

- [Cart86] M. W. Carter. A survey of practical applications of examination timetabling. *Operations Research*, 34:193–202, 1986.
- [Cart94] M. W. Carter, G. Laporte, and J.W.Chinneck. A general examination scheduling system. *Interfaces*, 11:109–120, 1994.
- [Cart95] M. W. Carter, G. Laporte, and S. Y. Lee. Examination timetabling: Algorithmic strategies and applications. Working Paper 94-03, University of Toronto Dept of Industrial Engineering, January 1995. to appear in *Journal of Operational Research Society*.
- [Cart96] M. W. Carter and G. Laporte. Recent developments in practical examination timetabling. In Edmund Burke and Peter Ross, editors, *The Practice and Theory of Automated Timetabling: Selected Papers from the 1st International Conference, Lecture Notes in Computer Science 1153*, pages 3–21. Springer-Verlag, Berlin, 1996.
- [Corn94] D. Corne, P. Ross, and H. Fang. Fast practical evolutionary timetabling. In T.C.Fogarty, editor, *Lecture Notes in Computer Science 865*, pages 250–263. Springer-Verlag, Berlin, 1994.
- [Davi91] L. Davis. *Handbook of Genetic Algorithms*. Van Nostrand Reinhold, 1991.
- [Dawk76] R. Dawkins. *The Selfish Gene*. Oxford University Press, 1976.
- [Foxl68] E. Foxley and K. Lockyer. The construction of examination timetables by computer. *The Computer Journal*, 11:264–268, 1968.
- [Hert91] A. Hertz. Tabu search for large scale timetabling problems. *European Journal of Operations Research*, 54:39–47, 1991.
- [Mosc91] P. Moscato and M. G. Norman. A ‘memetic’ approach for the travelling salesman problem - implementation of computational ecology for combinatorial optimisation on message-passing systems. In *Proceedings of the International Conference on Parallel computing and Transputer Applications*. IOS Press (Amsterdam), 1991.
- [Paec95] Ben Paechter, Andrew Cumming, and Henri Luchian. The use of local search suggestion lists for improving the solution of timetable problems with evolutionary algorithms. In T. C. Fogarty, editor, *Lecture Notes in Computer Science 993 (AISB Workshop on Evolutionary Computing)*, pages 86–93. Springer-Verlag, Berlin, 1995.
- [Ross94] Peter Ross, Dave Corne, and Hsiao-Lan Fang. Improving evolutionary timetabling with delta evaluation and directed mutation. In Y. Davidor, H-P. Schwefel, and R. Manner, editors, *Parallel Problem Solving in Nature*, volume III. Springer-Verlag, Berlin, 1994.
- [Thom96] J. M. Thompson and K. A. Dowsland. General cooling schedules for a simulated annealing based timetabling system. In Peter Ross Edmund Burke, editor, *The Practice and Theory of Automated Timetabling: Selected Papers from the 1st International Conference, Lecture Notes in Computer Science 1153*, pages 345–364. Springer-Verlag, Berlin, 1996.
- [Wear95] R. F. Weare. *Automated Examination Timetabling*. PhD dissertation, University of Nottingham, Department of Computer Science, 1995.

In attempting to determine the best subset size to use with this method, there appears to be an optimal size that is perhaps related to the number of events in the problem, though there may be other considerations. For instance when using saturation degree the smaller **kfu** and **carf92** problems seem to work best with the smaller subset size of 50 with a look ahead set, the larger **nott** data set produces roughly equivalent results at sizes of 50 and 100, both with look ahead sets, while a subset size of 100 with a look ahead set produces the best results on the **pur** data set. As mentioned earlier we might have expected the optimal subset size for the **pur** data set to be higher due to the huge nature of the problem. This could be due to the evolutionary algorithm perhaps having an upper limit on the amount of data it can efficiently handle (which could perhaps be overcome with larger population sizes, though this would inevitably lead to even higher run times), or maybe because of the low density of the conflict matrix. Throughout the trials however, the use of a look ahead set shows obvious benefits. Although there is substantial extra time involved in doing this, it is not quite so important when we use subset sizes that are relatively small such as those which seem to produce the best results.

Even with a good choice of heuristic and an optimal subset size and look ahead set it is possible that on some more difficult problems placements made earlier in the process could well prevent some events being scheduled in earlier periods. To prevent this from happening we could borrow another idea from heuristic sequencing methods and allow a backtracking operator to be applied at the end of a phase if there are any events still unscheduled. Such an operator would have to be allowed to move events that have previously been fixed in order to be effective. This would inevitably mean that the quality of the timetable is lowered to some degree but should mean that the process will have as much chance of finding a feasible timetable as any other method.

This approach could possibly be adapted to other scheduling problems, especially those where some sort of heuristic sequencing approach exists. Early work is being carried out in the Automated Scheduling and Planning group at the University of Nottingham into employing a similar approach for power maintenance scheduling problems.

References

- [Burk95] E. K. Burke, D. G. Elliman, and R. F. Weare. A hybrid genetic algorithm for highly constrained timetabling problems. In Larry J. Eshelman, editor, *Genetic Algorithms: Proceedings of the 6th International Conference*, pages 605–610, San Francisco, 1995. Morgan Kaufmann.
- [Burk96a] E. K. Burke, D. G. Elliman, P. H. Ford, and R. F. Weare. Examination timetabling in british universities - a survey. In Edmund Burke and Peter Ross, editors, *The Practice and Theory of Automated Timetabling: Selected Papers from the 1st International Conference, Lecture Notes in Computer Science 1153*, pages 76–90. Springer-Verlag, Berlin, 1996.
- [Burk96b] E. K. Burke, J. P. Newall, and R. F. Weare. A memetic algorithm for university exam timetabling. In Edmund Burke and Peter Ross, editors, *The Practice and Theory of Automated Timetabling: Selected Papers from the 1st International Conference, Lecture Notes in Computer Science 1153*, pages 241–250. Springer-Verlag, Berlin, 1996.
- [Burk97] E. K. Burke and J. P. Newall. Investigating the benefits of utilising problem specific heuristics within a memetic timetabling algorithm. Working Paper NOTTCS-TR-97-6, University of Nottingham, UK, Dept of Computer Science, 1997.
- [Cart83] M. W. Carter. A decomposition algorithm for practical timetabling problems. Working Paper 83-06, Industrial Engineering, University of Toronto, April 1983.

Considering the results when using colour degree we see little difference in the individual results when compared to largest degree apart from a few minor variations in solution quality. However if we instead compare with saturation degree we can discern a notable improvement. While there is no substantial variation for the **kfu** and **carf92** problem we can see that the approach works completely differently on the **nott** data set where the employment of saturation degree results in the algorithm consistently finding feasible solutions even with very small subset sizes. A probable reason for this is that the **nott** data set has variable length exams and periods resulting in fewer available timeslots for longer exams. This is something that neither colour degree or largest degree consider but saturation degree does. For this problem, subset sizes of 50 and 100, both with look ahead sets, produce the best quality solutions. While the general trend for the **pur** problem remains roughly the same there is a noticeable increase in quality for smaller subset sizes, indicating that saturation degree is also a much better heuristic for this problem.

Taking all the results into account it would seem that least saturation degree first is the most reliable heuristic to use for this method. Though other choices of heuristic may be more suitable for other problems it is notable that least saturation degree first outperforms the other heuristics on the **nott** problem, which is more constrained in terms of seats and less dense in terms of its conflict matrix, while still achieving equivalent performance on the more densely conflicting problems. In terms of which subset size to use a size of 50 seems to be appropriate for smaller problems and a size of 100 for larger problems. In both cases a look ahead set should be used. Using these sizes, and given a suitable heuristic, it is possible to produce substantially better results in a fraction of the time than if the memetic algorithm were applied.

Data	Penalty
carf92	2915
kfu	2700
nott	918
pur	97521

Table 5. Results when using heuristic backtracking

For the purposes of comparison with a more established method Table 5 shows the results when using a heuristic backtracking method similar to that described by Carter[Cart95] which is well established and has been implemented at several universities. These results represent the lowest penalty found for these data sets by a published algorithm. The heuristic used here was least saturation degree first, which was found to be the most effective heuristic in [Cart95]. This also provides the best comparison as saturation degree certainly appears to be the most effective heuristic for our method. This method requires very little run time, even with the large **pur** problem requiring no longer than a few minutes. Comparing the results in Table 5 with our results when using saturation degree we see a quite uniform reduction in penalty of roughly 40%.

5 Conclusions

We have seen that by applying an algorithm to a problem in phases we can not only drastically reduce the amount of time taken to find that solution, relative to the time taken to apply an algorithm to the entire problem, but also considerably improve the quality of that solution. In essence this method is a hybrid of heuristic sequencing and evolutionary methods, which (as we have shown) outperforms either method on its own. This hybrid improves on its components by utilising knowledge of the problem to order events according to expected difficulty but instead of taking the single most difficult event and perhaps placing it in the best available period, we take a number of the most difficult events and apply the EA to find the best placements with respect to each other, as well as with those events already scheduled.




Figure 9. Results when run on **nott** without a look ahead set

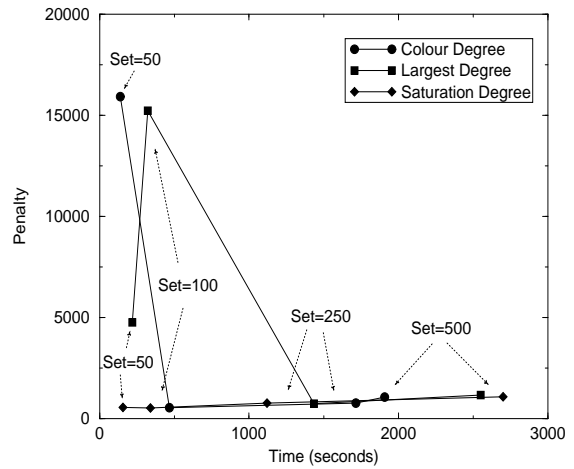


Figure 10. Results when run on **nott** with a look ahead set




Figure 11. Results when run on **pur** without a look ahead set

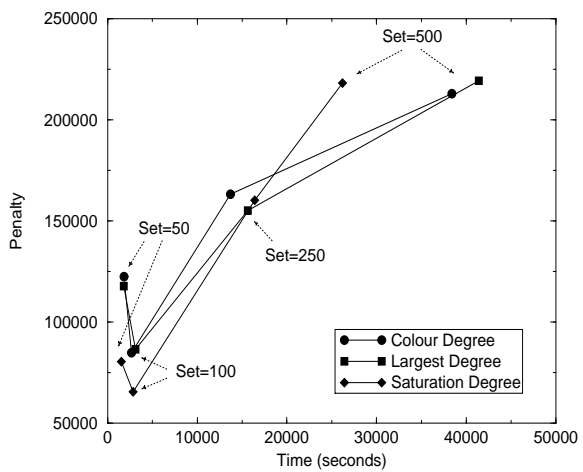


Figure 12. Results when run on **pur** with a look ahead set

tic. When considering the results for the **pur** problem we must take into account that “two in one [first order] conflicts are unavoidable”[Cart95], by this the authors mean that some conflicting exams will have to be scheduled in the same period. This situation in practice involves quarantining a student until they can sit an exam after the main sitting. Therefore we must reduce the amount of infeasibility in the timetable rather than find fully feasible timetables. Bearing this in mind it is clear that a larger subset size of 100 with a look ahead set finds the best solution. We might have expected this subset size to be larger given the huge number of events but this does of course depend on the structure of the problem.

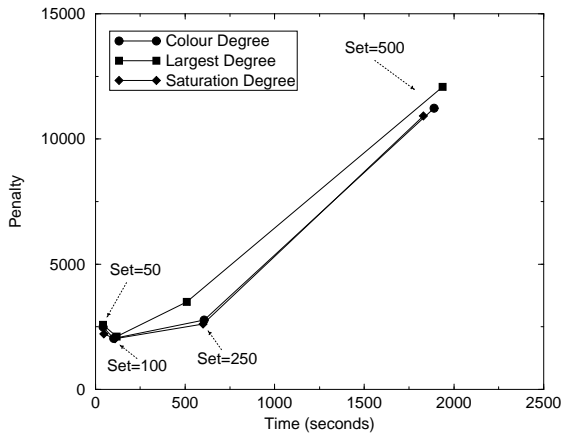


Figure 5. Results when run on **carf92** without a look ahead set

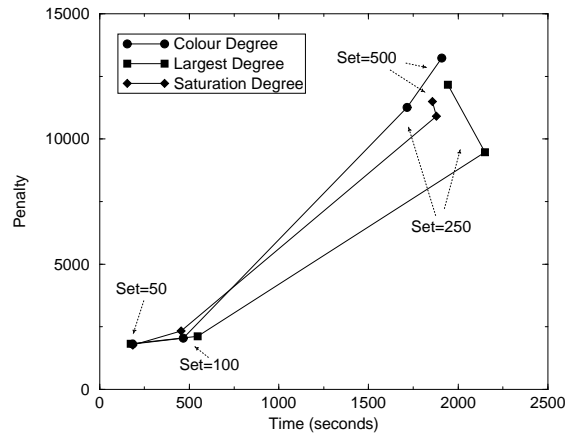


Figure 6. Results when run on **carf92** with a look ahead set

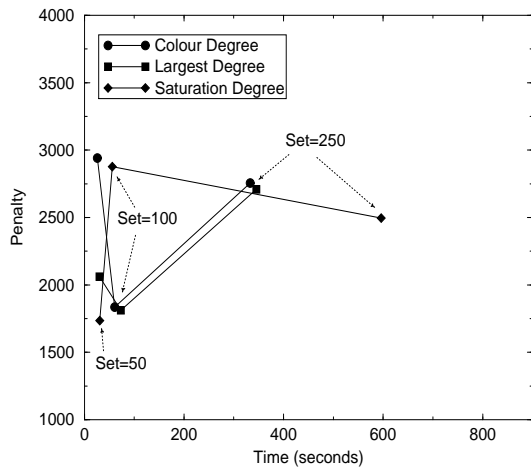


Figure 7. Results when run on **kfu** without a look ahead set




Figure 8. Results when run on **kfu** with a look ahead set

Looking at the results when using Largest degree, shown in Table 2, we see quite a mixed picture. On the **carf92** problem the best result is obtained with a subset size of 50 and a look ahead set, taking less than one tenth of the amount of time to execute, with the larger subset sizes failing to find feasible solutions. Similarly for the **kfu** problem the best result is also found with the subset size of 50 and a look ahead set, but the difference with the higher subset sizes is not quite as great. The situation with the **nott** data is quite different. In this case, only the larger subset sizes are able to consistently find good quality timetables. This is most probably due to high room utilisation (nearly 100%) and the low density of the problem, making it more of a bin-packing problem than a graph colouring one and therefore lessening the effects of the heuris-

Data	Subset Size	Look Ahead	Penalty	Time	Results from Best		
					Unscheduled exams	2nd Order (Same Day)	2nd Order (Overnight)
carf92	50	no	2211	47	0	433	792
		yes	1765	186	0	363	576
	100	no	2023	109	0	393	768
		yes	2328	455	0	433	783
	250	no	2613	600	0	463	931
		yes	10910	1878	1	694	917
500	no	10917	1830	1	697	911	
	yes	11495	1857	1	693	841	
kfu	50	no	1735	31	0	228	786
		yes	1626	105	0	307	589
	100	no	2878	56	0	372	636
		yes	2688	233	0	471	890
	250	no	2495	596	0	468	794
		yes	3883	690	0	455	1078
nott	50	no	907	41	0	133	420
		yes	552	156	0	65	324
	100	no	889	78	0	127	320
		yes	524	340	0	76	282
	250	no	639	336	0	80	316
		yes	771	1120	0	130	328
500	no	754	954	0	128	326	
	yes	1076	2700	0	147	404	
pur	50	no	130228	995	17	3116	3562
		yes	80323	1522	10	3514	3282
	100	no	108248	904	14	3835	3089
		yes	65461	2850	9	2890	3576
	250	no	107556	2655	16	2679	2643
		yes	160292	16411	28	1972	2294
500	no	161776	10240	28	2057	2480	
	yes	218192	26193	38	1727	2118	

Table 4. Results using Saturation Degree

Tables 2 to 4 show the results achieved when using the relevant heuristics. The value given in the time column is the execution time in CPU seconds on a DEC Alpha computer for that particular run. Where the number of exams implies that the last subset of a run is incomplete the final phase of the process is run with a reduced subset size. The breakdown of results for the best of the 5 runs is given as the number of unscheduled exams and the violations of the second order conflicts. The results are also shown graphically in figures 5 to 12.

Data	Subset Size	Look Ahead	Penalty	Time	Results from Best		
					Unscheduled exams	2nd Order (Same Day)	2nd Order (Overnight)
carf92	50	no	2495	42	0	515	884
		yes	1812	184	0	302	804
	100	no	2034	102	0	359	773
		yes	2043	467	0	376	700
	250	no	2766	607	0	501	847
		yes	11260	1715	0	1250	1056
	500	no	11232	1889	0	1244	1050
		yes	13236	1908	2	673	803
kfu	50	no	2942	26	0	359	728
		yes	1627	99	0	250	802
	100	no	1836	61	0	260	999
		yes	2355	254	0	464	724
	250	no	2755	333	0	431	910
		yes	3077	878	0	468	1258
nott	50	no	50897	31	1	259	520
		yes	15923	138	0	77	396
	100	no	15621	102	1	106	330
		yes	544	467	0	67	289
	250	no	5602	607	0	100	255
		yes	759	1715	0	128	293
	500	no	757	1889	0	131	322
		yes	1061	1908	0	208	398
pur	50	no	198097	1426	28	3579	2917
		yes	122499	1830	20	4667	3209
	100	no	166370	1203	22	3958	3212
		yes	84846	2661	11	2577	3068
	250	no	105955	2518	17	2566	3079
		yes	163145	13710	22	1987	3167
	500	no	163405	9616	27	1957	2591
		yes	212961	38394	41	1976	2033

Table 3. Results using Colour Degree

Data	Subset Size	Look Ahead	Penalty	Time	Results from Best		
					Unscheduled exams	2nd Order (Same Day)	2nd Order (Overnight)
carf92	50	no	2577	42	0	524	849
		yes	1822	173	0	313	766
	100	no	2104	119	0	469	614
		yes	2123	547	0	375	837
	250	no	3493	510	0	460	938
		yes	9476	2149	1	777	898
	500	no	12078	1937	1	777	900
		yes	12167	1942	1	812	837
kfu	50	no	2060	30	0	401	707
		yes	1608	105	0	228	704
	100	no	1811	73	0	222	838
		yes	2525	265	0	368	1008
	250	no	2708	345	0	683	839
		yes	2823	762	0	541	836
nott	50	no	31097	31	5	173	450
		yes	4762	219	0	128	350
	100	no	14984	83	0	155	416
		yes	15227	323	0	71	285
	250	no	5833	340	0	99	297
		yes	736	1434	0	94	316
	500	no	766	1265	0	121	397
		yes	1168	2551	0	216	358
pur	50	no	197635	1463	32	3385	3068
		yes	117730	1802	19	3401	3318
	100	no	139741	1097	21	3477	3473
		yes	86558	3097	13	2479	3100
	250	no	93110	2786	15	2765	2907
		yes	155096	15630	28	1939	2522
	500	no	155557	11331	28	2191	2541
		yes	219371	41422	39	1581	2370

Table 2. Results using Largest Degree

Code	Institution	Number of Exams	Enrolment	Seats per Period	Number of Periods	Density of Conflict Matrix
carf92	Carleton University (1992), Ottawa	543	55,552	2,000	36	0.14
kfu	King Fahd University, Dharan	461	25,118	1,955	21	0.06
nott	Nottingham University, UK	800	34,265	1,550	23	0.03
pur	Purdue University, Indiana	2,419	120,690	5,000	30	0.03

Table 1. Data used for testing

For all of the problems, a typical timetable layout was used that was comprised of three periods per day from Monday to Friday and one period on Saturday morning. There were no periods on Sunday. This layout was repeated until the given number of periods for the problem was reached.

The primary task of the algorithm was scheduling all events within the given number of periods and the secondary task was minimising the number of back-to-back conflicts with adjacent periods, with adjacent conflicts on the same day being penalised more than those occurring overnight. For the purposes of solution evaluation and quality the function as shown in EQ 1 was used.

4.2 Results

The method was tested on all the above data sets with subset sizes of 50, 100, 250 and 500 events both with and without a look ahead set. Each of these configurations was tested with all of the given heuristics five times each (and the average result given) in order to provide a reasonable statistically relevant result.

events’). Due to the way records of conflicts are handled, outlined in section 3.3.1, there is no more overhead for checking conflicts with a ‘super’ event as there is for a single event.

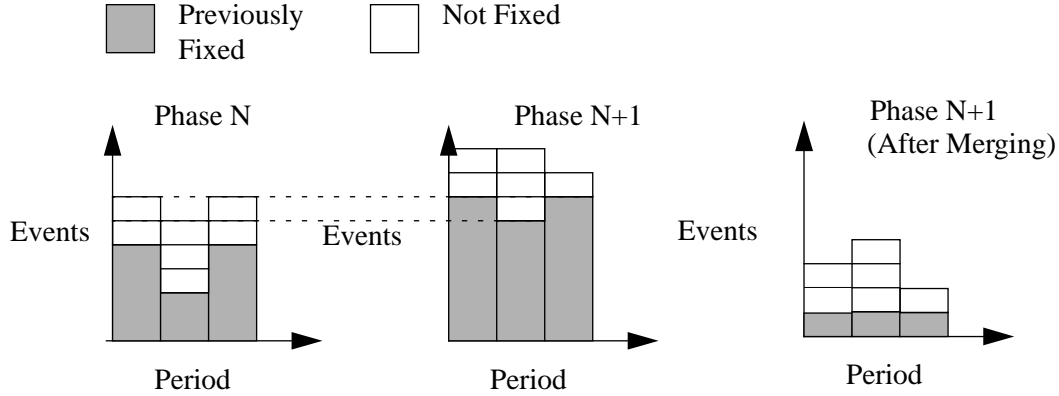


Figure 4. Merging fixed events to improve performance

Figure 4 illustrates how the timetable might build up phase by phase using this approach. The actual process of building a ‘super-event’ is basically a case of summing up the conflict matrices of the component events and the size of each event. For instance say we wanted to create a new super event e_p from a number of events e_1, \dots, e_n scheduled in period p , the event e_p would have the properties shown in equation 12. Note that as we must have a super event for each period the number of real events would now be given by $E - P$.

$$c_{e_p i} = \sum_{j=1}^n c_{e_j i}, \quad \forall i \in \{1, \dots, E\} \quad s_{e_p} = \sum_{j=1}^n s_{e_j} \quad (\text{EQ 12})$$

While we treat the combined ‘super-event’ as having size equal to the sum of its component events here the process is a little more complicated when room allocation is part of problem. Unfortunately only one of the data sets used for testing (the University of Nottingham data) includes information on rooms. In this case the ‘super-event’ preserves the room allocation of its component events, which cannot then be altered.

4 Results

4.1 Experimental Data

In order to evaluate the effectiveness of this approach, several real enrollment data sets were used for testing. Table 1 lists the data used together with the characteristics of each data set such as the total number of exams and the total number of student exam enrollments. The density of the conflict matrix is calculated as the average number of other exams that each exam conflicts with divided by the total number of exams. For example, a conflict density of 0.5 indicates that each exam conflicts with half of the other exams on average. As most of this data does not include data on rooms available, a simple upper limit on the number of seats available each period is used instead. The problems range from the smaller but more densely conflicting **kfu** and **carf92** problems to the fairly large but relatively sparsely conflicting **nott** data set and the huge **pur** data set. All of these data sets can be obtained over the internet from URLs:

- <ftp://ftp.cs.nott.ac.uk/ttp/Data>
- <ftp://ie.utoronto.ca/mwc/testprob>

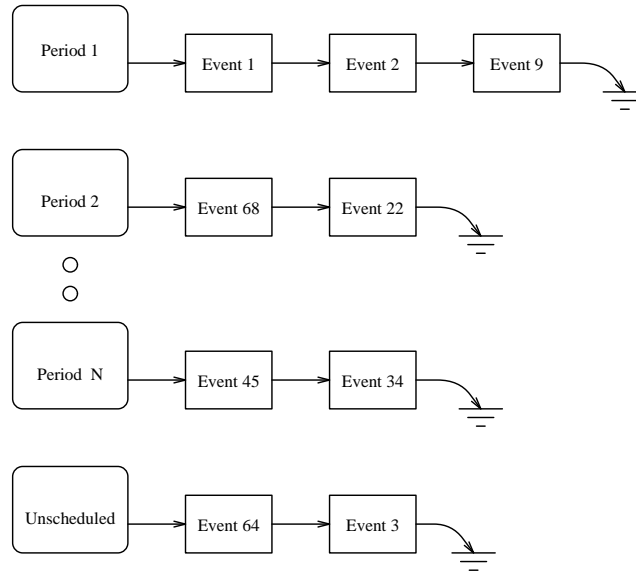


Figure 3. The encoding used for a solution timetable

schedule in sets of 100 events with no look ahead. Now consider that we are in the third phase. 200 events have already been fixed in the previous phase and we are trying to schedule the next 100 events. Even though we are only placing those 100 events we must consider the 200 events already scheduled when checking for first and second order conflicts, or other constraints. Similarly in the final phase we must still consider the 300 events that will have been fixed by then. In an algorithm such as ours where the majority of the execution time is spent either evaluating solutions, portions of solutions or possible moves, it becomes clear that the time taken for each phase will increase substantially later on in the process when, ideally, each phase that considers an equal number of events should require an equal amount of time.

In the implementation used for the experiments here we *merged* events that had already been fixed into a number of virtual “super” events, one of these for each period in the timetable. These super events could be considered as what would happen if all the students from all the events scheduled in a period so far were taken and were all registered on a new single event instead. We can then forget about those initial events and remove them from the timetable, letting our new super events be used for evaluation purposes. This will have obvious complexity advantages as the number of events involved in evaluations will now be the set size we have chosen plus the number of periods in the timetable, regardless of the phase we are currently on (with the obvious exception of the first phase, where there is no need for any of these ‘super

3.2.7 Using Within the Multi-Stage Framework

The operation of this algorithm varies very little when used within the proposed multi-stage framework. There are basically three differences at any particular phase of execution:

1. Any events fixed at the end of an earlier phase cannot be moved by *any* operator in the memetic algorithm, though they are still used for calculating conflicts, seat usage etc.
2. The algorithm ignores all events that will be added to the problem in a later phase.
3. When generating the initial population for a particular phase the set of events for this phase is added in the same random order *over* the fixed events from previous phases.

3.3 Implementation Matters

The following information is given for anyone reimplementing this approach so equivalent performance is achieved.

3.3.1 The Memetic Algorithm

Conflict Checking

Due to the number of times conflicts between exams need to be checked it is impractical to obtain the number of conflicts each time by comparing the lists of enrolled students for those exams. Fortunately as we are more interested in the *number of students involved in the conflict*, rather than the *actual students involved* in the conflict we need only perform the comparisons once. Having calculated the conflicts we can then create a *conflict matrix* C . This matrix C will have dimensions E by E , the definition c_{ij} from section 2.2 being the (i, j) th element of this matrix. Storing this matrix within the algorithm enables rapid conflict checking, and also eliminates the number of students as a factor in the complexity of the problem. For instance let us say that 4 students must attend both exams a and b . In this case the (a, b) th element of the matrix will have the value 4.

Data Structures

Due to the hill-climbing operator being responsible for a large part of the operation of the memetic algorithm, the encoding of solution timetables is designed to reduce the time taken to evaluate and make small steps. Figure 3 illustrates the link list style encoding used. Each period, plus the unscheduled list, has an associated object in the solution that knows the first event scheduled in that period. Similarly each event also has an associated object that knows the next event scheduled in that period, or has a null pointer as it is the last event. By representing a timetable this way we can easily move through the events in a period to check conflicts etc. It also facilitates rapid movement of single events from one period to another, making it highly suitable for the hill climbing operator. Aside from this however, we may also need to quickly establish which period an event is scheduled in. For this, we make the event objects keep knowledge of which period its associated event has been scheduled in. By then storing all the event objects in an ordered array we can quickly access the information about any one event. Altogether, this structure provides for efficient manipulation and stable memory usage.

3.3.2 The Multi-Stage Decomposition Method

This method has to be implemented very carefully, otherwise we could find that using this method produces results little or no faster than the traditional single phase approach. The authors found that the main reason that this might happen was the time taken to perform evaluations (and delta evaluations) at each stage of the algorithm. For instance, consider a problem with 400 events and suppose we have decided to

- **ELSE IF** disrupt **THEN**
 - disrupt this period
- **ELSE**
 - do not disrupt this period

A limitation of this approach is that the first period will never be disrupted, as there is no previous period by which to evaluate the spread properties. However this period is not immune to changes made by either the light mutation operator or the hill-climbing operator, and thus should change substantially over the course of evolution.

3.2.5 The Hill-Climbing Operator

A simple efficient hill-climbing operator is applied after each mutation operator to restore solutions to local optima. This utilises delta evaluation[Ross94] to avoid the time consuming process of performing a full evaluation at every step. Where equation 1 gives us our fitness function, our delta evaluation function when moving event i from a period p to a period q is defined in EQ 11.

$$\Delta fitness(i, p, q) = \sum_{j=1}^E ((t_{ip}t_{j(p+1)}c_{ij}d_{p(p+1)} + t_{ip}t_{j(p-1)}c_{ij}d_{p(p-1)}) - (t_{i(q,q+1)}t_{j(q+1)}c_{ij}d_{q(q+1)} + t_{iq}t_{j(q-1)}c_{ij}d_{q(q-1)})) \quad (\text{EQ 11})$$

The actual process used by the hill-climber is as follows:

- **REPEAT**
 - **FOR** each period p (in some random order)
 - **FOR** each event i scheduled in period p
 - Schedule event i in the valid period causing least penalty (this includes the original period)
 - Try and schedule any unscheduled events
- **UNTIL** no improvement can be made

While the hill-climbing routing does occasionally check if it can schedule events that are currently unscheduled its main aim is to improve the quality of already feasible solutions, rather than produce low quality feasible solutions from infeasible solutions. A more complicated local search procedure might achieve this, but would inevitably increase the total time spent by the algorithm.

3.2.6 Selection

Selection is achieved by using a simple rank based selection method, with each candidate being given a probability of being selected in order of quality. The process functions by ordering candidate solutions in order of fitness. The fittest solution (that with the lowest penalty) is given probability p of being selected. If it fails this probability the second fittest solution is offered the same chance. This continues until either a solution wins its chance or the final solution is reached. The probability p in this case was set to be 0.08.

3.2.2 Initial Population Generation

The initial population is generated by taking events in a random order and scheduling in the first valid period. The hill-climbing operator is then applied to each member of this population.

3.2.3 Light Mutation

As mentioned in Section 1.2 this operator selects a number of events (30 in this case) at random, and tries to place each in an alternative period, also picked at random. Moves that violate hard constraints are not allowed, and in this case another period is tried.

3.2.4 Heavy Mutation

This operator is one of the more complicated operators in the algorithm and we will describe it in some detail. It targets periods that appear to be causing “large” amounts of penalty for “disruption”. By “disruption” here we mean to say that all events contained within this period are temporarily unscheduled. All events from disrupted periods can then be randomly rescheduled in the timetable. The penalty for a period p is based on the evaluation function and is defined in EQ 8.

$$penalty_p = \sum_{i=1}^{E-1} \sum_{j=i+1}^E t_{ip} t_{j(p+1)} c_{ij} d_{p(p+1)} \quad (\text{EQ 8})$$

What actually constitutes a “large” amount of penalty is clearly relative so the algorithm also needs something with which to compare. For this it can calculate the average penalty per period of the *best* solution in our population as in EQ 9.

$$average = \frac{Fitness\ of\ best}{Number\ of\ periods} \quad (\text{EQ 9})$$

Having calculated these the algorithm can now calculate a probability of being disrupted for each period. EQ 10 shows how the probability of being disrupted is calculated, with *bias* being a definable value (0.1 here) to vary the probability of periods being disrupted. Periods causing greater than *average* penalty are automatically disrupted.

$$\begin{aligned} Probability(Disrupt\ period\ p) &= 1, && \text{if } penalty_p \geq average \\ &= \frac{penalty_p + (2 \times bias)}{2 \times average}, && \text{if } penalty_p < average \end{aligned} \quad (\text{EQ 10})$$

As in this case the penalty arising from a particular period is dependant on the events scheduled in the next period the algorithm must have special cases for disrupting two periods in a row. For instance, say the operator preserved period i because it had a low number of conflicts with period $i+1$ it does not make sense to then disrupt period $i+1$. Instead the operator should maybe disrupt the next period. This can be summarised formally as follows:

- **FOR** each period p (where $p > 1$) in the timetable in turn
 - decide if p is to be disrupted based on the probability given by EQ 10
 - **IF** decision is to disrupt but previous period was not disrupted **THEN**
 - do not actually disrupt this period, instead disrupt the next period

The multi-stage process when scheduling N exams at a time can be described by the following pseudo-code:

- $E' = 0$
- **REPEAT**
 - $E' := E' + N$
 - **IF** $E' > E$ **THEN**
 - $n = N - (E' - E)$
 - $E' := E$
 - **ELSE**
 - $n = N$
 - **FOR** each unscheduled exam
 - Calculate the desired degree of that exam
 - Pick the n exams with the largest/least degrees
 - Apply the memetic algorithm to schedule those n exams, keeping the $(E' - n)$ previously scheduled exams fixed in their current position.
- **UNTIL** $E' = E$

The point at which to exit the memetic algorithm presents a choice. Normally the memetic algorithm would exit when the population has fully converged but this is inappropriate with our mutation driven memetic algorithm. Two other options would be to either run the algorithm for a fixed number of generations, or alternatively to end the algorithm after a set number of generations have passed without finding a new “best so far” solution. In these experiments the algorithm was halted after 5 generations without a new “best so far” solution. While 5 generations may seem a little low, it has to taken into account that a generation typically achieves considerably more (while also consuming considerably more CPU cycles) than a typical genetic algorithm due to the use of the hill-climbing operator.

While the decomposition method is in itself independent of the particular technique used to solve each sub-problem the memetic algorithm presented in [Burk96b] has been shown to be effective at solving all but the very largest timetabling problems.

3.2 The Memetic Timetabling Algorithm

The memetic approach employs a simple evolutionary model with a population of solutions (of size 50 for all experiments described here). The algorithm applies mutation operators (one of the light or heavy forms, decided by a 50/50 probability), followed by a hill-climbing operator, to produce an oversize population (which is twice the normal population size). Solutions are then selected from this oversize set to form the new population for the next generation. Previous experiments [Burk97] with recombination operators for this algorithm have proved unfruitful and as such were not used for these experiments.

3.2.1 Quality and Evaluation

A simple linear weighted penalty function is used to assess the quality (and hence the fitness, or rather unfitness) of solutions. This function is identical to that shown in EQ 1.

- *Largest Degree.* We choose to schedule first those events with the greatest number of conflicts with other events. These would generally be considered to be more difficult.
- *Colour Degree.* Similar to largest degree except that we schedule those events first that have the greatest number of conflicts with other events that have already been scheduled. We would expect that events with conflicting events already scheduled would be more difficult to place than events with a large amount of conflicts, but little or none with any events already scheduled.
- *Saturation Degree.* Here we schedule first those events with fewer valid periods remaining in the timetable. We might expect that if we did not schedule these events sooner there may no valid periods at all remaining to schedule these events later on.

We can use our formal definition of the problem to define the above terms unambiguously. The *degree* of an exam is defined in EQ 5, where # gives the cardinality of a set and the dot notation represents “such that” relationships. As the degree of an exam is constant this need only be calculated once in any run of an algorithm.

$$degree_i = \# \{ (j \in \{1, \dots, E\}) \bullet c_{ij} > 0 \} \quad (EQ 5)$$

Similarly the *colourDegree* of an exam is defined in EQ 6. Unlike the degree, the colour degree of an exam changes every time one of its conflicting exams is moved from the unscheduled list to a valid period.

$$colourDegree_i = \# \{ (j \in \{1, \dots, E\}) \bullet \exists (p \in \{1, \dots, P\}) \bullet t_{jp} = 1 \wedge c_{ij} > 0 \} \quad (EQ 6)$$

Finally the *saturation degree* of an exam is defined in EQ 7. This will need to be recalculated every time an event is moved to a period, either from the unscheduled list or another valid period.

$$satDegree_i = \# \left\{ (p \in \{1, \dots, P\}) \bullet \sum_{j=1}^E t_{jp} c_{ij} = 0 \wedge \sum_{j=1}^E t_{jp} s_j \leq S - s_i \right\} \quad (EQ 7)$$

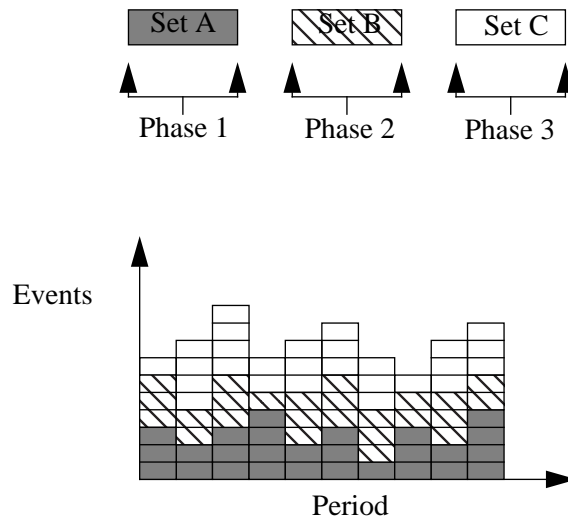


Figure 1. How a problem could be divided

This should substantially reduce the complexity of the problem but there will of course be the obvious pit-fall that by fixing events in periods in this fashion it may be impossible to schedule events later on in the process. Fortunately, there are a number of methods we can employ to reduce the chances of this happening. Firstly we could borrow an idea from heuristic sequencing methods and choose our subsets according to some heuristic ordering. This should help improve the quality of produced timetables by optimising first those events that are likely to cause most problems. Secondly the process could look ahead in the process by optimising two subsets at a time but only fixing the first of these subsets at the end of every stage. This will inevitably lead to the process taking longer than if a single set of the same size was considered, as the process is considering most of the events twice, but this could be justified if it leads to substantial improvements in quality. Figure 2 illustrates how the earlier 3 phase example would be approached using a look ahead set.




Figure 2. Using a look ahead set

While having this look ahead set will help, it is probably more important that a good heuristic is chosen with which to order our events. If the process concentrates on scheduling those more difficult events earlier on in the process and reducing the penalty caused by these events with respect to each other we should find that fewer problems are encountered later on in the process, regardless of whether or not we are using a look ahead set. For these experiments we will use three heuristics that are generally accepted to be suitable for the exam timetabling problem[Cart95]:

$$\sum_{i=1}^{E-1} \sum_{j=i+1}^E \left[\sum_{p=1}^P t_{ip} t_{j(p+1)} c_{ij} d_{p(p+1)} + t_{ip} t_{j(p-1)} c_{ij} d_{p(p-1)} \right] + 5000 t_{i(P+1)} \quad (\text{EQ 1})$$

subject to

$$\sum_{p=1}^{P+1} t_{ip} = 1, \quad \forall i \in \{1, \dots, E\} \quad (\text{EQ 2})$$

and

$$\sum_{i=1}^{E-1} \sum_{j=i+1}^E \sum_{p=1}^P t_{ip} t_{jp} c_{ij} = 0 \quad (\text{EQ 3})$$

and

$$\sum_{i=1}^E t_{ip} s_i \leq S, \quad \forall p \in \{1, \dots, P\} \quad (\text{EQ 4})$$

3 A Multi-Stage Memetic Algorithm

3.1 Basic Framework

While memetic algorithms show promise for timetabling problems[Burk96b, Paec95] the time involved in optimizing large problems (of say greater than 500 events) is much longer than we would find ideal. It would be much more preferable if the algorithm could produce timetables in a matter of minutes rather than hours. To achieve this it might make sense to use the original memetic algorithm but only apply it to a subset of the total events at a time. The algorithm can then fix those events in the timetable before considering the next subset of events and trying to schedule those on top of those already scheduled in the previous stage. The process of decomposition has been studied by Carter[Cart83], proposing a heuristic method of recursively splitting a large problem into smaller problems until each subproblem is small enough to be solved by conventional methods such as linear integer programming.

The idea has also been studied by Weare[Wear95] who applied such a technique to random data on flexible length timetables. The results of this indicated that as the number of events considered at one time is decreased the time taken to produce solutions also decreases but unfortunately so did the quality of solutions. A contributing factor to the lower quality could well have been the use of the flexible length timetable model and a preference for shorter timetables. This may have caused shorter timetables to be produced in the earlier phases (at the expense of second order conflicts) when the final timetable may be much longer anyway. Therefore for the purposes of our experiments we will be using a fixed length timetable model to avoid these problems. We will also experiment with modifications and enhancements for this method. Figure 1 shows how a set of events could be split into three subsets and then scheduled in three different phases, where the darker portions of the graph represent how early on in the process that portion was fixed in the timetable.

2 Problem Definition

The technique presented here was tested on an instance of the examination timetabling problem. At this point we will introduce the notion of *hard* and *soft* constraints. Hard constraints are considered essential to satisfy in terms of producing a practical timetable. Soft constraints on the other hand are regarded as desirable to satisfy, but not essential.

In this problem we have E exams that must be scheduled in P periods with S examination seats available for each period (where E , P and S are non-negative integers). There are 3 periods in a day. Any two of the exams may *conflict* with each other, that is they have a number of students enrolled for both exams. When scheduling the exams we can encounter two types of conflicts:

- *First Order Conflicts*. This term is used to describe situations where conflicting exams are scheduled in the same period. This is highly undesirable as it involves quarantining some students after one exam so they may sit for the other exam after the main sitting. In all but the most difficult of situations this is regarded as a basic hard constraint.
- *Second Order Conflicts* on the other hand are less important and represent situations where two conflicting exams are not scheduled in the period, but are scheduled in periods too near to each other. For instance we might not want students to have to sit for two exams in consecutive periods, or to sit for two exams in the same day. To totally satisfy these constraints is often not practical, in which case they are treated as soft constraints.

As well as avoiding these conflicts we must also adhere to the limitations on seating available for each period. We can formally specify the problem mathematically by first defining:

$t_{ip} = 1$ if exam i is scheduled in period p , 0 otherwise

c_{ij} as the number of students taking both exams i and j

d_{pq} is 3 if period p is on the same day as period q , 1 if they are on adjacent days and 0 otherwise. This weights same day conflicts higher as they are more important to satisfy.

s_i is the number of students taking exam i

To take account of the fact that a feasible timetable might not necessarily be possible we need an extra period, the $(P+1)$ th period, into which we can place any exams not yet scheduled in a valid period. There are no constraints to prevent scheduling in this period, but it will need to be heavily penalised. Now we need to minimise EQ 1 subject to EQ 2, EQ 3 and EQ 4 (all presented below). EQ 1 sums up occurrences where students must attend 2 exams in consecutive periods and weights the number of unscheduled exams by 5000 to strongly discourage incomplete timetables. If two consecutive periods are on the same day then any adjacent conflicts are weighted by 3. If there is a single night between them the conflicts are weighted by 1, otherwise any conflicts are ignored (for example when periods are split by a weekend). EQ 2 states that every event should be scheduled once, and only once in the timetable. EQ 3 specifies that no conflicting events should be scheduled within the same period and EQ 4 enforces the condition that the total number of seats required for any period is not greater than the number of seats available.

The problem can be presented formally as follows:

Minimise

be the case that this would not cause any problems. Unfortunately limits on time and rooms often mean that finding any feasible solution is a considerable task in itself and that no zero penalty solution will exist.

The timetable problem is known in general to be NP-complete [welsh67, karp72]. Of the many approaches that have been applied to the timetable problem the earliest methods are those that utilise heuristics in some way. A good example of this is *heuristic sequencing* [Fox168], which involves using a heuristic to estimate how difficult each event will be to schedule. If we then schedule the events we regard as most difficult before the other events we should produce a substantially better timetable than if we were merely using a random ordering. A useful addition to this process is *backtracking*, such as that used in [Cart94]. This deals with situations where there is no valid period available to schedule an event because of the placements that have already been made. It does this by unscheduling the conflicting events from a chosen period in order to allow the current event to be scheduled in that period. The unscheduled events are then placed back in the ordering for later scheduling.

In addition to these heuristic methods a number of general search methods have also been applied to the timetable problem. Such approaches include simulated annealing [Thom96], Tabu search [Hert91], and variants on genetic algorithms [Burk95, Burk96b, Corn94, Paec95]. While most of these have shown worth it is notable that many of the more successful methods employ some kind of ‘trick’ in the algorithm to improve results for timetable problems. It logically follows that specialized timetabling algorithms may function substantially better than those based on more general optimisation techniques.

For the interested reader more information on the examination timetabling problem and the methods that have been applied can be found in two excellent surveys by Carter [Cart86] and (more recently) by Carter and Laporte [Cart96] which go into considerable depth. For further information about the problem itself please see the survey by Burke *et al.* [Burk96a] of examination timetabling in universities throughout the United Kingdom.

1.2 Memetic Timetabling

In [Burk96b] an approach was proposed for timetabling problems based on memetic algorithms. In essence memetic algorithms are evolutionary algorithms that utilize local search to some extent. The concept of memetics originates from Dawkins [Dawk76] who describes the meme as an idea or concept that is passed around society. Individuals can then adapt ideas to suit their own environment, as opposed to genes which are passed down whole and cannot be altered by the recipient. In the same way that memes were related to genes by Dawkins, memetic algorithms were proposed as an alternative to genetic algorithms [Davi91] by Moscato and Norman [Mosc91]. The main motivation of this approach is that by applying a hill-climbing operator after each of the genetic operators we can then treat the process as a search through local optima, rather than the entire search space. Applying a local search operator like this will clearly cause each generation to take considerably longer but this can of course be justified if sufficiently more is achieved per generation than if local search were not used.

The memetic algorithm proposed in [Burk96b] uses a combination of mutation and local search to effectively search the solution space. The mutation process actually consists of two separate operators, light and heavy mutation. The light mutation operator merely shuffles a few individual events to other valid periods, the purpose of this being to ‘nudge’ solutions away from local optima in order to find a new solution. The heavy mutation operator, on the other hand, functions by targeting periods where large amounts of penalty are arising. The algorithm then randomly reschedules the events in these periods to produce a new solution that retains some of the (hopefully better) characteristics of the original solution. Neither of these operators on their own achieves any substantial improvements in solution quality but when followed by an application of a simple hill-climber the process becomes much more effective [Burk96b].

A Multi-Stage Evolutionary Algorithm for the Timetable Problem

E. K. Burke

J. P. Newall

Automated Scheduling and Planning Group

Department of Computer Science

University of Nottingham

Nottingham

NG7 2RD

ABSTRACT

It is well known that timetabling problems can be particularly difficult to solve, especially when dealing with particularly large instances. Finding near optimal results can prove to be extremely difficult, even when using advanced search methods such as evolutionary algorithms (EAs). In this paper we present a method of decomposing larger problems into smaller components, each of which is of a size that the EA can effectively handle. We will show various experimental results using this method to prove that not only can the execution time be considerably reduced but also that the presented method can actually improve the quality of produced solutions.

1 Introduction

1.1 The Timetable Problem

In essence the timetabling problem consists of allocating a number of events to a finite number of time slots (or periods) such that the necessary constraints are satisfied. In general the nature of the constraints varies between different instances of the timetable problem but most problems share the conditions that:

- No individual entity (e.g. person) should be required to attend two events simultaneously.
- For each period there should be sufficient resources (e.g. rooms) to service the events scheduled within that period.

These constraints are fundamental to any timetabling problem and generally form the basis for a *feasible* timetable but, as mentioned above, a number of other constraints may be introduced depending on the particular flavour of the timetable problem being considered. Often, these varying constraints can be at odds with each other. For instance, when scheduling university examinations it is often considered undesirable to have conflicting exams in adjacent periods. Alternatively when scheduling university lectures it is often actually preferred for students to have two or three lectures in a row. Putting aside these differences for our experiments we will be concentrating on an instance of the examination problem. While there are a number of possible side constraints for this problem the most difficult task to carry out (aside from producing a feasible timetable at all) is often to minimise the number of conflicts between adjacent periods in order to allow students time to recover between exams. If allocated a large number of periods it would most likely