

Continuity of care in community midwifery

John Bowers · Helen Cheyne · Gillian Mould ·
Miranda Page

Received: 17 June 2013 / Accepted: 13 May 2014
© Springer Science+Business Media New York 2014

Abstract Continuity of care is often critical in delivering high quality health care. However, it is difficult to achieve in community health care where shift patterns and a need to minimise travelling time can reduce the scope for allocating staff to patients. Community midwifery is one example of such a challenge in the National Health Service where postnatal care typically involves a series of home visits. Ideally mothers would receive all of their antenatal and postnatal care from the same midwife. Minimising the number of staff-handovers helps ensure a better relationship between mothers and midwives, and provides more opportunity for staff to identify emerging problems over a series of home visits. This study examines the allocation and routing of midwives in the community using a variant of a multiple travelling salesmen problem algorithm incorporating staff preferences to explore trade-offs between travel time and continuity of care. This algorithm was integrated in a simulation to assess the additional effect of staff availability due to shift patterns and part-time working. The results indicate that continuity of care can

be achieved with relatively small increases in travel time. However, shift patterns are problematic: perfect continuity of care is impractical but if there is a degree of flexibility in the visit schedule, reasonable continuity is feasible.

Keywords Home health care · Staff scheduling · Routing · Simulation

1 Background

a. Continuity of care and maternity services

Across many areas of health care, continuity of care is often cited as a major element in patient satisfaction, improving communications and helping to establish greater trust and a willingness to confide in the professional carer. Health professionals also value continuity of care; it can increase their sense of responsibility and tendency to act as an advocate for the patient [1]. Continuity of care can improve both the effectiveness of health care and its efficiency with fewer staff-handovers and more consistency in the care provided. While the benefits of continuity of care are dependent on the particular service, it can have particular value in community-based care. Greater continuity of care can reduce the demand for additional services, reduce hospital admissions and increase patient satisfaction [2, 3]. In maternal health services women receive care at various stages: antenatal, intrapartum and postnatal, both in the community and in hospital. This can lead to a disrupted pattern of care with a mother experiencing many different professional staff [4].

Continuity of care has many interpretations [1]. “Information continuity” emphasises the need to ensure effective communications between health care professionals, while “longitudinal” and “relational” continuity focus on the patient having interactions with as few staff as

J. Bowers (✉)
Stirling Management School, University of Stirling, Stirling FK9
4LA, UK
e-mail: j.a.bowers@stir.ac.uk

H. Cheyne
Nursing, Midwifery and Allied Health Professions’ Research Unit,
University of Stirling, Stirling FK9 4LA, UK
e-mail: h.l.cheyne@stir.ac.uk

G. Mould
Stirling Management School, University of Stirling, Stirling FK9
4LA, UK
e-mail: g.i.mould@stir.ac.uk

M. Page
Nursing, Midwifery and Allied Health Professions’ Research Unit,
University of Stirling, Stirling FK9 4LA, UK
e-mail: m.e.page@stir.ac.uk

possible, and sufficient time to develop a useful relationship. This study focuses on longitudinal continuity of care and in particular an objective of minimising the number of staff-handovers. In the United Kingdom's National Health Service (NHS) the delivery of maternity care varies in different areas. However, a mother usually attends community based antenatal clinics, receiving advice and care from a midwife before giving birth in hospital, or midwifery led unit. Most mothers are then discharged between 6 and 48 h after giving birth, followed by four or five postnatal care visits at home. Postnatal visits serve many purposes: specific medical checks of the mother and baby; parenting and feeding advice; identifying the need for additional support due to social conditions or the mother's mental health. Continuity of care helps develop a trusting relationship with the mother and the midwife can monitor the more tacit signs of changing behaviour over a number of visits.

Various models of care have been proposed to maximise continuity of care from pregnancy to 28 days following birth [4]. Team-midwifery provides co-ordinated care from a number of staff; information-continuity can be provided but relational-continuity is impaired. Caseload-midwifery, where women receive individualized care, ideally from a single midwife, offers comprehensive continuity but this model of care is very difficult to deliver in practice.

While antenatal care can be planned around predetermined stages of pregnancy the unpredictable nature of birth makes complete continuity through all stages of maternity care problematic. It was suggested that longitudinal, relational continuity of care was impractical even for just the antenatal and postnatal stages, given the difficulties of staff availability and the need to avoid excessive travel. The objective of the current study was to develop the understanding of the logistical and resource implications of delivering continuity of care, as part of a larger programme examining the design of postnatal care in the NHS in the United Kingdom. Even if 100 % continuity of care is an unrealistic target, is it feasible to achieve a reasonably high degree of continuity without incurring substantial extra travelling costs? The study examines a typical area in an attempt to determine whether continuity of care is feasible in community midwifery.

b. A home health care problem

The study is an example of the Home Health Care Problem (HHCP) [5, 6], which can be categorised as a variant of the Multiple Travelling Salesman Problem [7]. In most HHCP studies, the focus is to minimise the total staff travelling time while respecting time windows. The current study does not consider time windows: at present the common practice in community care in the NHS is to organise home visits assuming that the patients have no

significant time constraints. Instead this study emphasises the desire for continuity of care. Other HHCP studies have considered the allocation and routing of single sets of home care visits; the current study is concerned with providing a series of visits for each patient over a number of days, striving to maximise continuity of care, while minimising the total travel time. Thus a patient's preference for a particular member of staff on a particular visit i is determined by their allocation of staff for visit $i - 1$. Other HHCP studies have incorporated patient preferences, either as a constraint [8–10] or as a component of the objective function [11], but the preferences were modelled as being determined externally and part of the problem specification, independent of the allocation and routing solution. In the current study the problem specification for any one day is dependent on allocation decisions of the previous days, increasing the dimensionality of the HHCP. A study [9] of a more traditional vehicle routing problem developed a “record-to-record” algorithm to include a hard constraint that customers should have their orders delivered by the same driver over a series of days. However, in the current problem meeting the patient's preference is treated as a soft constraint, accepting that a need to avoid excessive journey times and to respect staff daily availability will often preclude perfect continuity of care. Patient preferences have been included as soft constraints in other studies, using penalty costs in the objective function. However, the problems examined had substantial hard constraints, notably strict time windows that severely limited the number of permutations, enabling the final routing and sequencing to be determined using enumeration [11].

The scale of the problem was further increased by the need to consider the stochastic nature of demand, over time and geography. The current study considers daily sets of requests, including both newly discharged mothers requiring their first postnatal visit and follow-up visits for other mothers. The daily problems can be classified as static and stochastic [12] and were explored with an allocation and routing algorithm embedded in a Monte-Carlo simulation framework, sampling the new mothers according to specified geographical distributions. In practice there are also some dynamic features to the routing problem [12] with mothers being absent and repeat visits having to be made. This additional complexity was omitted from the present study. An additional requirement for the algorithm was that the resultant simulation should provide a sufficiently rapid response to provide statistically robust results for a variety of experiments exploring various policy options and developing an appreciation of the critical trade-offs between efficiency and the quality of care as measured by the preference satisfaction. A similar

challenging requirement is found in designing operational systems. One system, LAPS CARE, which included continuity of care as a key quality objective, adopted a repeated matching algorithm, uses splitting to help avoid local optima [9]. As in other application-focussed HHCP studies [13], the choice of algorithm was strongly influenced by the requirements for speed of response and the flexibility to readily incorporate additional constraints, such as varying staff availability.

Other studies of the Multiple Travelling Salesmen Problem [14], and variants of the HHCP [13] in particular, have used algorithms based on the classic Clarke-Wright heuristics [15] to deliver speed and flexibility with only a small reduction in performance compared to optimal solutions [16]. This study adopted such an approach, developing an algorithm that was then incorporated in a discrete event simulation to explore the trade-offs between travel time and continuity of care. This provided a practical basis for examining policy issues, for example incorporating typical staff availability as dictated by their shift patterns when determining a possible target for continuity of care for community postnatal care. However, further research would be valuable, examining the viability of adapting more sophisticated branch-and-price algorithms [8] or genetic algorithms [17]. Even if such approaches cannot be readily incorporated in a simulation, they would provide a useful basis for comparison when assessing the performance of the modified Clarke-Wright algorithm, and identifying possible better choices.

The study just considers the staff travel time, though this will be strongly correlated to costs: the mileage expenses paid to midwives, their salaries and also the opportunity costs which could be significant in areas where there is a shortage of midwives.

c. Developing the model

The paper first considers a self-contained HHCP, developing a simple allocation and routing algorithm based on a modification of the Clarke-Wright [15] heuristics. This is illustrated in a small example, demonstrating the possible trade-off between travel time and satisfying preferences. The algorithm is embedded in a simulation, incorporating some of the basic characteristics of organising community care. This simulation is then refined to include more of the features of community postnatal care, notably that the preferences are generated by the desire to have the same midwife visiting a mother on every occasion, or at least striving for a high level of continuity of care. Later sections introduce more refinements reflecting the problems of shift patterns and part-time working, and also the value of flexibility in visits rather than having to keep to a strict sequence.

2 The travelling nurses' problem

a. A variant of the multiple travelling salesmen problem

The trade-offs between continuity of care and travelling time require a reasonably efficient staff allocation and routing algorithm. The problem could be described as a variant of the multiple travelling salesman problems with load balancing [7]. The multiple travelling salesmen problem considers the allocation of patients to staff selected from a team with the objective of minimising total travel time while considering the relative workloads. The current problem is distinguished by a patient's preference for specific staff, notably the staff that provided the care on the previous visit. It may be acceptable to have an anonymous salesman visit a customer but in health services it is important to have the same trusted staff: the travelling nurses' problem introduces additional allocation and routing challenges.

The objective of the current study was to inform the design of postnatal care. A model was required that could be used to explore various policies and undertake sensitivity analyses of key design parameters to develop a better understanding of the implications of continuity of care. The major requirement was for an algorithm that could be readily incorporated in a simulation model. While the algorithm had to provide good solutions to the travelling nurses' problem, speed and ease of use were also necessary. The approach adopted in this study was based on the classic multiple travelling salesmen problem algorithm [15]. The Clarke-Wright algorithm assumes an initial solution with simple journeys directly to and returning from each of the locations to the central depot. The possibility of introducing single new links is considered, identifying the potential savings in travel time. A "greedy" approach is adopted, selecting the new link offering the largest saving, repeating the process until all of the requests have been satisfied. While the Clarke-Wright algorithm does not always provide optimal solutions, it offers good solutions [16], speed and a flexibility that have been exploited in many variants of the capacitated vehicle routing problem [13, 14]. This study modified the Clarke-Wright algorithm, including an additional component in the "saving" to reflect the benefit of the chosen staff matching the preference of the particular patient. This additional "preference saving" (PS) reflects the relative priorities for minimising travel time or maximising preference satisfaction. When $PS=0$ there is no benefit in matching staff to patients, with a small value of PS , there is some reallocation of staff to patients, satisfying preferences as long as no major diversions are required. As the value of PS is increased, there is greater emphasis on matching staff to patients' preferences. With large values of PS , greater than the mean travel time between patients' homes, the large majority of patients' preferences will be satisfied. When $PS \gg$ mean travel time between patients, the problem

degenerates into a set of independent sub-problems minimising the total travel time of each member of staff separately. Although heuristics based on the Clarke-Wright savings usually perform well, it is possible to identify test cases in which the algorithm offers far from optimal solutions. In some applications this failing in specific cases could be critical but is less important in a simulation-based study. However, further study might usefully explore alternative solutions and the ranges of validity of the current simple algorithm. Branch and bound or genetic algorithms may not be practical when analysing multiple sets of visit requests, as required in simulation experiments, but they would provide a valuable comparison to determine the range of validity for the simple algorithm adopted in this study.

b. The travelling nurses' algorithm

Consider a set of requests for patient visits, with each patient j having a preferred member of staff n_j . The travel time from patient j to patient k is c_{jk} , and c_{j0} is the time to travel from patient j to the depot or base station. Extending a member of staff's route, as in the Clarke Wright algorithm, typically involves adding a new, previous unallocated request j , to a request k at the start or end of a chain of links defining a staff member's sequence of visits. Adding a new link to the start of a sequence implies travel from the depot to patient j and then to the patient k , producing the standard Clarke-Wright saving of:

$$s_{jk} = c_{j0} + c_{0k} - c_{jk} \tag{1}$$

Similar savings may be achieved by adding new links to the end, rather than the start of the evolving chain of links, though it cannot always be assumed that the travel times are symmetrical. In the travelling nurses' problem, the saving is dependent on the particular staff v allocated. If the new link involves employing the preferred staff to visit request j , then a greater saving may be achieved:

$$s_{jkv} = c_{j0} + c_{0k} - c_{jk} + PS \quad \text{if } v = n_j \tag{2}$$

When identifying the initial link in the development of a staff member's route, it is possible that both previously unallocated requests j and k have the same preference and the saving is greater with a preference saving associated with both patients:

$$s_{jkv} = c_{j0} + c_{0k} - c_{jk} + 2PS \quad \text{if } v = n_j = n_k \tag{3}$$

Thus when PS is large, compared to the mean travel time between patients, the greatest savings will be associated with patients having identical staff preferences,

living close to each other but at a large distance from the base depot: such links will be prioritized when constructing staff routes.

The complete set of potential links is identified, with separate entries for each option corresponding to Eqs 1–3, and then ordered to produce a list S' . Each entry in S' , beginning with the one offering the greatest potential savings, is considered in turn, checking whether the new link is feasible, given the state of the evolving allocation and routing solution. Visits are added to each staff's sequence developing routes in parallel until all requests have been allocated, or no more staff-time is available. The algorithm is illustrated in Fig. 1; it involves multiple iterations, reconsidering some potential links since they may only become relevant as the routes evolve. The algorithm constructs routes in parallel for each member of staff; a parallel approach is often recommended for multiple vehicle routing problems [18]. When considering a potential link in S' there are two main requirements:

- Have the request(s) j and k already been satisfied; if not is one of them available at the start, or end, of an existing chain of links defining a staff's route?
- If the saving associated with the potential link is dependent on a particular member of staff, are the staff available? This assessment compares the allocated workloads with the specified shift-times and may also include workload balancing to ensure a reasonably even distribution across the staff.

c. Applying the algorithm

The application of the algorithm is illustrated in an analysis of a set of 10 requests, distributed over an area surrounding a central depot, with each request having a preference for one of three staff, as in Fig. 2. A simple workload balancing requirement is incorporated with a constraint that no member of staff should be allocated more than 4 patients. This assumes that all patients require the same staff-time, excluding travel. A more sophisticated algorithm could consider the complete workloads and the likely time commitments associated with each patient-visit, including the time for direct patient contact and indirect administrative tasks.

The algorithm calculates the modified Clarke-Wright savings using Eqs 1–3; assuming a relatively low priority for patients' preferences ($PS=10$) the first 16 of the ordered possible links and their potential savings, S' , are noted in Table 1. In order to avoid unnecessary calculations, S' , records just the possible links with distinct savings, noting any specific staff needed to achieve the specified saving; staff = 0 implies that any staff may be used and the saving is just associated with efficient routing rather than satisfying a patient's preference. In the

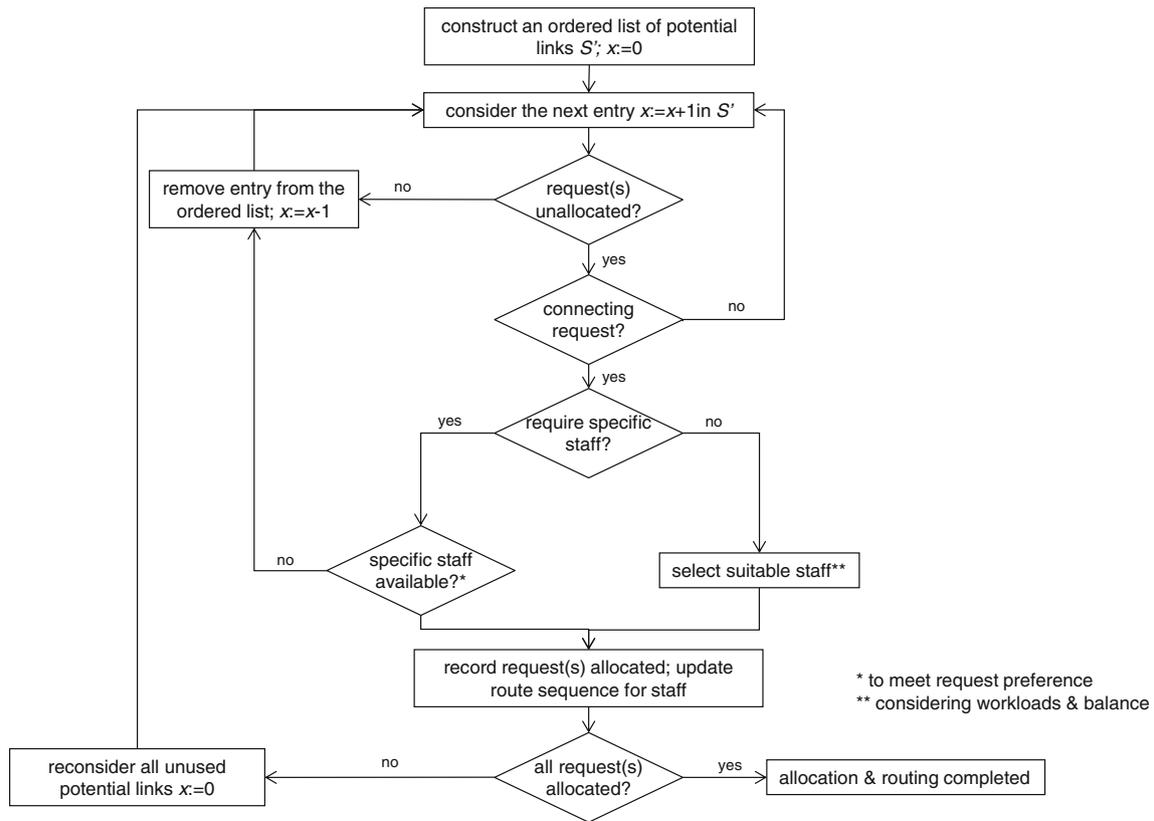


Fig. 1 Travelling Nurses' algorithm

example of Table 1, the most desirable link is that between patient 3 and 2, using staff 2. If this is an initial link for staff 2 with no connecting link required, i.e. as in Eq. (3), a total saving of 139 min is possible, including $2 \times 10 = 20$ min for satisfying the patients' preferences. Entry 3 relates to the possibility of adding the request from patient 3 to an existing sequence of visits for staff 2, connecting via patient 2. The additional saving, as in Eq. (2), is smaller since just one additional patients' preference is

satisfied. If the request from patient 3 is added to an existing sequence of visits for staff other than the one preferred (staff=0), as specified in entry 7, the saving is lower still, as in Eq. (1). Other entries 1–10 reflect the symmetrical possible links; in some examples the savings

Table 1 Extract of the order list of potential links S'

| | Request 1 | Request 2 | Connecting link | Staff | Savings |
|----|-----------|-----------|-----------------|-------|---------|
| 1 | 3 | 2 | 0 | 2 | 139 |
| 2 | 2 | 3 | 0 | 2 | 139 |
| 3 | 3 | 2 | 2 | 2 | 129 |
| 4 | 2 | 3 | 3 | 2 | 129 |
| 5 | 3 | 2 | 0 | 2 | 129 |
| 6 | 2 | 3 | 0 | 2 | 129 |
| 7 | 3 | 2 | 2 | 0 | 119 |
| 8 | 2 | 3 | 3 | 0 | 119 |
| 9 | 3 | 2 | 0 | 0 | 119 |
| 10 | 2 | 3 | 0 | 0 | 119 |
| 11 | 3 | 1 | 1 | 2 | 118 |
| 12 | 1 | 3 | 3 | 1 | 118 |
| 13 | 2 | 1 | 1 | 2 | 118 |
| 14 | 1 | 2 | 2 | 1 | 118 |
| 15 | 3 | 1 | 0 | 2 | 118 |
| 16 | 1 | 3 | 0 | 1 | 118 |

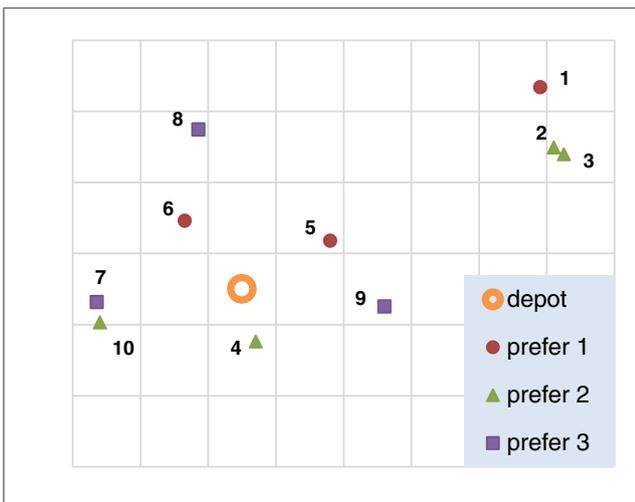


Fig. 2 Mapping a set of requests (10 requests; 3 staff)

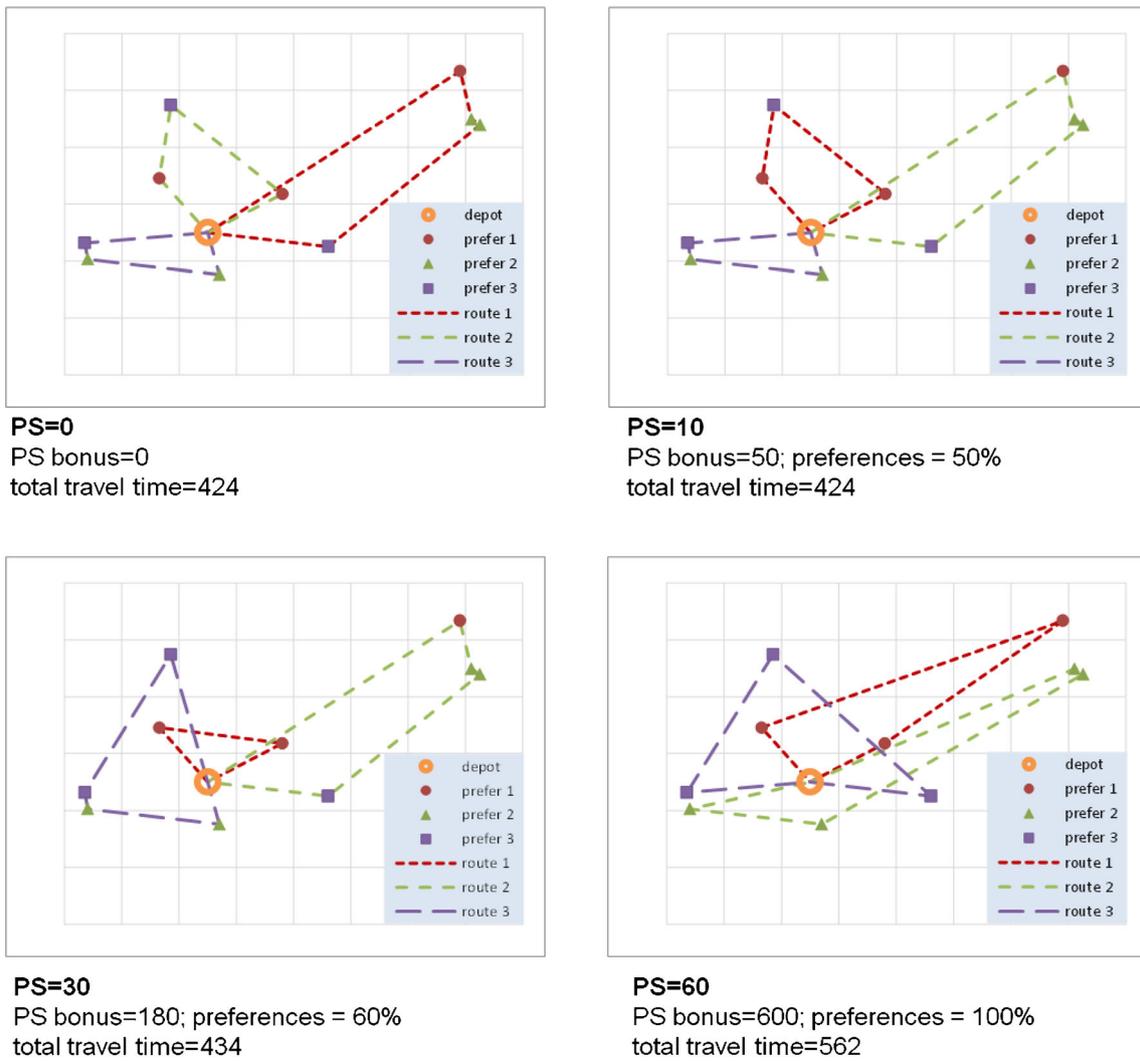


Fig. 3 Shifting the emphasis on patients' preferences (*PS* preference saving associated with satisfying a single patient's preference; *PS* bonus=total saving=no. patients with preferences satisfied×*PS*)

Fig. 4 Simulating postnatal care

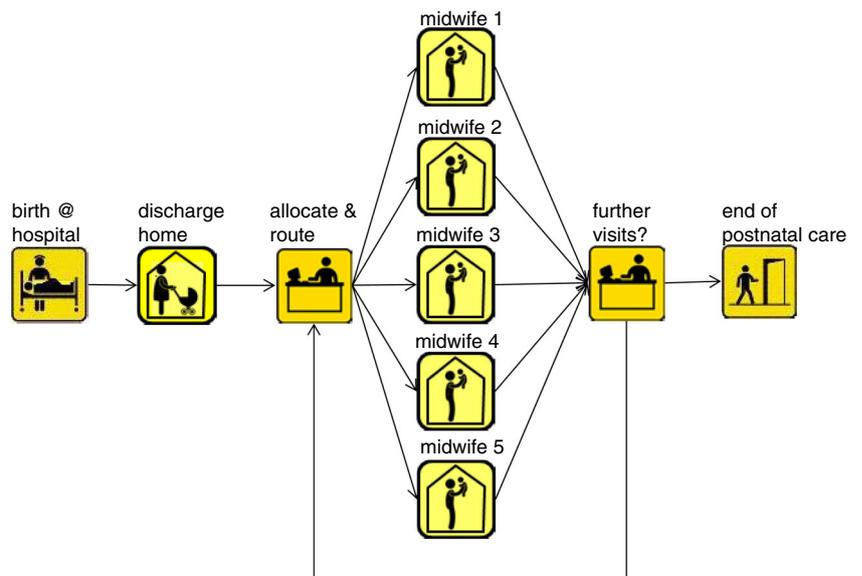


Table 2 Simulation experiments

| Experiment | Problem specification | Objective |
|----------------------|---|---|
| Travel & Preferences | Each mother requires a single home visit, with preferences generated randomly. | To understand the basic trade-off between satisfying preferences and total travel time. |
| Travel & Continuity | Each mother requires four home visits; the initial preference is random but the subsequent preferences are determined by the staff from the previous visit. | To assess the effect on the trade-off when the preferences are dictated by a desire for continuity of care. |
| Antenatal clinics | The initial preference is dictated by the mother's home location. | To understand how the natural geographical clustering of demand may ease the problem. |
| Shifts & p/t | Shift patterns are incorporated. | To determine the magnitude of the additional challenge of staff availability to achieving continuity of care. |
| Flexible visits | The timing of some home visits may be flexible. | To assess whether a small degree of flexibility might be sufficient to help deliver continuity of care. |

may not be identical if the travel times are asymmetric. Entries 7–10 all note the possible savings if preferred staff are not used ($PS=0$); such links would only be adopted if the preferred staff were unavailable.

d. The trade-offs between travel time and patient preferences

Applying the algorithm with no concern for patients' preferences ($PS=0$) produces an allocation and routing with a travel time of 424 min, as noted in Fig. 3. If patients' preferences are considered, but with just a low priority ($PS=10$), the same routes are adopted, with the same total travel time of 424 min, but a better allocation of staff such that 50 % of the preferences are satisfied. Placing greater emphasis on the preferences ($PS=30$) induces a change in the allocation, and routes, such that 60 % are satisfied though travel time increases a little to 434 min. Increasing the emphasis on patients' preferences still further ($PS=60$) ensures that all are satisfied, though the travel time is now 562 min.

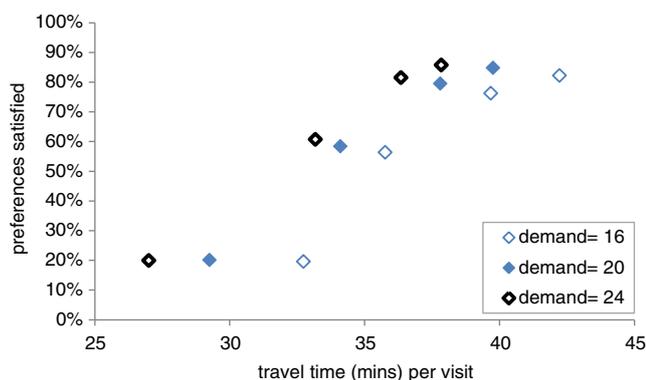


Fig. 5 Travel time: preferences trade-off (noting the results of experiment 1 with $PS=0,10,30,60$ for each specified mean demand)

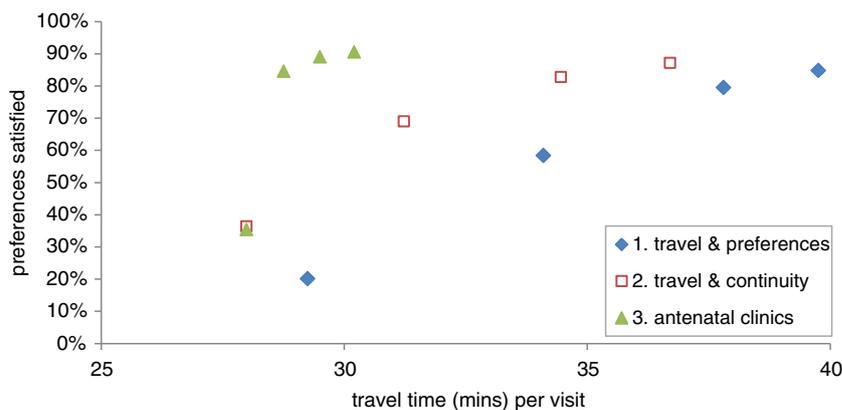
3 Simulating postnatal care in the community

a. Simulating the sequence of postnatal visits

The pattern of maternity care can vary, depending on local practice and the mother and baby's needs, but this paper considers the care provided by a small team of midwives in a typical area. Women attend their local antenatal clinic and begin to build a relationship with a midwife. After the birth in hospital, the mother and baby return home and receive a series of four visits over a period of 10–15 days. A standard pattern of visits might be days 1, 2, 6 and 12 though there may be some flexibility about this schedule, as discussed later. Ideally the first visit will be from the midwife who provided the mother's antenatal clinic, with all subsequent visits from the same midwife, minimising the number of handovers.

The simulation, illustrated in Fig. 4, was constructed in Simul8 to model the pattern of visits, incorporating the travelling nurses' algorithm to allocate mothers to staff and to construct efficient routes while striving to provide continuity of care. The simulation captures the critical stochastic nature of demand, with considerable variability over time and geography. A new set of mothers is generated each day, with a Poisson distribution corresponding to a specified mean demand. Each mother has a home postcode, randomly generated reflecting the geographical distribution of the population. The midwives are associated with GP practices and local health centres where the antenatal clinics are held. Thus the mother's location implies a relationship with a particular midwife which, ideally, will be maintained throughout her postnatal care. Every day in the simulation, the set of mothers requiring visits, both initial and follow-up, is analysed using the travelling nurses' algorithm to allocate a midwife to each mother and identify the sequences of visits.

Fig. 6 Travel time: continuity of care trade-off (mean daily demand=20; considering *PS*=0,10,30,60 for each scenario)



The results presented in this paper refer to a typical unit serving a specified area, though often units will co-operate when there are particular needs such as staff sickness. There were 6 midwives providing cover over 7 days per week, though given shift patterns and other work or training commitments there were typically just 3–5 midwives available on any 1 day visiting 10–40 mothers.

b. Simulation experiments

The simulation was used in a series of experiments, as summarized in Table 2, refining the problem specification and introducing additional complexities in a phased series of experiments to help develop understanding. Throughout the paper 95 % confidence intervals are quoted, based on 20 trials each composed of 30 simulated days and a warm-up period of 10 days reflecting the typical time for the completion of a mother’s series of postnatal visits.

4 Travelling time and satisfying preferences

The initial simulation experiment assumed that each patient required a single visit, with a randomly generated preference for one of five midwives, each of whom was available every

Table 3 Typical postnatal care shift pattern (hours available)

| | Staff | | | | | | Total |
|-----------|-------|----|----|----|----|----|-------|
| | 1 | 2 | 3 | 4 | 5 | 6 | |
| Monday | 0 | 8 | 8 | 0 | 4 | 0 | 20 |
| Tuesday | 8 | 4 | 0 | 4 | 4 | 0 | 20 |
| Wednesday | 8 | 0 | 4 | 0 | 0 | 8 | 20 |
| Thursday | 4 | 8 | 8 | 0 | 0 | 0 | 20 |
| Friday | 4 | 8 | 0 | 0 | 8 | 0 | 20 |
| Saturday | 4 | 0 | 0 | 8 | 8 | 0 | 20 |
| Sunday | 0 | 0 | 8 | 2 | 4 | 6 | 20 |
| Total | 28 | 28 | 28 | 14 | 28 | 14 | 140 |

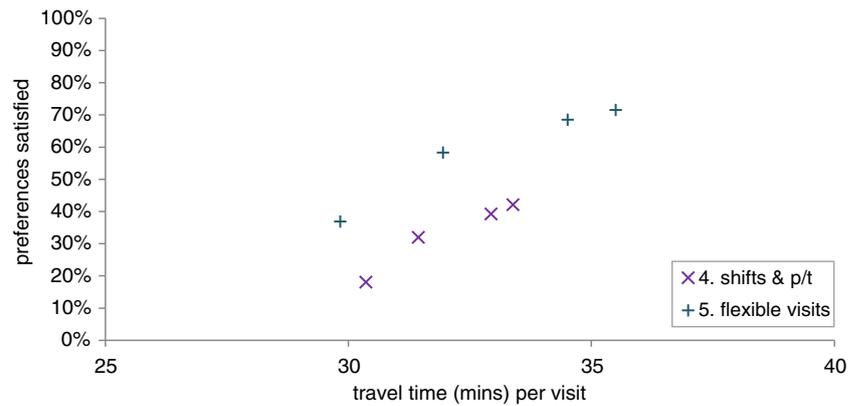
day. The experiment was designed to examine the relationship between travelling time and the satisfaction of preferences for particular staff. Later experiments explored more realistic scenarios involving more complex interrelationships. Three levels of demand were considered (mean daily demand=16, 20, 24 requests per day) and four degrees of priority for satisfying the preferences (*PS*=0, 10, 30, 60). When demand is relatively low (mean daily demand=16) and preferences are ignored (*PS*=0) the mean travel time per visit is 32.7±0.5 min and 19.6±0.9 % of preferences are satisfied, as noted in Fig. 5. When demand is higher (mean daily demand=24), the travel time per visit is reduced to 27.0±0.3 min, with 19.9±0.6 % of preferences satisfied. Some of this reduction, 2.6 min, is a consequence of each midwife having more visits per day and hence fewer legs per visit, allowing for the initial and final journeys to the base station. But 3.1 min of the reduction in travel time per visit may be attributed to exploiting the additional allocation and routing options inherent in a higher level of demand.

Even a small consideration of preferences (*PS*=10) produces a substantial increase in the proportion of satisfied preferences: with low demand (demand=16) 56.3±1.2 % of preferences are satisfied, though with a small increase in mean travel time to 35.8±0.5 min. As recorded in Fig. 5, further improvements in preference satisfaction might be achieved but at the expense of greater travelling time: 82.3±0.8 % of preferences could be met if a mean travel time of 42.2±0.6 min is accepted. Similar improvements in preference satisfaction are observed in the experiments exploring higher levels of demand, though the mean travel times are always lower when the demand is higher.

5 Travelling time and continuity of care

In the study of maternity care, the preferences were dictated by a desire to minimise the number of handovers, enhancing longitudinal continuity of care. In experiment 2, each mother received four home-visits: the preference for the first visit was

Fig. 7 Managing restricted staff availability off (mean daily demand=20; considering $PS=0,10,30,60$ for each scenario)



still random but on the three subsequent home-visits, mothers preferred to receive their care from the same midwife as the previous visit. Other preference functions might be explored, possibly modelling small team-working, though this would require a more complex allocation and routing method. The mean daily number of newly discharged mothers was specified as 5, hence the total daily mean demand for home-visits was 20, identical to that of experiment 1. In experiment 2 it was again assumed that all five midwives were available every day of the week; this assumption was explored in later experiments. Different priorities were considered again ($PS=0, 10, 30, 60$) and the resulting mean travel times and preference satisfactions are illustrated in Fig. 6 as experiment 2 “travel & continuity”. Travel times are reduced and the preference satisfactions increased, compared to the results of experiment 1 “travel & preferences” when all of the mothers’ preferences were random. In experiment 3 “antenatal clinics”, the initial preferences are determined by the mother’s home location, assuming that she attended the local clinic and established a relationship with the midwife associated with the area. The three subsequent visits’ preferences, reflecting the desire for continuity of care, tend to ensure that the midwife is usually allocated to the same set of locations, though workload balancing considerations may occasionally preclude such an obvious allocation. Hence the objectives of minimising travel time and maximising continuity of care are often complementary: high levels of preference satisfaction, 84.6 % can be achieved with just an increase of 0.8 min per visit ($PS=10$), compared to the minimum travelling time solution ignoring preferences ($PS=0$), as illustrated in Fig. 6.

Table 4 Visit flexibility

| Visit | Days after discharge | | | |
|----------------|----------------------|---|---|----|
| | 1 | 2 | 3 | 4 |
| Fixed schedule | 1 | 2 | 6 | 12 |
| Early schedule | 1 | 2 | 5 | 10 |
| Late schedule | 1 | 3 | 7 | 14 |

6 Shift patterns, part-time working and continuity of care

In practice, midwives are not available every day of the week and handovers will be generated by the shift patterns: staff work a 5 day week arranged such that the whole team provides 7 day coverage. However, this is further complicated by midwives having commitments to tasks other than postnatal care, and also part-time working. A typical shift pattern for a team of six staff, with two part-time midwives, is illustrated in Table 3. Adopting this shift pattern in the simulation, the results of experiment 4 “shifts & p/t”, as illustrated in Fig. 7, suggest that even when continuity of care is prioritized ($PS=30$), the proportion of visits satisfying the preferences is just 39.2 ± 1.2 % with a mean travel time of 33.0 ± 0.8 min per visit. However, there may be some flexibility in the visit schedule with larger time windows. Although it is essential to visit the mother on the first day at home, even if the preferred midwife is not available, the precise days of the subsequent visits may be adjusted, as noted in Table 4. For example the third visit might occur on 5–7 days after discharge. Assuming such flexibility the handovers can be reduced substantially, as noted in experiment 5 “flexible visits” in Fig. 7, increasing the proportion of satisfied preferences to 68.5 ± 0.9 % with a mean travel time of 34.5 ± 0.8 min per visit ($PS=30$). Given the sequence of 4 home visits and assuming a binomial distribution, this implies that only 22.0 % of mothers would have the same midwife throughout their postnatal care though 62.5 % should have no more than one change of staff. This degree of continuity is achieved at a relatively low cost of an additional 4.7 min, or 16 %, of travelling per visit compared to the minimum suggested by this analysis (when $PS=0$, mean travel time=29.8 min).

7 Further developments

The current study has a number of limitations which deserve further work:

- Comparing alternative measures of continuity of care

- The possibility of team working, i.e. second preferences for midwives such that any handovers are restricted to a small, coherent team who can share information about mothers, facilitating information continuity even if longitudinal and relational continuity are impaired
- Effect of changing system parameters to reflect local conditions, e.g. varying the number of home visits received by each mother or the amount of part-time working; this would enable more specific recommendations that might be incorporated in practical allocation and routing guidance
- Exploring variants of the Clarke-Wright algorithm to improve performance of the travelling nurses' algorithm
- Comparisons of the travelling nurses' algorithm with more rigorous approaches; branch and bound or genetic algorithms may not be practical when analysing multiple sets of visit requests as in the simulation experiments but they would provide a valuable comparison to determine the range of validity for the simple algorithm adopted in this study.

8 Conclusions

The results of the simulation experiments suggest that high degrees of continuity of care can be achieved at a relatively low cost in terms of additional travelling time. However, shift patterns and part-time working disrupt continuity of care considerably. Even then a high degree of longitudinal continuity of care is achievable in post-natal services if there is reasonable flexibility in the schedule of home visits. 100 % continuity is impractical but this study suggests that 70 % of visit preferences might be satisfied, though specific targets for continuity of care should reflect the local conditions.

The study has potential relevance beyond healthcare to other services where travel time is important and the staff should not be viewed as anonymous "service deliverers". Simulation modelling of flows of activity usually focus on measures such as utilisation, process and waiting times. Quantitative measures of continuity of service provide a mechanism to relate simulation to other aspects of quality which so often depends on the relationships between staff and their patients or clients.

Acknowledgments This study was funded by the Royal College of Midwives, Scottish Government Health Department and ESRC.

References

1. Parker G, Corden A, Heaton J (2010) Synthesis and conceptual analysis of the SDO's programme's research on continuity of care, national institute for health research evaluations. Trials and Studies Coordinating Centre, Southampton
2. Cabana MD, Jee SH (2004) Does continuity of care improve patient outcomes? *J Family Pract* 52(12):974–980
3. Dreier J, Comaneshter DS, Rosenbluth Y, Battat E, Bitterman H, Cohen AD (2012) The association between continuity of care in the community and health outcomes: a population-based study, *Israel J Health Policy Research* 1(21) available at <http://www.ijhrp.org/content/1/1/21> accessed 27.11.13
4. Sandall J, Hatem M, Daven D, Soltani H, Gates S (2009) Discussions of findings from a Cochrane review of midwife-led versus other models of care for childbearing women: continuity, normality and safety. *Midwifery* 25:8–13
5. Kergosien Y, Lenté C, Billaut J-C (2009) Home health care problem: an extended multiple Traveling Salesman Problem, Multidisciplinary International Conference on Scheduling : Theory and Applications. Dublin, Ireland 85–92.: <http://www.schedulingconference.org/previous/publications/displaypub.php?key=2009-085-092-P&filename=mista.bib>. Accessed 12 June 2013
6. Steeg J, Schroder M (2008) A hybrid approach to solve the periodic home health care problem, *Operations Research Proceedings 2007* http://link.springer.com/content/pdf/10.1007%2F978-3-540-77903-2_46. Accessed 12 June 2013
7. Bektas T (2006) The multiple traveling salesman problems: an overview of formulations and solution procedures. *OMEGA* 34(3):209–219
8. Rasmussen MS, Justesen T, Dohn A, Larsen J (2012) The home care crew scheduling problem: preference-based visit clustering and temporal dependencies. *Eur J Opl Res* 219:598–610
9. Eveborn P, Flisberg P, Rönnqvist M (2006) LAPS CARE – an operational system for staff planning of home care. *Eur J Opl Res* 171:962–976
10. Groër C, Gloden B, Wasil E (2009) The consistent vehicle routing problem. *Manuf Serv Oper Manag* 11(4):630–643
11. Bertels S, Fahle T (2005) A hybrid setup for a hybrid scenario: combining heuristics for the home health care problem. *Comput Oper Res* 33(10):2866–2890
12. Pillac V, Gendreau M, Guéret C, Medaglia AL (2013) A review of dynamic vehicle routing problems. *Eur J Opl Res* 225(1):1–11
13. Begur SV, Miller DM, Weaver JR (1997) An integrated spatial DSS for scheduling and routing home-health-care nurses. *Interfaces* 27(4):35–48
14. Altinel I, Öncan T (2005) A new enhancement of the Clarke and Wright savings heuristic for the capacitated vehicle routing problem. *J Opl Res Soc* 56:954–961
15. Clarke G, Wright JW (1964) Scheduling of vehicles from a central depot to a number of delivery points. *Oper Res* 12(4):568–581
16. Ballou RH, Agarwal YK (1988) A performance comparison of several popular algorithms for vehicle routing and scheduling. *J Bus Logist* 9(1):51–65
17. Mutingi M, Mbohwa C (2013) Home healthcare worker scheduling: a group genetic algorithm approach, *Proceedings of the World Congress on Engineering 2013 Vol I, WCE 2013, July 3–5, 2013, London, UK*
18. Laporte G, Semet F (2001) Classical heuristics for the vehicle routing problem. In: Toth P, Vigo D (eds) *The vehicle routing problem*. SIAM monographs on discrete mathematics and applications. SIAM Publishing, Philadelphia, pp 109–128