Stochastic approximation based methods for computing the optimal thresholds in remote-state estimation with packet drops

Jhelum Chakravorty

Joint work with Jayakumar Subramanian and Aditya Mahajan

McGill University

American Control Conference
May 24, 2017

Motivation

- Sequential transmission of data
- Zero delay in reconstruction

Motivation

Applications?

- Smart grids

Motivation

Applications?

- Environmental monitoring, sensor network

Motivation

Applications?

- Internet of things

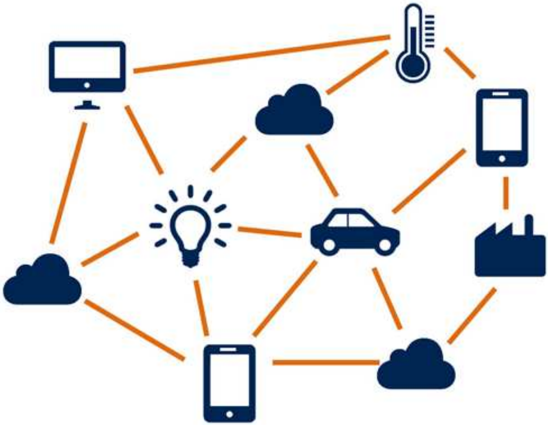

Motivation

Applications?

- Smart grids
- Environmental monitoring, sensor network
- Internet of things

Salient features

- Sensing is cheap
- Transmission is expensive
- Size of data-packet is not critical

Motivation

We study a stylized model.

Characterization of the fundamental trade-off between estimation accuracy and transmission cost!

The remote-state estimation setup

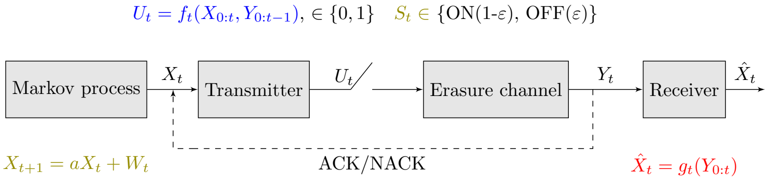

Source model $X_{t+1} = aX_t + W_t,\ W_t$ i.i.d.

The remote-state estimation setup

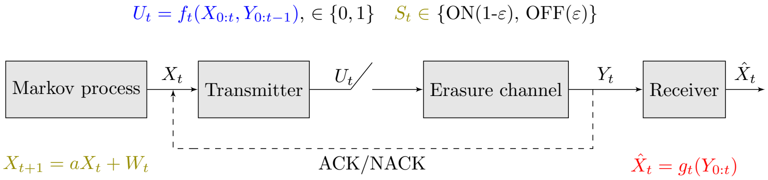

Source model $X_{t+1} = aX_t + W_t,\ W_t$ i.i.d.

- $a, X_t, W_t \in \mathbb{R}$, pdf of W_t: $\phi(\cdot)$ - **Gaussian**.

The remote-state estimation setup

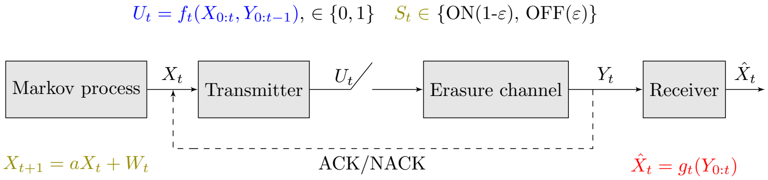

Source model $X_{t+1} = aX_t + W_t$, W_t i.i.d.

Channel model S_t i.i.d.; $S_t = 1$: channel ON, $S_t = 0$: channel OFF
Packet drop with probability ε.

The remote-state estimation setup

Transmitter $U_t = f_t(X_{0:t}, Y_{0:t-1})$ and $Y_t = \begin{cases} X_t, & \text{if } U_t S_t = 1 \\ \mathfrak{E}, & \text{if } U_t S_t = 0. \end{cases}$

Receiver $\hat{X}_t = g_t(Y_{0:t})$
Per-step distortion: $d(X_t - \hat{X}_t) = (X_t - \hat{X}_t)^2$.

Communication strategies
Transmission strategy $f = \{f_t\}_{t=0}^{\infty}$
Estimation strategy $g = \{g_t\}_{t=0}^{\infty}$

The optimization problem

Discounted setup: $\beta \in (0,1)$

- $D_\beta(f,g) := (1-\beta)\mathbb{E}^{(f,g)}\left[\sum_{t=0}^{\infty} \beta^t d(X_t - \hat{X}_t) \,\middle|\, X_0 = 0\right]$
- $N_\beta(f,g) := (1-\beta)\mathbb{E}^{(f,g)}\left[\sum_{t=0}^{\infty} \beta^t U_t \,\middle|\, X_0 = 0\right]$

Long-term average setup: $\beta = 1$

- $D_1(f,g) := \limsup_{T\to\infty} \frac{1}{T}\mathbb{E}^{(f,g)}\left[\sum_{t=0}^{T-1} d(X_t - \hat{X}_t) \,\middle|\, X_0 = 0\right]$
- $N_1(f,g) := \limsup_{T\to\infty} \frac{1}{T}\mathbb{E}^{(f,g)}\left[\sum_{t=0}^{T-1} U_t \,\middle|\, X_0 = 0\right]$

The optimization problem

Constrained performance: The Distortion-Transmission function

$$D_\beta^*(\alpha) \coloneqq D_\beta(f^*, g^*) \coloneqq \inf_{(f,g): N_\beta(f,g) \le \alpha} D_\beta(f, g), \ \beta \in (0, 1]$$

Minimize expected distortion such that expected number of transmissions is less than α

The optimization problem

Constrained performance: The Distortion-Transmission function

$$D^*_\beta(\alpha) := D_\beta(f^*, g^*) := \inf_{(f,g):N_\beta(f,g)\le\alpha} D_\beta(f,g),\ \beta \in (0,1]$$

Minimize expected distortion such that expected number of transmissions is less than α

Costly performance: Lagrange relaxation

$$C^*_\beta(\lambda) := \inf_{(f,g)} D_\beta(f,g) + \lambda N_\beta(f,g),\ \beta \in (0,1]$$

Decentralized control systems

Team: Multiple decision makers to achieve a common goal

Decentralized control systems

Pioneers:

Theory of teams

- Economics: Marschak, 1955; Radner, 1962
- Systems and control: Witsenhausen, 1971; Ho, Chu, 1972

Decentralized control systems

Pioneers:

Theory of teams

- Economics: Marschak, 1955; Radner, 1962
- Systems and control: Witsenhausen, 1971; Ho, Chu, 1972

Remote-state estimation as Team problem

- No packet drop - Marshak, 1954; Kushner, 1964; Åstrom, Bernhardsson, 2002; Xu and Hespanha, 2004; Imer and Basar, 2005; Lipsa and Martins, 2011; Molin and Hirche, 2012; Nayyar, Başar, Teneketzis and Veeravalli, 2013; D. Shi, L. Shi and Chen, 2015
- With packet drop - Ren, Wu, Johansson, G. Shi and L. Shi, 2016; Chen, Wang, D. Shi and L. Shi, 2017;
- With noise - Gao, Akyol and Başar, 2015–2017

Remote-state estimation - Steps towards optimal solution

- Establish the structure of optimal strategies (transmission and estimation)
- Computation of optimal strategies and performances

Step 1 - Structure of optimal strategies: Lipsa-Martins 2011 & Molin-Hirsche 2012 - no packet drop

Optimal estimator

Time homogeneous!

$$\hat{X}_t = g_t^*(Y_t) = g^*(Y_t) = \begin{cases} Y_t, & \text{if } Y_t \neq \mathfrak{E}; \\ a\hat{X}_{t-1}, & \text{if } Y_t = \mathfrak{E}. \end{cases}$$

Step 1 - Structure of optimal strategies: Lipsa-Martins 2011 & Molin-Hirsche 2012 - no packet drop

Optimal estimator

Time homogeneous!

$$\hat{X}_t = g_t^*(Y_t) = g^*(Y_t) = \begin{cases} Y_t, & \text{if } Y_t \neq \mathfrak{E}; \\ a\hat{X}_{t-1}, & \text{if } Y_t = \mathfrak{E}. \end{cases}$$

Optimal transmitter

$X_t \in \mathbb{R}$; U_t is threshold based action:

$$U_t = f_t^*(X_t, U_{0:t-1}) = f^*(X_t) = \begin{cases} 1, & \text{if } |X_t - a\hat{X}_t| \geq k \\ 0, & \text{if } |X_t - a\hat{X}_t| < k \end{cases}$$

Step 1 - Structure of optimal strategies: Lipsa-Martins 2011 & Molin-Hirsche 2012 - no packet drop

Optimal estimator

Time homogeneous!

$$\hat{X}_t = g_t^*(Y_t) = g^*(Y_t) = \begin{cases} Y_t, & \text{if } Y_t \neq \mathfrak{E}; \\ a\hat{X}_{t-1}, & \text{if } Y_t = \mathfrak{E}. \end{cases}$$

Optimal transmitter

$X_t \in \mathbb{R}$; U_t is threshold based action:

$$U_t = f_t^*(X_t, U_{0:t-1}) = f^*(X_t) = \begin{cases} 1, & \text{if } |X_t - a\hat{X}_t| \geq k \\ 0, & \text{if } |X_t - a\hat{X}_t| < k \end{cases}$$

Similar structural results for channel with packet drops.

Step 2 - The error process E_t

$\tau^{(k)}$: the time a packet was last received successfully.

$E_t := X_t - a^{t-\tau^{(k)}} X_{\tau^{(k)}}$, $\hat{E}_t := \hat{X}_t - a^{t-\tau^{(k)}} X_{\tau^{(k)}}$;

Step 2 - The error process E_t

$\tau^{(k)}$: the time a packet was last received successfully.

$E_t := X_t - a^{t-\tau^{(k)}} X_{\tau^{(k)}}$, $\hat{E}_t := \hat{X}_t - a^{t-\tau^{(k)}} X_{\tau^{(k)}}$;

$$d(X_t - \hat{X}_t) = d(E_t - \hat{E}_t).$$

Step 2 - The error process E_t

$\tau^{(k)}$: the time a packet was last received successfully.

$E_t := X_t - a^{t-\tau^{(k)}} X_{\tau^{(k)}}$, $\hat{E}_t := \hat{X}_t - a^{t-\tau^{(k)}} X_{\tau^{(k)}}$;

$$= X_t - a(\hat{X}_{t-1} - \hat{E}_{t-1})$$

$$= \begin{cases} aE_{t-1} + W_{t-1}, & \text{if } Y_t = \mathfrak{E} \\ W_t, & \text{if } Y_t \neq \mathfrak{E} \end{cases}$$

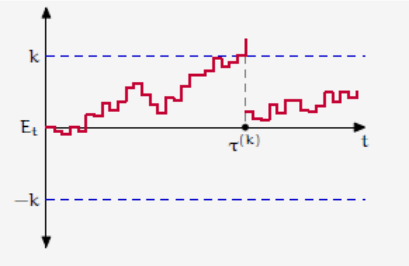

Performance evaluation - JC-AM TAC '17, NecSys '16

$$f^{(k)}(e) = \begin{cases} 1, & \text{if } |e| \geq k \\ 0, & \text{if } |e| < k \end{cases}$$

Till first successful reception

$$L_\beta^{(k)}(0) := \mathbb{E}\Big[\sum_{t=0}^{\tau^{(k)}-1} \beta^t d(E_t) \;\Big|\; E_0 = 0 \Big]$$

$$M_\beta^{(k)}(0) := \mathbb{E}\Big[\sum_{t=0}^{\tau^{(k)}-1} \beta^t \;\Big|\; E_0 = 0 \Big]$$

$$K_\beta^{(k)}(0) := \mathbb{E}\Big[\sum_{t=0}^{\tau^{(k)}} \beta^t U_t \;\Big|\; E_0 = 0 \Big]$$

Performance evaluation - JC-AM TAC '17, NecSys '16

$$f^{(k)}(e) = \begin{cases} 1, & \text{if } |e| \geq k \\ 0, & \text{if } |e| < k \end{cases}$$

E_t is regenerative process

Renewal relationships

$$D^{(k)}_{\beta}(0) := D_{\beta}(f^{(k)}, g^*) = \frac{L^{(k)}_{\beta}(0)}{M^{(k)}_{\beta}(0)}$$

$$N^{(k)}_{\beta}(0) := N_{\beta}(f^{(k)}, g^*) = \frac{K^{(k)}_{\beta}(0)}{M^{(k)}_{\beta}(0)}$$

Computation of D, N

$$L_\beta^{(k)}(e) = \begin{cases} \varepsilon\Big[d(e) + \beta \displaystyle\int_{n\in\mathbb{R}} \phi(n-ae)L_\beta^{(k)}(n)dn\Big], & \text{if } |e| \geq k \\ d(e) + \beta \displaystyle\int_{n\in\mathbb{R}} \phi(n-ae)L_\beta^{(k)}(n)dn, & \text{if } |e| < k, \end{cases}$$

Computation of D, N

$$L_\beta^{(k)}(e) = \begin{cases} \varepsilon\Big[d(e) + \beta \int_{n\in\mathbb{R}} \phi(n-ae)L_\beta^{(k)}(n)dn\Big], & \text{if } |e| \geq k \\ d(e) + \beta \int_{n\in\mathbb{R}} \phi(n-ae)L_\beta^{(k)}(n)dn, & \text{if } |e| < k, \end{cases}$$

$M_\beta^{(k)}(e)$ and $K_\beta^{(k)}(e)$ defined in a similar way.

Computation of D, N

$$L_\beta^{(k)}(e) = \begin{cases} \varepsilon\left[d(e) + \beta \int_{n\in\mathbb{R}} \phi(n - ae) L_\beta^{(k)}(n)dn\right], & \text{if } |e| \geq k \\ d(e) + \beta \int_{n\in\mathbb{R}} \phi(n - ae) L_\beta^{(k)}(n)dn, & \text{if } |e| < k, \end{cases}$$

- $\varepsilon = 0$: Fredholm integral equations of second kind - bisection method to compute optimal threshold
- $\varepsilon \neq 0$: Fredholm-like equation; discontinuous kernel, infinite limit - analytical methods difficult

Optimality condition (JC & AM: TAC'17, NecSys '16)

$D_\beta^{(k)}$,$N_\beta^{(k)}$, $C_\beta^{(k)}$ - differentiable in k

Theorem - costly communication

If (k, λ) satisfies $\partial_k D_\beta^{(k)} + \lambda \partial_k N_\beta^{(k)} = 0$, then, $(f^{(k)}, g^*)$ optimal for costly comm. with cost λ.

Optimality condition (JC & AM: TAC'17, NecSys '16)

$D_\beta^{(k)}$, $N_\beta^{(k)}$, $C_\beta^{(k)}$ - differentiable in k

Theorem - costly communication

If (k, λ) satisfies $\partial_k D_\beta^{(k)} + \lambda \partial_k N_\beta^{(k)} = 0$, then, $(f^{(k)}, g^*)$ optimal for costly comm. with cost λ.

$C_\beta^*(\lambda) := C_\beta(f^{(k)}, g^*; \lambda)$ is continuous, increasing and concave in λ.

Optimality condition (JC & AM: TAC'17, NecSys '16)

$D_\beta^{(k)}, N_\beta^{(k)}$, $C_\beta^{(k)}$ - differentiable in k

Theorem - costly communication

If (k, λ) satisfies $\partial_k D_\beta^{(k)} + \lambda \partial_k N_\beta^{(k)} = 0$, then, $(f^{(k)}, g^*)$ optimal for costly comm. with cost λ.

$C_\beta^*(\lambda) := C_\beta(f^{(k)}, g^*; \lambda)$ is continuous, increasing and concave in λ.

Theorem - constrained communication

$k_\beta^*(\alpha) := \{k \,:\, N_\beta^{(k)}(0) = \alpha\}$. $(f^{k_\beta^*(\alpha)}, g^*)$ is optimal for the optimization problem with constraint $\alpha \in (0, 1)$.

$D_\beta^*(\alpha) := D_\beta(f^{(k)}, g^*)$ is continuous, decreasing and convex in α.

Main results

Computation of optimal thresholds

Difficulty

- Numerically compute $L_\beta^{(k)}$, $M_\beta^{(k)}$ and $K_\beta^{(k)}$; use renewal relationship to compute $C_\beta^{(k)}$ and $D_\beta^{(k)}$.
- Need to solve Fredholm-like integral - computationally difficult.

Computation of optimal thresholds

Difficulty

- Numerically compute $L_\beta^{(k)}$, $M_\beta^{(k)}$ and $K_\beta^{(k)}$; use renewal relationship to compute $C_\beta^{(k)}$ and $D_\beta^{(k)}$.
- Need to solve Fredholm-like integral - computationally difficult.

Simulation based approach

- Two main approaches - Monte Carlo (MC) and Temporal Difference (TD)
  - MC - High variance due to one sample path; low bias
  - TD - Low variance due to *bootstrapping*; high bias

Computation of optimal thresholds

Difficulty

- Numerically compute $L_\beta^{(k)}$, $M_\beta^{(k)}$ and $K_\beta^{(k)}$; use renewal relationship to compute $C_\beta^{(k)}$ and $D_\beta^{(k)}$.
- Need to solve Fredholm-like integral - computationally difficult.

Simulation based approach

- Two main approaches - Monte Carlo (MC) and Temporal Difference (TD)
  - MC - High variance due to one sample path; low bias
  - TD - Low variance due to *bootstrapping*; high bias
- Exploit regenerative property of the underlying state (error) process
- Renewal Monte Carlo (RMC) - low variance (independent sample paths from renewal) and low bias (since MC)

Computation of optimal thresholds

Difficulty

- Numerically compute $L_\beta^{(k)}$, $M_\beta^{(k)}$ and $K_\beta^{(k)}$; use renewal relationship to compute $C_\beta^{(k)}$ and $D_\beta^{(k)}$.
- Need to solve Fredholm-like integral - computationally difficult.

Key idea

- **Renewal Monte Carlo**
  - Pick a k, compute sample values L, M, K till first successful reception
  - Sample average to compute $L_\beta^{(k)}$, $M_\beta^{(k)}$, $K_\beta^{(k)}$.
- **Stochastic approximation** techniques to compute optimal k.

Computation of optimal thresholds

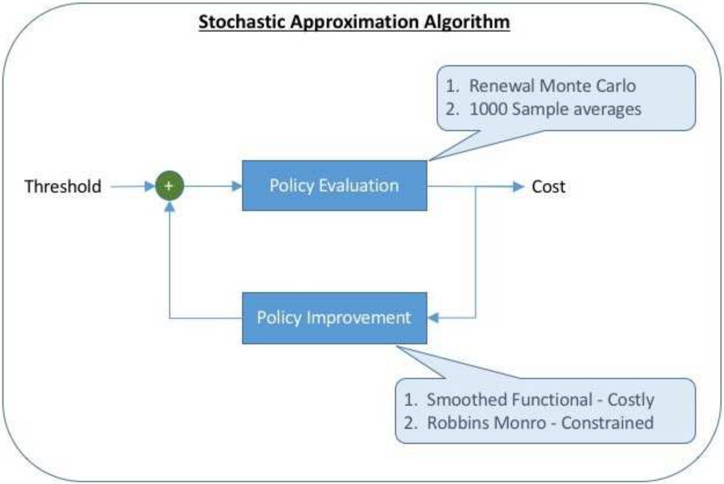

Computation of optimal thresholds

Key steps of the algorithms

- **Noisy policy evaluation** - MC until successful reception: constitutes one episode; sample average over few episodes to find $\hat{L}$, $\hat{M}$, $\hat{K}$ and hence $\hat{C}$ and $\hat{D}$.

- **Policy improvement - Smoothed Functional**

$$\hat{k}_{i+1} = \hat{k}_i - \gamma_i \frac{\eta}{2\tilde{\beta}} \left(\hat{C}(\hat{k}_i + \tilde{\beta}\eta) - \hat{C}(\hat{k}_i - \tilde{\beta}\eta) \right)$$

- **Policy improvement - Robbins-Monro**

$$\hat{k}_{i+1} = \hat{k}_i - \gamma_i(\alpha \hat{M} - \hat{K}).$$

Validation of simulation results

Results validated by comparing with analytical results of no packet-drop case: JC-AM, TAC '17.

- Costly performance - Error in k^*: $10^{-2} - 10^{-3}$; Error in C^*: $10^{-4} - 10^{-5}$
- Constrained performance - Error in k^*: 10^{-3}; Error in D^*: $10^{-3} - 10^{-5}$

Optimal thresholds from simulations

Costly performance:

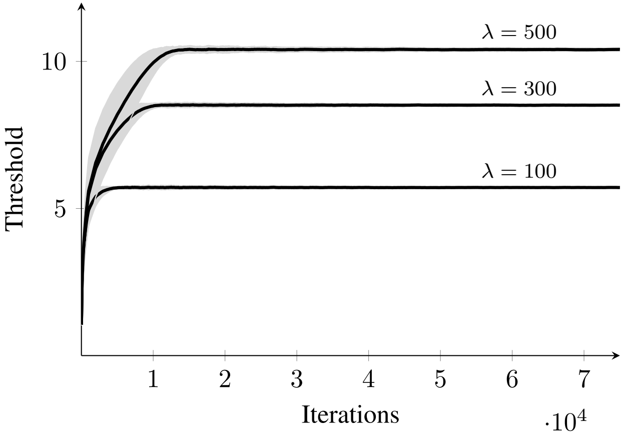

Figure: Costly communication: $\beta = 0.9$, $\varepsilon = 0.3$.

Optimal thresholds from simulations

Constrained performance:

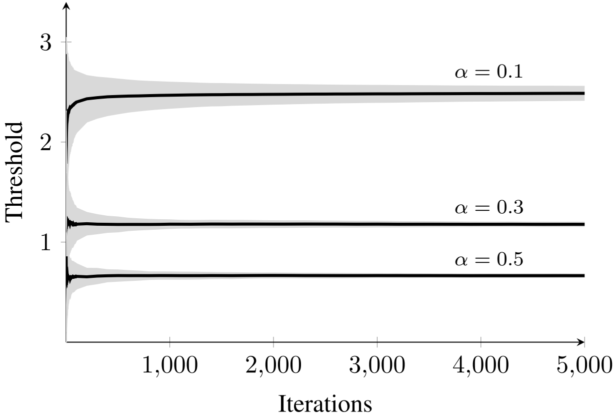

Figure: Constrained communication using RM: $\beta = 0.9$, $\varepsilon = 0.3$.

Optimal trade-off between distortion and communication cost

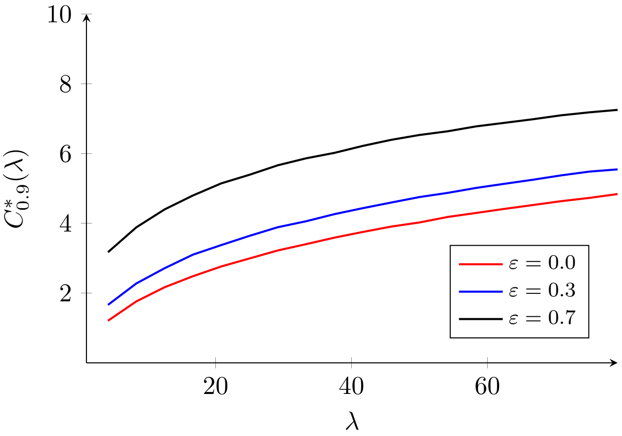

Figure: Costly communication: $\beta = 0.9$, $\varepsilon \in \{0, 0.3, 0.7\}$.

Optimal trade-off between distortion and communication cost

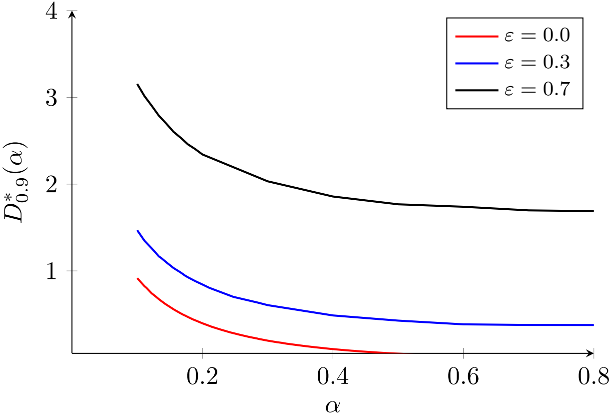

Figure: Constrained communication: $\beta = 0.9$, $\varepsilon \in \{0, 0.3, 0.7\}$.

Future work

- Markovian erasure channel -
  - Thresholds at t are function of channel-state at $t-1$
- Higher dimension -
  - $X_t \in \mathbb{R}^m$ is ASU $\overset{?}{\Longrightarrow}$ $AX_t + W_t$ is ASU
  - Notion of stochastic dominance in higher dimension

Thank you!