FLETCH

IBM

Washington Systems Center

Technical Bulletin

Attached Processing (AP) System Performance Characteristics and Considerations

By W. W. White, SPD

Published by T. E. Giblin, WSC

DAPS Code 0910/0911

GG22-9004-00

May 1977

Washington Systems Center

Gaithersburg, Maryland

Technical Bulletin

Attached Processing (AP) System

Performance Characteristics and Considerations

by W. W. White, SPD

This Technical Bulletin is being made available to IBM and customer personnel. This presentation was given at GUIDE 44 in San Francisco, May, 1977. It has not been subject to any formal review and may not be a total solution. The exact organization and implementation of the functions described will vary from installation to installation and must be individually evaluated for applicability.

A form is provided in the back for comments, criticisms, new data, and suggestions for future studies, etc. This Bulletin is being published by Mr. T. E. Giblin of the Washington Systems Center.

DAPS Code 0910/0911

GG22-9004-00

May, 1977

SPECIAL NOTE TO ALL READERS

ON EACH CHART SHOWING PERFORMANCE GRAPHS OR INDICATING PERFORMANCE RESULTS, THE PERFORMANCE RESULTS WERE ACHIEVED IN A CONTROLLED LABORATORY ENVIRONMENT AND MAY NOT BE APPLICABLE TO AN INDIVIDUAL USERS ENVIRONMENT BECAUSE OF DIFFERENCES IN JOB MIX, CONFIGURATION, ETC.

IF YOU PLAN TO USE THIS BULLETIN AS A FOIL PRESENTATION, YOU SHOULD ENSURE THAT THIS INFORMATION IS CONVEYED TO ALL VIEWERS.

Attached Processing (AP) System Performance

Dr. William W. White
IBM Corporation

1. Introduction (Foil 1)

The IBM Attached Processing Systems are recent additions to the IBM product line. The 168 APS is currently available, while the 158 APS, already announced, will be available shortly. For both these systems, the internal performance is quoted as 1.5 to 1.8 times the internal performance of the corresponding uniprocessor under the MVS release current at ship time, and using identical configurations and programs. Throughput is approximately that of the corresponding asymmetric multiprocessor under MVS. Actual results achieved will, of course, depend upon the

- .multiprogramming capability of the jobstream
- .multiprocessing suitability of the jobstream
- .adequacy of the I/O configuration
- .adequacy of the storage configuration

in use (Foil 2).

This presentation will discuss the performance of IBM's Attached Processing Systems to provide the background behind the above performance statements, as well as to develop an appreciation for the performance characteristics of the general family of asymmetric tightly coupled multiprocessors to which the 158 APS and 168 APS belong. After reviewing some general terms and definitions which will be used throughout the presentation, some of the performance considerations of tightly coupled processing will be discussed, and the implications of these considerations on overall system performance. Since the APS is a special case of tightly coupled processing, these general considerations and their implications will carry over to APS as well (Foil 3).

Some of the special characteristics of asymmetric processing performance will be presented, following which some results specific to APS performance will be reviewed. There will then be some discussions of the general points one might consider when going to an APS (principally from a uniprocessor environment), and a short summary will complete the presentation.

The orientation of this presentation is to provide a higher level understanding of APS performance. The statements and conclusions are based on results obtained in a laboratory environment, and on analysis of those results. The extrapolation to any given user environment, particularly at the detail level, will have varying degrees of validity, depending upon the particular user environment, although the general thrust of the presentation remarks should carry over.

2. Some Terms and Definitions

During the course of this presentation, we will make use of some terms and definitions specific for this presentation, among them being (Foils 4, 5):

UP (Uniprocessor): a single processor with associated I/O and storage

MP (Multiprocessor): two tightly coupled processors sharing the
- .storage of both processors
- .I/O of both processors

AMP (Asymmetric MP): two tightly coupled processors sharing the
- .storage of <u>only</u> <u>one</u> of the processors
- .I/O of <u>that</u> <u>same</u> processor

AP (Attached processor): two tightly coupled processors where one processor shares the
- .storage
- .I/O.

belonging to the second processor.

Logically, an AP and AMP are similar in function and behavior. There can be performance differences, however, since an AMP is basically a MP configured to be asymmetric (e.g., by enabling only one processor's storage, and by varying the other processor's channels off line), while an AP is designed and built to operate asymmetrically, thereby achieving cost economies.

In addition, the terms HD (half duplex, or one half of an MP, including processor, main storage and (I/O), and cross-configured HD (a half duplex with some main storage of both processors enabled to the single HD processor) may occur. While both of these configurations operate logically like a UP, there can be performance differences due to details of implementation.

Furthermore, it will be necessary on occasion to distinguish between each side of an MP, AMP or AP configuration. To facilitate this differentiation, we will employ the term "Base" to refer to that side of an MP, AMP or AP which includes a processor, main storage, and I/O (channels, and the phrase "Attached" to refer to the side of AMP or AP which includes just the single processor. Thus, an MP consists of two "Bases's", while an AMP or AP have a "Base" side and an "Attached" side each.

3. Tightly Coupled Processing Performance Considerations

Since an APS is a special case of the more general tightly coupled multiprocessing, many of the performance characteristics of MP will carry over to asymmetric multiprocessing. Among these considerations are those arising from hardware, system programming and the applications themselves.

One of the principal hardware factors (Foil 6) is contention for storage, where the two processors or a processor and channels compete for access to main storage. This contention is eased through the use of high speed buffer storage, interleaving of storage (although this is not implemented on the 158), and through the use of storage "selection" algorithms to establish priority between competing components for storage. For this, as well as for other hardware factors, no new design is needed for AP; the current design carries over directly.

Similarly, many of the system software factors are also shared by AP and MP. For example, mechanisms for interprocessor communication have already been designed and implemented in MP (Foil 7). Since for RAS reasons, it may be necessary for an MP to function asymmetrically, the AP just employs these same capabilities directly. In particular, either processor can be awakened by the other if it is currently in wait state but the other processor notes that there is some work which is dispatchable. Similarly, an AP, for I/O purposes, functions like an MP whose channels on one side are always busy, but where alternate pathing is available via the other side.

A second software factor is interprocessor synchronization, needed in any tightly coupled processing configuration, for instance, to prevent the simultaneous modification/access of vital system information (Foil 8). The design and implementation of synchronization procedures, e.g., via instructions such as Compare And Swap, or via Locks, is as valid for AP as for MP, and works as naturally for AP as it does for MP.

A requirement for good performance of any tightly coupled processing configuration is that the workload be structured so that parallelism can be exploited (Foil 9). This is more than just a multiprogramming consideration, in that multiprogramming switches between dispatchable units, whereas multiprocessing actually executes dispatchable units in parallel. The structure for parallelism exists within MVS via its use of dispatchable units such as System Request Blocks (SRB's) for system work, and Task Control Blocks (TCB's) for user work. While MVS itself uses SRB's to accomplish parallelism, it is necessary that application code employ TCB's so that MVS can exploit the dispatching of parallel units of work. This is naturally accomplished in batch by the use of multiple initiators, and in TSO in which each user has his own address space. Other subsystems, such as IMS (at the appropriate release level), are structured so parallelism can be exploited, but occasionally there will be a subsystem for which this is not so--in this case it will be difficult to achieve the performance potential at either MP or AP unless other dispatchable units of work can be added to the system as well.

4. Tightly Coupled Processing Performance Implications

The considerations described in the preceding section carry some implications as to the performance of tightly coupled processing (Foil 10). In particular, while the internal performance of an AP or MP is 50% to 80% better than that of the corresponding UP, the overhead (both hardware and software) in managing parallelism prevents one from achieving full 100% internal performance betterment. This internal performance effect carries over to throughput as follows.

The busy time for any single job will be longer on an MP or AP than on the corresponding UP (although parallelism saves the day on a system basis). In particular, in the case where busy time tracks internal performance, a single job could run 18% longer on AP or MP than a UP for a ratio of 1.7 (where 18% comes from this 70% factor by the reciprocal of one-half of 1.7).

As noted, parallelism enables the jobstream to be processed substantially faster in MP or AP than in UP. There can, however, be a wide variation in processing time ratios, depending both on the amount of parallelism inherent in the workload as well as on what bottlenecks and utilizations are present on the UP or the AP/MP.

5. Performance Characteristics of Asymmetric Processing

There are a number of characteristics of asymmetric tightly coupled processing, be it on an AP or an AMP, which derive specifically from the asymmetric nature of the processing configuration.

In particular, there are certain functions which must be performed on the "Base" side of the configuration. One of these, for example, is the fielding of I/O interrupts (Foil 11). Since these functions are naturally "reserved" for the "Base" side, the "Attached" side is free to take a larger share of the remaining workload. Since almost all such "reserved" functions execute in supervisor state, the result is that there is generally a higher proportion of supervisor state code executing on the "Base" side than on the "Attached" side and conversely, a higher proportion of problem program code executing on the "Attached" side. The actual amount of shift depends on the particular workload and configuration in question. In any case, however, this is just a shifting of activity, as the total supervisor state for both sides is about what it would be for a symmetric MP.

A second effect of asymmetricity on internal activity comes from the fact that the "Base" side has more interrupts than the "Attached" side, since it has the I/O (Foil 12). Thus, there is a longer residency (before being interrupted) of dispatchable units on the "Attached" side than the "Base" side. The lower frequency of interrupts on the "Attached" side concomitantly results in a more efficient use of the high speed buffer, i.e., a higher buffer utilization (or BHR--buffer hit ratio) on the "attached" side than on the "Base" side, although again the amount of this effect is dependent on which workloads and configurations are in operation at the time in question. Once more, the shift in activity is internal to the system--as a system, the total amount of activity is roughly the same on an AP/AMP (summed over both sides) as on an MP.

A third effect of asymmetricity on internal activity is specifically a property of 158-based models (Foil 13). This is because, for a 158, "cycle stealing" by the channels from the processor takes place for I/O operations. For an AP (since there are no channels on the "Attached" side) or an AMP (since the channels on the "Attached" side are inoperative), no channel interference takes place on the "Attached" side, and all the cycle stealing takes place on the "Base" side. While this does provide more execution time on the "Attached" side, it also causes less execution time on the "Base" side. However, once again, total interference is conserved when comparing asymmetric to symmetric processing. There is no additional interference on an AP/AMP than on an MP--it just appears on one side.

The asymmetric nature of the I/O processing suggests that, since I/O's issued from the "Attached" side must be handled by the "Base" side, there might be some internal limiting factors to the volume of I/O traffic which can be supported in tightly coupled asymmetric processing (Foil 14). However, some stress tests were performed on a 168 environment which show that factors such as SIGP's for I/O and internal queueing for I/O are not significant. In these tests, the configuration and workload was selected so that, if a bottleneck (and resulting lowered I/O rates) occurred, it would be due to limiting factors within the processors. A sustained I/O rate of over 900 EXCP's a second was achieved on a 168 AMP. This was within 3% of the peak rate achieved on a symmetric MP, and was well over the highest known I/O requirements for MVS installations for real workloads. The natural adaptation of MVS for asymmetric processing carries no significant internal limitations vis a vis high I/O requirements.

On a 158, high I/O activity can have other effects, specifically in terms of data rate and channel capabilities (Foil 15). In particular, the increased "horsepower" of a 158 AP compared to a 158 UP can generate heavier data rate requirements compared to the UP. A stress test was performed comparing an AMP to a UP on a configuration which had an aggregate data rate of 4.8 MB/sec, including 4 channels of DASD and one of tape. This stress test is such that 99% of the processing power of the UP configuration is devoted to driving I/O. The AMP sustained a SIO rate and an average data rate 50% higher than the UP for page-size blocks, but, as expected, also had a higher rate of overruns, with one overrun every 800 or so SIO's, compared to the UP's one overrun every 5000 SIO's. However, no overruns occurred on the tape channel, and since retry is performed at the channel and the control unit for these DASD devices, the effect on an overrun was a missed revolution, or a "response time" effect, during which the system was busy performing other activity. The net result was that, for each second of processing .0047 seconds were spent for missed revolutions, or less than 1/2 of one percent of the time on an AMP. The overruns were minimal and the system degradation was not significant in this instance. Of course, different effects could appear in different environments, and could depend on various factors such as different aggregate data rates, peak channel utilization and how often and for how long at a time it occurs, and on the amount of chaining, particularly data chaining, present.

6. AP Performance

We have seen how AP performance can depend both on the general characteristics of tightly coupled processing, and on the specific characteristics of asymmetric processing. However, the AP is an even more specific form of asymmetric processing, and one can make specific performance statements about AP.

In terms of internal performance, the stated aim of an AP is 1.5 to 1.8 times the internal performance of the corresponding UP, under the release of MVS available at ship time, using identical configurations and programs (Foil 16). Internal performance, for these purposes, is measured in MIPS, or Millions of Instructions Per Second. For an AP (as for an MP or AMP), the MIPS is the sum of the MIPS for each processor in the configuration. Laboratory benchmarks provide support for these claims: it has held true for the 168 APS, and recent tests on the engineering model 158 APS have actually provided MIPS ratios (AP MIPS divided by UP MIPS) of 1.6 to 1.9 in the laboratory environment. These laboratory benchmarks include a spectrum of batch only runs, from COBOL-type environments to FORTRAN environments, as well as more general TSO/BATCH and IMS/BATCH environments.

It is, of course, not necessarily true that system throughput tracks internal performance (Foil 17). For throughput, however, we have seen that an AP should behave similarly to an AMP, both being asymmetric tightly coupled processing configurations, under MVS using identical configurations and workloads. Laboratory experiments have indeed substantiated this statement, as does one's intuition. This has held true for 168 APS and, once more, recent tests on the engineering model 158 APS shows that this holds true for the tested environments. These tests included on extensive comparison in a TSO/BATCH environment, in which the 158 APS and a Model 3 158 AMP track very closely while the proportion of batch and TSO activity was varied by bringing on different numbers of terminals (it should be noted that the AMP in this case required the use of one meg of storage of "Attached" side to be enabled to both processors; such cross configuration of storage can provide a performance variation from a "pure" AMP). An implication here is that AMP benchmarks can indeed provide guidance as to AP throughput behavior.

However, one might ask whether or not AP system throughput follows that of a symmetric MP. This, of course, need not be the case, since the AP cannot have the full main storage or channel configuration of an MP. However, if the I/O and main storage configurations are adequate for an AP, then AP throughput may be similar to that of a symmetric MP under MVS using identical programs and configurations (here, configurations for I/O refer to control units and devices, as if one varies the channels on one side of the MP offline). This has held true for 168 APS experiments, and some recent testing on the engineering model 158 APS shows that this is true here as well (Foil 18). Laboratory experiments have shown this both for an IMS/BATCH workload and for a more extensive TSO/BATCH comparison (as described in the above paragraph--but here, cross-configuration is not a factor). Basically, what this shows is that as long as any bottlenecks, if they occur, are those of processing power (and not channels or main storage) then, as one would expect, AP throughput is not substantially different from MP throughput.

One might also ask how AP throughput relates to UP throughput. Here, the question is in terms of making use of the increased processing power of the AP. Obviously, if the UP configuration is close to its limitations in terms of I/O or main storage, then the full AP potential may not be realized without a concomitant "scaling up" of these resources as well. Laboratory tests on certain benchmark environments have shown that the throughput potential can largely be achieved (Foil 19). For example, on an engineering model 158 APS compared to a cross-configured 158 (HD) where one meg of storage on the second tightly coupled processor was enabled to the

HD processor--with an associated degradation in internal performance), laboratory tests in both an IMS/BATCH and a TSO/BATCH environment shows a throughput increase somewhat similar to the increase in internal performance, keeping the same proportion of workload mix on both the HD and the AP. Here, both HD and AP used 4 megs of main storage and 3 DASD channels, with the processor utilization in the high 90% range on all configurations--response times were about the same in the TSO/BATCH environments, while, for the IMS/BATCH environments, the response time was 1.8 seconds for the AP and 1.3 for the HD (but it was also 1.8 for the corresponding MP test, indicating that the difference was not due solely to asymmetric processing).

7. Considerations in Going to AP

Clearly, the AP provides a significant increase in processing power as compared to a UP. However, to achieve the throughput benefits from this increased processor power, it may be necessary to scale upward the other components of a system (Foil 20). The identification of which components, if any, and the degree of scaling needed, if any, will depend, vis a vis a UP configuration, where the current operating points of the UP are (e.g., current workload mix, utilization of system resources, etc.), and in which direction growth is anticipated. In particular,

- will the workload mix be the same as it is now (with proportionate increase in system needs)?
- will more background batch be added (with the possible increase in processing power needed, but smaller increase in other system resources)?
- will more interactive work be added (with a possible requirement for more storage or I/O capability)?

Individual adjustments will differ depending on many factors, including the answers to the above questions. With this in mind, there are some points which are worth mentioning.

First of all, is the workload itself suitable for AP (Foil 21)? We have seen that an AP is a special case of a tightly coupled processor. However, if starting from a UP, the required parallelism might not be present. Here MVS itself should be of assistance--if dispatchable units of work (TCB's and SRB's) are present, then MVS itself will handle the execution of parallel activities, and will naturally dispatch the work to take advantage of multiprocessing, asymmetric or not. In fact, due to the level of granularity at which MVS can dispatch units of work, it is probably better to let MVS make its own decisions rather than to attempt to put one's own estimations of processor performance and specialization into effect (e.g., by using affinity), at least until the behavior of any particular workload in any particular environment can be assessed. Since, as we have seen, MVS can handle I/O activity effectively via internal I/O queueing and SIGP's, asymmetricity should not be a user concern in terms of work dispatching in most cases.

Secondly, an assessment should be made as to the I/O adequacy of the envisioned configuration and its anticipated workload (Foil 22). It has been noted in laboratory experiments that, if the same workload mix is to be maintained, the total channel utilization (summed over all channels) will increase in about the same proportion as throughput. Thus, if 5 channels in a UP configuration have a utilization averaging 25% each, which sums to 125%, then, keeping the same workload mix, an AP may need to have channels operating at 190-225% total utilization, and it would take 2 or 3 additional channels to bring average channel utilization below 30%. Of course, if the proportion of activity will be more compute bound, there will be less of a channel requirement, while the converse will be true if the direction is toward more heavy I/O activity.

If an increase in overall channel utilization is forseen, it may be desirable to add channels (so as to more equitably spread the increase between channels, maintaining a lower average channel utilization) or to take steps to reduce overall channel utilization. This could be effected by workload adjustment in terms of its I/O load, by moving to devices which are more efficient (thus lowering the length of time needed to transfer data, e.g., 3350's as compared to 3330's), or by making sure one is operating at the most recent software levels (recent MVS improvements to result in lowered channel utilizations). Exactly which, if any, of these suggestions is appropriate must depend on the conditions of the given configuration and workload under consideration, and must be evaluated in such a context. In some instances, a realistic solution may not be available, as could be the case when one is operating his 158 UP with high utilizations on all 5 block multiplex channels--here, a 158 MP may be more appropraite than an AP if one wishes to go to a workload with higher I/O activity. Similarly, if a spreading of I/O across channels is desired for response time considerations, an MP with its higher complement of channels may be attractive.

It should also be noted that other I/O adjustments may be desirable. This could include both physical considerations, e.g., adjustments to or increases in control unit and device configurations, and structural considerations such as pack of data set placement or catalog splitting, etc.

Similarly, an assessment could also be made with regard to the adequacy of the main storage, since, as with I/O, main storage requirements may also increase (Foil 23). Note that, however, since only a single copy of the SCP is needed for the AP, the increase, if any, is better estimated with respect to the current "user" requirements, e.g., an increase from 4 to 5 megs may be equivalent to a 59% increase in storage if one assumes the SCP needs 2 megs (before and after the change). Again, for workloads with singificant main storage requirements, the limitation in AP may make an MP more appropriate.

8. Recent Activity in Support of Tightly Coupled Processing

As the diverse usage of tightly coupled processing, and the bottlenecks which contribute to performance degradation, continue to be better understood, it can be expected that both AP and MP will be made applicable to a broader range of user environments and that there will be improvements in the performance of MVS in tightly coupled processing. The 158 APS, soon to be developed, illustrates this, as does the Engineering Change EC 717728, which is introduced to create better implementation compatibility between 168 AP and MP.

For software, the performance improvements benefit MP and AP alike, as shown by SU's 4, 5 and 7--and, the reduction of channel utilization usually seen with these SU's is a specific benefit to AP. Further improvements are anticipated as a result of the recently announced MVS/System Extensions Program Product, in combination with the System/370 Extended feature, where reductions will take place both in interprocessor interference and in lock contention for certain locks.

These activities illustrate the continued search for opportunities to improve the effectiveness of tightly coupled processing.

9. Summary

In this presentation we have reviewed some performance aspects of tightly coupled processing in general, and with this perspective, discussed a number of performance aspects to AP (Foil 25). The following points are worth reiterating:

.MP performance is a good indicator of AP performance

.MVS tightly coupled processing design naturally supports AP

.AP internal performance is 1.5 to 1.8 times that of the corresponding UP

.System throughput performance of AP requires

-Parallelism in workload
-Adequate I/O
-Adequate main storage

.Asymmetric I/O is well supported

-No internal bottlenecks
-High SIO activity can be supported
-High data rate can be supported

.Recent announcements improve tightly coupled processing performance potential

AP is being recognized as an effective solution for many users with increasing throughput requirements.

ATTACHED PROCESSING (AP) SYSTEM

PERFORMANCE

158 ⟶ 158 APS

168 ⟶ 168 APS

W. W. WHITE
IBM CORPORATION
P.O. BOX 390
POUGHKEEPSIE, N.Y. 12602

MAY 6, 1977

TOPICS TO BE COVERED

- Performance Statements
- Some Terms and Definitions
- Tightly Coupled Processing Performance Considerations
  - Hardware
  - Programming
  - Applications
- Implications of Tightly Coupled Processing Performance
- Performance Characteristics of Asymmetric Processing
  - Internal Migration of Activity
  - I/O Characteristics
- AP Performance
  - Internal Performance
  - Throughput
- Considerations In Going To AP
- Summary

APS PERFORMANCE

INTERNAL PERFORMANCE

1.5 TO 1.8 TIMES THAT OF THE CORRESPONDING UNIPROCESSOR

- UNDER MVS
- USING IDENTICAL CONFIGURATIONS AND PROGRAMS

THROUGHPUT

APPROXIMATELY THAT OF THE CORRESPONDING ASYMMETRIC MULTIPROCESSOR UNDER MVS

ACTUAL RESULTS ACHIEVED DEPEND UPON

- MULTIPROGRAMMING CAPABILITY OF THE JOBSTREAM
- MULTIPROCESSING SUITABILITY OF THE JOBSTREAM
- ADEQUACY OF THE I/O CONFIGURATION
- ADEQUACY OF THE STORAGE CONFIGURATION

THE PERFORMANCE CHARACTERISTICS OF AN APS ARE BASICALLY MULTIPROCESSING CHARACTERISTICS

SOME TERMS FOR THIS PRESENTATION

UP (UNIPROCESSOR)

- A SINGLE PROCESSOR WITH ASSOCIATED I/O AND STORAGE
- CONCEPTUALLY:

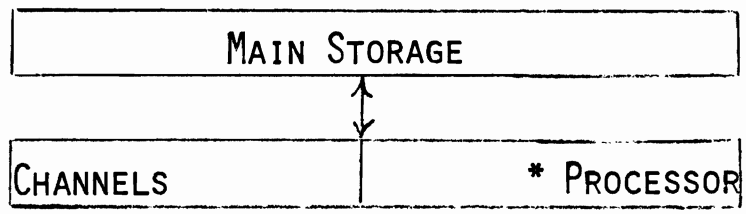

MP (MULTIPROCESSOR)

- TWO TIGHTLY COUPLED PROCESSORS SHARING THE
  - STORAGE OF BOTH PROCESSORS
  - I/O OF BOTH PROCESSORS
- CONCEPTUALLY:

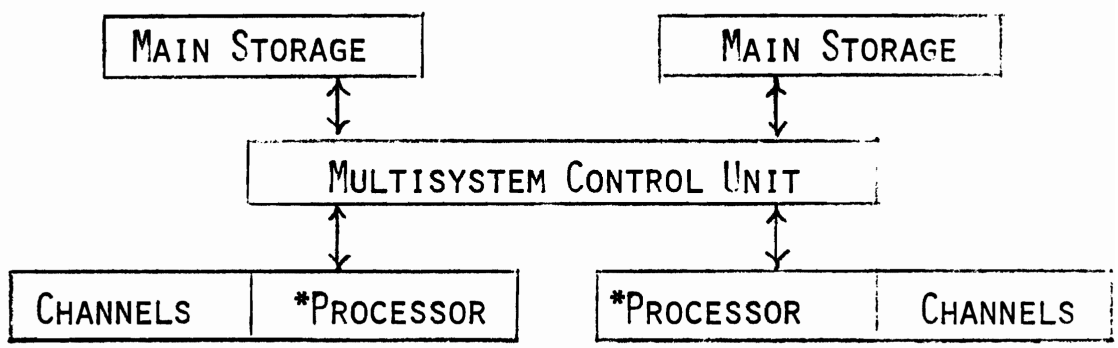

*ALL OTHER FUNCTIONS OF PROCESSOR EXCEPT THOSE SHOWN SEPARATELY

SOME TERMS (CONTINUED)

AMP (ASYMMETRIC MULTIPROCESSOR)

- TWO TIGHTLY COUPLED PROCESSORS SHARING THE
  - STORAGE OF ONLY ONE OF THE PROCESSORS
  - I/O OF THAT SAME PROCESSOR
- CONCEPTUALLY

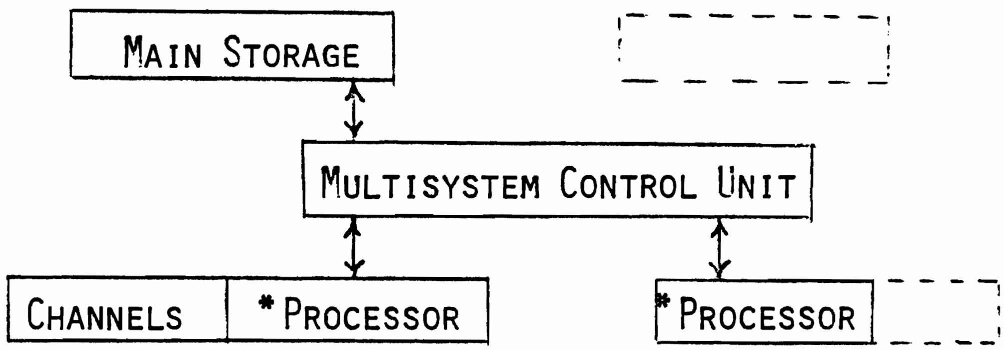

AP (ATTACHED PROCESSOR)

- TWO TIGHTLY COUPLED PROCESSORS WHERE ONE PROCESSOR SHARES THE
  - STORAGE
  - I/O

  BELONGING TO THE SECOND PROCESSOR
- CONCEPTUALLY

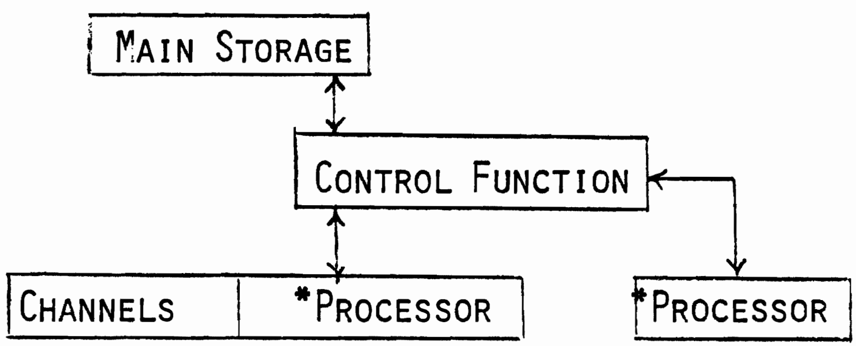

AN AMP IS CONFIGURED TO OPERATE ASYMMETRICALLY

AN AP IS DESIGNED AND BUILT TO OPERATE ASYMMETRICALLY

TIGHTLY COUPLED PROCESSING PERFORMANCE CONSIDERATIONS

HARDWARE FACTORS

STORAGE CONTENTION

- PROCESSOR VS PROCESSOR, PROCESSOR VS CHANNEL

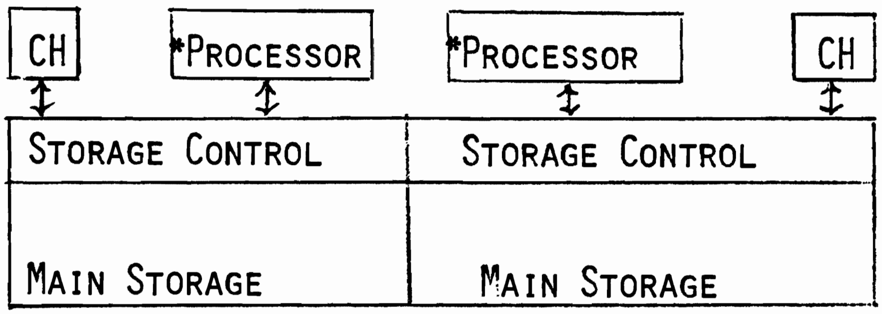

- EASED BY
  - BUFFER STORAGE
  - INTERLEAVING (168)
  - STORAGE 'SELECTION' ALGORITHMS

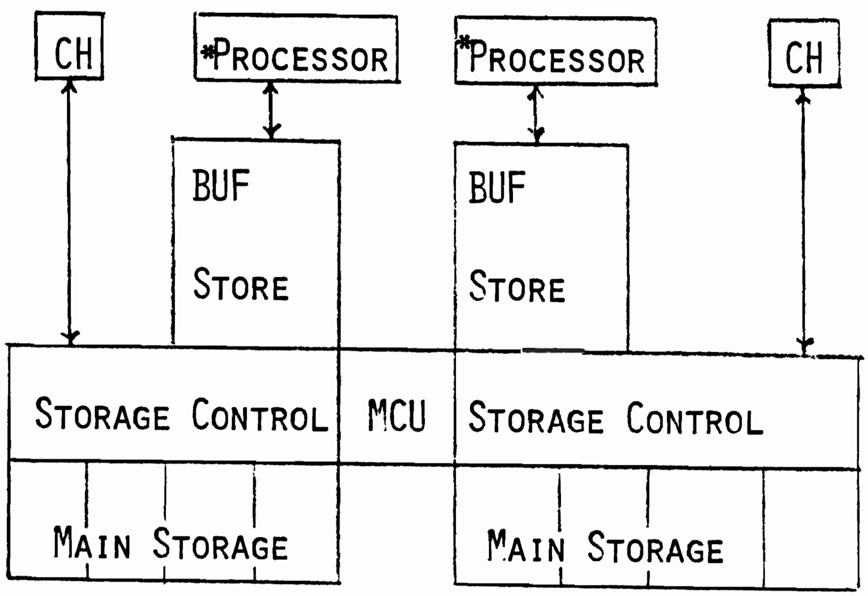

NO NEW DESIGN NEEDED FOR AP; CONCEPTUALLY

- REMOVE ONE SET OF CHANNELS
- REMOVE MAIN STORAGE ON SAME SIDE

TIGHTLY COUPLED PROCESSING PERFORMANCE CONSIDERATIONS

PROGRAMMING FACTORS

INTERPROCESSOR COMMUNICATION

- DISPATCHING/PROCESSOR ACTIVATE

IN WAIT STATE

IN RUN STATE - FINDS WORK WAITING

- I/O

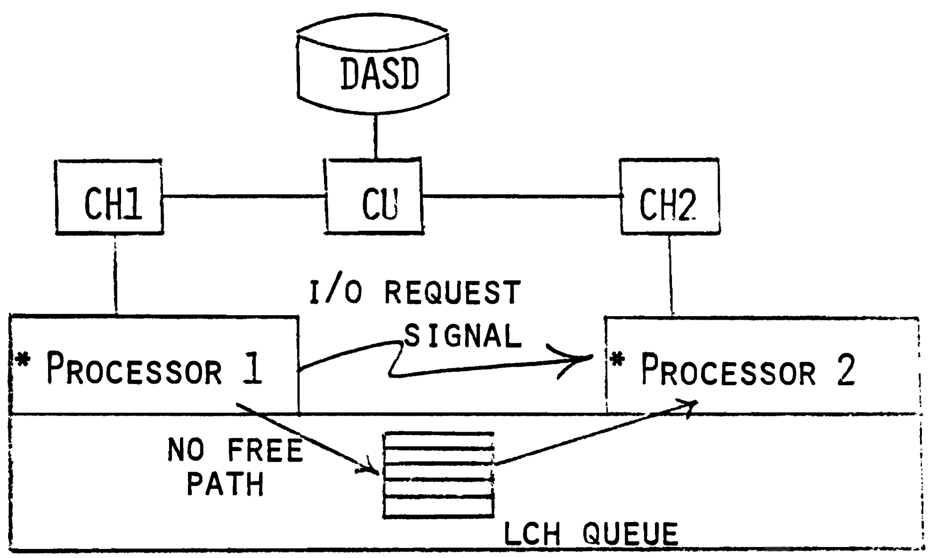

PROCESSOR 1 WANTS TO DO I/O BUT CH 1 IS BUSY

THESE CAPABILITIES EMPLOYED DIRECTLY BY AP

- E.G., CH 1 IS ALWAYS UNAVAILABLE OR "BUSY"

*ALL OTHER FUNCTIONS OF PROCESSOR EXCEPT THOSE SHOWN SEPARATELY

TIGHTLY COUPLED PROCESSING PERFORMANCE CONSIDERATIONS

PROGRAMMING FACTORS (CONTINUED)

INTERPROCESSOR SYNCHRONIZATION: PREVENT SIMULTANOUS MODIFICATION/ACCESS OF VITAL SYSTEM INFORMATION VIA

- INSTRUCTIONS, SUCH AS COMPARE & SWAP
- LOCKS, E.G.

PROCESSOR 1	PROCESSOR 2
SET DISPATCHER LOCK	
SEARCH WORK QUEUE	ENTER DISPATCHER
INSERT WORK ELEMENT IN QUEUE	SPIN ON DISPATCHER LOCK
RESET DISPATCHER LOCK	SET DISPATCHER LOCK
ENTER DISPATCHER	REMOVE TOP ELEMENT

THIS WORKS NATURALLY IN AP

- NO CHANGES ARE NECESSARY

TIGHTLY COUPLED PROCESSING PERFORMANCE CONSIDERATIONS

APPLICATIONS

PARALLELISM IN WORKLOAD VIA MULTIPLE DISPATCHABLE UNITS

- TASK CONTROL BLOCKS (TCB) FOR USER WORK

SINGLE TASK
ADDRESS SPACE

MULTI TASK
ADDRESS SPACE

MULTIPROGRAMMED
BATCH

MULTIPLE USER
TSO

IMS

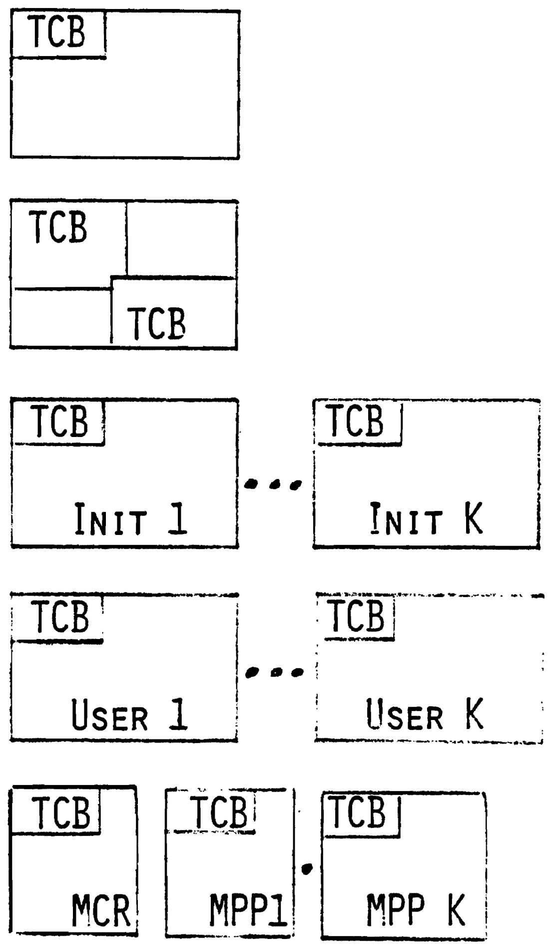

NOTE

- MULTIPROGRAMMING SWITCHES BETWEEN DISPATCHABLE UNITS
- MULTIPROCESSING EXECUTES DISPATCHABLE UNITS IN PARALLEL

AP NATURALLY EXPLOITS PARALLELISM

IMPLICATIONS OF TIGHTLY COUPLED PROCESSING PERFORMANCE

Internal Performance

- Is 1.5 to 1.8 times that of corresponding UP
- Overhead in managing parallelism prevents achieving twice UP performance

Busy Time for a Single Job

- Is longer on MP/AP than on UP
- If busy time tracks internal performance, this can be 18% elongation, for example, for a 1.7 ratio

System Throughput

- substantially greater than that of a UP
- Variation higher or lower, depends on
  - Workload parallelism
  - UP utilization, bottlenecks

Conceptual Example (Batch):

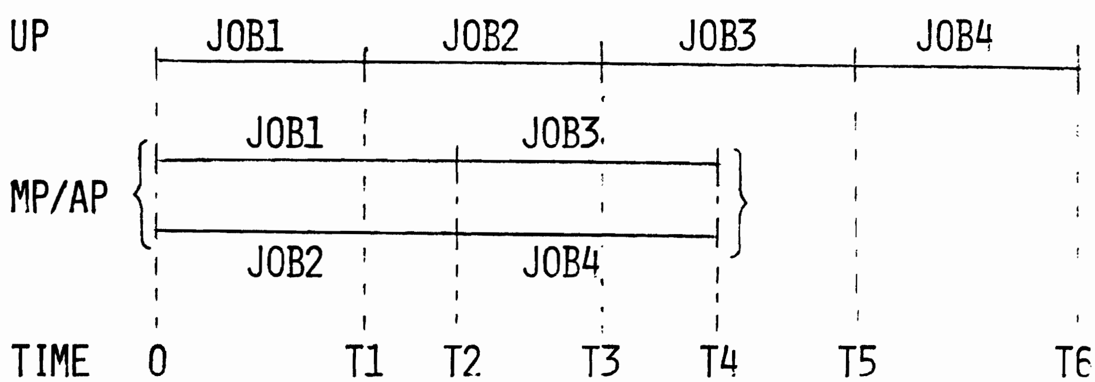

If time tracks internal performance, and we use a 1.7 ratio,

- THEN T6 is 70% longer than T4, "System View"
- AND T2 is 18% longer than T1, "Individual View

PERFORMANCE CHARACTERISTICS OF ASYMMETRIC (AP OR AMP) PROCESSING

INTERNAL MIGRATION OF ACTIVITY

THE SIDE WITH I/O ("BASE" SIDE) MUST PERFORM CERTAIN FUNCTIONS, E.G., FIELDING I/O INTERRUPTS.

CONSEQUENTLY:

- "BASE" SIDE MUST SPEND TIME EXECUTING THESE FUNCTIONS
- THEREFORE SIDE WITHOUT I/O ("ATTACHED" SIDE) HAS MORE TIME FOR OTHER FUNCTIONS
- SO THAT
  - ° MORE SUPERVISOR STATE ON "BASE" SIDE'
  - ° MORE PROBLEM PROGRAM STATE ON "ATTACHED" SIDE
- EXAMPLE OF PROCESSING TIME SHIFT

SYMMETRIC PROCESSING ASYMMETRIC PROCESSING

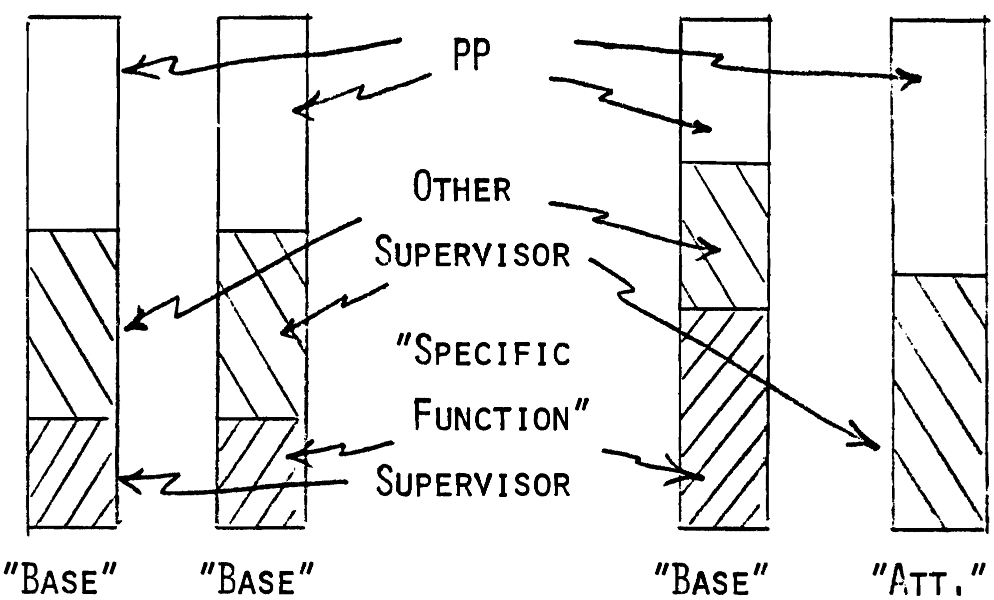

AMOUNT OF SHIFT DEPENDS ON WORKLOAD AND CONFIGURATION

TOTAL SUPERVISOR STATE IS ABOUT THE SAME FOR EQUIVALENT WORKLOADS AND CONFIGURATIONS FOR BOTH SYMMETRIC AND ASYMMETRIC PROCESSING.

Performance Characteristics Of Asymmetric (AP or AMP) Processing

Internal Migration of Activity (cont'd)

The "BASE" side has more interrupts, since it has the I/O. Thus:

- longer residency of dispatchable units on "attached" side
- better utilization of high speed buffer on "attached" side (an improved "buffer hit ratio" - BHR)
- example: high speed buffer efficiency (BHR, as a %) for selected laboratory benchmarks

	For 158				For 168			
	Symmetric	Asymmetric			Symmetric	Asymmetric		
	"Base"	"Base"	"Att"	"Base"	"Base"	"Base"	"Att"	"Base"
Batch 1	97	95	98	93	99	99	99	98
Batch 2	85	81	88	79	97	97	98	96
Batch 3	93	90	96	87	95	95	96	94
Tso/Batch	84	82	86	80	97	97	98	96
Ims/Batch	84	80	89	77	96	95	97	94

Amount of shift of activity depends on workload and configuration.

Total amount of activity is about the same for equivalent workloads and configurations for both symmetric and asymmetric processing

Performance Characteristics Of Asymmetric (AP or AMP) Processing

Internal Migration of Activity (cont'd)

Specific to a 158, cycle stealing for I/O (total channel interference - "TCI") takes place only on the "BASE"

- "attached" side has more time available for instruction execution, since no TCI
- overall TCI is about the same for equivalent workloads and configurations for both symmetric and asymmetric processing
- example: TCI as % of aggregate busy time for selected laboratory batch benchmarks (158)

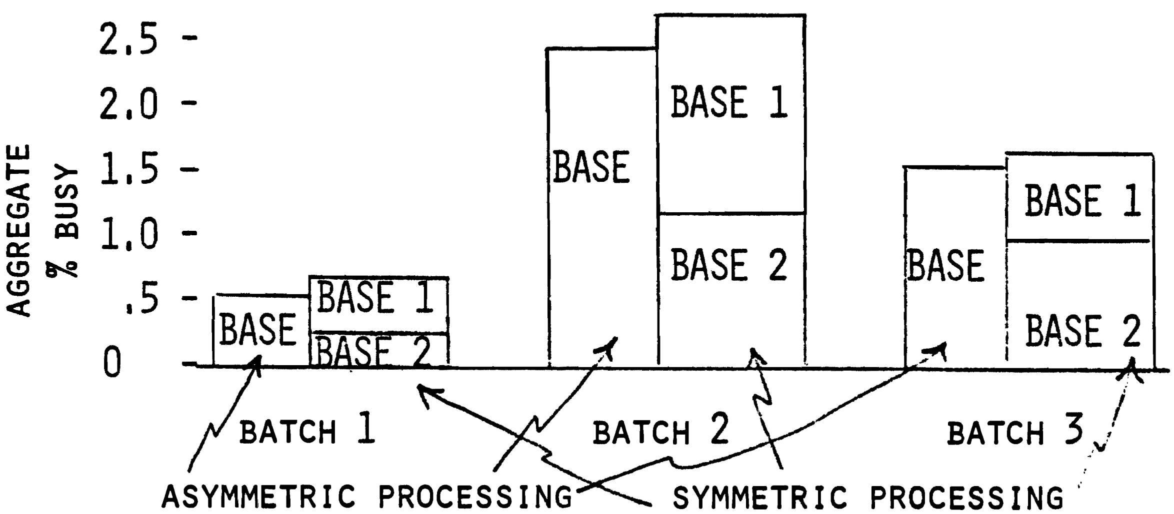

There is not additional impact of TCI in asymmetric processing as opposed to symmetric processing

PERFORMANCE CHARACTERISTICS OF ASYMMETRIC (AP OR AMP) PROCESSING

I/O CHARACTERISTICS

IN A HIGH I/O ENVIRONMENT WITH PROCESSORS RUNNING AT HIGH UTILIZATION, ASYMMETRIC PROCESSING CAN APPROXIMATE SYMMETRIC PROCESSING

- NEITHER
  - SIGP'S FOR I/O
  - INTERNAL QUEUEING FOR I/O

  CAUSE SIGNIFICANT DEGRADATION IN I/O PROCESSING
- EXAMPLE: I/O STRESS TEST ON 168
  - PROGRAM - 1. WRITE A FILE ON EACH OF 2 DASD DEVICES
    2. READ FROM EACH DASD DEVICE (EXCP)
    3. WAIT FOR READS TO COMPLETE
    4. GO TO 2
  - ENVIRONMENT - MULTIPLE COPIES WERE EXECUTED TO LOAD THE SYSTEM
  - RESULTS

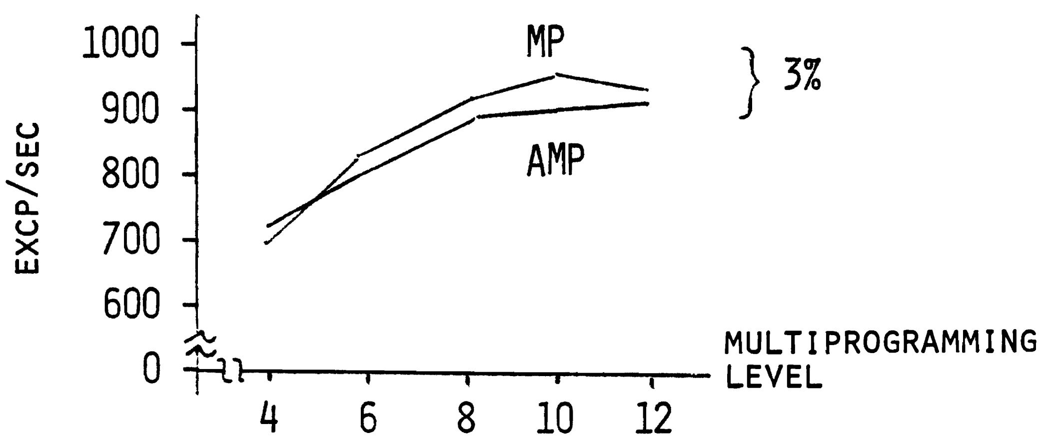

MVS CAN HANDLE ASYMMETRIC I/O PROCESSING ALMOST AS WELL AS SYMMETRIC I/O PROCESSING IN THIS ENVIRONMENT

PERFORMANCE CHARACTERISTICS OF ASYMMETRIC (AP OR AMP) PROCESSING

I/O CHARACTERISTICS (CONT'D)

SPECIFIC TO A 158, ASYMMETRIC PROCESSING CAN SUPPORT A AGGREGATE DATA RATE COMPARED TO UP, WITH MINIMAL OVERRUN

- EXAMPLE: I/O STRESS TEST ON 158 UNDER MVS
  - PROGRAM - AS BEFORE, EXCEPT WRITE TO DEVICES (4K BLOCKS)
    - 99% OF PROCESSING CAPABILITY DEVOTED TO I/O DRIVING FOR THE UP
    - IN ADDITION, A DISK TO TAPE DUMP
  - CONFIGURATION - 2 CHANNELS OF 3350
    - 2 CHANNELS OF 3330
    - 1 CHANNEL WITH 3420-6
    - AGGREGATE DATA RATE OF 4.8 MB/SEC
  - RESULTS

RESULTS	UP	AMP
AVG DATA RATE (MB/SEC)	.98	1.44
SIO RATE (SEC)	159	233
OVERRUN RATE (SEC)	.029	.284
OVERRUNS/SIO	.0002	.0012

  - IMPACT OF OVERRUNS - HERE, OVERRUNS ARE HANDLED AT THE CHANNEL AND THE CONTROL UNIT
    - EFFECT IS A MISSED REVOLUTION
    - PROCESSORS REMAIN BUSY DOING OTHER WORK

    AT 16.7 MS/REV, THE INCREASE IN I/O TIME IS
    - .05% FOR UP
    - .47% FOR AMP

HERE, FOR AMP
- OVERRUNS WERE MINIMAL
- SYSTEM DEGRADATION WAS NOT SIGNIFICANT

AND, SINCE THE UP WAS COMPUTE BOUND,
- AMP SUPPORTED A HIGH AGGREGATE DATA RATE, WITH HIGHER SIO RATE AND AVERAGE DATA RATE THAN UP

AP Performance

Internal Performance

An AP is 1.5 to 1.8 times the internal performance of the corresponding UP

- under MVS
- using identical configurations and programs

Measured in MIPS

- MIPS: millions of instructions per non-wait second
- for MP/AP/AMP, total MIPS is sum of MIPS for each processor

Examples: MIPS ratios (AP ÷ UP) for selected laboratory benchmarks on 158-3

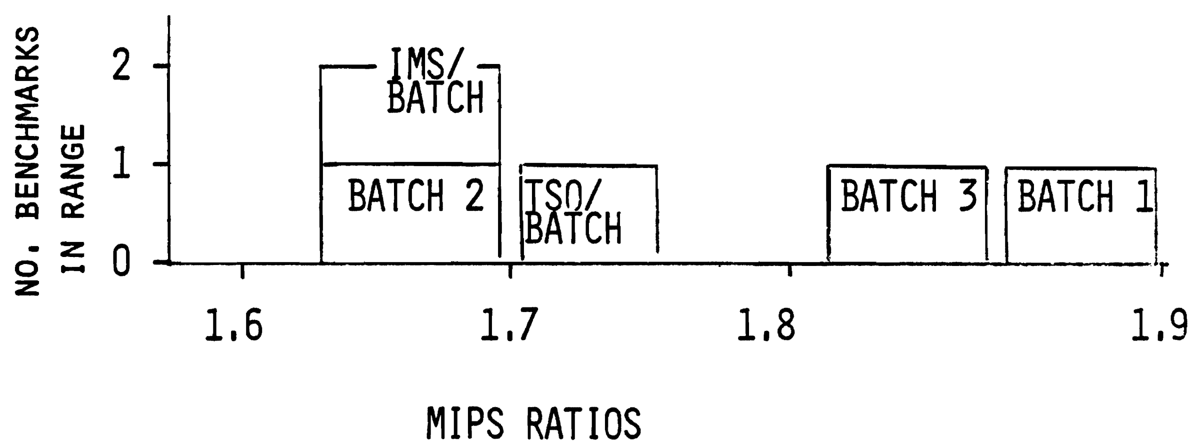

System Throughput

AP provides approximately the same throughput as the corresponding AMP

- under MVS
- using identical configurations and programs

Example: TSO/batch on 158-3, varying workload mix by using different numbers of terminals

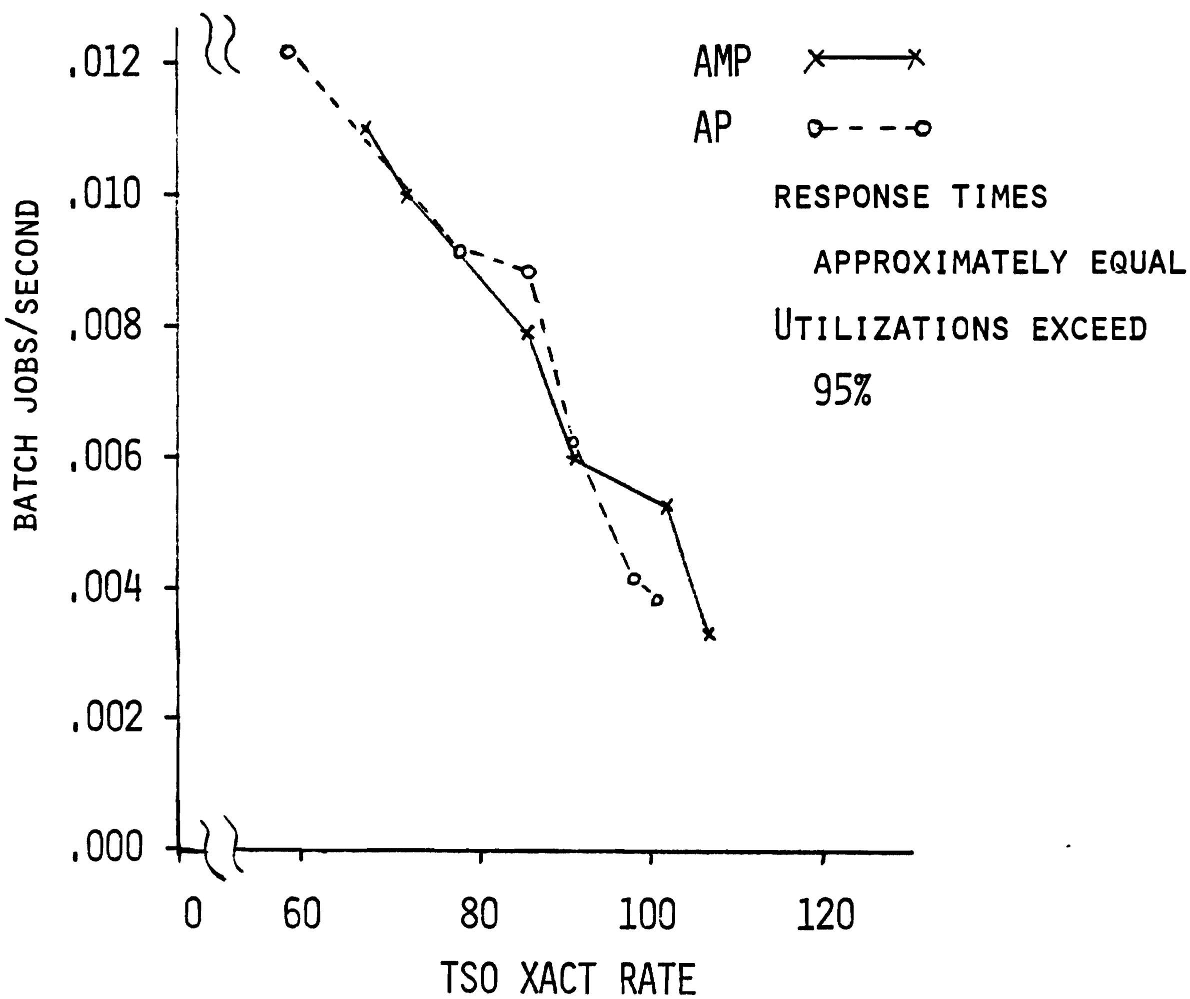

SYSTEM THROUGHPUT (CONT'D)

IF THE I/O AND STORAGE CONFIGURATIONS ARE ADEQUATE FOR THE AP, THEN AP THROUGHPUT MAY BE SIMILAR TO THAT OF A SYMMETRIC MP

- UNDER MVS
- USING IDENTICAL PROGRAMS AND, SUBJECT TO PHYSICAL LIMITATIONS, CONFIGURATIONS

EXAMPLES: TSO/BATCH AND IMS/BATCH ON 158-3 (WITH SAME NUMBERS OF CONTROL UNITS AND DEVICES)

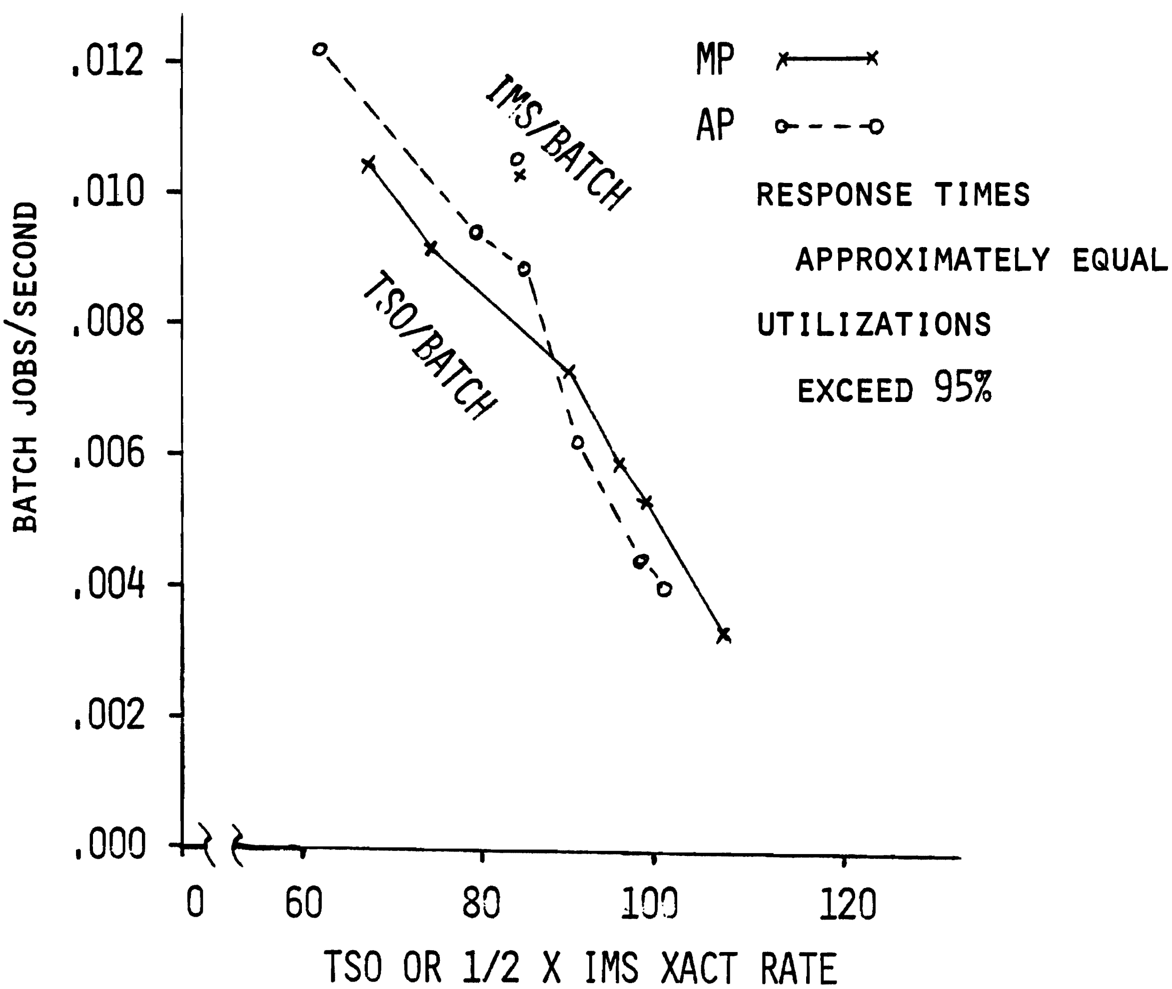

SYSTEM THROUGHPUT (CONT'D)

EXAMPLES OF AP THROUGHPUT COMPARED TO CROSS-CONFIGURED HD 158, FOR LABORATORY BENCHMARKS

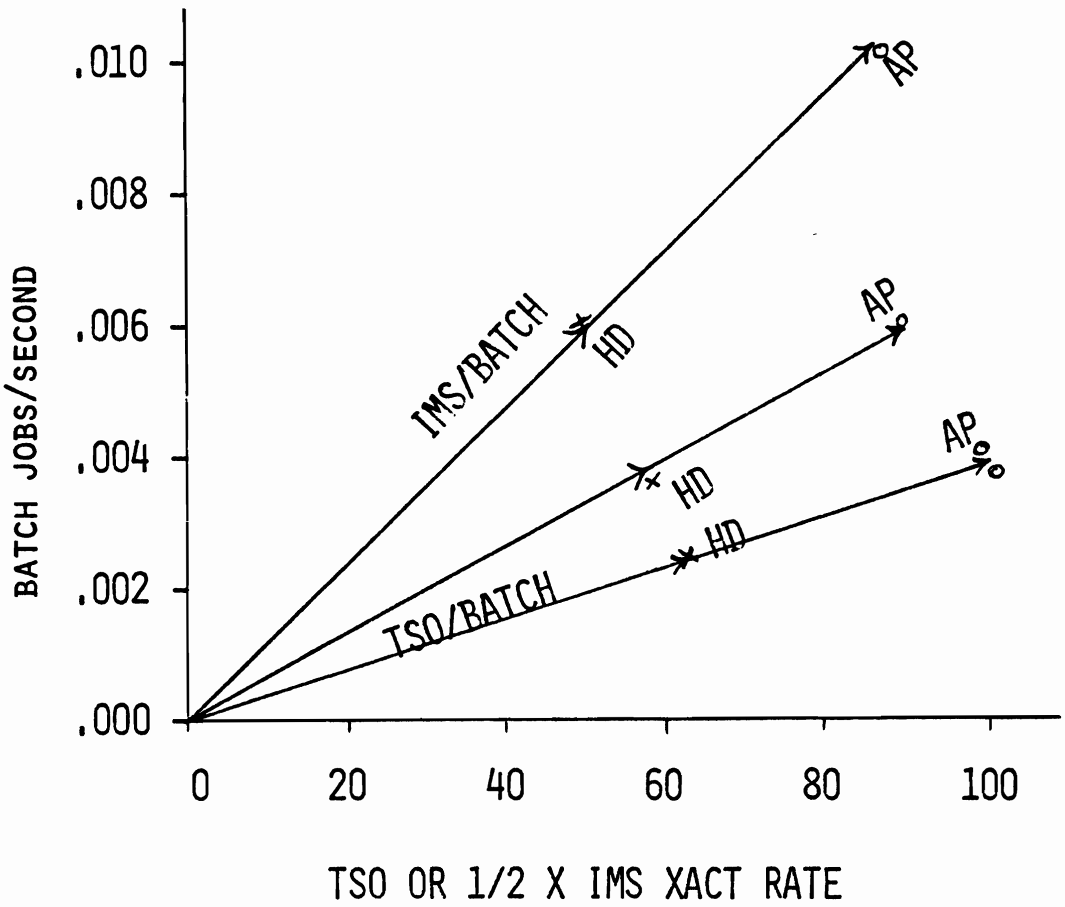

RESULTS FOR <u>THESE</u> <u>EXAMPLES</u>:

IMS/BATCH: AP IS ABOUT 1.7 TIMES HD THROUGHPUT

TSO/BATCH: AP IS ABOUT 1.6 TIMES HD THROUGHPUT

Considerations in Going to AP

To achieve potential throughput increase from increased processor power over UP, it may be necessary to scale upward other system components.

The identification of

- which components, if any
- degree of scaling, if any

depends upon, in relation to UP,

- where the current operating points are now
  - e.g., current workload mix, system resources
- in which directions growth is anticipated
  - will workload mix be the same
  - will more background batch be added
  - will more interactive work be added, etc.

Individual adjustments will vary depending on the above

With this in mind, some points to consider will be mentioned

Considerations in Going to AP

Workload and Workload Processing

Is parallelism present, in terms of MVS dispatchable units?

In most cases, asymmetry should <u>not</u> be a user concern in connection with workload assignment internal to the processor complex

- MVS works at its own level of granularity
- Let MVS handle the dispatching of dispatchable units
- MVS will naturally assign the work where MVS thinks it can best be done
- Asymmetry is treated naturally in MVS at MVS's own level

CONSIDERATIONS IN GOING TO AP

ASSESSMENT OF I/O ADEQUACY

IF SAME WORKLOAD MIX IS TO BE MAINTAINED, CHANNEL UTILIZATIONS WILL BE INCREASED

- EXPERIENCE WITH LABORATORY BENCHMARKS INDICATES THAT TOTAL CHANNEL UTILIZATION (SUMMED OVER ALL CHANNELS) INCREASES IN ABOUT THE SAME PROPORTION AS THROUGHPUT
- IF MIX IS MORE COMPUTE BOUND, THIS IS LESS IMPORTANT
- IF MIX WILL HAVE MORE I/O ACTIVITY, THIS WILL BE MORE IMPORTANT

TO KEEP INDIVIDUAL CHANNEL UTILIZATION AT AN INSTALLATION-DEPENDENT ACCEPTABLE LEVEL, IT MAY BE DESIRABLE TO HAVE

- MORE CHANNELS
- IMPROVED DEVICES (TO PUT LESS BURDEN ON CHANNELS)
- LATEST SOFTWARE LEVELS (TO PUT LESS BURDEN ON CHANNELS)

OUTBOARD I/O PROBABLY NEEDS TO BE INCREASED, APPROPRIATE TO INCREASED I/O ACTIVITY

- CONTROL UNITS
- DEVICES
- PACK, DATA SET PLACEMENT
- RESTRUCTURING, E.G., CATALOG SPLITTING

IF PROJECTED CHANNEL UTILIZATIONS ON THE AVAILABLE CHANNELS EXCEED DESIRED OPERATING POINTS, MP MAY BE MORE APPROPRIATE

Considerations in Going to AP

Assessment of Main Storage Adequacy

If same workload mix is to be maintained, main storage requirements may increase

Since single copy of SCP is used, just consider added needs

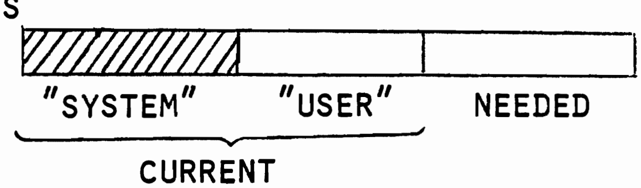

NEEDED = FACTOR X "USER"

Example: TSO/batch on 158 AMP, 4 vs 6 megs

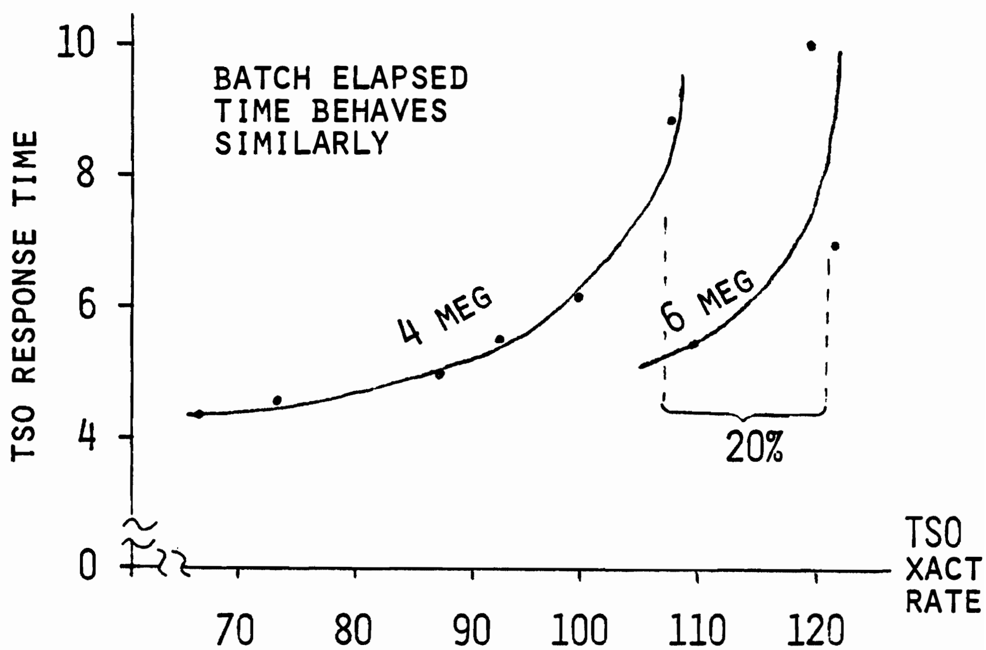

If storage requirements exceed availability, MP may be more appropriate

Recent Announcements Supporting Tightly Coupled Processing

EC 717728 for 168 AP, MP

- provides better implementation compatibility between AP and MP

Improvements in MVS

- in TSO/batch 168 AP test comparing SU 0, 6 with SU 0, 4, 5, 6, 7 indicate
  - improved throughput
  - decreased channel utilization
- improvements in channel utilization provide specific benefit for AP

MVS/SE PP, in combination with the System/370 Extended Feature

- reduced interprocessor interference
- reduced lock contention
  - dispatcher lock
  - UCB lock
  - SALLOC lock

SUMMARY

MP PERFORMANCE IS A GOOD INDICATOR OF APS PERFORMANCE

MVS TIGHTLY COUPLED PROCESSING DESIGN NATURALLY SUPPORTS APS

APS INTERNAL PERFORMANCE IS 1.5 TO 1.8 TIMES THAT OF THE CORRESPONDING UP UNDER MVS

SYSTEM THROUGHPUT PERFORMANCE OF APS REQUIRES

- PARALLELISM IN WORKLOAD
- ADEQUATE I/O
- ADEQUATE MAIN STORAGE

ASYMMETRIC I/O IS WELL SUPPORTED

- NO INTERNAL BOTTLENECKS
- HIGH SIO ACTIVITY CAN BE SUPPORTED
- HIGH DATA RATE CAN BE SUPPORTED

RECENT ANNOUNCEMENTS IMPROVE TIGHTLY COUPLED PROCESSING PERFORMANCE POTENTIAL

Reader's Comments

Title: Attached Processing (AP) System
Performance Characteristics and Considerations
Washington Systems Center
Technical Bulletin GG22-9004-0

Please state your occupation:___________________________

Comments:

Please mail to:

Mr. T. E. Giblin
IBM Corporation
Building 2
18100 Frederick Pike
Gaithersburg, Maryland 20760