exact particle flow for nonlinear filters

Fred Daum

16 August 2010

nonlinear filtering problem

$$dx/dt = f(x, t) + G(x,t)\, dw/dt$$

continuous time dynamics

x = d-dimensional state vector
t = time
$w(t)$ = process noise vector

discrete time measurements

$$z(t_k) = h(x(t_k), t_k, v_k)$$

$z(t_k)$ = m-dimensional measurement vector
t_k = time of k^{th} measurement
v_k = measurement noise vector

$p(x, t \mid Z_k)$ = probability density of x at time t given Z_k

Z_k = set of all measurements up to & including time t_k

Curse of dimensionality for classic particle filter

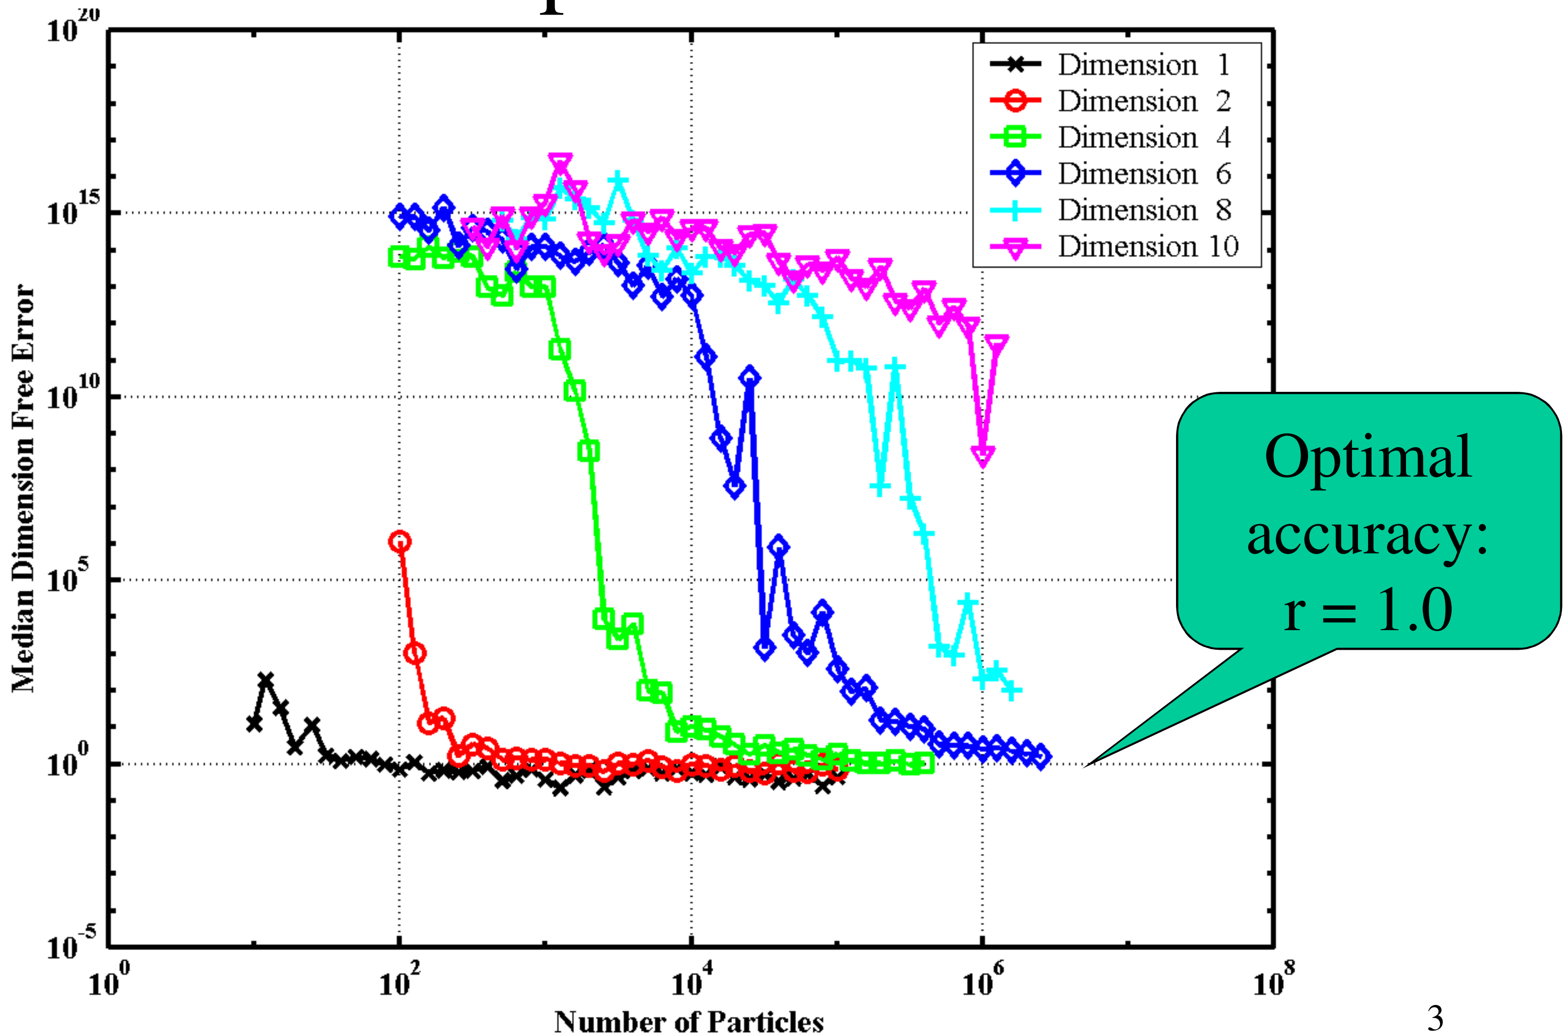

Exact flow PF is many orders of magnitude faster per particle than classic PF and inexact log-homotopy PF

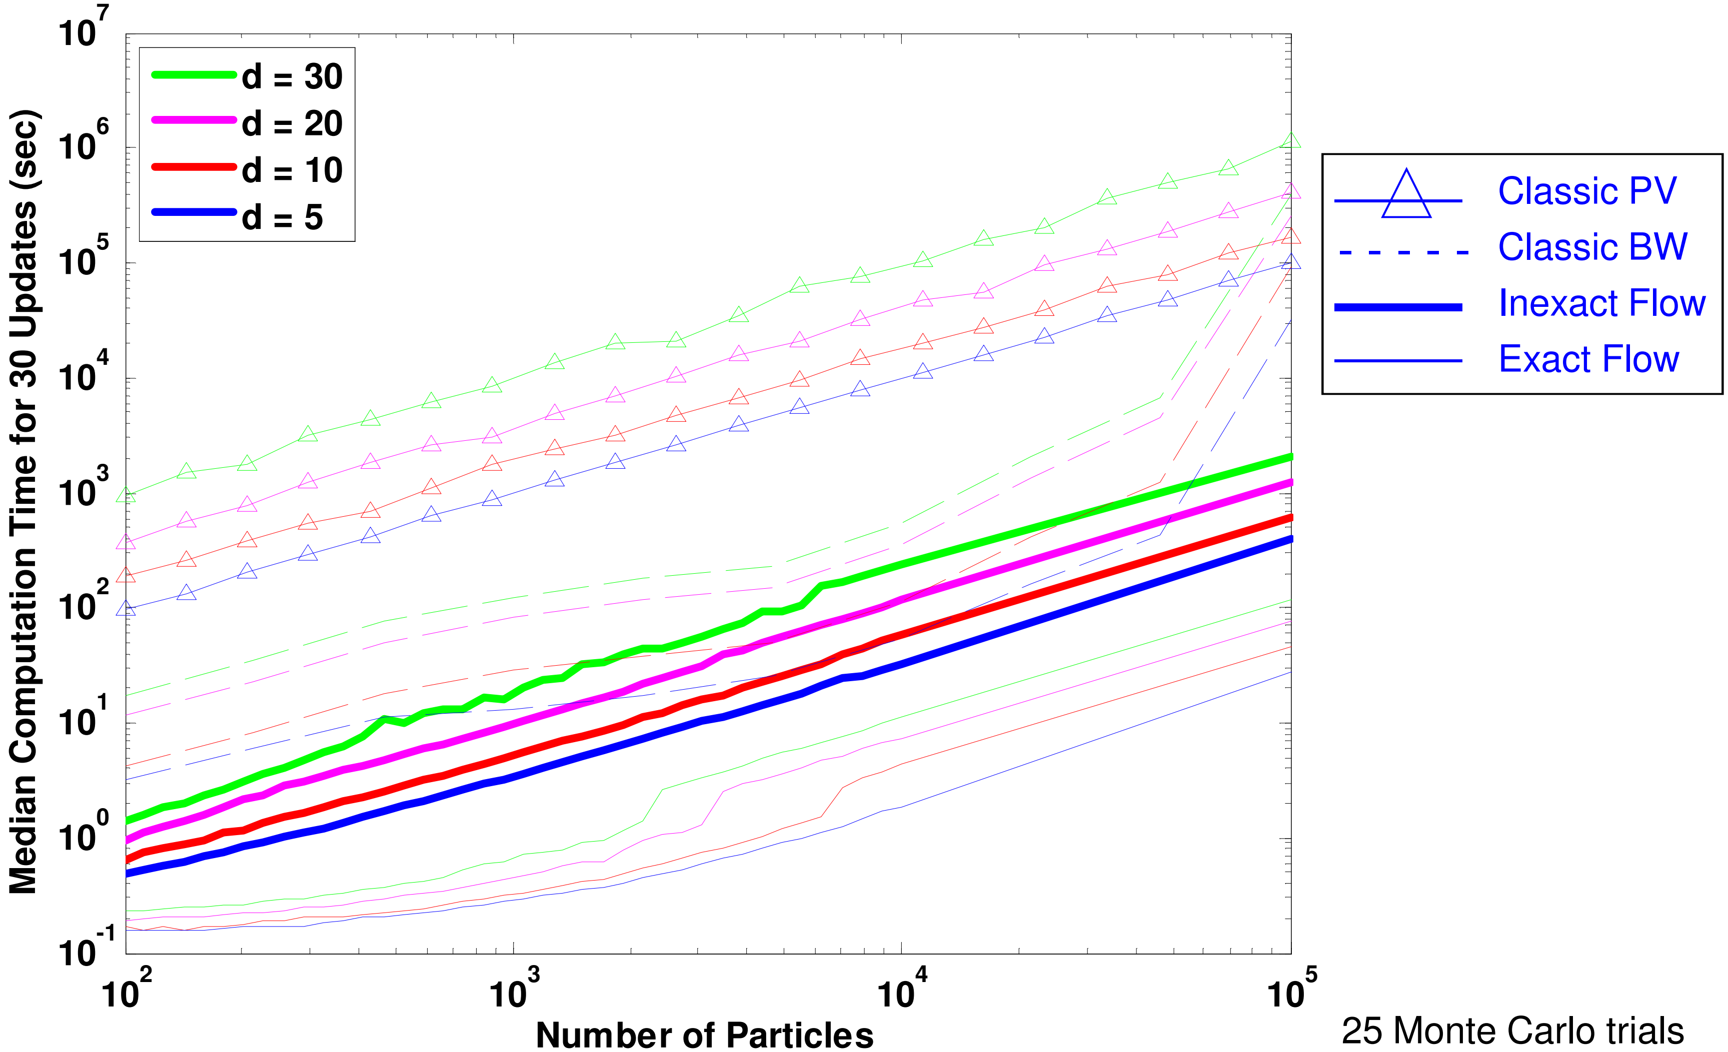

25 Monte Carlo trials

* Intel CoreI 2 CPU, 1.86GHz, 0.98GB of RAM, PC-MATLAB version 7.7

Nonlinear Filter

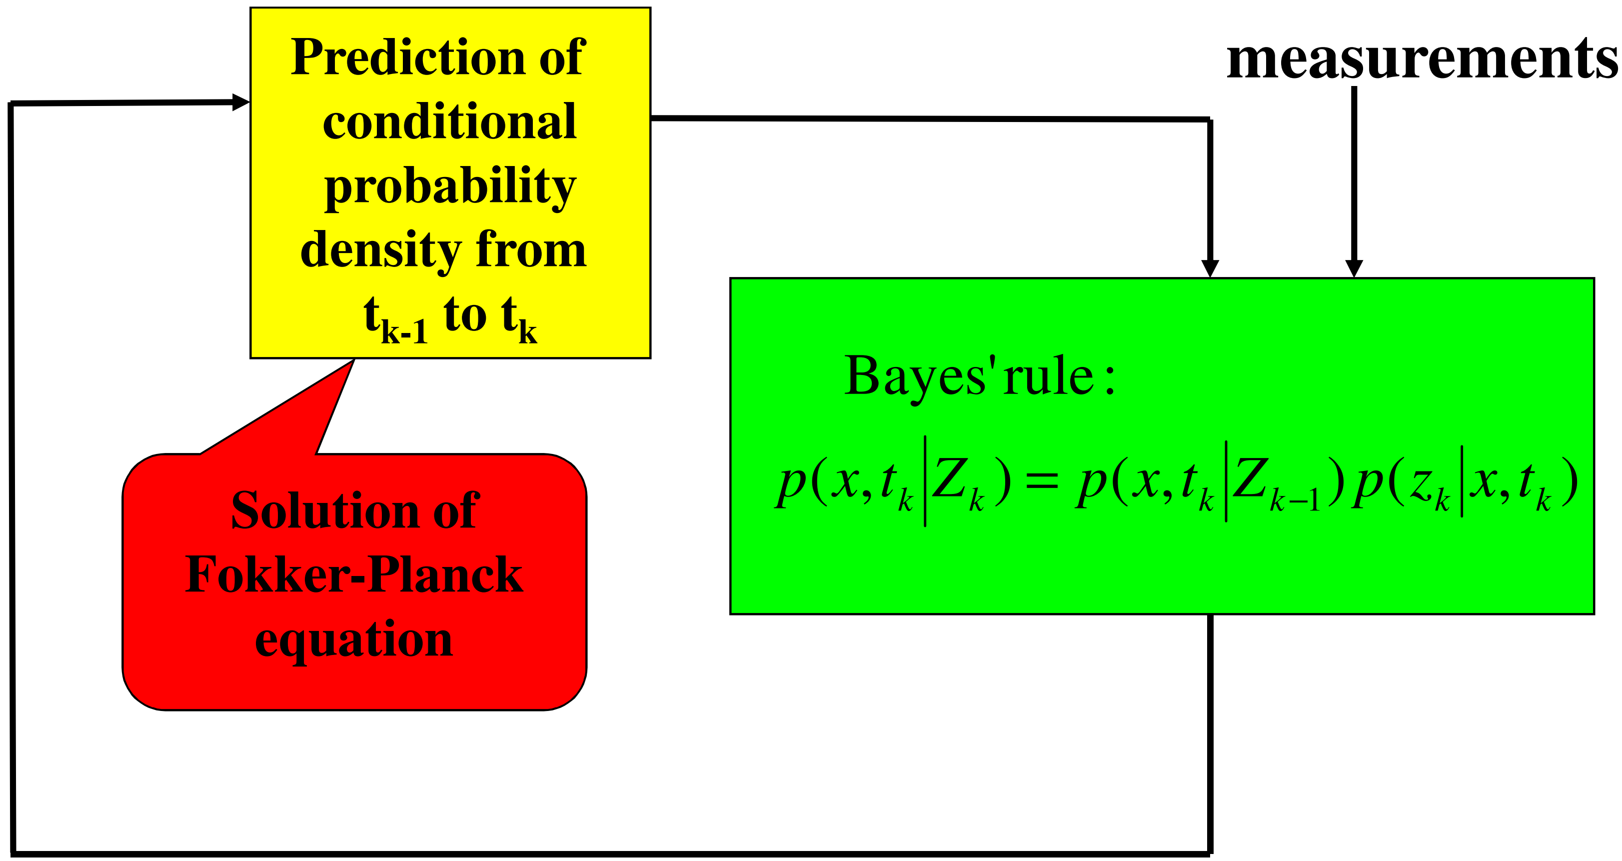

Particle degeneracy*

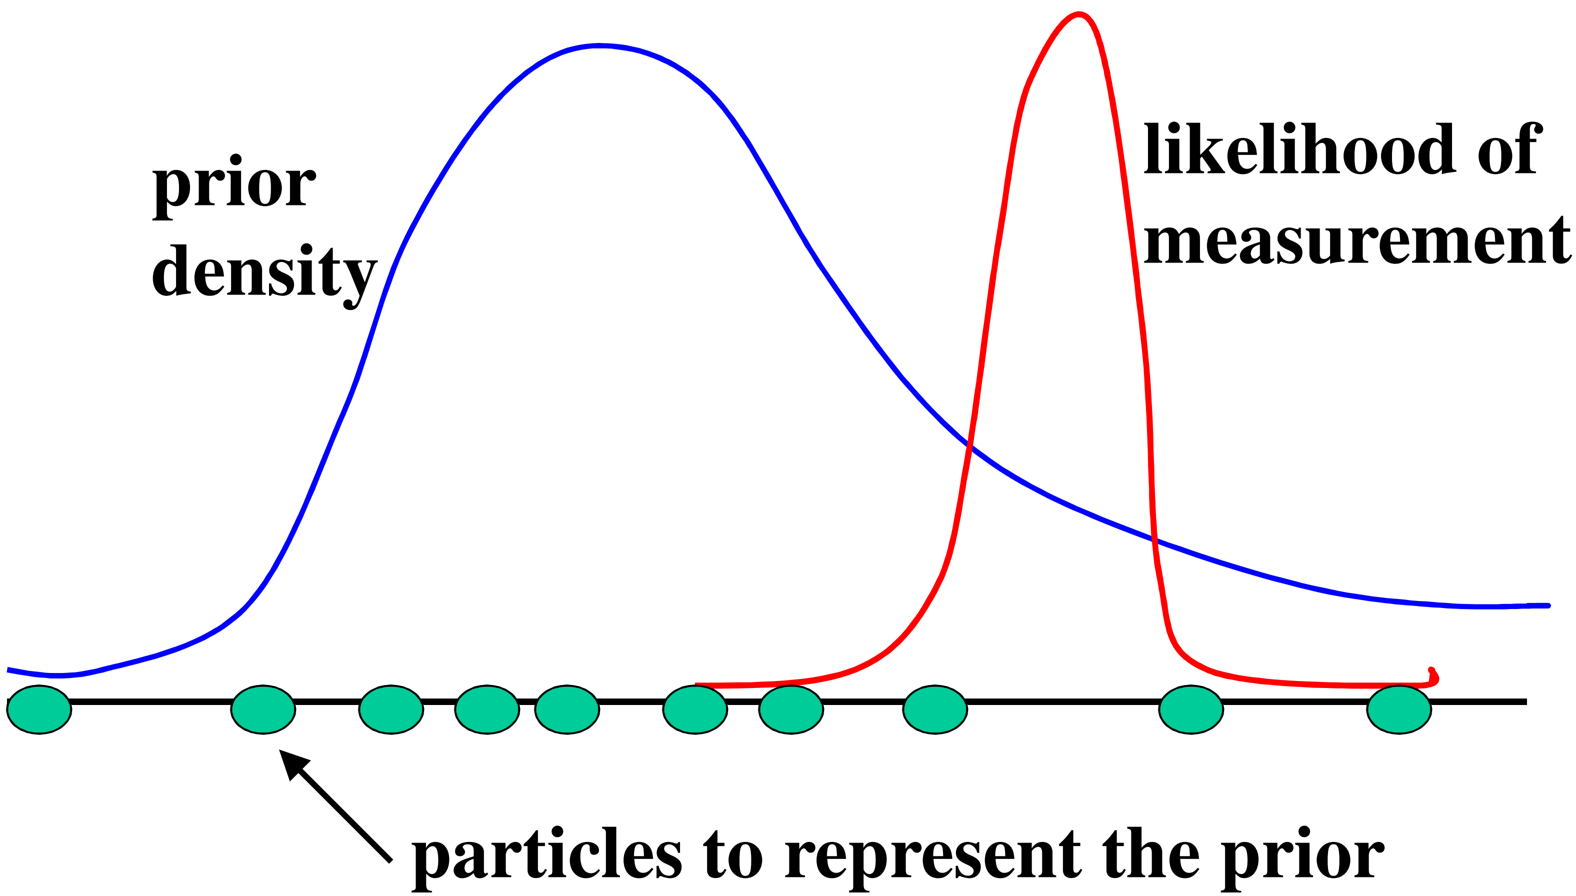

*note two simple ideas for low dimension (more particles & sample from likelihood)

Chicken & egg problem

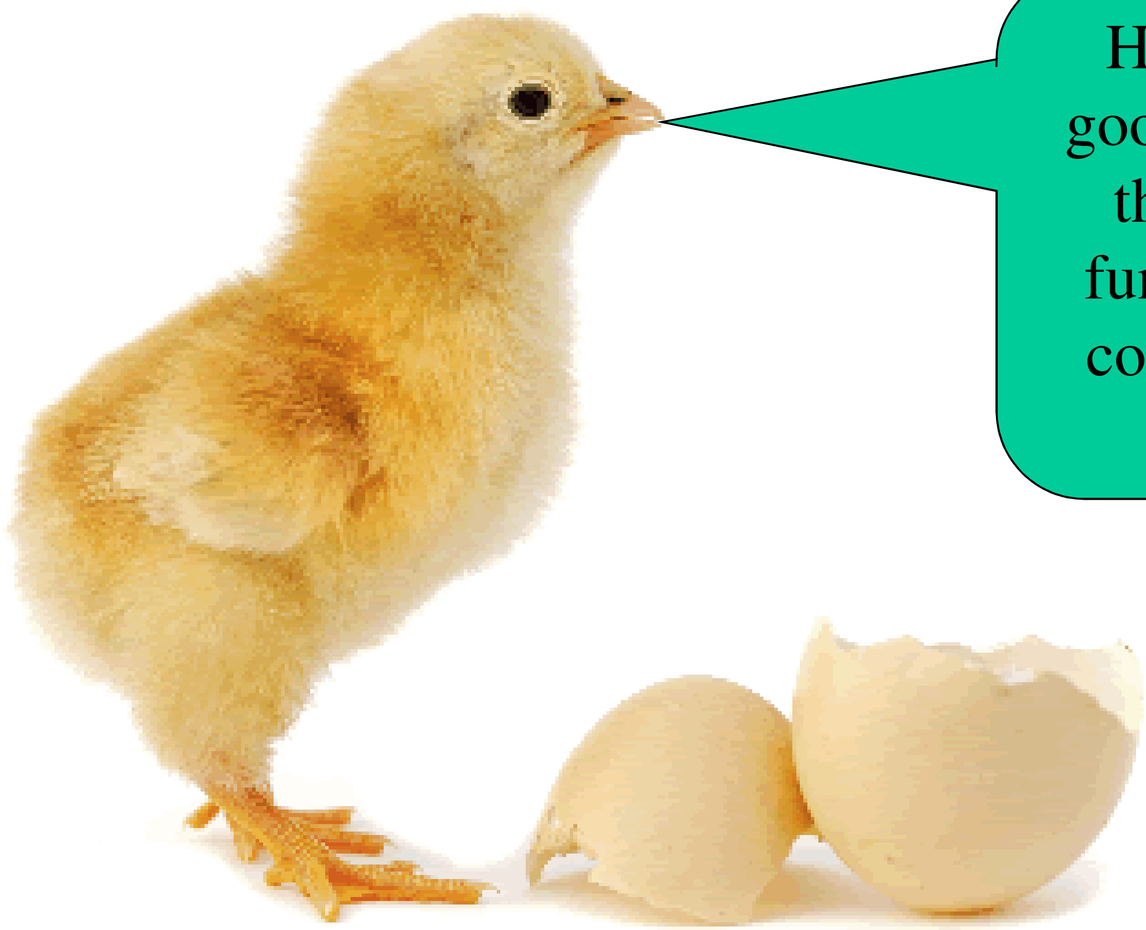

How do you pick a good way to represent the product of two functions before you compute the product itself?

Induced flow of particles

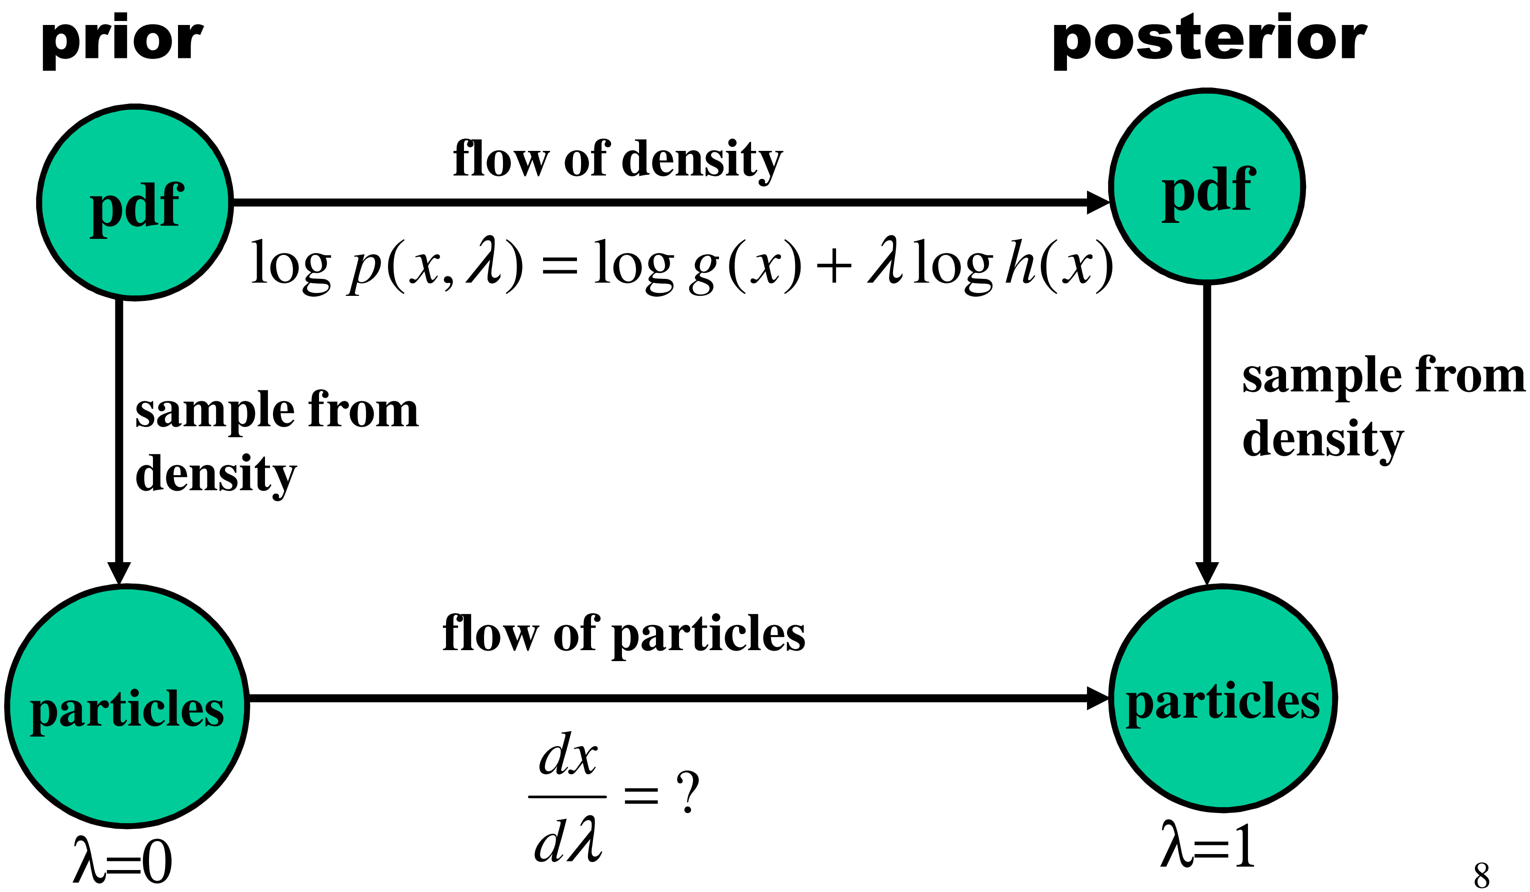

Fundamental PDE for exact particle flow:

$$\frac{dx}{d\lambda} = f(x,\lambda)$$

$$\frac{\partial p(x,\lambda)}{\partial \lambda} = -Tr\left[\frac{\partial (pf)}{\partial x}\right]$$

Fokker-Planck equation

$$\frac{\partial \log p(x,\lambda)}{\partial \lambda} p(x,\lambda) = -Tr\left[\frac{\partial (pf)}{\partial x}\right]$$

assume log-homotopy

$$\log p(x,\lambda) = \log g(x) + \lambda \log h(x)$$

$$\log h(x) p(x,\lambda) = -p(x,\lambda) Tr\left[\frac{\partial f}{\partial x}\right] - \frac{\partial p}{\partial x} f$$

$$\log(h) = -div(f) - \frac{\partial \log p}{\partial x} f$$

First order linear underdetermined PDE in f(x,λ)

method	uniqueness	computational complexity	accuracy
1. generalized inverse of differential operator	minimum norm*	fast Poisson solver in d-dimensions or homotopy for generalized inverse	?
2. Poisson's equation	gradient of potential* (assume irrotational flow)	fast Poisson solver in d-dimensions	?
3. generalized inverse of gradient of log-homotopy	assume incompressible flow	fast (but need to compute the gradient from random points)	good to excellent depending on the problem
4. most general solution	most robustly stable filter or minimize obstructions & singularities in flow or random pick	fast Poisson solver in d-dimensions	good to excellent depending on the problem
5. separation of variables (Gaussian)	specific form (analytic)	extremely fast (formula)	excellent for many examples; automatic stability of flow!
6. separation of variables (exponential family)	specific form	very fast (formula)	should be excellent for certain examples
7. variational formulation (Gauss & Hertz)	convex function minimization	ODEs	?
8. optimal control formulation	convex functional minimization	Euler-Lagrange PDEs (or maybe ODES for nice problem)	?
9. direct integration (of first order linear PDE in divergence form)	choice of d-1 arbitrary functions	d decoupled one-dimensional integrals	?
10. method of characteristics	more conditions (e.g., curl = 0 & chain rule) like Jacobi's method	ODEs from chain rule	??
11. another homotopy (like Gromov's h-principle)	initial condition of ODE & uniqueness of sol. to ODE	ODEs from homotopy	???

Exact particle flow for Gaussian densities:

$$\frac{dx}{d\lambda} = f(x,\lambda)$$

$$\log(h) = -div(f) - \frac{\partial \log p}{\partial x} f$$

for g & h Gaussian, we can solve for f exactly:

$$f = Ax + b$$

$$A = -\frac{1}{2} PH^T \left[\lambda HPH^T + R\right]^{-1} H$$

$$b = (I + 2\lambda A)\left[(I + \lambda A)PH^T R^{-1} z + A\bar{x}\right]$$

Exact particle flow & Poisson's equation:

$$\frac{dx}{d\lambda} = f(x,\lambda)$$

$$\frac{\partial p(x,\lambda)}{\partial \lambda} = -Tr\left[\frac{\partial (pf)}{\partial x}\right]$$

$$\frac{\partial \log p(x,\lambda)}{\partial \lambda} p(x,\lambda) = -Tr\left[\frac{\partial (pf)}{\partial x}\right]$$

$$\log p(x,\lambda) = \log g(x) + \lambda \log h(x)$$

$$\log h(x) p(x,\lambda) = -Tr\left[\frac{\partial q(x,\lambda)}{\partial x}\right] = -div(q)$$

divergence form of PDE

obviously there is no unique solution,

so pick the unique minimum norm solution:

$$\mathrm{q}(\mathrm{x},\lambda) = \frac{\partial \mathrm{V}(\mathrm{x},\lambda)}{\partial \mathrm{x}}$$

$$Tr\left[\frac{\partial^2 V(x,\lambda)}{\partial x^2}\right] = -\log h(x) p(x,\lambda)$$

Poisson's equation

$$\frac{dx}{d\lambda} = f(x,\lambda) = \frac{\partial V(x,\lambda)}{\partial x} / p(x,\lambda)$$

Solving Poisson's equation:

$$\frac{dx}{d\lambda} = f(x,\lambda) = \left[\frac{\partial V(x,\lambda)}{\partial x}\right]^T / p(x,\lambda)$$

$$Tr\left[\frac{\partial^2 V(x,\lambda)}{\partial x^2}\right] = -\log h(x)\,p(x,\lambda)$$

Poisson's equation

$$V(x,\lambda) = \int \log h(y)\,p(y,\lambda)\frac{c}{\|x-y\|^{d-2}}dy$$

in which

$$\mathrm{c} = \Gamma(\frac{\mathrm{d}}{2}-1)/4\pi^{d/2}$$

integration by parts yields:

$$\frac{\partial \mathrm{V}(\mathrm{x},\lambda)}{\partial \mathrm{x}} = -\int \frac{\partial \log h(y)\,p(y,\lambda)}{\partial y}\frac{c}{\|x-y\|^{d-2}}dy$$

or without integration by parts:

$$\frac{\partial \mathrm{V}(\mathrm{x},\lambda)}{\partial \mathrm{x}} = \int \log h(y)\,p(y,\lambda)\frac{c(2-d)(x-y)^T}{\|x-y\|^{d}}dy$$

direct integration of fundamental PDE:

use the divergence form of the PDE :

$$div(q(x,\lambda)) = Tr\left[\frac{\partial q(x,\lambda)}{\partial x}\right] = \eta(x,\lambda)$$

$$\eta = \frac{\partial q_1}{\partial x_1} + \frac{\partial q_2}{\partial x_2} + \ldots + \frac{\partial q_d}{\partial x_d}$$

$$\frac{\partial q_j}{\partial x_j} = \eta - \sum_{k \neq j}^{d} \frac{\partial q_k}{\partial x_k}$$

$$q_j(x) = \int^{x_j} \left[\eta(x) - \theta_j(x)\right] dx_j$$

pick for best stability of particle flow

$$\theta_j(x) = \sum_{k \neq j}^{d} \frac{\partial q_k}{\partial x_k} = \text{arbitrary function (except for compatibility conditions)}$$

Assuming regularity conditions on η and Ω, a solution for $q(x, \lambda)$ exists

iff $\int_{\Omega} \eta(x) dx = 0.$

more details of direct integration:

$$div(q) = \eta(x)$$

$$q_k = \int^{x_k} \left[\eta(x) - \rho(x_k) \int_{\Omega_k} \eta(x) dx_k \right] dx_k \text{ for k} \geq 2$$

in which

$\rho(\mathrm{x_k}) =$ arbitrary function such that

$$\int_{\Omega_k} \rho(\mathrm{x_k}) dx_k = 1$$

assuming smooth functions with compact support,

and Ω is bounded, open, connected smooth set,

a necessary & sufficient condition for the existence

of a solution q(x) is that $\int_{\Omega} \eta(\mathrm{x})\mathrm{dx} = 0$

Most general solution for exact flow:

$$\frac{dx}{d\lambda} = f(x, \lambda)$$

$$\log(h) = -div(f) - \frac{\partial \log p}{\partial x} f$$

fundamental PDE for exact particle flow

the most general solution is:

$$f = -C^{\#} \log h + (I - C^{\#} C) y$$

in which C is a linear differential operator:

$$C = \frac{\partial \log p}{\partial x} + div$$

$C^{\#}$ = generalized inverse of C

y = arbitrary d-dimensional vector

Could pick y to robustly stabilize the filter or avoid obstructions to particle flow*

Effect of divergence of f:

$$\frac{dx}{d\lambda} = f(x, \lambda)$$

$$\log(h) = -div(f) - \frac{\partial \log p}{\partial x} f$$

suppose that div(f) is given

$$f = -\left[\frac{\partial \log p}{\partial x}\right]^{\#} \left[\log h + div(f)\right]$$

if div(f) = 0 we get incompressible flow

effect of div(f) is to change the speed of particle flow

Nonlinear filter performance (accuracy wrt optimal & computational complexity)

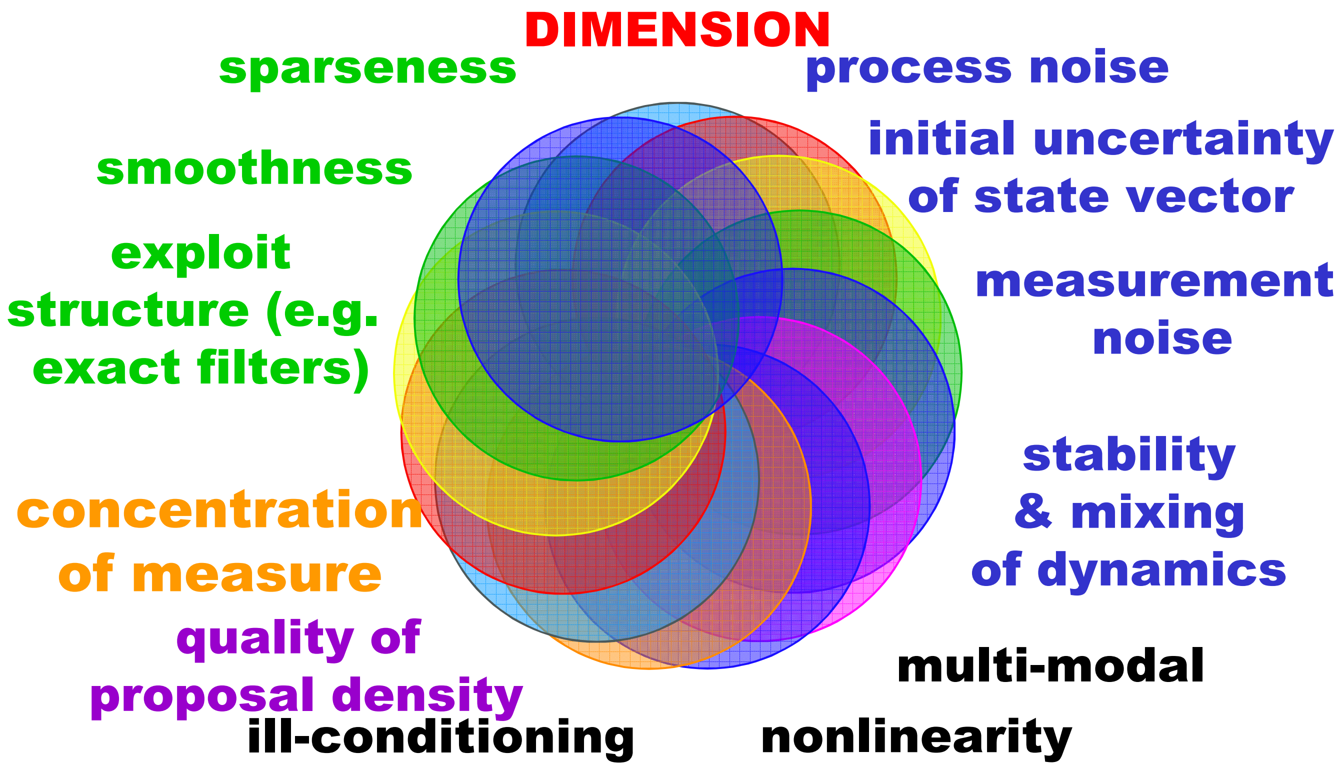

Variation in Initial Uncertainty

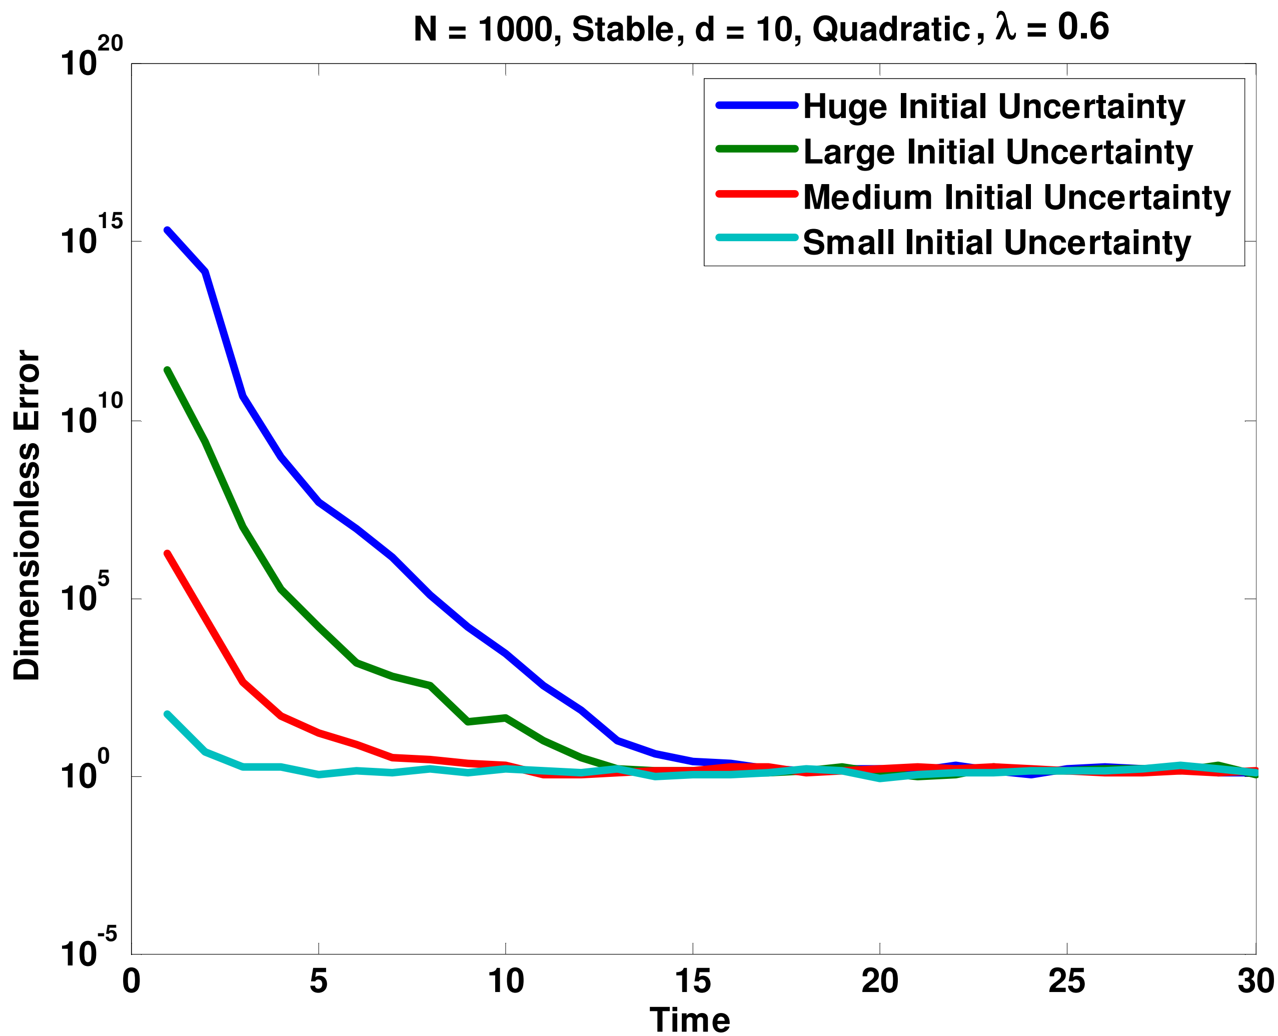

25 Monte Carlo Trials

Variation in Eigenvalues

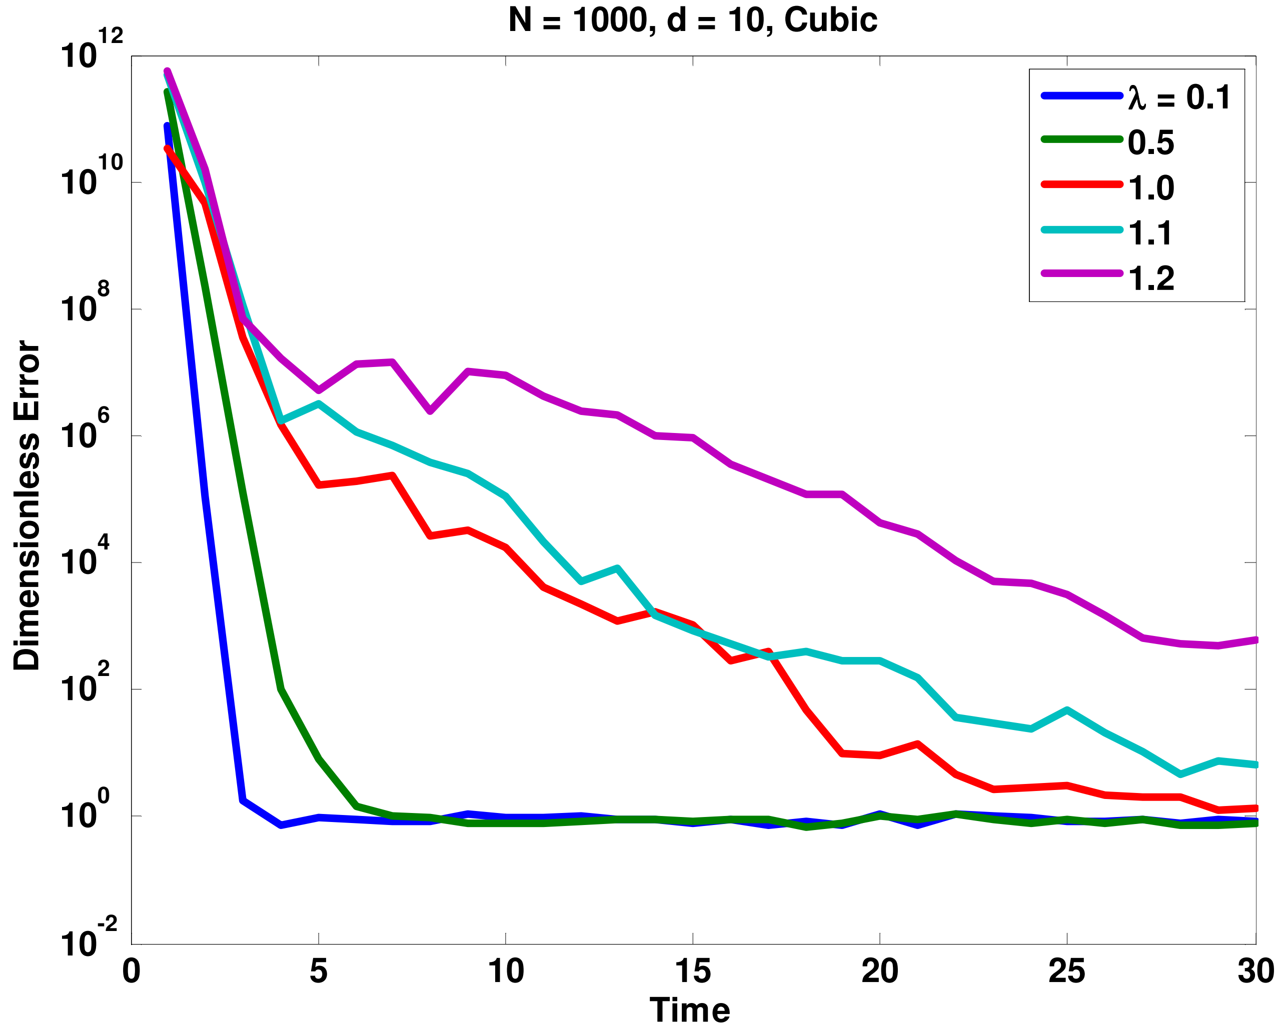

25 Monte Carlo Trials

Variation in Dimension

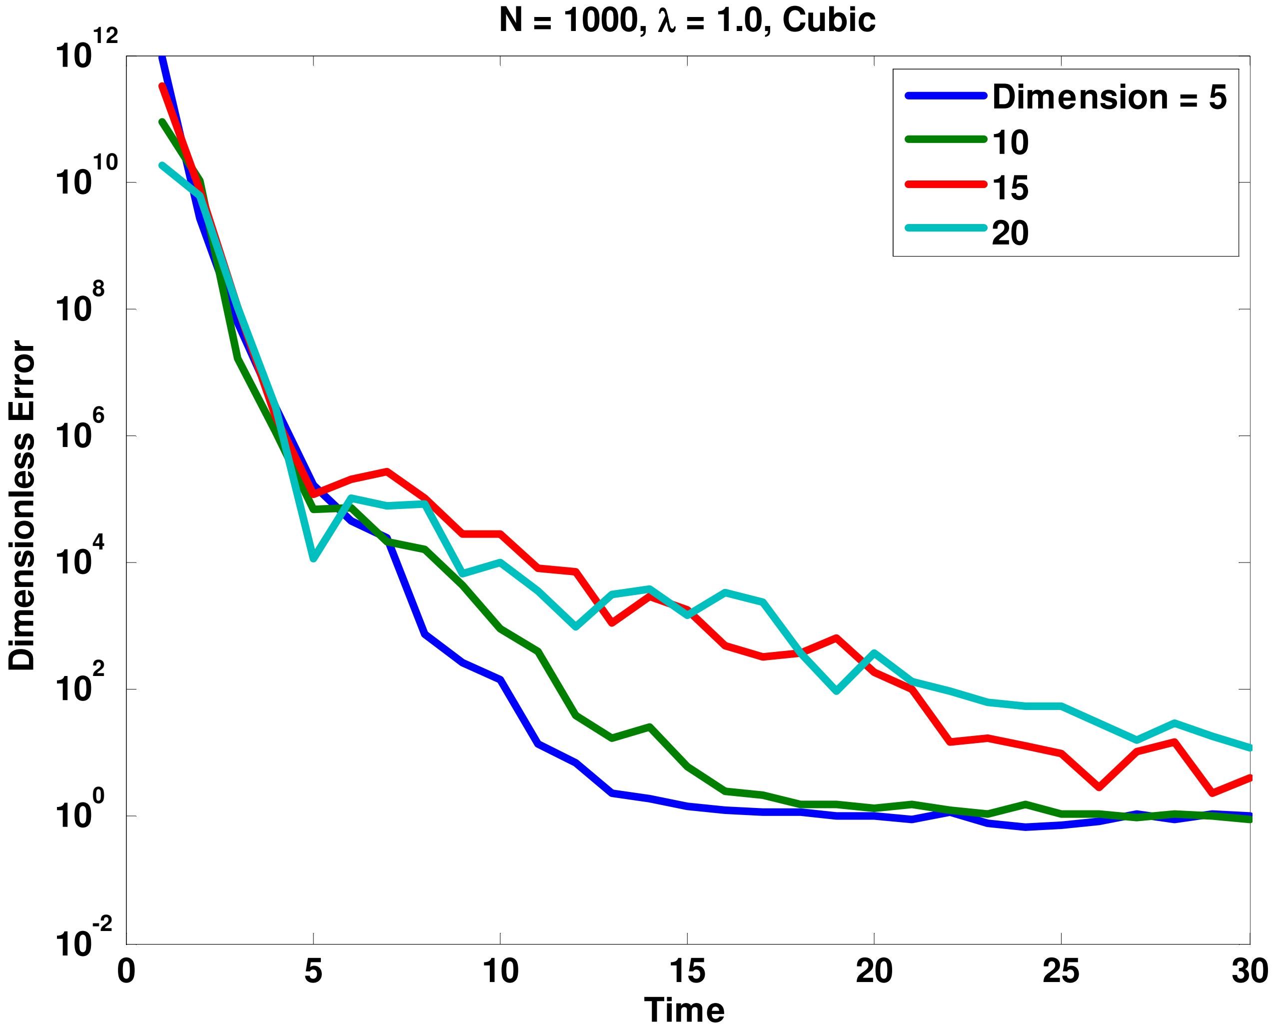

25 Monte Carlo Trials

Applications of particle flow induced by log-homotopy

- Particle filters
- Bayesian decision problems
- Bayesian estimation problems
- Multiplication of functions in high dimensions

History of Mathematics

1. Creation of the integers

2. Invention of counting

3. Invention of addition as a fast method of counting

4. Invention of multiplication as a fast method of addition

5. Invention of particle flow as a fast method of multiplication*

exact particle flow

- orders of magnitude faster than standard particle filters
- orders of magnitude more accurate than the extended Kalman filter for difficult nonlinear problems
- solves particle degeneracy problem using exact particle flow induced by log-homotopy (new theory)
- no resampling of particles
- no proposal density
- no importance sampling
- unnormalized log probability density
- exploits smoothness & regularity of densities

BACKUP

Type of nonlinear filter	Statistics computed	Computational Complexity	Estimation accuracy	Representation of probability density
Extended Kalman filters	mean vector & covariance matrix	d^3	sometimes good but often highly suboptimal	mean vector & covariance matrix
Unscented Kalman filters	mean vector & covariance matrix	d^3	sometimes better than EKF but sometimes worse	mean vector & covariance matrix
Batch least squares	mean vector & covariance matrix	d^3	sometimes better than EKF but sometimes worse	mean vector & covariance matrix
Numerical solution of Fokker-Planck PDE	full conditional probability density of state	curse of dimensionality*	optimal*	points in state space and/or smooth functions
Particle filters	full conditional probability density of state	curse of dimensionality**	optimal**	particles
Exact recursive filters (Kalman, Beneš, Daum, Wonham, Yau)	full conditional probability density of state	polynomial in d	optimal for special problems	sufficient statistics

difficulties for exact finite dimensional filters vs. particle filters

	Bayes' update of conditional density of x	prediction of conditional density of x with time
1. exact filters (e.g., Daum 1986)	easy	hard
2. particle filters	hard	easy
3. hybrid of exact & particle filters	?	?

Oh's Formula for Monte Carlo errors

$$\sigma^2 \approx \left\{ \left[\frac{1+k}{\sqrt{1+2k}} \right] \exp\left[\frac{\varepsilon^2}{1+2k} \right] \right\}^d / N$$

Assumptions:

(1) Gaussian density (zero mean & unit covariance matrix)

(2) d-dimensional random variable

(3) Proposal density is also Gaussian with mean ε and covariance matrix kI, but it is not exact for $k \neq 1$ or $\varepsilon \neq 0$

(4) N = number of Monte Carlo trials

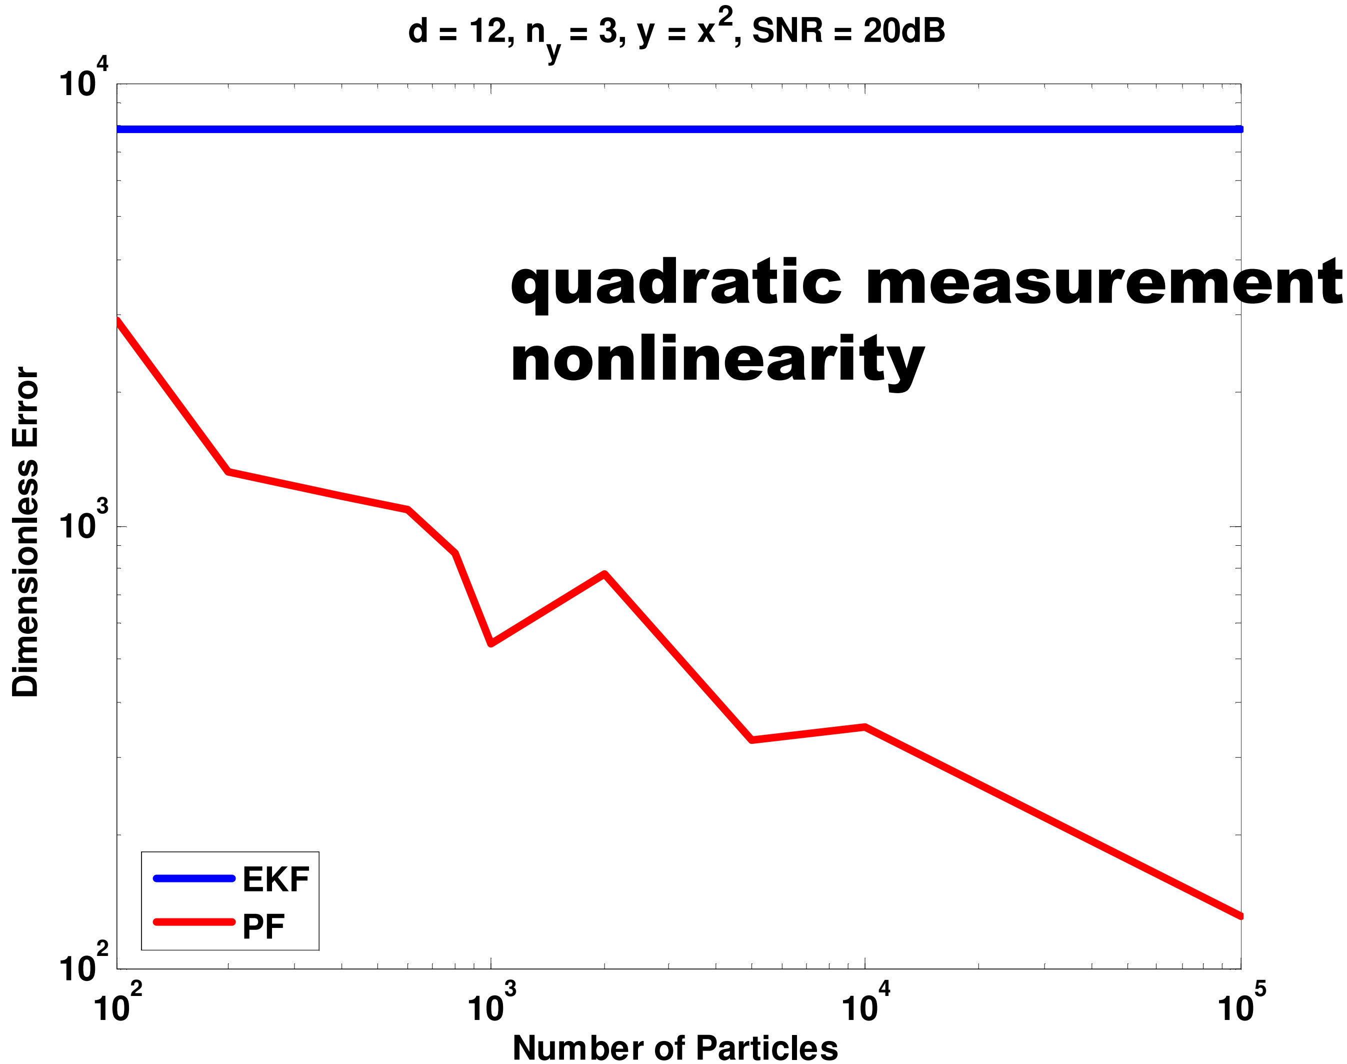

d = 12, n_y = 3, y = x^2, SNR = 20dB
Dimensionless Error
Number of Particles
EKF
PF
quadratic measurement nonlinearity

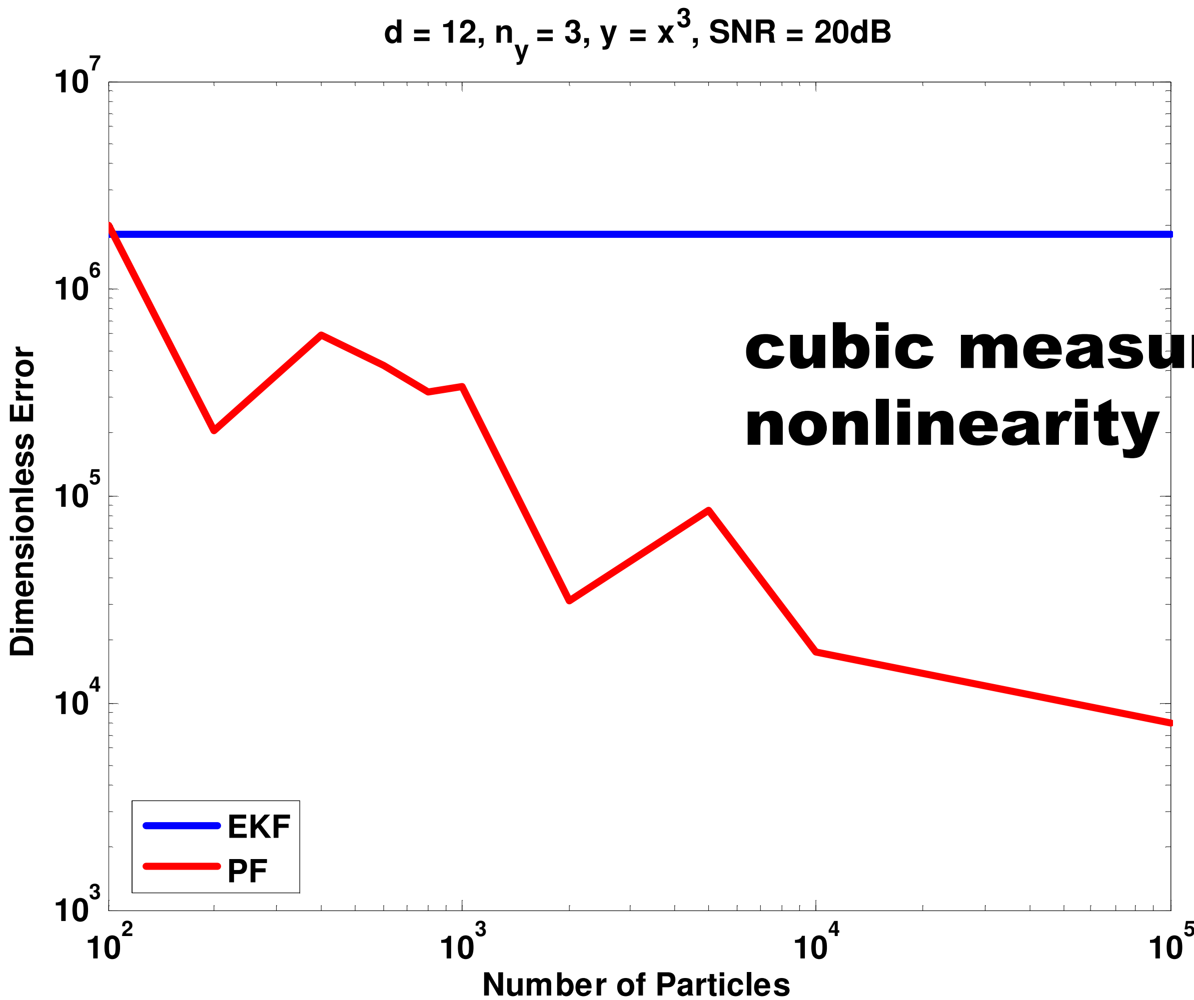

cubic measurement nonlinearity

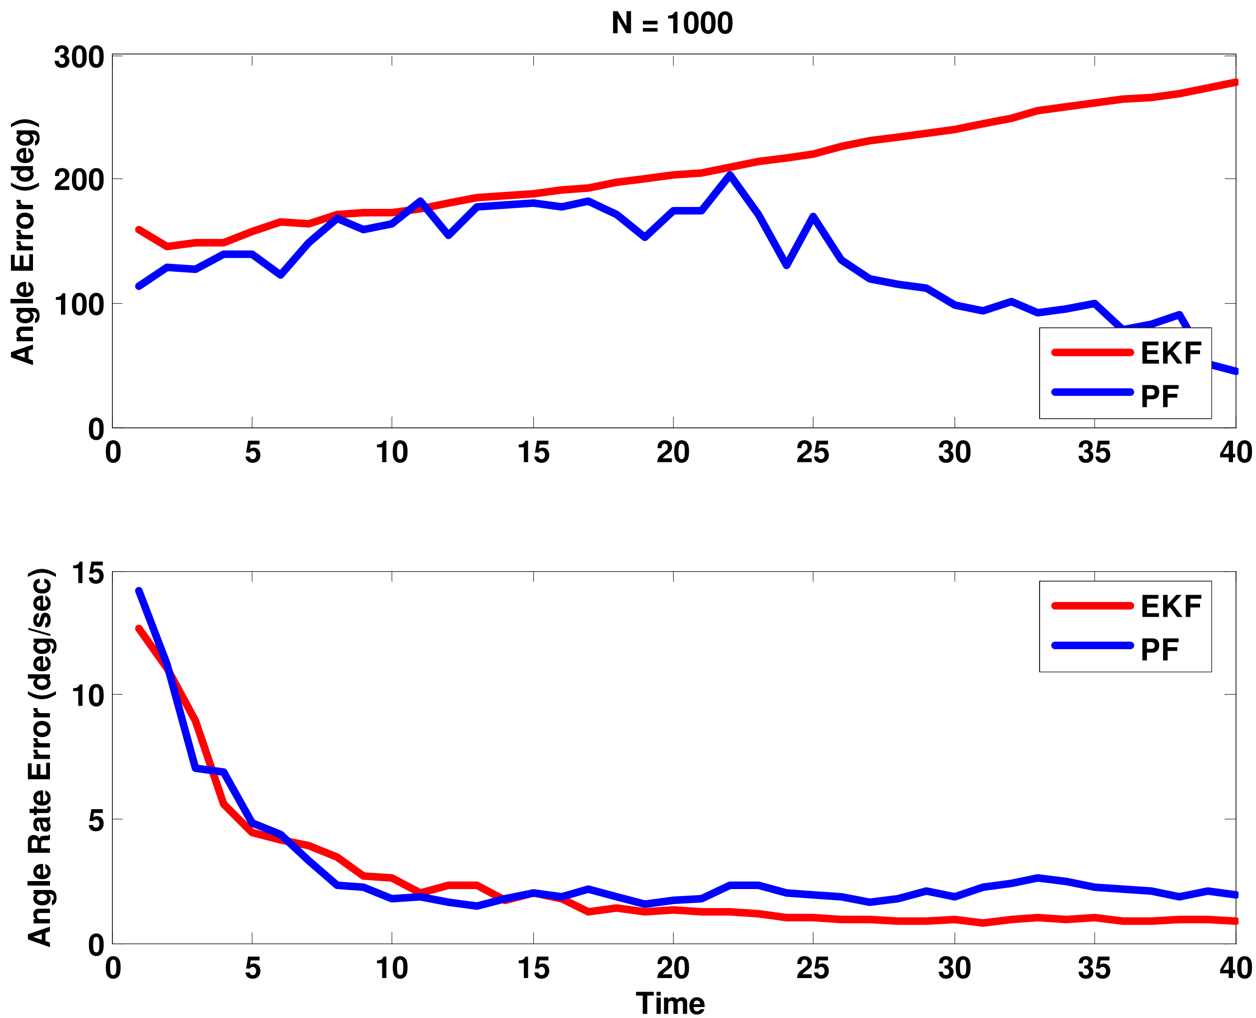
N = 1000
Angle Error (deg)
300
200
100
0
0
5
10
15
20
25
30
35
40
EKF
PF
Angle Rate Error (deg/sec)
15
10
5
0
0
5
10
15
20
25
30
35
40
Time
EKF
PF

Variation in SNR

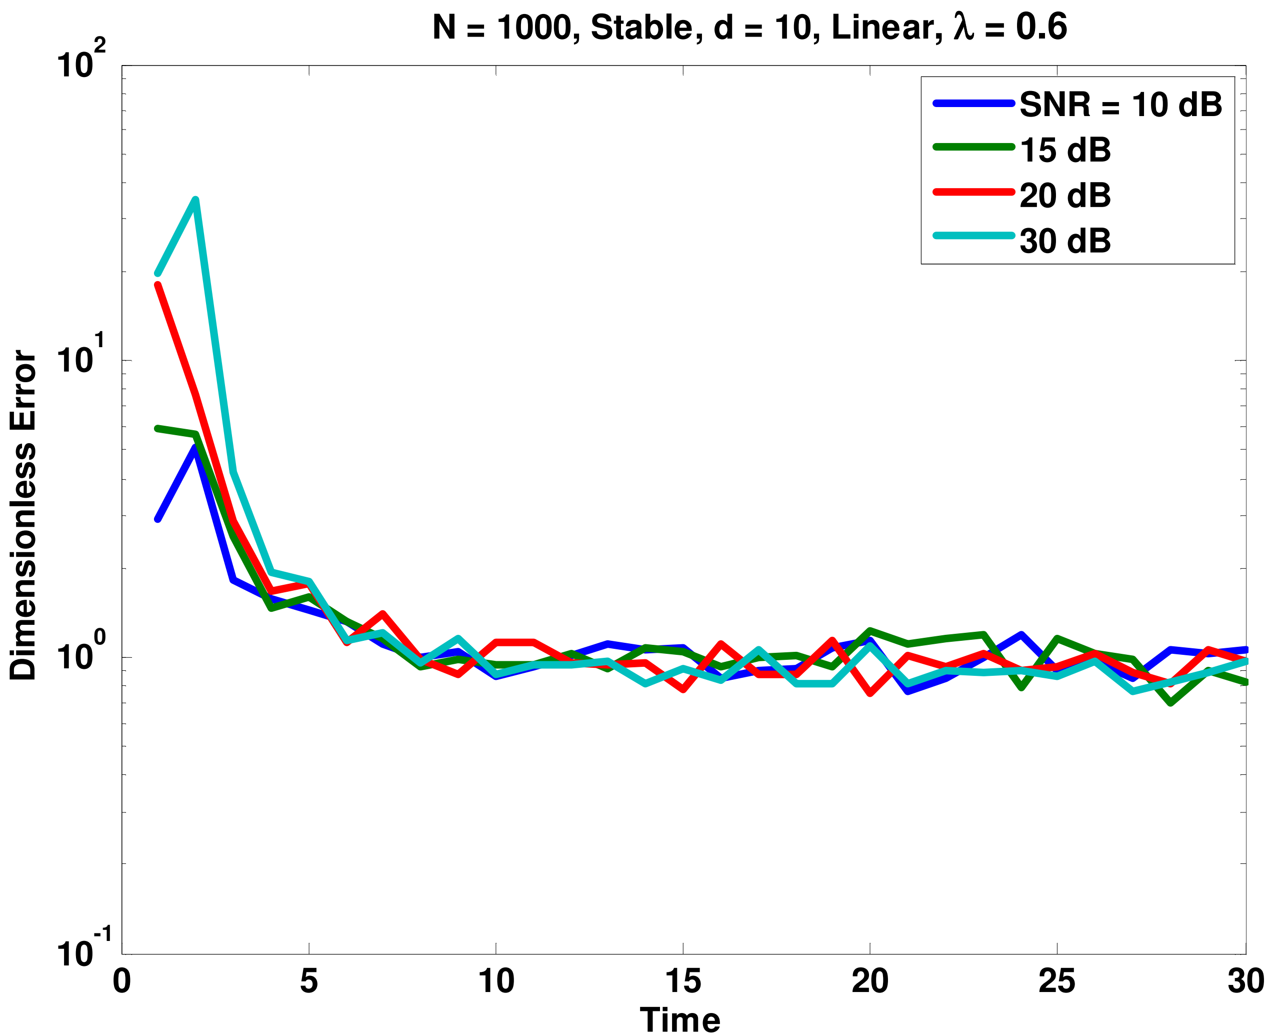

25 Monte Carlo Trials

Variation in Process Noise

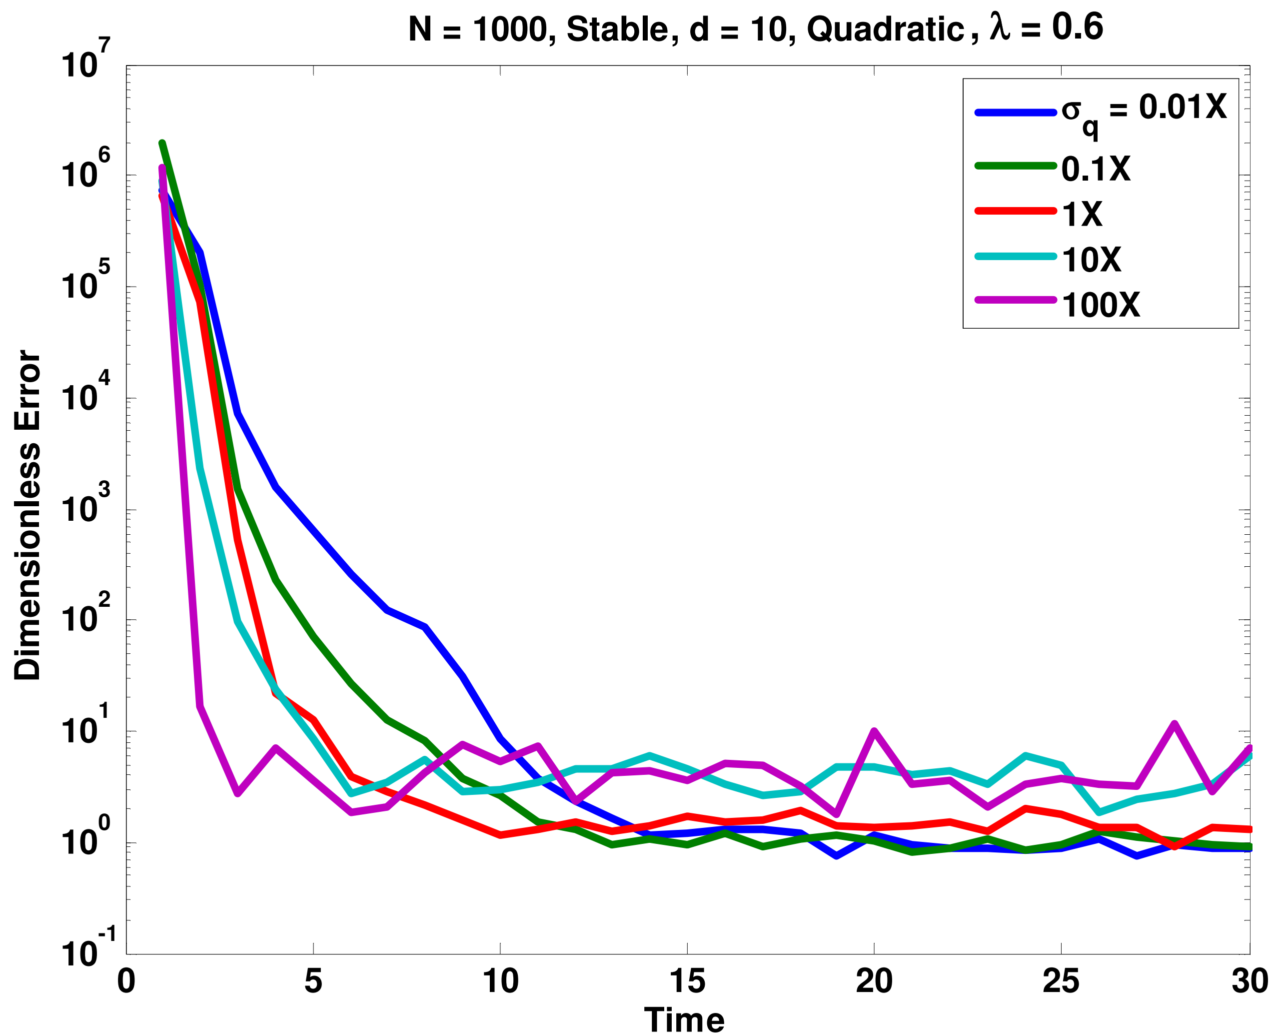

25 Monte Carlo Trials