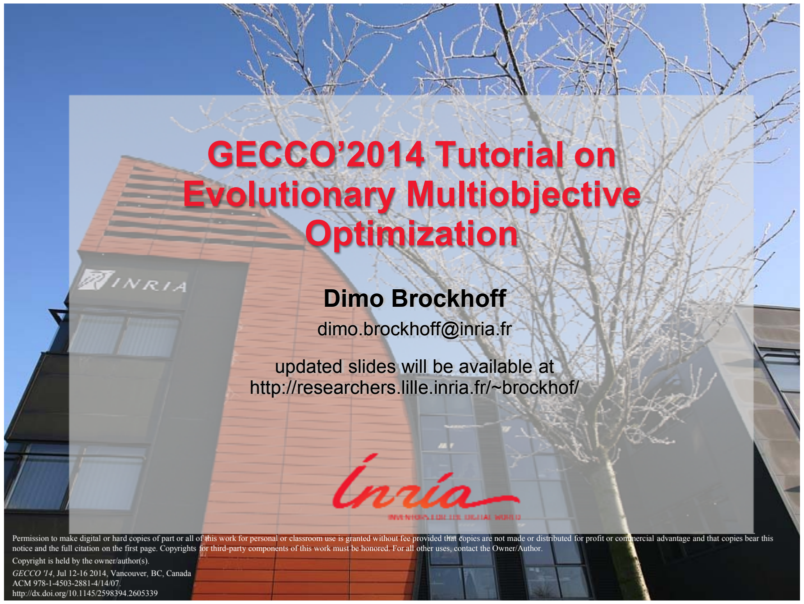

GECCO'2014 Tutorial on Evolutionary Multiobjective Optimization
Dimo Brockhoff
dimo.brockhoff@inria.fr
updated slides will be available at
http://researchers.lille.inria.fr/~brockhoff/
INRIA
Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted 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.
Copyright is held by the owner/author(s).
GECCO '14, Jul 12-16 2014, Vancouver, BC, Canada
ACM 978-1-4503-2881-4/14/07.
http://dx.doi.org/10.1145/2598394.2605339

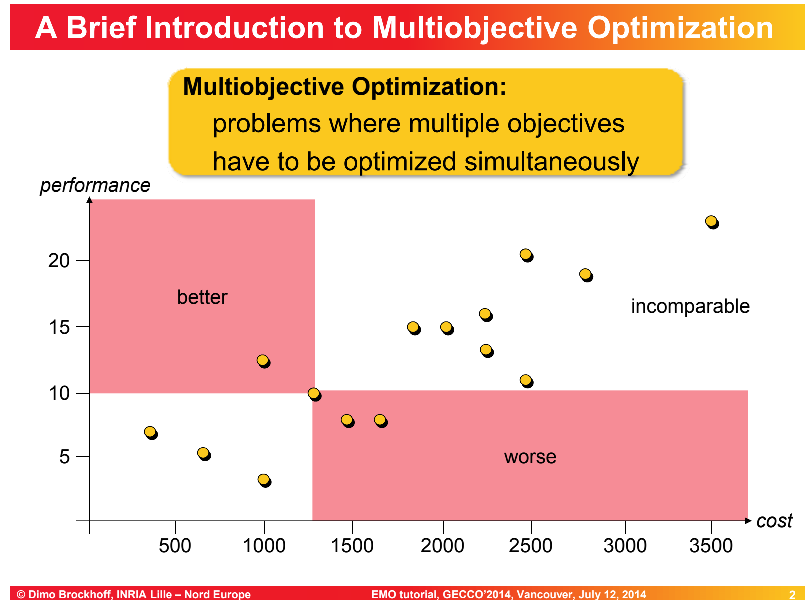

A Brief Introduction to Multiobjective Optimization
Multiobjective Optimization:
problems where multiple objectives have to be optimized simultaneously
performance
20
15
10
5
better
incomparable
worse
500 1000 1500 2000 2500 3000 3500
cost
© Dimo Brockhoff, INRIA Lille – Nord Europe
EMO tutorial, GECCO'2014, Vancouver, July 12, 2014
2

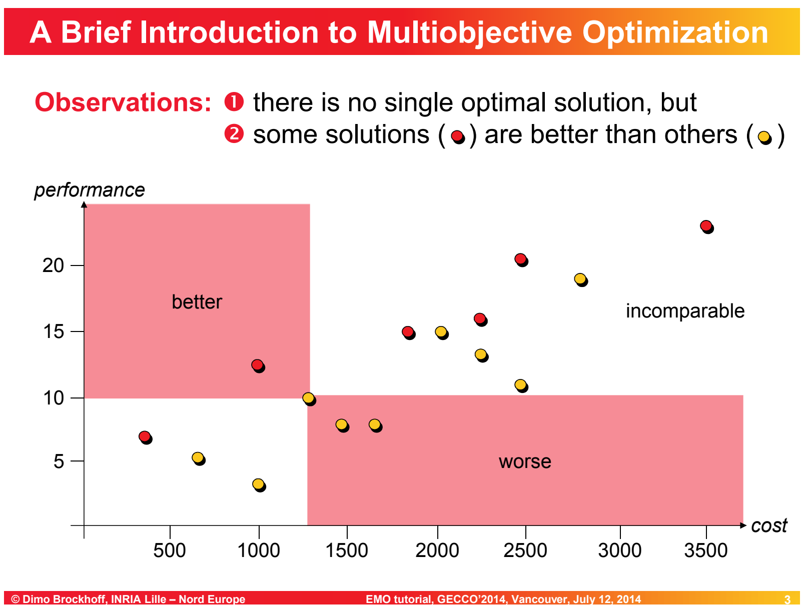

A Brief Introduction to Multiobjective Optimization
Observations: ❶ there is no single optimal solution, but
❷ some solutions ( ) are better than others ( )
performance
20
15
10
5
better
incomparable
worse
500 1000 1500 2000 2500 3000 3500
cost
© Dimo Brockhoff, INRIA Lille – Nord Europe
EMO tutorial, GECCO'2014, Vancouver, July 12, 2014
3

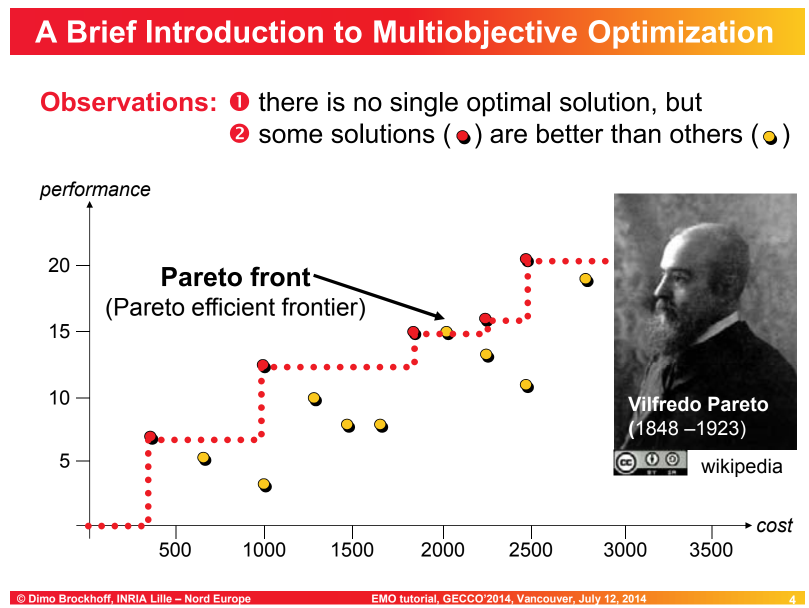

A Brief Introduction to Multiobjective Optimization
Observations: ❶ there is no single optimal solution, but
❷ some solutions ( ) are better than others ( )
performance
20
15
10
5
Pareto front
(Pareto efficient frontier)
Vilfredo Pareto
(1848 –1923)
wikipedia
500 1000 1500 2000 2500 3000 3500
cost
© Dimo Brockhoff, INRIA Lille – Nord Europe
EMO tutorial, GECCO'2014, Vancouver, July 12, 2014
4

A Brief Introduction to Multiobjective Optimization

Observations: ❶ there is no single optimal solution, but
❷ some solutions (●) are better than others (●)

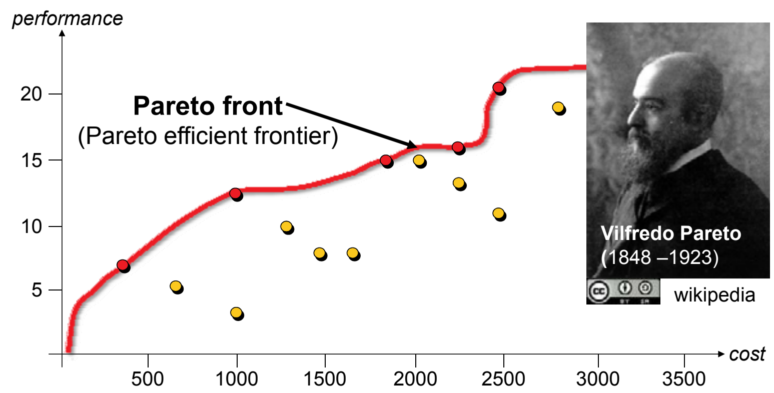

A Brief Introduction to Multiobjective Optimization

Observations: ❶ there is no single optimal solution, but
❷ some solutions (●) are better than others (●)

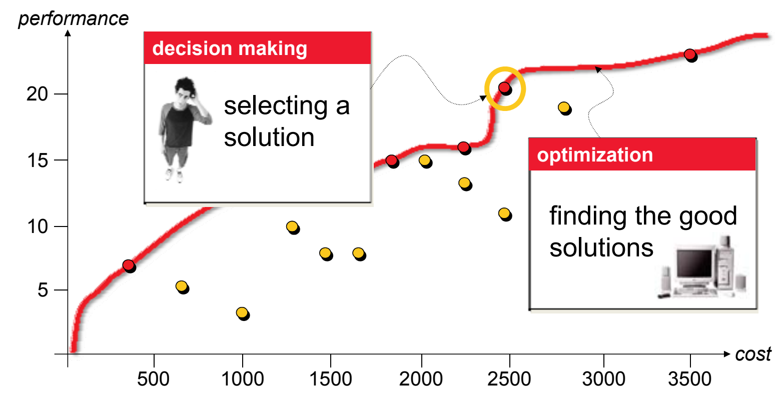

Selecting a Solution: Examples

Possible Approaches: ❶ **ranking:** performance more important than cost

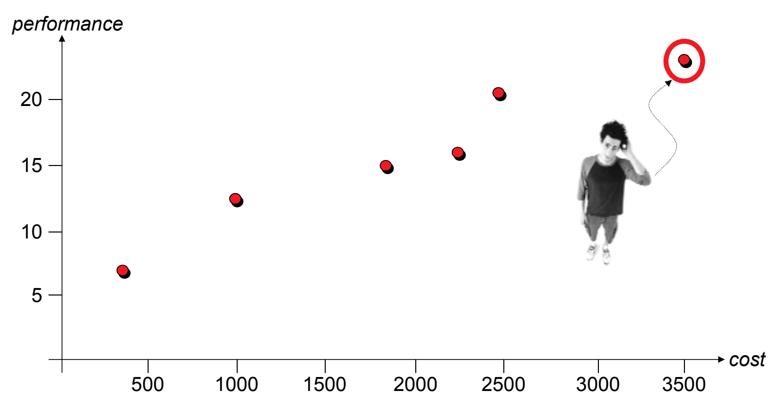

Selecting a Solution: Examples

Possible Approaches: ❶ ranking: performance more important than cost
❷ **constraints:** cost must not exceed 2400

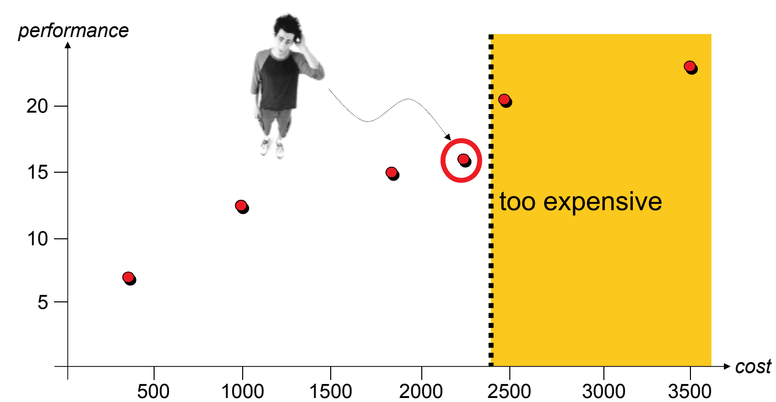

When to Make the Decision

Before Optimization:

rank objectives,
define constraints,…

search for one
(good) solution

When to Make the Decision

Before Optimization:

rank objectives,
define constraints,…

search for one
(good) solution

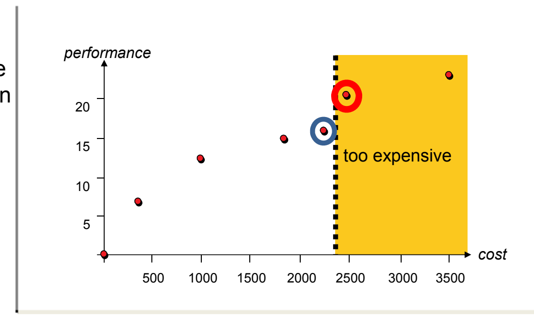

When to Make the Decision

Before Optimization:

rank objectives,
define constraints,…

search for one
(good) solution

After Optimization:

search for a set of
(good) solutions

select one solution
considering
constraints, etc.

When to Make the Decision

Before Optimization:

rank objectives,
define constraints,…

search for one
(good) solution

After Optimization:

search for a set of
(good) solutions

select one solution
considering
constraints, etc.

Focus: learning about a problem

- trade-off surface
- interactions among criteria
- structural information
- also: interactive optimization

Two Communities...

- beginning in 1950s/1960s
- bi-annual conferences since 1975
- background in economics, math, management science
- both optimization and decision making

- quite young field (first papers in mid 1980s)
- bi-annual conference since 2001
- background evolutionary computation (applied math, computer science, engineering, ...)
- focus on optimization algorithms

...Slowly Merge Into One

- MCDM track at EMO conference since 2009
- special sessions on EMO at the MCDM conference since 2008
- joint Dagstuhl seminars since 2004

One of the Main Differences

Blackbox optimization

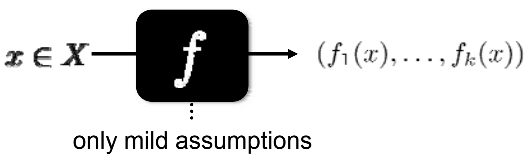

→ EMO therefore well-suited for real-world engineering problems

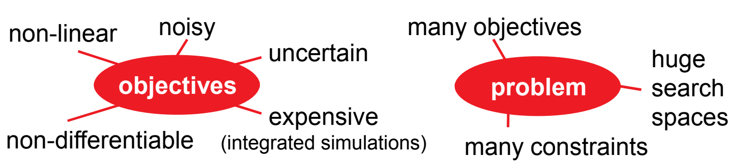

The Other Main Difference

Evolutionary Multiobjective Optimization
- set-based algorithms
- therefore possible to approximate the Pareto front in one run

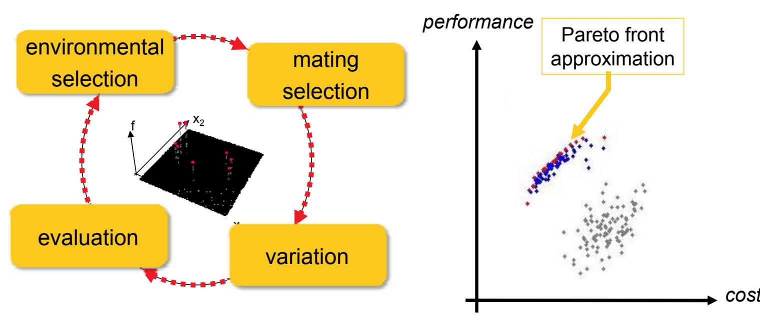

Multiobjectivization

Some problems are easier to solve in a multiobjective scenario

example: TSP
[Knowles et al. 2001]

$\pi \in S_n \rightarrow f(\pi)$ $\pi \in S_n \rightarrow (f_1(\pi, a, b), f_2(\pi, a, b))$

Multiobjectivization

by addition of new "helper objectives" [Jensen 2004]

job-shop scheduling [Jensen 2004], frame structural design [Greiner et al. 2007], VRP [Watanabe and Sakakibara 2007], ...

by decomposition of the single objective

TSP [Knowles et al. 2001], minimum spanning trees [Neumann and Wegener 2006], protein structure prediction [Handl et al. 2008a], ...

also backed up by theory e.g. [Brockhoff et al. 2009, Handl et al. 2008b]

Innovization

Often innovative design principles among solutions are found

example:
clutch brake design
[Deb and Srinivasan 2006]

min. mass +
stopping time

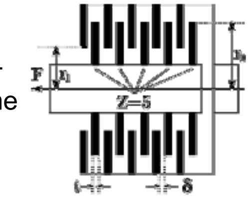

Innovization

Often innovative design principles among solutions are found

example:
clutch brake design
[Deb and Srinivasan 2006]

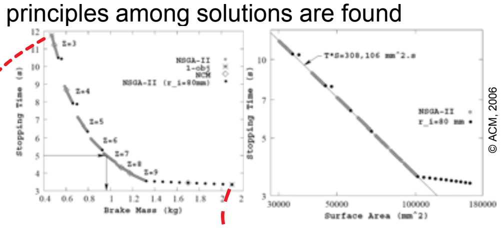

Solution	x_1	x_2	x_3	x_4	x_5	f_1	f_2
Min. f_1	70	90	1.5	1000	3	0.4704	11.7617
Min. f_2	80	110	1.5	1000	9	2.0948	3.3505

Innovization

Often innovative design principles among solutions are found

example:
clutch brake design
[Deb and Srinivasan 2006]

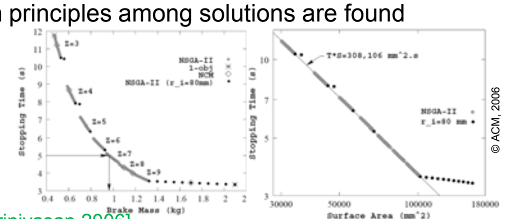

Innovization [Deb and Srinivasan 2006]

= using machine learning techniques to find new and innovative design principles among solution sets

= learning about a multiobjective optimization problem

Other examples:

- SOM for supersonic wing design [Obayashi and Sasaki 2003]
- biclustering for processor design and KP [Ulrich et al. 2007]

The History of EMO At A Glance

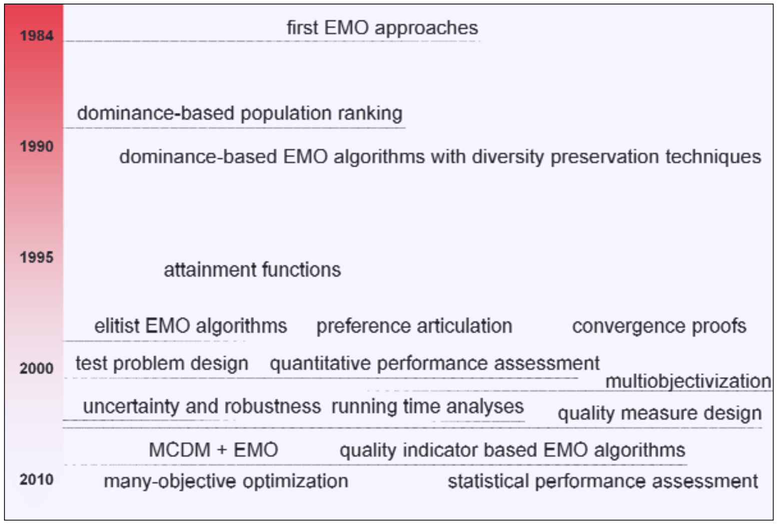

The History of EMO At A Glance

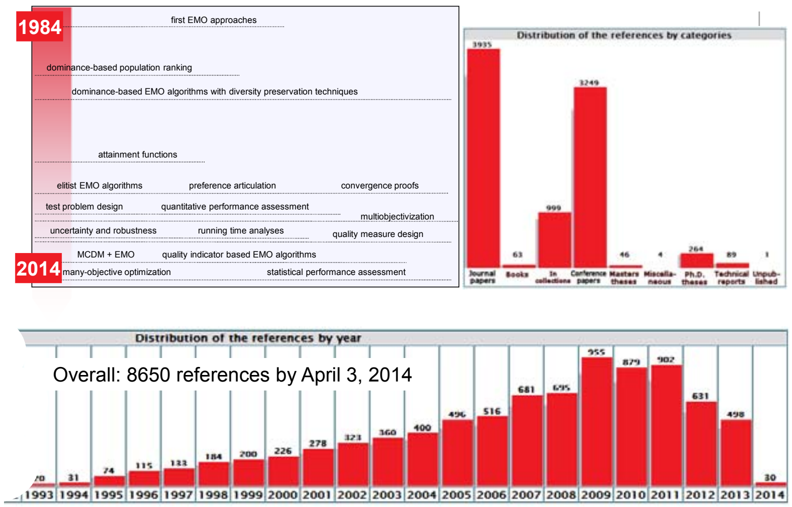

The EMO Community

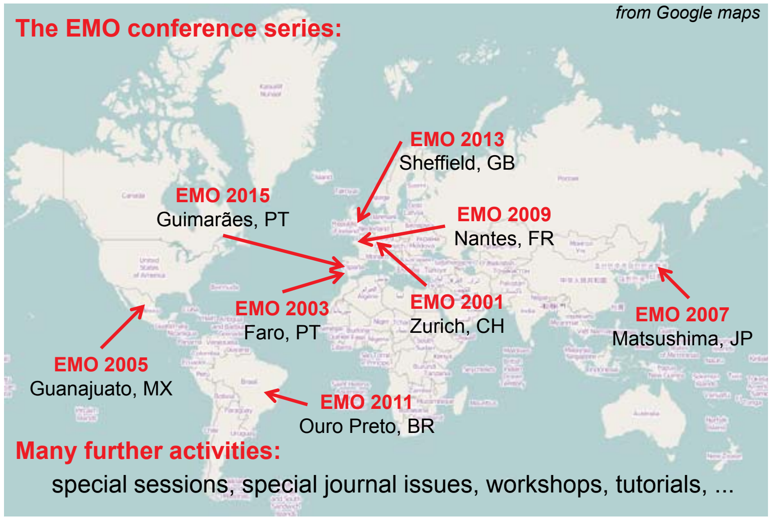

Overview

The Big Picture

Basic Principles of Multiobjective Optimization
- algorithm design principles and concepts
- performance assessment

Selected Advanced Concepts
- indicator-based EMO
- preference articulation

A Few Examples From Practice

Starting Point

What makes evolutionary multiobjective optimization different from single-objective optimization?

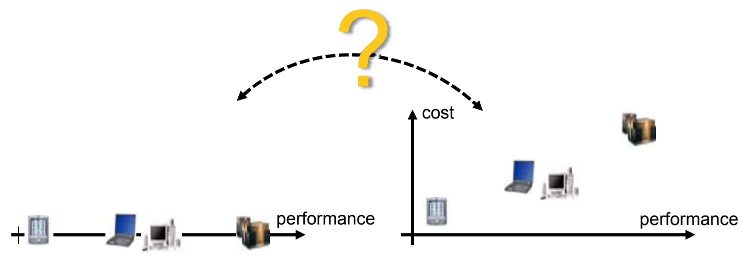

single objective **multiple objectives**

Starting Point

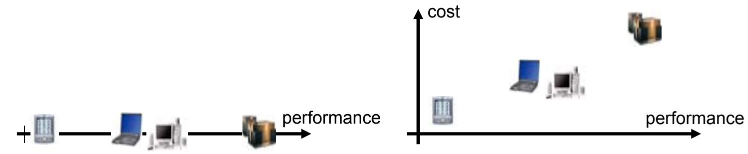

single objective **multiple objectives**

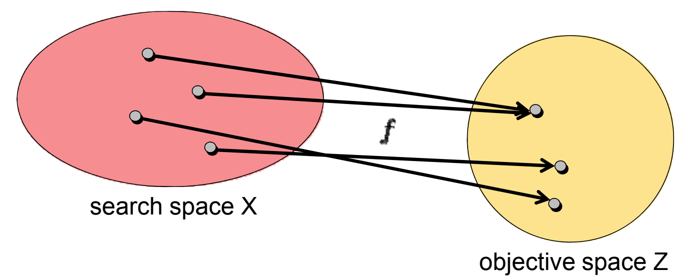

The Main Difference

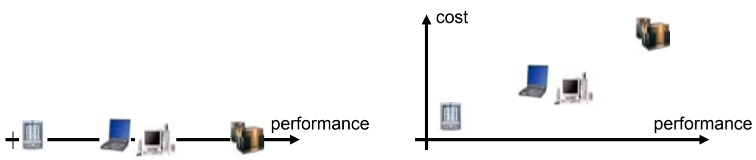

single objective	**multiple objectives**
total order on $f(X) \subseteq \mathbb{R}$	partial order on $f(X) \subseteq \mathbb{R}^k$
total (pre-)order on X	preorder on X
where a better than b if $f(a) < f(b)$	where a better than b if $f(a)$ prefrel $f(b)$

- Pareto dominance
- weak Pareto dominance
- ε-dominance
- cone dominance

The Main Difference

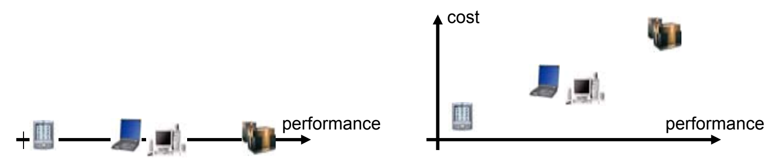

single objective	**multiple objectives**
total order on $f(X) \subseteq \mathbb{R}$	partial order on $f(X) \subseteq \mathbb{R}^k$
total (pre-)order on X	preorder on X
where a better than b if $f(a) \leq f(b)$	where a better than b if $f(a)$ prefrel $f(b)$

even more complicated: sought are **sets**!

Most Common Example: Pareto Dominance

u weakly Pareto dominates v ($u \preceq_{par} v$): $\forall 1 \leq i \leq k : f_i(u) \leq f_i(v)$

u Pareto dominates v ($u \prec_{par} v$): $u \preceq_{par} v \; \wedge \; v \not\preceq_{par} u$

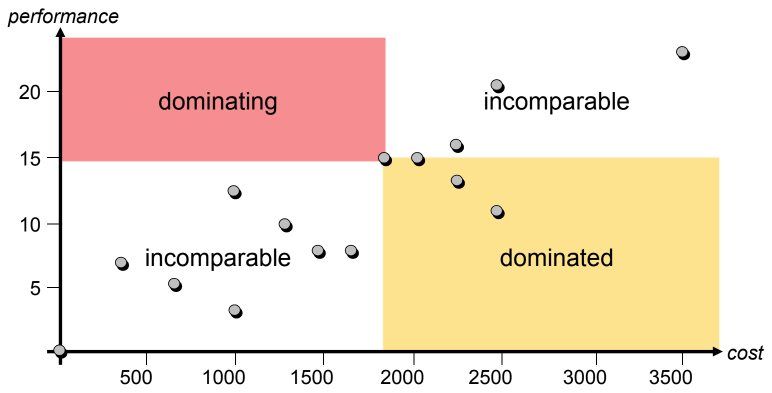

Different Notions of Dominance

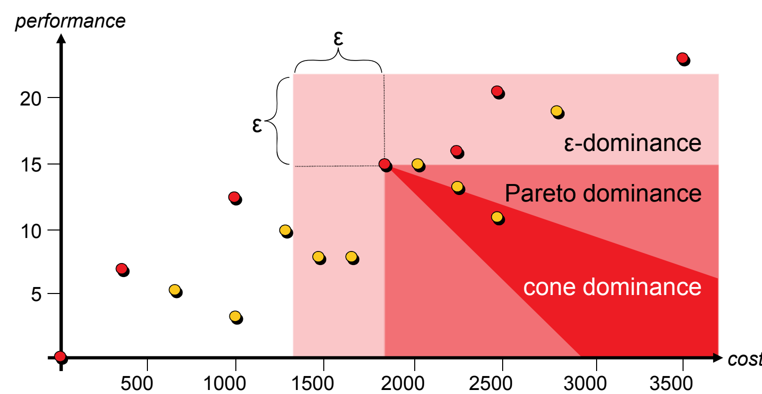

Visualizing Preference Relations

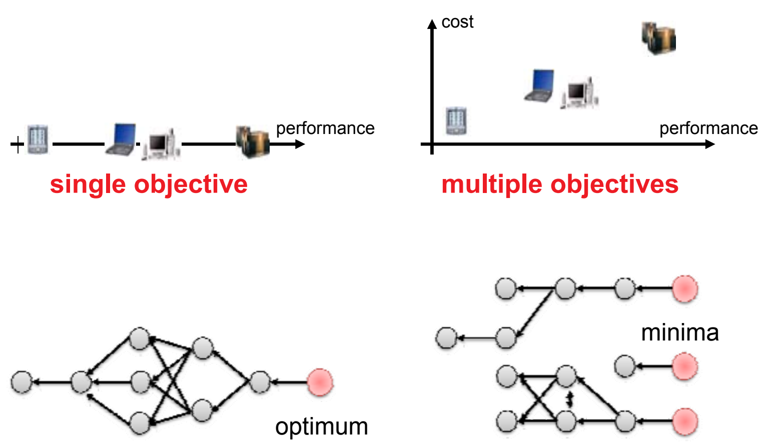

arrow from *a* to *b* if $f(a) < f(b)$

arrow from *a* to *b* if *a* weakly dominates *b*

Visualizing Preference Relations

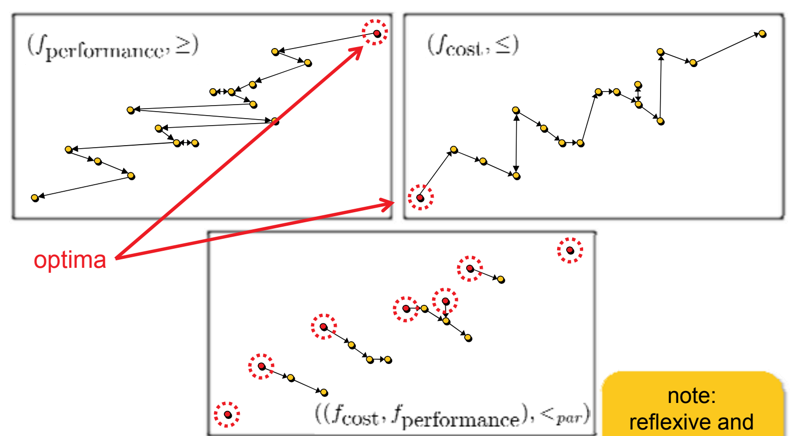

note:
reflexive and
transitive edges
not shown

Pareto-optimal Set and Pareto(-optimal) Front

The *minimal set* of a preordered set $(Y, \leqq)$ is defined as

$$Min(Y, \leqq) := \{a \in Y \mid \forall b \in Y : b \leqq a \Rightarrow a \leqq b\}$$

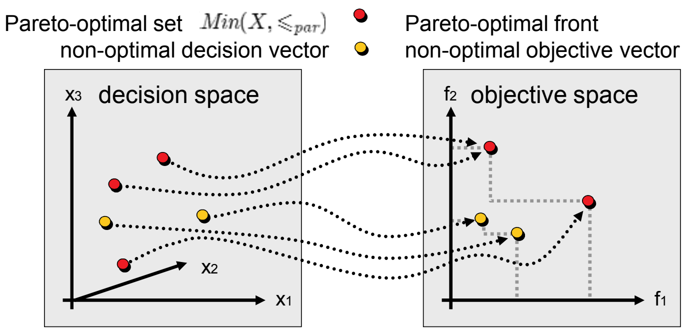

Other Related Definitions

Computational complexity for discrete problems:
multiobjective variants can become NP- and #P-complete

Size: Pareto set can be exponential in the input length
(e.g. shortest path [Serafini 1986], MSP [Camerini et al. 1984])

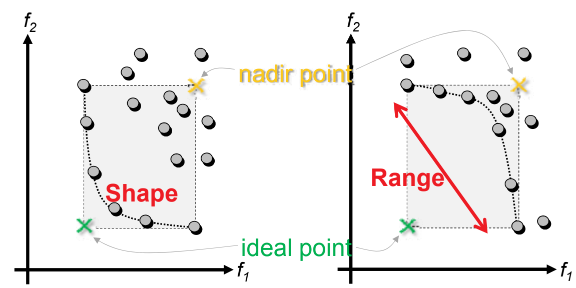

Approaches To Multiobjective Optimization

A multiobjective problem is as such underspecified
…because not any Pareto-optimum is equally suited!

Additional preferences are needed to tackle the problem:

Solution-Oriented Problem Transformation:
Classical approach: Induce a total order on the decision space, e.g., by aggregation

Set-Oriented Problem Transformation:
Recent view on EMO: First transform problem into a set problem and then define an objective function on sets [Zitzler et al. 2010]

Preferences are needed in bases cases, but the latter are weaker!

Solution-Oriented Problem Transformations

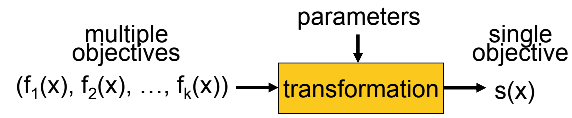

A scalarizing function s is a function $s: Z \to \mathbb{R}$ that maps each objective vector $u = (u_1, \ldots, u_n) \in Z$ to a real value $s(u) \in \mathbb{R}$

Solution-Oriented Problem Transformations

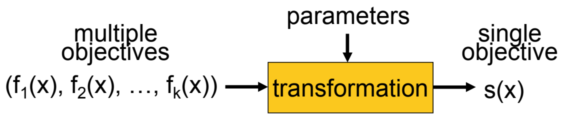

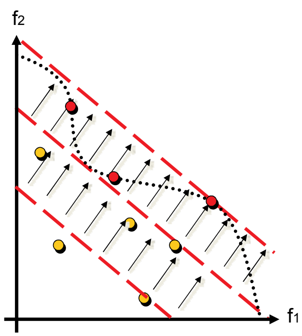

Example 1: weighted sum approach

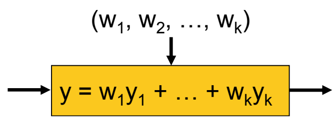

Disadvantage: not all Pareto-optimal solutions can be found if the front is not convex

Solution-Oriented Problem Transformations

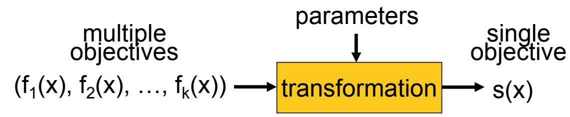

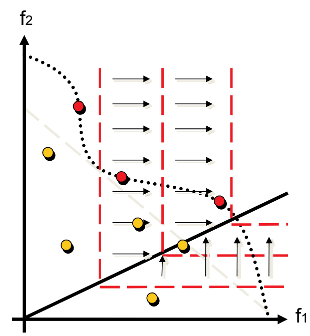

Example 2: weighted Tchebycheff

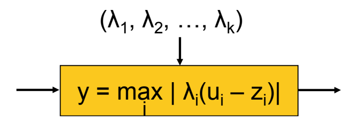

Several other scalarizing functions are known, see e.g. [Miettienen 1999]

Set-Oriented Problem Transformations

For a multiobjective optimization problem $(X, Z, \mathbf{f}, \mathbf{g}, \leqslant)$, the associated *set problem* is given by $(\Psi, \Omega, F, \mathbf{G}, \preccurlyeq)$ where

- $\Psi = 2^X$ is the space of decision vector sets, i.e., the powerset of X,
- $\Omega = 2^Z$ is the space of objective vector sets, i.e., the powerset of Z,
- F is the extension of $\mathbf{f}$ to sets, i.e., $F(A) := \{\mathbf{f}(\mathbf{a}) : \mathbf{a} \in A\}$ for $A \in \Psi$,
- $\mathbf{G} = (G_1, \ldots, G_m)$ is the extension of $\mathbf{g}$ to sets, i.e., $G_i(A) := \max\{g_i(\mathbf{a}) : \mathbf{a} \in A\}$ for $1 \leq i \leq m$ and $A \in \Psi$,
- $\preccurlyeq$ extends $\leqslant$ to sets where $A \preccurlyeq B :\Leftrightarrow \forall \mathbf{b} \in B\ \exists \mathbf{a} \in A : \mathbf{a} \leqslant \mathbf{b}$.

Pareto Set Approximations

Pareto set approximation (algorithm outcome) = set of (usually incomparable) solutions

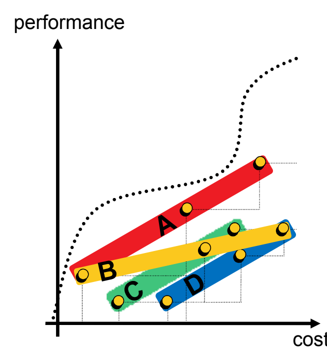

A weakly dominates B
= not worse in all objectives and sets not equal

C dominates D
= better in at least one objective

A strictly dominates C
= better in all objectives

B is **incomparable** to **C**
= neither set weakly better

What Is the Optimization Goal of a Set Problem?

- Find all Pareto-optimal solutions?
  - Impossible in continuous search spaces
  - How should the decision maker handle 10000 solutions?
- Find a representative subset of the Pareto set?
  - Many problems are NP-hard
  - What does representative actually mean?
- Find a good approximation of the Pareto set?
  - What is a good approximation?
  - How to formalize intuitive understanding:
    - ❶ close to the Pareto front
    - ❷ well distributed

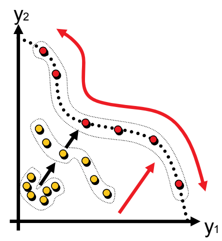

Most common: use of **quality indicators**

Quality of Pareto Set Approximations

A (unary) quality indicator I is a function $I : \Psi = 2^X \mapsto \mathbb{R}$ that assigns a Pareto set approximation a real value.

well-known examples:

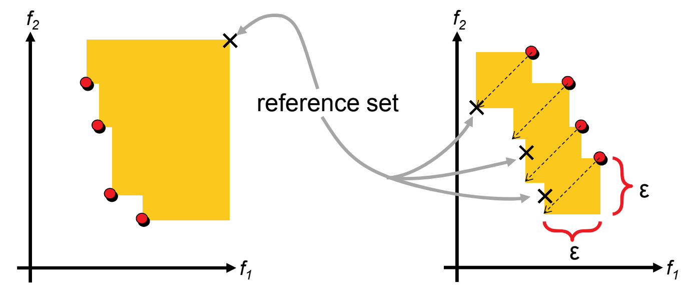

hypervolume indicator   epsilon indicator

Problem Transformations and Set Problems

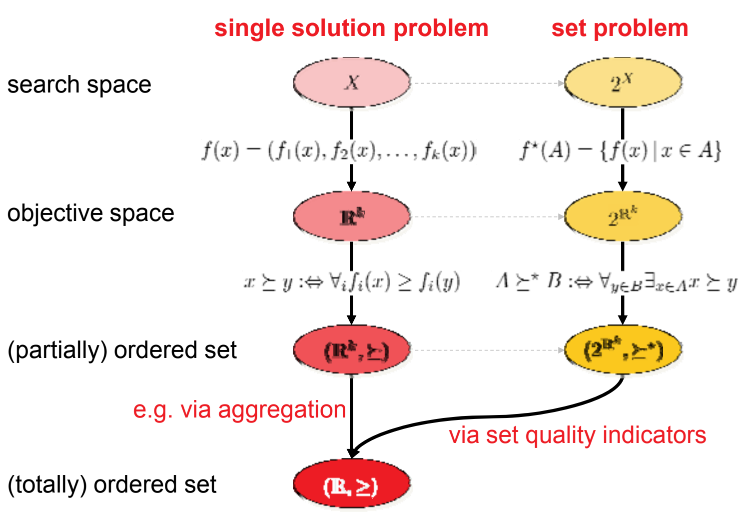

General Remarks on Problem Transformations

Main Goal:

Transform a preorder into a total preorder on X

Methods:

- Define single-objective function based on the multiple criteria *(e.g. via aggregation)*
- Define total preorder on sets by using a quality indicator *(e.g. via hypervolume indicator)*

Question:

Is any total preorder okay or are there any requirements concerning the resulting preference relation?

⇒ Underlying dominance relation should be reflected!

Refinements and Weak Refinements

❶ $\overset{\text{ref}}{\preccurlyeq}$ **refines** a preference relation $\preccurlyeq$ iff

$$A \preccurlyeq B \wedge B \not\preccurlyeq A \Rightarrow A \overset{\text{ref}}{\preccurlyeq} B \wedge B \overset{\text{ref}}{\not\preccurlyeq} A \qquad \textit{(better} \Rightarrow \textit{better)}$$

$\Rightarrow$ fulfills requirement

❷ $\overset{\text{ref}}{\preccurlyeq}$ **weakly refines** a preference relation $\preccurlyeq$ iff

$$A \preccurlyeq B \wedge B \not\preccurlyeq A \Rightarrow A \overset{\text{ref}}{\preccurlyeq} B \qquad \textit{(better} \Rightarrow \textit{weakly better)}$$

$\Rightarrow$ does not fulfill requirement, but $\overset{\text{ref}}{\preccurlyeq}$ does not contradict $\preccurlyeq$

! sought are total refinements… [Zitzler et al. 2010]

Example: Refinements Using Indicators

$A \overset{\text{ref}}{\preccurlyeq} B :\Leftrightarrow I(A) \geq I(B)$

I(A) = volume of the weakly dominated area in objective space

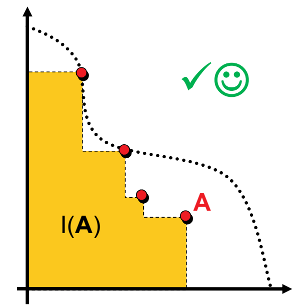

unary hypervolume indicator

$A \overset{\text{ref}}{\preccurlyeq} B :\Leftrightarrow I(A,B) \leq I(B,A)$

I(A,B) = how much needs A to be moved to weakly dominate B

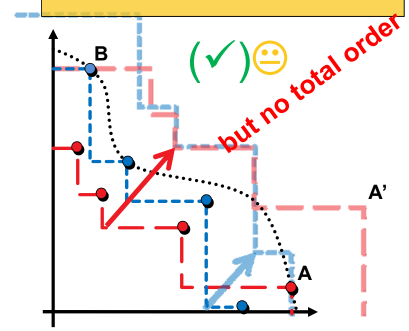

binary epsilon indicator

Example: Weak Refinement / No Refinement

$A \overset{\text{ref}}{\preccurlyeq} B :\Leftrightarrow I(A,R) \leq I(B,R)$

I(A,R) = how much needs A to be moved to weakly dominate R

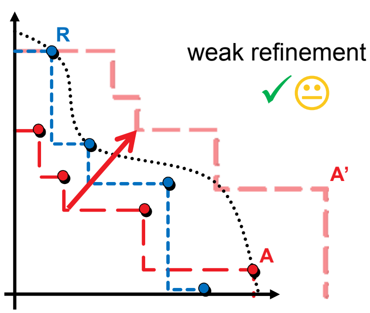

unary epsilon indicator

$A \overset{\text{ref}}{\preccurlyeq} B :\Leftrightarrow I(A) \leq I(B)$

I(A) = variance of pairwise distances

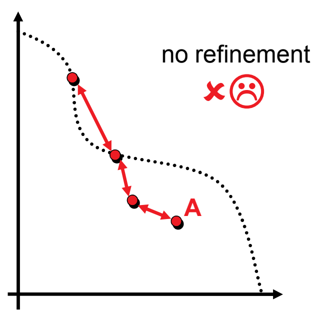

unary diversity indicator

Overview

The Big Picture

Basic Principles of Multiobjective Optimization
- algorithm design principles and concepts
- performance assessment

Selected Advanced Concepts
- indicator-based EMO
- preference articulation

A Few Examples From Practice

Algorithm Design: Particular Aspects

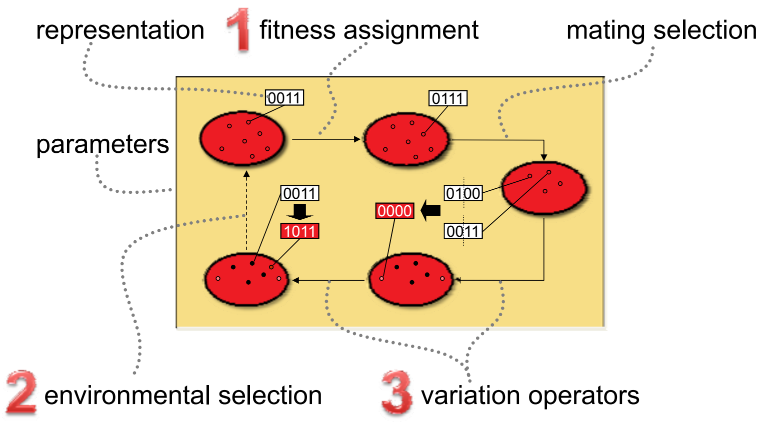

Fitness Assignment: Principal Approaches

aggregation-based
weighted sum
but also decomposition -based EMO

criterion-based
VEGA

dominance-based
SPEA2, NSGA-II
"modern" EMOA

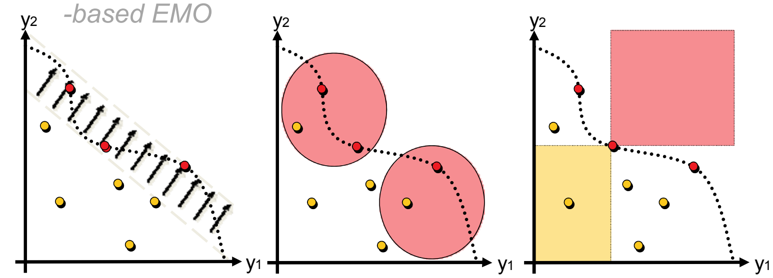

parameter-oriented
scaling-dependent

set-oriented
scaling-independent

Criterion-Based Selection: VEGA

[Schaffer 1985]

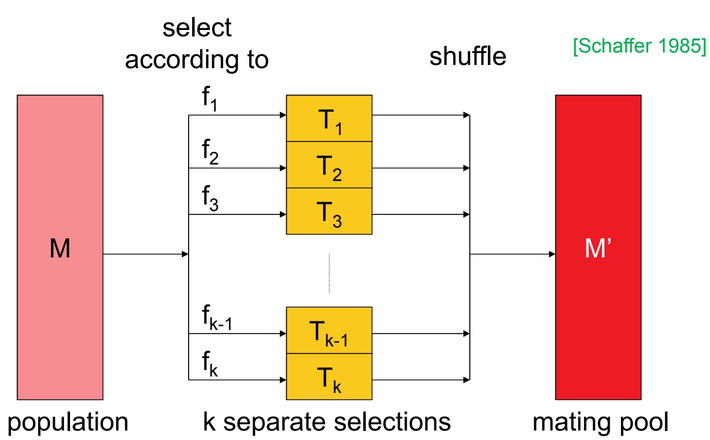

Drawback: only allows to find extremes of the Pareto front

Aggregation-Based: Multistart Constraint Method

Underlying concept:

- Convert all objectives except of one into constraints
- Adaptively vary constraints

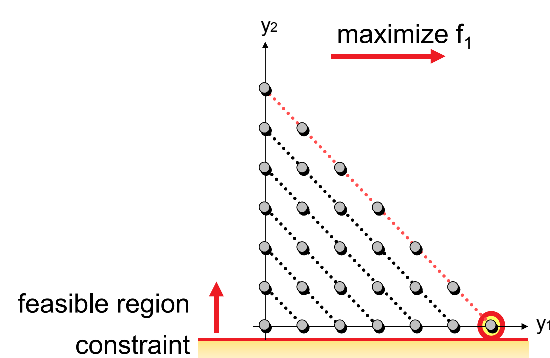

Aggregation-Based: Multistart Constraint Method

Underlying concept:

- Convert all objectives except of one into constraints
- Adaptively vary constraints

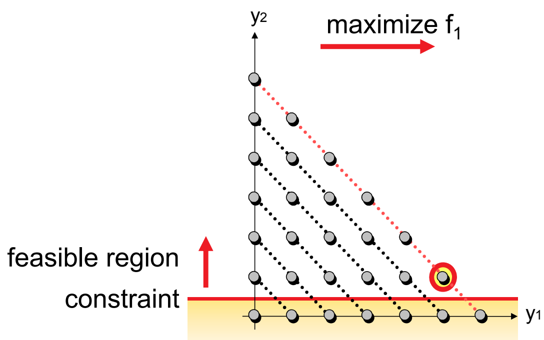

Aggregation-Based: Multistart Constraint Method

Underlying concept:

- Convert all objectives except of one into constraints
- Adaptively vary constraints

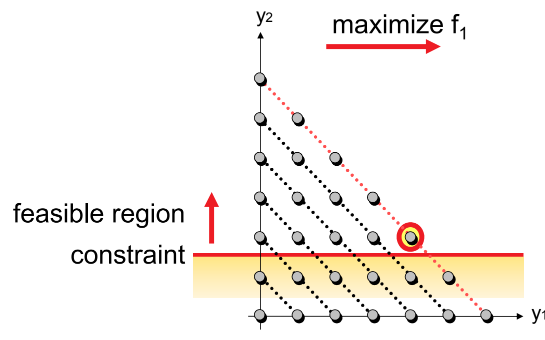

General Scheme of Most Dominance-Based EMO

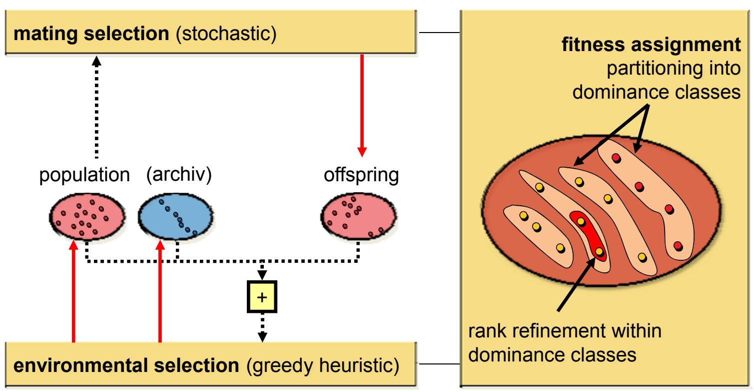

Note: good in terms of set quality = good in terms of search?

Ranking of the Population Using Dominance

... goes back to a proposal by David Goldberg in 1989.

... is based on pairwise comparisons of the individuals only.

- **dominance rank:** by how many individuals is an individual dominated?
  MOGA, NPGA
- **dominance count:** how many individuals does an individual dominate?
  SPEA, SPEA2
- **dominance depth:** at which front is an individual located?
  NSGA, NSGA-II, most of the recently proposed algorithms

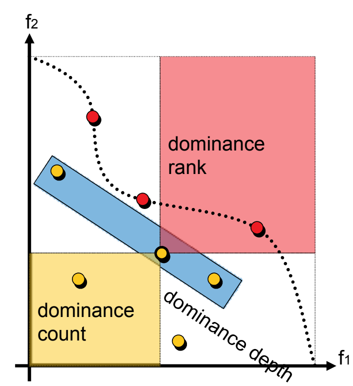

Illustration of Dominance-based Partitioning

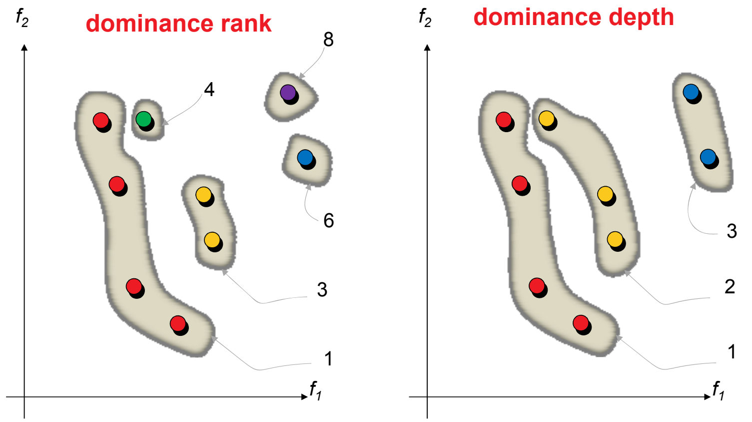

Refinement of Dominance Rankings

Goal: rank incomparable solutions within a dominance class

❶ Density information (good for search, but usually no refinements)

Kernel method	**k-th nearest neighbor**	**Histogram method**
density = function of the distances	density = function of distance to k-th neighbor	density = number of elements within box

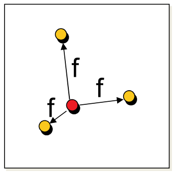

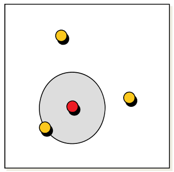

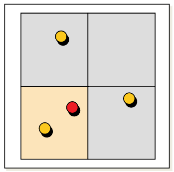

❷ Quality indicator (good for set quality): soon...

Example: NSGA-II Diversity Preservation

Density Estimation

crowding distance:

- sort solutions wrt. each objective
- crowding distance to neighbors:

$$d(i) - \sum_{\text{obj. } m} |f_m(i-1) - f_m(i+1)|$$

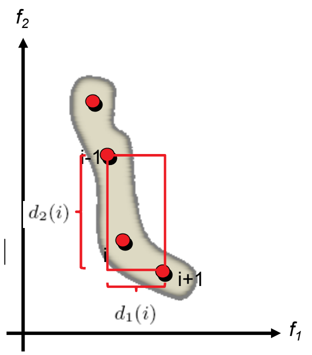

SPEA2 and NSGA-II: Cycles in Optimization

Selection in SPEA2 and NSGA-II can result in *deteriorative* cycles

non-dominated solutions already found can be lost

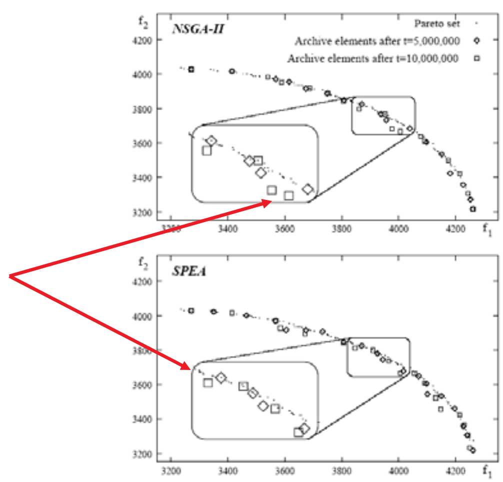

Hypervolume-Based Selection

Latest Approach (SMS-EMOA, MO-CMA-ES, HypE, …)
use hypervolume indicator to guide the search: refinement!

Main idea
Delete solutions with the smallest hypervolume loss $d(s) = I_H(P)-I_H(P / \{s\})$ iteratively

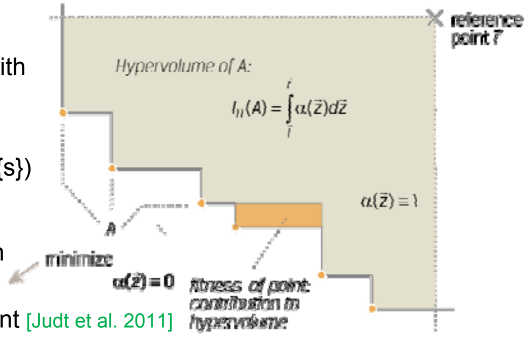

But: can also result in cycles if reference point is not constant [Judt et al. 2011]
and is expensive to compute exactly [Bringmann and Friedrich 2009]

Moreover: HypE [Bader and Zitzler 2011]
Sampling + Contribution if more than 1 solution deleted

Decomposition-Based Selection: MOEA/D

MOEA/D: Multiobjective Evolutionary Algorithm Based on Decomposition [Zhang and Li 2007]

Ideas:

- Optimize N scalarizing functions in parallel
- Use best solutions of "neighbored scalarizing function" for mating
- keep the best solutions for each scalarizing function
- eventually replace neighbors
- use external archive for non-dominated solutions
- several improved versions recently

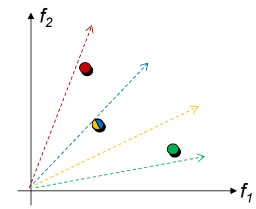

Scalarizing Approaches

Open Questions:

- how to choose "the right" scalarization even if the direction in objective space is given by the DM?
- combinations/adaptation of scalarization functions
- independent optimization vs. cooperation between single-objective optimization

Variation in EMO

- At first sight not different from single-objective optimization
- Most algorithm design effort on selection until now
- But: convergence to a set ≠ convergence to a point

Open Question:

- how to achieve fast convergence to a set?

Related work:

- multiobjective CMA-ES [Igel et al. 2007] [Voß et al. 2010]
- set-based variation [Bader et al. 2009]
- set-based fitness landscapes [Verel et al. 2011]

Overview

The Big Picture

Basic Principles of Multiobjective Optimization

- algorithm design principles and concepts
- performance assessment

Selected Advanced Concepts

- indicator-based EMO
- preference articulation

A Few Examples From Practice

Once Upon a Time...

... multiobjective EAs were mainly compared visually:

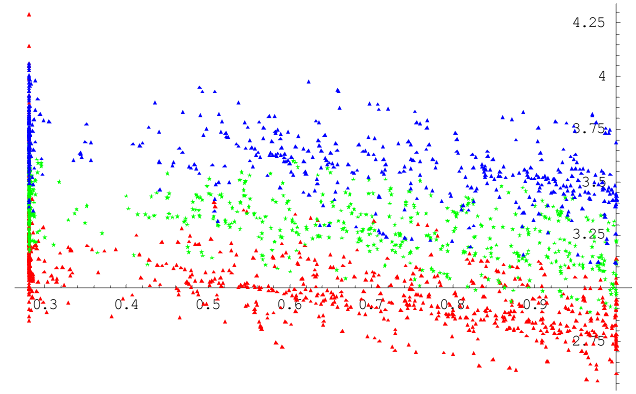

ZDT6 benchmark problem: IBEA, SPEA2, NSGA-II

Two Approaches for Empirical Studies

Attainment function approach:

- Applies statistical tests directly to the samples of approximation sets
- Gives detailed information about how and where performance differences occur

Quality indicator approach:

- First, reduces each approximation set to a single value of quality
- Applies statistical tests to the samples of quality values

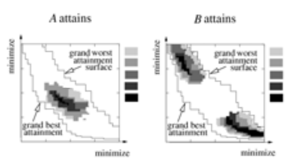

Indicator	A	B
Hypervolume indicator	6.3431	7.1924
ϵ-indicator	1.2090	0.12722
R_2 indicator	0.2434	0.1643
R_3 indicator	0.6454	0.3475

see e.g. [Zitzler et al. 2003]

Empirical Attainment Functions

three runs of two multiobjective optimizers

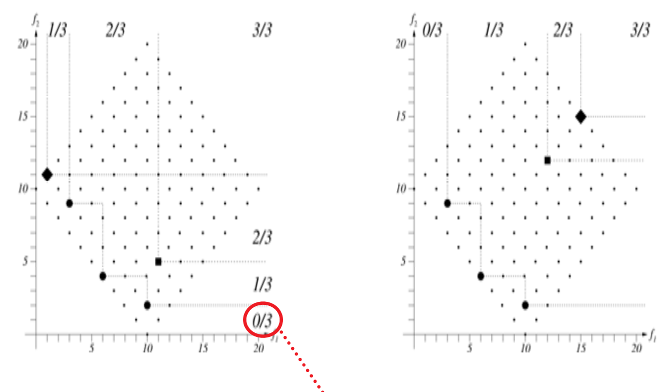

frequency of attaining regions

Attainment Plots

50% attainment surface for IBEA, SPEA2, NSGA2 (ZDT6)

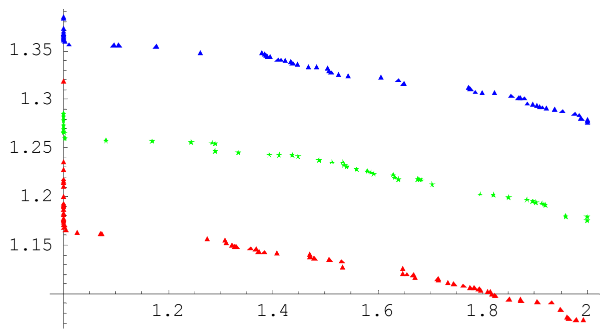

latest implementation online at
http://eden.dei.uc.pt/~cmfonsec/software.html
see [Fonseca et al. 2011]

Quality Indicator Approach

Goal: compare two Pareto set approximations A and B

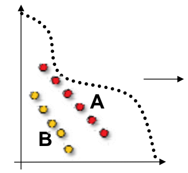

	A	B
hypervolume	*432.34*	*420.13*
distance	*0.3308*	*0.4532*
diversity	*0.3637*	*0.3463*
spread	*0.3622*	*0.3601*
cardinality	*6*	*5*

→ "**A** better"

Comparison method C = quality measure(s) + Boolean function

A, B → (quality measure) → $\mathbb{R}^n$ → (Boolean function) → statement

reduction **interpretation**

Example: Box Plots

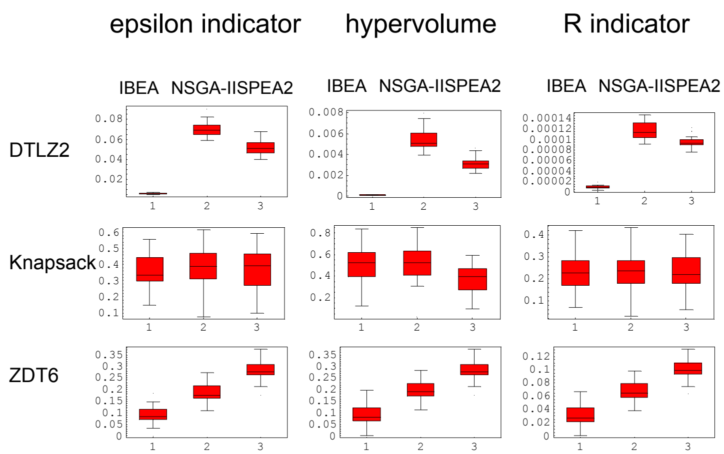

Statistical Assessment (Kruskal Test)

ZDT6
Epsilon

is better than →

	IBEA	NSGA2	SPEA2
IBEA		~0 ☺	~0 ☺
NSGA2	1		~0 ☺
SPEA2	1	1	

Overall p-value = 6.22079e-17.
Null hypothesis rejected (alpha 0.05)

DTLZ2
R

is better than →

	IBEA	NSGA2	SPEA2
IBEA		~0 ☺	~0 ☺
NSGA2	1		1
SPEA2	1	~0 ☺	

Overall p-value = 7.86834e-17.
Null hypothesis rejected (alpha 0.05)

Knapsack/Hypervolume: H_0 = No significance of any differences

Problems With Non-Compliant Indicators

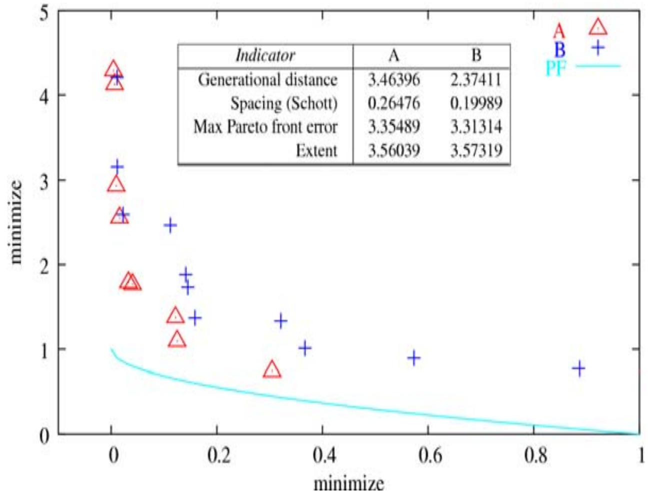

Indicator	A	B
Generational distance	3.46396	2.37411
Spacing (Schott)	0.26476	0.19989
Max Pareto front error	3.35489	3.31314
Extent	3.56039	3.57319

What Are Good Set Quality Measures?

There are **three aspects** [Zitzler et al. 2000]

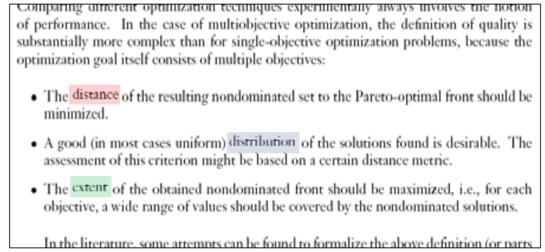
Comparing different optimization techniques experimentally always involves the notion of performance. In the case of multiobjective optimization, the definition of quality is substantially more complex than for single-objective optimization problems, because the optimization goal itself consists of multiple objectives:

- The distance of the resulting nondominated set to the Pareto-optimal front should be minimized.
- A good (in most cases uniform) distribution of the solutions found is desirable. The assessment of this criterion might be based on a certain distance metric.
- The extent of the obtained nondominated front should be maximized, i.e., for each objective, a wide range of values should be covered by the nondominated solutions.

In the literature, some attempts can be found to formalize the above definition (or parts

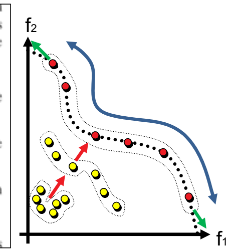

Wrong! [Zitzler et al. 2003]

An infinite number of unary set measures is needed to detect in general whether A is better than B

Set Quality Indicators

Open Questions:

- how to design a good benchmark suite?
- are there other unary indicators that are (weak) refinements?
- how to compute indicators efficiently (enough for practice)?
- how to achieve good indicator values?

Overview

The Big Picture

Basic Principles of Multiobjective Optimization

- algorithm design principles and concepts
- performance assessment

Selected Advanced Concepts

- indicator-based EMO
- preference articulation

A Few Examples From Practice

Indicator-Based EMO: Optimization Goal

When the goal is to maximize a unary indicator…

- we have a single-objective set problem to solve
- but what is the optimum?
- important: population size μ plays a role!

Multiobjective Problem —Indicator→ Single-objective Problem

Optimal μ-Distribution:

A set of μ solutions that maximizes a certain unary indicator I among all sets of μ solutions is called optimal μ-distribution for I. [Auger et al. 2009a]

Optimal μ-Distributions for the Hypervolume

Hypervolume indicator refines dominance relation

$\Longrightarrow$ most results on optimal μ-distributions for hypervolume

Optimal μ-Distributions (example results)

[Auger et al. 2009a]:

- contain equally spaced points iff front is linear
- density of points $\propto \sqrt{-f'(x)}$ with f' the slope of the front

[Friedrich et al. 2011]:

optimal μ-distributions for the hypervolume correspond to ε-approximations of the front

OPT	$1 + \frac{\log(\min\{A/a, B/b\})}{n}$
HYP	$1 + \frac{\sqrt{A/a} + \sqrt{B/b}}{n-4}$
logHYP	$1 + \frac{\sqrt{\log(A/a)\log(B/b)}}{n-2}$

! (probably) does not hold for > 2 objectives

Indicator-Based EMO

Open Questions:

- How do the optimal μ-distributions look like for >2 objectives?
- how to compute certain indicators quickly in practice?
  - several recent improvements for the hypervolume indicator [Yildiz and Suri 2012], [Bringmann 2012], [Bringmann 2013]
- how to do indicator-based subset selection quickly?
- what is the best strategy for the subset selection?

further open questions on indicator-based EMO available at
`http://simco.gforge.inria.fr/doku.php?id=openproblems`

Overview

The Big Picture

Basic Principles of Multiobjective Optimization

- algorithm design principles and concepts
- performance assessment

Selected Advanced Concepts

- indicator-based EMO
- preference articulation

A Few Examples From Practice

Articulating User Preferences During Search

What we thought: EMO is preference-less

> given by the DM.
>
> **Search before decision making:** Optimization is performed without any preference information given. The result of the search process is a set of (ideally Pareto-optimal) candidate solutions from which the final choice is made by the DM.
>
> **Decision making during search:** The DM can articulate preferences during

[Zitzler 1999]

What we learnt: EMO just uses weaker preference information

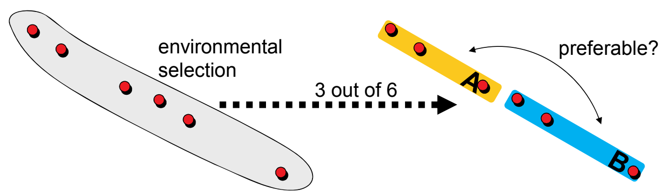

Incorporation of Preferences During Search

Nevertheless...

- the more (known) preferences incorporated the better
- in particular if search space is too large

[Branke 2008], [Rachmawati and Srinivasan 2006], [Coello Coello 2000]

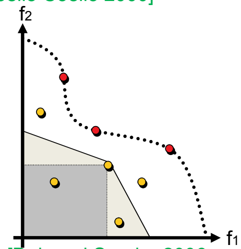

❶ **Refine/modify dominance relation, e.g.:**

- using goals, priorities, constraints [Fonseca and Fleming 1998a,b]
- using different types of cones [Branke and Deb 2004]

❷ **Use quality indicators, e.g.:**

- based on reference points and directions [Deb and Sundar 2006, Deb and Kumar 2007]
- based on binary quality indicators [Zitzler and Künzli 2004]
- based on the hypervolume indicator (now) [Zitzler et al. 2007]

Example: Weighted Hypervolume Indicator

[Zitzler et al. 2007]

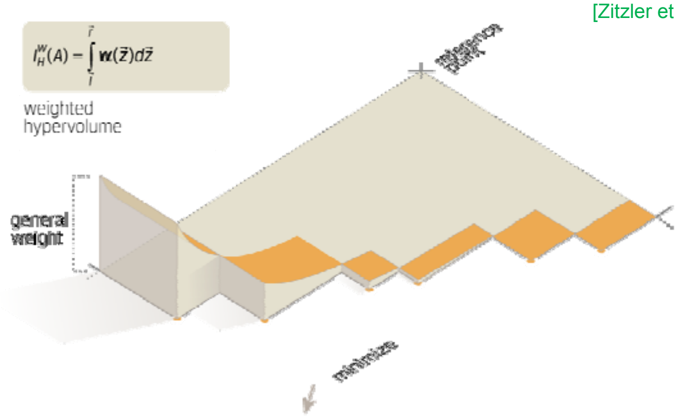

Weighted Hypervolume in Practice

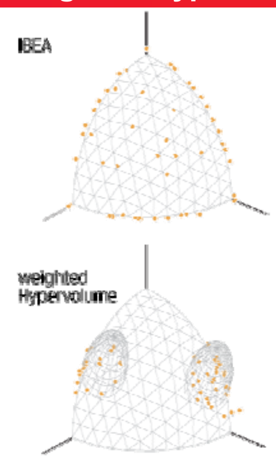

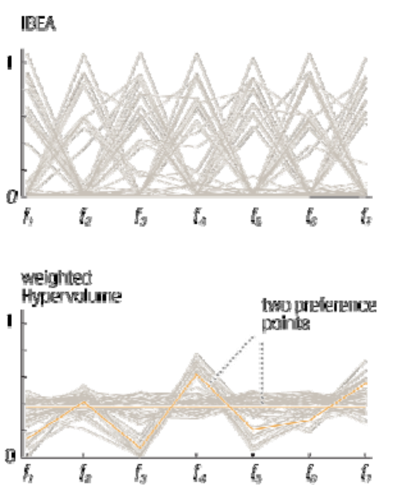

[Auger et al. 2009b]

Overview

The Big Picture

Basic Principles of Multiobjective Optimization

- algorithm design principles and concepts
- performance assessment

Selected Advanced Concepts

- indicator-based EMO
- preference articulation

A Few Examples From Practice

Application: Design Space Exploration

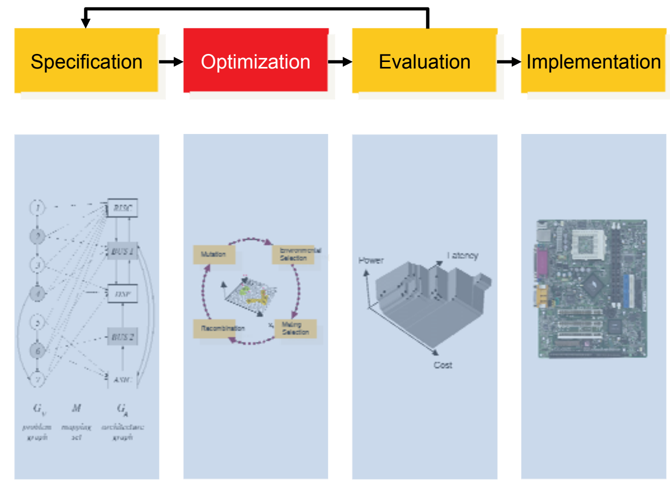

Application: Design Space Exploration

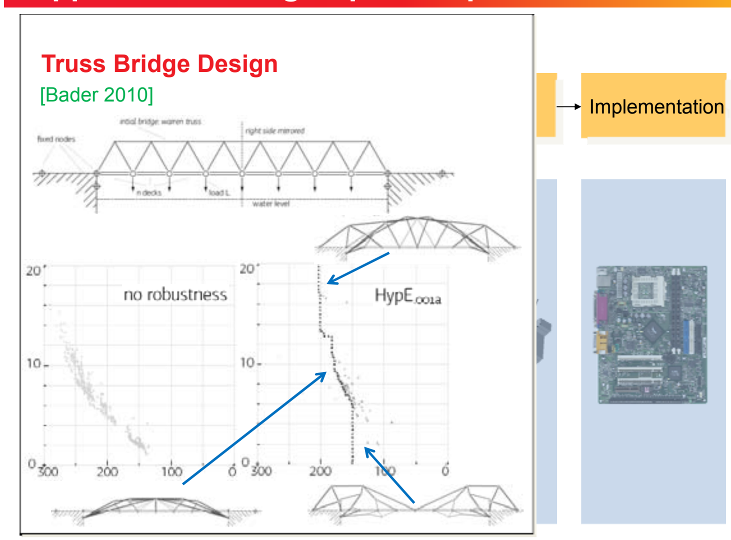

Application: Design Space Exploration

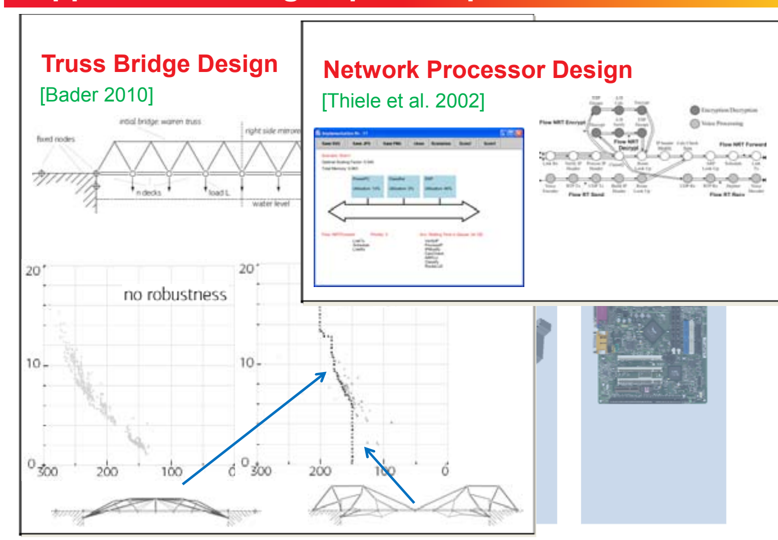

Application: Design Space Exploration

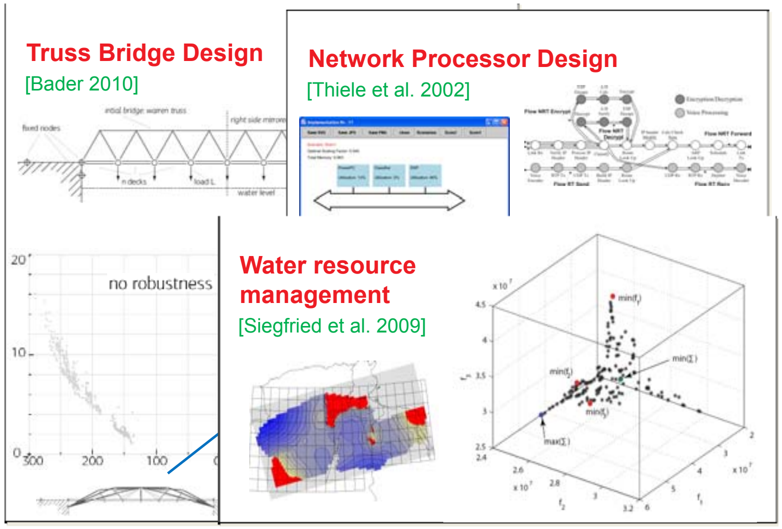

Application: Trade-Off Analysis

Module identification from biological data [Calonder et al. 2006]

Find group of genes wrt different data types:

- similarity of gene expression profiles
- overlap of protein interaction partners
- metabolic pathway map distances

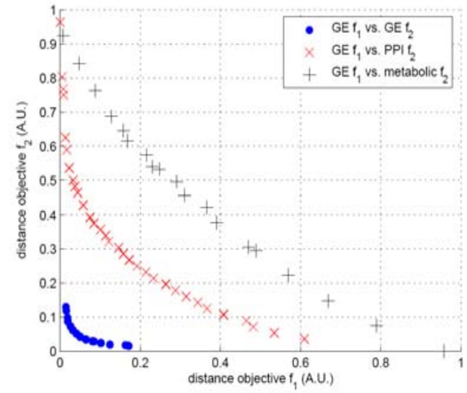

Conclusions: EMO as Interactive Decision Support

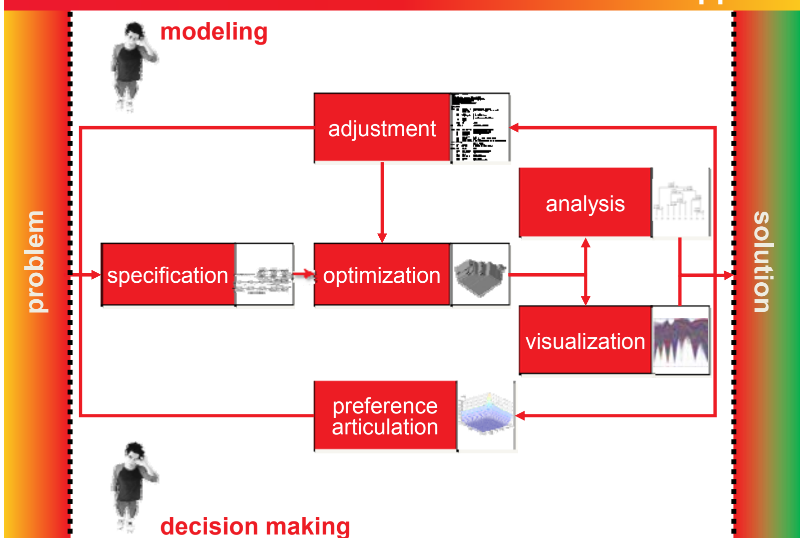

The EMO Community

Links:

- EMO mailing list: *https://lists.dei.uc.pt/mailman/listinfo/emo-list*
- MCDM mailing list: *http://lists.jyu.fi/mailman/listinfo/mcdm-discussion*
- EMO bibliography: *http://www.lania.mx/~ccoello/EMOO/*
- EMO conference series: *http://www.dep.uminho.pt/EMO2015/*

Books:

- ***Multi-Objective Optimization using Evolutionary Algorithms*** Kalyanmoy Deb, Wiley, 2001
- ***Evolutionary Algorithms for Solving Multi Evolutionary Algorithms for Solving Multi-Objective Problems Objective Problems***, Carlos A. Coello Coello, David A. Van Veldhuizen & Gary B. Lamont, Kluwer, 2nd Ed. 2007
- **Multiobjective Optimization—Interactive and Evolutionary Approaches**, J. Branke, K. Deb, K. Miettinen, and R. Slowinski, editors, volume 5252 of *LNCS*. Springer, 2008 [(still) many open questions!]
- and more…

PISA: http://www.tik.ee.ethz.ch/pisa/

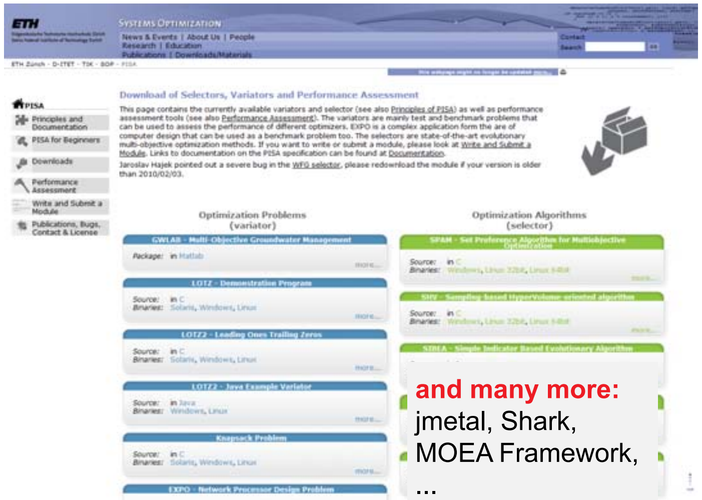

and many more:
jmetal, Shark,
MOEA Framework,
...

Perspectives

Challenging Open (Research) Directions

- Benchmarking
  - comparison with classical approaches
  - where are real strengths of EMO (how much better?)
  - algorithm recommendations for practice
- Many-objective Optimization
- growing EMO and MCDM to one field

Questions?

Additional Slides

Instructor Biography: Dimo Brockhoff

Dimo Brockhoff

INRIA Lille - Nord Europe
DOLPHIN team
Parc scientifique de la Haute Borne
40, avenue Halley - Bât B - Park Plaza
59650 Villeneuve d'Ascq
France

After obtaining his diploma in computer science (Dipl.-Inform.) from University of Dortmund, Germany in 2005, Dimo Brockhoff received his PhD (Dr. sc. ETH) from ETH Zurich, Switzerland in 2009. Between June 2009 and October 2011 he held postdoctoral research positions---first at INRIA Saclay Ile-de-France in Orsay and then at Ecole Polytechnique in Palaiseau, both in France. Since November 2011 he has been a junior researcher (now CR1) at INRIA Lille - Nord Europe in Villeneuve d'Ascq, France. His research interests are focused on evolutionary multiobjective optimization (EMO), in particular on many-objective optimization, benchmarking, and theoretical aspects of indicator-based search.

References

[Auger et al. 2009a] A. Auger, J. Bader, D. Brockhoff, and E. Zitzler. Theory of the Hypervolume Indicator: Optimal μ-Distributions and the Choice of the Reference Point. In Foundations of Genetic Algorithms (FOGA 2009), pages 87–102, New York, NY, USA, 2009. ACM.

[Auger et al. 2009b] A. Auger, J. Bader, D. Brockhoff, and E. Zitzler. Articulating User Preferences in Many-Objective Problems by Sampling the Weighted Hypervolume. In G. Raidl et al., editors, Genetic and Evolutionary Computation Conference (GECCO 2009), pages 555–562, New York, NY, USA, 2009. ACM

[Bader 2010] J. Bader. Hypervolume-Based Search For Multiobjective Optimization: Theory and Methods. PhD thesis, ETH Zurich, 2010

[Bader and Zitzler 2011] J. Bader and E. Zitzler. HypE: An Algorithm for Fast Hypervolume-Based Many-Objective Optimization. Evolutionary Computation 19(1):45-76, 2011.

[Bader et al. 2009] J. Bader, D. Brockhoff, S. Welten, and E. Zitzler. On Using Populations of Sets in Multiobjective Optimization. In M. Ehrgott et al., editors, Conference on Evolutionary Multi-Criterion Optimization (EMO 2009), volume 5467 of LNCS, pages 140–154. Springer, 2009

[Branke 2008] J. Branke. Consideration of Partial User Preferences in Evolutionary Multiobjective Optimization. In Multiobjective Optimization, volume 5252 of LNCS, pages 157-178. Springer, 2008

[Branke and Deb 2004] J. Branke and K. Deb. Integrating User Preferences into Evolutionary Multi-Objective Optimization. Technical Report 2004004, Indian Institute of Technology, Kanpur, India, 2004. Also published as book chapter in Y. Jin, editor: Knowledge Incorporation in Evolutionary Computation, pages 461–477, Springer, 2004

[Bringmann 2012] K. Bringmann. An improved algorithm for Klee's measure problem on fat boxes. Computational Geometry: Theory and Applications, 45:225–233, 2012.

[Bringmann 2013] K. Bringmann. Bringing Order to Special Cases of Klee's Measure Problem. arXiv preprint arXiv:1301.7154 (2013).

References

[Bringmann and Friedrich 2009] K. Bringmann and T. Friedrich. Approximating the Least Hypervolume Contributor: NP-hard in General, But Fast in Practice. In M. Ehrgott et al., editors, Conference on Evolutionary Multi-Criterion Optimization (EMO 2009),pages 6–20. Springer, 2009

[Brockhoff et al. 2009] D. Brockhoff, T. Friedrich, N. Hebbinghaus, C. Klein, F. Neumann, and E. Zitzler. On the Effects of Adding Objectives to Plateau Functions. IEEE Transactions on Evolutionary Computation, 13(3):591–603, 2009

[Calonder et al. 2006] M. Calonder, S. Bleuler, and E. Zitzler. Module Identification from Heterogeneous Biological Data Using Multiobjective Evolutionary Algorithms. In T. P. Runarsson et al., editors, Conference on Parallel Problem Solving from Nature (PPSN IX), volume 4193 of LNCS, pages 573–582. Springer, 2006

[Camerini et al. 1984] P. M. Camerini, G. Galbiati, and F. Maffioli. The complexity of multi-constrained spanning tree problems. In Theory of algorithms, Colloquium PECS 1984, pages 53-101, 1984.

[Coello Coello 2000] C. A. Coello Coello. Handling Preferences in Evolutionary Multiobjective Optimization: A Survey. In Congress on Evolutionary Computation (CEC 2000), pages 30–37. IEEE Press, 2000

[Deb and Kumar 2007] K. Deb and A. Kumar. Light Beam Search Based Multi-objective Optimization Using Evolutionary Algorithms. In Congress on Evolutionary Computation (CEC 2007), pages 2125–2132. IEEE Press, 2007

[Deb and Srinivasan 2006] K. Deb and A. Srinivasan. Innovization: Innovating Design Principles through Optimization. In Genetic and Evolutionary Computation Conference (GECCO 2006), pages 1629–1636. ACM, 2006

[Deb and Sundar 2006] K. Deb and J. Sundar. Reference Point Based Multi-Objective Optimization Using Evolutionary Algorithms. In Maarten Keijzer et al., editors, Conference on Genetic and Evolutionary Computation (GECCO 2006), pages 635–642. ACM Press, 2006

References

[Fonseca and Fleming 1998a] C. M. Fonseca and Peter J. Fleming. Multiobjective Optimization and Multiple Constraint Handling with Evolutionary Algorithms—Part I: A Unified Formulation. IEEE Transactions on Systems, Man, and Cybernetics, 28(1):26–37, 1998

[Fonseca and Fleming 1998b] C. M. Fonseca and Peter J. Fleming. Multiobjective Optimization and Multiple Constraint Handling with Evolutionary Algorithms—Part II: Application Example. IEEE Transactions on Systems, Man, and Cybernetics, 28(1):38–47, 1998

[Fonseca et al. 2011] C. M. Fonseca, A. P. Guerreiro, M. López-Ibáñez, and L. Paquete. On the computation of the empirical attainment function. In Conference on Evolutionary Multi-Criterion Optimization (EMO 2011). Volume 6576 of LNCS, pp. 106-120, Springer, 2011

[Friedrich et al. 2011] T. Friedrich, K. Bringmann, T. Voß, C. Igel. The Logarithmic Hypervolume Indicator. In Foundations of Genetic Algorithms (FOGA 2011). ACM, 2011. To appear.

[Greiner et al. 2007] D. Greiner, J. M. Emperador, G. Winter, and B. Galván. Improving Computational Mechanics Optimium Design Using Helper Objectives: An Application in Frame Bar Structures. In Conference on Evolutionary Multi-Criterion Optimization (EMO 2007), volume 4403 of LNCS, pages 575–589. Springer, 2007

[Handl et al. 2008a] J. Handl, S. C. Lovell, and J. Knowles. Investigations into the Effect of Multiobjectivization in Protein Structure Prediction. In G. Rudolph et al., editors, Conference on Parallel Problem Solving From Nature (PPSN X), volume 5199 of LNCS, pages 702–711. Springer, 2008

[Handl et al. 2008b] J. Handl, S. C. Lovell, and J. Knowles. Multiobjectivization by Decomposition of Scalar Cost Functions. In G. Rudolph et al., editors, Conference on Parallel Problem Solving From Nature (PPSN X), volume 5199 of LNCS, pages 31–40. Springer, 2008

[Igel et al. 2007] C. Igel, N. Hansen, and S. Roth. Covariance Matrix Adaptation for Multi-objective Optimization. Evolutionary Computation, 15(1):1–28, 2007

References

[Judt et al. 2011] L. Judt, O. Mersmann, and B. Naujoks. Non-monotonicity of obtained hypervolume in 1-greedy S-Metric Selection. In: Conference on Multiple Criteria Decision Making (MCDM 2011), abstract, 2011.

[Knowles et al. 2001] J. D. Knowles, R. A. Watson, and D. W. Corne. Reducing Local Optima in Single-Objective Problems by Multi-objectivization. In E. Zitzler et al., editors, Conference on Evolutionary Multi-Criterion Optimization (EMO 2001), volume 1993 of LNCS, pages 269–283, Berlin, 2001. Springer

[Jensen 2004] M. T. Jensen. Helper-Objectives: Using Multi-Objective Evolutionary Algorithms for Single-Objective Optimisation. Journal of Mathematical Modelling and Algorithms, 3(4):323–347, 2004. Online Date Wednesday, February 23, 2005

[Miettienen 1999] K. Miettinen. Nonlinear Multiobjective Optimization. Kluwer, Boston, MA, USA, 1999

[Neumann and Wegener 2006] F. Neumann and I. Wegener. Minimum Spanning Trees Made Easier Via Multi-Objective Optimization. Natural Computing, 5(3):305–319, 2006

[Obayashi and Sasaki 2003] S. Obayashi and D. Sasaki. Visualization and Data Mining of Pareto Solutions Using Self-Organizing Map. In Conference on Evolutionary Multi-Criterion Optimization (EMO 2003), volume 2632 of LNCS, pages 796–809. Springer, 2003

[Rachmawati and Srinivasan 2006] L. Rachmawati and D. Srinivasan. Preference Incorporation in Multi-objective Evolutionary Algorithms: A Survey. In Congress on Evolutionary Computation (CEC 2006), pages 962–968. IEEE Press, July 2006

[Schaffer 1985] J. D. Schaffer. Multiple Objective Optimization with Vector Evaluated Genetic Algorithms. In John J. Grefenstette, editor, Conference on Genetic Algorithms and Their Applications, pages 93–100, 1985.

[Serafini 1986] P. Serafini. Some considerations about computational complexity for multi objective combinatorial problems. In: Recent advances and historical development of vector optimization, number 294 in Lecture Notes in Economics and Mathematical Systems. Springer, 1986.

References

[Siegfried et al. 2009] T. Siegfried, S. Bleuler, M. Laumanns, E. Zitzler, and W. Kinzelbach. Multi-Objective Groundwater Management Using Evolutionary Algorithms. IEEE Transactions on Evolutionary Computation, 13(2):229–242, 2009

[Thiele et al. 2002] L. Thiele, S. Chakraborty, M. Gries, and S. Künzli. Design Space Exploration of Network Processor Architectures. In Network Processor Design 2002: Design Principles and Practices. Morgan Kaufmann, 2002

[Ulrich et al. 2007] T. Ulrich, D. Brockhoff, and E. Zitzler. Pattern Identification in Pareto-Set Approximations. In M. Keijzer et al., editors, Genetic and Evolutionary Computation Conference (GECCO 2008), pages 737–744. ACM, 2008.

[Verel et al. 2011] S. Verel, C. Dhaenens, A. Liefooghe. Set-based Multiobjective Fitness Landscapes: A Preliminary Study. In Genetic and Evolutionary Computation Conference (GECCO 2011). ACM, 2010. To appear.

[Voß et al. 2010] T. Voß, N. Hansen, and C. Igel. Improved Step Size Adaptation for the MO-CMA-ES. In J. Branke et al., editors, Genetic and Evolutionary Computation Conference (GECCO 2010), pages 487–494. ACM, 2010

[Yildiz and Suri 2012] H. Yildiz and S. Suri. On Klee's measure problem for grounded boxes. Proceedings of the 2012 symposuim on Computational Geometry. ACM, 2012.

[Watanabe and Sakakibara 2007] S. Watanabe and K. Sakakibara. A multiobjectivization approach for vehicle routing problems. In Conference on Evolutionary Multi-Criterion Optimization (EMO 2007), volume 4403 of LNCS, pages 660–672 Springer, 2007.

[Zhang and Li 2007] Q. Zhang and H. Li. MOEA/D: A Multiobjective Evolutionary Algorithm Based on Decomposition. IEEE Transactions on Evolutionary Computation, 11(6):712--731, 2007

[Zitzler 1999] E. Zitzler. Evolutionary Algorithms for Multiobjective Optimization: Methods and Applications. PhD thesis, ETH Zurich, Switzerland, 1999

References

[Zitzler and Künzli 2004] E. Zitzler and S. Künzli. Indicator-Based Selection in Multiobjective Search. In X. Yao et al., editors, Conference on Parallel Problem Solving from Nature (PPSN VIII), volume 3242 of LNCS, pages 832–842. Springer, 2004

[Zitzler et al. 2010] E. Zitzler, L. Thiele, and J. Bader. On Set-Based Multiobjective Optimization. IEEE Transactions on Evolutionary Computation, 14(1):58–79, 2010.

[Zitzler et al. 2003] E. Zitzler, L. Thiele, M. Laumanns, C. M. Fonseca, and V. Grunert da Fonseca. Performance Assessment of Multiobjective Optimizers: An Analysis and Review. IEEE Transactions on Evolutionary Computation, 7(2):117–132, 2003

[Zitzler et al. 2000] E. Zitzler, K. Deb, and L. Thiele. Comparison of Multiobjective Evolutionary Algorithms: Empirical Results. Evolutionary Computation, 8(2):173–195, 2000