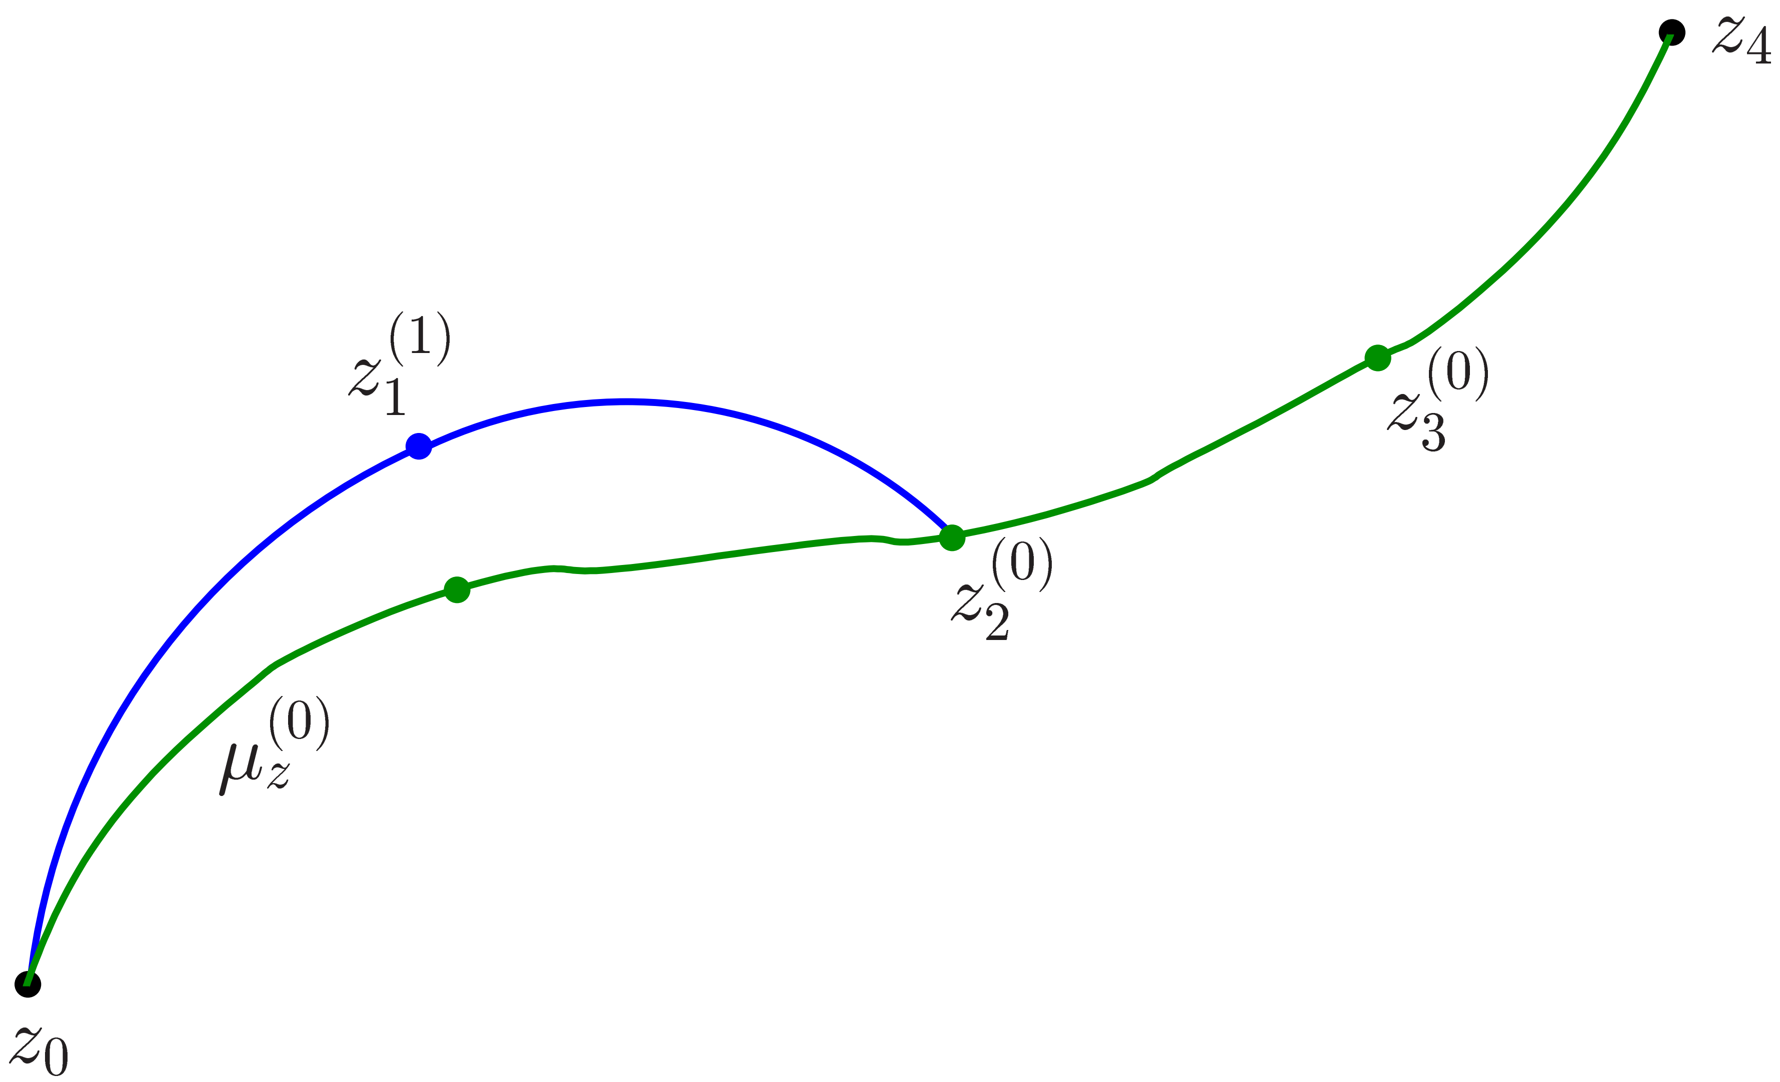

Leapfrog Algorithm

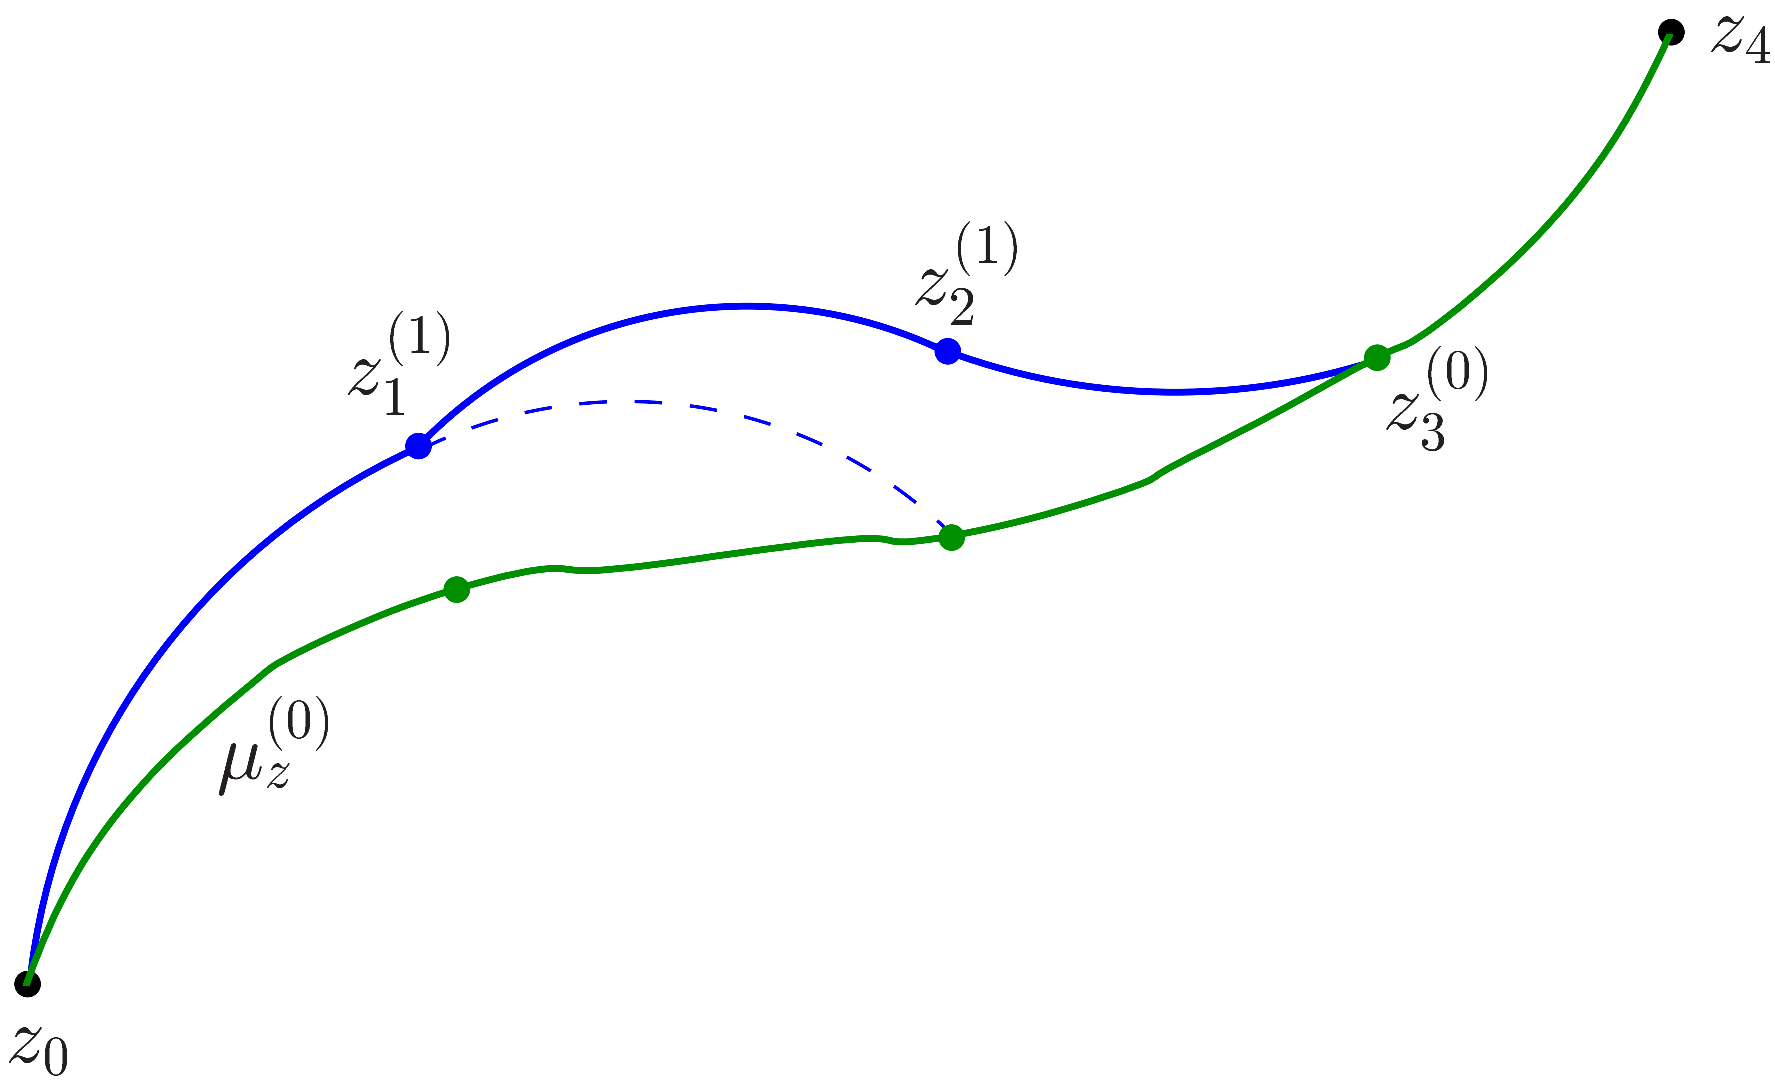

Leapfrog Algorithm

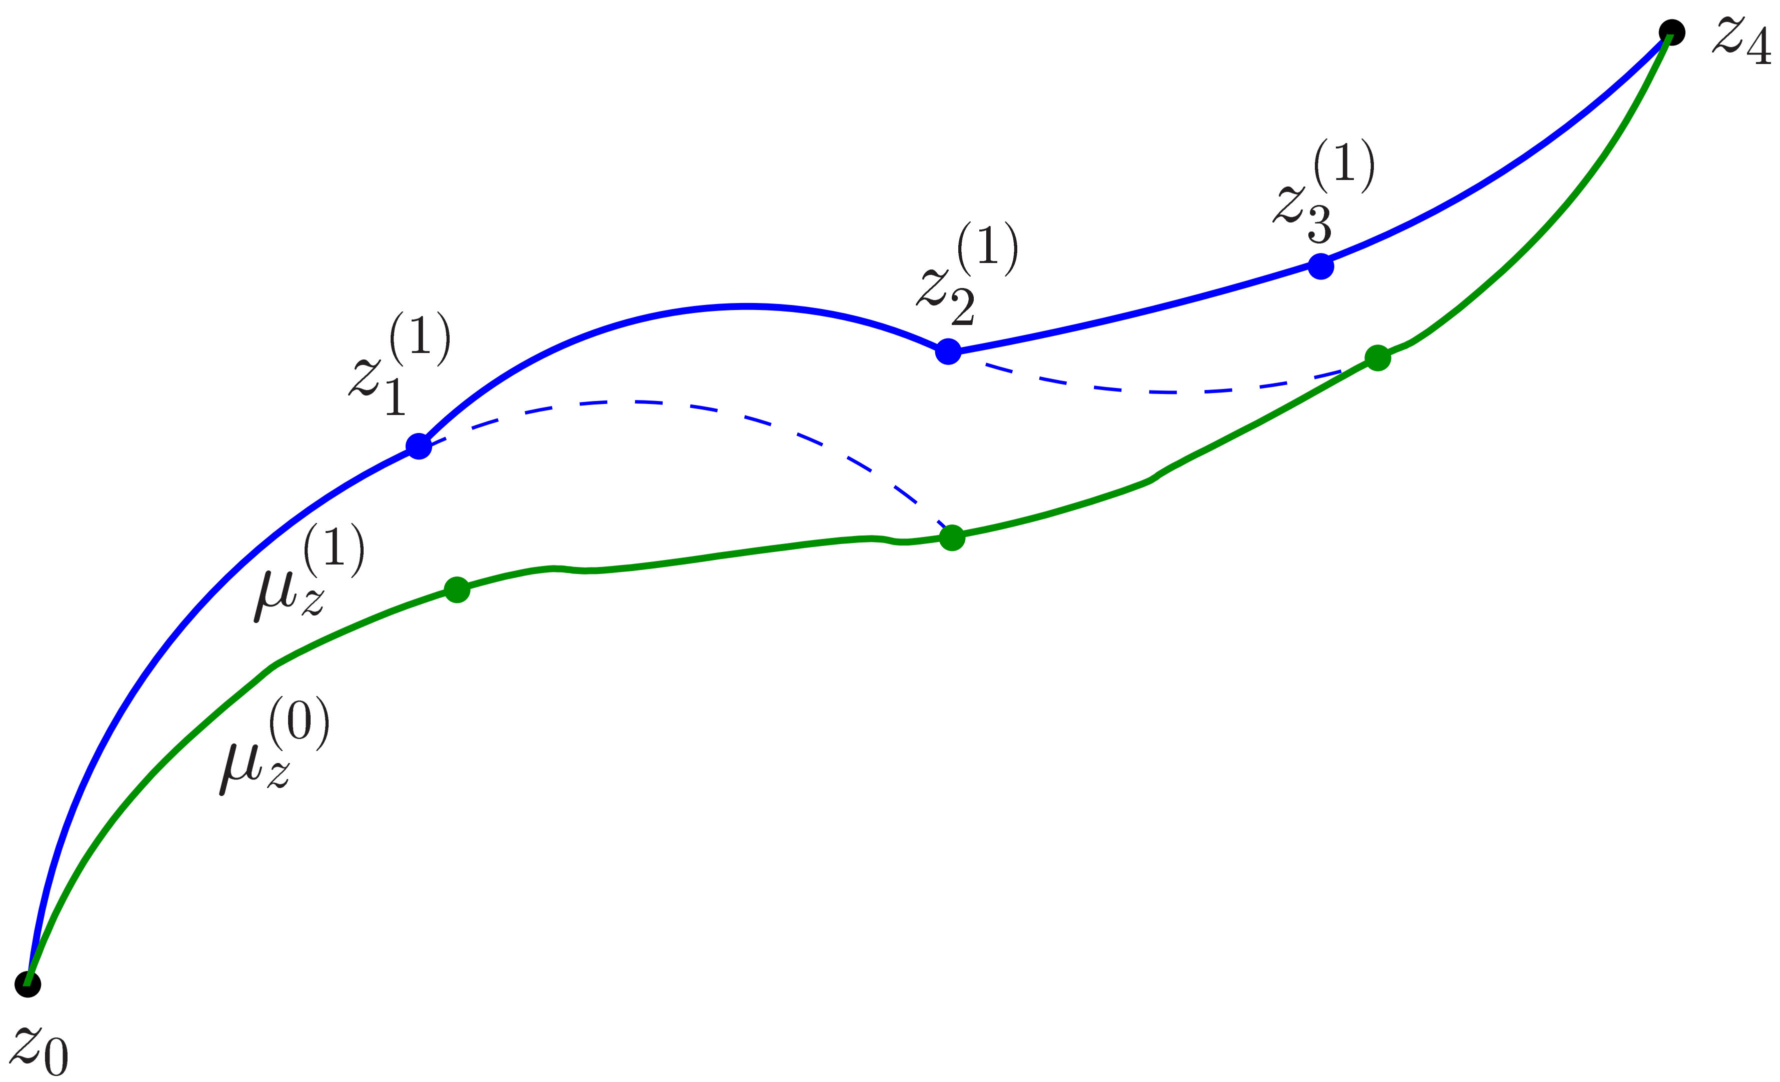

Leapfrog Algorithm

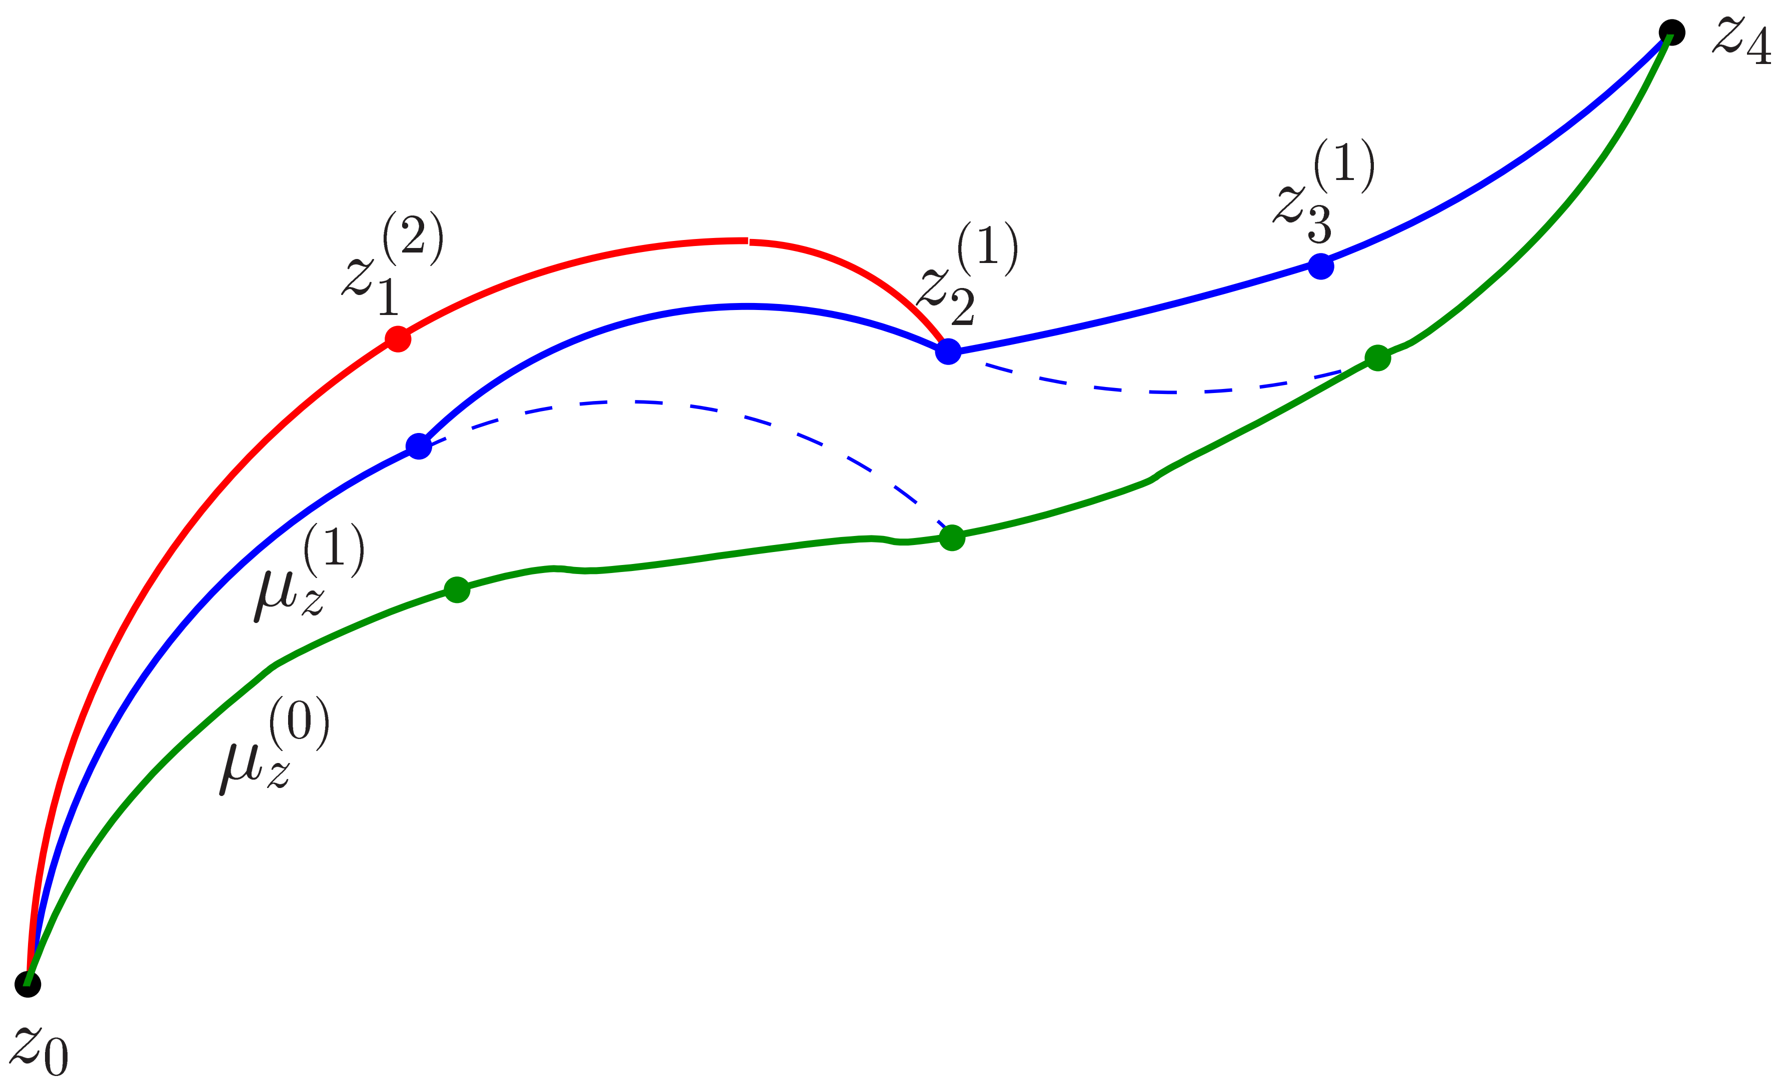

Leapfrog Algorithm

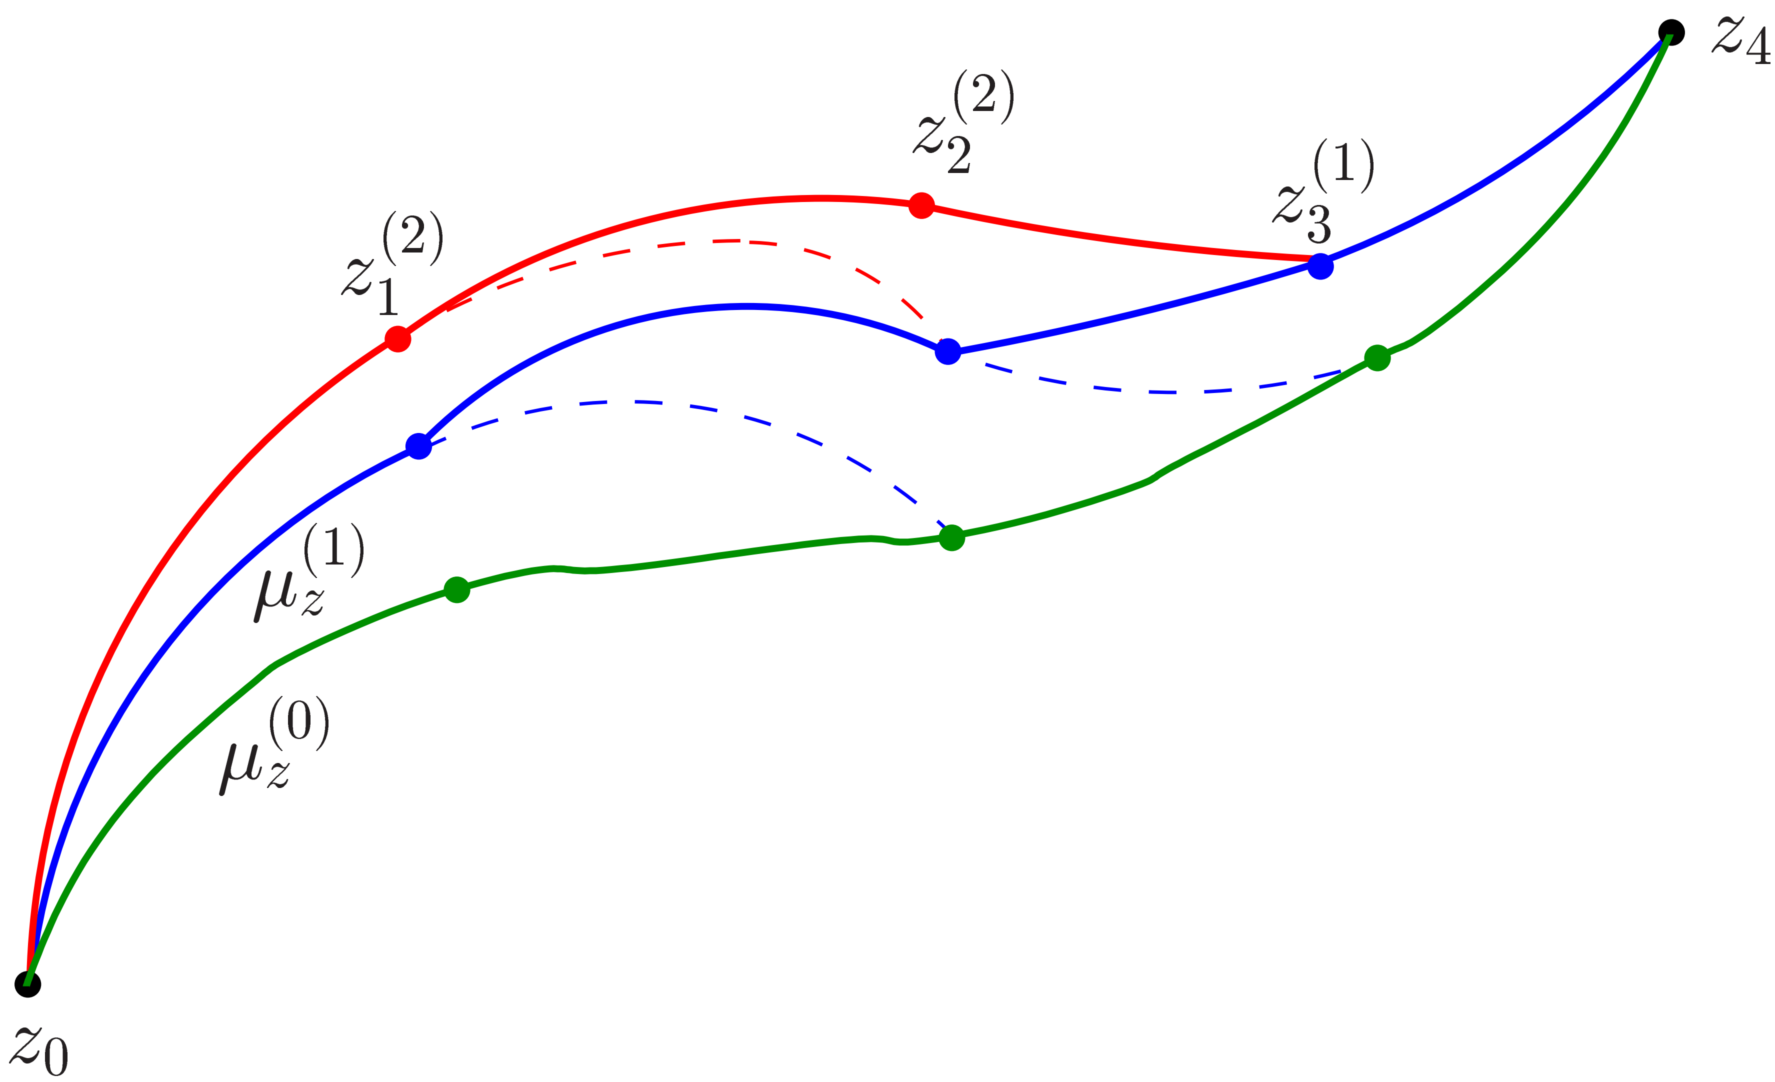

Leapfrog Algorithm

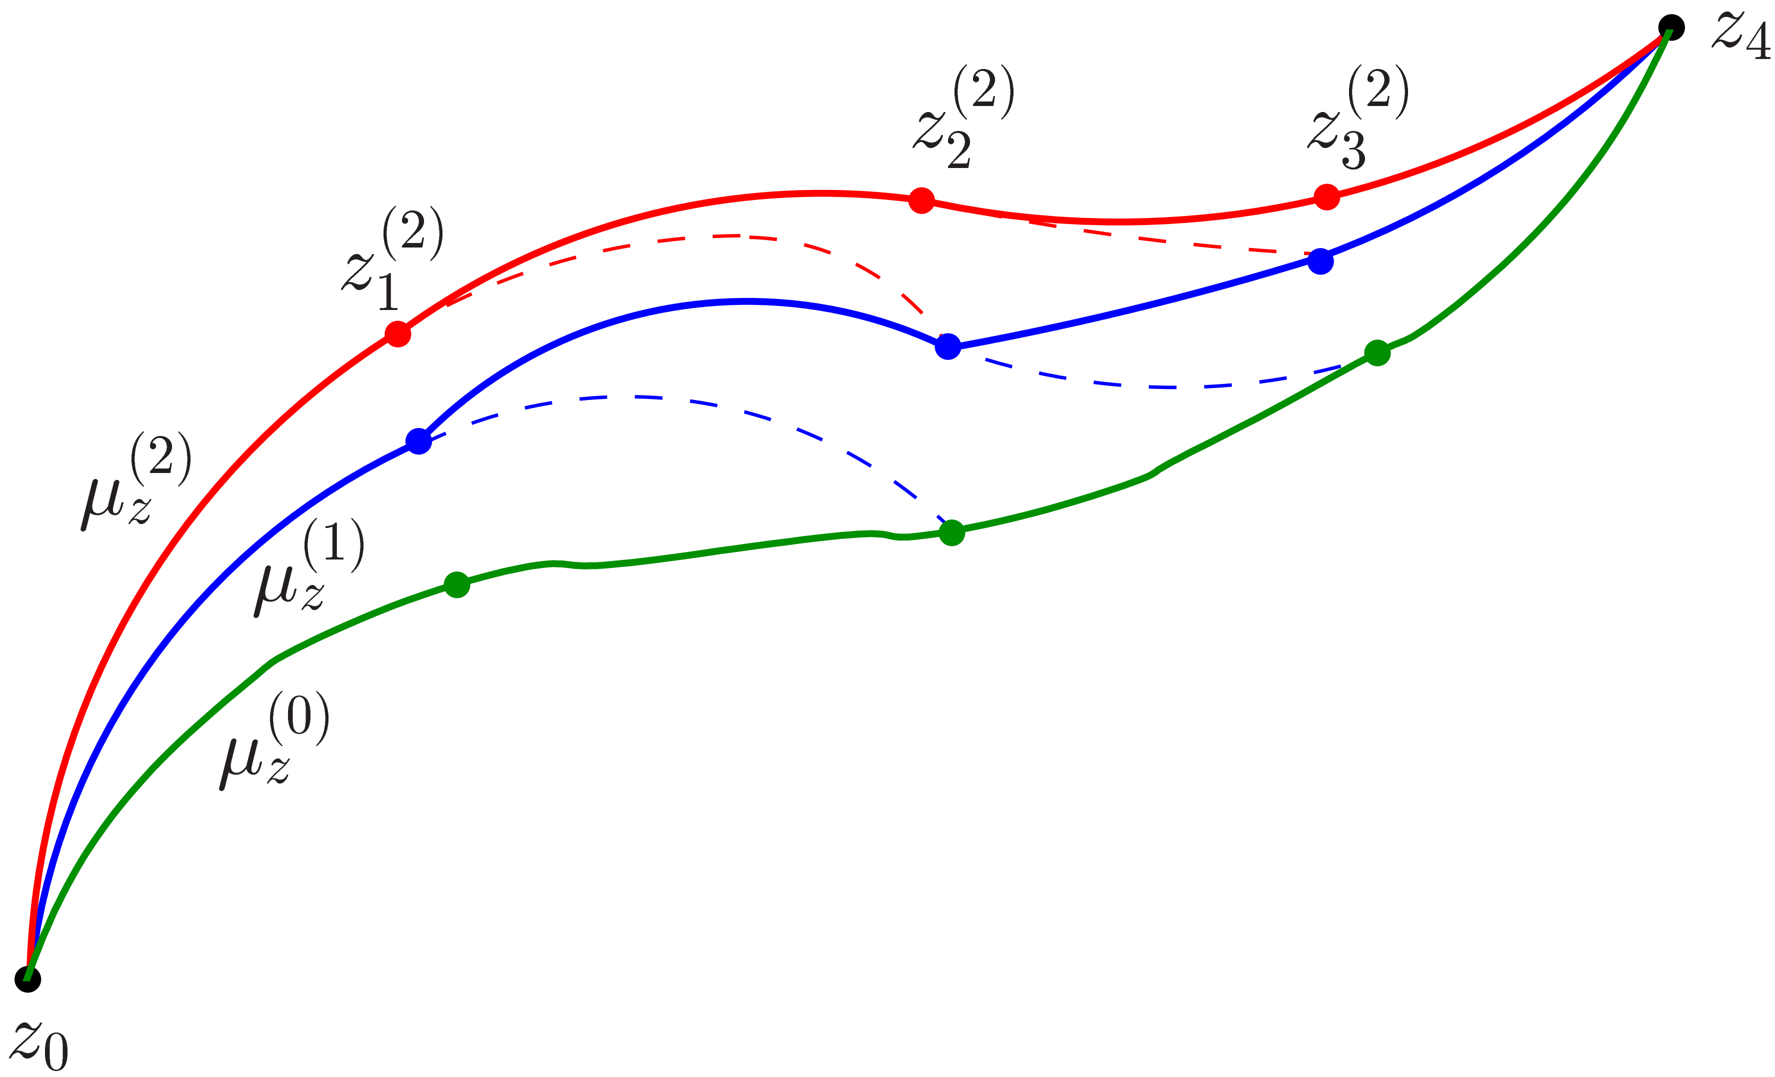

- We set

$$\mu_z^{(k)}(t) := x^{(k)}(t), \qquad t_{i-1}^{(k)} \le t \le t_i^{(k)} \, ,$$

where $t_i^{(k)}$ is the time point when half the optimal cost is reached in Subproblem (P_i).

- This generates a sequence

$$\mathcal{M}_z = \{\mu_z^{(k)} : [t_0, t_f] \longrightarrow \mathbb{R}^n : k > 0\}$$

of piecewise-optimal trajectories.

- Along each consecutive trajectory $\mu_z^{(k)}$ the cost decreases.

- However, to study convergence to a critical trajectory we need to consider the trajectory pair

$$\mu^{(k)}(\cdot) := (x^{(k)}(\cdot), \psi^{(k)}(\cdot))$$

for each subproblem.

- Then we construct the sequence

$$\mathcal{M} = \{\mu^{(k)} : [t_0, t_f] \longrightarrow \mathbb{R}^{2n} : k > 0\} \ .$$

- $\mathcal{M}$ will be shown to have a subsequence uniformly convergent to a critical trajectory.

Updates in Leapfrog

- After solving each (P$_i$) we update

$$t_i, \quad z_i = x(t_i), \quad \lambda_{i-1} := \psi(t_{i-1}) \ .$$

- $y_i := (z_i, \lambda_i) = (x(t_i), \psi(t_i)) \ .$

- Note that y_i, $i = 1, \ldots, q-2$, are updated by solving (P$_i$) and (P$_{i+1}$) consequtively. Also note y_{q-1} is updated only by solving (P$_{q-1}$).

Leapfrog Algorithm and Multiple Shooting

- Leapfrog updates only $2n$ variables (states and costates) at a time. Multiple Shooting updates $[(2q-1)n]$ variables, with q subdivisions.
- State trajectories generated by Leapfrog are feasible at every step of the iteration. $\longrightarrow$ Can be useful in real-time (or on-line) implementations.
- Leapfrog can achieve convergence to a critical trajectory, under some general assumptions.

Uniqueness of controls and costates

Assumptions:

- There exists a (locally) unique optimal control from z_{i-1} to z_{i+1}.
- The costate vector is determined uniquely by the control $u(t)$ and the initial point z_{i-1} of the local optimal trajectory.

Two immediate examples that verify these assumptions are the problem of finding geodesics and minimum energy control of LTI systems.

Local cost function

Define the set

$$D = \left\{ (z_{i-1}, z_{i+1}) \in \mathbb{R}^n \times \mathbb{R}^n : z_{i+1} \in \mathcal{R}(z_{i-1}, \leq 2\delta) \right\} .$$

The local cost function

$$\tau : D \longrightarrow \mathbb{R}$$

is the minimum cost to get from z_{i-1} to z_{i+1}. Note that $\tau(z_{i-1}, z_{i+1}) \leq 2\delta$.

Continuity of the local cost function

Reachability assumption:

Given $\epsilon > 0$ and $z_i, \widetilde{z}_i \in \mathbb{R}^n$, there exists $\kappa > 0$ such that, if $\|z_i - \widetilde{z}_i\| < \kappa$ then, given a locally optimal trajectory (with fixed end points) of cost less than $\epsilon/4$ containing z_i, there exists w in that trajectory reachable with cost less than $\epsilon/4$ from $\widetilde{z}_i$.

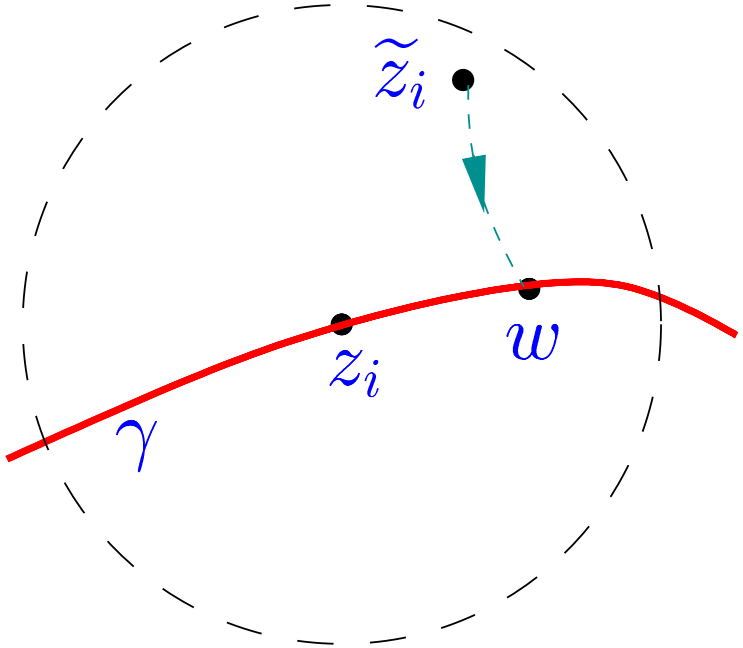

Lemma 1. The local cost function τ is continuous.

Midpoint maps and total cost

- Local optimal trajectory: $\gamma_{z_{i-1},z_{i+1}} : [t_{i-1}, t_{i+1}] \longrightarrow \mathbb{R}^n$.
- Midpoint: $\widehat{z}_i := \gamma_{z_{i-1},z_{i+1}}(\widehat{t}_i)$, $t_{i-1} \leq \widehat{t}_i \leq t_{i+1}$, where $\tau(z_{i-1}, \widehat{z}_i) = \tau(z_{i-1}, z_{i+1})/2$.
- Midpoint map: $M : \mathbb{R}^n \times \mathbb{R}^n \longrightarrow \mathbb{R}^n$ such that

$$M(z_{i-1}, z_{i+1}) = \gamma_{z_{i-1},z_{i+1}}(\widehat{t}_i) = \widehat{z}_i \ .$$

- State-costate space: Let $X \subseteq \mathbb{R}^{2n}$. Y is the set of all

$$y = (y_0, y_1, \ldots, y_q) \in X^{q+1} \ ,$$

with $y_i = (z_i, \lambda_i)$, such that $\tau(z_i, z_{i+1}) \leq \delta$.

- Midpoint iteration: $G_p : Y \longrightarrow X^{q+1}$ such that

$$G_p(y) = (y_0, \ldots, y_{p-1}, \widehat{y}_p, y_{p+1}, \ldots, y_q)$$

where $\widehat{y}_p = (\widehat{z}_p, \widehat{\lambda}_p)$.

- Composite iteration: $F : Y \longrightarrow Y$ such that

$$F = G_{q-1} \circ G_{q-2} \circ \cdots \circ G_1 \, .$$

- Total cost: $\alpha : Z \longrightarrow \mathbb{R}$ such that

$$\alpha(z) = \sum_{i=1}^{q} \tau(z_{i-1}, z_i) \ ,$$

where $z = (z_0, z_1, \ldots, z_q) \in Z$.

- Extended optimal cost: $\tilde{\tau}(z_0, z_q)$ is the minimum cost to get from z_0 to z_q.

Lemma 2.

(a) $G_p(y) \in Y$.

(b) $\alpha(\pi_1 \circ G_p(y)) \leq \alpha(z)$.

(c) $\alpha(\pi_1 \circ F(y)) \leq \alpha(z)$.

Compactness of the state-costate set

Define for $r \geq 0$

$$C(x_0, r) := \{(x(t), \psi(t)) \in \mathbb{R}^{2n} : t \in [t_0, t_1], \widetilde{\tau}(x_0, x(t)) \leq r\}\,.$$

Boundedness assumption:

Let α_0 be the cost along the initial feasible trajectory. Then $C(x_0, \alpha_0)$ is bounded.

Lemma 3.

(a) $\lim_{j \longrightarrow \infty} y^{(k_j)} = y^{(\infty)}$.

(b) Y is compact.

(c) $\alpha(z^{(\infty)}) = \alpha^{(\infty)} \in [\widetilde{\tau}(z_0, z_q), q\delta]$.

Extreme points and Leapfrog splicing

- $y_i = (z_i, \lambda_i)$ is between y_{i-1}, y_{i+1}, if z_i lies in the image of $\gamma_{z_{i-1},z_{i+1}}(\cdot)$.
- $y \in Y$ is extreme when y_i is between y_{i-1}, y_{i+1} for all $1 < i < q$.
- Let $\mu_i(\cdot) = (\gamma_{z_{i-1},z_i}(\cdot), \psi(\cdot))$. Then we write

$$\mu_y = \mu_1 \smile \mu_2 \smile \cdots \smile \mu_q \,.$$

Lemma 4. If $y \in Y$ is extreme, then μ_y is a critical trajectory.

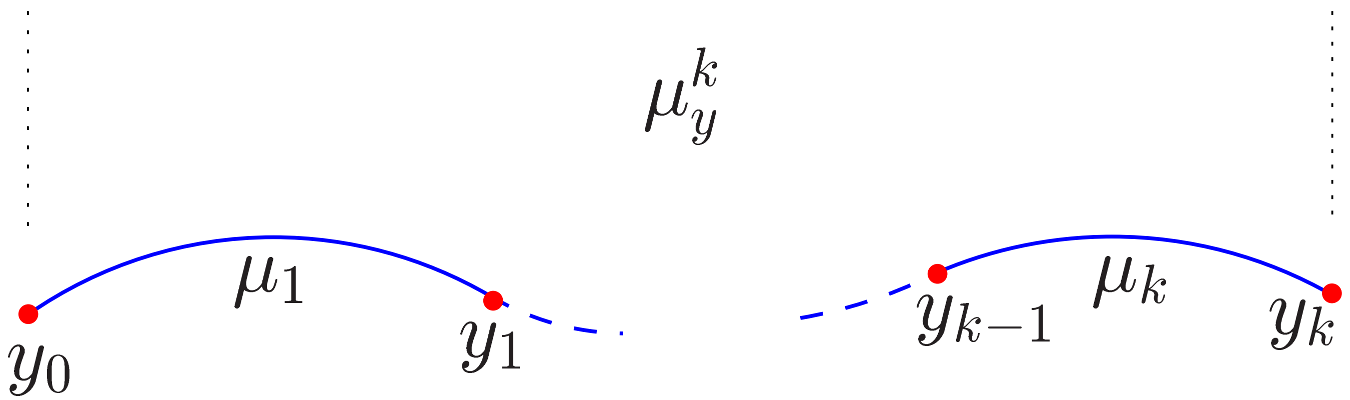
μ_y^k
μ_1
μ_k
y_0
y_1
y_{k-1}
y_k

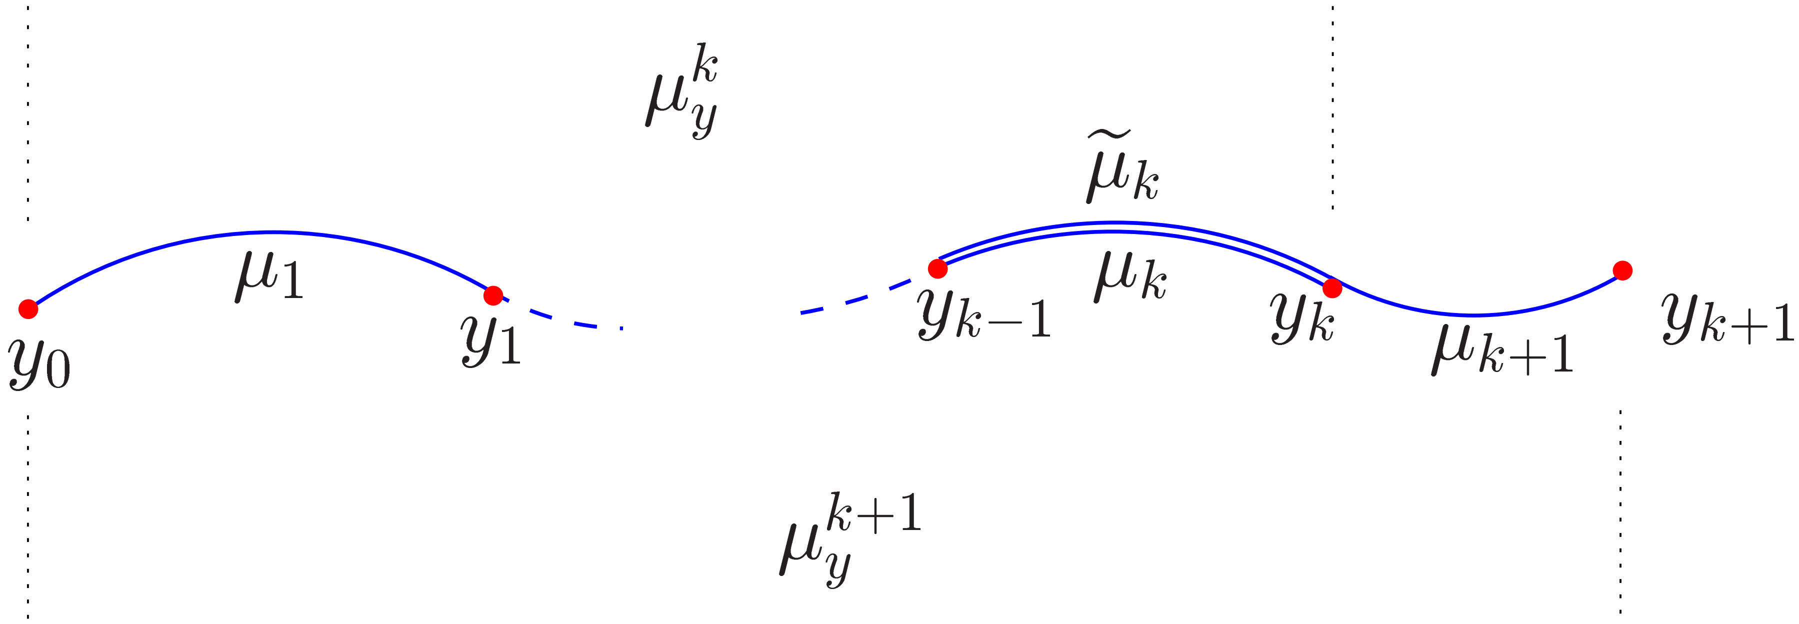
μ_y^k
$\widetilde{\mu}_k$
μ_1
μ_k
y_0
y_1
y_{k-1}
y_k
y_{k+1}
μ_{k+1}
μ_y^{k+1}

Lemma 4. If $y \in Y$ is extreme, then μ_y is a critical trajectory.

Lemma 5.

(a) If $F(y) = y$ then y is extreme.

(b) y is extreme if and only if $\mu_{F(y)} = \mu_y$.

Lemma 6. $F(y^{(\infty)}) = y^{(\infty)}$.

Convergence to a critical trajectory

Theorem 1. $\mu_{y(\infty)}$ is a critical trajectory.

Corollary 1. $\mu_{F(y(\infty))} = \mu_{y(\infty)}$; thereby convergence to a critical trajectory $\mu_{y(\infty)}$ is ensured.