KHE26: Progress towards an Efficient, Practical Nurse Rostering Solver (Abstract)

Jeffrey H. Kingston

School of Information Technologies, The University of Sydney, Australia
jeff@it.usyd.edu.au
http://jeffreykingston.id.au

Abstract. This paper presents ongoing work on a solver for the nurse rostering problem which is oriented towards practice, in that it aims to find good solutions quickly for a wide range of instances. This abstract shows that the current version of the solver, KHE26, has improved significantly on the previous version on one important data set, of four-week instances from the Second International Nurse Rostering Competition. Future work will present results for this as well as several other data sets, and will explain how these improvements have been obtained without fundamental changes to the solver.

Keywords: Automated timetabling · Nurse rostering

1 Introduction

Nurse rostering, the problem of assigning nurses to shifts in hospital wards, is one of the most-studied problems in automated timetabling.

In a previous paper [7], this author presented KHE24, an open source solver [5] for nurse rostering. KHE24 reads and writes instances and solutions in the XESTT data format [6], in which all of the well-known nurse rostering data sets have been expressed, plus solutions from various sources which are among the best known.

This paper presents KHE26, the current version of that solver. Rather than trying to find new best solutions, KHE24 and KHE26 try to find solutions whose cost is within 10% of the best known, given a 5 minute run on ordinary desktop hardware, and within 5% of the best known, given a 60 minute run on ordinary desktop hardware. If this can be done for a wide range of real-world instances, then the solver will be more useful in practice than solvers that are targeted to specific kinds of instances, or that run for many hours.

The work reported by the previous paper was avowedly incomplete. KHE24 performed well on two data sets: the Curtois original instances [3] and the First International Nurse Rostering Competition [4] instances. But it performed poorly on a standard set of four-week instances from the Second International Nurse Rostering Competition [1, 2]. Table 4 from [7] showed that, given 5 minutes per instance, KHE24 produced solutions for those instances which were on average 34% worse than the best known solutions.

KHE26, the version of the solver presented here, is much improved. Table 1 presents KHE26's results for the same set of four-week instances from the Second International Nurse Rostering Competition. It shows that, given 5 minutes per instance, KHE26 is producing solutions which are on average 10% worse than the best known solutions. While this is not mission accomplished, because 10% on average is not the same as 10% for each instance, it is a significant improvement on 34% and is approaching the desired outcome.

In addition to presenting results on other data sets, future work will explain how this improvement was obtained from a solver which, in outline at least, has not greatly changed. But, in brief, it is the result of detailed analysis of the solutions found by the solver, aiming to identify weaknesses in the algorithm. As these are corrected, small improvements in solution quality accumulate.

References

1. Sara Ceschia, Nguyen Thi Thanh Dang, Patrick De Causmaecker, Stefaan Haspeslagh, and Andrea Schaerf, Second international nurse rostering competition (INRC-II), problem description and rules. oRR abs/1501.04177 (2015). http://arxiv.org/abs/1501.04177
2. 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/.
3. Tim Curtois, Employee Shift Scheduling Benchmark Data Sets, http://www.schedulingbenchmarks.org/ (2019)
4. Stefaan Haspeslagh, Patrick De Causmaecker, Martin Stølevik, and Andrea Schaerf, First international nurse rostering competition website, URL: http://www.kuleuven-kortrijk.be/nrpcompetition (2010)
5. Jeffrey H. Kingston, KHE web site, http://jeffreykingston.id.au/khe (2014).
6. Jeffrey H. Kingston, XESTT web site, http://jeffreykingston.id.au/xestt (2017)
7. Jeffrey H. Kingston, KHE24: Towards a Practical Solver for Nurse Rostering (Abstract), PATAT 2024 (Fourteenth International Conference on Practice and Theory of Automated Timetabling, Copenhagen, Denmark, August 2024).
8. A. Legrain, J. Omer, and S. Rosat, A rotation-based branch-and-price approach for the nurse scheduling problem (2017). Working paper, available at https://hal.archives-ouvertes.fr/hal-01545421.

Table 1. KHE26 results on 20 standard four-week instances from the Second International Nurse Rostering Competition [1, 2], obtained on 26 February 2026. For each instance, column LOR17 shows the cost of the best known solution, as reported in [8] and obtained from the authors of that paper by this author. The running time is reported in [8] to conform to the competition rules but is not present in the file obtained from its authors, so does not appear here; it is about 20 minutes for the largest instances. Column KHE26x12 shows the cost of the best solution obtained by KHE26 over 12 simultaneous 5-minute runs, each with a different random seed, on the author's 12-core desktop computer (the same computer as the one used for testing in the author's KHE24 paper); the ratio of that cost to the cost of the best known solution; and the running time in seconds of the whole run of 12 simultaneous solves (that is, the wall clock time). Choosing the best of several solutions is an integral part of the KHE26 approach, not something that is done just for testing. It is a simple way to exploit the power of multi-core computers, which are now normal on desktops. The XESTT archive file from which this table is derived may be found online, at `http://jeffreykingston.id.au/tt_solns/KHE26-2026-02-26-INRC2-4.xml.gz`. This file contains the actual solutions, not just their costs.

Instances (20)	**LOR17**			**KHE26x12**		
	Cost	**Rel.**	**Time**	**Cost**	**Rel.**	**Time**
INRC2-4-030-1-6291	**1695**	1.00	-	1780	1.05	283.8
INRC2-4-030-1-6753	**1890**	1.00	-	1960	1.04	294.4
INRC2-4-035-0-1718	**1425**	1.00	-	1525	1.07	300.1
INRC2-4-035-2-8875	**1155**	1.00	-	1270	1.10	285.6
INRC2-4-040-0-2061	**1685**	1.00	-	1845	1.09	300.2
INRC2-4-040-2-6106	**1890**	1.00	-	2030	1.07	300.1
INRC2-4-050-0-0487	**1505**	1.00	-	1615	1.07	300.4
INRC2-4-050-0-7272	**1500**	1.00	-	1520	1.01	300.2
INRC2-4-060-1-6115	**2505**	1.00	-	2830	1.13	300.4
INRC2-4-060-1-9638	**2750**	1.00	-	3105	1.13	300.2
INRC2-4-070-0-3651	**2435**	1.00	-	2700	1.11	300.5
INRC2-4-070-0-4967	**2175**	1.00	-	2475	1.14	300.5
INRC2-4-080-2-4333	**3340**	1.00	-	3665	1.10	300.6
INRC2-4-080-2-6048	**3260**	1.00	-	3660	1.12	300.4
INRC2-4-100-0-1108	**1245**	1.00	-	1410	1.13	300.5
INRC2-4-100-2-0646	**1950**	1.00	-	2230	1.14	300.4
INRC2-4-110-0-1428	**2440**	1.00	-	2640	1.08	300.3
INRC2-4-110-0-1935	**2560**	1.00	-	2820	1.10	300.4
INRC2-4-120-1-4626	**2170**	1.00	-	2525	1.16	300.3
INRC2-4-120-1-5698	**2220**	1.00	-	2575	1.16	300.3
Average	**2090**	1.00		2309	1.10	298.5