

KHE18: A Solver for Nurse Rostering (Working version, 26 August 2018)

Jeffrey H. Kingston

Received: date / Accepted: date

Abstract The problem of assigning nurses to the shifts of a hospital ward, known as *nurse rostering*, has been studied for many years. This paper presents work in progress on the KHE18 nurse rostering solver. A time sweep algorithm is used to make an initial assignment, followed by repair using several methods, including ejection chains. Results are presented for several standard data sets.

Keywords Nurse rostering · Time sweep · Ejection chains · XESTT

1 Introduction

The problem of assigning nurses to the shifts of a hospital ward, known as *nurse rostering*, has been studied for many years. It is an NP-complete problem, and exact algorithms are out of reach in general, although many smaller instances have recently been solved to optimality with integer programming [4,30].

Many inexact methods have been tried. Even very recent work covers a wide range: integer programming [14,24,27,30], variable neighbourhood search [32], simulated annealing [9], weighted maxSAT [10], hyper-heuristics [2,15,28], and constraint programming [29]. For less recent work, consult [33].

The solver presented here, KHE18, is the 2018 version of the main solver built by the author on his KHE timetabling platform [19]. It runs in polynomial time and aims to find a good but not optimal solution quickly. It uses a time sweep algorithm for constructing an initial solution (Section 3.1). This assigns nurses to the first day of the cycle, then to the second, and so on. After that, it tries several repair methods, including ejection chains (Section 3.2).

KHE18 is work in progress. It has been tested on three data sets (Section 4). It runs quickly and has found some solutions as good as, or nearly as good as, the best in the literature; but overall its results are not yet competitive.

2 The nurse rostering problem and its XESTT formulation

Nurse rostering is the problem of assigning shifts to the nurses of a hospital ward. Hospitals operate 24 hours a day, so there are usually at least three types of shifts: morning, afternoon, and night. Each shift demands a certain number of nurses, often with specific skills. There may be some flexibility in how many nurses to assign, and the number typically changes from day to day.

Perhaps the most characteristic feature of the problem is the large array of requirements that each nurse's timetable must satisfy. In addition to workload limits, there are rules such as 'a nurse must have a day off after a sequence of night shifts', 'a nurse may work at most four days in a row', and so on.

Instead of the usual formulas, this paper's formal definition of the nurse rostering problem is supplied by the XESTT [21] nurse rostering data format. XESTT is an XML format which is capable of representing the instances found in all the well-known data sets. It is based on the XHSTT high school timetabling format [25,26]; the name 'XESTT' was chosen to be reminiscent of 'XHSTT', with 'employee scheduling' replacing 'high school'. Full details of XESTT appear online [16] and will not be repeated here. Instead, this section offers an overview, and explains the importance of XESTT to the present work.

An XESTT instance consists of the *cycle* (the sequence of *times* that events may be assigned); a set of *resources* (entities that attend events); a set of *events* (meetings); and a set of *constraints*, specifying conditions that solutions should satisfy, and penalties to impose when they don't.

Each event contains a *starting time*, which may either be preassigned a time or left open for a solver to assign; a *duration*, which is a fixed positive integer giving the number of consecutive times, starting at the starting time, that the event is running; an optional *workload*, which is a fixed non-negative integer representing the workload of the event in arbitrary units, for example in minutes; and any number of *event resources*, each specifying one resource which attends the event for the full duration, which may either be preassigned a resource or left open for a solver to assign.

In nurse rostering instances, each event represents one shift. Each event has duration 1; its actual duration in minutes can be expressed as a workload, if needed. Its starting time is preassigned to a time unique to the shift. For example, if on each day there is a morning, afternoon, and night shift, then each day will contain three times, one for each shift. This arrangement is somewhat artificial, but, as [21] explains, most nurse rostering constraints concern shifts, not workload in minutes, and this works best in practice. Within an event, each event resource represents a demand for one nurse.

Sets of times, resources, and events may be defined, called *time groups*, *resource groups* and *event groups*. Each resource has one *resource type*, saying what type of resource it is. In nurse rostering there is just one type, **Nurses**.

XESTT offers 18 constraint types, but 9 are not used in nurse rostering, mainly because all the events have preassigned times and duration 1. Of the 9 types that are used, 3 are *cover constraints*, specifying the number of resources that should attend each event, and the skills those resources should have. The

other 6, called *resource constraints* here, constrain the timetables of individual resources, specifying unavailable times, workload limits, unwanted patterns (such as a day shift immediately following a night shift), and so on.

Each constraint contains a Boolean *required* flag indicating whether it is *hard* or *soft*, and an integer *weight*. When a constraint is violated, the degree of violation is multiplied by the weight to give a cost. Algorithms aim to minimize firstly the total cost of hard constraints (the *hard cost*), and secondly the total cost of soft constraints (the *soft cost*). In nurse rostering, solutions with non-zero hard cost are usually considered to be *infeasible*, that is, useless.

XESTT is important here for two reasons. First, it makes it easy to test KHE18 on a wide range of instances, because all these instances have been converted from their original formats to XESTT.

Second, XESTT uses just 9 types of constraints to represent all of the constraints found in other models. An algorithm like the ejection chain repair algorithm in this paper, which handles each constraint type explicitly, has only 9 types to handle. Without XESTT or something like it, the number of constraint types would be much larger, and an approach that handles each constraint type explicitly would hardly be feasible.

3 The KHE18 solver

The KHE18 solver presented here is built on the author’s KHE solve platform and is available from the KHE web site [19]. It is descended from the KHE14 high school timetabling solver [18]. It is in fact a general timetabling solver, but because in nurse rostering the times of all events are preassigned, its time assignment part does nothing here except to convert the time preassignments in the instances into assignments in the solutions, taking almost no time.

After time assignment comes resource assignment—the assignment of nurses to shifts. KHE18 first constructs an initial assignment using time sweep, then continues with a number of repairs which improve that initial assignment. Space precludes a complete presentation, although that is available online, in the KHE documentation [19]. Instead, this section focuses on the main points.

3.1 Time sweep assignment

Because of the high density of constraints in nurse rostering, it is often easier to avoid introducing a problem during construction of the initial solution than to remove it later. So it makes sense to try hard to produce a very good solution during the initial construction phase.

Many nurse rostering constraints concern what happens over consecutive days within nurses’ timetables. This suggests that the initial solution should be constructed day by day—the shifts of the first day assigned first, then the shifts of the second day, and so on. As each day is assigned, these kinds of constraints can usually be satisfied. This we call the *time sweep* method. The

assignment of resources to the shifts of one day can often be made in a locally optimum manner in polynomial time using weighted bipartite matching. The rest of this section describes KHE18's implementation of this approach.

The algorithm first checks that a hard constraint limits each nurse to at most one shift per day (if not, it applies itself to individual shifts, rather than to days). Then, for each day it builds a weighted bipartite graph. The graph contains one *demand node* for each event resource of each event (shift) on that day, since each event resource is a demand for one nurse. It also contains one *supply node* for each nurse available for assignment. An edge joins a demand node to a supply node when the nurse can be assigned to the shift (which is practically always; the few exceptions need not detain us here). The cost of the edge is the cost of the solution containing all the assignments on previous days, plus just this one assignment on the current day. There is also one supply node for each demand node, representing non-assignment of a nurse to the demand node's event resource. The cost of this edge is just the current solution cost. This will usually include a penalty for not assigning a resource, which will be absent from the edges which do assign one. Assuming that the assignments represented by the edges have an independent effect on solution cost, a minimum-cost maximum matching in this graph represents a locally optimal assignment of nurses to the event resources of the day.

When are these assignments independent, in fact? All costs are produced by constraint violations, so this question can be answered by examining each of the 9 constraint types, to see whether the costs it generates are independent. And indeed most of them are. Constraints on the timetable of an individual nurse, for example, are affected only by assignments of that nurse, and there is at most one of those each day. XESTT has just one kind of constraint whose cost depends on multiple edges: the *limit resources constraint*. For example, a constraint of this type could require at least one senior nurse to be on duty at all times of the day, which is not the same as requiring one such nurse on every shift, because shifts may overlap in time within one day.

Rather than abandoning matching when limit resources constraints are present, the algorithm approximates them by adjusting edge weights. Thus, to encourage one senior nurse to be present, the weights on the edges from one demand node to the supply nodes for non-senior nurses are increased. Local optimality is often preserved, but not always. Full details appear in [19].

To encourage an even distribution of workload, the edge costs are adjusted slightly to break ties in favour of assigning resources with more unused workload. This has been observed to significantly improve solution quality. Two other adjustments are also included: one which favours assignments that do not bring constraints from below their upper limits to their upper limits, and another that favours shorter runs of consecutive busy days (reasoning that they will offer more flexibility when repairing the timetable later). Both have been observed to give marginal further improvements. These observations need to be tested by large-scale experiments, which have not yet been carried out.

Constraints with minimum limits often produce spurious costs during time sweep (costs that will disappear as the time sweep moves on). For example, if

busy days should come in groups of at least 2 consecutive days, then the first assignment after a free day will attract a spurious cost. KHE18 avoids this by informing constraints that the days after the current day should be treated like days after the end of the cycle often are: the constraint should not assume that they are either assigned or not assigned, and make its best estimate of cost (a lower bound on the true cost) accordingly. This turns off most minimum limits. A full description, with formulas, is given elsewhere [22].

The great weakness of time sweep—its inability to look ahead—is mitigated by including repairs in the construction stage. After each day is assigned, a few preceding days are individually unassigned and reassigned, again using weighted bipartite matching. This may produce small changes that coordinate better with the assignments now present on following days.

The author is unable to cite previous work using time sweep. The general idea is well known, although whether the details described above have been tried is not known. They make a substantial difference.

3.2 Ejection chain repair

After constructing an initial solution using time sweep, KHE18 repairs it using several methods. The most significant, in the amount of improvement attained and in running time, is *ejection chain repair*.

A *defect* in a solution is one violation of a constraint. For example, if nurse N2 should work at most 2 weekends but in fact works 3, that is a defect.

A *repair* is a change to the solution which removes a defect. For example, unassigning one of nurse N2's busy weekends repairs the defect just described.

An ejection chain is a sequence of repairs. Starting from some defect, the first repair repairs that defect but introduces one new defect. The second repair repairs that defect, but introduces another new defect, and so on. If some repair repairs a defect without introducing a new defect, then the chain ends and the solution has been improved. Or if the repair of one defect introduces two or more new defects, the chain has to be undone, with no improvement.

There are usually several ways to repair a defect. For example, nurse N2's defect can be repaired by unassigning any one of the three busy weekends. So finding a successful chain involves a search tree: if the first repair does not begin a successful chain, then the second is tried, and so on recursively.

The main loop of the ejection chain repair algorithm visits each defect of the current solution and attempts to remove it by searching a tree of ejection chains, stopping at the first successful chain, if any. It cycles around the defects until a complete cycle of attempts has failed, at which point no further progress is possible and the algorithm terminates. Alternatively, it terminates when one of its periodical checks reveals that the soft time limit has been reached.

There are several strategies for limiting the method to polynomial time. The usual one, which KHE18 uses, is to refuse to visit the same part of the solution twice while trying to remove a given defect [17,18].

Each of the 9 kinds of constraints gives rise to one type of defect (or one might say two types, because a maximum limit violation is repaired differently from a minimum limit violation). For each defect, a set of repairs tailored to its type is tried, making the chains *polymorphic*. Space precludes a full description of these sets of repairs, which in any case are still evolving. They will be defined at [19] when settled. Meanwhile, here is a general overview.

The basic operations underlying all repairs are *shift assignment*, which assigns resource r to unassigned shift s , *shift unassignment*, which unassigns resource r from shift s , and *shift move*, which changes the assignment of shift s from resource r_1 to resource r_2 . One repair is a set of these basic operations.

For example, if the defect is day shift s_2 immediately following night shift s_1 in the timetable of resource r , then the obvious repairs to try are shift moves of s_1 from r to any other resource, and shift moves of s_2 to any other resource. If the problem is 6 consecutive busy days in the timetable of resource r when the maximum limit is 4, the most likely repairs are moving the first two shifts to some other resource, and moving the last two shifts. And so on.

One general idea has emerged. Suppose that some shift move, of shift s from r_1 to r_2 say, seems worth trying for any reason. (The idea also applies to shift assignments and unassignments.) In high school timetabling, one would move that shift only, but that is not likely to work well in nurse rostering, owing to the interdependence of shift assignments on adjacent days. Instead, it seems better to always move a set of shifts from adjacent days together, to the same new value, as in the works by other authors cited below. So the author's code expands the proposed shift move into a set of alternative repairs, each of which moves several shifts on adjacent days, always including the original move. The number of days ranges from 1 up to some small constant (currently 4). Often the full range cannot be used. For example, if r_1 is already free on the day before s , then no repairs involving shifts on that day or earlier days are tried.

In instances where the demand for nurses is such that all nurses need to work at, or very close to, their workload limits, moving one or more shifts from one nurse to another is almost certain to create a workload limit defect in the second nurse. In these instances, therefore, KHE18 also tries repairs which include the move of an equal number of shifts in the other direction, to keep workloads constant.

The author is aware of three previous uses of ejection chains in nurse rostering. The first [11] is older than the data sets in use today, making its results hard to evaluate. It includes a repair which swaps one week's work between two resources. The second [3,4,7] is recent. Its basic repair swaps a variable number of consecutive days' work between two resources. It allocates equal running time to each top-level repair. Its results are compared with ours in Section 4. The third [23] uses ejection chains as individual moves within a tabu search framework. Each chain is generated at random in a way that preserves coverage, but otherwise without checking cost until the whole chain has been generated. It cites [1] as a source for these kinds of chains—a very early use.

This paper's ejection chain algorithm derives from recent work in high school timetabling [17,18]. Its polymorphism seems new to nurse rostering.

The works cited above bundle defects together by resource and search for repairs which reduce the total cost of one bundle. It is hard to say *a priori* whether this is advantageous or not: on the one hand, it reduces the precision with which repairs can address specific defects; on the other, it is less affected by high constraint density. One paper [3] adds precision by selecting repairs which focus on those parts of the resource's timetable where defects lie.

3.3 Grouping event resources

The KHE platform allows event resources to be *grouped*. Assigning a resource to one element of a group assigns it to all. This was included to support high school timetabling: the lessons of one course spread through the week should be assigned a common teacher. It has also proved useful in nurse rostering.

For example, in instance `COI-GPost` (Section 4.1), a nurse who takes a Friday night shift should also take the following Saturday and Sunday night shifts, because constraints penalize Friday night shifts before free weekends, incomplete weekends (working on Saturday or Sunday but not both), and day shifts after night shifts. Then the Monday and Tuesday night shifts can be grouped, as can the Wednesday and Thursday night shifts, owing to limits (minimum 2 and maximum 3) on the number of consecutive night shifts.

KHE18's method of finding such combinations is as follows. For each shift s , find all combinations of shifts and free time over 2 and 3 consecutive days starting with s . If only one combination has zero cost, group its shifts so that they must be assigned the same resource. Only include costs that depend only on what happens on the days in question and are the same for all resources.

One grouping may lead to another. For example, if an instance requires complete weekends and also that a day shift cannot follow a night shift, then the Saturday and Sunday night shifts can be grouped, but only after that does it become possible to group the Saturday and Sunday day shifts.

If nurses must have a certain number of weekends off, that may only be achievable if all their working weekends are complete, even without an explicit 'complete weekends' constraint. KHE18 detects this and similar cases and uses them to reduce the number of combinations, making grouping more likely.

Of course, grouping might turn out to be unadvisable for some unexpected reason. So KHE18 applies the groupings during the initial time sweep and first repair cycle, but then removes them, so that if something else is needed there is a chance to find it in a later repair cycle.

3.4 Randomization

Randomization is not stressed in algorithmic solvers like it is in, for example, metaheuristic ones. Still, including some randomization allows the algorithm to be run with different seeds to obtain different results. Doing this and keeping the best solution is a simple way to trade off solution quality and running time.

The time sweep algorithm randomizes by choosing a random resource as the first to be given a supply node, cycling through the resources from there. This causes ties in edge weights to be broken differently, at least at the start when many resources have equal unused workload. The ejection chain repair algorithm randomizes by trying alternative repairs starting at a random point. For example, when a shift is to be moved to some other resource, it traverses the list of resources from a random point. These methods are not deep, but they seem to work, judging from the spread of solution costs they produce.

4 Results

This section presents the results of running KHE18 on several well-known sets of instances, after conversion to XESTT by the NRConv program [20,21], with comparisons with previous results. The converted instances and results are available, in the form of XESTT archive files, at [20]. These results are work in progress and are likely to be superseded as KHE18 is improved.

Two versions of the solver are tested. The first, KHE18, runs the solver once and produces one solution. The second, KHE18x8, runs the solver 8 times, with a different random seed on each run, and keeps the best solution.

All solves are run on the author's machine, a quad-core 3.6GHz Intel Core i7-4790. All runs have a time limit of 300 seconds per instance. The limit is *soft*: instead of being interrupted, each potentially slow part of the solver consults wall clock time periodically and does as little as possible after the time limit.

The KHE18 tests run in parallel in 4 threads, solving a different instance in each thread. The average running time per instance is about one second longer than when the instances are solved sequentially. The reason for this difference is not clear; it might be contention in the memory allocator, or in the bus. It is small enough to ignore for the sake of getting the tests done quickly.

The KHE18x8 tests also run in parallel in 4 threads. For each instance, four solves are started initially, one per thread. New solves are started in these threads as old ones end, until all solves have been started or the whole run reaches the time limit. So KHE18x8 could start as few as 4 solves. Once started, a solve runs until it stops itself. After all solves are complete, wall clock time is recorded and the four threads start again on the next instance.

Each result table was generated by the author's HSEval program from an XESTT archive file and included with no hand editing. In each table, a blank entry indicates that there is no solution for the instance of its row in the solution group of its column. If there are multiple solutions, the solution with minimum running time among all solutions with minimum cost is shown.

The costs are not reported on trust; they are calculated from the solutions in the archive file, and hence verified, by HSEval. Only soft costs are shown; any solutions with non-zero hard cost are shown as 'inf.', meaning infeasible.

Running times in seconds are reported where present in the archive. HSEval cannot verify them. KHE18 and KH18x8's running times are wall clock times.

Table 1 Results for the Curtois original instances. Column Misc shows the best solutions from XESTT archive file `COI.xml`. The other columns show the results from the two versions of the solver described in this paper. Here and elsewhere, running times are in seconds. This table is derived from XESTT archive file `KHE18-2018-08-26-COI.xml`, available at [20].

Instance	Misc		KHE18		KHE18x8	
	Cost	Time	Cost	Time	Cost	Time
COI-Ozkarahan	0	-	400	0.0	0	0.1
COI-Musa	175	-	175	0.4	175	0.8
COI-Millar-2.1	0	1.0	0	0.0	0	0.0
COI-Millar-2.1.1	0	-	0	0.0	0	0.1
COI-LLR	301	10.0	301	0.3	301	0.7
COI-Azaiez	0	600.0	0	0.5	0	1.1
COI-GPost	5	-	12	0.2	12	1.2
COI-GPost-B	3	-	10	0.3	9	1.2
COI-QMC-1	13	-	30	0.6	26	1.2
COI-QMC-2	29	-	35	0.9	32	1.7
COI-WHPP	5	-	4008	3.0	4008	4.7
COI-BCV-3.46.2	894	17840.0	895	2.2	895	4.7
COI-BCV-4.13.1	10	-	10	0.1	10	0.3
COI-SINTEF	0	-	7	0.4	4	0.7
COI-ORTEC01	270	105.0	410	2.9	361	7.9
COI-ORTEC02	270	-	500	3.1	500	6.8
COI-ERMGH	779	124.0	1424	300.6	1359	300.9
COI-CHILD	149	-	1938	108.6	1713	160.2
COI-ERRVH	2001	-	9815	300.7	9817	300.8
COI-HED01	136	-	250	1.7	184	4.6
COI-Valouxis-1	20	-	180	0.4	180	0.9
COI-Ikegami-2.1	0	13.0	124	0.7	19	1.5
COI-Ikegami-3.1	2	21600.0	345	1.7	58	3.7
COI-Ikegami-3.1.1	3	2820.0	152	1.4	50	3.7
COI-Ikegami-3.1.2	3	2820.0	258	1.6	161	4.0
COI-BCDT-Sep	100	-	560	1.1	440	2.6
COI-MER	7081	36002.7	16177	300.6	14738	301.1
Average	454		1408	38.3	1298	41.4

At the foot of each table is a row of averages. This is meaningful for running times, but for costs it is sensitive to the scale of the constraint weights. For example, in instance `COI-Millar-2.1`, the weights are all multiples of 100; and constraints that measure workload in minutes rather than shifts may produce costs in the thousands. So average cost is useful mainly for instances with similar weights, as in the First International Nurse Rostering Competition.

4.1 The Curtois original instances

The Curtois original instances are the instances available online at [8] under the heading ‘Original instances.’ Their XESTT versions appear in XESTT archive file `COI.xml` at [20], along with the solutions posted with the instances.

This author does not know where these solutions came from originally. Most of them have the same costs as the solutions reported for the branch and

Table 2 Results for the Long and Medium instances of the First International Timetabling Competition. The GOAL column shows the solutions from the GOAL research group web site [31], which the GOAL group has shown to be virtually optimal. The other columns show the results from the two versions of the author’s solver. This table is derived from XESTT archive file KHE18-2018-08-26-INRC1-Long-And-Medium.xml, available at [20].

Instance	GOAL		KHE18		KHE18x8	
	Cost	Time	Cost	Time	Cost	Time
INRC1-L01	197	-	206	10.0	201	30.2
INRC1-L02	219	-	238	8.6	228	20.6
INRC1-L03	240	-	241	9.6	241	19.6
INRC1-L04	303	-	306	13.1	306	27.3
INRC1-L05	284	-	288	17.6	287	39.8
INRC1-LH01	346	-	465	14.7	405	67.5
INRC1-LH02	89	-	111	36.2	103	61.9
INRC1-LH03	38	-	48	5.6	48	21.8
INRC1-LH04	22	-	41	7.2	39	17.7
INRC1-LH05	41	-	94	8.3	69	24.9
INRC1-LL01	235	-	332	19.3	316	33.9
INRC1-LL02	229	-	322	14.0	305	36.7
INRC1-LL03	220	-	312	20.3	305	36.8
INRC1-LL04	222	-	350	15.0	311	41.5
INRC1-LL05	83	-	133	7.4	128	20.0
INRC1-M01	240	-	249	3.5	245	8.3
INRC1-M02	240	-	248	3.2	245	7.6
INRC1-M03	236	-	244	3.9	240	9.7
INRC1-M04	237	-	244	3.6	244	8.0
INRC1-M05	303	-	317	5.6	313	11.3
INRC1-MH01	111	-	190	5.2	141	10.5
INRC1-MH02	221	-	278	3.8	244	15.7
INRC1-MH03	34	-	42	3.1	42	7.4
INRC1-MH04	78	-	105	5.3	86	13.8
INRC1-MH05	119	-	156	4.5	150	9.9
INRC1-ML01	157	-	192	2.0	182	5.3
INRC1-ML02	18	-	33	1.1	32	3.6
INRC1-ML03	29	-	42	1.3	42	3.2
INRC1-ML04	35	-	55	2.1	49	5.7
INRC1-ML05	107	-	162	3.4	147	11.3
Average	164		201	8.6	190	21.1

price algorithm in Table 3 of [4], which also contains lower bounds showing that nearly all of them are optimal. However, their running times, where recorded in the solution files, are different. They may come from several solvers, in which case caution is needed in comparing their running times with KHE18’s.

Caution is also needed with costs. For example, analysis shows that in every solution of COI-WHPP with cost less than 1000, some nurses must take night shifts only, and the rest non-night shifts only—hardly a real-world scenario.

Table 1 compares KHE18’s solutions with those from [8]. It shows that, although KHE18x8 could not yet be said to be competitive, there are hopeful signs: the optimal or near-optimal solutions for about ten of the instances, and the moderate running times. More work is needed.

Table 3 Results for the Sprint instances of the First International Timetabling Competition. This table is derived from XESTT archive file `KHE18-2018-08-26-INRC1-Sprint.xml`, available at [20]. Other details as for Table 2.

Instance	GOAL		KHE18		KHE18x8	
	Cost	Time	Cost	Time	Cost	Time
INRC1-S01	56	-	58	0.9	58	1.9
INRC1-S02	58	-	60	0.7	60	2.0
INRC1-S03	51	-	54	1.0	54	1.6
INRC1-S04	59	-	65	0.8	62	1.7
INRC1-S05	58	-	60	0.6	59	1.9
INRC1-S06	54	-	57	0.9	56	1.5
INRC1-S07	56	-	60	0.5	59	1.7
INRC1-S08	56	-	60	1.2	60	1.8
INRC1-S09	55	-	59	0.7	59	1.9
INRC1-S10	52	-	55	0.4	55	1.5
INRC1-SH01	32	-	38	0.4	35	1.1
INRC1-SH02	32	-	38	0.4	37	0.9
INRC1-SH03	62	-	69	0.5	67	1.9
INRC1-SH04	66	-	79	1.2	74	2.0
INRC1-SH05	59	-	67	0.7	66	1.2
INRC1-SH06	130	-	178	0.3	178	1.2
INRC1-SH07	153	-	170	0.3	170	1.3
INRC1-SH08	204	-	237	1.5	219	2.8
INRC1-SH09	338	-	359	1.2	359	2.4
INRC1-SH10	306	-	320	0.9	319	1.5
INRC1-SL01	37	-	45	0.8	43	1.4
INRC1-SL02	42	-	49	0.6	47	1.2
INRC1-SL03	48	-	52	0.9	52	2.0
INRC1-SL04	73	-	89	0.5	89	1.6
INRC1-SL05	44	-	49	0.9	49	1.8
INRC1-SL06	42	-	45	0.6	44	1.3
INRC1-SL07	42	-	53	0.4	52	1.0
INRC1-SL08	17	-	21	0.2	19	0.5
INRC1-SL09	17	-	20	0.1	19	0.4
INRC1-SL10	43	-	56	0.3	50	0.9
Average	78		87	0.7	86	1.5

4.2 The First International Nurse Rostering Competition

The First International Nurse Rostering Competition instances are available from the competition web site [12]. Their converted versions appear in files `INRC1-Long-And-Medium.xml` and `INRC1-Sprint.xml` [20], with the GOAL research group's solutions [31], proved by GOAL to be virtually optimal.

The results appear in Tables 2 and 3. These tables do not give running times for the GOAL solutions, because no running times appear in the solution files downloaded from the GOAL site. However, the site itself contains a table that does give running times. About 10 of the instances, from the Long and Medium sets, have running times of about 4 hours. Many others, including all the Sprint instances, have running times under one minute, often well under.

Table 4 Solution costs for the instances published in 2014 by Curtois and Qu [7,8]. The CQ-EJ10, CQ-EJ60, CQ-BP, and CQ-GR columns show the four sets of solutions from XESTT archive `CQ14.xml`, corresponding to the ejection chain (10 mins), ejection chain (60 mins), Branch and Price, and Gurobi 5.6.3 columns of Table 2 of [7]. The last two columns show the results from the two versions of the author’s solver. This table is derived from XESTT archive file `KHE18-2018-08-26-CQ14.xml`, available at [20].

Instance	CQ-EJ10	CQ-EJ60	CQ-BP	CQ-GR	KHE18	KHE18x8
CQ14-01	1114		607	607	608	608
CQ14-02	1019	837		828	937	834
CQ14-03	1001	1003		1001	1111	1105
CQ14-04	1716	1718	1716	1716	2327	1924
CQ14-05	1144	1358	1160	1143	1651	1551
CQ14-06	1952	2258	1952	1950	inf.	2659
CQ14-07	1056	1269	1058	1056	1800	1390
CQ14-08	2133	1314	1308	1323	2581	2581
CQ14-09	1449	439		439	667	572
CQ14-10	4870	4631		4631	5198	4998
CQ14-11	3443	3661		3443	inf.	4196
CQ14-12	5476	4040	4046	4040	4747	4448
CQ14-13	11432	1486		1388	2595	2595
CQ14-14	2286	1300		1280	2114	2114
CQ14-15	12050	4378		4039	6916	6916
CQ14-16	3926	3225	3323	3233	inf.	inf.
CQ14-17	7801	5872		5851	8347	8347
CQ14-18	10002	4969		4760	inf.	inf.
CQ14-19	14788	3715		3218	5514	5514
CQ14-20					inf.	inf.
CQ14-21					inf.	inf.
CQ14-22					inf.	inf.
CQ14-23		44819		17428	inf.	inf.
CQ14-24				48777	inf.	inf.
Average						

Another source of virtually optimal solutions to these instances is the branch and price algorithm whose results are reported in Table 5 of [4]. The author has not tried to obtain these solutions. Their reported running times are better than the GOAL ones, never exceeding about 10 minutes.

The results on these instances are quite good. For example, the KHE18x8 results on INRC1-L01 to INRC1-L05 have near-optimum costs, with quite moderate running times. This is competitive with [4] and [31].

4.3 The 2014 Curtois and Qu instances

The instances tested here are the 24 instances published in 2014 by Curtois and Qu [7,8]. Converted versions appear in file `CQ14.xml`, which also holds four sets of solutions received from Curtois via private correspondence.

These four sets of solutions were made by the solvers reported on in Table 2 of [7], namely ejection chains (10 mins), ejection chains (60 minutes), branch and price, and Gurobi. However, they differ from the solutions presented in

Table 5 Running times in seconds for the instances published in 2014 by Curtois and Qu. Details as for Table 4, only showing running times instead of solution costs.

Instance	CQ-EJ10	CQ-EJ60	CQ-BP	CQ-GR	KHE18	KHE18x8
CQ14-01	0.0		0.4	-	0.1	0.3
CQ14-02	30.0	3600.0		-	0.4	1.2
CQ14-03	2.6	3600.0		-	0.6	1.8
CQ14-04	8.2	3600.0	1.5	-	0.6	1.5
CQ14-05	76.8	3600.0	25.6	-	1.5	2.5
CQ14-06	88.8	3600.0	10.5	-	2.1	4.7
CQ14-07	262.1	3600.0	93.7	-	4.7	9.7
CQ14-08	600.0	3600.0	11831.1	-	6.9	15.5
CQ14-09	0.6	635.2		-	24.6	50.7
CQ14-10	600.0	863.2		-	6.6	15.9
CQ14-11	171.4	3600.0		-	8.4	19.4
CQ14-12	600.0	1526.1	1336.4	-	44.3	87.5
CQ14-13	30.1	3600.1		-	116.9	215.9
CQ14-14	600.0	3600.3		-	9.5	32.1
CQ14-15	30.1	3600.0		-	33.5	65.8
CQ14-16	30.0	2965.7	265.0	-	62.1	98.9
CQ14-17	600.0	3600.1		-	96.0	216.6
CQ14-18	30.0	3600.0		-	120.5	258.9
CQ14-19	30.0	3600.0		-	84.6	152.0
CQ14-20					300.2	300.4
CQ14-21					300.6	301.0
CQ14-22					300.4	301.0
CQ14-23		3601.0		-	301.1	301.9
CQ14-24				-	303.0	306.7
Average					88.7	115.1

Table 2 of [7], for reasons explained in [21]: the solutions received were not exactly the ones reported on in the table, and some had to be omitted owing to extreme overstaffing of shifts, which is strictly speaking legal but which this author chose not to allow in the XESTT versions of the instances.

Another source of solutions to these instances is [10]. They are not competitive, so adding them to CQ14.xml has not been a priority.

The results appear in Tables 4 and 5. The author has only just begun to tackle these instances, and the results are not competitive, although there is promise in the generally modest running times, and in the fact that feasible solutions have been found for 17 of the 24 instances which are often competitive with the Curtois ejection chains (10 mins) solutions. The later instances are very large, and even finding a feasible solution is challenging.

5 Conclusion

This paper has presented KHE18, a solver for nurse rostering which aims to produce very good but not optimal solutions quickly across a wide range of instances. Polynomial-time methods are used, including time sweep for the initial assignment, and ejection chains for repair.

This paper reports on work in progress. At present, KHE18 is producing some very good solutions, but overall its results are not competitive, except in running time. What is needed now is more analysis of its behaviour, leading, hopefully, to new insights which can be incorporated into the algorithm.

One can see from the leading recent papers [4,30] that provably optimal solutions to realistic instances are becoming available using methods based on integer programming. There is still a role for algorithmic methods, however, in finding good solutions for large instances with reliably fast running times. It is not clear to this author that *very* large instances like the last few of the 2014 Curtois and Qu instances really do need to be solved in the real world. However, if they do, that might be another role for algorithmic methods.

References

1. J. L. Arthur and A. Ravindran A multiple objective nurse scheduling model, *AIIE Transactions* 13(1), pages 55–60 (1981).
2. Shahriar Asta and Ender Özcan, A tensor-based approach to nurse rostering, In Ender Özcan, Edmund Burke, and Barry McCollum (eds.), *PATAT 2014 (Tenth international conference on the Practice and Theory of Automated Timetabling, York, UK, August 2014)*, 442–445 (2014)
3. Edmund K. Burke, Tim Curtois, Rong Qu, and Greet Vanden Berghe, A time predefined variable depth search for nurse rostering, *INFORMS Journal on Computing*, 25(3), 411419 (2013), accessed via <http://eprints.nottingham.ac.uk/28283/1/JOC12vds.pdf>
4. Edmund K. Burke and Tim Curtois, New approaches to nurse rostering benchmark instances, *European Journal of Operational Research* 237, 71–81 (2014)
5. Sara Ceschia, Nguyen Thi Thanh Dang, Patrick De Causmaecker, Stephaan Haspeslagh, and Andrea Schaerf, Second international nurse rostering competition (INRC-II), problem description and rules. oRR abs/1501.04177 (2015). URL <http://arxiv.org/abs/1501.04177>
6. Sara Ceschia, Nguyen Thi Thanh Dang, Patrick De Causmaecker, Stephaan Haspeslagh, and Andrea Schaerf, Second international nurse rostering competition (INRC-II) web site, URL <http://mobiz.vives.be/inrc2/>.
7. Tim Curtois and Rong Qu, Computational results on new staff scheduling benchmark instances URL <http://www.cs.nott.ac.uk/~psztc/NRP/> (2014)
8. Tim Curtois, Employee Shift Scheduling Benchmark Data Sets, <http://www.cs.nott.ac.uk/~psztc/NRP/> (2016)
9. Nguyen Thi Thanh Dang, Sara Ceschia, Andrea Schaerf, Patrick De Causmaecker, and Stefaan Haspeslagh, Solving the multi-stage nurse rostering problem, *PATAT 2016 (Eleventh international conference on the Practice and Theory of Automated Timetabling, Udine, Italy, August 2016)*, 473–475 (2016)
10. Emir Demirović, Nysret Musliu, and Felix Winter, Modeling and solving staff scheduling with partial weighted maxSAT, *PATAT 2016 (Eleventh international conference on the Practice and Theory of Automated Timetabling, Udine, Italy, August 2016)*, 109–125 (2016)
11. Kathryn A. Dowsland, Nurse scheduling with tabu search and strategic oscillation, *European Journal of Operational Research* 106, 393–407 (1998)
12. Stefaan Haspeslagh, Patrick De Causmaecker, Martin Stlevik, and Andrea Schaerf, First international nurse rostering competition website, URL: <http://www.kuleuven-kortrijk.be/nrppcompetition> (2010)
13. Stefaan Haspeslagh, Patrick De Causmaecker, Martin Stlevik, and Andrea Schaerf, The first international nurse rostering competition 2010, *Annals of Operations Research*, 218, 221–236 (2014)
14. Han Hoogeveen and Tim van Weelden, Personalized nurse rostering through linear programming, In Ender Özcan, Edmund Burke, and Barry McCollum (eds.), *PATAT 2014 (Tenth international conference on the Practice and Theory of Automated Timetabling, York, UK, August 2014)*, 476–478 (2014)

15. Ahmed Kheiri, Ender Özcan, Rhyd Lewis, and Jonathan Thompson, A sequence-based selection hyper-heuristic: a case study in nurse rostering, PATAT 2016 (Eleventh international conference on the Practice and Theory of Automated Timetabling, Udine, Italy, August 2016), 503–505 (2016)
16. Jeffrey H. Kingston, The HSEval High School Timetable Evaluator, URL <http://jeffreykingston.id.au/cgi-bin/hseval.cgi> (2010)
17. Jeffrey H. Kingston, Repairing high school timetables with polymorphic ejection chains, *Annals of Operations Research*, DOI 10.1007/s10479-013-1504-3
18. Jeffrey H. Kingston, KHE14: An algorithm for high school timetabling, In Ender Özcan, Edmund Burke, and Barry McCollum (eds.), PATAT 2014 (Tenth international conference on the Practice and Theory of Automated Timetabling, York, UK, August 2014), 269–291
19. Jeffrey H. Kingston, KHE web site, <http://jeffreykingston.id.au/khe> (2014)
20. Jeffrey H. Kingston, XESTT web site, <http://jeffreykingston.id.au/xestt> (2017)
21. Jeffrey H. Kingston, Gerhard Post, and Greet Vanden Berghe, A unified nurse rostering model based on XHSTT, PATAT 2018 (Twelfth international conference on the Practice and Theory of Automated Timetabling, Vienna, August 2018)
22. Jeffrey H. Kingston, Modelling history in nurse rostering, PATAT 2018 (Twelfth international conference on the Practice and Theory of Automated Timetabling, Vienna, August 2018)
23. M. J. Louw, I. Nieuwoudt, and J. H. Van Vuuren, Finding good nursing duty schedules: a case study. Technical report, Department of Applied Mathematics, Stellenbosch University, South Africa (2005). Received from Tim Curtois and held by this author.
24. Antonn Novák, Roman Václavk, Premysl Sucha, and Zdenek Hanzálek, Nurse rostering problem: tighter upper bound for pricing problem in branch and price approach, MISTA 2015 (7th Multidisciplinary International Conference on Scheduling: Theory and Applications), Prague, August 2015, 759–763
25. Gerhard Post, XHSTT web site, <http://www.utwente.nl/ctit/hstt/> (2011)
26. Samad Ahmadi, Sophia Daskalaki, Jeffrey H. Kingston, Jari Kyngäs, Cimmo Nurmi, Gerhard Post, David Ranson, and Henri Ruizenaar, An XML format for benchmarks in high school timetabling, *Annals of Operations Research*, 194, 385–397 (2012)
27. Florian Mischek and Nysret Musliu, Integer programming and heuristic approaches for a multi-stage nurse rostering problem, PATAT 2016 (Eleventh international conference on the Practice and Theory of Automated Timetabling, Udine, Italy, August 2016), 245–262 (2016)
28. Christopher Rae and Nelishia Pillay, Investigation into an evolutionary algorithm hyper-heuristic for the nurse rostering problem, In Ender Özcan, Edmund Burke, and Barry McCollum (eds.), PATAT 2014 (Tenth international conference on the Practice and Theory of Automated Timetabling, York, UK, August 2014), 527–532 (2014)
29. Erfan Rahimian, Kerem Akartunali, and John Levine, A Hybrid Constraint Integer Programming Approach to Solve Nurse Scheduling Problems, MISTA 2015 (7th Multidisciplinary International Conference on Scheduling: Theory and Applications), Prague, August 2015, 429–442
30. Haroldo G. Santos, Túlio A. M. Toffolo, Rafael A. M. Gomes, and Sabir Ribas, Integer programming techniques for the nurse rostering problem, *Annals of Operations Research* 239, 225–251 (2016)
31. Haroldo G. Santos, Túlio A. M. Toffolo, Rafael A. M. Gomes, and Sabir Ribas, <http://www.goal.ufop.br/nrp/>
32. Pieter Smet and Greet Vanden Berghe, Variable neighbourhood search for rich personnel rostering problems, MISTA 2015 (7th Multidisciplinary International Conference on Scheduling: Theory and Applications), Prague, August 2015, 928–930
33. J. Van den Bergh, J. Belien, P. De Bruecker, E. Demeulemeester and L. De Boeck, Personnel scheduling: a literature review, *European Journal of Operational Research*, 226(3), 367–385, (2013).