

Real Evolutionary Scheduling: some observations and examples

Peter Ross
P.Ross@napier.ac.uk

1

What this talk is about

Lots of published academic papers on scheduling...
... but most of them are from another planet.

Some experiences from 'real' life whatever that is:

- what do the users *really* want?
- does the software *really* work?
- who is *really* going to use the software?
- what's the *real* problem?

2

What do users really want?

The problem:

- two east-coast factories
- chicken farms all over Scotland
- catching squads go from farm to farm
- lorries go from factory to farm & back
- varied weekly requirements



3

What do users really want?

Real problem (we thought):

One (skilled) man was solving the weekly problem - on a whiteboard

Bad news: holidays, looming retirement

Wanted: optimal results

Real problem (in fact):

Wanted: solutions just like his!

Outcome: two coupled GAs
- one chunks the work
- other schedules vehicles

.. and solutions just like his

4

Chicken-catching

Successful thanks to spending a lot of time modelling what the human scheduler did.

His heuristics \Rightarrow penalties for various features in task decomposition and in lorry scheduling

Fortunately not too sensitive to penalty weights

5

Does the software really work?

EAs are stochastic – therefore, acceptance testing is tricky. Faulty implementation or faulty concept can still run smoothly and improve quality.

Example: exam timetabling:

- students take differing sets of subjects
- exams must not clash for any student
- exam hall capacities constrain choices
- many idiosyncratic constraints

6

Does it really work? continued

GATT: the representation used:

Exam	1	2	3	4	5	...	n
Slot	23	7	19	7	28	...	3

Used by lots of universities/departments

... but only works for lightly-constrained problems

It seems most users were grateful for any answer!

7

Who will really use the software?

Tatties II (Paechter et al): timetabling all of Napier University

Eg: 2000+ events, 45 weekly timeslots, 183 rooms, 700+ staff, 1000+ student groups

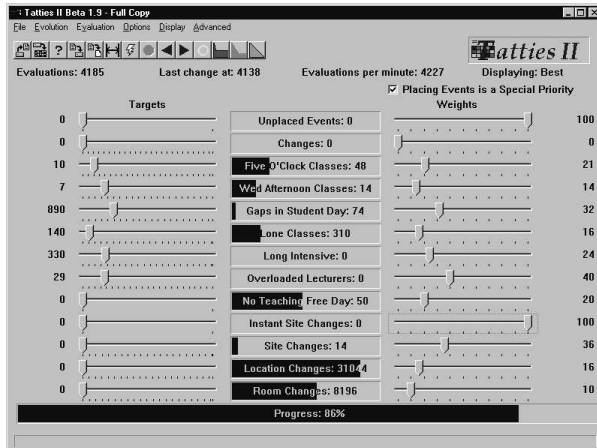
Twelve objectives: interactively set/adjust targets/weights for each.

Representation: permutation of events, plus suggested slots for each

No hard-constraint violations, but maybe unplaced events

8

Who will really use...?



9

Who will really use..?

Very successful: generated web timetables for all staff, students, rooms

However: running delegated to (changing) admin staff
Frequent training sessions, but use is not fire-and-forget

Replaced in 2003 by a standard commercial package..

10

What's the real problem?

A 'blending'-type problem (commercial, but anonymous)

The task: plan a sequence of combinations of ingredients

Turned into a kind of optimisation problem for a single step

11

What's the real..?

The maths:

$$\begin{aligned}
 &\text{Maximise} && \sum c_i x_i \\
 &\text{subject to} && \sum \sum a_{ij} x_i x_j \leq 0.06 \\
 &&& \sum_{i \in A} x_i = 0.5 \\
 &&& \sum_{i \in B} x_i = 0.5 \\
 &&& \dots \text{etc} \dots
 \end{aligned}$$

Looked like a straightforward MIP problem

12

What's the real..?

Non-EA methods could find very high-quality answers

The real problem:

thousands of variables

non-EA methods: 1-2 hours

good answers needed in *1-2 minutes*

The EA: an evolution strategy, plus analysis of matrix $\{a_{ij}\}$ to find tractable sub-problems

13

Other worries about EAs

- all those choices, parameters... needs skill
- no guarantees; when to stop?
- *black-box*: where did that answer come from?
- continuity: change the problem slightly \Rightarrow changes the solution a lot
- ongoing maintenance? consultancy?

14


Countering such worries

Rather than using an EA to solve the problem, use an EA to find an algorithm tailored for the family of problems:

- good worst-case behaviour
- familiar/trustworthy ingredients
- non-stochastic
- fast

15

Hyper-heuristics



- eg: applied to big class and exam timetabling
- generated algorithm finds good, feasible solutions
- ..without any search or backtracking

16

In conclusion...

EAs can work very well

Understanding the real nature of the problem is often the main hurdle

Textbooks and academic papers often *don't* tell it like it is