



THE UNIVERSITY OF WARWICK

#### **Example: Traffic**

Street networks



#### Warwick Business School

We live in a complex world

GECCO '18 Companion, July 15-19, 2018, Kyoto, Japan © 2018 Copyright is held by the owner/author(s).

ACM ISBN 978-1-4503-5764-7/18/07. https://doi.org/10.1145/3205651.3207887

**Simulation Optimisation** 

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author

- Large number of interacting elements
- Can not be understood by analysis of components

• Emergence

Tutorial

Warwick Business School

Prof Juergen Branke

• Simulation can capture emergent phenomena



#### Warwick Business School



### stochastic processing times, complex

Simulate machine breakdowns,

**Example: Manufacturing** 

scheduling rules, etc.



Warwick Business School

### **Example: Engineering**

Simulation can replace physical testing

## Simulation as knowledge generation tool

Great tool to understand and analyse complex real-world systems!



# Knowledge Data Analysis Simulation Theory Warwick Business School

# **Example: Stock market** Simulation allows taking into account

- bounded rationality ۲
- learning agents ۲
- heterogeneous agents ۲
- network effects  $\bigcirc$



Famous:

Santa Fe Stock Market:

Expectations of learning agents lead to technical trading

## The next step: **Simulation optimisation**

• Automatically search vast spaces of parameter settings to find "optimal" settings



- Model calibration
- Automated design and optimisation of complex systems

#### **Simulation optimisation examples**

- Traffic: Optimise traffic light controller
- Manufacturing: Find optimal dispatching rules
- Engineering: Find optimal wing design
- Finance: Find better investment strategies

#### Outline

- Strategies to deal with expensive evaluations
  - Parallelisation
  - Surrogate models
- Strategies to deal with noise
  - Selecting the best system
  - Simulation optimisation
- Applications
  - Design of traffic light controller
  - Design of dispatching rules
  - Design of caching strategies

Warwick Business School

wbs.ac.uk

#### Challenges

Warwick Business School

- 1. Simulations are mostly black boxes
- 2. Simulations are computationally expensive
- 3. Simulations are often stochastic





#### **Dealing with expensive evaluations**

#### How long can you wait?

- EAs typically require 100,000 function evaluations
- If every simulation takes 1 minute...
- ... this is 70 days runtime!



### Surrogate models

- "Substitute"
- Often also called "Metamodel"



- Much cheaper, but not necessarily as accurate
- Replace some of the expensive function evaluations by surrogate-model based evaluations

For survey, see e.g. [Jin 2011]

Parallellisation

Warwick Business School



- a. Parallel evaluations
  - Multi-core/Multi-processor
  - Graphics Processing Units
- b. Parallel selection
  - Grid/cloud computing

e.g. [Nedjah et al. 2006]



wbs.ac.uk

# Where does the surrogate model come from?

Simplified:

Warwick Business School

- More coarse grained simulation
- Smaller simulation
- $\odot$  Abort simulation early

#### Learned:

From data systematically sampled from search space

• From data collected during the run

#### **Metamodel design questions**

- Type of model, or ensemble
  - Linear/quadratic regression, Gaussian Process, Artificial Neural Network, etc.
- Training data

Warwick Business School

- Global vs. local models
- Which individuals to evaluate based on metamodel, which on full model

#### Learning vs. optimisation

From an optimisation point of view, we want to

- Fully evaluate the best solutions
- Fully evaluate where we are most uncertain
- Ensure the selection works accurately

From a modelling point of view, we want to

- Evaluate where we can most improve model accuracy
- Evaluate where we are most uncertain
- EAs evaluate many solutions in promising areas, so these areas can be modelled accurately
- Model does not need to accurately predict fitness, only accurately predict ranking

#### Which solutions to evaluate?

- Promising solutions?
- Representative solutions?
- Solutions where surrogate model is uncertain?
- Solutions that improve accuracy of surrogate model?
- Fixed or flexible budget?

## Most typical uses of metamodels

Pre-selection

Warwick Business Schoo

Locally optimise each solution

wbs.ac.uk

#### **Pre-selection**

- Generate an abundance of children
- $\odot$  Pre-select  $\lambda$  children based on metamodel
- Fully evaluate pre-selected children



#### **Trust region method**

For each individual

- Repeat at most k times, or until no better solution found
  - Build local surrogate model
  - Perform local search on surrogate model within Trust region Evaluate best found solution
  - Replace individual with best found solution if better Adapt Trust region

Warwick Business School

wbs.ac.uk

#### Benefit of pre-selection [Branke&Schmidt 2004]



## **Multi-fidelity optimisation**

- Sometimes, multiple surrogate models with different trade off between accuracy and running time
- Use fast, rough models to approximate good region
- Use slower, more accurate models to refine the best solution

## Use of partially converged simulation

#### [Branke et al. 2017]

- Simulation can be stopped, and later continued
- Evaluate all solutions using short runs
- Clearly good solutions survive, clearly poor solutions are discarded, remaining are run for longer



#### **Efficient Global Optimisation (EGO)**



## **Efficient Global Optimisation (EGO)**

[Jones, Schonlau, Welch 1998]

- Fit a Gaussian Process (GP) to data
- Response model provides information about
  - expected value
  - uncertainty
- Use response model to determine next data point (replaces genetic operators)
- Expected improvement makes explicit trade-off between exploration and exploitation

#### **Example: GP in 1 dimension**



#### Max expected improvement principle



#### Outline

- Strategies to deal with expensive evaluations
  - Parallelisation
  - Surrogate models
- Strategies to deal with noise
  - Selecting the best system
  - Simulation optimisation
- Applications

Warwick Business School

- Design of traffic light controller
- Design of dispatching rules

## **EGO** algorithm

Take initial  $n_0$  samples

Build GP model

WHILE stopping criterion not met DO

Estimate hyperparameters using maximum likelihood estimation

Take additional sample at position with maximum El

Update GP model

Return best found solution

#### **Selecting the Best System**



wbs.ac.uk

#### **Ranking and selection problem**

- Select, out of *k* systems, the one with best mean performance
- Let  $X_{ij}$  be output of *j*th replication of *i*th system  $\{X_{ij} : j = 1, 2, \ldots\} \stackrel{i.i.d.}{\sim} \operatorname{Normal}(w_i, \sigma_i^2,) \quad i = 1, \ldots, k$
- Sample statistics:  $\bar{x}_i$  and  $\hat{\sigma}_i^2$  based on  $n_i$  observations seen so far
- Order statistics:  $\bar{x}_{(1)} \leq \bar{x}_{(2)} \leq \ldots \leq \bar{x}_{(k)}$
- Correct selection if selected system (*k*) is the true best system [*k*]

#### **Variance Reduction Techniques**

- Try to reduce variance without additional runs, but instead by influencing the settings of the experiments
  - Common Random Numbers
  - Antithetic Variates
  - Control Variates
  - ...

Warwick Business School

wbs.ac.uk

## **Standard: Equal allocation**

- Sample each system *n* times
- Reduces standard error by  $\frac{1}{\sqrt{n}}$

#### **Common Random Numbers**

Intuition:

- Compare two alternatives under similar conditions
- Keep track of performance differences in identical environments
- The observed differences are more likely attributable to the actual system differences, rather than to the differences in environmental conditions

Var(X-Y)=Var(X)+Var(Y)-2Cov(X,Y)

Hope: Cov(*X*,*Y*)>0

Warwick Business School

#### **Common Random Numbers**

How to generate similar conditions

- Conditions are influenced by random number generator
- But: random number generator cannot produce random numbers
- Note: random numbers are used in different contexts
  - arrival rate of customers
  - processing times
  - action selection by agents
  - etc.
- It is necessary to ensure that the random numbers are used for the same purposes in the simulations of the two systems -> synchronization
- Best way to maintain synchronization: Use separate random number streams to corresponding sources of randomness

Warwick Business School

## Common Random Numbers

- CRN works only if random numbers influence performance in the same way
- Can backfire, if that is not the case (rare)
- -> Run some initial experiments with and without CRN to test the influence of CRN



## **Common Random Numbers**







Comparison of the effect of

- increasing the number of samples (n) and
- using CRN

Warwick Business Schoo

Warwick Business School



But: the effect very much depends on the model properties!Possible problem: may invalidate (or at least complicate) statistical analysis methods (ranking/selection, ANOVA)

#### **Design of experiments**

Stratified sampling

 Latin Hypercube sampling





#### **Performance criteria**

Probability of correct selection (PCS)
 PCS = P(w<sub>(k)</sub> = w<sub>[k]</sub>)
 Probability of good selection (PGS)

Probability of good selection (PGS 
$$PGS = P(w_{(k)} \geq w_{[k]} - \delta)$$

• Expected opportunity cost (EOC)  $EOC = E(w_{[k]} - w_{(k)})$ 

Warwick Business School

## **Comparison of m>2 alternatives**

- Allocate samples sequentially
- Maximise the value of information

- Collect more information on promising solutions
- Collect more information where uncertainty is high

# Myopic approach to maximize probability of correct selection

[Chick, Branke, Schmidt: J. of Computing, 2010]

- Assume we can take only one more sample
- If the sample doesn't change selected solution
   -> information had no value
- PCS: Expected value of information is probability of a change in the index of the individual with the best mean
- EOC: Expected value of information is expected change in the value of the selected individual

wbs.ac.uk

Warwick Business School

## **Expected value of information (PCS)**

Change of best system if

- system (i)  $\neq$  (k) is evaluated and becomes new best system
- system (k) is evaluated and becomes worse than second best

$$EVI_{(i)} = \begin{cases} \Phi_{n_{(i)}-1} \left( \frac{\bar{x}_{(i)} - \bar{x}_{(k)}}{\sqrt{\frac{\hat{\sigma}_{(i)}^2}{n_{(i)}(n_{(i)}+1)}}} \right) & \text{if } (i) \neq (k) \\ \Phi_{n_{(k)}-1} \left( \frac{\bar{x}_{(k-1)} - \bar{x}_{(k)}}{\sqrt{\frac{\hat{\sigma}_{(k)}^2}{n_{(k)}(n_{(k)}+1)}}} \right) & \text{if } (i) = (k) \end{cases}$$

#### Stopping rule [Branke, Chick, Schmidt, Mngmt Sci, 2007]

- So far: Fixed budget
- Now: Estimate Probability of Correct Selection (PCS)

$$\begin{aligned} \operatorname{PCS}_{\operatorname{Bayes}} &= \operatorname{Pr}(W_{(k)} \geq \max_{j \neq (k)} W_{(j)}) \mid \Xi \\ &\geq \prod_{j:(j) \neq (k)} \operatorname{Pr}(W_{(k)} > W_{(j)}) \mid \Xi \\ &\approx \prod_{j:(j) \neq (k)} \Phi_{\nu_{(j)(k)}}(d_{jk}^{*}) \\ &\text{with } d_{jk}^{*} = (\bar{x}_{(k)} - \bar{x}_{(j)}) \left( \frac{\hat{\sigma}_{(k)}^{2}}{n_{(k)}} + \frac{\hat{\sigma}_{(j)}^{2}}{n_{(j)}} \right)^{-1/2} \end{aligned}$$

#### Algorithm

Warwick Business School

Sample each alternative  $n_0$  times Determine sample statistics  $\bar{x}_i$  and  $\sigma_i^2$  and order statistics  $\bar{x}_{(1)} \leq \ldots \leq \bar{x}_{(k)}$ WHILE stopping criterion not reached DO Take additional sample of system *i* with maximal EVI Update sample and order statistics Pick solution with maximal  $\bar{x}_i$ 

#### Empirical evaluation (find best out of 10 systems)



wbs.ac.uk

#### Similar approaches

- Optimal Computing Budget Allocation (OCBA)
  - Asymptotic assumption [Chen&Lee 2011]
- Racing [Birattari et al. 2010]
  - In each iteration, allocate one sample to each alternative "still in the race"
  - F-test to detect whether there is significant difference
  - Eliminate alternatives that are significantly worse than best alternative
  - Stop when budget has been used up or only one alternative is left
  - Version that runs with fixed budget [Branke&Elomari 2012]

```
Warwick Business School
```

## If the number of alternatives is large

- If the number of possible alternatives is large, it is no longer possible to evaluate each alternative a few times
- We need an optimization heuristic
- Typical: Simulated annealing, evolutionary algorithm
- If computational budget is very limited, dimensionality small and variables continuous: Bayesian optimisation

#### **Bayesian Optimisation**

- Extends naturally to the noisy case
- Sequential Parameter Optimisation [Bartz-Beielstein et al. 2005], Stochastic Kriging Optimization [Huang et al. 2006], Knowledge Gradient [Frazier et al. 2009]

## **Simulation Optimisation**



wbs.ac.uk

Warwick Business School



#### **Simulated Annealing**



#### **Simulated Annealing**

- Stochastic local search
- Inspired by physical annealing processes



#### **Effect of noise**





## Simulated Annealing in Noisy Environments (SANE) [Branke et al 2008]



## Simulated Annealing in Noisy Environments (SANE) [Branke et al. 2008]

#### Idea:

- Always accept seemingly better solution
- Number of samples depends on temperature and probability to accept worse solution
- Keep sampling until the probability to erroneously select the worse solution is smaller than the acceptance probability for the worse solution

#### **Empirical comparison**

- TSP with normally distributed noise
- Comparison with Alkhamis et al. (AAT) [Alkhamis, Ahmed, Tuan, 1999]



## **Optimal Stochastic Annealing (OSA)**

[Ball et al. 2017]

Warwick Business School

- Assumes known Gaussian noise
- Uses sequential sampling
- At every stage, decision to accept, reject or continue
- Acceptance criterion modified to maintain detailed balance
- Acceptance criterion has optimal efficiency (acceptance probability per sample)

## **Benchmark algorithms**

- SANE [Branke et al. 2007]
  - Sequential sampling and adjusted acceptance criterion
  - Current state-of-the-art, shown to outperform several other methods
- CD1 [Ceperley&Dewing 1999]
  - Adjusted acceptance criterion, obeys detailed balance
- CD10[Ceperley&Dewing 1999]
  - As CD1, but with 10 samples per move decision

#### **OSA acceptance rule**

• Based on sum of samples taken so far

$$c_n = \sum_{i=1}^n \delta_i$$

• Acceptance probability at current stage:

$$A(c_n, c_{n-1}) = \begin{cases} 1 & c_n < -\beta\sigma^2/2 \\ e^{-2(c_n + \beta\sigma^2/2)(c_{n-1} + \beta\sigma^2/2)} & \text{otherwise} \end{cases}$$

- If not accepted, reject if  $c_n > 0$
- Continue otherwise

#### Efficiency (σ/T=10)

Warwick Business School



wbs.ac.uk

Warwick Business School



#### **Optimization performance** (TSP, σ<sup>2</sup>=3200)

#### **Populations are robust to noise**

- Implicit averaging over the neighbourhood
- With infinite populations, fitness proportional selection is not affected by noise [Miller & Goldberg 1996]
- Theory for optimal population sizes in simplified Cases [Arnold&Beyer 2000]
- Black-box Optimization Benchmark competitions show advantages of EAs in noisy environments

#### **Evolutionary algorithm**





#### **Explicit averaging**

Reduce noise by factor sqrt(n)

88

ŧ

222 10

1000000011

Warwick Business School

## **CRN and Evolutionary Algorithms**

#### [Branke 2001]

Warwick Business School

- Use CRN for all individuals to be compared within a generation
  - may drastically improve probability of correct ranking
  - risk of optimizing for one random seed
- Change random number seeds from generation to generation
  - Only individuals that work on a wide range of scenarios will survive for a long time
- Re-evaluate elite individuals
  - A "lucky" individual should be prevented from surviving forever

# Optimal distribution of samples over run [Branke 2001]



#### Change sample size over the run

- Increase sample size over the run [Aizawah&Wah 1993]
- Optimise distribution of samples over the run [Branke 2001]
- Clean up after optimisation [Boesel et al. 03]

#### Use metamodels – average over space



#### **Benefit**



#### The relevant comparisons

 Steady-State-EA with 2-Tournament Population size: 9, offspring: 1

- Replacement: Worst individual
- Stopping criterion: Best individual
- Selection: Best out of {3, 7} and {2, 5}
- (5,10)-Evolution strategy Population size:5, offspring: 10
  - Replacement: 5 best individuals
  - Stopping criterion : Best individual







#### Integrating Ranking&Selection

[Schmidt et al. 2006]



## **Integrating OCBA and EA**

#### Procedure OCBAEA

- 1. Evaluate each new individual  $n_0$  times. Estimate the ranks
- 2. Determine set of relevant comparisons C
- 3. WHILE evidence is not sufficient
  - a) allocate new sample to individual according to modified OCBA rule
  - b) if ranks have changed, update C

#### Benefits over the run



#### Outline

- Strategies to deal with expensive evaluations
  - Parallelisation
  - Surrogate models
- Strategies to deal with noise
  - Selecting the best system
  - Simulation optimisation
- Applications

Warwick Business School

- Design of traffic light controller
- Design of dispatching rules
- Design of caching strategies

## Integrating KG and CMA-ES

- $\odot$  CMA-ES only needs to identify top  $\mu$  individuals
- Uses this to adapt the mutation step size
- $\odot\,$  Which individual, if re-evaluated, has biggest

rror

- potential impact on resulting mutation distribution?
- See paper at this GECCO



wbs.ac.uk

# **Example: Optimisation of a traffic light controller**



#### **Results**



#### **Dispatching rule based scheduling**

- Determine job priorities based on job and machine attributes
- Whenever a machine becomes idle, process the job with highest priority next
- Popular examples: FIFO, SPT, EDD, CoverT

Advantages:

Warwick Business School

- · Always take latest information into account
- Easy to implement and to compute

#### Example 2: Evolution of Job Shop Dispatching Rules [Pickardt et al. 2012]

- Complex wafer manufacturing system
  - 31 work centres (35 machines)
  - 10 batching machines, 2 machines with setup times
  - 7 different products
  - 20-100 operations per job (cycles)





wbs.ac.uk

# Automatic generation of dispatching rules

- Genetic Programming can generate Lisp expressions
- Evaluation of a dispatching rule via stochastic simulation



wbs.ac.uk

#### **Terminals**

- Processing time
- Processing time on next machine
- Number of operations remaining
- Remaining processing time
- Work in next queue
- Time in queue
- Time in system
- Slack
- Time until deadline
- Weight
- Setup time
- Number of compatible jobs for batching

Warwick Business School

wbs.ac.uk

Results (2)

#### Comparison with best rules from literature

Util 93.8%;	Product mix 30/70	Util <mark>85</mark> %; Pr	85%; Product mix 30/70		
Rule	WeightedTardiness	Rule	WeightedTardiness		
ATCS/MBS(5)	2336	ATCS/MBS(4)	451		
GP98	782	GP98	47		

Util <mark>85</mark> %; Produ	ct mix 70/30		Util 93.8%; Product mix <mark>70/30</mark>			
Rule	WeightedTardiness		Rule	WeightedTardiness		
WMOD/MBS(1)	216		WMOD/MBS(3)	1245		
GP98	51		GP98	206		

Warwick Business School

wbs.ac.uk

#### Results

- Rule of length 9: w/max(L,P)-s+b
- Rule of length 98:

$$\begin{split} & \text{ifte}(\max(1,r) - \max(1,r,L),w,b) * b * \max(r/L + \max(-\text{ifte}(b-L,w,b) + s + b, S + b * ifte(\max(1,r) - \max(L,d),w,b) - s - \max(1,r,L) + \max(1,r) + 1) * \text{ifte}(b-L,w,b) - s, S + b * \text{ifte}(\max(1,r) - L,w,b) * (2*r/L - s) + r/L - s + 1) \end{split}$$

#### **Example 3: Evolving Caching Strategies**

The WWW: A huge distributed database Documents are relayed by a sequence of routers

[Branke et al. 2007]



#### The same document is sent many times

Problems: congestion, delays, timeouts, ...



#### With En-Route Caching

Second request can be serviced from nearby replica



#### **Potential solution: Caching**

- Storing replicas of frequently requested documents at nearby nodes
- Possible because requested servers and documents have powerlaw-distribution
- Common on browser level and proxy level
- New idea: En-route web caching

#### Goals:

- Reduce Internet traffic
- Reduce load on highly requested servers
- Reduce latencies

#### **En-Route Caching**

- All nodes/routers involved in relaying an object have an opportunity to keep a cached copy
- Potential for huge resource savings
- Problem: limited memory



#### Caching Policy = Decision Rule

• Question: Which documents to delete?

#### **Challenges**

- Large, dynamically changing distributed system
- No global information
- No global authority
- How good a strategy is depends on what neighbors do
- Symmetry problem
- How can global efficiency and coordination emerge from local caching rules?

#### State of the art:

- LRU (default rule, ignores network aspects)
- RANDOM (often better than LRU, ignores everything)
- GDSF [Cherkasova 2001]

 $priority = rac{access\ count \cdot distance}{access\ count \cdot distance}$ + aging factorsize

#### Goal:

Warwick Business School

Automated design of better en-route caching strategies Use evolutionary algorithms to explore the space of caching rules



#### **Evolving caching strategies**

Main challenges:

- Definition of search space
  - Identify relevant document attributes
  - Use GP to form priority rule
- How to evaluate a caching strategy
  - Only simulation possible due to emergent behavior
  - Requires parallelisation

Warwick Business School

۲

wbs.ac.uk

#### **GP Evolves Caching Strategies**

Priority rules encoded by GP

Inputs = information about the object

- a) Time of document creation
- b) Document size
- c) Access count
- d) Time of last access
- e) Distance from sender
- f) Frequency (no. accesses / second)
- g) Random Constant

#### Functions:

+ \* - / sin cos exp iflte

**Output** = priority

#### Warwick Business School

## **Network Simulator**

- 1. Create network topology
  - Find paths from all to all
- 2. Create random set of objects
  - Random size
  - Power-law distributed set of demands
  - Distributed among hosts
- 3. Poisson process for each host
  - Generate requests for documents
- 4. Routing
  - Break object in packets, send them along shortest path
- 5. Bandwidth
  - Each network link is a queue of requests
- 6. Caching
  - Always send first replica found down request path

## Internet-like Networks (scale-free)

- Internet-like random networks
  - 100 nodes
  - Scale-free topology (Bu & Towsley 02)
- Noisy fitness
  - Different random topologies and request patterns lead to large differences in latency
  - Test in 3 different random nets per generation
  - Change test scenario in every generation
  - Evaluate results on many networks

Warwick Business School

wbs.ac.uk

wbs.ac.uk

wbs.ac.uk

## RUDF

• Often, resulting rules are very complex

(\* (\* b (+ (+ (\* b (+ (+ a (- 0.994 b)) (\* b (+ (\* b e) (exp (exp e)))))) (\* d e)) (exp e))) (iflte (iflte (+ (exp f) (\* (iflte 0.694 f d a) (- b b))) (exp (exp (\* (\* (+ a f) f) (exp c)))) (iflte b (\* c (- 0.139 (\* 1.616 a))) 0.444 (iflte e (iflte d a 0.601 (- f (- a e))) c f)) b) e 0.507 (exp (\* c (+ (+ (exp d) d) (\* a b)))))

• One of the runs yielded a short, powerful rule

#### lastTimeAccessed (distance + accessCount)

• RUDF outperforms all the comparison rules, including GDSF.

#### **Performance on test networks**

	Caching Strategy	Ø rank(latency)
	RUDF	$1.13\pm0.06$
	GDSF	$2.80\pm0.15$
30 nodes	DISTANCE	$3.10\pm0.28$
	LRU	$3.70\pm0.16$
	RANDOM	$4.27\pm0.14$
	RUDF	$1.03\pm0.03$
	DISTANCE	$2.23\pm0.16$
300 nodes	GDSF	$3.20\pm0.11$
	RANDOM	$4.20\pm0.16$
	LRU	$4.33\pm0.13$
		•





#### **Conclusion and additional resources**

#### **Linear Networks**

Warwick Business School

Node	0-	1)-	2	-3)-	(	₽)-	-(5)-	-6)-	7	8	-(9	
Original	0	1	2	3	2	1	5	6	7	8	9	
Replica	5	6	7	8	Ş	9	0	1	2	3	4	
						Са	aching	strategy	A	vg. late	ncy	
Result: Evolved GP policy yields near-optimal performance, far better than comparison algorithms						OPTIMAL			31	31.58 ± 0.03		
						BESTGP			31	$31.98 \pm 0.06$		
						RUDF			35	$35.80 \pm 0.43$		
						GDSF			47	47.67 ± 1.17		
						DISTANCE			50	$50.40 \pm 1.24$		
						RANDOM			61	$61.65 \pm 0.14$		
						LRU				74.77 ± 0.19		

#### • We can find optimal caching analytically

#### Conclusion

Warwick Business School

- Combining simulation and optimisation is the next step in the design of complex systems
- Metaheuristics hold great promise
- Challenges of runtime and noise can be tackled

wbs.ac.uk

#### What I have not talked about

- Stochastic Approximation
- Handling of multiple objectives
- Combinatorial problems
- Warm-up period

Warwick Business School

• Worst-case optimisation



wbs.ac.uk

#### **Further resources**

- The Winter Simulation Conference always has a stream on simulation optimisation
- GECCO tutorial on cloud computing
- Library of simulation optimisation problems <u>http://www.simopt.org</u>
- Available solvers:
  - OptQuest (<u>http://http://www.opttek.com/OptQuest</u>)
  - COMPASS (<u>http://www.iscompass.net</u>)
  - SPOT (<u>https://cran.r-project.org/web/packages/SPOT/</u>)
  - irace (http://iridia.ulb.ac.be/irace/)

#### References

- Aizawa, A.N. and Wah, B. W. (1993). Dynamic control of genetic algorithms in a noisy environment. In International Conference on Genetic Algorithms, pp. 48–55
- Arnold, D. V.; Beyer, H.-G. (2000). Efficiency and mutation strength adap- tation of the -ES in a noisy environment. In Parallel Problem Solving from Nature. LNCS 1917, Springer, pp. 39–48.
- Ball, R.; Branke, J.; Meisel, S. (2017) Optimal Sampling for Simulated Annealing Under Noise. INFORMS Journal on Computing 30(1):200-215
- Bartz-Beielstein, T.; Lasarczyk, C.; Preuß, M. (2005) Sequential parameter optimization. In: McKay B, et al (eds) Congress on Evolutionary Computation, IEEE Press, vol 1, pp. 773–780
- Birattari, M.; Yuan, Z.; Balaprakash, P.; Stützle, T. (2010). F-Race and Iterated F-Race: An overview. In: Experimental Methods for the Analysis of Optimization Algorithms, pp. 311-336
- Boesel, J.; Nelson, B. L.; Kim, S.-H. (2003). Using ranking and selection to clean up after simulation optimization. Operations Research 51(5):814-825
- Branke, J. (2001) Reducing the sampling variance when searching for robust solutions, Genetic and Evolutionary Computation Conference, Morgan Kaufmann, pp. 235-242

Warwick Business School

- Branke, J. (2001). Evolutionary optimization in dynamic environments. Kluwer
- Branke, J.; Chick, S.; Schmidt, C. (2007). Selecting a selection procedure. Management Science 53(12):1916-1932
- Branke, J.; Elomari, J. (2012). Meta-optimization for parameter tuning with a flexible computing budget. Genetic and Evolutionary Computation Conference, ACM, pp. 1245-1252
- Branke, J.; Funes, P.; Thiele, F. (2007). Evolving en-route caching strategies for the Internet. Applied Soft Computing Journal 7(3):890-898
- Branke, J.; Meisel, S.; Schmidt, C.(2008). Simulated annealing in the presence of noise. Journal of Heuristics 14:627-654
- Branke, J.; Asafuddoula, M.; Bhattacharjee, K.; Ray, T. (2017) Efficient Use of Partially Converged Simulations in Evolutionary Optimization, IEEE Transactions on Evolutionary Computation 21(1):52-64
- Chen, H.-C.; Lee, L.-H. (2011). Stochastic simulation optimization: an optimal computing budget allocation. World Scientific
- Chick, S; Branke, J; Schmidt, C.(2010). Sequential sampling to myopically maximize the expected value of information". Informs Journal on Computing 22(1):71-80
- Frazier, P.; Powell, W.; Dayanik, S.: The knowledge gradient policy for correlated normal beliefs. INFORMS Journal on Computing 21(4):599-613

```
Warwick Business School
```

- Fu, M. (2002): Optimization for simulation: Theory vs. practice. Informs Journal on Computing 14(3):192-215
- Fu, M. (ed., 2015): Handbook of Simulation Optimization. Springer
- Hong, L.J.; Nelson, B.L. (2007). Selecting the best system when systems are revealed sequentially. IIE Transactions, 39:723-734
- Huang, D.; Allen, TT.; Notz, WI.; Zeng, N. (2006) Global optimization of stochastic black-box systems via sequential kriging meta-models. Journal of Global Optimization 34(3):441–466
- Jin, Y.; (2011) Surrogate-assisted evolutionary computation: Recent advances and future challenges. Swarm and Evolutionary Computation 1(2):61-70
- Jin, Y.; Branke, J. (2005) Evolutionary optimization in uncertain environments A survey. IEEE Transactions on Evolutionary Computation 9(3):303-318
- Jones, D. R.; Schonlau, M.; Welch, W. J. (1998). Efficient global optimization of expensive black-box functions. Journal of Global Optimization 13:455-492
- Law, A.; Kelton, W. D. (2001). Simulation Modeling and Analysis. McGraw Hill
- Miller, B. L.; Goldberg, D. E. (1996). Genetic algorithms, selection schemes, and the varying effects of noise. *Evolutionary Computation*, 4(2):113–131
- Schmidt, C.; Branke, J.; Chick, S. (2006). Integrating techniques from statistical ranking into evolutionary algorithms. Applications of Evolutionary Computation, Springer, LNCS 3907, pp. 753-762