THE AIM OF THE PROPOSED TITAN SUPERVISOR

1.0 INTRODUCTION

The Proposed Titan Supervisor differs in certain respects from the Atlas Supervisor. This document therefore discusses why the proposed system is recommended and shows why the differences are necessary. A companion document "A Brief Description of the proposed Titan Supervisor" describes the system in outline.

The hardware of Titan uses much of the engineering work already done for Atlas. In addition, Titan is meant to be compatible with Atlas from the user's point of view, and one of the functions of the "built-in" programmes is to achieve this as far as possible. Furthermore, Titan is a large and expensive machine and hence there seems good reason to supply, as with Atlas, a Supervisor which will enhance the effective performance over and above that of rival computers with the same basic instruction-rate. On the face of it therefore one might expect the Atlas Supervisor to be adopted for Titan.

However, there are certain fundamental hardware differences between Atlas and Titan, for instance:-

(a) no large main-store on Titan

(b) different lock-out facilities

(c) no Fixed Store on Titan

(d) fewer tape-channels on Titan

The Atlas Supervisor therefore needs considerable alteration to work at all on Titan; and if such alterations are essential, one wonders if further alterations would not make it more suitable - for the Atlas Supervisor was designed for Atlas, and it does not follow necessarily that it is the system most suited to Titan.

Therefore before we opt for any particular system, we consider with an open mind all the possible alternatives (they are discussed in this document), and if we are to make valid comparisons, we must first of all establish proper criteria. The arguments presented here about what one should offer to customers were presumably followed in the case of Atlas as well; if they are repeated here, it is because hardware differences force us to review them again in the new context of Titan.

- continued -

2.0 SUMMARY OF THE ARGUMENTS

Firstly there is a point which is peculiar to Titan and which applies to all possible systems (even to "no system at all"). Atlas has 9K of Fixed and Subsidiary Store in lieu of conventional hardware. Titan has neither. Hence on Titan some of the Core-Store must be regarded as part of the hardware and sold as such (later arguments show that about 8K is needed). This applies to any system; the only way round it is to add either a separate store of 8K or some conventional hardware, and the only way to sell it is to demonstrate that it is the better way.

The systems considered fall into 4 broad categories, which can be represented thus:-

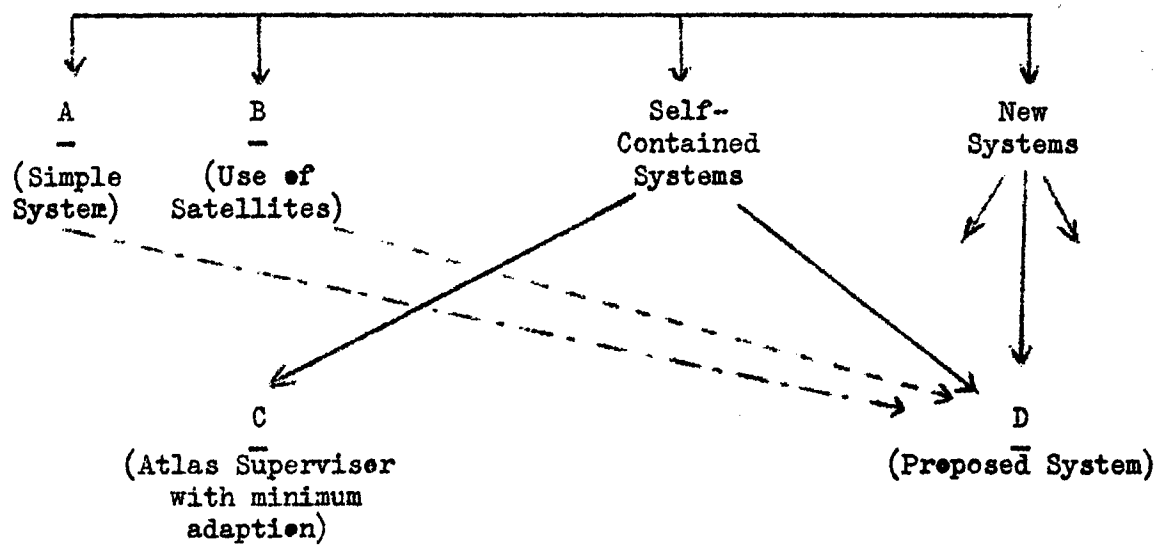

A, B and C are rejected, as general systems for Titan, for the following reasons. A is hopelessly inefficient; B is expensive and of limited efficiency - though there are certain special circumstances in which it has its uses; C is too expensive in mainstore (= core-store on Titan).

A new approach is made, therefore, to find the system most suited to Titan. The system proposed, D, is in fact generally similar to the Atlas System; this is because many of the programming techniques of the Atlas System are either found to be suitable for Titan in their own right, or are necessary to preserve compatibility.

- continued -

D is a comprehensive system - potentially it is adaptable to any situation, and in each situation it claims to give the best possible value for money. Though it can contract to A or to B for those special circumstances in which A or B are desirable, it is normally a self-contained system, handling and time-sharing all activities within the central machine, and using as much equipment (to increase efficiency) as is appropriate to the situation. When it contracts to A, it uses less core-space than would be required if there were no system at all.

3.0 A CONSIDERATION PECULIAR TO TITAN

Coming now to detailed arguments, we note first that in comparing systems there are two distinct points at issue, and failure to keep them separate will result in misleading comparisons. They are:-

a) The use of software in lieu of hardware

b) Handling documents and programmes (i.e. the "System")

We discuss the first point in this section and the second in the remainder of the document.

On both Atlas and Titan, software is used in lieu of hardware, both to control peripherals and to provide extra functions (and also in lieu of library storage if it be considered that library subroutines are the proper alternative to some of the extracode functions). The cost of thus reducing hardware is:

*(x% of computer time handling peripherals.
(
(nK of storage and working space for the routines.

Whether the decision to save hardware in this way was correct is not discussed here, but clearly it is most important that the Atlas-Titan method be demonstrated to be cheaper and more reliable (if in fact it is - if it is not, then the hardware should be increased until the best solution is reached; we note that a move in this direction is already being considered for Titan, but it will only reduce x, not n).

* On Titan, (a) x (10 on a 5-/usec machine) per 1000 characters per
( )
(5 on a $2\frac{1}{2}$-/usec machine) sec. read as binary.

reducible to about $2\frac{1}{2}$ by extra hardware now proposed

(b) n ≤ 8

These figures are extremely tentative. n is discussed in the appendix and x in a separate document.

- continued -

Where Atlas and Titan differ is that Atlas has a Fixed and a Subsidiary Store for the built-in routines, whereas on Titan they can only be in core-store. Thus, with any system on Titan, however simple, 8K of core-store must be regarded as part of the hardware (and is justified by the relative cheapness of the machine).

The immediate result of this is that to compare Titan with a 32K, 64K or 128K machine is invalid, and will only result in misleading comparisons. We must regard Titan as a 24K, 56K or 120K machine from the user's point of view, and to do otherwise is to admit defeat - for if the Atlas-Titan method be demonstrated convincingly to be sufficiently better, then it will be accepted as the better approach and programmes will be tailored in future to the limitations inherent in the better approach.

If a Supervisor exists (as opposed to a set of routines handling peripherals) then it is possible to arrange for only those parts of the 8K as are required at any one time to be in mainstore - though it is doubtful whether one could usefully reduce to below about $5\frac{1}{2}$K (see the Appendix). However to do this one must either have at least one system tape-deck, or one must accept the inconvenience of re-inputting the routines through slow peripherals whenever the Supervisor requests them. In other words, hardware considerations impose a minimum Supervisor requirement of:-

either approximately 8K of core-store

or approximately $5\frac{1}{2}$K (increasing to 8K) of core-store plus

(either at least one system-deck

(or considerable inconvenience

4.0 CRITERIA FOR COMPARING SYSTEMS

Having discussed the separate issue of software used in lieu of hardware, and which applies to all systems on Titan, we turn to the criteria by which individual systems for Titan should be compared.

4.1 3 Broad Considerations

The main function of a system is to control the flow of documents and jobs to the satisfaction of both the individual user and the installation as a whole. For simplicity of argument we subdivide this into 3 broad considerations. A system can offer:-

- continued -

a) user-convenience (e.g. library-routines, monitoring facilities, use of multiple data and output streams into and out of a single programme, ability to dump and recall later, ability to "crash" an urgent programme through, etc.)

and b) increased efficiency (by time-sharing to overlap peripheral delays, and by initiating automatically any engineer's routines required)

at a cost of:

c) extra equipment: mainstore (=core-store on Titan) for (system programme, bufferage, time-sharing)
decks and/or discs
tape/disc channels
satellites (if used)

Each of the 3 aspects contains so many variables that we could easily be confused when trying to compare systems. In order to simplify the issue therefore we first of all assume that user-convenience can be provided on any system if at least one deck (or disc) is available for a library or a dump tape: where systems will differ is in the price, in terms of equipment or loss of efficiency, that must be paid for the extra facilities.

Hence to compare systems we consider efficiency and equipment, bearing in mind the effect on these when user-convenience is required.

4.2 Efficiency and Equipment

The relationship between efficiency (which we have not yet defined precisely) and equipment is complicated, for many variables are involved. Attempts to generalise in terms of such concepts as "Value for Money" do not seem to lead anywhere useful, so instead we start from particular situations and work as far towards the general as we need.

- continued -

We start with a definition of "efficiency" which is deliberately limited and relative to the particular situation.

(a) Let a given set of jobs be run on a given set of equipment

(b) Let T be the total time that they actually take

(c) Let t be the total time that they would have taken if only the machine-orders had been obeyed - i.e. with no waits for peripherals or for the obeying of System-Programmes.

(d) Let t/T be called the "efficiency" of the System for this set of jobs on this set of equipment. (In other words, the efficiency is the average rate at which those particular object-programme instructions are obeyed, expressed as a fraction of the basic instruction rate quoted for the machine).

(e) Let the definition of efficiency now be broadened to apply not to a given set of jobs but to a given type of job-mix though still on the given equipment. (The definition of "Job-Mix" is elusive, but the concept is useful).

This is the (limited) yardstick by which we compare systems. If it be accepted as an extremely important criterion, then no more need be said - but in case it is not, we show how relevant it is to two particular situations (which we hope are sufficiently general to cover most cases), then proceed to a more general approach.

The first situation is an installation with given equipment (i.e. given cost) and a given set of jobs to execute; and a System is required which gets these jobs executed as quickly as possible. Thus t is fixed and T is to be minimised. Hence t/T, or the "efficiency", is to be maximised. The System attempts to achieve this by scheduling the time-sharing of the various activities. Clearly T can never quite equal t, because of the time taken to obey System-programme. Hence the efficiency can never quite reach unity - nevertheless it still requires to be maximised. Similarly if there are insufficient peripherals this will imply an inevitable "waiting-time" - but good scheduling can cut it to a minimum.

We note that this situation effectively covers the case of the customer with given financial resources who wants as powerful a machine as possible (in the sense of doing as many jobs as possible per day) provided that the type of job-mix is defined.

- continued -

The second situation is the tender for an installation to execute a given set of jobs in a given time (e.g. per shift or per day) with the minimum equipment. Here T is fixed, and for a given central machine t and hence the "efficiency" are fixed, but the equipment (i.e. the cost) is to be minimised. The better System achieves the required efficiency with less equipment because it schedules the time-sharing more effectively.

In both cases, the System is required to maximise the efficiency/cost ratio. This ratio is a second yardstick for comparing systems - it is still relative to the particular instruction rate and to the particular type of job-mix, but it is sufficiently general to be useful to us.

A slightly more general yardstick, which potential customers with a given job-mix in mind would do well to apply, is:

$$\frac{\text{instruction-rate} \times \text{efficiency}}{\text{cost}}$$

Thus the instruction-rate of a computer has little meaning until one knows the efficiency of the system offered with it for the job-mix envisaged. This is a point which rival manufacturers offering poorer systems will try to conceal - and they will find it easy to conceal unless we exploit it, for the average customer knows little about systems and about efficiency in this sense. It is a point of weakness in rival machines which could be exploited with great advantage.

To return to given central machines, the most practicable approach to system-designing would seem to be:

a) to confine oneself to the type of job-mix envisaged.

b) for each possible configuration, to maximise the efficiency (By implication we minimise the configuration required for a given job-rate)

c) By comparing the results for the various configurations, to study the effect of adding equipment, to see if:

   (i) it improves the efficiency (for the customer who wishes to increase the power of his machine)

   (ii) it improves the efficiency/cost ratio (for if it does then it has increased the machine's "Value for money")

- continued -

4.3 Levels of Efficiency

Broadly speaking, there are 3 main levels of efficiency:

a) No time-sharing - except built-in overlap of central and peripheral instructions.
At Titan speeds this is extremely inefficient (except in those special cases - if they exist at all - where the ratio of computing to input and output is so high that it is acceptable to wait for slow peripherals). It is doubtful whether machines offering only this would compete seriously with Titan, except in the special case mentioned.

b) Time-Sharing of Slow Peripherals
To Titan this is probably the most important form of time-sharing. If it is achieved successfully, then the only delays are those caused by private tape-transfers and by the loading-time of a job from tape to core. Those machines which achieve it successfully could seriously rival Titan.

c) Complete Time-Sharing
Programmes are time-shared with each other, hence all peripherals can be time-shared. On Atlas, with its drums, this is important; on Titan, with only tape-transfer overlap to achieve, it is not so important - though clearly there will be situations, occurring quite frequently in some installations, where programme-changing to avoid tape-delays will significantly improve efficiency.

On some systems (the Atlas type in particular) time-sharing of programmes can improve slow-peripheral time-sharing.

5.0 COMPARISON OF ALTERNATIVE SYSTEMS

5.1 A: Simple Systems

The first type of system we consider is the Simple System.

It is difficult to see how any time-sharing can be achieved without either satellites or some subtleties in the System; therefore we can define the Simple System as one which attempts no time-sharing. Such a system is so inefficient, except in special circumstances, that it is hardly worth considering for the general case.

If, however, at little extra programming cost, it could be a special mode of the general system, then some advantage would be gained.

- continued -

We note that because of lack of both Fixed Store and conventional hardware, even a simple system must occupy at least:

(8K of core
(or $5\frac{1}{2}$K (optionally increasing to 8K) + (library - tape
(or inconvenience
+ separate decks for compilers

User-convenience would cost at least one deck (it could be the deck mentioned above), and a certain amount (say $\frac{1}{2}$K) of core space for system-programme.

5.2 B: Satellite System

The second system we consider is the Satellite System. This transfers the bulk of input/output from the central computer to a satellite (or satellites), and so in effect time-shares the slow peripherals. Therefore it must be considered seriously for Titan. However it does not lend itself to complete time-sharing and hence it will be inefficient with tape-limited jobs, and it is expensive in equipment.

In ideal circumstances it would need a minimum equipment of:-

a) On Titan 1 data deck
1 output deck
(say 9K of core (plus space for engineer's tests)
(or say $5\frac{1}{2}$K of core (plus space for a minimum
(.. of engineer's tests).
plus (library deck
(or inefficiency and inconvenience.
separate decks for compilers.

b) A Satellite with at least one deck.

In fact, it is very doubtful if the ideal circumstances ever exist: when they do not, the equipment needed increases considerably. For instance

(a) Any reasonable data and output rate will require several satellites - this point is taken up again in 7.0, but in any case it would repay further investigation.

(b) If more than one input and output medium are to be available, then there needs to be either one deck per peripheral-type on Titan, or a Satellite System powerful enough to merge input streams suitably and to sort mixed output streams -

- continued -

unless one is prepared to allow grossly inefficient utilisation of slow peripherals, which again is expensive.

(When several satellites are needed in any case to deal with the data rate, it would be possible to attach one type of peripheral to each and swop the tapes around whenever a new document was reached - but this is a very inconvenient method of operating unless it be automatic - again at extra expense).

(c) whenever jobs are tape-limited, the system will be inefficient unless complete time-sharing is adopted. This however relies, on any system, on a dynamic approach to scheduling: with a satellite system, dynamic scheduling would require several extra decks on the central machine, to sort documents satisfactorily.

So far we have not considered user-convenience. This, too, would be expensive:

(a) Multi-stream input and output can only be provided if Titan has either one deck per stream or several extra decks to sort the documents, or (in the case of output only) the Satellites are powerful enough to sort mixed streams (in this case the streams will overlap each other on the tape)

(b) library and monitoring require an extra deck (whether this be classed as a system or a compiler deck is irrelevant - they cost the same!); and extensive use of it would be inefficient because of the tape-search time.

(c) the turn-round time for a job is long, measured probably in hours (though admittedly this will not always matter).

(d) the "crashing" through of an urgent job is either ineffective or expensive:

either (i) one waits for the present job to finish or

(ii) with 1 extra deck one waits for tape-rewinding and deck-reloading time, or

(iii) one needs 2 or possibly 3 extra decks.

(e) dumping and later recalling require an extra deck (whether this be a system or a private deck is irrelevant)

- continued -

5.3 C: Atlas Supervisor with Minimum Adaption

The third alternative we consider is a minimum adaption of the Atlas Supervisor (which on Atlas gives complete time-sharing and as much user-convenience as required).

However, because of hardware differences, it would need considerable alteration to work at all on Titan; and even then it would be unsuitable, for it was designed for machines with large and relatively cheap main-store. On any Titan it would be wasteful of core-store (= mainstore on Titan), and on a 32K Titan it would leave little room for object programmes.

5.4 The Search for D: Possible New Systems

We have seen that Simple Systems are hopelessly inefficient, Satellite Systems are expensive for what they give, and the adapted Atlas System will not run on a 32K Titan.

However, Titan has basic hardware which is suitable for a self-contained system; and in addition many of the ideas and programming techniques of the Atlas Supervisor are available for use on Titan, for instance:

(a) general approach to time-sharing and programme-scheduling

(b) a method of handling independant branches of the Supervisor (as embodied in the Co-ordinator)

(c) the philosophy of the Input and Output well

(d) the "creeping-swing" method of using Magnetic Tape.

We therefore consider whether it is not possible to make use of this basic work (and the available hardware) to devise a system suitable for Titan, which is:

(a) more economical and/or more efficient than the Satellite method

i.e. (i) time-shares slow peripherals with less equipment.

(ii) is more efficient and/or convenient with the same amount of equipment.

(iii) offers even greater efficiency and convenience for those prepared to allocate more equipment (if this be necessary)

- continued -

(b) more economical of mainstore than the Atlas Supervisor

We cannot expect to be more efficient than Atlas (or even as efficient in many cases) but we can (and have to) try to use very considerably less mainstore. We also have to use fewer channels (though this does not affect any general scheme) and it will benefit the installation if we use as few decks as possible (it would be reasonable to expect an increase in the number of tape-decks, in lieu of mainstore not used; but by more exhaustive use of tapes, we try to keep the number down as much as possible).

As a general point, efficiency and convenience cost equipment, and different customers will be able to afford (or will wish to pay for) different degrees of it. Ideally therefore one might wish to offer a range of complexity, at the lower end a small Titan little more than a conventional machine, at the higher end a machine equivalent in every practical respect to a small Atlas.

6.0 THE PROPOSED SYSTEM

6.1 Basic Principles

Historically, therefore, work on the Titan Supervisor started as a search for the system which gives the greatest efficiency with the least equipment, an open mind being kept both about the use of Satellites or Atlas-type methods and about the use of fundamental Atlas thinking; and the basic question was posed:

"With nK words of store and m tape-decks available to the Supervisor, what is the most efficient system?"

The question was asked for various values of n and m; and with maximum efficiency postulated in each case one could well have arrived at a large number of very different systems. Fortunately, however, the most efficient systems for each of the various situations are remarkably alike and from the programming point of view fit into one general pattern. Thus, assuming that in each case the best solution has in fact been derived, one has a comprehensive system which:

(i) gives the best value for money for each individual installation

(ii) with alterations to parameters only, adapts itself to new situations (again giving the best solution).

- continued -

The same question can be posed in two other forms, and it follows that, each way, the proposed system gives the best solution so far suggested:

(a) "For a given degree of efficiency, and nK words of store, how few decks are needed?"

(b) "For a given degree of efficiency, and m decks, how little core-store is required?"

The question could also be generalised:

"With nK words, n decks and p access-channels to a disc with qK capacity.....?" Where the addition of a disc gives a more economic answer, there is a strong case for including it, though whether the change from no discs to one disc is as simple as from m decks to (m+1) remains to be seen.

The basic question was posed originally with scientific installations in mind (that is, with a relatively low data-rate, and with the majority of documents relatively small). The result was a system like the Atlas one except that:

(a) the allocation of mainstore space is different

(b) greater use is made of off-channel tape-movement (theTitan hardware will facilitate this)

(c) the Atlas method of using system tapes is extended to the concept of a "generalised Well".

Since then, further thought has been given to installations handling higher data-rates and larger documents, and the result is a composite system with 2 basic modes of operation which work in parallel, the bias being towards one or other according to circumstances. The second mode, for large documents, is nearer to the satellite method (in that documents are copied direct to a "Bulk-Tape") but Titan acts as its own satellite, time-sharing this activity with normal execution. Being nearer to a satellite method, this could in certain cases be expensive in decks, but no better way has yet been devised. The comprehensive system is described in the companion document.

The question remains: "Has the most efficient system been found in every case?" In a sense this can never be proved- and it may be disproved by the discovery of a better system. The superiority of the proposed system over the satellite method is discussed more fully in 7.0; but it should be mentioned here that there is one situation where the basic question may not have been answered satisfactorily, namely when less than 10K of store is available to the Supervisor for the <u>majority</u> of the time but plenty of decks are available (say at least 6) - however it is argued in 7.0 that in this case it is better to rephrase the basic question.

- continued -

6.2 Static Flexibility

Because the proposed system is adaptable to different situations, each customer can be given the system most suitable for his job-mix and as efficient a system as he is prepared to pay for.

In particular this means that until his final daily load is known, he can manage with a small installation (with a Supervisor which is less efficient - but if the machine is underloaded this does not matter). But provided he leaves space for more core or more decks, he can add either or both when the load increases and in doing so will effectively speed up the machine.

This static flexibility complicates the planning of an installation, for only experience will show exactly how many decks are needed in given situations. But potentially this flexibility is available, and further study should show how to make use of it (an initial study is presented in the companion document).

In effect, a new degree of freedom is introduced into the planning stage - one need not make use of it, but if one does a more efficient machine will result.

6.3. Dynamic Flexibility

The proposed system not only gives the most economic solution so far suggested, for each installation or situation, but it also adapts itself smoothly and rapidly from one mode to another.

Thus:

Core: Supervisor Core-Space can be relinquished (at a loss of efficiency), or increased (to gain efficiency) at very short notice.

Decks: Except in the case of decks currently reading or writing large documents, all system tapes can be relinquished or accepted in a period comparable to the external operation of mounting a deck. That is to say, the internal change-over from one mode to the next takes usually a few seconds and rarely more than a minute; to this must be added tape rewind time simultaneous with the settling-down time for the System. When tapes are switched from one system use to another (e.g. a change from scientific to commercial bias) remounting is not normally necessary.

- continued -

Channels: The Supervisor shares all channels with the object programmes. When it is the only user, it gains in efficiency; when it is not, it will normally wait its turn for each transfer (with consequent loss of efficiency).

It is suggested that this dynamic flexibility (which can be made automatic, in response to job-descriptions) will help to cut the cost of an installation, for:

(a) the maximum use is made of all three items - any item not required by object programmes is immediately appropriated by the Supervisor, to increase the throughput of the machine, until required again by programmes.

(b) in many practical cases it is probably more economic to have a smaller machine which is less efficient when the bigger jobs go through and more efficient with the smaller ones, than a larger machine which takes the bigger job with ease and has idle equipment when the smaller jobs run.

It will be seen that the Atlas philosophy of using peripherals to their best advantage, by scheduling to keep them evenly loaded, is extended to main-store, decks and channels.

Earlier it was stated that, because Titan lacks both Fixed Store and conventional hardware, 8K of core must be regarded as part of the hardware and only 24K, 56K or 120K as object programme core-store. If the proposed system is adopted in full, then this statement should be extended:

(a) 8K of core is in lieu of conventional peripheral hardware

(b) nK of core is in lieu of faster instruction rates

and (c) Titan should be regarded, from the user's point of view, as a (24-n)K, (56-n)K or (120-n)K machine.

Owing to dynamic flexibility, however, n is variable and hence not easy to define. Perhaps the most meaningful simplification would be to say that n = 8, nominally, but that the system arranges for core-space to be borrowed indiscriminately between the three users. Thus for instance, when programmes do not need their full share of space, the System will borrow the remainder to attempt to increase still further the effective instruction rate; and conversely programmes are enabled to borrow from the system or even the peripheral share, at the possible cost of lowering the effective instruction-rate. Alternatively, one could say that even the mean value of n can be set to the customer's wishes - according to whether he prefers core-space or efficiency.

- continued -

6.4 Tapes and Discs

As was mentioned in 6.1, there is no reason why discs should not be included in the general philosophy, and hence used by the Supervisor when they are available.

However the answer to the question "Should £x be spent on discs or tape-decks?" is not immediately obvious and almost certainly depends on the situation. There are probably situations where a disc is better than the equivalent (cost-wise) number of decks; on the other hand a tape-based system seems to be more flexible than a disc-based system, especially at the lower end of the scale (the tape-based Supervisor will often be able to manage temporarily, and in some situations permanently, with 2 or even 1 deck).

Three alternatives appear to be open, in descending order of programme-complexity:

(i) A Supervisor which can use tapes, discs or both, as they are available to it. This would give a choice:

a) (statically) to the individual installation

b) dynamically, to the Supervisor

(ii) Two separate versions of the Supervisor, one using discs (plus optional tapes), the other using tapes (plus possibly an optional disc). This would still give the static choice (and a "semi-dynamic" choice also, if a change-over period initiated by the operator be acceptable).

(iii) One only, of these two versions.

7.0 THE PLACE OF SATELLITES

Because Satellites are an established rival to the Atlas-Titan approach, and because there are certain situations where they are useful, we now discuss in more detail the relative merits of Satellites and the proposed system.

7.1. Where Satellites could be useful

(a) Satellites are useful as remote stations. Whether or in which circumstances they are better than data-links is not discussed here.

- continued -

(b) Adding satellites can increase the effective speed of an existing machine - though this can also be done, in the case of Titan, by adding core-store or decks, which in most cases will be the cheaper and more effective way (further study of the point is recommended).

The only case where satellites are undoubtedly the cheaper way (because they are the only way!) is when a Titan already has 128K of core and as many decks as it can handle - and it will only be effective in this case on large documents which do not require sorting. In this situation, it would have been better to have bought an Atlas in the first place!.

(c) The Titan system must be allowed to expand to at least the 12K range for much of the time if it is to be efficient. Therefore when a potential customer -

(i) already has a large stock of programmes all written for 24K, and hence has good reason to require the Supervisor to restrict itself to 8K for most of the time.

and (ii) is prepared to have plenty of system-decks

then some of those system-decks could with advantage be on a set of satellites.

However, a cheaper solution would be to add 32K of store and reduce the number of system-decks. If much of the extra 32K were regarded as Supervisor Space, then the number of system-decks could be reduced considerably and the extra space would still be available for occasional programmes. In any case the efficiency of the system would be increased. (More detailed study of relative costs is suggested, in order to convince customers).

7.2 A Comparison of Costs

In all other cases the proposed system seems to be cheaper than the satellite method. The cost of both, for a given efficiency, varies with the number and the size of documents, so the only fair comparison is to take typical and particular situations.

(a) Large documents: the proposed system is very similar to a satellite system, except that the function of the satellite is time-shared, internally, with execution. Though there will often be a saving in decks it is unwise to assume so in the general case, hence the main advantage is that x% of central computer time to control peripherals is cheaper than the basic cost of the complete set of satellites (one or more satellite according to the data rate). This point requires further study, for it is linked with the question of whether the Titan peripheral hardware should be increased.

- continued -

(b) Small documents

Assuming that no sorting is required (i.e. each job has one input and one output stream, and only one external medium is in use) then the use of satellites for input and output costs:

2 decks on Titan	£25K
1 satellite per 1000 ch/sec	£100K per 1000 ch/sec

(these figures are to an order of magnitude)

Thus one gets the following very approximate comparison:

	Proposed System	Satellite
at 500 ch/sec	1 deck plus 1½% of computer time	£30K	£125K
at 1000 ch/sec	2 " " 3% " " "	£55K	£125K
at 2000 ch/sec	3 " " 6% " " "	£100K	£225K
at 3000 ch/sec	4 " " 9% " " "	£140K	£325K

(assuming that the proposed additional hardware is added

This comparison, however, assumes no complications. If external media are mixed, or multi-stream input and output is required, or tape-limited programmes imply a need for complete time-sharing, then the figures for the proposed system remain, but the satellite cost increases by the addition of several extra decks on the central machine.

7.3 Miscellaneous Disadvantages of Satellites

(a) The Satellite system is less flexible, for a self-contained system can when the need arises devote itself completely to an extra large job. In other words, for the same capital cost, the self-contained system has more core-store and decks available for the occasional large object programme.

(b) Though the proposed system is potentially adaptable to the satellite mode, work needs to be done to implement this:

(i) engineering work on, say, Argus
plus programming work on Argus and Titan.

and/or

(ii) extra programming on Titan to allow the System to accept or produce tapes handled by existing satellites. (e.g. the IBM 1401)

- continued -

The extra programming on Titan may not be implemented unless requested. To accept a low average rate of data from a distant station employing Atlas-compatable decks requires virtually no programming, for the data off the tape can be absorbed into the system as if it came from a slow peripheral. If, however, large quantities of data arrive or depart on magnetic tapes, then it will be necessary to handle the documents in the most efficient manner. This may need more effort - especially if efficiency is required in cases where satellites are not the recommended mode. Also, to accept tapes from say an IBM 1401 requires routines both to handle on IBM deck and to interpret the format of the information when it is in core.

7.4 Owners of Existing Satellites

In a new installation there is little question that a self-contained system should be offered except where remote stations are involved. However, customers who already own a satellite present another problem. If the satellite mode of the proposed system is implemented, then it would be tempting to condone the retention of the satellite, and its integration into the system, in order to make a sale. However -

(a) this is not the best way to use Titan : the customer would do better to sell his satellite if he can get a good price for it, (a further study of prices and throughput is suggested).

(b) manufacturers who have not sunk money into developing elaborate systems will clearly do their utmost to play down the merits of the Atlas-Titan approach. Every case of a customer persuaded to sell or to return his satellite will be one more psychological victory in the struggle between rival systems.

7.5 Further study of the Satellite System

The Satellite System has been studied sufficiently to show that as a general system it is inferior, but that there may be situations where the satellite mode would be worth implementing. Any further study is irrelevant to the writing of the Supervisor, and should therefore be regarded as a separate effort. However, such a study is strongly recommended, for:

(a) we need to know more exactly the conditions, if any, for which the Satellite mode should be offered.

- continued -

(v) detailed study of its implications would strengthen the evidence against it. This is important, for the "battle of the Systems" concerns the future even more than it does the present. The Atlas-Titan approach will have an even clearer advantage when central computers are faster still, when discs are cheaper and more reliable, and if and when core-prices come down. Hence it is to the advantage of rivals to oppose its general acceptance (by fair and unfair means!), and to our advantage not only to exploit it but more important still to maintain the lead we now have.

Points worth following up have been indicated in the text; and the companion document is intended to provide a starting point for any detailed comparison of systems.

8.0 SUMMARY OF THE PROPOSED SYSTEM

The Titan Supervisor is adapted from the Atlas Supervisor, to which it owes many of its ideas and features. The main new emphasis is on flexibility - from very modest time-sharing, requiring very little equipment, in a small installation, to the elaborate time-sharing that is justified in a bigger installation - and also between different types of load.

It is designed to use whatever equipment (core-space, tape decks etc) that is available at any instant of time, in the way which will best speed the throughput of the machine; moreover it can adapt itself smoothly and quickly to changing situations. The advantages of this are threefold:-

1) All parts of the computer are always being used to the best advantage - any piece of equipment that would otherwise be idle is used to increase the throughput.

2) An installation need only be planned for the Supervisor to be modest when the biggest job is going through - but once this is through the Supervisor will automatically expand and give all the benefits of time-sharing until another big job arrives.

3) An installation can be planned modestly at first, assuming that peripheral delays (which cannot be avoided on a conventional machine, and can be so costly if the central computer is a fast one) can be tolerated. But when the load increases, the effective throughput in the machine can be progressively increased merely by adding tape-decks or core-store.

- continued -

Titan can accept magnetic tapes loaded with input documents, and produce tapes loaded with output - either for use with satellites, or for communication with other machines. But satellites are an expensive means of time-sharing input and output, and the internal System on Titan gives the same result at lower cost.

Whatever the size of the installation, all the features (except use of widely-scattered addresses) which make Atlas so easy to use are present in Titan; it is only the effective speed of the machine which suffers when the Supervisor is modest.

CRS/JT
11th January, 1963.

APPENDIX : MINIMUM CORE REQUIREMENTS

The following are estimates of the minimum core-space required by the Supervisor, and they are to be taken as extremely tentative.

	Simple System		Satellite System		Proposed System
	A	B	A	B	B
Co-ordinator	1 K	1 K	1 K	1 K	1 K
Extracodes and Programme-control (jump-table routines monitor	2½K	1½K	2½K	1½K	1½K
Magnetic Tape (control organisation monitor	2 K	2 K	2½K	2 K	2 K
Slow Peripherals (control organisation	2 K	½K	2 K	½K	½K
Operator-communication	½K	½K	1 K	½K	½K
Total (less Engineer's Tests)	8 K	5½K	9 K	5½K	5½K

A : Assumes routines cannot be fetched from tape or slow-peripheral.

B : Minimum, assuming that the remaining routines can be fetched from tape or slow-peripheral.

The last column shows the minimum to which the proposed system can contract. To quote an uncontracted figure would be irrelevant : the System occupies "5½K + whatever space is not currently occupied by object programmes" and for efficiency it is recommended that for the greater part of the time it be allowed at least, say, 12K (though this figure, like the others, is extremely tentative).