

3.1. Construct a timetable.

The method of constructing a timetable is to make successive decisions about the assignment of activities to periods. Each decision can be made in one of two ways

- either (a) to assign an activity ^{at} to a period
- or (b) not to assign it

Choice (b) is known as a cancellation of the activity at the period

taken
 at this stage

We construct a procedure "progress" which jumps to print out the timetable if and when it is successfully completed, and exits normally if there is no way of completing a timetable on the basis of previously decisions. Each activation of progress takes one decision, and then enters itself recursively. If normal exit from the recursive call takes place, the decision must be reversed, and another recursive call is made. If this also exits normally, this proves that no decision can lead to a successful timetable, and thus the given activation should also exit normally, so that further backtracking may take place.

Thus the basic structure of the procedure is:

progress recursive procedure

```

begin  if complete then go to printout;
        choose appropriate (a,p);
        try assignment (a,p);
        try cancellation (a,p);
end

```

(both assignments and cancellations)

In order to detect completion of the timetable, we keep a count of all decisions taken so far, and compare it with the number of decisions which need to be taken.

```

count: Integer initially 0
to be decided: Integer constant initially  $\sum_a \text{size}(\text{possible}_a - \text{timetable}_0(a))$ 

```

size (Period) x size (Activity)

In constructing a backtracking algorithm to tackle a problem of any size it is most important to avoid as far as possible the pursuit of decision sequences which can be readily detected to be inconsistent, in the sense that they can never lead to success. Thus before taking each decision, the consistency of the previous decision should be rigorously checked, and if inconsistency is detected, an immediate exit from the current activation of "progress" should be made. We thus derive the following program:

construct the timetable:

begin progress recursive procedure

begin check consistency;

if inconsistent then go to impossible;

~~if count = *count_max* then go to printout;~~ *to be decided* stop

select suitable (a,p);

try assignment(a,p);

try cancellation (a,p);

impossible: ~~for~~

end;

← progress;

← print failure message; stop;

printout: *end*

end

3.2. Assignment and cancellation.

This method requires that we keep a record of decisions ^(as they are) made ~~previously~~. This may be done by two mappings:

$T, P: \text{Activity} \rightarrow \text{Period set}$

where $T(a)$ is the set of periods which have been assigned to a

~~(initially set equal to timetable)~~

initially empty

and $P(a)$ is the set of periods which have not been cancelled for a

~~(initially set equal to possible)~~

initially the full set.

Obviously, if our decisions are non-contradictory, the following will always hold:

$$\forall a. T(a) \overset{\text{space}}{\cap} P(a)$$

This is T

If $p \in P(a) - T(a)$, this means that no decision has yet been made about assigning or cancelling a ~~from~~ ^{at} p . An assignment may then be made by

$$T(a) := +p$$

and a cancellation may be made by

$$P(a) := -p.$$

We can now write the programs for assignment and cancellation:

Thus each decision either increases the number of periods in T or decreases the number of periods in P, until (eventually) they are equal, at which stage to each other, and therefore to the final timetable. The printout will output the value of T

try assignment(a,p):

begin T(a):+p;count:+1;progress;count:-1;T(a)~~#~~-pend

space

try cancellation(a,p):

begin P(a):-p;count:+1;progress;count:-1;P(a):+p end

3.3. Check Consistency.

It can be seen that the efficiency of the program is critically dependent on the success of the consistency check in ensuring that only those decisions are pursued which lead to a successful timetable. If an absolute test of consistency were available, it would never be necessary to backtrack over more than one decision, since the futility of a decision which failed to lead to a successful timetable would be immediately detected.

Unfortunately, it is too much to hope for a 100% test of consistency, since this would require a guarantee of the existence of a complex object like a timetable before it had been constructed. So all that can be done is to find a set of necessary conditions for the existence of a timetable based on current decisions; for falsity of a necessary condition will then be a clear indication that failure is inevitable. A condition N is shown to be necessary if it can be proved that:

$$C \ \&\forall a \ (T(a) \subseteq \text{timetable}(a) \supseteq CP(a)) \supset N.$$

This means, close up the space

The following necessary conditions can readily be proved:

- $\forall a \ \text{size}(T(a)) \leq \text{times}(a) \leq \text{size}(P(a))$ N1
- $\forall i,p. \ \text{busy}(i,p) \leq \text{lives}(i) \leq \text{possbusy}(i,p)$ N2
- where $\text{busy}(i,p) = \text{size}\{a \in \text{users}(i) \ \& \ p \in T(a)\}$ I1
- and $\text{possbusy}(i,p) = \text{size}\{a \in \text{users}(i) \ \& \ p \in P(a)\}$ I2
- $\forall a \ \text{a} \in \text{spread} \supset \text{times}(a) \leq \text{size}(\text{possdays}(a)) \times \text{length}(a)$
- where $\text{possdays}(a) = \{d \mid P(a) \wedge \text{periods in } (d) \neq \text{empty}\}$ I3 N3

insert N3, N4, N5, N6 here. (from next page)

In writing the program to check consistency, much time can be saved if it can be assumed that consistency held before the most recent decision, and it is known what that decision was. It also pays to store the values of busy and possbusy rather than recomputing them each time.

We therefore introduce variables

olda: Activity; oldp: Period; action: {arriv, cancel}

after T(a):+p do {olda:=a; oldp:=p; action:=arriv}

after P(a):-p do {olda:=a; oldp:=p; action:=cancel}

The following necessary conditions can be ~~derived~~ derived from the corresponding constraints, together with the observation that T approaches the final timetable "from below" and P approaches it "from above." Thus ~~derivation is so~~ the necessity of the conditions is sufficiently obvious to require no proof.

$$\forall a: \quad \text{size}(T(a)) \leq \text{times}(a) \leq \text{size}(P(a)) \quad N1$$

$$\forall a, i: \quad \text{busy}(i, p) \leq \text{lives}(i) \leq \text{possbusy}(i, p) \quad N2$$

$$\text{where } \text{busy}(i, p) = \text{size} \{a \mid a \in \text{users}(i) \ \& \ p \in T(a)\}$$

$$\text{and } \text{possbusy}(i, p) = \text{size} \{a \mid a \in \text{users}(i) \ \& \ p \in P(a)\}$$

$$\forall a, d: \quad a \in \text{spread} \supset \text{size}(T(a) \wedge \text{periods in } (d)) \leq 1 \quad N3$$

$$\forall a, d: \quad a \in \text{spread} \vee \text{length}(a) > 1 \supset$$

$$\text{possdays}(T(a)) \times \text{length}(a) \leq \text{times}(a) \leq \text{possdays}(P(a)) \times \text{length}(a) \quad N4$$

~~forall~~

$$\forall a, d: \quad \text{length}(a) > 1 \supset \text{times}(a) \leq \text{size} \{d \mid \text{posstuples}(a, d) \neq \text{empty}\} \times \text{length}(a)$$

$$\text{where } \text{posstuples}(a, d) = \{ps \mid T(a) \subset ps \subset P(a) \wedge \text{periods in } (d) \quad N5$$

$$\& \text{ps} = \text{tuple}(\text{length}(a), \text{first}(ps))$$

$$\& \text{first}(ps) \in \text{starts}(a)\}$$

The conjunction of these five conditions is known as N.

The correctness of the program depends on the fact that ~~it~~ when the ~~transition~~ jump to printout is made, the consistency check which ^{guarantees} ~~the~~ truth of N will also guarantee that ~~truth~~ ^{T satisfies} of the constraints C . When $\text{count} = \text{size}(\text{Period}) \times \text{size}(\text{Activity})$, every decision will have been taken, and T will necessarily equal P . Substitution of T for P in ~~the~~ N will imply the truth of $C1$ to $C5$. $C6$ we ignore; and $C7$ and $C8$ can be guaranteed by ~~making all preassignments, before any other assign~~ (first).

The construction of a program to check N and jump to impossible if false may be easily derived from the definition of N , and may be postponed to a later stage.

length(a) > 1 \supset countuples (length(a), P(a) x length(a) \geq times(a) N4
 a \in spread \supset define postuples(a, d) = { p \in period set | T p \in start length(a)
 $\forall d$ (countuples(length(a), T(a) \wedge periods in(d)) \leq 1
 & size {d | countuples(length(a), P(a)) \geq 1} x length(a) \geq times(a)) N5

$\forall a, a', p, p'$ a' \in tie(a) & p \in T(a) & p' \in T(a) \supset dayof(p) \neq dayof(p') N6
p' \in dayof(p) p \in T(a)

The conjunction of these ^{five} conditions is known as N.

~~The program to check the truth of N may be fairly easily derived from the definition of N, and its detailed coding might well be postponed.~~
 3.4. Select suitable (a, p).

Another factor which is critical to the efficiency of timetable construction is the judicious selection of the next decision to take. An obviously sensible strategy is to take first those decisions which are forced, in the sense that they can only be taken one way; because if such a decision leads to inconsistency, backtracking is comparatively cheap.

(1) A condition FA is said to be sufficient for forcing assignment of ~~p to a~~ ^(a at p) if it can be proved that

$$FA(a, p) \supset \neg N^P_{P'}$$

where $P' = (P \cap a: P(a) - p)$, ie, to cancel p from (a) leads to immediate inconsistency.

(2) A condition FC is said to be sufficient for cancelling ~~p for~~ ^{a at p} if it can be proved that

$$FC(a, p) \supset \neg N^T_{T'}$$

where $T' = (T \cap a: T(a) + p)$, ie, to assign ~~a to p~~ ^{p to T(a)} leads to immediate inconsistency.

(3) An easy way to ensure that all preassignments are carried out first is to treat them as forced decisions.

Since a single decision may give rise to several further forced decisions, it is necessary to store these decisions for future execution in two sets

forced assign: Activity \rightarrow period set initially timetable.

~~forced assign, forced cancel: (Activity x Period) set initially empty~~

forced cancel: Activity \rightarrow period set initially forbidden.

Decisions are placed in these sets as soon as a sufficient condition for forcing is detected; thus

(a, p) \in forced assign \equiv FA(a, p) \checkmark p \in timetable(a)
 (a, p) \in forced cancel \equiv FC(a, p) \checkmark p \in forbidden(a)

We can now split selection of the next decision into two parts:

```

select suitable(a,p);
{select forced(a,p); select unforced(a,p)}

```

where

```

select forced(a,p):
  if forced assign ≠ empty then
    {(a,p) from forced assign; try assignment(a,p); go to impossible}
  else if forced cancel ≠ empty then
    {(a,p) from forced cancel; try cancellation(a,p); go to impossible}

```

Sufficient conditions for forcing can be readily derived from the corresponding consistency check.

~~The detection of forced decisions should be made during the check of consistency. Sufficient conditions for forcing can readily be derived from the corresponding consistency check.~~

~~Each of the following is a sufficient condition for assignment of a to p where $p \in P(a) - T(a)$:~~

- ~~size (P(a)) = times(a) FA1~~
- ~~$\exists i (i \in \text{requirement}(a) \ \& \ \text{lives}(i) = \text{possbusy}(i,p))$ FA2~~
- ~~day determined(a) & size(P(a) \wedge periods in(day of(p))) = length(a) FA3~~

~~where day determined(a) = $\exists a \in \text{spread} \ \& \ \text{times}(a) = \text{size}(\text{possdays}(a))$~~

~~$T_{\text{day}}(a) = \exists a \in \text{spread} \ \& \ \text{times}(a) = \text{size}(\text{possdays}(a)) \times \text{length}(a)$ then $\text{possdays}(a)$ else $\{d | T(a) \wedge \text{periods in}(d) \neq \text{empty}\}$ I3 FA4~~

~~length(a) > 1 & a \in spread & p \in \wedge posstuples(a,p)~~

~~where posstuples(a,p) =~~

~~$\{ps: \text{Period set} \mid T(a) \in ps \in P(a) \wedge \text{periods in}(\text{day of}(p))$
 $\& \ ps = \text{tuple}(\text{length}(a), \text{first}(ps))$
 $\& \ \text{first}(ps) \in \text{starts}(a)\}$~~

~~Each of the following is a sufficient condition for cancellation of p from a , where $p \in P(a) - T(a)$ ^{forced}~~

- ~~size(T(a)) = times(a) FC1~~
- ~~$\exists i (i \in \text{requirement}(a) \ \& \ \text{lives}(i) = \text{busy}(i,p))$ FC2~~
- ~~a \in spread & size(T(a) \wedge periods in(day of(p))) = length(a) FC3~~
- ~~length(a) > 1 & a \in spread & p \in \forall posstuples(a,p) FC4~~
- ~~$\exists a' (a' \in \text{tie}(a) \ \& \ T(a') \wedge \text{periods in}(\text{day of}(p)) \neq \text{empty})$ FC5~~
- ~~$\exists a' (a' \in \text{tie}(a) \ \& \ \text{day determined}(a') \ \& \ \text{day of}(p) \in \text{possdays}(a'))$ FC5~~

~~As before, the ^{program to} detection of these conditions can be easily derived, and the ^{its} detailed coding may ^{therefore} be postponed.~~

copy FA1 to FA4 here

copy FC1 to FC5 here

The following are sufficient ^(10a) conditions for cancellation of a at p , where $p \in P(a) - T(a)$:

$times(a) = size(T(a))$

FC1

$\exists i. i \in requirement(a) \ \& \ \text{lwes}(i) = \text{busy}(i, p)$
 $\text{daydetermined}(a) \ \& \ \text{dayof}(p) \in \text{posdays}(T(a))$
 where $\text{daydetermined}(a) \equiv (a \in \text{spread} \vee \text{length}(a) > 1) \ \& \ \text{times}(a) = \text{posdays}(P(a)) \times \text{length}(a)$.

FC2

FC4

FC4

$\exists p' a \in \text{spread} \ \& \ p' \in T(a) \ \& \ p \in \text{periodsin}(\text{dayof}(p'))$

FC3

$\text{length}(a) > 1 \ \& \ p \notin \bigcup \text{posstuples}(a, \text{dayof}(p))$

FC5

The following are sufficient conditions for assignment of a at p , where $p \in P(a) - T(a)$:

$size(P(a)) = times(a)$

FA1

$\exists i. i \in requirement(a) \ \& \ \text{possbuzy}(i, p) = \text{lwes}(i)$
 (no corresponding condition)
 $\text{daydetermined}(a)$

FA2

FA3

$\& \ size(P(a)) \wedge \text{periodsin}(\text{dayof}(p)) = \text{length}(a)$

FA4

$\text{length}(a) > 1 \ \& \ p \in \bigcap \text{posstuples}(a, \text{dayof}(p))$

FA5

close up ()

When an inconsistency has been detected, there may still be some forced decisions in the set. These should be removed before taking the next unforced decision. Therefore, between try assignment and try cancellation there should be

$\text{forced_assign} := \text{forced_cancel} := \text{empty}$.

The construction of a program to detect sufficient conditions for forcing can again be easily derived; and will again be postponed until after more difficult decisions have been taken.

35, select unforced (a,p);

In selecting an unforced decision, it is a good idea to select a decision which is most likely to lead to ~~small~~ ^{early} revelation of a latent inconsistency if there is one. The most suitable candidates will be those decisions which are most likely to lead to the longest chains of consequential forced decisions; ~~These decisions will be found~~

~~in areas which the human timetable~~
Such decisions will occur ⁱⁿ for ~~these are~~ those areas in which ^{offer} there is least freedom of ~~decision~~ choice in making decisions, - the areas which the human timetable would regard as "sticky", and ^{on} which he ^{also} would concentrate his early attention.

~~More precisely, for each (a,p) we can estimate roughly how many forces would result from ~~can~~ assignment and from cancellation; we choose the decision for which one of these is the greatest; and then, to give best chance of success, we ~~take the~~ choose first the other decision.~~

$$\cancel{\text{times}(a) - \text{size}(T(a)) - 1} \vee \text{size}(P(a)) - \text{times}(a) = 1 \quad \text{ST1}$$

$$\bar{I}i \in \text{requirement}(a) \ \& \ \left(\begin{array}{l} \cancel{\text{times}(i) - \text{busy}(i,p) - 1} \\ \text{possbusy}(i,p) - \text{times}(i) = 1 \end{array} \right) \quad \text{ST2}$$

(no corresponding condition) ST3

$$\text{daydetermined}(a) \ \& \ \text{size}(P(a) \wedge \text{periodsin}(\text{dayof}(p))) \quad \text{ST4}$$

$$\left(-\text{length}(a) = 1 \right) \quad \text{ST4}$$

$$\text{length}(a) > 1 \ \& \ T(a) \wedge \text{periodsin}(\text{dayof}(p)) \neq \text{empty}. \quad \text{ST5}$$

ST5 ensures that a partially assigned multiple period is always regarded as stiff.

We therefore introduce a set:

Activity \rightarrow Period set
stiff: (Activity x Period) set

into which decisions are placed when they are recognised to be stiff. The main criterion of stiffness is that the decision ~~is~~ is likely to result in further forced decisions and this can conveniently be detected ~~is~~ during ^(consistency check, at the same time as) the search for forced decisions. ~~are decided from~~

The most obvious criteria are for stiffness of (a, p) , where $p \in \mathbb{K} P(a) - T(a)$ are:

copy from

- size $(P(a)) - \text{times}(a) = 1$ ST1
- $\exists i (i \in \text{requirement}(a) \ \& \ \text{posbusy}(i, p) - \text{lives}(i) = 1)$ ST2
- $\exists d. \text{pred} \ \text{daydetermined}(a, d) \ \& \ \text{size}(P(a) \cap \text{periods in } (d)) - \text{length}(a) = 1$ ST3
- $\text{times}(a) - \text{size}(T(a)) = 1$ ST4
- $\exists i \ i \in \text{requirement}(a) \ \& \ \text{lives}(i) - \text{busy}(i, p) = 1$ ST5

The conjunction of these conditions will be known as ST.

~~These conditions may most readily be detected during the check on consistency; ~~is~~ again, the detailed coding may be postponed.~~ ^{programming}

It remains to decide how to select non-stiff decisions (4)

← The easiest way of doing this is ~~on "stiffness"~~ to allocate to each activity a score, indicating its "a priori" difficulty (and to ~~not~~ maintain a sorted list (from sequence) which the activity with highest score will always be ~~selecting~~ selected. A scoring formula will ~~also~~ give highest scores to ~~multiple~~ multiperiod activities; ~~and to tied activities; the score of an actor~~ a medium score to activities that have to be spread, ^{a low score to activities with only a few mits required,} and a zero score to free periods. The exact weighting coefficients are somewhat arbitrary, and may need to be adjusted in the light of experience.

~~sorted stiffness: \rightarrow size(Activity) \rightarrow Activity sequence~~

~~In addition to a static allocation of priority, (it is worth while to keep a record of activities and periods which during the course of timetabling, become ~~stiffer~~ stiffer during the while to maintain ~~two~~ several grades of stiffness, from the most stiff (N) to the least stiff (1), where N might be 3 or 4.~~

~~stiff: $0..N \rightarrow$ (Activity \times Period) set~~

~~stiff, verystiff: (Activity \times Period) set~~

From within

Thus we introduce a sequence.
~~This requires a sequence~~

5

sorted: Activity sequence,
which holds all activities in order of their a priori difficulty

Thus we can code.

select until $(a, p) \neq \perp$:

~~if sorted = empty
if sorted \neq empty then { a := ~~first~~^{head}(sorted);
if $P(a) - T(a) = \text{empty}$ then { sorted := tail(sorted)
while $p \in P(a) - T(a)$ & sorted \neq empty do
{ if sorted = empty then go to success
{ sorted := tail(sorted)
if 1~~

repeat begin
if sorted = empty then go to success; printout;
a := head(sorted);
if $P(a) - T(a) = \text{empty}$ then { sorted := tail(sorted);
else p := any($P(a) - T(a)$)

until $p \in P(a) - T(a)$;

3.5. Premature Termination

~~A very rough~~

The program as designed so far will always find a timetable if there is one, and report failure if there is not. However, a very cursory calculation shows - that ~~the backtracking~~ ^{may} take a wholly impractical amount of time. ~~to solve it~~ ^{it} would be better when the difficulty of making further progress is too great, it would be better to stop and print out the contents of $T(x)$ as it is so far. On the other hand, it would be a shame to stop at a point when there was every ~~a~~ ^a good chance of ~~proceeding to a successful conclusion,~~ ^{(making good progress towards a successful conclusion.} This suggests that a measure be taken of the inherent difficulty of the current situation, and when some predetermined limit is reached, the timetable will be printed out.

A good measure of the difficulty of the situation is ~~(a count of)~~ ^{indicated by} ~~the number of backtracking steps~~ ^{in the report part.} ~~has taken place~~ ^{of the number} ~~of decisions made~~ ^{can} ~~This~~ ^{can} ~~be~~ ^{estimated} ~~by~~

comparing the current value of the count with the highest value it has ever achieved. Thus we need a variable

countmax: Integer initially 0

~~and~~ after

after count: +1 do if count > countmax then countmax := count

after count: -1 do if countmax - count < limit do

{print T; stop}

~~The Detailed Coding Programming of Consistency Check.~~

4. Details of Consistency Check

The time has now come to carry out the ~~combined tasks~~ of the tasks of detailed ^(programming) coding which have been postponed into the previous sections. ~~The task of this coding~~ with the claim that ~~it was trivial.~~ We must therefore program the consistency check which is to guarantee the truth of (N) and also to take the necessary steps to ensure the truth of conditions FA, FC, and S, and also ~~I1, I2, and I3.~~ ~~All these tasks are~~ to be carried out

during the check of consistency which follows each decision ~~made~~ ^(perhaps 100,000) ~~the check will be made~~ ^{in view of the very large number of decisions involved in a complete timetable} it is most important that this coding program be highly efficient.

The best way to ensure efficiency is to take cognisance of the fact that only ~~the~~ one decision has been taken since the last time the conditions were checked, and so it is necessary only to examine those data which have been changed since the last time (might have been affected by this ^{one} change).

The nature of the change can be discovered by examining the current values of a and p ; and if $p \in T(a)$ the decision was an assignment; if $p \in P(a)$ it was a cancellation.

~~In the program which follows~~

Since the derivation of the program follows so closely from the invariants which it is attempting to ^{establish} ~~restore~~, it is not worth explaining ~~it in detail~~ or annotating it in detail; all that is necessary is to reference the ~~prop~~ relevant invariant at the appropriate place where it becomes true again.

check consistency(a: Activity, p: Period):

begin if $p \in T(a)$ then note the most recent decision was $T(a):+p$;

begin if $p \in P(a)$ then goto impossible;

case times(a) - size(T(a)) of

{ < 0: goto impossible;

= 0: for $p' \in P(a) - T(a)$ do forcedcancel(a):+p' }

N1

FC1

for $i \in$ requirement(a) do

case lives(i) - busy(i,p) of

{ < 0: goto impossible;

= 0: for $a' \in$ users(i) do if $p \in P(a) - T(a)$ then forcedcancel(a):+p;

if $a \in$ spread then for $p \in$ periods in (day of (p)) do if $p \neq p$ then forcedcancel(a):+p

N2

FC2

N3

FC3

(and N3)

if (a in spread & length(a) > 1) then

case times(a) - length(a) - possdays(T(a)) of

{ < 0: goto impossible;

= 0: for $d \in$ possdays(T(a)) do for $p' \in P(a) \cap$ periods in (d) do forcedcancel(a):+p' }

N4

FC4

else note the most recent decision was $P(a):-p$;

begin case size(P(a)) - times(a) of

{ < 0: goto impossible;

= 0: for $p' \in P(a) - T(a)$ do forcedassign(a):+p' ;

= 1: for $p' \in P(a) - T(a)$ do stuff(a):+p' ;

N1

FA1

ST1

for $i \in$ requirement(a) do

case possbusy(i,p) - lives(i) of

{ < 0: goto impossible;

= 0: for $a' \in$ users(i) do if $a' \neq a$ & $p \in T(a)$ then forcedassign(a'):+p' ;

= 1: for $a' \in$ users(i) do if $a' \neq a$ & $p \in T(a)$ then stuff(a'):+p' }

N1

FA2

ST2

if (a in spread & length(a) > 1) then

case possdays(a) - times(a) - length(a) of

{ < 0: goto impossible;

= 0: for $d \in$ possdays(P(a)) do note daydetermined

case size(P(a)) - periods in (d) - length(a) of

{ < 0: goto impossible; comment never happens.

= 0: for $p' \in P(a) \cap$ periods in (d) - T(a) do forcedassign(a):+p' ;

= 1: for $p' \in P(a) \cap$ periods in (d) - T(a) do stuff(a):+p' ;

N4

daydetermined(a,p)

FA4

ST4

if $p \in$ spread & size(P(a) - periods in (day of (p))) = 1 then stuff(a):+p' ;

ST3

end check consistency;

if length(a) > 1 & T(a) & periods in (day of (p)) != empty then

for $p' \in P(a) - T(a)$ do stuff(a):+p' }

ST4

check consistency:

begin

if length(a) > 1 then
begin intersection := ^{periods in} (day of (p)); union := empty;
for p' ∈ starts(length(a) ∧ day of (p)) do
 { ps := tuple(p', length(a));
 if T(a) ⊂ ps ⊂ P(a) then { intersection := ∩ ps;
 union := ∪ ps }
 }

for if p' ∈ intersection - T(a) do forcedassign := +(a, p')

FA4

for p' ∈ P(a) - union do forcedcancel := +(a, p')

FC5

~~(for p' ∈ P(a) ∧ periods in (day of (p)) T(a) do stiff := +(a, p')~~

ST

5. Data Representation.

Now that all the general decisions have been taken, and the program itself written in outline, the time has come to design representations of the data. This design must be made in the light of a knowledge of the most frequent operations on the data, and its likely size, and it seems that at this stage we have ^{most of} the necessary information. But it must also be made in the light of a knowledge of the storage characteristics of the computer on which the program will be run. We will therefore assume a 24-bit word length, and attempt to accommodate all the data in about ten thousand words.

One of the characteristics of this problem is the wide variability of some of the data sizes. For example, most activities require ^{only} one, two, or three items, but some activities may require twenty or more. Thus it will

5.1 Permanent Data.

The storage of data that does not vary during the execution of the main part of the program in general presents little difficulty. However, it does seem worth while to use packed representations to economise on storage. Where several ~~map~~ packed mappings share the same domain, it is usually worth while to combine them into a single mapping onto a Cartesian Product. // The ^(permanent) mappings which have activities as their domain can be represented

A: Activity \rightarrow (times: [0..31]; length [0..3]; spread: Boolean; requirements: (first: [0..255]; if first = 0 then long: Poolpointer else short: [1..4] \rightarrow [0..255]))

- Where
- (1) we will add 1 to length (before using them and times)
 - (2) S is selected so that the short case of users fills the remainder of a computer word. (eg. $S=4$ on a 48-bit word length)
 - (3) ^{by the short case} The zero item-number will stand for the end of the sequence.
 - (4) In the long case, the poolpointer will index the pool Δ at the point where the sequence of users actually begins.
 - (5) Each element of A requires two words, making 1500 words in all
 - (6) Most activities will have less than five items in their requirement, so little space in the pool will be needed.

often: ^(padding) space is to allocate enough storage for the common case; and use a pointer to a common pool of storage in the exceptionally long case. We thus declare // type Poolpointer = 1..2047, and Pool: Poolpointer Δ word where word is just a normal computer word.

The mappings which have Item as their domain can similarly be represented:

L: Item \rightarrow (users: Poolpointer;
lines: [1..31]; ~~used of~~
~~word, busy, possbusy: 1..48~~ \rightarrow [0..15], long: (Poolpointer))

(1) users points to the first of a ^(contiguous) sequence of 12-bit items in the pool, which list the activities which require this item. It is terminated by a zero; and will be typically about five words long. — 1250 in all.

~~(2) busy and possbusy occupy some 16 words for each item. Add one word for lines and users, and the total size of the table will be 4250 words.~~
~~(3) An item with more than fifteen users will present special problems.~~

Since there are less than 48 periods in the week, the obvious way of storing a Period set is as a bit-map representation, using two words. Thus, the arrays T and P can readily be represented, as arrays of these word pairs:

T, P: 1..750 \rightarrow two words.

and will occupy 1500 locations each

The mapping starts will occupy about eight words, periods in about 10 words, day of, 48 words

The mappings forced assign, forced cancel, and stiff could be stored in the same way as simple contiguous stacks. The two forced sets ^{together} are very unlikely to exceed a hundred or two, and although and if they ever do, it would be extremely likely that one of them would generate a contradiction. Thus we would be quite justified in backtracking immediately in case either of these overflow. The stiff list may also get quite long; but if it overflows, a "scavenging" routine can eliminate those entries which ~~are also~~ represent decisions which are already taken or forced; and in the last resort, entries can simply be discarded from the list. ~~the~~

~~forced assign, forced cancel, stiff: 1..100 \rightarrow~~

~~(a: 1..1023, p: 1..48)~~

~~# factop, fctop, stiffop, integer~~

The two forced lists ^{sets} can be combined into a single ^{stack} list, by including a marker with each element to indicate from which set each it comes from. Thus we may ~~allocate~~ use a single array, with ~~to hold~~ ^{one} ~~two~~ ends of a single array to hold both the forced stack and the ~~stiff stack~~ other end to hold the stiff stack. This ~~takes advantage~~ ^{caters for} of the fact that the stiff stack will be largest when the forced stack is empty.

S: 1..300 → (a: 1..1023; p: 1..48),
d: {assign, cancel})

forcedtop, stifftop: Integer
forcedtop := 0;
stifftop := 301;

Thus the total number of data words required is:

Pool	² 2000
A	1500
L	4250
T	1500
P	1500
others:	400
	<hr/>
	11200

Which is just within our target of 12000 words.

The next problem arises from consideration of the depth of the recursion. A complete timetable may require $750 \times 40 = 30000$ decisions, each of them involving a recursion. In most automatic implementations of this would require an excessively large amount of storage. A solution arises from the realisation that the only reason for maintaining the stack is in the event of a backback. But we have deliberately set a limit (perhaps 100) on the amount of backtracking we shall perform. Thus after a hundred recursions, the inactive end of the stack may be overwritten.

This means that the recursion in the program will have to be implemented by explicit stack reference manipulation on a stack which is treated as a "cyclic buffer." Each element of the stack needs to record the activity and period which is affected, and also whether the decision was forced or not.

stack: 1..limit \rightarrow (a: Activity; p: Period; forced: Boolean)

4. Tight sets.

As mentioned above, the feasibility of the timetabling method depends critically on very early detection of inconsistency; for if an inconsistent decision is detected only after ~~n~~ n subsequent assignments, it may take 2^n backtracking operations before the error is corrected. Furthermore, very powerful detection methods for forced decisions are vitally important, since latent inconsistencies can often be detected after a chain of forced decisions, which can comparatively cheaply be backtracked.

We are therefore interested in strengthening the conditions for consistency and forcing; and welcome a suggestion made in [2], namely the search for tight sets. We first note that the following is a necessary condition of consistency:

$$\forall i, \forall as: \text{Activity set } (as \subseteq \text{users}(i)) \text{ size} \left(\bigcup_{a \in as} P(a) \right) \times \text{lives}(i) \geq \sum_{a \in as} \text{times}(a) \quad \text{N6}$$

Proof. The activities in as will require a total of $\sum_{a \in as} \text{times}(a)$ unit periods of item i ; and these must be taken during the periods of $\bigcup_{a \in as} P(a)$. But if there are too few such periods, this will be impossible. N7

If equality in N6 obtains, the set as is said to be tight in i .

$$as \in \text{tight in}(i) \stackrel{\text{def}}{=} as \subseteq \text{user}(i) \ \& \ \text{size} \left(\bigcup_{a \in as} P(a) \right) \times \text{lives}(i) = \sum_{a \in as} \text{times}(a)$$

Note that the empty set and the full set ($\text{users}(i)$) are both trivially tight.

Now it is clear that if as is tight in i , all the period-units of item i will be occupied during $\bigcup_{a \in as} P(a)$ in satisfying the needs of the activities in as ; and none can be spared during these periods for any activity outside as . We can thus derive a sufficient condition for forced cancellation of a from p , where $p \in P(a)$:

$$\exists as, i \quad as \in \text{tight in}(i) \ \& \ a \in \text{users}(i) - as \ \& \ p \in \bigcup_{a' \in as} P(a') \quad \dots \quad \text{FC7}$$

Note that FC7 can be true only if $as \neq \emptyset$ and $as \neq \text{users}(i)$; thus no cancellations are forced by these trivially tight sets.

4.1. Example.

In order to develop a deeper understanding of the nature of a tight set, we shall give an example of a tight set search. We shall initially confine attention to an item with only one life, and suppose that each of its user activities is to occur only one time. For the sake of illustration, we assume size (Period) = 10. Now each value of P(a) may be regarded as a Boolean vector of length 10, with 1 corresponding to each $p \in P(a)$, and 0 for each $p \notin P(a)$:

eg $P(a) = 1\ 0\ 1\ 1\ 0\ 1\ 0\ 0\ 0\ 0$

signifies that periods 1, 3, 4 and 6 are members of P(a). Now the requirement that an item be kept fully busy during 10 periods implies that it must have exactly 10 users (since our simplification states that each user uses the item exactly once). Consequently the rows for each of the users may be extracted to form a square Boolean matrix.

Periods users	p1	p2	p3	p4	p5	p6	p7	p8	p9	p10
a1	1	0	1	1	0	1	0	0	0	0
a2	1	0	1	1	0	1	0	0	0	0
a3	0	1	0	1	1	0	0	0	0	0
a4	1	0	0	1	0	1	0	0	0	0
a5	1	0	1	0	0	1	0	0	0	0
a6	1	1	1	0	1	1	0	0	0	0
a7	1	1	1	1	0	1	1	1	0	0
a8	1	1	1	1	0	1	1	1	0	0
a9	1	1	0	0	0	1	1	0	1	0
a10	0	0	0	1	1	1	0	1	1	1

The first fact to note is that a_1 to a_6 form a tight set, with $U P(a)$ comprising the first six periods (columns). This may most clearly be recognised by noticing that there is a solid 6×4 rectangle of zeroes on its right of the 6×6 square on the major diagonal. The rule FC7 now permits cancellation of all 1's in the corresponding bottom left hand rectangle, leaving the following:

1	0	1	1	0	1						
1	0	1	1	0	1						
0	1	0	1	1	0					0	
1	0	1	1	0	1						
1	0	1	1	0	1						
1	1	1	0	1	1						
						1	1	0	0		
						1	1	0	0		
						1	0	1	0		
						0	1	1	1		

Now it is clear that the ~~complement~~ complement of a tight set in the matrix is also tight - after the cancellation has taken place.

But a_1 to a_6 is not the only tight set in this matrix. For example in the bottom right hand square the first two activities form a tight subset, and the bottom left corner of that square should be blanked out.

p7p8p9p10;

a7	1	1	0	0
a8	1	1	0	0
a9	0	0	1	0
a10	0	0	1	1

Now a_9 has only one possible time when it can be assigned; this is in fact a special case of a tight set, and justifies yet another cancellation:

		p9	p10
a9	1	0	
a10	0	1	

Of course, it cannot in general be expected that the rows or the columns of a tight set will be contiguous, as they were in the cases described above. However, if there is a tight set, its rows and columns could be made contiguous by suitable interchange. For example, in the top left square of the original matrix, activities

a1, a2, a4 and a5 form a tight set, with a union containing p1, p3, p4, p6. By inter change of rows and columns we obtain:

	p1, p3, p4, p6	p2, p5
a1	1 1 1 1	0 0
a2	1 1 1 1	0 0
a4	1 0 1 1	0 0
a5	1 1 0 1	0 0
a3	0 0 1 0	1 1
a6	1 1 0 1	1 1

Thus four ones in the bottom left-hand corner of this square should be cancelled.

This reasoning applies also if some activities are intended to occur more than once. An activity which is intended to occur n times may be regarded as equivalent to n identical rows, each of which is to occur once (thus a4 and a5 in our example might be a single activity; also a7 and a8).

If an item has multiple lives, say n lives, the Boolean matrix will be n times as long as it is wide; and the "squares" along the diagonal will be rectangles, also n times as long as wide. Apart from this, the reason given above applies also to this case.

Since searching for tight sets is an extremely laborious business, we do not wish to do it too often. In particular, we can avoid doing it in the case of forced decisions; instead we wait until all forced decisions have been taken and the next unforced decision is due; consequently the tight set search should occur just before "select unforced(a,p)", and if one or more forced decisions have been uncovered by the tight set search, one of these should be selected.

before select unforced(a,p) do {tight set search; select forced(a,p)}

4.3 Reduction of Inefficiency.

In view of the extremely time-consuming nature of the "scan" procedure, it is necessary to take firm steps to reduce the frequency with which it is called, and the size of the sets on which it is to operate. Suppose the set of users of an item i have suffered no cancellation since they were last scanned for tight subsets. Then there is no point in making a further scan. If one or more user activities have suffered cancellation since the last scan, then any tight subsets within users(i) must contain at least one of those cancelled activities. This suggests that we keep an account of all activities cancelled since the last scan, together with all items which need rescanning.

changed: Activity set initially empty
 needscan: Item set initially empty.

whenever $P(a): \neg P$ do {changed: + a; needscan: + requirement(a)}

Now we can code:

tight set search:

```

begin   for  $i \in$  needscan do
          begin   ts := users(i);
                  scan ts;
                  needscan: -i
          end;
  changed := empty;
  success: end

```

where scan ts: for $a \in$ ts \wedge changed do {scan(unitset(a), P(a), times(a));
ts: -a }

Further efficiency can be gained if we remark that:

- (1) cancellation can never cause a tight set to become nontight
- (2) any new tight set must be wholly contained in some previously existing tight set.

Point (2) may be established by visualising the diagonal form of tight sets. Since it is much more efficient to scan two separate sets than their union, it will pay us to record any tight sets as we discover them, and confine future searches to these individual tight sets. For this purpose we introduce a variable

tss: Item \rightarrow Activity-set sequence initially tss(i) = unitsequence(users(i))

which maps each item onto a sequence ^{of sets which are} ~~those elements~~ are tight in that item.

Whenever a tight set is discovered, both it and its ~~relative~~ complement with respect to the set being searched must be added to the appropriate tight set sequence

in scan before go to success do

tss(i): ^ as; tss(i): ^ (ts-as)

Whenever backtracking occurs, any tight sets discovered as a result of a changed decision must be removed.)

Thus the item for which a tight set has been discovered must be recorded in a variable local to progress

tightset found: Boolean initially false

item with tightset: Item

in scan before go to success do

tightset found:=true; item with tight sets:=i

before impossible do

if tightset found then remove two elements from tss(item with tight set)

Now the tight set search for item i involves scanning through all sets in tss(i) and selecting those which contain at least one changed activity. However, once a tight set containing a given activity has been scanned, there is no point in scanning a subsequent tight set containing that activity. We therefore keep in a variable "changed" that subset of the users(i) which have been changed but not yet dealt with.

Now we can code a more efficient version of

```
tightset search:
begin for i ∈ need scan do
  begin new ichanged initially changed ∧ users(i);
    for ts in tss(i) while ichanged ≠ empty do
      for a ∈ ts ∧ ichanged do
        {scan ts; ichanged:-a} ;
      needscan:-i
    end;
  changed:=empty;
success: end
```