USE OF AVAILABLE STORAGE TO IMPROVE SCHEDULING IN AN AUTOMOBILE ASSEMBLY PLANT

The Final Assembly Plant at General Motors–Holden's Automotive Ltd converts painted body shells into drive-away vehicles. It encounters difficulties when processing certain sequences of vehicles with high work contents, so GMHAL wishes to schedule its input to reduce or eliminate such undesirable sequences. GMHAL has a set of empirical rules for delineating undesirability.

The Painted Body Storage (PBS), which precedes Final Assembly, has 4 lanes that can be used to partially reschedule a vehicle sequence. Information on vehicle work content is available prior to arrival at the PBS, and GMHAL wants advice on using this data and the PBS to achieve a more satisfactory input to Final Assembly.

The Study Group devised three approaches.

1. Use the rules to show which short sequences are desirable and devise input and output strategies for the PBS to achieve these consistently. Choice between the strategies requires further investigation.

2. Model Final Assembly to produce an optimality criterion for vehicle sequences and use combinatorial optimization methods to optimise it over possible PBS outputs. A characterization of these outputs was derived.

3. Suggest that the initial production be suitably scheduled, which may substantially reduce the difficulties at the PBS stage.

1. Introduction

The production of a motor vehicle involves a systematic and highly integrated production line which literally builds vehicles from the floor up. Work on the line is divided into a large number of separate tasks all of roughly equal duration. Each task is assigned to a workstation on the line, which performs it repetitively on the vehicles as they come through the station in a steady stream. Workstations are staffed by one or more persons, or occasionally are fully automated. There are sub-assembly lines to the main production line, each operating in a similar fashion.

The GMHAL plant at Elizabeth is the sole producer in Australia of the Commodore range of vehicles. The range includes a variety of models, from a utility or basic sedan up to the luxury Statesman/Caprice. Each of these in turn can be fitted with a number of optional extras (options), such as power accessories or manual transmission. The result is that different vehicles will have quite substantially different work contents in their production. This means that any particular workstation may be faced with a large variability in the amount of work it needs to provide to different vehicles as they pass through it, and clearly the station may get into difficulties if it encounters a succession of vehicles requiring a high work content from it.

How the workstations cope with this problem at present is discussed below, but it is obvious that one important aspect of the problem is whether the vehicles can be scheduled, or sequenced, so as to minimise the occurrence of such events.

The question brought to the Study Group was how to decide on, and implement, an appropriate schedule in the Final Assembly Plant. This innocuous name in fact covers a considerable part of the production process; a vehicle enters the Final Assembly as a painted body shell and leaves it as a drive-away product. So this is the part of the production line where the variety of models and options has its major impact on workstation loads. There are 145 workstations in the Final Assembly Plant main line, divided into 9 operating groups, with 3 major sub-assembly lines.

Vehicle bodies approach the Final Assembly from the Paint Shop. En route they pass through the Painted Body Storage (PBS) (Figure 1), which consists of four parallel storage lanes that rejoin to feed the Delivery Accumulator. The PBS functions in part purely as a storage buffer for the Final Assembly, but it also provides an opportunity for a limited reordering of vehicles on the production line to achieve a more satisfactory schedule in the Final Assembly Plant. Exactly what can be achieved in this way, and how to achieve it, are matters central to consideration of the problem.

2. More details

(i) **Operation of a workstation.** The fact that a workstation needs to provide substantially different amounts of work to different types of vehicles has already been mentioned as a catalyst for the scheduling problem. Typically the line runs at 10ft/min, and a workstation occupies 20ft of the line, so nominally a station has 2 min. work on a vehicle. The reality is rather different. For example, the first group of 31 stations in Final Assembly (concerned with wiring, removing the doors and headlining) has to provide a total of 76 minutes work on a basic utility, but 117 minutes on a Statesman; options may add further work requirements. So a Statesman requires, on average, nearly 4 min. work per station. How is this managed?

First, stations may have more than one worker, providing 2 min of work *each*. But this is not the whole answer; the group above, for instance, usually

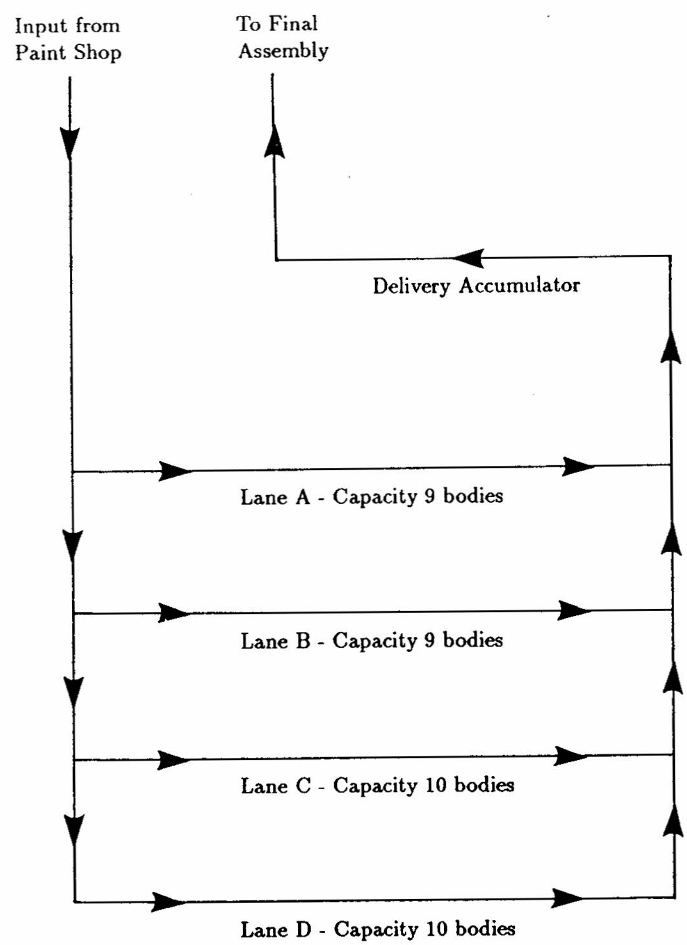

Figure 1: The painted body storage area (PBS)

has 42 workers so they must still average about 2 3/4 min. work per person on a Statesman. An important step to achieving this is to allow workers to 'float' along the line with a vehicle, possibly as far as the end of the following station, combined with an arrangement of work that allows personnel in that following station to work on the vehicle simultaneously. In fact it is a critical design problem when planning the original operation to ensure task scheduling that permits the necessary amount of synchronous working.

As a backup, the foreperson of most groups has a pool of emergency labour that can be directed to actual or potential troublespots to provide a temporary increase in work input for a station. The line can be stopped as a last resort to prevent a major emergency, but this is to be avoided if at all possible. A typical troublespot could be the arrival of two high work-content vehicles in succession, so that a worker floating down line with the first vehicle does not have time to return to the start before encountering the second.

At present, the backup labour assists with this difficulty. However, Holden's would like to keep staff levels as low as possible and hence to minimise the amount of such backup. Scheduling of vehicles so as to avoid as far as possible such troublesome conjunctions would help in achieving this goal.

(ii) **Current operation of the PBS.** At some point early in its production, a vehicle body acquires an electronic identification 'tag' containing information about its model and the options to be fitted. This can be read by various sensors around the plant. Currently, a sensor near the start of the PBS relays this information to a controller, which converts it into a PBS code. This is an integer between 1 and 10 broadly related to model type; for example 1 is a utility, 2 an executive/Berlina sedan, 3 an S-pack and so on. The controller then refers to a preassigned input table which specifies which code goes into which PBS storage lane; codes 2, 3 and 5 go to lane A for instance. At the other end of the PBS vehicles are removed from lanes by the controller in a cycle of length 11, namely D A B D A B D A C D A and repeat. There are also default options for each end to handle full and empty lanes.

This procedure is a first attempt at rescheduling by using the PBS. It is not sufficiently successful at avoiding undesirable sequences of high work-content vehicles, not surprisingly considering its inflexible methodology. For instance, in a randomly chosen run of 38 vehicles out of the PBS there were 4 major, 8 medium and 11 minor instances of undesirable vehicle conjunctions (see (iv)). The question asked is whether, and if so how, this performance can be improved.

(iii) **New vehicle classifications.** To aid in the improvement of vehicle rescheduling procedures, among other reasons, Holden's are proposing a more detailed coding system than the current 1 - 10 PBS code. Vehicles will be classified by model and (high work content) options. This new code will also be carried in the identification tag.

Specifically, there will be

- 4 models: Sedan; Luxury; Wagon; Utility; abbreviate them as S, L, W, U;
- 6 options: Power Accessories; V8 motor; Manual Transmission; Independent Rear Suspension; Bench Seats; Colour Bumper Facias.

This appears to give $4 \times 2^6 = 256$ vehicle types. In fact, not all options are available on all models, and the total number of types, hence codes, will be about 100.

(iv) **Separation rules.** To help decide what constitutes a satisfactory schedule, line workers were asked what sequences they found troublesome in practice. From their answers, a set of heuristic separation rules was formulated, 8 primary rules related to model separation, 6 secondary rules to option separation. They are given in Appendix A.

Although vehicle sequences conforming closely or wholly to these rules have not yet been run through the Final Assembly, Holden's believe that, if implementable, these rules ought to produce satisfactory scheduling.

Note that the rules are not very tidy at present and are not of equal importance. More will be said on these points later.

(v) **Objectives.** To draw together the motivations for Holden's to raise this problem with the Study Group, most of which have already been mentioned, the objective given to the Study Group is quoted here.

- To effectively use PBS to sort incoming vehicles from paint into a sequence that effectively separates high work value features from each other.
- To ensure that production operators are given the best possible time frame for completion of their work assignments on all types of vehicles.
- To reduce the need to over supply labour in at least 7 production line operating groups to cover periods of high (excessive) work content in certain strings of vehicle types.
- To assist in the achievement of improved quality goals.

3. A general approach

At present, it is possible to find out the type of a vehicle about 100 stations before it arrives at the PBS. So at any moment there is a known type sequence of length about 100 which is the current input to the PBS. We ignore the fact that in practice this sequence may change slightly en route, for example due to defective bodies being removed for repair and reinserted later.

Given this input sequence, we need to decide what permutations of it can possibly be generated by the PBS, and which of these is optimal in some sense. Thus stated, the problem looks like a (very large) combinatorial optimization exercise. A possible solution method is:

- Find an *objective function* which gives a measure of the acceptability of a sequence of vehicles as input to the Final Assembly.
- Optimize this function over all *possible sequences* generated by the PBS for the given input.
- Input or output a vehicle, and repeat.

As stated, this looks like a major computational task. Its complexity depends in part on how the term 'objective function' is interpreted. Most of the Study Group's work was centred on the two key ingredients of this method, the objective function and the possible sequences, principally the latter.

4. The objective function

There are two possible bases for the construction of an objective function in this problem; the separation rules (§2(iv)) or a model of the Final Assembly Plant. Some combination of these approaches may also be useful, of course.

(i) **Rule-based.** In principle this is straightforward. To assign a numerical measure to a sequence of vehicles, decide on a weight or score to be attached to each rule. This will quantify the relative seriousness of violating the rule, so the primary rules will typically have larger scores. For any given sequence the total score of all its rule violations can then be calculated as the value of the objective function.

Holden's believe that such scores can be assigned. There are various refinements that can be added to this basic scheme, such as a series of scores relating to the degree of violation of primary rules 2 or 8. One advantage of this approach is that it is simple to include other rules, and their scores, as felt necessary.

A more qualitative version of the approach is to use the rules in a less formal way to provide general guidance on what are 'good' or 'bad' sequences. Much of the work reported below on generating satisfactory schedules of vehicles *via* the PBS is based on this idea.

In more technical language, the former version looks for an optimal sequence and the latter for only a feasible sequence, treating the rules as constraints.

(ii) **Model-based.** This approach requires a more detailed study of what happens when a sequence of vehicles passes through the Final Assembly, to quantify the work content demands placed by the sequence on the workstations. These can then be used to measure how often and to what extent stations will become overloaded. There are various possibilities here.

- Suppose station i can provide an average amount of work/vehicle of α_i, and that vehicle j in the sequence requires work w_{ij} from station i. After n vehicles, the total discrepancy between work required and work provideable at station i is

$$d_i(n) = \sum_{j=1}^{n} w_{ij} - n\alpha_i. \tag{1}$$

  This could be incorporated in an objective function in a number of ways. Broadly speaking, a station will be overloaded if $d_i(n)$ ever gets too large; more generally, it is desirable to keep the fluctuations in the sequence $d_i(n)$ as small as possible. These ideas lead to the suggestions:

  1. Score a penalty if $d_i(n) > \delta_i$, for some suitably chosen δ_i (possibly $\delta_i \equiv \delta$), and add up penalties over n and some or all i.
  2. Take $\mu_i \equiv \max_n n^{-1}|d_i(n)|$, and minimise some function of these, such as their sum over some or all i.

  Note that this concept does not at present incorporate the important idea of synchronous work by adjacent stations, though it would not be difficult to build this in where it was allowed by treating several adjacent stations as a new 'station'. This raises the question of whether it is possible to produce a satisfactory model based just on the 9 groups rather than individual stations. Since this would clearly lead to a much simpler model, it is worth serious study.

- Calculate the total throughput time for a sequence of vehicles. This is like the process scheduling, or sequencing, problem (Moder and Elmaghraby, 1978, I-9), for which there are heuristic algorithms to minimise the throughput time. The important question here is whether a small total throughput time for a sequence is necessarily related to a satisfactory distribution of work at individual stations. It seems likely they are connected, but the matter would need further investigation before such an objective function could be used confidently.

- Develop a program to simulate the flow of vehicles through the Final Assembly, and use the output to derive an objective function. Both previous suggestions are, of course, candidates here; more generally, the program could include options to supply extra labour as necessary, or stop the line in extreme crises, and costs could be incurred by these events. This method gives the greatest flexibility but also requires a very detailed study of the whole Final Assembly. Of course, if a treatment of groups rather than individual stations is possible the process is greatly simplified and so more feasible.

5. Input and output sequences

This section will examine two important issues. The first is heuristic strategies for generating satisfactory input and output sequences at the PBS. The second is the more general question of exactly which output sequences can be produced by the PBS from a given input sequence. The answer to this is, of course, a prerequisite for development of more precise strategies to replace the heuristics.

(a) Strategies for sequence generation.

This is essentially an empirical exercise. It initially uses the primary rules as the principal guideline, with the secondary rules entering to resolve ambiguities, although a more refined version is undoubtedly possible.

In the more general approach described in §6, the output sequence is the one of critical importance, and the input strategy is largely determined by the output sought. Here, however, it seems more practical to consider the strategies separately.

A basic concept is that of a *good* sequence of vehicles. This is a sequence that satisfies the primary, and perhaps also the secondary, rules; a distinction between good (primary only) and very good (primary and secondary) sequences may be useful here. It is certainly desirable to try and make at least the output a good sequence.

Using this concept suggests the following strategies:

Strategy I1. Add the next input to a PBS lane so as to continue a good sequence if possible. When there is a choice of such lanes, add it to the one with fewest vehicles, breaking ties at random. If no such continuation is possible, choose between all 4 lanes in the same way.

Strategy O1. Add a vehicle from the head of a lane to the output so as to continue a good sequence if possible. When there is a choice of such lanes, add from the one with the most vehicles, breaking ties at random. If no such continuation is possible, choose between all 4 lanes in the same way.

Since there are 4 model types and 4 PBS lanes, it is possible to implement another input strategy.

Strategy I2. Reserve each PBS lane for a single model. If overflow occurs, add to another lane so as to make a good sequence, if possible.

Comments

(i) There are obviously many variants of these rules, such as deliberately spreading around sedans in I2 to try and avoid overflow in the sedan lane.

(ii) If the secondary rules have not been used in judging whether a sequence is good, they could be used to break ties.

(iii) I2 has the advantage of nearly always providing a model of each type for output at every stage. However, it has the disadvantage that usually there will be exactly one of each type available for output, when further choice using secondary rules is not possible. Also I1 distributes vehicles more evenly between lanes, reducing overflow problems.

(iv) An empirical good sequence that is worth noting is SWSUSLSWS, repeated (see §2(iii) for this notation). Note that sedans can be added, and other models removed without affecting the goodness. It could be used as a target in either I1 or O2, in the latter case especially if I2 is used for input.

There are, clearly, many other possible strategies. Three of these follow:

Strategy I3. Add the next input to a PBS lane to maximise its separation from the last vehicle of the same model in that lane. If necessary, use shortest queues and random tie breaking as before.

Strategy O2. This attempts to reproduce the input model frequencies in the output. Use the 100+ vehicles in the input queue and PBS to determine current input model frequencies, hence average model spacings. For the next output, choose a model whose spacing to the most recent output vehicle of the same model exceeds its current input spacing; if there is more than one such, choose the one with the largest spacing. If there are none, choose so as to make the output spacing as large as possible. Break ties at random. Update the input frequencies and repeat.

Strategy O3. Consider choosing the next n output vehicles from the PBS; this gives at most 4^n choices. Select an objective function (see §4), evaluate it for each possible output sequence and choose the optimal one. Output the first vehicle in this sequence from the PBS and repeat.

Comments

(v) O3 is more in the spirit of the approach of §6, but the optimization here can probably be done by enumeration in a practical time.

(vi) It would be interesting to take $n < 8$ in O3 and see how often Rule 3* is violated.

(vii) In O3, it is possible, and perhaps desirable to include part of the immediate past output sequence in the optimization. It is the only strategy mentioned here that attempts to plan ahead and not just take a vehicle at a time. If it is adopted, then it is possible to take $n < 8$ in O3 and still satisfy Rule 3*.

(viii) O2 and I3 do not use the rules explicitly, though their use of separation criteria is obviously in the spirit of the rules. With I3 it may be more reasonable to maximise spacings only for non-S models.

We give the results of a small simulation study, undertaken by Dr Gary Newsam, in Table 1. Note that I2 gives by far the best output characteristics, suggesting that putting good sequences into lanes does not necessarily produce good output. However, I2 does overflow lanes, though as this is usually the sedan lane it is not a critical problem. There is little to choose between O1 and O2 here, though other evidence (and commonsense) suggest O1 will often be the better.

Table 1

	No. of failures		Maximum queue length	
	O1	O2	O1	O2
I1	10	8	7	7
I2	0	0	10	12
I3	9	12	9	9

Based on a run of 200 vehicles, with random inputs in the proportions (S, W, U, L) = (.55, .20, .14, .11). It gives the numbers of failures, *i.e.* violations of the primary rules, and the maximum queue lengths in any lane.

(b) **Characterization of output sequences.**

Whenever a strategy linking input and output is considered, such as the one in §6, it is important to be able to decide what outputs are achievable for a given input. In principle this is possible, as shown below.

Assume that (i) PBS has 4 lanes
(ii) each lane has infinite capacities
(iii) output only occurs when the appropriate vehicle reaches the head of its lane

Suppose vehicles labelled $1, \ldots, N$ in sequence are permuted by the PBS into an output $\sigma = (a_1, \ldots, a_N)$. Say we have a *partition* of σ into subsequences $\sigma_1, \ldots, \sigma_T$ if:

- each a_i belongs to exactly one σ_t $(t = 1, \ldots, T)$
- the elements of each σ_t have the same ordering as they had in σ.

Proposition 1. σ is a possible output sequence if and only if we can find a partition of σ with T $\leq$ 4 and each σ_t an increasing sequence.

For if we have such a partition, we can achieve σ by assigning each σ_t to one of the PBS lanes. Conversely, if σ is a possible output the sequence of vehicles moved down each of the lanes gives an appropriate partition.

Now this is not a *constructive* proof; it does not say how to get suitable σ_t (not unique in general, of course), though for short σ's it is easy to check suitability directly. Here is one possible construction which always generates *a* partition of any σ into increasing σ_t.

1. To generate σ_1, choose the longest possible increasing subsequence from σ starting at a_1.
2. To generate σ_2, choose the longest possible increasing subsequence from σ - σ_1, starting from the first element of σ - σ_1.
3. Repeat to generate increasing subsequences $\sigma_3, \ldots, \sigma_T$, T $\leq$ N.

Example: N = 20
σ = (3,1,6,7,4,2,9,11,10,5,17,20,8,18,12,13,19,16,14,15)
σ_1 = (3,6,7,9,11,17,20)
σ_2 = (1,4,10,18,19)
σ_3 = (2,5,8,12,13,16)
σ_4 = (14,15)

If T$\leq$4, this partition satisfies the conditions of Proposition 1, so σ is a possible output sequence. Further, if $T > 4$, σ is not a possible output sequence. For choose l_T = index of 1st element of σ_T and then for $j = 1, \ldots, T-1$ choose l_j as the largest index less than l_{j+1} for a's in σ_j. So $l_1 < \ldots < l_T$, but $a_{l_1} > \ldots > a_{l_T}$, since if $a_{l_{j+1}} > a_{l_j}$ and $a_{l_j} \in \sigma_j$ then $a_{l_{j+1}}$ would belong to σ_j also, by construction. So these T a's have been completely reversed in order between input and output, and this needs T lanes. This has proved the following result.

Proposition 2. σ is a possible output sequence if and only if the partition constructed above has T$\leq$4.

The practical drawback of this result is its reliance on assumptions (ii) and (iii), (ii) being the bigger problem. It can be overcome as follows. Suppose lane t has capacity $C_t(t = 1, \ldots, 4)$. The situation to be avoided is moving an input element 'too far forward' in the output. Specifically, suppose $a_j = i$ in σ. If $j < i$ (so i is moved forward in the output) there must *not* be any σ_t $(t = 1, \ldots, 4)$ which has a run of length $C_t + 1, a_{l_0} < \ldots < a_{l_{C_t}}$, say, such that $j < l_0$ and $i > a_{l_{C_t}}$. Otherwise, there will be a queue of length $C_t + 1$ in lane t waiting for $i = a_j$ to arrive (after $a_{l_{C_t}}$) and depart (before a_{l_0}). This proves

Proposition 3. Under assumptions (i) and (iii), σ is a possible output sequence if and only if we can find a partition $\sigma_1, \ldots, \sigma_T$ of σ satisfying the conditions of Proposition 1 and such that for *any* run $a_{l_0}, \ldots, a_{l_{C_t}}$ of σ_t $(t = 1, \ldots, T)$ all inputs $i > a_{l_{C_t}}$ have outputs a_j with $j > l_0$.

Unfortunately, there does not appear to be a simple constructive method here analogous to Proposition 2.

6. The constrained optimization approach

In §3 a combinatorial optimization method was outlined and in §4 possible objective functions were discussed. A method for checking whether a vehicle sequence can be an output is given in §5(b).

The challenge is to devise a practical method of carrying out the optimization. One possibility is *simulated annealing* (van Laarhoven and Aarts, 1987) which is entirely suitable in principle but not easy to implement in any particular case.

Basically, we begin with an objective function $f(\sigma)$ over outputs σ, and form $g_T(\sigma)=\exp\{T^{-1}f(\sigma)\}$. As $T \downarrow 0, g_T(\sigma)$ becomes concentrated at the σ which maximises f. Fix T initially, start at some (hopefully reasonable) σ and move to a σ^1 close to σ, in some sense. If $g_T(\sigma^1) > g_T(\sigma)$, accept σ^1 as the new σ; otherwise accept it with a certain probability based on g_T (this is to avoid entrapment at a local maximum). This process is repeated whilst progressively reducing T, the so-called *cooling schedule*. Hopefully, if the schedule is suitably chosen, the process will converge to the global optimum. However this is difficult to check in any particular case.

In this problem, we want to find the optimum output from the PBS for a given input which could be up to 100 vehicles long. This is a huge task, and remains so even if only a subset of this input is used initially. To make the method practical we need, among other things:

(i) An f which can be rapidly updated as we move from σ to σ^1;

(ii) A sensible idea for getting σ^1's close to σ which are still possible output sequences e.g. those produced by swapping one or more pairs of elements of σ;

(iii) A cooling schedule tuned to this problem.

Some progress has been made towards achieving these goals. If they can be achieved we would have an initial optimal output. Since the calculation must now operate in real time, we need a rapid ($\leq$ 2 minutes) method of updating the solution as the first output leaves and the next input arrives.

7. Comments on the separation rules

As they stand, the primary rules are not at all disjoint, and some are formally redundant. Observe that:

- Rule 2 $\Rightarrow$ Rule 1 and Rule 6 $\Rightarrow$ Rule 5.
- Rules 1, 3, 4 $\Rightarrow$ S/Wags must have a sedan on each side.
- Rules 3, 5, 6, 7 $\Rightarrow$ Utilities must have a sedan on each side.

In the notation of §2(iii), W's must occur in triplets SWS, U's in triplets SUS. Further, by Rule 2 adjacent wagon triplets can not share S's although a utility triplet could share an S with an adjacent triplet of either type, *i.e.* $\underline{SW\overline{S}}\overline{US}$ is acceptable, $\underline{SW\overline{S}}\overline{WS}$ is not. So, the primary rules can be reduced to:

1.* S/Wags must occur in triplets SWS, which cannot merge with each other

2.* Utilities must occur in triplets SUS which can merge with other triplets

3.* Rule 8

For scoring purposes (see §4), some version of the original rules may still be more useful of course.

It is clear that the absolute implementability of the rules depends on the production percentages of the different models and options. For example, demand for coloured bumper facias usually runs between 35% and 60%. If it exceeds 50%, secondary rule 6 must be violated eventually. Similarly, the percentage of wagons produced is usually between 15% and 25%; should it exceed 33% then primary rule 2 cannot always hold.

8. Scheduling the initial production

At present, daily production schedules for the Body Shop, where the building of a vehicle begins, are set by a combination of orders on hand and anticipated demand. So a decision is made to commence production of so many sedans, wagons *etc.*, each having a particular set of options. However, the sequence in which these vehicles are actually commenced seems to be nearly arbitrary.

It would, therefore, be advantageous if the initial production schedule could be made to generate, as nearly as possible, a sequence of vehicles satisfying the rules. One major complication is that sedans and wagons are commenced on one line, utilities and luxury on another, with the two lines merging at some point. The merging is done manually and under the constraint of keeping the vehicle flow going, so it is probably unreasonable currently to expect the workers at that point to ensure that the merging gives all appropriate separations. (The wagon/sedan rules will not be affected by this of course). Further, vehicle bodies are subsequently removed and reinserted at various stages before the PBS is reached, so the merging sequence is unlikely to be totally maintained. However, if only a small proportion of vehicles was reordered in an otherwise satisfactory sequence, the task of using the PBS to generate the original or another satisfactory sequence would be made much easier.

9. Some related ideas

If we consider successive pairs of vehicles passing a workstation, it is possible to associate a total workstation idle time for the combination, defined to be $+\infty$ if in fact the station cannot handle the pair in the allotted time. Add these times over all pairs within a fixed vehicle sequence, so that the totals of types are known, and minimize over permutations of the sequence. This gives a schedule that maximizes productivity without creating crises, at least pairwise, and it can be solved as a standard transportation problem (Hillier and Lieberman, 1974, Ch.3). Ideally, the method would now be extended to successive n-tuples with $n \geq 8$ (*cf.* Rule 8), but apparently a solution to this does not yet exist.

Two recent references bear on the problem. In Parrello *et al.* (1986), there is a nice discussion of the principles of construction of an optimal sequence, using penalties for violating rules rather like the GMHAL ones and a strategy akin to O1 or O3. The strategy explicitly includes the idea of placing 'difficult' cars in the sequence whenever possible (difficulty is roughly proportional to high work content), a notion mentioned in the Study Group but not otherwise incorporated in this report. McCormack *et al.* (1989) develop a version of the sequencing problem (see §4(ii)) applied to cyclic schedules, which are probably not sufficiently practical in this context.

10. Conclusions

In principle, the combinatorial optimization approach *via* simulated annealing will provide a solution to the problem. It may be, however, that the practical difficulty of its implementation will prove too formidable for input sequences of a useful length; more work is needed to examine this point. If so, the more empirical methods of §5 are suggested. Again, a good deal of study will be necessary to decide which input and output strategies are most appropriate for cases encountered in practice. For example, although I2 worked well in one simulation study it is probably too risky to use without some modification or a lot more proving. Use of O3 would introduce some measure of the actual performance of the Final Assembly into the process. Scheduling of the initial production run should also be advantageous.

Appendix. The proposed separation rules

Primary (model) separation rules

1. Station wagons (S/Wags) not together.
2. S/Wags to be separated by at least two other models.
3. S/Wags and utility not together.
4. S/Wags not with luxury.
5. Utilities not together.
6. Utilities to be separated by one sedan.
7. Utilities not with luxury.
8. Luxury to be separated by at least 8 other models.

Secondary (option) separation rules

1. Electronic power options not together.
2. V8's not together.
3. Manual transmission not together.
4. Independent Rear Suspension not together.
5. Bench seats not together.
6. Coloured bumper facias not together.

Acknowledgements

The moderator for this problem (Mark Westcott) would like to thank Murray Thompson of Holden's for his expert sponsorship of the problem, for his patience and his clarity of exposition, and for organizing a most fascinating visit to the Elizabeth plant. There was a large and enthusiastic group of participants working on the problem, including Rod Bell, David Gates, Geena Gwan, Phil Howlett, Huu-Nhon Huynh, Graham Mills, Bernhard Neumann, Gary Newsam, Duncan Roper, Peter Trudinger and Tony Watts.

References

F.S. Hillier and G.J. Lieberman, *Operations Research*, 2nd ed. (Holden-Day, San Francisco, 1974).

P.J.M. van Laarhoven and E.H.L. Aarts, *Simulated Annealing: Theory and Applications* (Reidel, Dordrecht, 1987).

S.T. McCormack, M.L. Pinedo, S. Shenker and B. Wolf, "Sequencing in an assembly line with blocking to minimise cycle time", *Operations Research* 37 (1989), 925-935.

J.J. Moder and S.E. Elmaghraby (eds) *Handbook of Operations Research Vol 2. Models and Applications*, (van Nostrand, New York, 1978).

B.D. Parrello, W.C. Kabat and L. Wos, "Job-shop scheduling using automated reasoning: A case study of the car-sequencing problem", *J. Automated Reasoning* 2 (1986), 1-42.