An explicit proof of “Vinter’s” theorem

Milan Korda, Corbinian Schlosser

May 17, 2021

“The Vinter program”

“The Winter program”

Setting

$$
\begin{aligned}
V(s, x_0) := \quad & \inf_{u(\cdot), x(\cdot)} \quad \int_s^T l(t, x(t), u(t))\, dt + l_{\mathrm{T}}(x(T)) \\
\text{s.t.} \quad & \dot{x}(t) = f(t, x(t), u(t)), t \in [s, T] \\
& x(s) = x_0 \\
& x(t) \in X, t \in [s, T] \\
& x(T) \in X_{\mathrm{T}} \\
& u(t) \in U, t \in [s, T].
\end{aligned}
$$

for X, X_T, U compact, l, l_{T} continuous and f Lipschitz.

Setting

$$
\begin{aligned}
V(s,x_0) := \quad & \inf_{u(\cdot),x(\cdot)} \quad \int_s^T l(t,x(t),u(t))\,dt + l_{\mathrm{T}}(x(T)) \\
\text{s.t.} \quad & \dot{x}(t) = f(t,x(t),u(t)), t \in [s,T] \\
& x(s) = x_0 \\
& x(t) \in X, t \in [s,T] \\
& x(T) \in X_{\mathrm{T}} \\
& u(t) \in U, t \in [s,T].
\end{aligned}
$$

for X, X_T, U compact, l, l_{T} continuous and f Lipschitz.

Remark

Even though the setting is friendly (we have enough regularity and compactness) but still V is not smooth in general.

Rough idea and take away message

1. Setting:

Rough idea and take away message

1. Setting: Nonlinear optimal control problem (OCP)

Rough idea and take away message

1. Setting: Nonlinear optimal control problem (OCP)
2. Linearize:

Rough idea and take away message

1. Setting: Nonlinear optimal control problem (OCP)
2. Linearize: By embedding into a measure setting

Rough idea and take away message

1. Setting: Nonlinear optimal control problem (OCP)
2. Linearize: By embedding into a measure setting
3. Solving:

Rough idea and take away message

1. Setting: Nonlinear optimal control problem (OCP)
2. Linearize: By embedding into a measure setting
3. Solving:
    i. What happened by transforming the nonlinear ocp to a linear programming problem on measures? Did we relax the problem, does a solution exist, how does a solution of the embedding problem correspond to a solution of the original problem, etc.
    ii. What are properties of the solution (regularity), how can I check that a given point is a solution?
    iii. How to solve the linearized problem?

Rough idea and take away message

1. Setting: Nonlinear optimal control problem (OCP)
2. Linearize: By embedding into a measure setting
3. Solving:
   i. What happened by transforming the nonlinear ocp to a linear programming problem on measures? Did we relax the problem, does a solution exist, how does a solution of the embedding problem correspond to a solution of the original problem, etc.
   ii. What are properties of the solution (regularity), how can I check that a given point is a solution?
   iii. How to solve the linearized problem?

Rough idea and take away message

1. Setting: Nonlinear optimal control problem (OCP)
2. Linearize: By embedding into a measure setting
3. Solving:
   i. What happened by transforming the nonlinear ocp to a linear programming problem on measures? Did we relax the problem, does a solution exist, how does a solution of the embedding problem correspond to a solution of the original problem, etc.
   ii. What are properties of the solution (regularity), how can I check that a given point is a solution?
   iii. How to solve the linearized problem?

Linearize

Embedding the problem into the space of measures via occupation measures leads to an infinite dimensional linear programming problem (Rubio, Vinter, Lasserre,...). An occupation measure μ with respect to $x(\cdot), u(\cdot)$ is defined by

$$\langle g, \mu \rangle := \int_0^T g(t, x(t), u(t))\, dt.$$

Linearize

Embedding the problem into the space of measures via occupation measures leads to an infinite dimensional linear programming problem (Rubio, Vinter, Lasserre,...). An occupation measure μ with respect to $x(\cdot), u(\cdot)$ is defined by

$$\langle g, \mu \rangle := \int_0^T g(t, x(t), u(t))\, dt.$$

In particular we can express the cost by
$\int_0^T l(t, x(t), u(t))\, dt = \langle l, \mu \rangle.$

Linearize

Embedding the problem into the space of measures via occupation measures leads to an infinite dimensional linear programming problem (Rubio, Vinter, Lasserre,...). An occupation measure μ with respect to $x(\cdot), u(\cdot)$ is defined by

$$\langle g, \mu \rangle := \int_0^T g(t, x(t), u(t))\, dt.$$

In particular we can express the cost by $\int_0^T l(t, x(t), u(t))\, dt = \langle l, \mu \rangle$.

Remark

This procedure is similar to how the Perron-Frobenius operator can be obtained – where the state constraint X is embedded into $M(X)$ by $x \mapsto \delta_x$ and the dynamics $x \mapsto \varphi_t(x)$ induce the Perron-Frobenius operator $\delta_x \mapsto \delta_{\varphi_t(x)} = P_t \delta_x$.

Primal and dual LP

The primal problem is given by

$$
\begin{array}{ll}
p^* := & \inf\langle (l, l_{\mathrm{T}}), (\mu, \nu)\rangle \\
\text{s.t.} & \mu \in M([0, T] \times X \times U), \nu \in M(X_{\mathrm{T}}) \\
 & \int \frac{\mathrm{d}}{\mathrm{dt}} g + \nabla_x g \; d\mu - \int g(T, \cdot) \; d\nu = -g(0, x_0) \quad \text{for all } g \in \mathcal{C}^1
\end{array}
$$

Primal and dual LP

The primal problem is given by

$$
\begin{array}{ll}
p^* := & \inf \langle (l, l_{\mathrm{T}}), (\mu, \nu) \rangle \\
\text{s.t.} & \mu \in M([0, T] \times X \times U), \nu \in M(X_{\mathrm{T}}) \\
& \int \frac{\mathrm{d}}{\mathrm{dt}} g + \nabla_x g \; d\mu - \int g(T, \cdot) \; d\nu = -g(0, x_0) \quad \text{for all } g \in \mathcal{C}^1
\end{array}
$$

The dual problem has the following form

$$
\begin{array}{ll}
d^* := \sup\limits_{\phi \in \mathcal{C}^1(\mathbb{R}^{n+1})} \quad \phi(0, x_0) & \\
\text{s.t. } \phi(T, x) \leq l_{\mathrm{T}}(x) & x \in X_T \\
\frac{\mathrm{d}}{\mathrm{dt}} \phi(t, x) + \nabla_x \phi(t, x) f(t, x, u) + l(t, x, u) \geq 0 & \text{on } [0, T] \times X \times U
\end{array}
$$

Primal and dual LP

The primal problem is given by

$$
\begin{array}{ll}
p^* := & \inf \langle (l, l_{\mathrm{T}}), (\mu, \nu) \rangle \\
\text{s.t.} & \mu \in M([0,T] \times X \times U), \nu \in M(X_{\mathrm{T}}) \\
& \int \frac{\mathrm{d}}{\mathrm{dt}} g + \nabla_x g \; d\mu - \int g(T, \cdot) \; d\nu = -g(0, x_0) \quad \text{for all } g \in \mathcal{C}^1
\end{array}
$$

The dual problem has the following form

$$
\begin{array}{ll}
d^* := \sup_{\phi \in \mathcal{C}^1(\mathbb{R}^{n+1})} \quad \phi(0, x_0) & \\
\text{s.t. } \phi(T, x) \leq l_{\mathrm{T}}(x) & x \in X_T \\
\frac{\mathrm{d}}{\mathrm{dt}} \phi(t, x) + \nabla_x \phi(t, x) f(t, x, u) + l(t, x, u) \geq 0 & \text{on } [0, T] \times X \times U
\end{array}
$$

Remark

The constraint in the primal problem is called Liouville equation, we refer to the constraint in the dual by "Hamilton-Jacobi-Bellman" (HJB) inequality.

Well posedness

Lemma

There exists feasible points for the primal and dual LPs.

Proof.

For the dual problem choose $\phi(t, x) := ct - m$ constant for a c, m large enough. For the primal choos any occupation measure induced by a feasible $(x(\cdot), u(\cdot))$ for the OCP. $\square$

Relaxation gap

Proposition

We have $d^* \leq p^* \leq V(0, x_0)$.

Proof.

Since the primal is a relaxation of the OCP we have $p^* \leq V(0, x_0)$ and $d^* \leq p^*$ follows from weak duality. □

The optimal cost function V and smoothness

We state two ways of how smoothness of the optimal cost function affects the LPs.

Proposition

If V is smooth then there is no relaxation gap and $d^* = p^* = V(0, x_0)$.

We state two ways of how smoothness of the optimal cost function affects the LPs.

Proposition

If V is smooth then there is no relaxation gap and $d^* = p^* = V(0, x_0)$.

Proof.

Via the dual problem. We already know $d^* \leq V(0, x_0)$. If V is smooth it satisfies the HJB equality by Bellman's optimality principle. Hence V is feasible and the cost of the dual problem is (by definition) $V(0, x_0)$, i.e. $d^* \geq V(0, x_0)$. $\square$

The optimal cost function V and smoothness

Proof.

For the primal. We use V as a test function for the primal problem and use that it satisfies the HJB equation $\partial_t V + \nabla V \cdot f \leq l$, and $V(T, \cdot) = l_{\mathrm{T}}$ to get for any feasible (μ, ν)

$$\begin{aligned} V(0, x_0) &\overset{(\mu,\nu)\text{ feasible}}{=} \int \partial_t V + \nabla V \, d\mu + \int V(T, \cdot) \, d\nu \\ &\overset{\text{HJB}}{\leq} \int l \, d\mu + \int l_{\mathrm{T}} \, d\nu = \langle (l, l_{\mathrm{T}}), (\mu, \nu) \rangle. \end{aligned}$$

□

Vinter's theorem

(A1) Well posedness: $V(0, x_0) < \infty$

(A2) Compactness: X, X_T and U are compact.

(A3) Regularity: $f : \mathbb{R} \times \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^n$ is Lipschitz,
$l : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}$, $l_\mathrm{T} : [0, \bar{T}] \times \mathbb{R}^n \to \mathbb{R}$ are continuous

(A4) Convexity in the dynamics: $f(t, x, U)$ is convex for all (t, x)

(A5) Convexity in the cost: For all x the following map is convex

$$v \mapsto g_{t,x}(v) := \inf_{u \in U} \{l(t, x, u) : f(t, x, u) = v\}$$

Theorem (Vinter's theorem)

Under assumptions (A1)-(A5) we have $d^* = p^* = V(0, x_0)$ for $V(0, x_0)$ being the optimal value of the OCP, p^* the optimal value of the primal (occupation measure) LP and d^* the optimal value of the Vinter LP.

Problem: V is not smooth in general

In general V is not smooth.

Problem: V is not smooth in general

In general V is not smooth.

The state constraints are problematic.

Problem: V is not smooth in general

In general V is not smooth.

The state constraints are problematic.

Lemma (If there are no state constraints then ϕ is Lipschitz)

If f, l, l_T are Lipschitz and there are no state constraints. Then V is Lipschitz and solves almost everywhere

$$l(t,x,v) + \partial_t V(t,x) + \partial_x V(t,x) f(t,x,v) \geq 0.$$

Proof.

Gronwall's inequality, Lipschitz continuity of l and l_{T}; Bellman's optimality principle. □

Outline of the proof

1. We smoothen the OCP by replacing the state constraints by smooth penalties in the cost.
2. We show that the penalized setting (almost) satisfies the HJB inequality.
3. We show that as the penalty blows up the optimal values of the penalized problems converge to the optimal value $V(0, x_0)$ of the original problem. Here the convexity assumptions will be needed.
4. Mollifying the (almost smooth) optimal cost functions for the penalized problems leads to sequence of feasible points for Vinter's LP with cost converging to $V(0, x_0)$, i.e. we have found a minimizing sequence and $d^* \geq V(0, x_0)$.

Almost candidate functions

How do we treat the state constraints?

Almost candidate functions

How do we treat the state constraints?
Idea: We add them as penalties to the cost.

Almost candidate functions

How do we treat the state constraints?
Idea: We add them as penalties to the cost.

For $k \in \mathbb{N}$ we define our (almost) candidate function $\phi_k(s, x_0)$ by

$$\begin{aligned} \inf_{u(\cdot),x(\cdot)} \; & \int_0^T l(t, x(t), u(t)) + k\psi_X(x(t))\, dt + l_{\mathrm{T}}(x(T)) + k\psi_{X_{\mathrm{T}}}(x(T)) \\ \text{s.t.} \quad & \dot{x}(t) = f(t, x(t), u(t)), t \in [0, T] \\ & x(s) = x_0 \\ & u(t) \in U, t \in [0, T]. \end{aligned}$$

where ψ_X and $\psi_{X_{\mathrm{T}}}$ are smooth indicators of X, X_{T}.

Corollary

The functions ϕ_k are Lipschitz continuous and solve the HJB inequality almost everywhere.

Corollary

The functions ϕ_k are Lipschitz continuous and solve the HJB inequality almost everywhere.

Proposition (ϕ_k converge to the optimal value function)

As $k \to \infty$ the penalty candidate functions ϕ_k converge from below and monotonously to the optimal value function V.

Proof idea; selecting a limit object

a) For each k choose optimal $(u_k(\cdot), x_k(\cdot))$

Proof idea; selecting a limit object

a) For each k choose optimal $(u_k(\cdot), x_k(\cdot))$

b) Selecting convergent subsequences:

1. For x: By Arzela-Ascoli there exists a subsequence $x_k \to x$.
   Penalties: x respects the state constraints.
2. For u: More subtle because we don't have regularity for u.
   Occupation measures μ_k

$$\langle g, \mu_k \rangle := \int\limits_{[0,T]\times U} g \, d\mu_k = \int\limits_0^T g(r, u_k(r)) \, dr$$

   Banach-Alaoglu: $\exists \ \mu$ with $\mu_k \rightharpoonup \mu$.

c) Limit properties:

   i. Cost: $\phi_k(s, x_0) \to \int\limits_{[s,T]\times U} l(r, x(r), u)\, d\mu + l_T(x(T))$

   ii. ODE: disintegrating $d\mu = d\nu_r dr$

$$
\begin{aligned}
x(t) - x_0 \quad \leftarrow \quad x_k(t) - x_0 &= \int_0^t f(r, x_k(r), u_k(r))\, dr \\
&= \int\limits_{[0,t]\times U} f(\cdot, x_k(\cdot), \cdot)\, d\mu_k \\
&\to \int\limits_{[0,t]\times U} f(\cdot, x(\cdot), \cdot)\, d\mu = \int_0^t \int_U f(r, x(r), u)\, d\nu_r(u)\, dr
\end{aligned}
$$

How good is the limit? Did we relax to much?

Need for convexity

By passing to the limit for the occupation measures we allow a simultaneous mix of controls. Because the disintegration measures ν_r are probability measures we have

$$\dot{x} = \int_U f(r, x(r), u)\, d\nu_r(u) \in \mathrm{conv}(f(r, x(r), U)) = f(r, x(r), U)$$

Need for convexity

By passing to the limit for the occupation measures we allow a simultaneous mix of controls. Because the disintegration measures ν_r are probability measures we have

$$\dot{x} = \int_U f(r, x(r), u)\, d\nu_r(u) \in \operatorname{conv}(f(r, x(r), U)) = f(r, x(r), U)$$

Therefore we choose u^* by

$$u^*(r) \in \operatorname{argmin} \left\{ l(t, x(t), v) : f(t, x(t), v) = \int_U f(t, x(t), u)\, d\nu_r(u) \right\}.$$

Need for convexity

By passing to the limit for the occupation measures we allow a simultaneous mix of controls. Because the disintegration measures ν_r are probability measures we have

$$\dot{x} = \int_U f(r, x(r), u)\, d\nu_r(u) \in \operatorname{conv}(f(r, x(r), U)) = f(r, x(r), U)$$

Therefore we choose u^* by

$$u^*(r) \in \operatorname{argmin} \left\{ l(t, x(t), v) : f(t, x(t), v) = \int_U f(t, x(t), u)\, d\nu_r(u) \right\}.$$

The second convexity condition is used to apply Jensen's inequality to the the probability measures ν_r.

Before we proof Vinter's theorem let us recall what we did so far and what we have shown:

1. Easy part: $d^* \leq V(0, x_0)$
2. Smoothing: We approximated the OCP by replacing the state constraints by penalty costs. This leads to Lipschitz continuous optimal cost functions ϕ_k.
3. Constraints: By Bellman's optimality principle these functions satisfy the HJB equation almost everywhere
4. Convergence/Minimizing sequence: The functions ϕ_k converge to the optimal cost function V of the original problem (convexity needed in order not to relax).

Proof of Vinter's theorem

To obtain $d^* \geq V(0, x_0)$ we show that a slight modification of ϕ_k gives a sequence with costs converging to $V(0, x_0)$.

Proof of Vinter's theorem

To obtain $d^* \geq V(0, x_0)$ we show that a slight modification of ϕ_k gives a sequence with costs converging to $V(0, x_0)$.
We have for all k and a.e. (t, x) and all $v \in U$

$$\begin{array}{rcl} \phi_k(T, x_{\mathrm{T}}) & \leq & l_{\mathrm{T}}(x_{\mathrm{T}}) + k\psi_{X_{\mathrm{T}}}(x_{\mathrm{T}}) \\ 0 & \leq & l(\cdot, \cdot, v) + k\psi_X + \partial_t \phi_k + \partial_x \phi_k f(\cdot, \cdot, v). \end{array}$$

Proof of Vinter's theorem

To obtain $d^* \geq V(0, x_0)$ we show that a slight modification of ϕ_k gives a sequence with costs converging to $V(0, x_0)$.
We have for all k and a.e. (t, x) and all $v \in U$

$$\begin{array}{rcl} \phi_k(T, x_{\mathrm{T}}) & \leq & l_{\mathrm{T}}(x_{\mathrm{T}}) + k\psi_{X_{\mathrm{T}}}(x_{\mathrm{T}}) \\ 0 & \leq & l(\cdot, \cdot, v) + k\psi_X + \partial_t \phi_k + \partial_x \phi_k f(\cdot, \cdot, v). \end{array}$$

Mollifying ϕ_k to $\tilde{\phi}_k$ in order to make it $\mathcal{C}^1$ implies that $\tilde{\phi}_k$ solves the mollified version of of the above HJB constraint.

Proof of Vinter's theorem

To obtain $d^* \geq V(0, x_0)$ we show that a slight modification of ϕ_k gives a sequence with costs converging to $V(0, x_0)$.
We have for all k and a.e. (t, x) and all $v \in U$

$$
\begin{aligned}
\phi_k(T, x_\mathrm{T}) &\leq l_\mathrm{T}(x_\mathrm{T}) + k\psi_{X_\mathrm{T}}(x_\mathrm{T}) \\
0 &\leq l(\cdot, \cdot, v) + k\psi_X + \partial_t \phi_k + \partial_x \phi_k f(\cdot, \cdot, v).
\end{aligned}
$$

Mollifying ϕ_k to $\tilde{\phi}_k$ in order to make it $\mathcal{C}^1$ implies that $\tilde{\phi}_k$ solves the mollified version of of the above HJB constraint.Since everything is continuous this means that the originial HJB is only slightly violated (ε). Thanks to the specific structure of the HJB we can consider the function $\tilde{\phi}_k + \varepsilon t - 2T\varepsilon$ to handle this error.

A cherry on the cake

The Liouville equation has the form

$$\langle \mu, \partial_t g + \partial_x g \cdot f \rangle - \langle \nu, g(T, \cdot) \rangle = -\delta_{(0,x_0)} g := -g(0, x_0) \qquad (1)$$

Let the set W be set of measures satisfying the Liouville equation, then

$$W := \{(\mu, \nu) : \mu \geq 0, \nu \geq 0, \langle \mu, \mathbf{1} \rangle = T, \langle \nu, \mathbf{1} \rangle = 1 \\ \text{and } (\mu, \nu) \text{ solves } (1)\}$$

and the set S of generalized arcs/occupation measures

$$S := \{(\mu, \nu) : (\mu, \nu) \text{ occupation measure} \\ \text{induced by some} (x(\cdot), u(\cdot)) \text{ that solve the ODE}\}.$$

Corollary (Vinter's superposition principle)

$W = \overline{\mathrm{conv}(S)}$ *with respect to the weak* topology.*

Proof.

Occupation measures satisfy the Liouville equation, further W is convex and weak* closed, hence we have $W \supset \overline{\mathrm{conv}(S)}$. Assume we have strict inclusion. By Hahn-Banach there exist $(\mu_0, \nu_0) \in W$, $\alpha \in \mathbb{R}$ and functions (l, l_{T}) with for all $(\mu, \nu) \in \overline{\mathrm{conv}(S)}$

$$\langle (l, l_{\mathrm{T}}), (\mu_0, \nu_0) \rangle \geq \alpha > \langle (l, l_{\mathrm{T}}), (\mu, \nu) \rangle. \tag{2}$$

Let $V(\cdot, \cdot)$ denote the optimal value function of the corresponding OCP (with Young measure control $\rightsquigarrow$ no convexity assumption needed). By passing to a smooth minimizing sequence we can assume that V is smooth and solves the HJB

$$\begin{aligned} V(0, x_0) &\overset{(\mu_0,\nu_0)\text{ feasible}}{=} \int \partial_t V + \nabla V \, d\mu_0 + \int V(T, \cdot)\, d\nu_0 \\ &\overset{\text{HJB}}{\leq} \int l \, d\mu_0 + \int l_{\mathrm{T}} \, d\nu_0 \\ &= \langle (l, l_{\mathrm{T}}), (\mu_0, \nu_0) \rangle. \end{aligned}$$

□

“The Winter program”