Evaluation heuristics for tug fleet optimisation algorithms

A computational simulation study of a receding horizon genetic algorithm

Robin T. Bye   Hans Georg Schaathun

Faculty of Engineering and Natural Sciences
Aalesund University College, Norway
Email: {roby,hasc}@hials.no
Web: www.robinbye.com

ICORES 2015
Lisbon, Portugal, 10–12 January 2015

Outline

Abstract

A fleet of tugs along the northern Norwegian coast must be dynamically positioned to minimise the risk of oil tanker drifting accidents. We have previously presented a receding horizon genetic algorithm (RHGA) for solving this tug fleet optimisation (TFO) problem. Here, we first present an overview of the TFO problem, the basics of the RHGA, and a set of potential cost functions with which the RHGA can be configured. The set of these RHGA configurations are effectively equivalent to a set of different TFO algorithms that each can be used for dynamic tug fleet positioning. In order to compare the merit of TFO algorithms that solve the TFO problem as defined here, we propose two evaluation heuristics and test them by means of a computational simulation study. Finally, we discuss our results and directions forward.

Ship traffic along the northern Norwegian coast

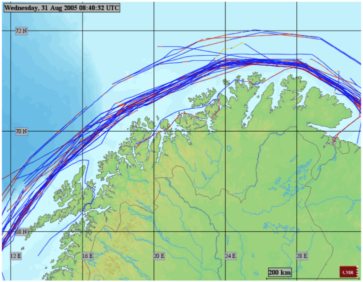

- Thousands of ships pass each year
- 2013: 186 drifting vessels, 29 groundings, 36 pollution incidents, 10 fires, 7 shipwrecks [1]
- Oil tankers high environmental risk
- Steering or propulsion failures → drift → grounding → oil spill

How to reduce the risk of drift grounding accidents?

Answer: Laws, regulations, tax, incentives, attitude campaigns, improved ship design, better nautical education, . . .

. . . and an actively patrolling tug fleet!

Norwegian Coastal Administration (NCA)

Administration of Vessel Traffic Service (VTS) centres, tug fleet, and much more

- VTS centres monitor all ship traffic in Norway
- Use sensory data fusion and integration technology, e.g.,
  - Automatic Identification System (AIS)
  - ship databases
  - electronic maps
  - present and predicted weather and ocean conditions
- VTS Vardø responsible for northern Norwegian region
  - commands a fleet of 3 patrolling tug vessels


Normand Jarl


AHTS Beta


NSO Crusader

How to position tugs such that risk is minimised?

Vardø (NOR) VTS and region of interest

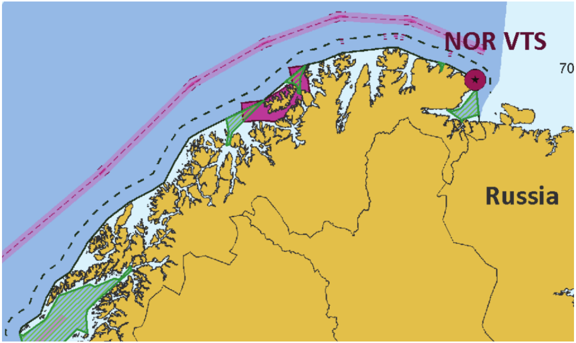

- Solid —: geographical baseline
- Stapled - - -: border of Norwegian territorial waters (NWS)
- Pink - - -: corridor for Traffic Separation Scheme (TSS)

The tug fleet optimisation (TFO) problem

1D problem description where tankers and tugs move along parallel lines

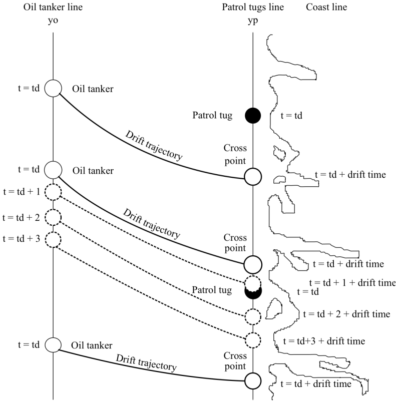

- Example scenario: 3 oil tankers (white) and 2 patrol tugs (black)
- Tankers may begin drift at some time from now into future
- Drift trajectories will intersect patrol line at crosspoints
- Fast drift times: 8–12 hours (typically much slower)
  - . . . but drift detection delay can be significant!

Where should tugs move to optimise some desired criterion?

A method for solving the TFO problem

A combination of optimisation, an intelligent algorithm, and modern control theory

1. Design cost function based on current and predicted data
   - Data can be positions and speeds of tugs and tankers, drift trajectories, crosspoints, ocean currents, fuel, time, etc.
   - Cost must be a function of future tug positions s.t. minimisation finds optimal tug trajectories
   - How to define the cost function?
2. Calculate future tug positions that minimise cost function
   - Genetic algorithm (GA): Fast, (sub)optimal solution
   - Mixed integer programming (MIP): Slow, optimal solution
3. Use receding horizon control (RHC) for closed-loop control
   - Feedback ensures adaptation to dynamic and uncertain environment
   - Plan for duration T_h into future (*how far?*)
   - Execute only first step of plan
   - Repeat and update plan

Earlier work and cost function design

Receding horizon algorithms using GA or MIP for minimisation of cost function

- Earlier work and absolute distance metric [2, 3, 4]:
  - cost is sum of the distances between all crosspoints and the *nearest* patrol point (position of tug) for all times from start of drift at time t_d and for a prediction horizon T_h ahead
  - equivalent to minimum rescue time if all tugs same max speed
- Recent work suggests other metrics [5]:
  - square of distances (penalise large distances more)
  - safe zone r (no cost inside)
  - detection delay δ and drift-from-alarm (DFA) time
  - number of unsalvageable tankers

Illustration of original cost function

Cost is accumulated for each crosspoint

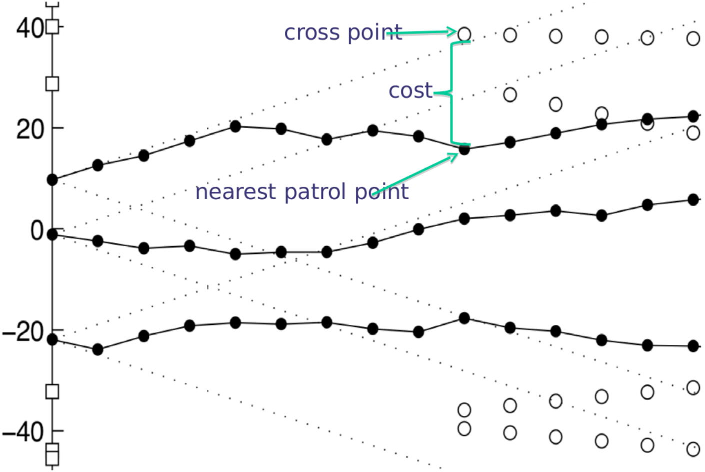

A flaw in the original cost function

Ignores drift alarm delay and assumes tugs continue executing plan despite alarm

- Inevitable detection delay δ from drift start at t_{d} until drift alarm at t_{a} ($\delta = t_{\mathrm{a}} - t_{\mathrm{d}} = 3$ hours, say)
- Define drift-from-alarm (DFA) time $\hat{\Delta}_{\mathrm{a}}$ as drift time from tugs receive alarm at t_{a} until crosspoint $\Rightarrow$ should replace entire drift time with shorter $\hat{\Delta}_{\mathrm{a}}$ for planning
- Original cost function assumes tugs continue original plan after alarm $\Rightarrow$ instead tugs should abandon plan and intercept drifting tanker

HØGSKOLEN I ÅLESUND
Aalesund University College

Illustration of DFA time and modified cost function

Rectification of flaw in original work

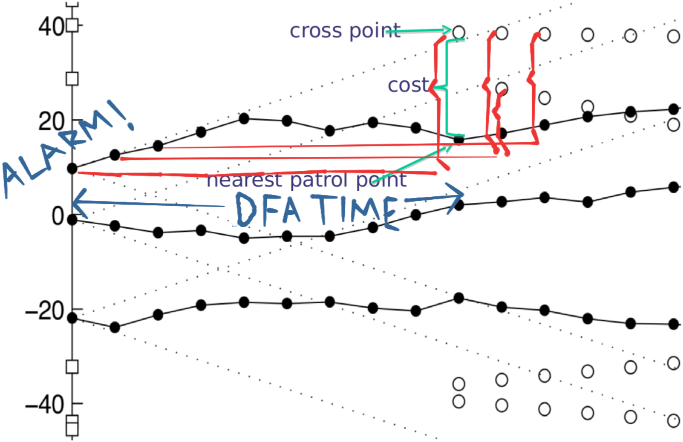

Cost functions in this study

Three cost functions f_1, f_2, and f_3 with parameters e and r

$$f_1(t) = \sum_{t=t_\mathrm{d}}^{t_\mathrm{d}+T_\mathrm{h}} \sum_{o\in O} \max\left\{0, \min_{p\in P} \left|y_t^c - y_t^p\right|^e - r\right\} \tag{1}$$

$$f_2(t) = \sum_{t=t_\mathrm{a}}^{t_\mathrm{a}+T_\mathrm{h}} \sum_{o\in O} \max\left\{0, \min_{p\in P} \left|y_{t+\hat{\Delta}_\mathrm{a}}^c - y_t^p\right|^e - r\right\} \tag{2}$$

$$f_3(t) = \sum_{t=t_\mathrm{a}}^{t_\mathrm{a}+T_\mathrm{h}} \sum_{o\in O} g\left(\min_{p\in P} \left|y_{t+\hat{\Delta}_\mathrm{a}}^c - y_t^p\right| - r\right) \quad \text{where} \tag{3}$$

$$g(x) = \begin{cases} 1, & x > 0 \quad (\text{outside } r) \\ 0, & x \leq 0 \quad (\text{inside } r) \end{cases} \tag{4}$$

Cost function configurations and static strategy

- Cost function configurations:
  - particular choices of e and r in f_1, f_2, f_3
  - $e \in \{1, 2\}$ and $r \in \{0, 50, 100\}$ km yields 14 configurations
- Static strategy:
  - add static "cost function" f_0
  - tugs stationary at base stations uniformly spread out
  - cheaper than actively patrolling coast
- Exact cost function optimisation such as MIP is slow $\Rightarrow$ use RHGA implemented with each cost function configuration
- RHGA + configuration $\equiv$ unique TFO algorithm

Table of RHGA configurations

Cost function f_i	Power e	Safe region r	#
0	0	0	1
1	1	0	2
		50	3
		100	4
	2	0	5
		50	6
		100	7
2	1	0	8
		50	9
		100	10
	2	0	11
		50	12
		100	13
3	0	50	14
		100	15

Evaluation and comparison of TFO algorithms

- Different cost functions are generally not directly comparable
- Cost functions may not reflect the “real” cost of the solution
- Many stochastic elements without known probability models
- Incorporation of these elements may cause too high complexity
- TFO algorithms generate tug fleet control solutions
- TFO algorithms may not use even use cost functions

How can we evaluate and compare the performance of different TFO algorithms?

Simulation framework

- Use Monte Carlo simulations
- Employ 2 pseudo-random algorithms:
  1. Generate problem to be solved by TFO algorithm
  2. Generate event (drifting tanker) where solution is tested
- TFO algorithms can plan solutions (tanker movements) based on *a priori* knowledge, e.g.,
  - Tanker positions, speeds, directions
  - Tug positions
  - Weather and ocean conditions now and in the future
  - Geographical concerns
- TFO algorithms don't know in advance which events will occur
- Evaluation heuristics can evaluate the cost of a particular TFO solution (tug trajectories) when some event occurs

Simulation model

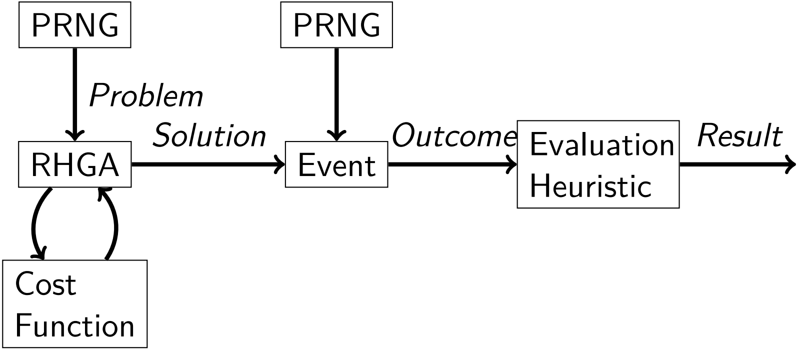

Steps of evaluation method

1. Randomly generate a deterministic and reproducible simulation scenario
2. Run the RHGA (or another TFO algorithm) for a given number of planning steps
3. Considering each oil tanker separately, assume each tanker begins drifting and count the number of salvageable tankers
4. For the same simulation scenario, repeat (2) and (3) with a different cost function configuration in the RHGA (or a different TFO algorithm)
5. Repeat steps (1)–(4) for a number of different simulation scenarios and find the accumulated evaluation cost for each RHGA configuration (or TFO algorithm)

Simulation scenarios

- Choose time period of interest, e.g., 24 hours
- Randomly generate oil tanker movements with corresponding drift trajectories and cross points
- Randomly create large number of scenarios offline
- Use scenarios as input data for testing TFO algorithms
- Testing must use some well-designed evaluation heuristic
- In future real-world application, use actual oil tanker movements and predicted drift trajectories and cross points, in real-time (prediction requires models of movement and drift)

Evaluation heuristic h_1

- Similar to f_3 counting unsalvageable tankers
- Assume each patrol tug p can save any ship with cross points inside safe region $r = v^p_{\mathrm{max}}\hat{\Delta}_{\mathrm{a}}$, where
  - $\hat{\Delta}_{\mathrm{a}}$ is the DFA time
  - v^p_{max} is the pth tug's maximum speed
- Safe region r is the maximal reach of a tug upon a drift alarm

$$h_1(t_{\mathrm{a}}) = \sum_{o \in O} g\left(\min_{p \in P}\left|y^c_{t_{\mathrm{a}}+\hat{\Delta}_{\mathrm{a}}} - y^p_{t_{\mathrm{a}}}\right| - r\right) \tag{5}$$

$$r = v^p_{\mathrm{max}}\hat{\Delta}_{\mathrm{a}} \tag{6}$$

$$g(x) = \begin{cases} 1, & x > 0 \quad (\text{outside } r) \\ 0, & x \leq 0 \quad (\text{inside } r) \end{cases} \tag{7}$$

Evaluation heuristic h_2

- Max tug speed may vary significantly depending on weather
- Hookup time not take into account in h_1
- Suggested changes:
  - Reduce safe region to area reachable for any tug with some minimum speed $v^p_{\min}$
  - Assume tugs always able to attain this speed in any weather
  - Squaring to punish larger distances more

$$h_2(t_\mathrm{a}) = \sum_{o \in O} \left(\max \left\{ 0, \min_{p \in P} \left| y^c_{t_\mathrm{a}+\hat{\Delta}_\mathrm{a}} - y^p_{t_\mathrm{a}} \right| - r \right\} \right)^2 \tag{8}$$

$$r = v^p_{\min} \hat{\Delta}_\mathrm{a}, \tag{9}$$

Other possible evaluation measures

Some examples of possible evaluation measures for cost functions (may also be weighted and combined):

- Total fuel consumption
- Continuous probabilities of not saving drifting tankers
- Estimated probabilistic financial cost of grounding accidents
- Various time measures, e.g.,
  - time to reach drifting tankers
  - time left before tankers will ground

Simulation parameters, settings, and units

Parameters	Settings	Units
Patrol zone (south-north line)	$Y = [-750, 750]$	km
Tanker zone (south-north line)	$Z = [-750, 750]$	km
Number of oil tankers	$N_{\mathrm{o}} = 6$	-
Set of oil tankers	$O = \{1, 2 \ldots, N_{\mathrm{o}}\}$	-
Number of tugs	$N_{\mathrm{p}} = \{1, \ldots, 6\}$	-
Set of tugs	$P = \{1, 2 \ldots, N_{\mathrm{p}}\}$	-
Initial tug positions (base stations)	Uniformly distributed	km
Random initial tanker positions	$y^{\mathrm{o}} \in Z, \forall o \in O$	km
Maximum speed of tugs	$v^p_{\mathrm{max}} = 20, \ \forall p \in P$	km/h
Minimum speed of tugs	$v^p_{\mathrm{min}} = 5, \ \forall p \in P$	km/h
Random speed of oil tankers	$v^o \in [20, 30], \ \forall o \in O$	km/h
Initial simulation time	$t_{\mathrm{i}} = 0$	h
Simulation step	$t_{\mathrm{s}} = 1$	h
Final simulation time	$t_{\mathrm{f}} = 24$	h
Prediction horizon	$T_{\mathrm{h}} = 24$	h
Time of start of drift	$t_{\mathrm{d}} \in \{t_{\mathrm{i}}, t_{\mathrm{i}} + 1, \ldots, t_{\mathrm{f}}\}$	h

HØGSKOLEN I ÅLESUND Aalesund University College

Simulation parameters, settings, and units (cont'd)

Parameters	Settings	Units
Detection delay	$\delta = 3$	h
Alarm time	$t_{\mathrm{a}} = t_{\mathrm{d}} + \delta \in \{t_{\mathrm{i}} + \delta, t_{\mathrm{i}} + \delta + 1, \ldots, t_{\mathrm{f}} + \delta\}$	h
Drift direction	Eastbound	-
Estimated drift times	$\hat{\Delta} \in \{8, 9, \ldots, 12\}$	h
Drift-from-alarm (DFA) times	$\hat{\Delta}_{\mathrm{a}} = \hat{\Delta} - \delta \in \{5, 6, \ldots, 9\}$	h
Static strategy	$y_t^p = y_{t_i}^p, \ \forall t$	km
Cost functions	$F = \{f_1, f_2, f_3\}$	-
Distance power	$e = \{1, 2\}, \quad \mathrm{in} \quad f_1, f_2$	-
Safe region	$r = \begin{cases} \{0, 50, 100\}, & \mathrm{in} \quad f_1, f_2 \\ \{50, 100\}, & \mathrm{in} \quad f_3 \\ v_{\max}^p \hat{\Delta}_{\mathrm{a}} = [100, 180], & \mathrm{in} \quad h_1 \\ v_{\min}^p \hat{\Delta}_{\mathrm{a}} = [25, 45], & \mathrm{in} \quad h_2 \end{cases}$	km
TFO algorithms	Configurations of RHGA($f_i, e, r, N_{\mathrm{p}}$)	-
Number of RHGA($f_i, e, r, N_{\mathrm{p}}$) configurations	$N_{\mathrm{conf}} = 15$	-
Number of scenarios	$N_{\mathsf{sc}} = 1600$	-
Total number of simulations	$N_{\mathsf{sim}} = N_{\mathsf{conf}} \times N_{\mathsf{sc}} \times \dim N_{\mathsf{p}} = 144,000$	-

HØGSKOLEN I ÅLESUND
Aalesund University College

Obtaining the results

- Randomly generated set of 1600 unique simulation scenarios
- Tested performance of 15 RHGA configurations for $N_\mathrm{p} = \{1, \ldots, 6\}$ tugs on the same set of scenarios
- Calculated h_1 and h_2 at end of each simulation
- Grand total of 140,000 simulations
- Found statistics for each configuration and number of tugs:
  - sample mean $\bar{h}_1$
  - standard deviation
  - coefficient of variance (relative standard deviation)
  - standard error (standard deviation of the sample mean)
  - relative standard error
- Focus on sample mean of active schemes and compare with static strategy as a low performance benchmark

Results of h_1

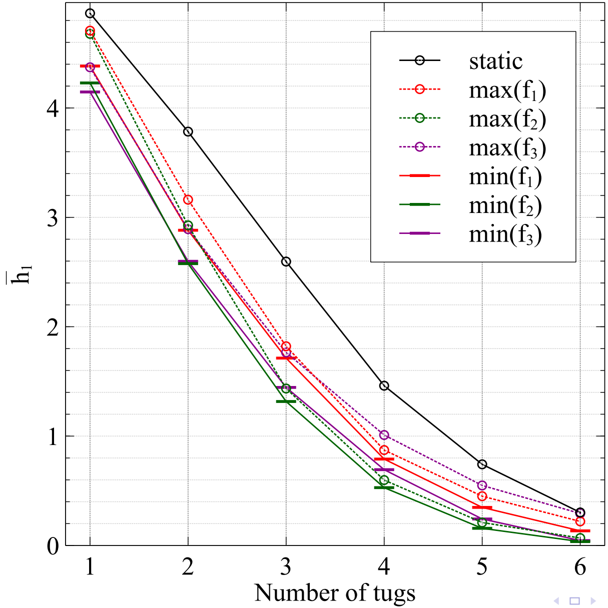
static
max(f1)
max(f2)
max(f3)
min(f1)
min(f2)
min(f3)
h1
0
1
2
3
4
1
2
3
4
5
6
Number of tugs

Selected results from h_1 evaluation

- h_1 is a measure of number of unsalvageable tankers
- Size of tug fleet strongly affects h_1
- RHGA configured with f_1–f_3 outperforms static strategy
- Except for using a single tug, the best f_2 and f_3 configurations outperform the static strategy with one less tug in fleet
- Best cost functions (smallest min/max) for number of tugs:
  - 1 tug: f_3
  - 2 tugs: f_2 and f_3
  - 3–6 tugs: f_2
- f_1 similar to f_2 but consistently worse for all tug fleet sizes
- Very small standard error (0.005 to 0.032)

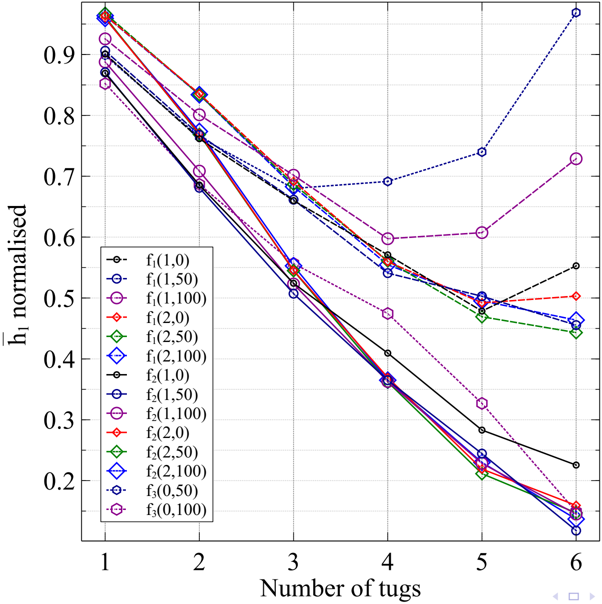

$\bar{h}_1$ normalised
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
1
2
3
4
5
6
Number of tugs
$f_1(1,0)$
$f_1(1,50)$
$f_1(1,100)$
$f_1(2,0)$
$f_1(2,50)$
$f_1(2,100)$
$f_2(1,0)$
$f_2(1,50)$
$f_2(1,100)$
$f_2(2,0)$
$f_2(2,50)$
$f_2(2,100)$
$f_3(0,50)$
$f_3(0,100)$

Observations when $\bar{h}_1$ normalised by static strategy

- Best settings for e and r for f_1 and f_2
  - $e = 1$ for 1–3 tugs
  - $e = 2$ for 4–6 tugs
  - $r = 50$ km is best overall
- Best settings for r for f_3
  - $r = 50$ km performs badly
  - $r = 100$ km performs well
  - difference is particulary big with many tugs in fleet
- Only f_2 (all configurations) and $f_3(r = 100)$ improves monotonically vs static strategy with increasing number of tugs

Results of h_2

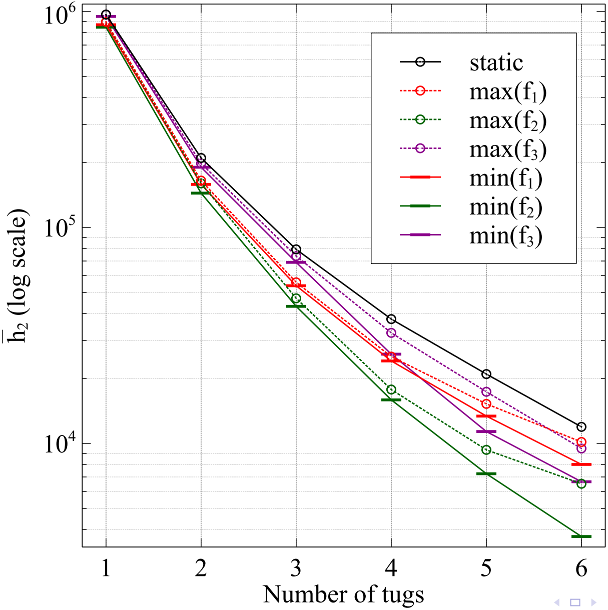
10^6
10^5
10^4
$\overline{h}_2$ (log scale)
1
2
3
4
5
6
Number of tugs
static
max(f_1)
max(f_2)
max(f_3)
min(f_1)
min(f_2)
min(f_3)

Selected results from h_2 evaluation

- h_2 is a measure of the sum of squared distances to cross points of unsalvageable tankers
- Use log scale due to orders of magnitude difference due to squaring
- Size of tug fleet strongly affects h_2
- Results similar to those of h_1 ...
- ... except f_3 performs worse with 1–4 tugs
- Very small standard error (typically about 1% of mean)

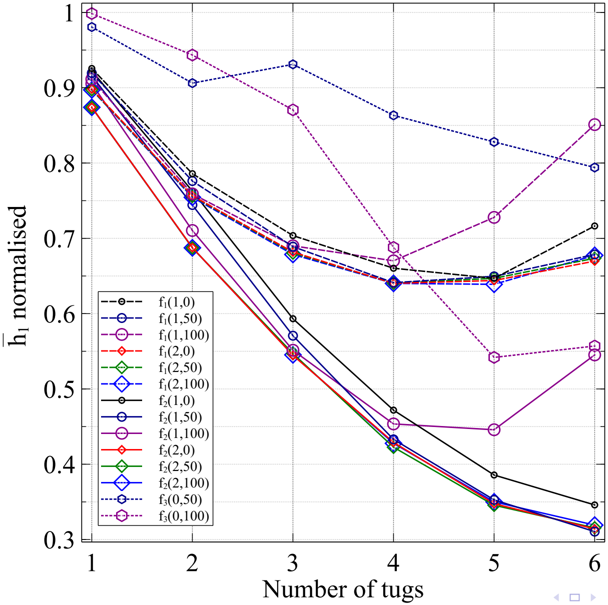

$\overline{h}_1$ normalised
Number of tugs
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
1 2 3 4 5 6
$f_1(1,0)$
$f_1(1,50)$
$f_1(1,100)$
$f_1(2,0)$
$f_1(2,50)$
$f_1(2,100)$
$f_2(1,0)$
$f_2(1,50)$
$f_2(1,100)$
$f_2(2,0)$
$f_2(2,50)$
$f_2(2,100)$
$f_3(0,50)$
$f_3(0,100)$

Observations when $\bar{h}_2$ normalised by static strategy

- Best settings for e and r for f_1 and f_2 was $e = 2$ for any number of tugs if $r = 0$ or $r = 100$
- Best settings for r for f_3
    - $r = 50$ km best for 1–2 tugs
    - $r = 100$ km best for 3–6 tugs
- Only f_2 (all configurations except $e = 1, r = 100$) improves monotonically vs static strategy with increasing number of tugs

Conclusions

- Both evaluation heuristics are able to quantify the performance of TFO algorithms designed to solve the TFO problem as defined here
- Small standard error means that the uncertainty in the means of h_1 and h_2 is small
- Static strategy more viable with increasing number of tugs ...
- ... yet f_2 increases its performance relative to the static strategy with number of tugs
- NCA not likely to use more than 2–3 tugs $\Rightarrow$ RHGA configured with f_2, $r = 50$, $e = 1$ (for h_1) or $e = 2$ (for h_2) is best choice
- f_1 should not be used (probably due to previously identified flaw)

Current and future work

- Test and verify RHGA in real-world systems with realistic conditions:
  - historical data of oil tanker traffic
  - realistic estimates of variable maximum tug speeds attainable under various conditions
  - realistic modelling of drift trajectories and cross points
  - downtime of tugs due to secondary missions or change of crew
- 2D modelling and fleet control
- Probabilistic modelling
- Real traffic and weather data (historic and real-time)
- Develop software prototype for NCA operators
- PhD project using MIP and examining these issues is well on its way with journal paper soon to be submitted [6]

Acknowledgements

The DRAMA research group is grateful for the support provided by Regionalt Forskningsfond Midt-Norge and the Research Council of Norway through the project *Dynamic Resource Allocation with Maritime Application* (DRAMA), grant no. ES504913.

References

[1] Vardø VTS. Annual report on petroleum transport statistics. Technical report, Norwegian Coastal Administration, March 2014.

[2] Robin T. Bye, Siebe B. van Albada, and Harald Yndestad. A receding horizon genetic algorithm for dynamic multi-target assignment and tracking: A case study on the optimal positioning of tug vessels along the northern Norwegian coast. In *Proceedings of the International Conference on Evolutionary Computation (ICEC 2010) — part of the International Joint Conference on Computational Intelligence (IJCCI 2010)*, pages 114–125, 2010.

[3] Robin T. Bye. A receding horizon genetic algorithm for dynamic resource allocation: A case study on optimal positioning of tugs. *Series: Studies in Computational Intelligence*, 399:131–147, 2012. Springer-Verlag: Berlin Heidelberg.

[4] Brice Assimizele, Johan Oppen, and Robin T. Bye. A sustainable model for optimal dynamic allocation of patrol tugs to oil tankers. In *Proceedings of the 27th European Conference on Modelling and Simulation*, pages 801–807, 2013.

[5] Robin T. Bye and Hans G. Schaathun. An improved receding horizon genetic algorithm for the tug fleet optimisation problem. In *Proceedings of the 28th European Conference on Modelling and Simulation*, pages 682–690, 2014.

[6] Brice Assimizele, Johannes O. Royset, Robin T. Bye, and Johan Oppen. Preventing Environmental Disasters from Grounding Accidents: A Case Study of Tugboat Positioning along the Norwegian Coast. *Manuscript soon to be submitted*, 2015.

To download this presentation and this paper, please visit the Publications page at www.robinbye.com.

Q & A