BTG Workshop 7

Enhanced Hospital Resource Management using Anticipatory Policies in Online Dynamic Multi-objective Optimization

Anke K. Hutzschenreuter[1,2], Peter A.N. Bosman[2], Han La Poutré[1,2]

[1] Technische Universiteit Eindhoven – TU/e
(Technical University Eindhoven)
Eindhoven, The Netherlands

[2] Centrum Wiskunde & Informatica – CWI
(Centre for Mathematics and Computer Science)
Amsterdam, The Netherlands

Optimization problem

- Hospitals need to reduce costs and improve quality of service
- Complex relationship between resources, utilization and patient throughput for different patient groups
- Resource usage at hospital unit is stochastic:
  - Arriving emergency patients
  - Unexpected complications
  - Uncertain treatment duration
- Analytical evaluation of objectives not feasible
- Evaluation of objectives by simulation

Optimization problem (ctd.)

- Parameter space
  - Resources: operating rooms & beds (fully staffed)
- Objective space
  - When is resource allocation efficient?
  - Multiple measures are important:
    - Throughput: number of patients discharged
    - Costs: total cost of using rooms and beds, including staff
    - Backup-use: use of extra beds or temporary patient transfers to other units in case of resource shortage
- Optimization problem is thus multi-objective

Multi-objective optimization

- Multiple objectives should be optimized simultaneously
- Conflicting objectives, no expression of weights
- Can't combine the objectives in a single scalar objective
- Example:
  Maximize the quality
  and minimize the
  production costs
  of a product

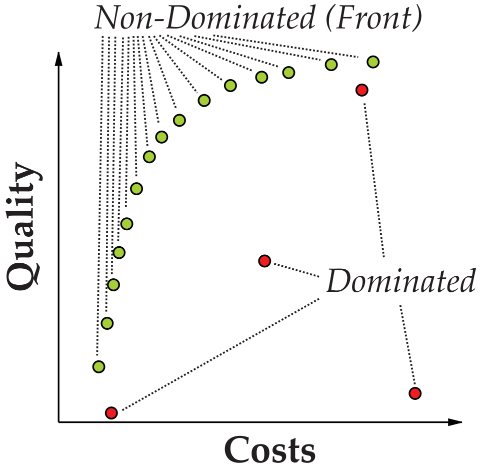

CWI

Dynamic multi-objective optimization

- Objectives change over time ($f_i(\mathbf{x})(t) \neq f_i(\mathbf{x})(t'), t \neq t'$).
- Goal is to optimize sum of objectives over time
- But can only take a single decision each timestep, i.e.:

$$\min_{\mathbf{x}(t)} \left\{ \left(\int_0^{t^{\text{end}}} f_0(\mathbf{x}(t), t)\, dt, \ldots, \int_0^{t^{\text{end}}} f_{m-1}(\mathbf{x}(t), t)\, dt \right) \right\}$$

Policy optimization

- Online human decision making infeasible
- Online dynamic optimization problem
- Offline approach:
  - Policy is a parameterized function
    - Input: the current situation
    - Output: setting for decision variables
  - Evaluate performance of a policy using simulation
  - Optimize policy parameters
  - Use optimized policy in real-life

Policy optimization - II

- Deals immediately with sub-optimality of sums of time-step-wise Pareto-optimal solutions

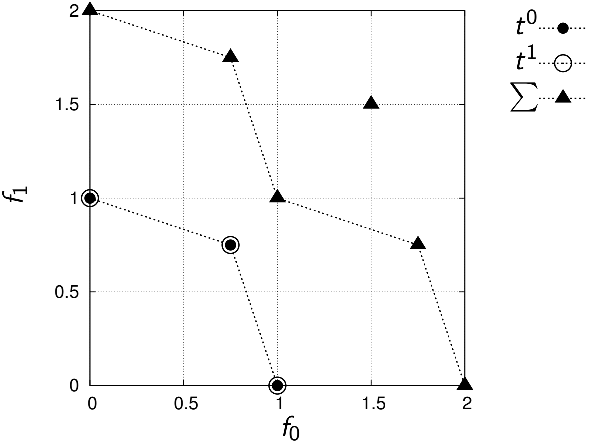

Anticipation

- Additional difficulty: decisions taken now may have future consequences (time-dependence)
- Decisions that seem suboptimal now may lead to situations where the reward is eventually better:

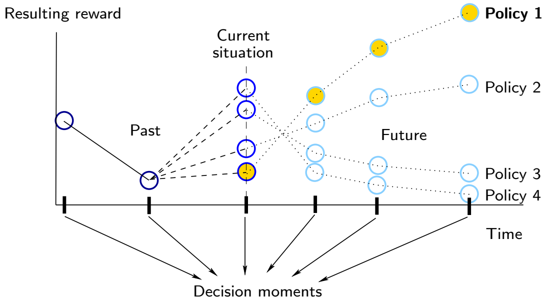

- Predict and optimize simultaneously
- How to combine with offline policy optimization?

Policy design

- Decision variables: #resources allocated to hospital units
- Decision moments: when resource allocations can be changed, e.g. once/day
- Compute utilization rate s of resources
- If s above threshold θ^{incr}, then increase allocated resources
- If s below threshold θ^{decr}, then deccrease allocated resources
- Otherwise, current allocation remains unchanged
- Policy parameters (per unit):
  - Default/starting resource allocation
  - Increase stepsize
  - Decrease stepsize
  - Increase threshold θ^{incr}
  - Decrease threshold θ^{decr}

Anticipatory policy design

- Policy looks at utilization rate
- Changes in beds made if utilization rate too low/too high
- Thresholds and change-size optimized using a MOEA
- Anticipation via utilization rate observation
  - look only at current utilization rate (no anticipation)
  - look at average over future 2 days (anticipation)
- Evaluate the future by forward simulation with current allocation

Multi-objective optimization algorithm (MOEA)

- Complex, non-analytic optimization problem (i.e. a simulation)
- Therefore employ metaheuristic algorithms
- Particularly: evolutionary algorithms (EAs)
- EAs well-suited for multi-objective optimization
  - Goal is to find a set
  - EAs use a set to perform optimization
- Typically reduces number of evaluations required
- We use specific type of EA:
  Estimation-of-Distribution Algorithm (EDA)
- EDAs discover and exploit interactions between problem variables using probability distributions
- Especially useful for multi-objective optimization:
  mixture distributions

Multi-objective mixture-based EDAs

- Power of mixture distributions: spatial separation.
- Assign each cluster same mixing coefficient.
- Distributes search power along front.

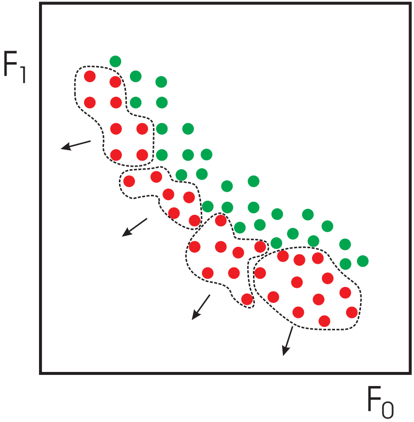

Experimental setup

- Can better resource allocations be found for CHE?
- Settings of the EDA (SDR-AVS-$\mathbb{M}$ID$\mathbb{E}$A)
  - 4 clusters
  - Normal distribution in each cluster
  - Real-valued representation: 35 problem variables (five per unit)
  - 1600 generations (no significant changes observed anymore)
- Settings of the simulation
  - Warming-up: simulate 8 weeks first
  - Compute average measures over 10 runs of 12 more weeks
- Perform evaluations in parallel on computer cluster
- Total time for single optimization run: 10h - 30h

Experimental results

- Pareto-fronts for non-anticipatory policy and current CHE practice
- F_0: Patient throughput (maximize)
- F_1: Resource costs (minimize)
- F_2: Backup-use (minimize)
- Without anticipation already outperforms current hospital practice

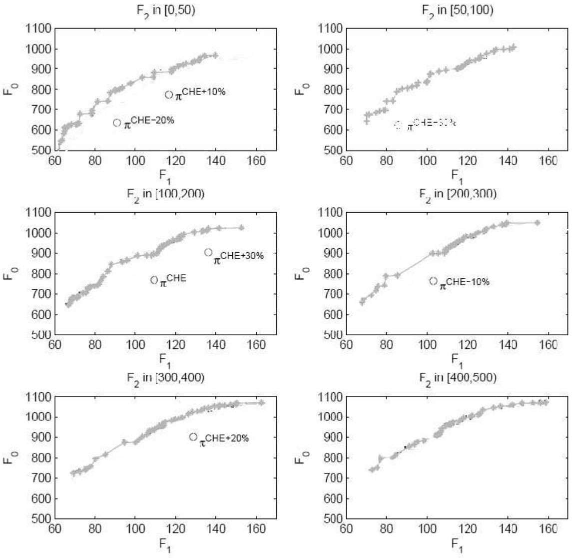

Experimental setup (ctd.)

- No significant difference with anticipatory policy
- Change simulation model:
  - increased availability of resource capacity
    $\longrightarrow$
    increased attraction of patients

Experimental results (ctd.)

- With anticipation, availability-capacity relation is exploited
- For higher resource cost, anticipatory policy improves throughput

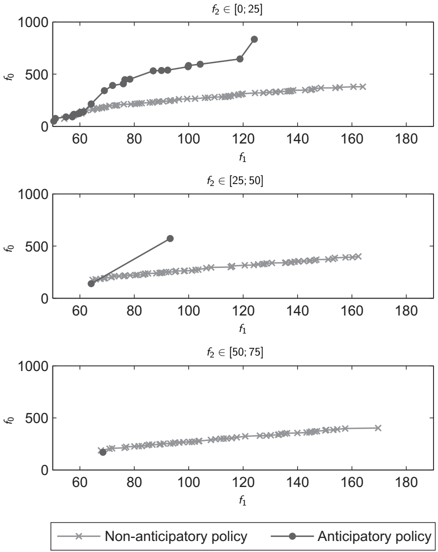

Experimental results (ctd.)

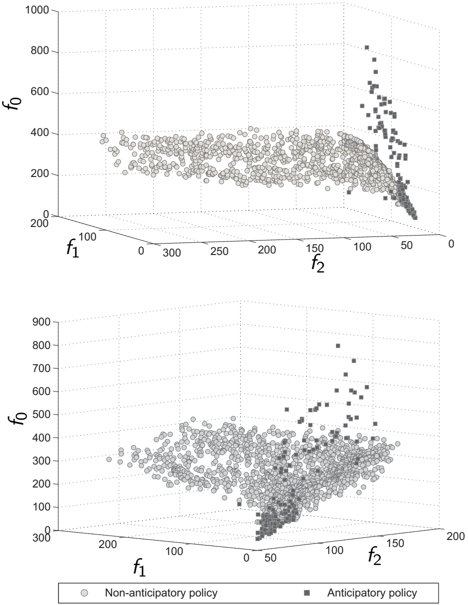

CWI

Conclusions

- Optimization for hospital resource management
- Online dynamic multi-objective optimization problem
- Offline approach: policy optimization
- Results outperform benchmarks (i.e. current practice)
- Future consequences of current decisions are important
- Should not only track optima, but also perform anticipation
- Optimize for current and future situations simultaneously
- Allows better results to be obtained in online dynamic optimization, even with multiple objectives