A Model Reference Adaptive Search Method for Global Optimization

Steve Marcus, University of Maryland, College Park
http://www.isr.umd.edu/~marcus
(based on joint work with Jiaqiao Hu and Michael Fu)

Carnegie Mellon University
October 22, 2010

.

The Institute for Systems Research

UNIVERSITY OF MARYLAND 18 56 A. JAMES CLARK SCHOOL OF ENGINEERING

Outline

- Problem Setting
- Instance-Based vs. Model-Based
- Model-Based Algorithms
  - Estimation of Distribution Algorithms (EDAs)
  - Cross-Entropy (CE) Method
  - Model Reference Adaptive Search (MRAS)
- Convergence of MRAS
- Numerical Examples
- Extension to Stochastic Optimization and MDPs
- A New Particle Filtering Framework (if time)

Problem Setting

- Solution space $\chi \subseteq \Re^n$
  - continuous or discrete (combinatorial)
- Objective function $H(\cdot): \chi \rightarrow \Re$
- Objective: find optimal $x^* \in \chi$ such that

$$x^* \in \arg\min_{x \in \chi} H(x)$$

  - Assumptions: existence, uniqueness
    (but possibly many local minima)

Overview of Global Optimization Approaches

- **Instance-based approaches:** search for new solutions depends directly on previously generated solutions
  - simulated annealing (**SA**)
  - genetic algorithms (**GA**s)
  - tabu search
  - nested partitions

Model-Based Search Methods

Main idea: new solutions generated via an intermediate probability model

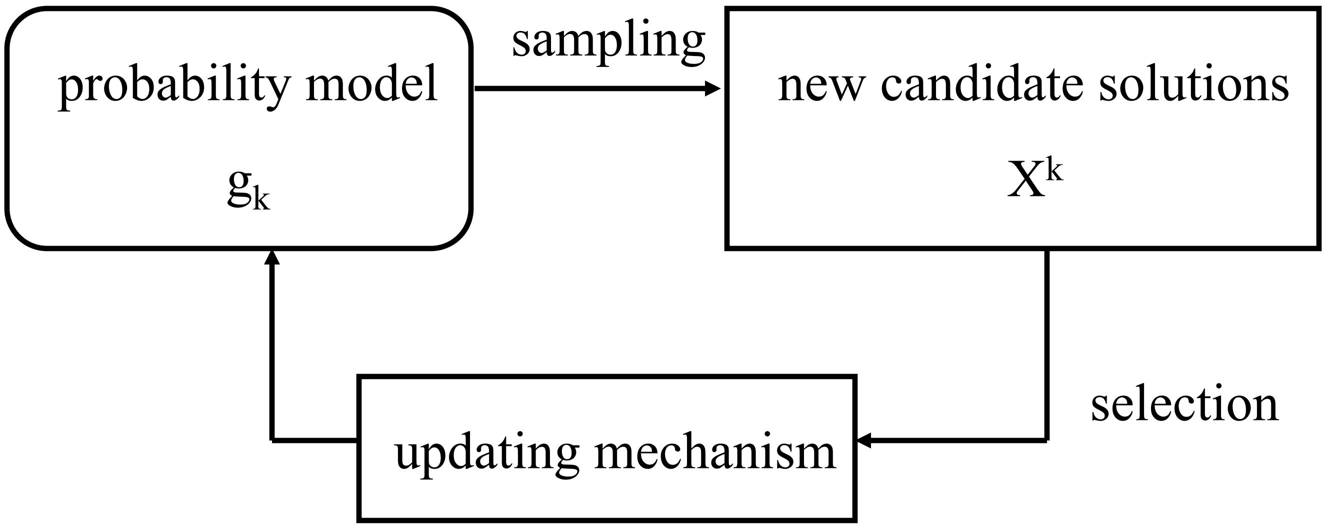

Model-Based Approach: Graphical Depiction

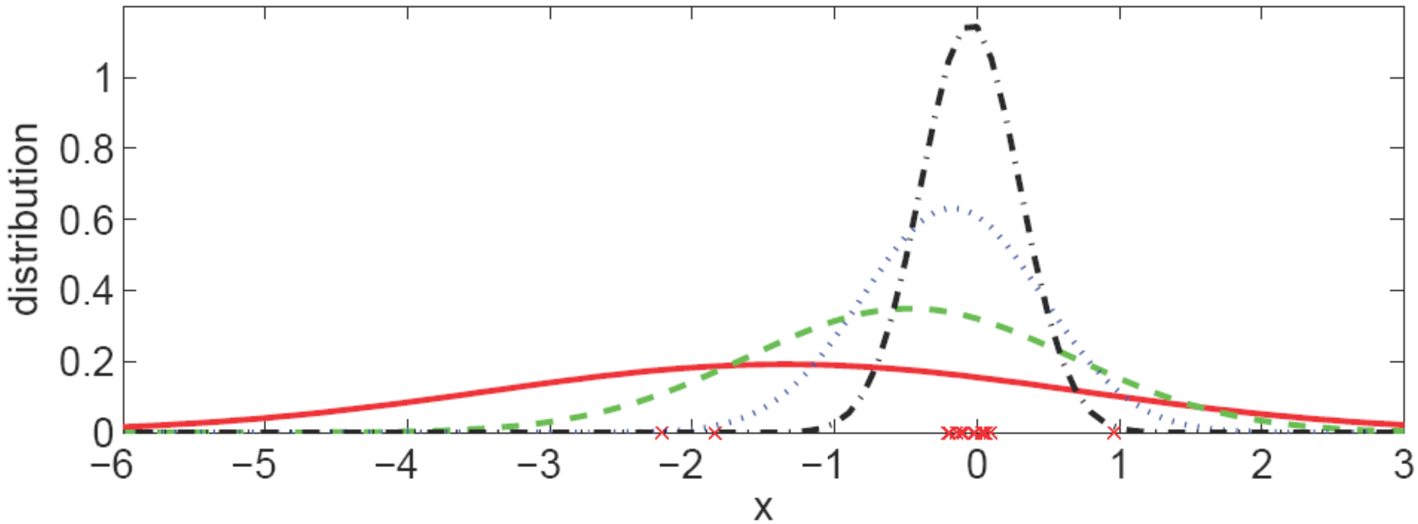

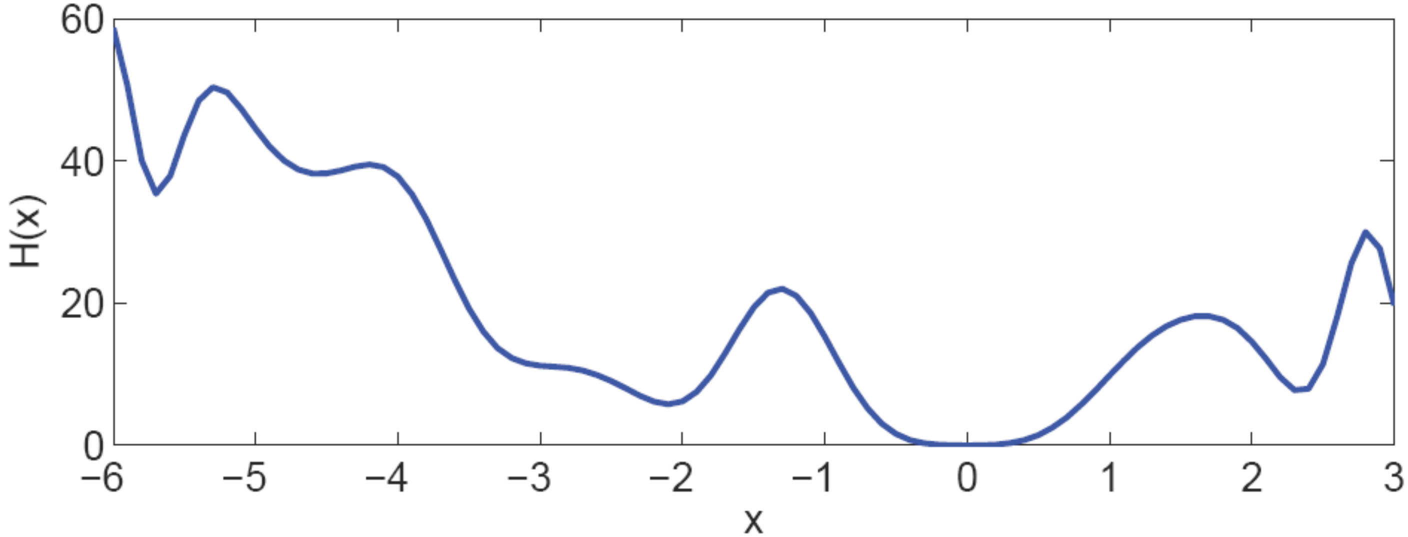

Combinatorial Optimization Example: TSP

How do we formulate this problem
to use a probability distribution?

- routing matrix of probability of arc i → j.
  - Example: four cities

0	0.5	0.4	0.1
0.2	0	0.6	0.2
0.4	0.4	0	0.2
0.3	0.3	0.4	0

- What is convergence?
  - single 1 in each row
  - single 1 in each column

Model-Based Methods

similarities to genetic algorithms

- uses a population
- selection process
- randomized algorithm,
  but uses “model” (distribution) instead of operators

Main Model-Based Methods

- estimation of distribution algorithms (**EDA**s)
  Muhlenbein and Paas (1996);
  book by Larranaga and Lozano (2001)
  [other names, e.g., probabilistic model-building GAs]
- cross-entropy method (**CE**)
  Rubinstein (1997, 1999) (www.cemethod.org);
  book by Rubinstein and Kroese (2004)
- probability collectives (Wolpert 2004)
- model reference adaptive search (**MRAS**)

Model-Based Methods (continued)

BIG QUESTION:

How to update distribution?

- traditional EDAs use an explicit construction,
  can be difficult & computationally expensive
- CE method uses single fixed target distribution (optimal importance sampling measure)
- MRAS approach:
  sequence of implicit model reference distributions

MRAS and CE Methods

- ALTERNATIVE: sample from a parameterized family of distributions, and update parameters by minimizing "distance" to desired distributions (reference distributions in MRAS)

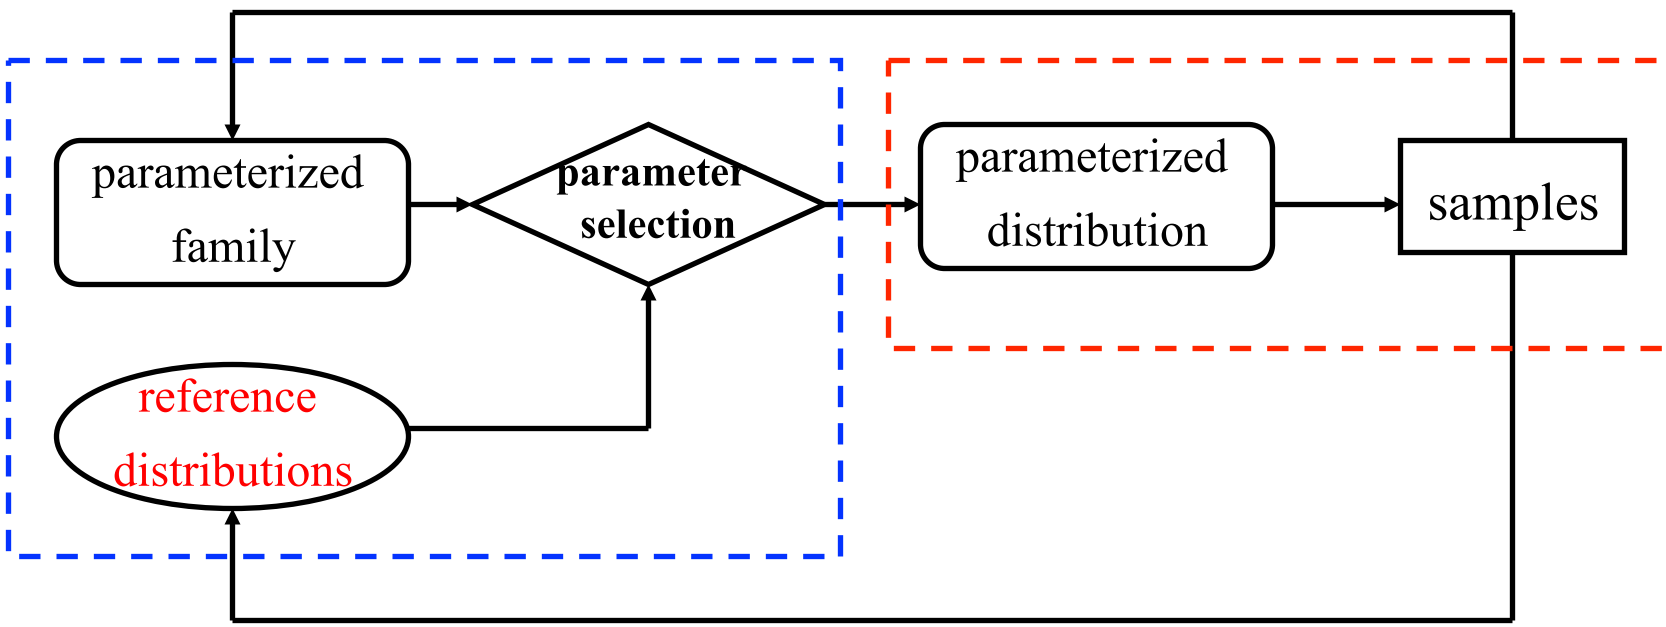

Model Reference Adaptive Search

- Main characteristics
  - Given sequence of reference distributions $\{g_k(\cdot)\}$
  - works with a family of parameterized probability distributions $\{f(\cdot,\theta)\}$ over the solution space
  - fundamental steps at iteration k :

    * generate candidate solutions according to the current probability distribution $f(\cdot, \theta_k)$

    * calculate θ_{k+1} using data collected in previous step to bias future search toward promising regions, by minimizing distance between $\{f(\cdot,\theta)\}$ and $g_{k+1}(\cdot)$
  - Algorithm converges to optimal if $\{g_k(\cdot)\}$ does

MRAS: Specific Instantiation

- reference distribution construction:

  Next distribution obtained by tilting previous

$$g_{k+1}(x) = \frac{S(H(x))g_k(x)}{E_{g_k}[S(H(X))]}, \ \forall x \in \chi,$$

  where S(.) is non-negative and strictly decreasing (increasing for max problems)

  Properties:

$$E_{g_{k+1}}[S(H(X))] \geq E_{g_k}[S(H(X))],$$

$$\lim_{k \to \infty} E_{g_k}[S(H(X))] = S(H(x^*)).$$

- selection parameter ρ determines the proportion of solutions used in updating θ_{k+1}

MRAS: Parameter Updating

- $(1-\ \rho)$−quantiles w.r.t. $f(\cdot,\theta_k)$

$$\gamma_{k+1} = \sup_l \{l : P_{\theta_k}(H(X) < l) \geq \rho\}$$

- update θ_{k+1} as

$$\theta_{k+1} = \arg\max_{\theta\in\Theta} \int_{x\in\chi} [S(H(x))]^k I\{H(x) < \gamma_{k+1}\} \ln f(x\ ,\ \theta)dx$$

Lemma : θ_{k+1} minimizes the Kullback-Leibler (KL) divergence between g_{k+1} and $f(\cdot,\theta)$, i.e.,

$$\theta_{k+1} = \arg\min_{\theta\in\Theta} D(g_{k+1} \mid f(\cdot\,,\theta)) := \arg\min_{\theta\in\Theta} E_{g_{k+1}}\left[\ln \frac{g_{k+1}(X)}{f(X,\theta)}\right]$$, where

$$g_{k+1}(x) = \frac{S(H(x)) I_{\{H(x)<\gamma_{k+1}\}} g_k(x)}{E_{g_k}[S(H(X)) I_{\{H(X)<\gamma_{k+1}\}}]},\quad g_1(x) := \frac{I_{\{H(x)<\gamma_1\}}}{E_{\theta_0}[I_{\{H(X)<\gamma_1\}} / f(X,\theta_0)]}$$

Restriction to Natural Exponential Family (NEF)

- covers broad class of distributions
- closed-form solution for θ_{k+1}
- global convergence can be established under some mild regularity conditions

  * multivariate Gaussian case

$$\lim_{k\to\infty} \mu_k = x^*, \quad \lim_{k\to\infty} \Sigma_k = 0_{n\times n}$$

  * independent univariate case

$$\lim_{k\to\infty} E_{\theta_k}[X] = x^*.$$

MRAS: Monte-Carlo version

- Changes from exact version
  - finite number of samples, say N_k, at each iteration
  - replace the true $(1-\rho)$-quantiles by sample quantiles
  - replace the integrals (expected values) by sample averages
  - ρ_k adaptively decreasing and N_k adaptively increasing

- Global convergence can be established
  - multivariate normal case

$$\lim_{k\to\infty} \hat{\mu}_k = x^*, \text{ and } \lim_{k\to\infty} \hat{\Sigma}_k = 0_{n\times n} \quad \text{w.p.1.}$$

  - independent univariate case

$$\lim_{k\to\infty} E_{\hat{\theta}_k}[X] = x^* \quad \text{w.p.1.}$$

Comparison of MRAS & CE

- MRAS provides general framework, and has general sequence of implicit reference models $\{g_k\}$. CE can be interpreted by defining appropriate $\{g_k\}$, but the sequence depends on $\{f(\cdot,\theta)\}$
- Global convergence results for MRAS use property of $\{g_k\}$ convergence to optimal distribution, not in general true for CE
- both use parameterized distributions, KL divergence
- CE generally easier to implement; preliminary computational results indicate no clear dominance of either

Numerical Examples (deterministic problems)

- Continuous optimization
  - * 20-D Rosenbrock function

$$\sum_{i=1}^{19} 100(x_{i+1} - x_i^2)^2 + (x_i - 1)^2$$

  - * 20-D Trigonometric function

$$1 + \sum_{i=1}^{20} 8\sin^2(7(x_i - 0.9)^2) + 6\sin^2(14(x_i - 0.9)^2) + (x_i - 0.9)^2$$

  - * 20-D Pinter function

- Combinatorial optimization
  - * various asymmetric traveling salesman problems (ATSP)

http://www.iwr.uniheidelberg.de/groups/comopt/software/TSPLIB95

Rosenbrock function, where $-5 \leq x_i \leq 5$, i=1,2

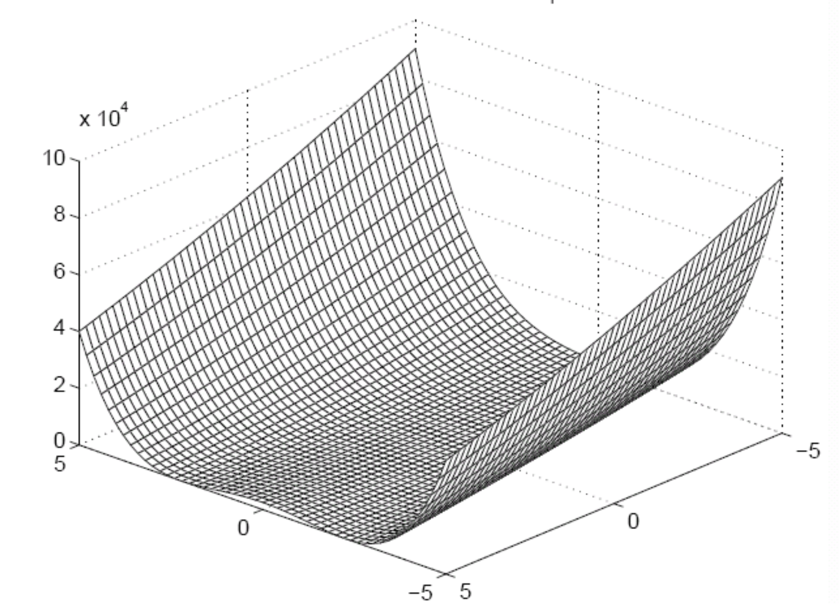

Trigonometric function, where $-3 \leq x_i \leq 3$, i=1,2

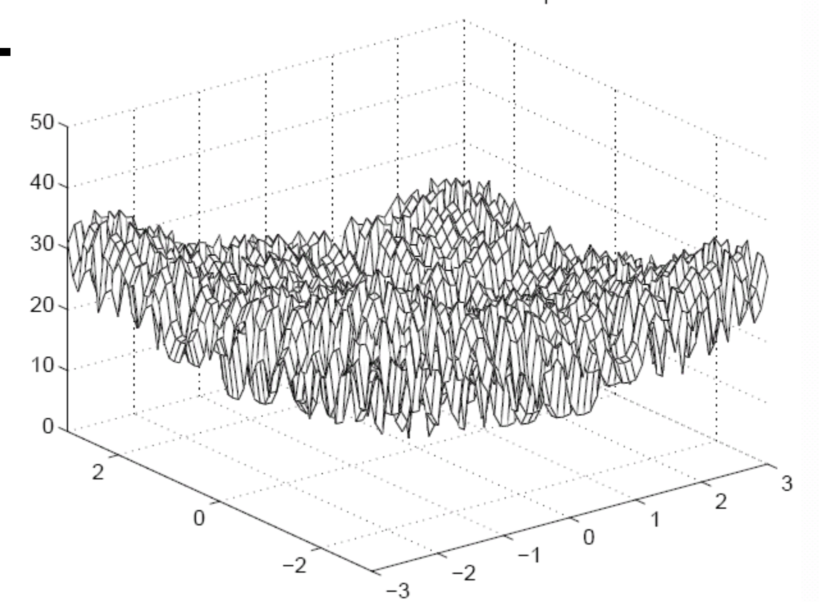

Pinter's function, where $-5 \leq x_i \leq 5$, i=1,2

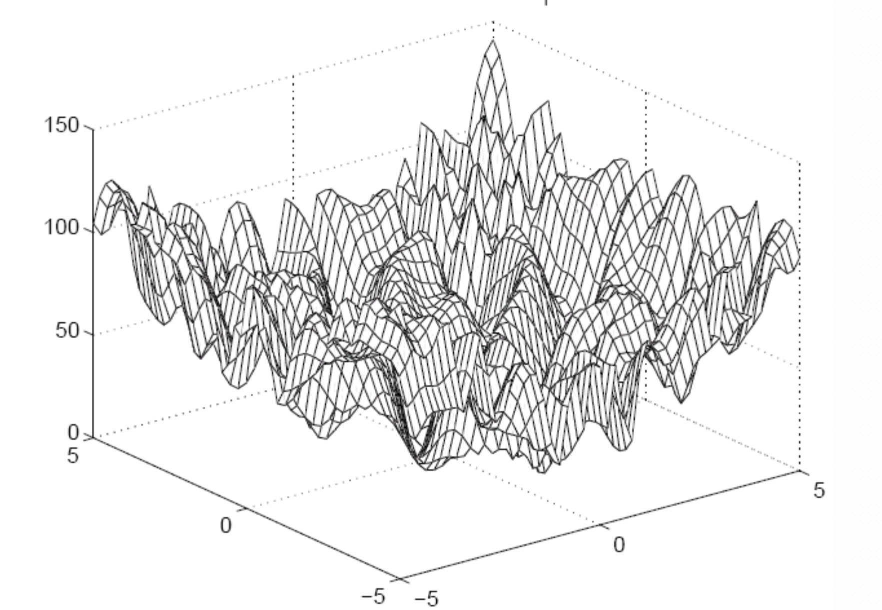

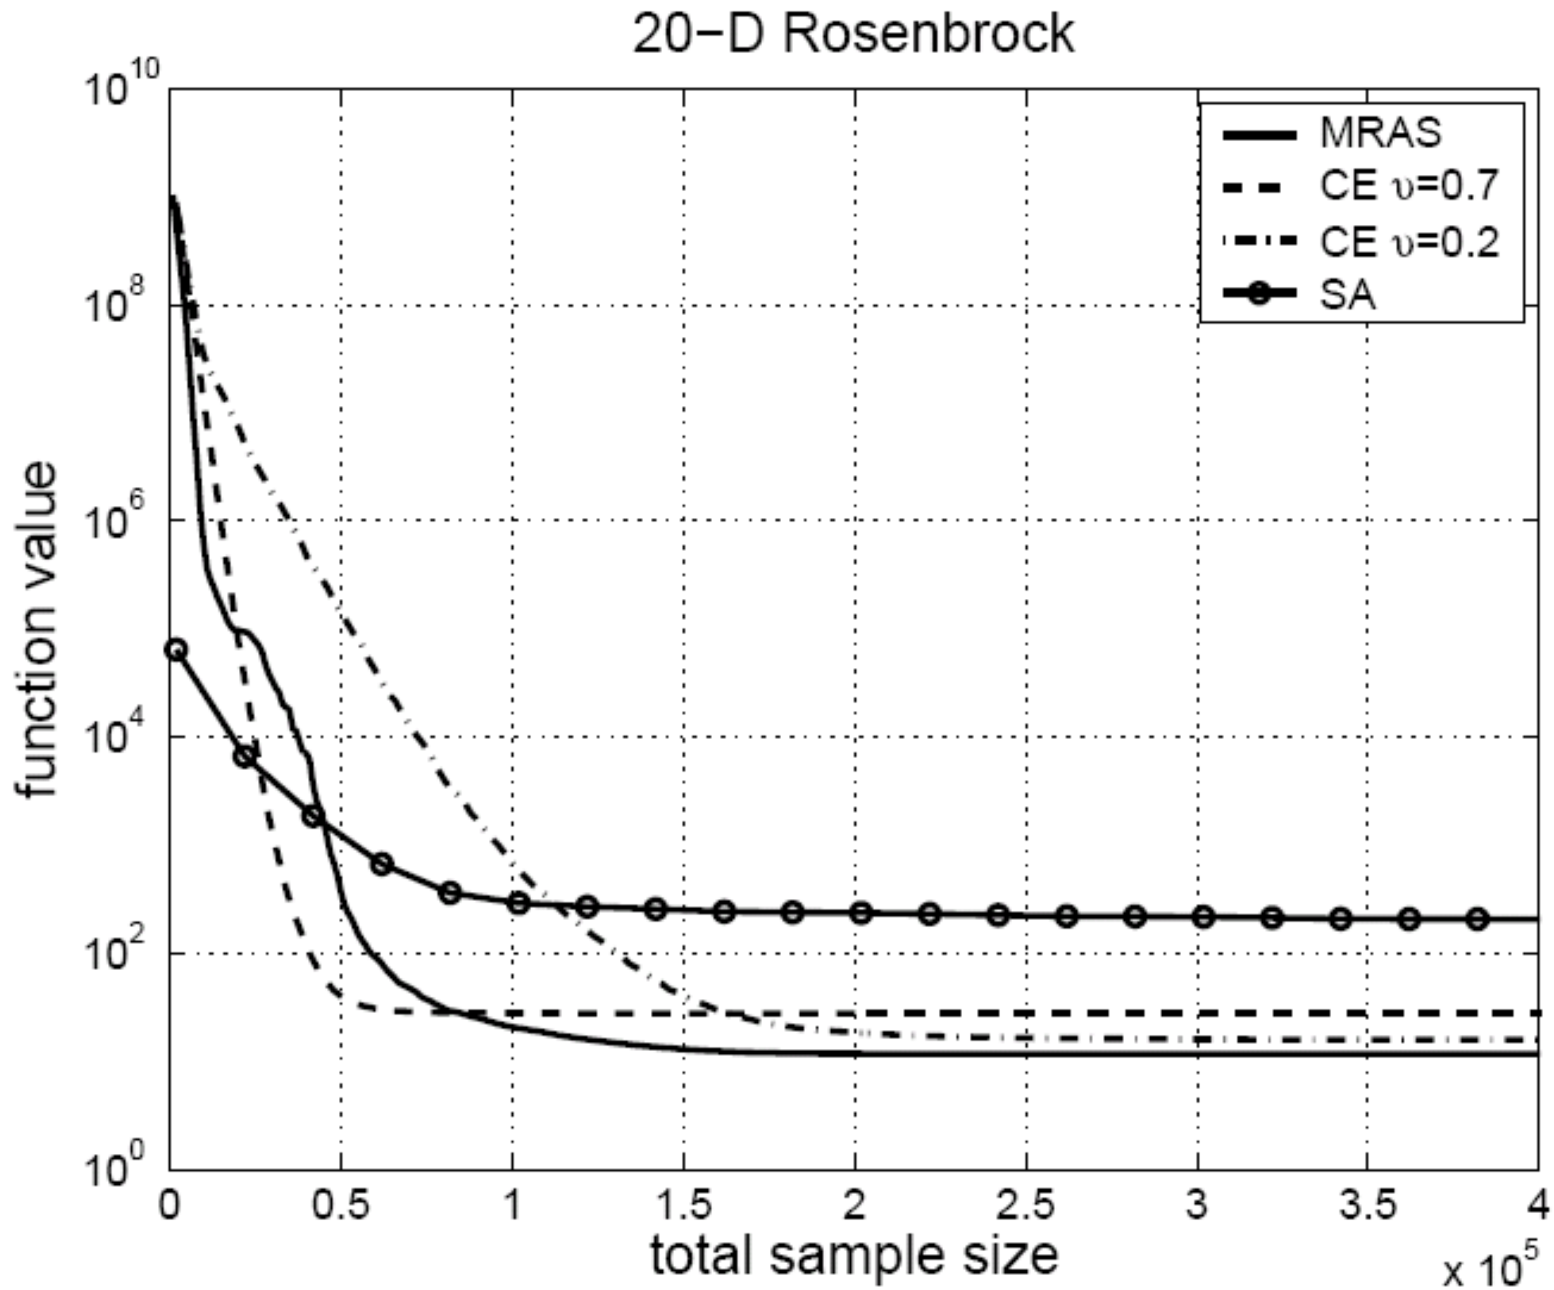
20−D Rosenbrock
function value
total sample size
x 10^5
MRAS
CE υ=0.7
CE υ=0.2
SA
10^10
10^8
10^6
10^4
10^2
10^0
0
0.5
1
1.5
2
2.5
3
3.5
4

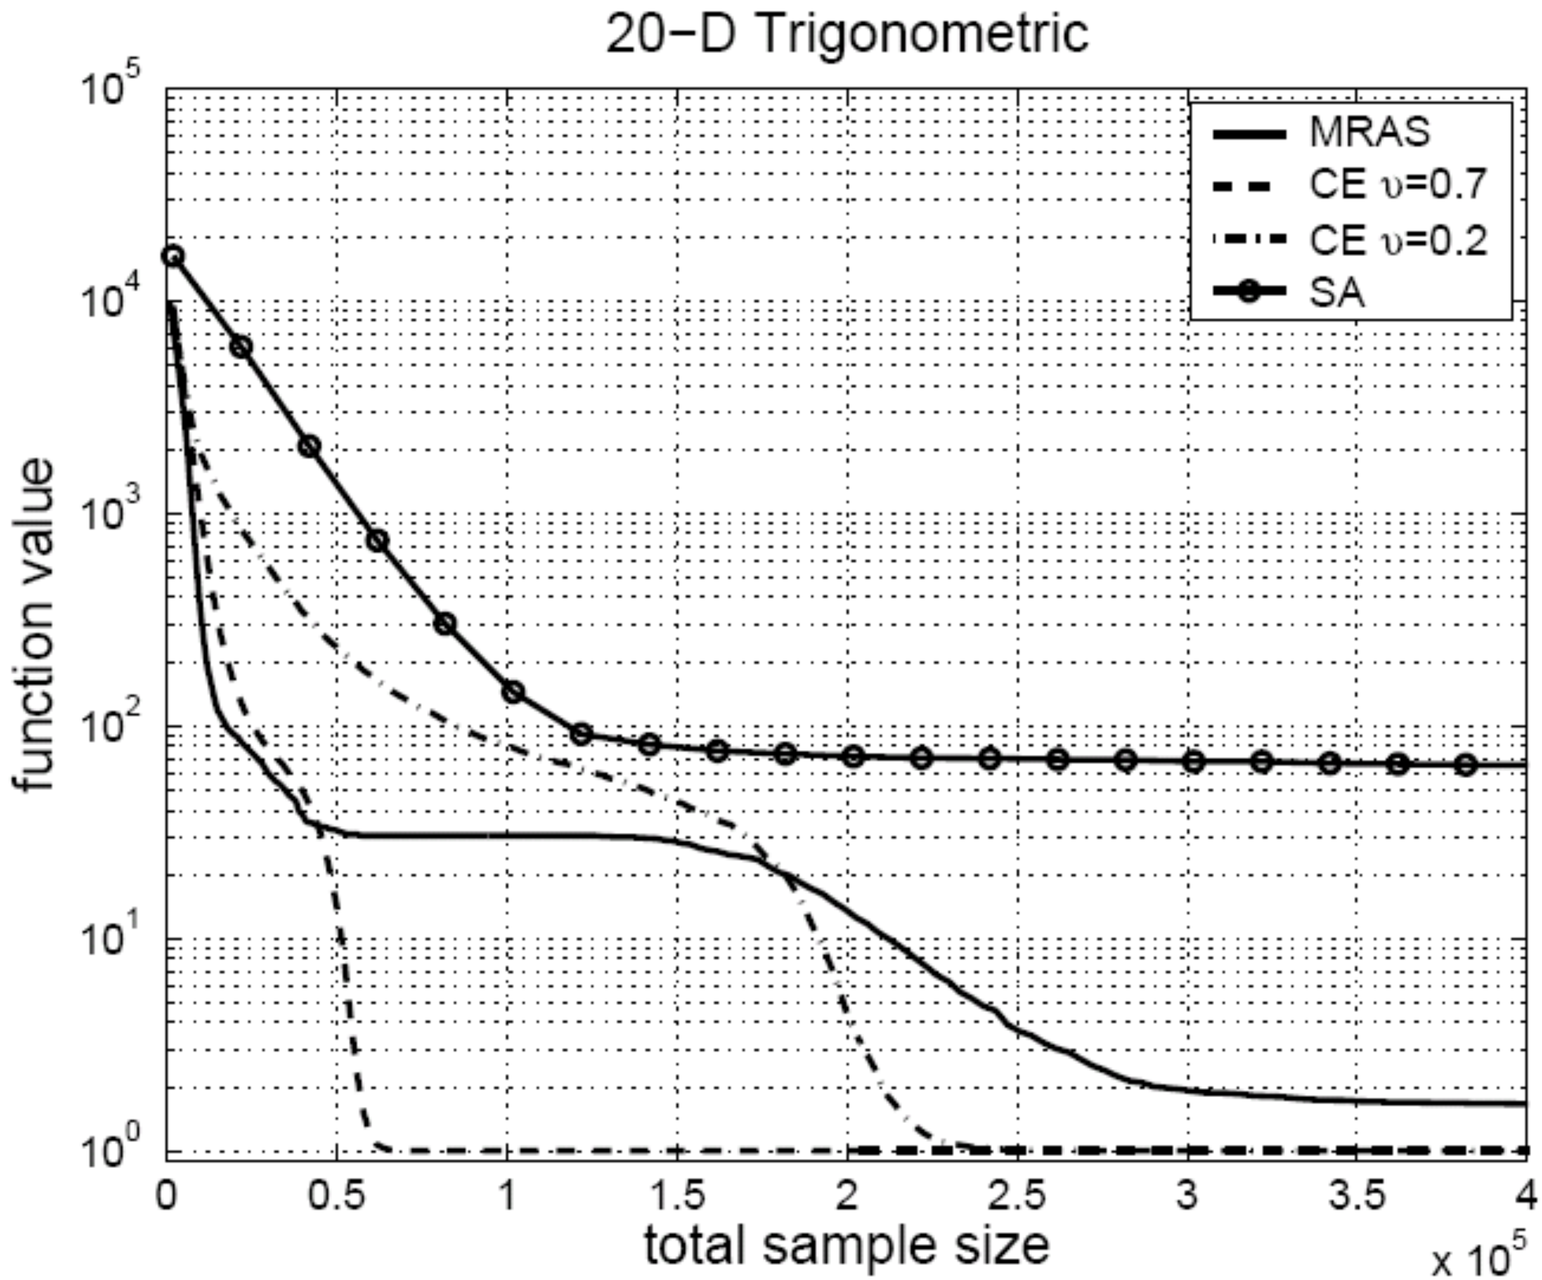

20−D Trigonometric
function value
total sample size
x 10^5
MRAS
CE υ=0.7
CE υ=0.2
SA
10^0
10^1
10^2
10^3
10^4
10^5
0
0.5
1
1.5
2
2.5
3
3.5
4

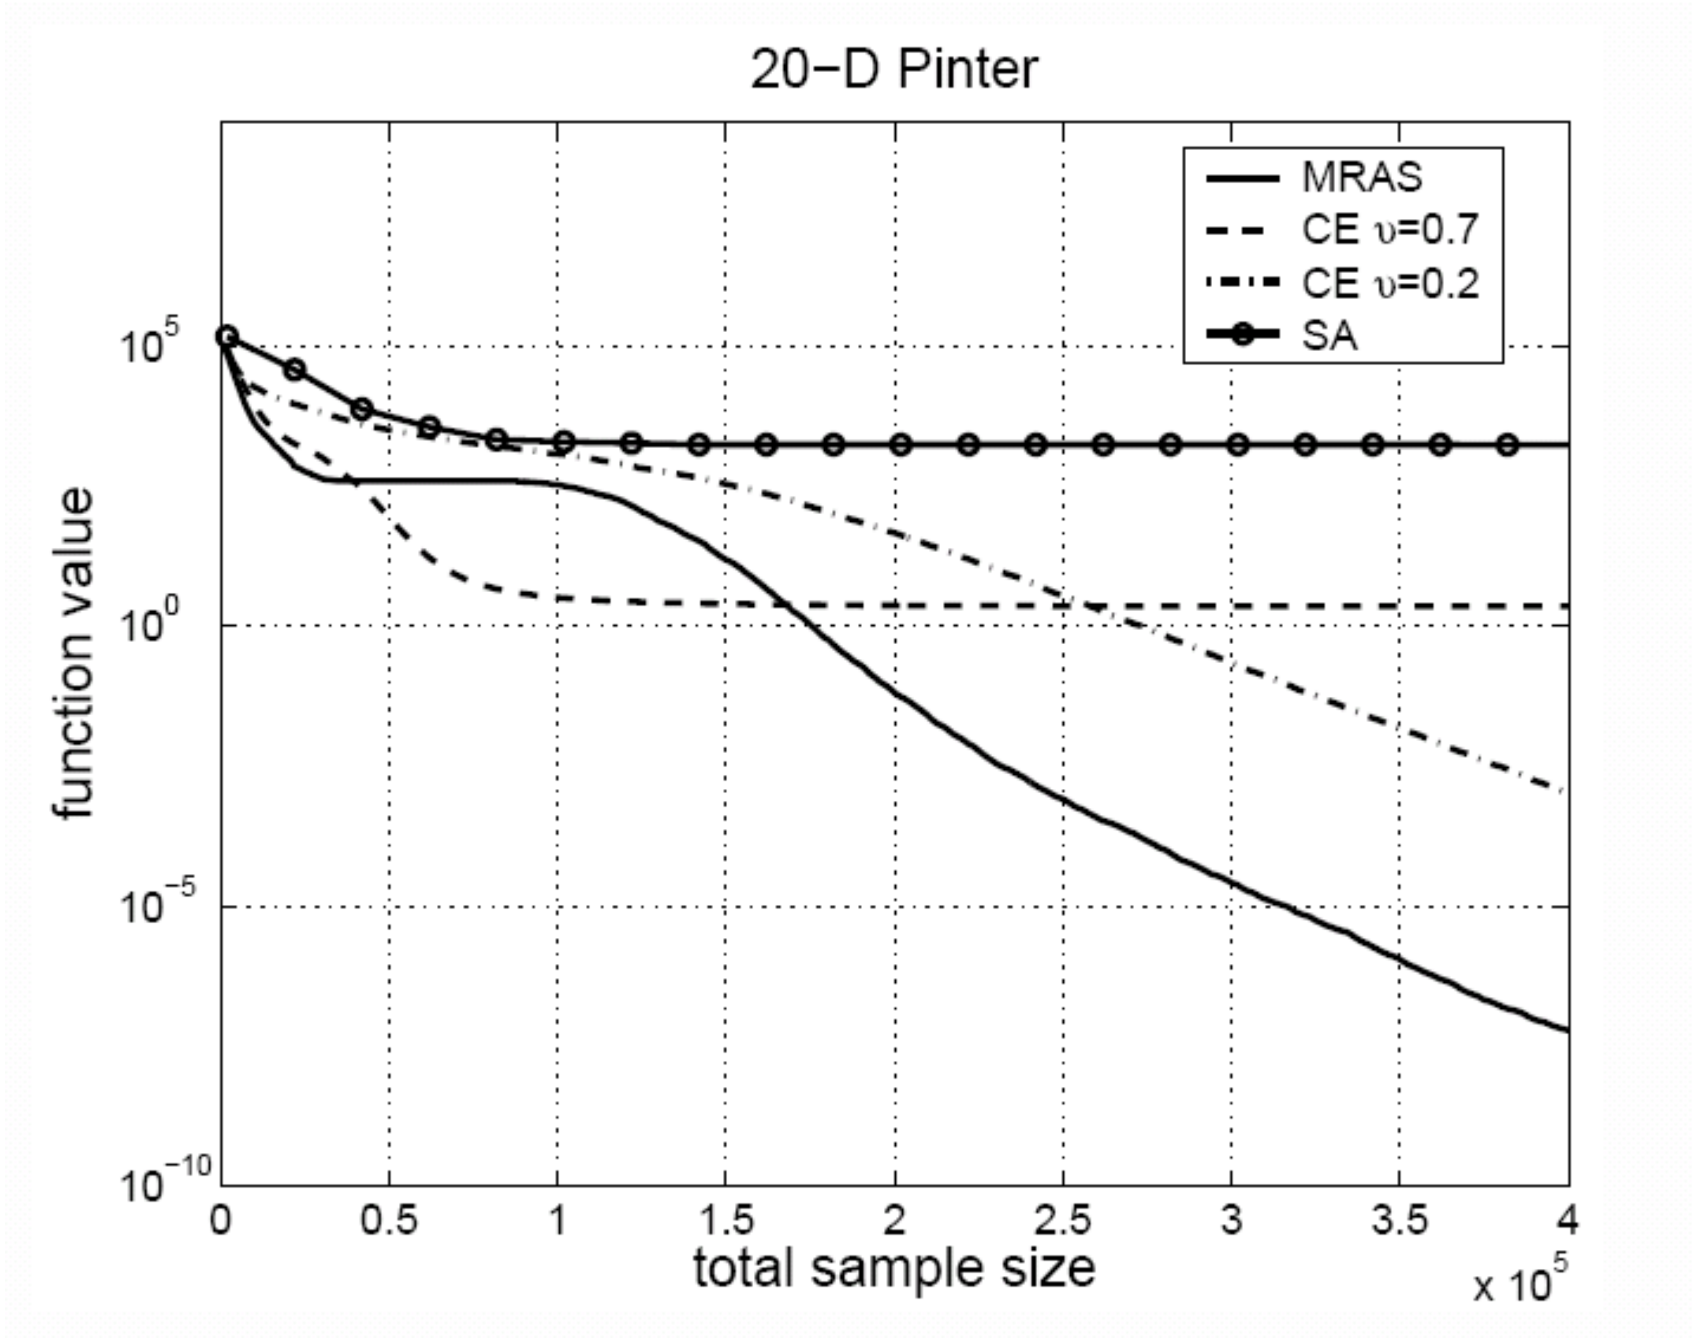
20−D Pinter
MRAS
CE υ=0.7
CE υ=0.2
SA
function value
10^5
10^0
10^−5
10^−10
0
0.5
1
1.5
2
2.5
3
3.5
4
total sample size
x 10^5

Numerical Examples (deterministic problems)

- Numerical results for ATSPs
  - DISCRETE distribution (matrix: probability i→j on tour)
  - Good performance with modest number of tours generated
  - ft70 case: total number of admissible tours = 70! ≈ 10^{100}

file	N_c	$N_{avg.}$ *(std err)*	H_*	H^*	H_{best}	δ_{avg} *(std err)*
ftv33	**34**	**7.95e+4(3.25e+3)**	**1364**	**1286**	**1286**	**0.023(0.008)**
ftv35	**36**	**1.02e+5(3.08e+3)**	**1500**	**1475**	**1473**	**0.008(0.002)**
ftv38	**39**	**1.31e+5(4.90e+3)**	**1563**	**1530**	**1530**	**0.008(0.003)**
p43	**43**	**1.02e+5(4.67e+3)**	**5637**	**5620**	**5620**	**0.001(2.5e-4)**
ry48p	**48**	**2.62e+5(1.59e+4)**	**14810**	**14446**	**14422**	**0.012(0.003)**
ft53	**53**	**2.94e+5(1.58e+4)**	**7236**	**6973**	**6905**	**0.029(0.005)**
ft70	**70**	**4.73e+5(2.91e+4)**	**39751**	**38744**	**38673**	**0.017(0.003)**

Extension to Stochastic Optimization

- Objective: find optimal $x^* \in \chi$ such that

$$x^* \in \arg\min_{x \in \chi} E_\omega[H(x,\omega)]$$

  - Assumptions: existence, uniqueness (but possibly many local minima)

- Idea: sample average approximation
  - At each iteration *k,* approximate $E_\omega[H(x,\omega)]$ by

$$\overline{H}_k(x) := \frac{1}{M_k} \sum_{i=1}^{M_k} H_{i,k}(x),$$

where $H_{i,k}(x)$ are i.i.d. random observations at x.

Extension to Stochastic Optimization

- Key convergence issue
  - $g_{k+1}(x) = \dfrac{S(\overline{H}_k(x)) I_{\{\overline{H}_k(x) < \bar{\gamma}_{k+1}\}} g_k(x)}{E_{g_k}[S(\overline{H}_k(X)) I_{\{\overline{H}_k(X) < \bar{\gamma}_{k+1}\}}]}$ may not converge.

    * $\overline{H}_k(x) \to E_\omega[H(x,\omega)]$, need $M_k \to \infty$ as $k \to \infty$

    * schedule of sample size M_k, restrictions on $S(\cdot)$.

- Practical efficiency
  - increase M_k adaptively, i.e., small M_k value initially, large M_k when precise estimates required
  - reuse of "good" samples for finite solution spaces

(s,S) Inventory Control Problem

- X_t : inventory position in period t.
- D_t : the i.i.d exponential demand in period t
- h : per period per unit holding cost; p: demand lost penalty cost ; c: per unit ordering cost; K: fixed set-up cost

$$X_{t+1} = \begin{cases} S - D_{t+1} & X_t < s, \\ X_t - D_{t+1} & X_t \geq s. \end{cases}$$

- The objective is to minimize the long run average cost per period:

$$J_t(s, S) := \frac{1}{t} \sum_{i=1}^{t} \left[I\{X_i < s\}(K + c(S - X_i)) + hX_i^+ + pX_i^- \right].$$

(s,S) Inventory Control Problem

Case 1: $c = h = 1$, p=10, K=100, E[D]=200

Case 2: $c = h = 1$, p=10, K=10000, E[D]=200

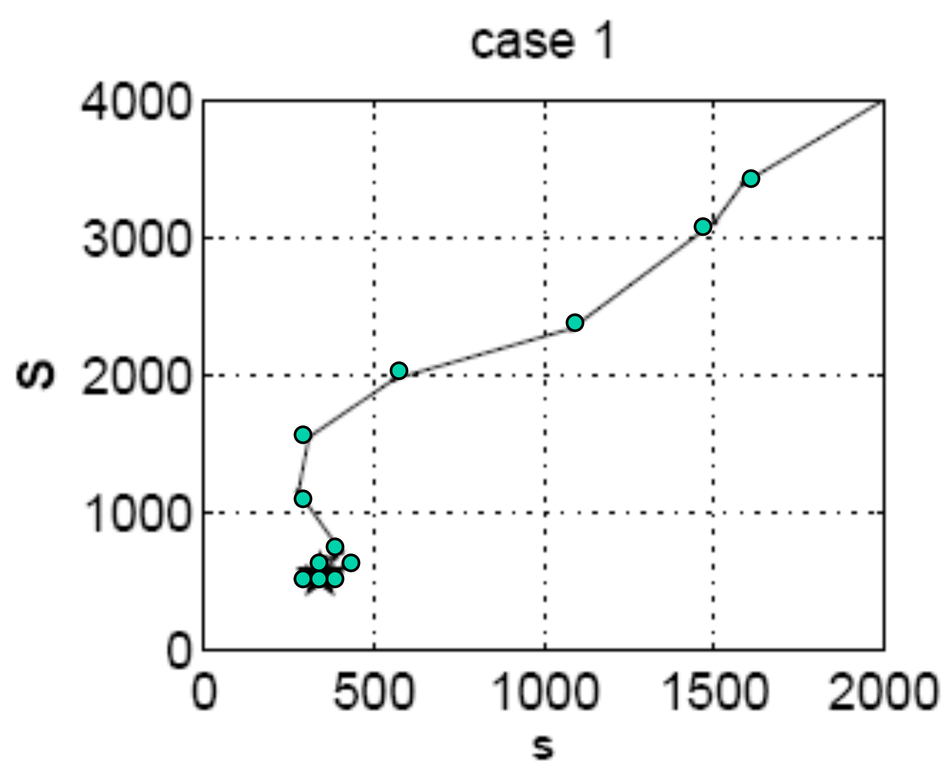

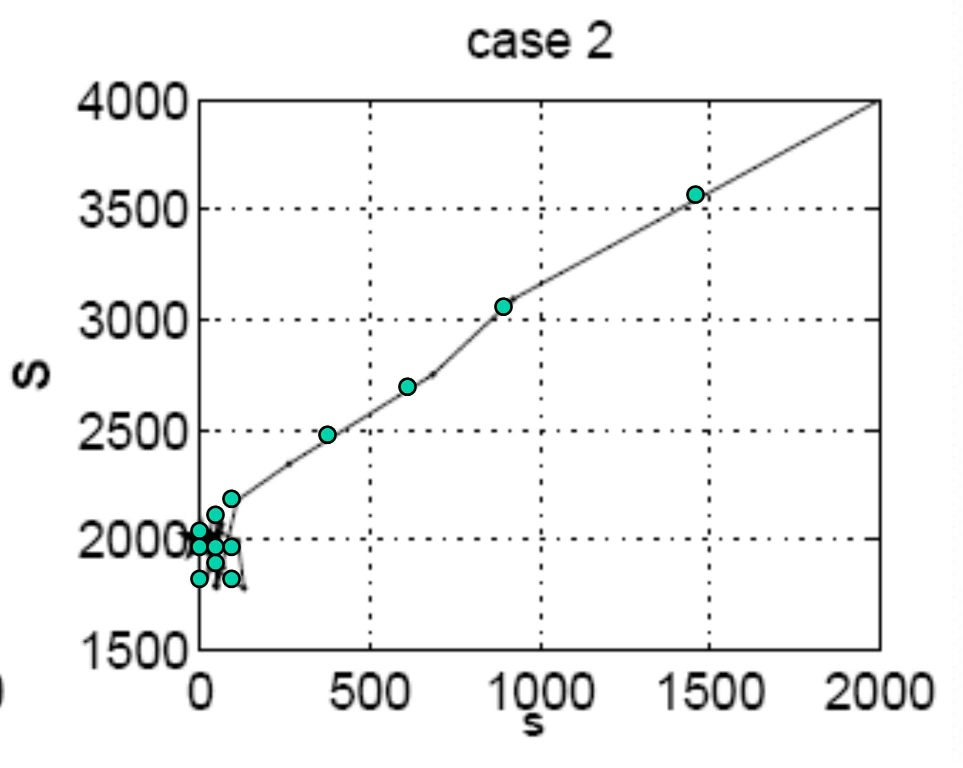

Buffer Allocation in Unreliable Production Lines

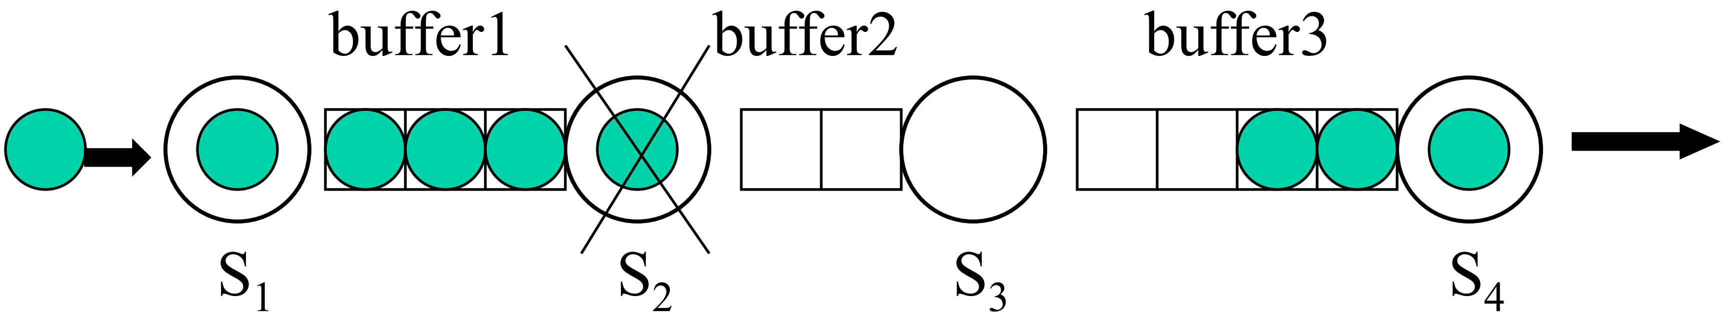

- Input:
  - μ_i: service rate of server i
  - f_i : failure rate of server i
  - r_i : repair rate of server i
  - n : total number of buffers available
- Let n_i be the number of buffers allocated to S_i satisfying $\Sigma n_i = n$, the objective is to choose n_i to maximize the steady-state throughput

Buffer Allocation in Unreliable Production Lines

$\mu_1 = 1,\ \mu_2 = 1.1,\ \mu_3 = 1.2,\ \mu_4 = 1.3,\ \mu_5 = 1.5,$

$f_i = 0.05$ and $r_i = 0.5$

n	$N_{avg}(std\ err)$	$\bar{\mathcal{T}}(std\ err)$	$\mathcal{T}^*$
1	1.02e+2(7.49)	0.523(6.79e-4)	0.521
2	1.29e+2(14.8)	0.555(3.86e-4)	0.551
3	1.75e+2(15.7)	0.587(4.57e-4)	0.582
4	2.51e+2(25.9)	0.606(1.20e-3)	0.603
5	3.37e+2(42.0)	0.626(6.57e-4)	0.621
6	4.69e+2(55.2)	0.644(1.10e-3)	0.642
7	4.56e+2(58.2)	0.659(1.10e-3)	0.659
8	4.45e+2(54.9)	0.674(1.10e-3)	0.674
9	5.91e+2(56.1)	0.689(1.39e-3)	0.689
10	5.29e+2(54.0)	0.701(1.10e-3)	0.701

Extension to MDPs

- book published by Springer: Chang, Fu, Hu, Marcus

Simulation-Based Approaches to Markov Decision Processes

- optimization over policy space
- population-based evolutionary algorithms (EPI/ERPS)

Filtering (with Enlu Zhou and M. Fu)

- State equation

$$x_{k+1} = f(x_k, u_k), k = 0, 1, \ldots .$$

- Observation equation

$$y_k = h(x_k, v_k), k = 1, 2, \ldots .$$

- Filtering:
  Estimate $b_k(x_k) = p(x_k | y_{0:k})$.

Optimization via Filtering

- Formulation as a Filtering Problem

$$
\begin{aligned}
x_k &= x_{k-1}, \;\; k = 1, 2, \ldots, \\
y_k &= H(x_k) - v_k, \;\; k = 0, 1, \ldots,
\end{aligned}
$$

where $x_k = x^*$ is the unobserved state, v_k has a p.d.f. $\varphi(\cdot)$.

- Interpretation: Observe the optimal function value $y^* = H(x^*)$ with some noise.

- We expect by suitable choice of $\{y_k\}$

$$
b_k(x_k) = p(x_k | y_{0:k}) \;\; \rightarrow \;\; \delta(x_k - x^*), \; a.s.
$$

Optimization via Filtering

- Transition density $p(x_k|x_{k-1}) = \delta(x_k - x_{k-1})$.
  Likelihood function $p(y_k|x_k) = \varphi(H(x_k) - y_k)$.

- Hence, the conditional density

$$b_k(x_k) = \frac{\varphi(H(x_k) - y_k)b_{k-1}(x_k)}{\int \varphi(H(x_{k-1}) - y_k)b_{k-1}(x_k)dx_k}.$$

- Interpretation: The conditional density is tuned by the performance of solutions at previous iteration.

Result: Using particle filtering (Monte Carlo simulation), EDAs, CE, MRAS can all be viewed in this framework.

Conclusions and Future Work

- **Summary**
  - new general framework for problems with little structure
  - guaranteed theoretical convergence
  - good experimental performance
- **Future Work**
  - incorporate known structure (e.g., local search)
  - convergence rate, computational complexity
  - more new algorithm instantiations in this framework
  - more comparisons with other algorithms