U.S. DEPARTMENT OF COMMERCE
NATIONAL OCEANIC AND ATMOSPHERIC ADMINISTRATION
NATIONAL WEATHER SERVICE
OFFICE OF SYSTEMS DEVELOPMENT
TECHNIQUES DEVELOPMENT LABORATORY

TDL OFFICE NOTE 89-7

SIMULATION OF AWIPS-90 TECHNIQUE PROCESSING LOAD FOR A WFO

Harry R. Glahn and Wendy L. Wolf

December 1989

SIMULATION OF AWIPS-90 TECHNIQUE PROCESSING LOAD FOR A WFO

Harry R. Glahn and Wendy L. Wolf

1. INTRODUCTION

Meteorological and hydrological techniques for Weather Forecast Offices (WFO's) are defined as those processes which, when programmed and run on a computer, perform functions that assist station personnel in their daily activities and which are not what are usually considered to be "system functions." For instance, the plotting of a chart with forecaster-designated data and the running of numerical or statistical models are techniques, while the processing of data directly related to receipt, distribution, or databasing are system functions. In preparing for AWIPS-90, the National Weather Service (NWS) has defined 14 classes of techniques (e.g., decoding and parameter derivation are two of the classes) and several techniques in each class for a total of about 140 techniques.

In order to make an assessment of the computer capacity needed at a WFO, the computer requirements for each technique were estimated. Whenever possible, the estimate was based on experience with a similar technique. Eventually, the estimates were put in terms of millions of instructions per second (MIPS) over some period of time. There are many opportunities for error in such estimates, including the degree to which the "experience" program fully meets the AWIPS requirement; the difficulty of knowing a MIPS rating for the computer the experience program was running on; the degree to which the computer was being used for the technique as opposed to system management, etc.; and the unknown efficiency of the program.

After these estimates were made, another set of questions was then addressed: What would be the frequency of running these techniques, how much would their use overlap, etc.? To provide a first approximation to the answer, a "worst case hour" (WCH) was fashioned. The first version of this appeared in a Techniques Development Laboratory (TDL) document "AWIPS-90 Requirements for Meteorological and Hydrological Techniques" dated June 5, 1984. A revision of this document is dated January 17, 1985 and another (with a slightly modified title) dated November 1986. Information from this last version was provided to AWIPS bidders for the Definition Phase contracts (see the AWIPS Systems Requirements Specification (SRS) dated November 11, 1986, Table G-1). Significant aspects of most classes of techniques were "scheduled" in 5-minute segments of an hour. Expected time of receipt of incoming data and dependencies of techniques were used in the scheduling whenever appropriate. Some attempt was made to "spread" the processing over the hour. For instance, objective analysis of surface data did not overlap objective analysis of upper air data. Insufficient information was available to provide data on the Decision Assistance Class (10.0 in Table G-1) and Pilot Weather Briefing (Technique 4.1.2.2.1 in Table G-1). In estimating the WCH, it was assumed that the WFO was in the Warning Mode of operation (hazardous weather existing or imminent--see the SRS, Appendix A for details).

Fig. 1 contains the estimated required applications processing capacity in MIPS in 5-minute segments during the WCH based on information in the SRS, Table G-1. The highest value is 8.5 MIPS, the lowest 3.6, and the average for the hour about 6.2. The implication is that the provision of 8.5 MIPS capacity maintained over a 5-min period would suffice for classes of techniques for which data were provided in Table G-1. Computer utilization cannot be maintained at 100% for a mix of user programs over periods of several minutes, and a system is heavily used at, say, 80% utilization. Therefore, one might imagine a computer processing capacity rated at 10.6 MIPS running at 80% utilization to meet the 8.5 MIPS requirement. Note that this does not include system functions such as database management, data ingest processing, communications processing, system overhead, system control monitoring, and loop loading.

As more analysis was done and more information generated, it became apparent that adequate thought had not been given to specifying response time requirements for forecaster-initiated requests. That is, a high level of computer processing might be required in bursts of a few seconds to provide information from a technique that the forecaster is waiting for. Therefore, a plan was devised in December 1988 for redefining the worst case hour. (See "Plan for Redefining Required Response Times for AWIPS-90 Applications at WFO's" for details available from TDL.)

This document describes the activities associated with carrying out that plan, and provides information that may be useful in assessing the impact of redefining the WCH and overlaying the workstation-initiated requests.

2. REVISION OF WCH

For each subclass of technique, it was estimated [primarily by the Office of Meteorology (OM) with assistance from TDL] what percent of a technique's use would be scheduled and what percent would be by requests from workstations. For any technique that would be initiated from a workstation, a realistic response time was estimated. This was usually on the order of one to a few seconds.

Then, together with a few improved estimates of technique computer processing requirements, Table G-1 was revised to reflect only the scheduled techniques. The result is shown in Fig. 1. [Note: all data presented herein were current as of the date this paper was written--September 1989. See the Postscript for a discussion of changes made prior to publication.] The highest 5-minute segment is 6.7 MIPS, the lowest 2.7, and the average about 4.8. Differences between the WCH as presently formulated and the one in the SRS are due primarily to five factors: A few requirements were dropped (e.g., decoding of TAF's), a few requirements were added or increased (e.g., contouring of NMC gridpoint data), some techniques were reclassified as system functions (e.g., transformations of gridpoint data from one map projection to another), some techniques were taken wholly or partly out of the scheduled WCH because all or part of their use would be forecaster-initiated (e.g., mathematical operations on an image), and some techniques were rescheduled either because of better computer estimates (e.g., the advective and extrapolative models) or because of response times provided by OM (e.g., some objective analyses that had been spread over 10 minutes were put into a 5-min period).

Table 1 provides the approximate change in scheduled computer processing requirements due to the five factors mentioned above. It can be seen that very little change resulted from reclassifying techniques as system functions (about 0.2 MIPS). The combination of dropping or adding requirements results in a net positive requirement of 0.9 MIPS averaged over the hour. Putting techniques into the non-scheduled category resulted in a significant decrease in the scheduled requirements--1.9 MIPS averaged over the hour. Rescheduling techniques into different 5-min periods or into a shorter period of time had a maximum change in any 5-min period of 1.2 MIPS; the net result of rescheduling was -0.2 MIPS averaged over the hour. Most changes in this category were due to changes in required response times (e.g., from 10 minutes to 5 minutes), which does not result in a net increase or decrease. The overall decrease in this category is almost entirely accounted for by having a better estimate for the extrapolative and advective models (requirement 7.1.1 in Table G-1 of the SRS).

In order to address forecaster-initiated requests, Jim Kaplan of OM devised three scenarios for such requests (see the Appendix). Two of them involve five workstations--thought to be the maximum number that would be in use at the same time at AWIPS Initial Operating Capability (IOC)--and contain the requests that would be made from each workstation scheduled to the second within a 5-minute period. The third allows for six busy workstations, a not impossible situation, and represents post-IOC conditions with quality controlling of the four-dimensional database. Possible dependencies of requests at different workstations were not considered. Each of these workstation requests is for a particular technique, which has associated with it a response time and a MIPS requirement. That is, if the response time is 2 seconds and the MIPS requirement is 3, this implies that 6 million instructions would be required to accommodate the request. It also means that if 3 MIPS were available for 2 seconds, the response time could be met.

The question remained--how would the workstation requirements be overlaid with the scheduled processing to obtain the total computer capacity? Many possibilities undoubtedly exist; we tried three of them, as described below.

3. SIMULATION MODELS FOR MERGING OF WORKSTATION-INITIATED PROCESSING WITH SCHEDULED PROCESSING

A. Model 1

This is probably the simplest simulation procedure--for a particular scenario, just add the workstations' loads to the scheduled load whereever they overlap. The maximum value would indicate that a processing capacity equal to that value would meet all requirements of the scheduled processing and the workloads of the workstations detailed in that scenario. However, this seems unrealistic from the point of view that it assumes the exact dependency between requests at different workstations as specified in the scenarios--scenarios fashioned without regard to possible dependencies. Therefore, it was decided to randomize each workstation scenario with regard to the others, leaving intact each individual workstation schedule. That is, a different pseudorandom number was used to determine the scenario start time for each workstation somewhere within the 5-minute period. Since each workstation's scenario covered the total period, starting it midway extends it past the end. To alleviate this problem, the overlap was wrapped around to the start. That is, a workstation's scenario might start at second 200, run through second 300, then from 1 through 199. (This is very much like duplicating the scenarios over a number of 5-minute periods, but the end would always present a closure problem.) This has the effect of "rescheduling" a forecaster's duties at the beginning of the 5-minute period to somewhere nearer the end of the period, so that the sequence of duties may not be reasonable. For instance, he/she may be preparing a user product before analysis of the data. However, for our purposes, this "rescheduling" has little or no overall effect.

This model is somewhat unrealistic because it does not allow effective use of computer resources through a priority system which would reschedule workload as needed to meet specific workstation-initiated requests. The models described below contain priority scheduling.

Fig. 2 shows the results of one randomization of the three scenarios with unlimited computer capacity. This figure matches the randomized start times shown in Tables 2 through 4.

B. Model 2

A more realistic model is to permit the scheduled processing to be delayed a few seconds at those times when workstation-initiated requests are being met, so long as the scheduled processing can be completed within the 5 minutes. The rules for this model are:

1. A portion of the scheduled processing is not delayable. This is to account for such things as decoding of observations and the processing of NEXRAD data which appear in the scheduled load but have a response time on the order of seconds.

2. The scheduled processing is delayed by the combined workstation requirements existing at that particular second. No attempt is made to run a technique faster than necessary to meet the response time. That is, a 2 MIPS need for 2 seconds is filled at that rate, provided the capacity is there. If the 2 MIPS can only be partially met, the technique will "get behind." In succeeding seconds, it is allowed to "catch up," if possible, but not to "get ahead."

3. Each workstation requirement gets the same portion of the processing capability available, up to the amount needed. (In the simulation model, the smallest requests are filled first to guarantee all capability is used for workstation requirements, if needed.)

4. A technique scheduled at the very end of the 5-min period might not have a chance to finish. That is, one requiring 5 seconds scheduled at second 299 could not complete. To give all a chance to complete, the randomization was modified to the extent each technique was provided time equal to twice its response time to complete. As an example, a technique requiring 5 seconds was not scheduled after second 291. If the randomization put it after 291, it was "backed up" to 291.

5. Even with the procedure described in No. 4 above, a technique might not complete within the 300-sec period. For purposes of calculation, it was assumed to complete in one more second. However, this assumption did not affect any of the statistics computed for this model.

6. Since workstation requests can backlog scheduled processing, the 5-min period might end with a backlog. This was not permitted. Workstation requirements were not filled beyond the time that the backlogged scheduled processing could be wiped out. This, for low values of maximum available processing capacity, caused several workstation-initiated techniques to not complete by the end of the period.

C. Model 3

Model 3 is identical to Model 2 except that the full computer processing capacity (except for that portion meeting the non-delayable processing requirement) is used to meet the workstation requests, as needed, thus allowing the techniques to "get ahead." For instance, a technique requiring 2 MIPS each second over a 2-sec period would use 2 MIPS for each of 2 seconds in Model 2, but in Model 3 would use the full computer processing capacity to complete as soon as possible. If the full capacity were 11 MIPS and 1 MIPS were reserved for uninterruptible scheduled work, then the technique could complete in 4/10 = 0.4 seconds rather than 2 seconds if not competing with other workstation requests. Rule No. 5 (described under Model 2) is necessary for one of the computations made for this model.

4. SIMULATION RESULTS

Each scenario was run on each of the three models with a total of 100 random relationships among the five workstations. Tables 5 through 7 show results of these runs. Also, one realization of the random process is shown in Figs. 3 through 5 for each of the models and scenarios; the input schedules of forecaster-initiated techniques are described in Tables 2 through 4, respectively, which embody Rule No. 4 listed under Model 2.

Fig. 2 shows the results of one randomization for each of the three scenarios for Model 1 for unlimited computer power. This figure is mainly to indicate that if priority scheduling were not employed (scheduled workload interrupted as necessary) and all response times were to be met, the computer power necessary to meet <u>all</u> workstation requests would be substantial for Scenarios 2 and 3 (17 and 34 MIPS, respectively). Note that this is for only one randomization, and other randomizations would produce both higher and lower maximum peaks. The major peaks for Scenario 3 are due to requirement No. 2.2 "Check Consistency of Entered Information" when quality controlling the four-dimensional database (see Table 4).

Fig. 3 shows the computer processing use profiles for Scenario 1 that matches the randomization in Table 2 for all three models for an 11 MIPS maximum computer processing capacity (hereafter referred to as MCPC). Note that the computer processing use profiles for Models 1 and 2 are the same. Model 3 makes use of the maximum computer processing power a greater portion of the time than do Models 1 and 2, and the processing settles back to scheduled workload more quickly.

Although Models 1 and 2 produce similar results in terms of use profile (Fig. 3a), they are quite different in meeting response times, as shown in Tables 5 and 6. For a scheduled workload of 6.7 MIPS (actually, 6.735, the best estimate we have at the moment), 1.5 MIPS of which is non-interruptible, an 8 MIPS MCPC meets the response times 93.0% of the time for Model 2 but only 15.6% for Model 1. The column headed "Increase in Response Times" should be interpreted in the following manner. For those techniques that did not finish by the required response time, the time required to complete was <u>increased</u> by the percentage shown. For example, for Scenario 1 and Model 2, for those 1.8% (0.3%) not meeting the response time for a 9 (10) MIPS MCPC, the <u>average</u> completion time was 198% (149%) of that "required." Note that this value is <u>slightly</u> dependent on Rule No. 5 in the previous section.

Table 7 shows that Model 3 is not much better than Model 2 for meeting response times (differences as small as 1 percent shouldn't be regarded as significant; simulations with different sets of random numbers show changes in response times met of 0.5%), but the "Decrease in Response Time" column shows that for a 9 MIPS MCPC, of those 98.7% of the techniques meeting the response time, the actual response time was on the average 71.3% of the "required" response time. So, for Model 3, the average response was 71.3% of the required time for 98.7% of the requests and 212% of the required response time for the remaining 1.3% of the requests. (Remember that these are averages and individual responses will vary.)

Fig. 4 shows the computer processing use profile for Scenario 2 that matches the randomization in Table 3 for all three models for an 11 MIPS MCPC. Notice the period of time nearly 1 minute long where the maximum processing capacity is used. Remember, this is for only one randomization of the workstations for this scenario. Tables 6 and 7 show that a 10 MIPS MCPC is necessary to get the percentage met into the 90's and an 11 MIPS MCPC is needed for the percentage met to exceed 95. For an 11 MIPS machine, simulation Model 3 shows that the response times were met 97.4% of the time; of those met, the actual response time averaged only 64.3% of that required; and of those not met the increase averaged only 14%.

Fig. 5 shows the computer processing use profile for Scenario 3 that matches the randomization in Table 4 for all three models for an 11 MIPS MCPC. It is immediately apparent that the computer processing is saturated much of the time (from Tables 6 and 7, 50.2% for Model 2 and 59.8% for Model 3). As with the high peaks in Fig. 2, this high use is connected with requirement No. 2.2 "Check Consistency of Entered Information," which relates to quality controlling grid fields. Scenario 3 also includes a request for a constant 0.1 MIPS for a hydrologic forecast model throughout the 5-min period and according to the scheduling algorithm of Model 2 was honored, so the quality control requirement could not be met with less than a 1.5 + 0.1 + 21.0 = 22.6 MIPS MCPC; Tables 5 and 6 verify this calculation. (The 10 quality control requests comprise 10/37 = 27.0% of the total requests, so only 73.0% could be met until a MCPC in excess of 22.6 MIPS was available.) For Model 3, the provision for techniques to "get ahead" allows many to finish early (e.g., for the randomization shown in Fig. 5, the hydrologic model finished in 3.3 seconds rather than 300 by delaying scheduled processing, which "caught up" about 6 seconds after that).

5. SUMMARY AND CONCLUSIONS

As indicated in the scenarios given in the Appendix, much of a forecaster's time is spent in looking at data, including loops of satellite and radar data. In view of that expectation, Scenarios 1 and especially 2 may be quite realistic for a "heavy" forecaster-initiated load for IOC. Accordingly, a 10 or 11 MIPS MCPC seems adequate--for Models 2 and 3, requests are met about 95% of the time (nearly all being met in double the time); of those not met, the increase is (on the average) less than double that desired; the maximum computer processing capacity is used only 15% to 33% of the time; and it is utilized less than 80%. The use of a "get ahead" algorithm like that employed in Model 3 allows the actual response time to be less than the "required" time--an average of less than 65% for Model 3 for a 10 or 11 MIPS MCPC. (This "improvement" of Model 3 over 2 is probably an overestimate. Some techniques' speeds will undoubtedly be paced by input and/or output, so that the maximum available capacity, in terms of cpu cycles, cannot be used. Input/output requirements were taken into account to some degree when estimating realistic response times. This was done partly through using processing experiences which, in turn, many times included input/output. However, the differences between Models 2 and 3 are minor except for the "finished early" aspect.)

Scenario 3 is more problematical for at least three reasons. First, there is practically no experience with quality controlling a four-dimensional database in the way it is envisioned at full capability AWIPS, so the processing estimate of 42 million instructions (21 MIPS for 2 seconds according to Table 4) is very soft. In addition, the requirement for extensive checking to be done in 2 seconds 10 times in a 5-min period may be more demanding than necessary. (If the 21 MIPS for 2 seconds is changed to 7 MIPS for 6 seconds, then 95.8% (98.4%) of the requests can be met with 14 MIPS MCPC for Model 2 (3), a considerable difference from the indications in Tables 6 and 7.) Finally, almost no effort has gone into specifying this extensive quality control capability, so the likelihood of it being in place even as a first major upgrade is doubtful. For these reasons, we believe the uncertainties associated with this scenario are too great for it to heavily influence computer sizing for IOC. There are many other uncertainties associated with techniques. For example, processing of NEXRAD data, either alone (because NEXRAD doesn't furnish data in all the forms needed by the forecaster) or in combination with other data, may not be adequately addressed because of lack of knowledge of just what is needed. These and other issues not even thought of yet, along with quality control of a four-dimensional database, will come into better focus 2 or 3 years from now and can be dealt with then.

The original WCH showed a maximum 8.5 MIPS 5-min period. It was never expected that <u>exactly</u> 8.5 MIPS would suffice, but that some "safety margin" would be necessary to provide the capacity for 8.5 MIPS to actually be used, on average, over a 5-min period. For instance, one might suppose an 80% utilization over periods as long as 5 minutes (this may even be high); a 10.6 MIPS MCPC could provide that capability with the extra capacity to assure the schedule could be met. Therefore, our conclusion reached above that a 10 or 11 MIPS MCPC would meet the new WCH (including workstation-initiated requests) indicates that the total technique requirements are still essentially the same as when the SRS was issued in November 1986.

One might wonder why the "tightening" of response times didn't appreciably affect the overall requirement, especially since more requirements were added than dropped. There are two main reasons: 1) the total processing for the workstation-initiated requests had been originally overestimated, if Scenarios 1 and 2 embody typical loads (the required MIPS averaged over the 5 minutes for Scenarios 1 and 2 are 0.7 and 1.2, respectively, whereas the non-scheduled processing dropped from the WCH is 1.9 MIPS), and 2) a priority scheduling algorithm allows the scheduled processing to be delayed a few seconds while the computer processing power of the system is brought to bear on the workstation requests. This <u>does</u> assume a priority system, and while Models 2 and 3 used only a 2-level system, a system of several levels is very desirable, if not necessary.

A possible flaw (only one!?) in our argument is in stating that the original 8.5 MIPS represented a, say, 80% utilized MCPC, yet the peaks of 11 MIPS for the 11 MIPS MCPC is a 100% utilization. Our defense is that the 11 MIPS peaks are usually for periods of only a few seconds rather than 5 minutes and it's more plausible that the MCPC could be nearly fully utilized for a few seconds than 300 seconds. For Model 2 and Scenario 1 (2), the peaks of 11 MIPS represent only 8.3% (15.3%) of the 5 minutes.

We also point out that while the SRS only omitted the technique class No. 10 "Decision Assistance," we can now say with assurance that nothing in this category will be specified for IOC, and that nothing significant will be specified for the first major upgrade. Therefore, our new estimates represent "all" requirements, while those in the SRS were partial because they omitted class No. 10. Also, reservation of 1.5 MIPS for uninterruptible workload is probably conservative, and 0.5 MIPS may be more realistic for a multi-level priority system. Reserving only 0.5 MIPS allows the responses met for a 10 MIPS MCPC for Model 2 and Scenario 2 to go from 93.7% to 96.4%.

It appears that we could require workstation-initiated response times be met 95% of the time and twice those response times 99% of the time without driving the overall computer requirements up from the original estimates by more than, say, 10 percent. The many other uncertainties in the total requirement/estimation process puts such a figure well into the noise level. A pessimistic estimate might be (10.2-8.5)/8.5 = 20% increase, the 8.5 MIPS being the old estimate and the 10.2 MIPS being the point at which 95% of the responses are met (also, 99% of the responses met in double the time) for Scenario 2 and Model 2. (Note that Scenario 2 is more demanding that Scenario 1, and Model 2 is more conservative that Model 3.)

Our analysis only addressed the aggregate of techniques being run; we did not compute statistics on <u>individual</u> techniques. It could be that the response time of one or more of the techniques in a scenario was seldom met even though the overall percentage was high. This situation would not strictly meet the intent, which is that the percent met is, say, 95% for <u>each</u> technique. We believe that a multilevel priority system would solve any problems that might arise with individual techniques tending to be "slow."

Our main conclusions are, then, 1) that the new technique processing requirements embodied in the new WCH (both scheduled and Scenarios 1 and 2) are quite similar to those in the original WCH, 2) that we can require workstation initiated response times be met 95% or the time and twice those times 99% of the time without significantly increasing the requirements from those in the SRS, 3) that the quality control of the four-dimensional database be left for a future upgrade when more is known about total processing requirements, including, but not limited to, the quality control, and 4) that a scheduling algorithm be provided for techniques in AWIPS to allow jobs to use resources according to their priorities.

ACKNOWLEDGMENTS

We thank Jim Kaplan for his provision of the scenarios and for his helpful comments.

POSTSCRIPT

This document was written in September 1989 to assist in the specification of requirements for hydrometeorological techniques to be included in the Request for Proposal (RFP) for the outphases (Development, Deployment, Operations) of AWIPS. Publishing of the paper was delayed until the draft RFP was released to the AWIPS contractors for comment (December 22, 1989). Between September and December, refinements were made to the estimates provided in this paper. A document is in preparation which will address the exact information provided in the draft RFP. A summary of the major changes and an estimate of their impact on the presented analysis is provided below.

Scenario No. 3 has been changed in the following manner. The Forecaster 2 sub-scenario was replaced with a sub-scenario of actions representative of a forecaster interacting with a full 4-dimensional gridded data base. The Forecaster 5 sub-scenario was replaced with the sub-scenario of Forecaster 4 in Scenario No. 2. In addition, the estimated processing needed to quality control the gridded data set (technique 2.2) was reduced from 21 MIPS to 10 MIPS. These changes decrease the high MIPS spikes in this scenario while increasing the number of lower MIPS spikes. Scenario No. 3 was included in the draft RFP as the second (post-IOC) scenario.

Due to requirements to receive NMC model guidance at a finer time and vertical resolution than previously required, the scheduled processing load increased by 1.2 MIPS during the worst case 5-minutes in order to produce contours of the gridpoint data (technique 4.1.3.5). This 1.2 MIPS would be added to the 6.7 MIPS scheduled workload value used within this paper.

Finally, the background scheduled processing load as specified in the draft RFP includes an additional 2 MIPS to account for applications processing in support of backup of a failed neighboring WFO. Again, this would be added to the 6.7 MIPS value used within this paper.

These additional 3.2 MIPS (1.2 for contouring and 2.0 for backup) represent an <u>interruptible</u> load. The additional processing power necessary to support the worst case 5-minute scenarios would likely not be increased by 3.2 MIPS, but rather by some smaller value, say 2.0 MIPS. A thorough assessment of the impact of these changes is in preparation and will be published when completed.

APPENDIX

This appendix contains three scenarios for activity at WFO workstations. Each one covers 5 minutes. The first two are for IOC; the third is for post IOC.

1. SCENARIO NO. 1

Location: An East Coast WFO, in warning mode.

Situation: There has been a lot of recent rainfall over part of the area. Flash flood guidance over much of the area shows only an inch or so needed to produce flash flooding.

River Flooding in Progress
Flash Flood Watch in Effect
Severe Thunderstorm Watch in Effect
Scheduled (Public/Aviation) product issuance due shortly
Other scheduled products need updating

Time: 8:45 PM

Workstation Configuration: 3 Type III, 2 Type II, 1 Type 1

Note - Animation will occupy a significant portion of any individual's time and includes the time to load the loops. Processing is not included for this function in the worst case 5-min processing estimates.

Forecaster 1 - WARNING METEOROLOGIST - WS Type III

Time	Action
0:00	Run the Interactive Skew-T (3.2.7) twice to determine Equilibrium Level with surface conditions at station x as surface input.
0:30	Load and animate a rapid scan combined VIS/IR satellite loop of 8 frames on Local scale.
1:15	Load and animate an 8-frame loop of NEXRAD composite reflectivity and overlaid echo-top contours on Radar scale.
1:55	Load and animate NEXRAD combined Z/V image of 12 frames and toggle several times between Z and V portions on Radar scale.
2:45	Overlay several NEXRAD algorithm graphics and continue animation of Z/V image on Radar scale.
3:50	Run read-cursor (9.2.2) to determine cloud top height from displayed IR image correlated with nearest sounding.
4:20	Initiate Warning Generation (11.2.2.1) to generate Severe Thunderstorm warning. (Warning area is interactively specified. Then product is formatted (11.2.1.5.1) at 4:45, followed by an update of the status display of watches, warnings, and alerts (12.4.2,3)).

Forecaster 2 - AVIATION/MARINE FORECASTER - WS Type III

Time	Action
0:00	Coordination & shared task with Public Forecaster. No work at his/her own workstation.

2:00	Plot h+30 surface winds (4.1.2.1.1). Examines output: offshore thunderstorms are producing strong outflows.
3:00	Updates marine forecasts (CWF,OFF) (11.2.1.4.1). Quickly examines and approves the output.
3:40	Issues Special Marine Warning (SMW) (11.2.2.1). [Warning area is interactively specified. Then, product is formatted at 4:00 (11.2.1.4.1) followed by an update of the status display of watches, warnings, and alerts (12.4.2,3).] Examines and performs minor editing of the output.
4:45	In another window, load, animate, and examine a 16-frame IR loop with Profiler 300-mb winds overlaid.

Forecaster 3 - FLASH FLOOD/THUNDERSTORM WARNING COORDINATOR - WS Type II

Time	Action
0:00	Run plot of LFWS reports (4.1.2.1.3).
0:05	Load and examine the NEXRAD precipitation algorithm/imagery (in another window).
0:40	Run Extract Information at Cursor (9.2.2) in order to estimate the location of maximum rainfall and get names of nearby gauges and cities. (This is done in a third window.) A listing is sent to the printer.
0:50	Examine plotted LFWS reports.
1:20	Examine list printed above.
2:00	Run program to Put Cursor (9.2.2) at "location x" on precipitation algorithm/imagery since storm spotter has just reported 1" of rain in 5 minutes.
2:20	Phone call to additional spotter in vicinity of above report to update rainfall information; discuss briefly with lead forecaster. A second spotter reports 1.5"/10 min.
4:05	Initiate Warning Generation Program (11.2.2.1) to generate Flash Flood Warning. [Warning area is interactively specified. Then the product is formatted (11.2.1.3.1) at 4:30, followed by automatic updating of the status display of watches, warnings, and alerts (12.4.2,3).]

Forecaster 4 - HYDROLOGIST - WS Type II

Time	Action
0:00	Initiates a dambreak model (7.2.3) to produce "what if" scenario on an unsafe dam which could soon be getting heavy rain just upstream.
0:15	In a second window, initiate plot of DCP river stages (4.1.1.1).
0:20	In a third window, initiate Hydro Product Formatter (11.2.1.3.1) to generate Flood Statement, or Warning if needed, for Basin X, Basin Y, and Basin Z, based on latest updated RFC guidance.
0:45	Examines output from product formatter task and makes manual changes to text.
2:15	Spends time retrieving and examining six image or graphic products.

Forecaster 5 - PUBLIC FORECASTER - WS Type III

Time	Action
0:00	Plot Profiler/rawinsonde combined winds at 300-, 500-, and 700-mb levels.

0:15	Load loop of last 4 frames of each of the above into separate windows with IR imagery overlaid on 300-mb winds, and examine for trough locations.
2:30	Run extrapolation (5.2.3) to determine time of passage of trough line at 300 mb. Same for 700 mb. (Features are interactively specified on the two sets of images. Actual extrapolation function is performed at 3:00 for 300-mb and 3:30 for 700-mb features.)
4:00	Runs extract point value from graphic four times (6.4.1) to determine the current temperatures at zone forecast points not taking hourly observations. He/she suspects some temperatures may already be at or below previous forecast minima.
4:20	Run zone formatter (11.2.1.5.1) to update flash flood watch in two additional zones.
4:50	Manually inspect/edit the updated Zones.

****

2. SCENARIO NO. 2

Location: An East Coast WFO, in warning mode.

Situation: There has been a lot of recent rainfall over part of the area. Flash flood guidance over much of the area shows only an inch or so needed to produce flash flooding. Scattered thunderstorms and a front approaching from the west.

River Flooding in Progress
Flash Flood Watch in Effect
Severe Thunderstorm Watch in Effect
Scheduled (Public/Aviation) product issuance due shortly
Other scheduled products need updating

Time: 8:45 PM

Workstation Configuration: 3 Type III, 2 Type II, 1 Type 1

Note - Animation will occupy a significant portion of any individual's time and includes the time to load the loops. Processing is not included for this function in the worst case 5 minute processing estimates.

Forecaster 1 - WARNING METEOROLOGIST - WS Type III

Time	Action
0:00	Runs surface contour (4.1.3.1) of half-hourly mass convergence for persistent areas of convergence.
0:25	Run the Interactive Skew-T (3.2.7) twice to determine Equilibrium level from nearby sounding with surface conditions at station x as surface input.
0:55	Load and animate an 8-frame loop of NEXRAD composite reflectivity and overlaid echo-top contours on Radar scale.
1:40	Load and animate NEXRAD combined Z/V image of 12 frames and toggle several times between Z and V portions on Radar scale.
2:40	Overlay several NEXRAD algorithm graphics and continue animation of Z/V image on Radar scale.

3:45	Run Read Cursor (9.2.2) to determine cloud top height from displayed IR image correlated with nearest sounding. Above actions confirm need for warning, so . . .
4:05	Run warning generation (11.2.2.1) to generate Severe Thunderstorm Warning. [Warning area is interactively specified. Then product is formatted (11.2.1.5.1) at 4:30, followed by an update of the status display of watches, warnings, and alerts (12.4.2,3).]

Forecaster 2 - AVIATION/MARINE FORECASTER - WS Type III

Time	Action
0:00	Run Extrapolate Image Feature (5.2.3) to determine speed of movement of back edge of low-mid cloud deck. Interactively defines feature by marking it (frame 1 0:40-0:50 and frame 2 1:30-1:40). Actual extrapolation begins at 1:45. While doing this, he receives (at 1:45) an alarm for a 40-kt wind gust at a data buoy so he switches from his aviation "hat" to his marine "hat".
2:00	Plot h+30 surface winds (4.1.2.1.1). Examines output: offshore thunderstorms are producing strong outflows.
2:45	Issues Special Marine Warning (SMW) (11.2.2.1). [Warning area is interactively specified. Then, product is formatted at 3:05 (11.2.1.4.1) followed by an update of the status display of watches, warnings, and alerts (12.4.2,3).] Examines and performs minor editing of the output.
3:50	Updates marine forecasts (CWF,OFF) (11.2.1.4.1). Quickly examines and approves the output.
4:30	Load, animate, and examine a 16-frame IR loop with Profiler 300-mb winds overlaid.

Forecaster 3 - FLASH FLOOD/THUNDERSTORM WARNING COORDINATOR - WS Type II

Time	Action
0:00	Run plot of Precip. accumulation (4.1.1.1).
0:03	Run plot of LFWS reports (4.1.2.1.3) in another window.
0:08	Load and examine the NEXRAD precipitation algorithm/imagery (in another window).
0:15	Initiate plot of DCP river stages (4.1.1.1).
0:40	Run Extract Information at Cursor (9.2.2) in order to estimate the location of maximum rainfall and get names of nearby gauges and cities. (This is done in another window.) A listing is sent to the printer.
0:50	Examine plotted LFWS/precip. reports and river stages.
2:10	Examine list printed above.
2:30	Run program to put cursor (9.2.2) at "location x" on precipitation algorithm/imagery since storm spotter has just reported 1 inch of rain in 5 minutes.
2:45	Run Extract Information at Cursor (9.2.2) to get names of nearby spotters.
3:00	Calls another spotter in vicinity of apparent maximum rainfall area. Report confirms very heavy rains.
4:00	Initiate Warning generation program (11.2.2.1) to generate Flash Flood Warning. [Warning area is interactively specified. Then, the product is formatted (11.2.1.3.1) at 4:30, followed by automatic updating of the status display of watches, warnings, and alerts (12.4.2,3).]

Forecaster 4 - PUBLIC FORECASTER - WS Type III

Time	Action
0:00	Run Zone formatter (11.2.1.5.1) to produce update for Flash Flood Watch in two additional zones. Includes quality control. He/she is finishing up an earlier task.
0:45	Plot Profiler/rawinsonde combined winds (4.1.2.1.4) at 300-, 500-, and 700-mb levels.
1:30	Load loop of last 4 frames of each of the above into separate windows, with IR imagery overlaid on the 300-mb winds, and examine for trough locations.
2:30	Run Extrapolate Image Feature (5.2.3) to determine time of passage of trough line at 300 mb. Same for 700 mb. (Features are interactively specified on the two sets of images. Actual extrapolation function is performed at 2:50 for 300 mb and 3:10 for 700 mb.
3:20	Initiates surface contour (4.1.3.1) from 1/2 hourly data.
3:30	Runs Extract Point Value from Graphic (6.4.1) to determine current temperatures at zone forecast points not taking observations. He/she suspects that some temperatures may already be at or below the previous forecast minimum due to rain-cooled air.
3:35	Examines output from above.
4:30	Run Zone formatter (11.2.1.5.1) to update for lower temperatures. These lower temperatures should be reflected in the most recent LAMP guidance input to the product formatter.
4:45	Examines output of above to see if assumption about LAMP was correct.

Forecaster 5 - MET TECH - WS Type II

Time	Action
0:00	He/she spends the entire time editing a Weather Summary. In addition the following is done:
0:30	Runs a surface streamline analysis (4.1.3.3) to update the previous hourly version of this graphic and examines it.
1:00	Runs time series of station plot (4.1.5.1) for six ASOS stations to determine how the weather has changed throughout the day. (Assume they are invoked by one operator action and are run consecutively for 18 seconds beginning at 1:10.)
1:28	Examine output of above and incorporate into Weather Summary.
3:35	Run three Profiler time-height cross-section plots (4.1.5.4) and examine them for features he/she may wish to describe in the Weather Summary.

****

3. SCENARIO NO. 3

Location: An East Coast WFO, in warning mode.

Situation: There has been a lot of recent rainfall over part of the area. Flash flood guidance over much of the area shows only an inch or so needed to produce flash flooding.

River Flooding in Progress
Flash Flood Watch in Effect
Severe Thunderstorm Watch in Effect
Scheduled (Public/Aviation) product issuance due shortly

Time: 8:45 PM

Workstation Configuration: 3 Type III, 2 Type II, 1 Type 1

Note - Animation will occupy a significant portion of any individual's time and includes the time to load the loops. Processing is not included for this function in the worst case 5 minute processing estimates.

Forecaster 1 - WARNING METEOROLOGIST - WS Type III

Time	Action
0:00	Run the Interactive Skew-T (3.2.7) twice to determine Equilibrium Level with surface conditions at station x as surface input.
0:30	Load and animate a rapid scan combined VIS/IR satellite loop of 8 frames on Local scale.
1:15	Load and animate an 8-frame loop of NEXRAD composite reflectivity and overlaid echo-top contours on Radar scale.
1:55	Load and animate NEXRAD combined Z/V image of 12 frames and toggle several times between Z and V portions on Radar scale.
2:45	Overlay several NEXRAD algorithm graphics and continue animation of Z/V image on Radar scale.
3:50	Run read-cursor (9.2.2) to determine cloud top height from displayed IR image correlated with nearest sounding.
4:20	Initiate Warning Generation (11.2.2.1) to generate Severe Thunderstorm warning. [Warning area is interactively specified. Then product is formatted (11.2.1.5.1) at 4:45, followed by an update of the status display of watches, warnings, and alerts (12.4.2,3).]

Forecaster 2 - AVIATION/MARINE FORECASTER - WS Type III

Time	Action
0:00	He/she spends the entire 5-min period editing gridfields (6.4.2) and quality controlling (2.2) the edited gridfields. Assume 5 gridfields are done in 5 minutes. (Actual processing for both techniques would be back to back in the last 7 seconds of each minute.)

Forecaster 3 - FLASH FLOOD/THUNDERSTORM WARNING COORDINATOR - WS Type II

Time	Action
0:00	Run plot of LFWS reports (4.1.2.1.3).
0:05	Load and examine the NEXRAD precipitation algorithm/imagery (in another window).
0:40	Run Extract Information at Cursor (9.2.2) in order to estimate the location of maximum rainfall and get names of nearby gauges and cities. (This is done in a third window.) A listing is sent to the printer.
0:50	Examine plotted LFWS reports.
1:20	Examine list printed above.

2:00	Run program to Put Cursor (9.2.2) at "location x" on precipitation algorithm/imagery since storm spotter has just reported 1 inch of rain in 5 minutes.
2:20	Phone call to additional spotter in vicinity of above report to update rainfall information; discuss briefly with lead forecaster. A second spotter reports 1.5 inch in 10 minutes.
4:05	Initiate Warning Generation Program (11.2.2.1) to generate Flash Flood Warning. [Warning area is interactively specified. Then the product is formatted (11.2.1.3.1) at 4:30, followed by automatic updating of the status display of watches, warnings, and alerts (12.4.2,3).]

Forecaster 4 - HYDROLOGIST - WS Type II

Time	Action
0:00	Initiates a dambreak model (7.2.3) to produce "what if" scenario on an unsafe dam which could soon be getting heavy rain just upstream.
0:15	In a second window, initiate plot of DCP river stages (4.1.1.1).
0:20	In a third window, initiate Hydro Product Formatter (11.2.1.3.1) to generate Flood Statement, or Warning if needed, for Basin X, Basin Y, and Basin Z, based on latest updated RFC guidance.
0:45	Examines output from product formatter task and makes manual changes to text.
2:15	Spends time retrieving and examining six image or graphic products.

Forecaster 5 - PUBLIC FORECASTER - WS Type III

Time	Action
0:00	He/she spends the entire 5-min period editing gridfields (6.4.2) and quality controlling (2.2) the edited gridfields. Assume 5 gridfields are done in 5 minutes. (Actual processing for both techniques would be back to back in the last 7 seconds of each minute.)

Forecaster 6 - MET TECH - WS Type I

Time	Action
0:00	Invokes Radar Coded Message editor (9.1.2). (Pre-processing occurs first; then 4 minutes of editing; then post-processing occurs.)

Table 1. Changes in Table G-1 in MIPS for each 5-min period and the average over the hour due to various factors.

Factor	Ending of 5-min Period												Avg.
	5	10	15	20	25	30	35	40	45	50	55	60	
Addition of New Requirements	1.4	1.4	1.4	1.4	1.4	1.6	1.7	1.6	2.3	1.4	1.4	1.4	1.5
Dropping of Requirements	-0.1	-0.3	-0.5	-0.4	-0.4	-1.0	-1.2	-1.1	-0.6	-0.2	-1.4	-0.3	-0.6
Reclassification as System	-0.2	-0.2	-0.2	-0.2	-0.2	-0.2	-0.2	-0.2	-0.2	-0.2	-0.2	-0.2	-0.2
Classification as Non-Scheduled	-1.8	-1.8	-2.4	-1.5	-2.0	-2.0	-1.6	-1.6	-2.7	-2.7	-0.8	-1.7	-1.9
Rescheduled	0.9	-0.8	-0.5	-0.1	-0.6	-0.8	-1.2	-0.9	0.2	0.2	0.5	-0.1	-0.2
Total	0.2	-1.7	-2.2	-0.8	-1.8	-2.4	-2.5	-2.2	-1.0	-1.5	-0.5	-0.9	-1.4

Table 2. Scenario No. 1. The scheduled start times are after the randomization that resulted in Fig. 2. The technique numbers are from the SRS, Appendix G as it is being updated. The last three columns are forecaster number, MIPS required, and required response time, respectively.

Start Time (Sec)	Technique		Fcst'r No.	MIPS Req.	Req. Resp. Time (Sec)
1	7.2.3	Hydrologic Forecast Model	4	0.1	300
4	3.2.7	Interactive Skew-T	1	0.9	5
19	3.2.7	Interactive Skew-T	1	0.9	5
69	11.2.2.1	Compose Graphical Watch/Warning	3	1.0	5
74	12.4.2	Display Status of Watches/Warnings	3	0.6	5
74	12.4.3	Display Extent of Watches/Warnings	3	0.6	5
99	4.1.2.1.3	Plot Mesonet Data	3	5.0	5
110	4.1.1.1	Plot Surface/Upper Air Data	4	3.8	7
115	11.2.1.3.1	Prepare Hydrologic Product	4	1.0	5
139	9.2.2	Retrieve Information at Cursor	3	1.0	2
142	4.1.2.1.1	Plot SAO Reports	2	2.3	15
160	5.2.3	Extrapolate Image Feature	5	2.9	1
190	5.2.3	Extrapolate Image Feature	5	2.9	1
202	11.2.1.4.1	Prepare Marine Product	2	1.0	5
204	11.2.1.4.1	Prepare Marine Product	2	1.0	5
219	9.2.2	Retrieve Information at Cursor	3	2.0	1
220	6.4.1	Extract Value from Graphic	5	3.0	1
221	6.4.1	Extract Value from Graphic	5	3.0	1
222	6.4.1	Extract Value from Graphic	5	3.0	1
223	6.4.1	Extract Value from Graphic	5	3.0	1
234	9.2.2	Retrieve Information at Cursor	1	1.0	2
240	11.2.1.5.1	Prepare Public Product	5	1.0	5
262	11.2.2.1	Compose Graphical Watch/Warning	2	1.0	5
267	12.4.2	Display Status of Watches/Warnings	2	0.6	5
267	12.4.3	Display Extent of Watches/Warnings	2	0.6	5
280	4.1.2.1.4	Plot Upper air Data	5	1.8	10
289	11.2.2.1	Compose Graphical Watch/Warning	1	1.0	5
291	12.4.2	Display Status of Watches/Warnings	1	0.6	5
291	12.4.3	Display Extent of Watches/Warnings	1	0.6	5

Table 3. Same as Table 2, except for Scenario No. 2.

Start Time (Sec)		Technique	Fcst'r No.	MIPS Req.	Req. Resp. Time (Sec)
4	4.1.3.1	Contour Surface Data	1	2.6	3
5	6.4.1	Extract Value from Graphic	4	3.0	1
6	6.4.1	Extract Value from Graphic	4	3.0	1
7	6.4.1	Extract Value from Graphic	4	3.0	1
8	6.4.1	Extract Value from Graphic	4	3.0	1
10	4.1.3.3	Display Streamlines	5	1.6	3
29	3.2.7	Interactive Skew-T	1	0.9	5
44	3.2.7	Interactive Skew-T	1	0.9	5
50	4.1.5.1	Plot Time Series of Surface Data	5	0.6	18
65	11.2.1.5.1	Prepare Public Product	4	1.0	5
69	11.2.2.1	Compose Graphical Watch/Warning	3	1.0	5
74	12.4.2	Display Status of Watches/Warnings	3	0.6	5
74	12.4.3	Display Extent of Watches/Warnings	3	0.6	5
95	11.2.1.5.1	Prepare Public Product	4	1.0	5
99	4.1.1.1	Plot Surface/Upper Air Data	3	3.8	30
102	4.1.2.1.3	Plot Mesonet Data	3	5.0	5
114	4.1.1.1	Plot Surface/Upper Air Data	3	3.8	7
127	5.2.3	Extrapolate Image Feature	2	2.9	1
139	9.2.2	Retrieve Information at Cursor	3	1.0	2
140	4.1.2.1.4	Plot Upper Air Data	4	1.8	10
142	4.1.2.1.1	Plot SAO Reports	2	2.3	15
205	4.1.5.4	Plot Time Series of Profiler Data	5	1.8	15
207	11.2.2.1	Compose Graphical Watch/Warning	2	1.0	5
212	12.4.2	Display Status of Watches/Warnings	2	0.6	5
212	12.4.3	Display Extent of Watches/Warnings	2	0.6	5
229	9.2.2	Retrieve Information at Cursor	1	1.0	2
249	9.2.2	Retrieve Information at Cursor	3	2.0	1
252	11.2.1.4.1	Prepare Marine Product	2	1.0	5
254	11.2.1.4.1	Prepare Marine Product	2	1.0	5
264	9.2.2	Retrieve Information at Cursor	3	1.0	2
265	5.2.3	Extrapolate Image Feature	4	2.9	1
274	11.2.2.1	Compose Graphical Watch/Warning	1	1.0	5
279	12.4.2	Display Status of Watches/Warnings	1	0.6	5
279	12.4.3	Display Extent of Watches/Warnings	1	0.6	5
285	5.2.3	Extrapolate Image Feature	4	2.9	1
295	4.1.3.1	Contour Surface Data	4	2.6	3

Table 4. Same as Table 2, except for Scenario No. 3.

Start Time (Sec)		Technique	Fcst'r No.	MIPS Req.	Req. Resp. Time (Sec)
1	7.2.3	Hydrologic Forecast Model	4	0.1	300
4	3.2.7	Interactive Skew-T	1	0.9	5
15	6.4.2	Determine Gridpoints from Contours	2	4.1	5
19	3.2.7	Interactive Skew-T	1	0.9	5
20	2.2	Check Consistency of Entered Information	2	21.0	2
33	6.4.2	Determine Gridpoints from Contours	5	4.1	5
38	2.2	Check Consistency of Entered Information	5	21.0	2
69	11.2.2.1	Compose Graphical Watch/Warning	3	1.0	5
74	12.4.2	Display Status of Watches/Warnings	3	0.6	5
74	12.4.3	Display Extent of Watches/Warnings	3	0.6	5
75	6.4.2	Determine Gridpoints from Contours	2	4.1	5
80	2.2	Check Consistency of Entered Information	2	21.0	2
93	6.4.2	Determine Gridpoints from Contours	5	4.1	5
98	2.2	Check Consistency of Entered Information	5	21.0	2
99	4.1.2.1.3	Plot Mesonet Data	3	5.0	5
110	4.1.1.1	Plot Surface/Upper Air Data	4	3.8	7
115	11.2.1.3.1	Prepare Hydrologic Product	4	1.0	5
135	6.4.2	Determine Gridpoints from Contours	2	4.1	5
139	9.2.2	Retrieve Information at Cursor	3	1.0	2
140	2.2	Check Consistency of Entered Information	2	21.0	2
153	6.4.2	Determine Gridpoints from Contours	5	4.1	5
158	2.2	Check Consistency of Entered Information	5	21.0	2
182	9.1.2	Edit Radar Coded Message	6	1.2	10
195	6.4.2	Determine Gridpoints from Contours	2	4.1	5
200	2.2	Check Consistency of Entered Information	2	21.0	2
213	6.4.2	Determine Gridpoints from Contours	5	4.1	5
218	2.2	Check Consistency of Entered Information	5	21.0	2
219	9.2.2	Retrieve Information at Cursor	3	2.0	1
227	9.1.2	Edit Radar Coded Message	6	4.0	15
234	9.2.2	Retrieve Information at Cursor	1	1.0	2
255	6.4.2	Determine Gridpoints from Contours	2	4.1	5
260	2.2	Check Consistency of Entered Information	2	21.0	2
273	6.4.2	Determine Gridpoints from Contours	5	4.1	5
278	2.2	Check Consistency of Entered Information	5	21.0	2
289	11.2.2.1	Compose Graphical Watch/Warning	1	1.0	5
291	12.4.2	Display Status of Watches/Warnings	1	0.6	5
291	12.4.3	Display Extent of Watches/Warnings	1	0.6	5

Table 5. Results of simulation Model No. 1 for the three workstation scenarios. The second column (MIPS) is the total computer processing capacity. The third column is the percent of the cases the required response time is met. The fourth column is the percent of the cases the technique completes in double the response time. The fifth column (Increase in Response Times) is the percent increase in response time for those techniques not meeting the requirement (i.e., those not included in column 3). The sixth column is the percent of the time the amount of computer processing used is at its rated maximum capacity. The final column is the percent utilization of the maximum computer processing capacity over the 5-minute period (i.e., the total instructions used in relation to the instructions available--MIPS X 300).

Scenario No.	MIPS	Response Times Met (%)	Twice Response Times Met (%)	Increase in Response Times (%)	Proc. Capacity Used to Max (%)	Util. (%)
1	8	15.6	43.6	337	46.8	93.1
	9	54.7	78.0	152	19.7	82.9
	10	82.1	95.4	97	8.3	74.7
	11	88.1	99.0	55	4.3	67.9
2	8	9.3	25.3	501	69.2	95.9
	9	37.8	67.8	202	41.7	87.7
	10	69.9	85.3	149	24.1	79.5
	11	81.4	92.8	108	15.3	72.3
	12	85.8	96.6	71	10.8	66.3
3	10	23.3	47.4	563	77.7	93.5
	11	46.5	59.1	439	50.2	85.6
	12	55.0	67.7	353	37.0	78.6
	14	61.6	72.7	222	23.2	67.5
	16	70.5	73.0	187	15.8	59.1
	18	72.1	89.9	133	11.8	52.6
	20	72.8	94.6	114	10.6	47.3
	22	73.0	97.9	68	7.7	43.0
	23	73.0	98.9	62	7.3	41.2

Table 6. Same as Table 5, except for Model 2. The uninterruptible scheduled workload was set at 1.5 MIPS.

Scenario No.	MIPS	Response Times Met (%)	Twice Response Times Met (%)	Increase in Response Times (%)	Proc. Capacity Used to Max (%)	Util. (%)
1	8	93.0	95.7	200	46.8	93.1
	9	98.2	99.0	98	19.7	82.9
	10	99.7	99.9	49	8.3	74.7
	11	100.0	100.0	--	4.3	67.9
2	8	67.3	73.1	1244	69.2	95.9
	9	87.3	93.8	272	41.7	87.7
	10	93.7	98.5	87	24.1	79.5
	11	97.9	100.0	15	15.3	72.3
	12	99.1	100.0	9	10.8	66.3
3	10	63.4	69.8	168	77.7	93.5
	11	68.5	71.6	202	50.2	85.6
	12	70.6	72.6	138	37.0	78.6
	14	72.0	92.7	95	23.3	67.5
	16	72.6	95.9	64	15.8	59.1
	18	72.9	98.9	41	11.8	52.6
	20	72.9	99.5	23	10.6	47.3
	22	73.0	99.7	44	7.7	43.0
	23	90.8	99.9	38	7.3	41.2

Table 7. Same as Table 5, except for Model 3. The uninterruptible workload was set at 1.5 MIPS. The last column is the average percent decrease (improvement) in response time for those techniques meeting the response time.

Scenario No.	MIPS	Response Times Met (%)	Twice Response Times Met (%)	Increase in Response Times (%)	Proc. Capacity Used to Max (%)	Util. (%)	Decrease in Response Times (%)
1	8	95.9	96.7	227	54.5	93.1	70.7
	9	98.7	99.2	112	29.6	82.9	71.3
	10	99.5	100.0	17	18.1	74.7	63.5
	11	99.9	100.0	8	14.4	67.9	64.1
2	8	67.9	73.3	1303	72.7	95.9	69.7
	9	87.7	93.8	339	48.7	87.7	70.4
	10	94.8	98.7	91	33.1	79.5	63.8
	11	97.4	100.0	14	25.5	72.3	64.3
	12	99.0	100.0	11	18.5	66.3	61.5
3	10	68.4	70.3	222	79.3	93.7	63.2
	11	71.1	72.1	172	59.8	85.6	65.9
	12	72.3	88.9	138	46.7	78.7	67.6
	14	72.9	94.8	91	32.3	67.6	69.8
	16	73.0	97.1	60	25.7	59.2	72.3
	18	73.0	98.4	38	18.8	52.6	69.7
	20	73.0	99.4	22	16.7	47.4	69.7
	22	73.0	99.5	46	13.0	43.0	75.6
	23	95.6	99.9	52	12.4	41.2	58.4

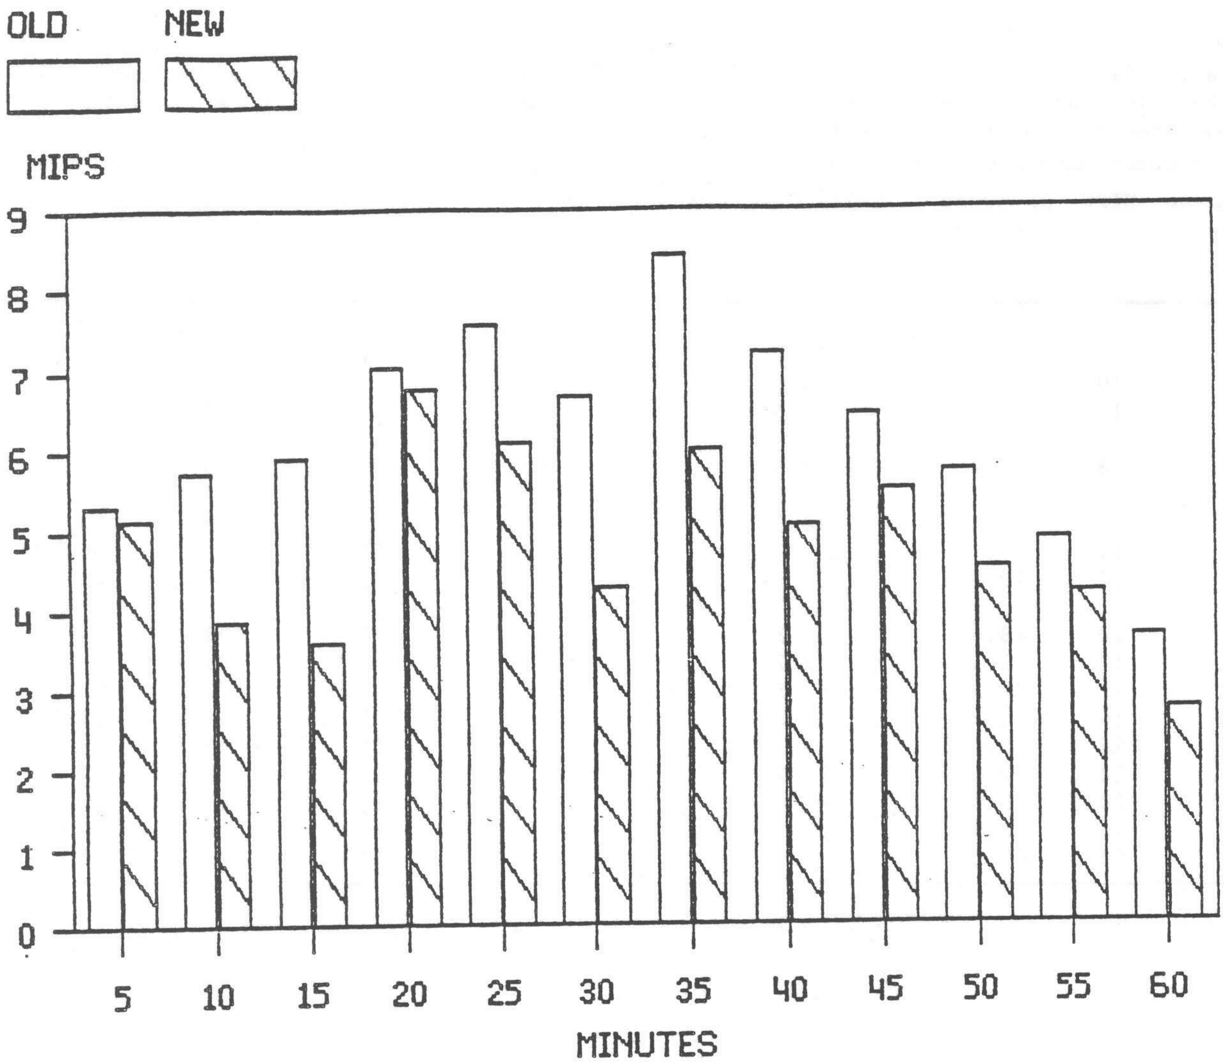

Figure 1. Estimated millions of instructions per second (MIPS) required to support meteorological and hydrological techniques in 5-min segments of a worst case hour. Both initial estimates contained in the AWIPS SRS (OLD) and estimates recently revised (NEW) are shown. The revised estimates are for only scheduled processing and explicitly exclude ad hoc requests made at workstations. The original estimates, although not addressing ad hoc requests explicitly, did purport to include all technique processing including that portion which is now considered to be non-scheduled.

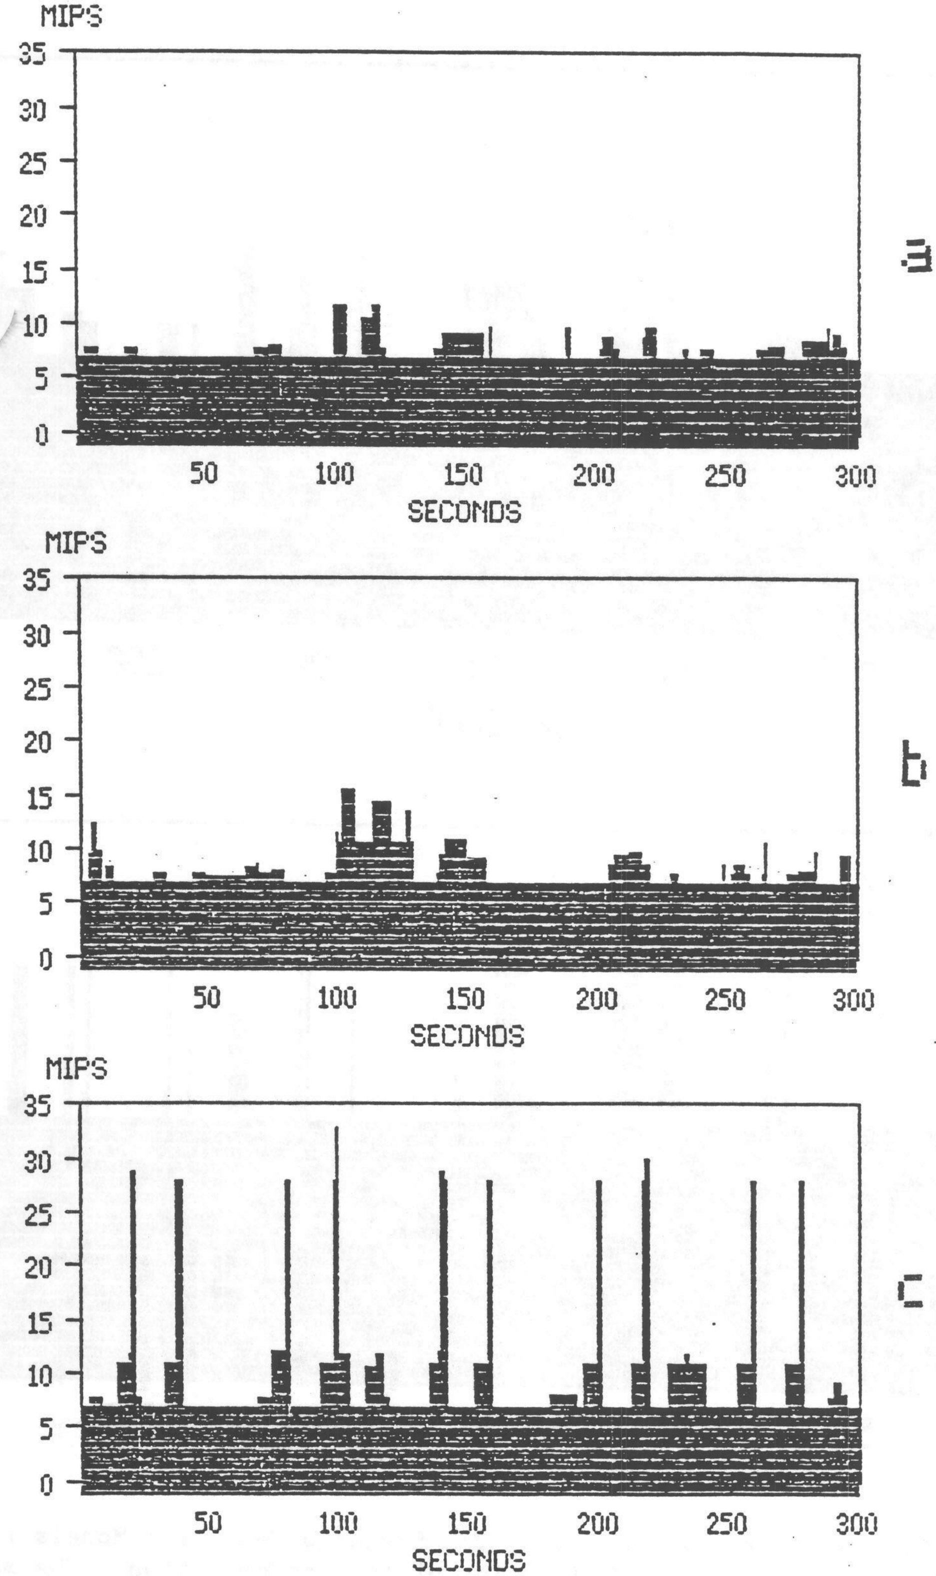

Figure 2. MIPS that would be used on a computer system of unlimited capacity over a 5-min period for Model 1 for Scenarios No. 1 (a), No. 2 (b), and No. 3 (c) for one realization of the randomization.

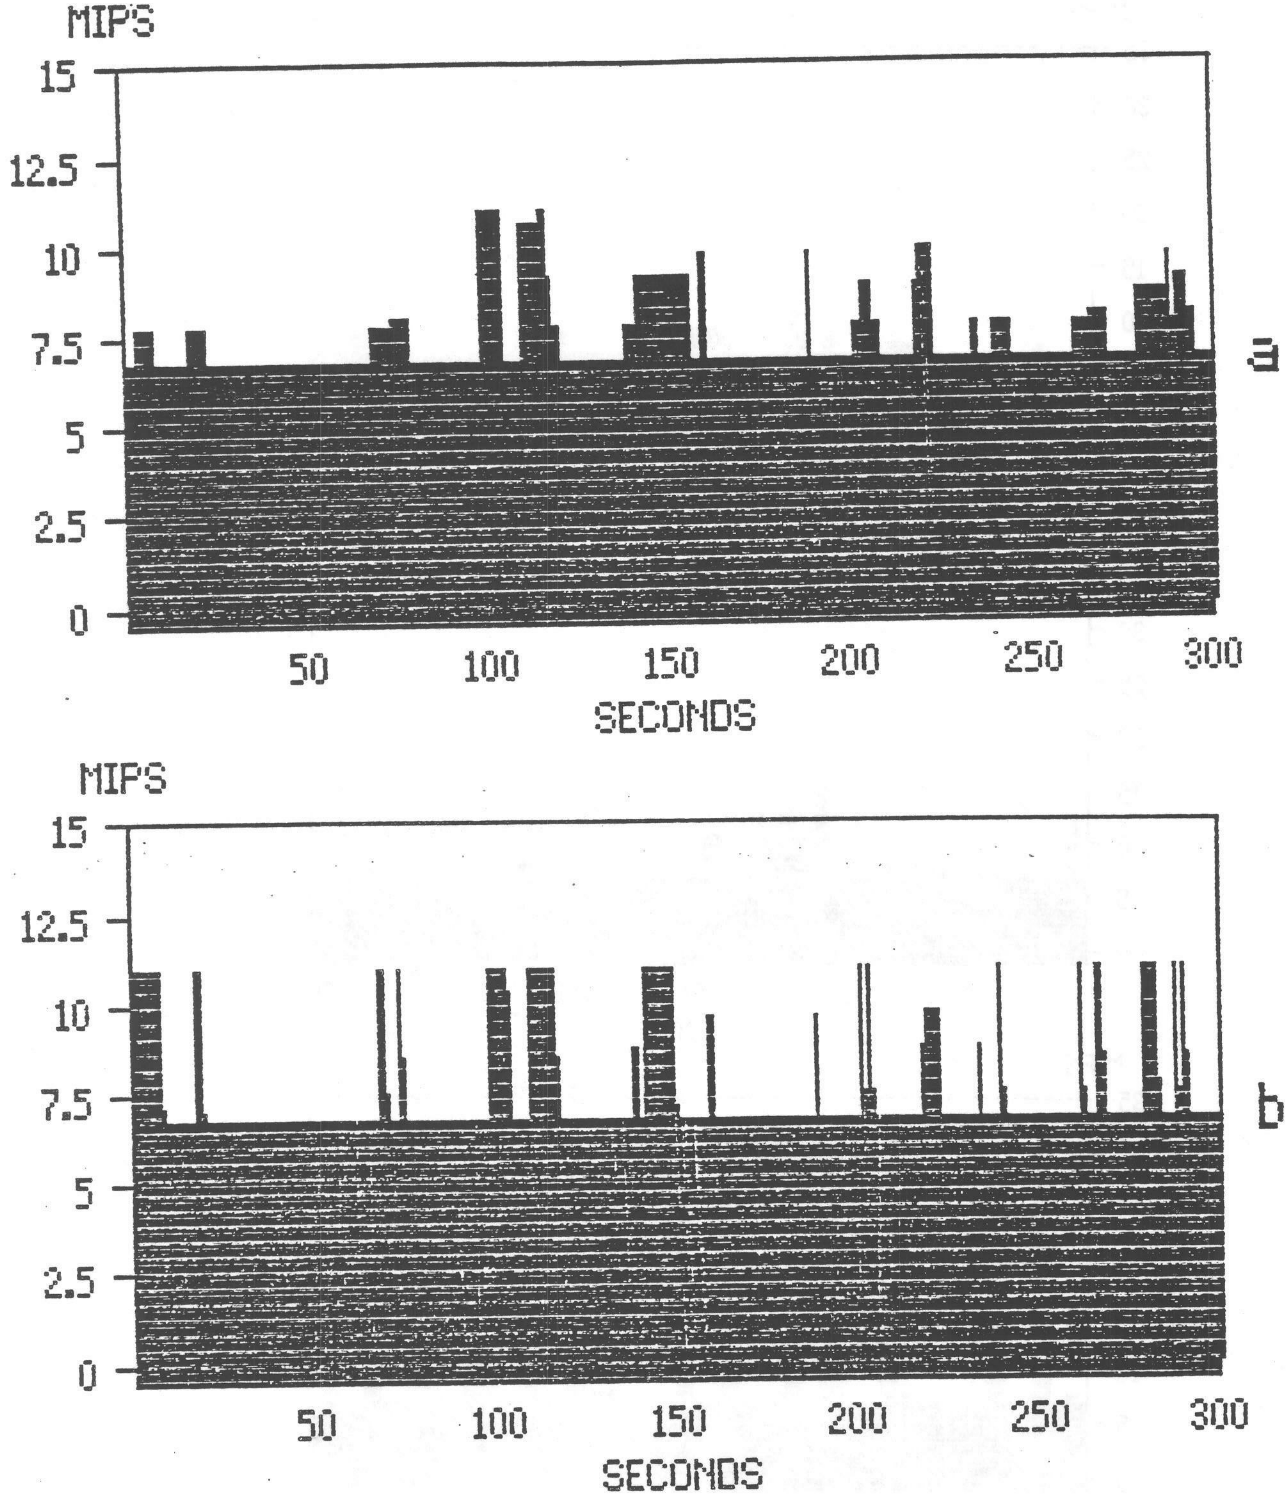

Figure 3. MIPS used over a 5-min period for Scenario No. 1 for Models 1 and 2 (a) and Model 3 (b) for one realization of the randomization. The maximum computer capacity was 11.0 MIPS, and 1.5 MIPS of the scheduled load of 6.7 MIPS were reserved for scheduled (non-interruptable) processing.

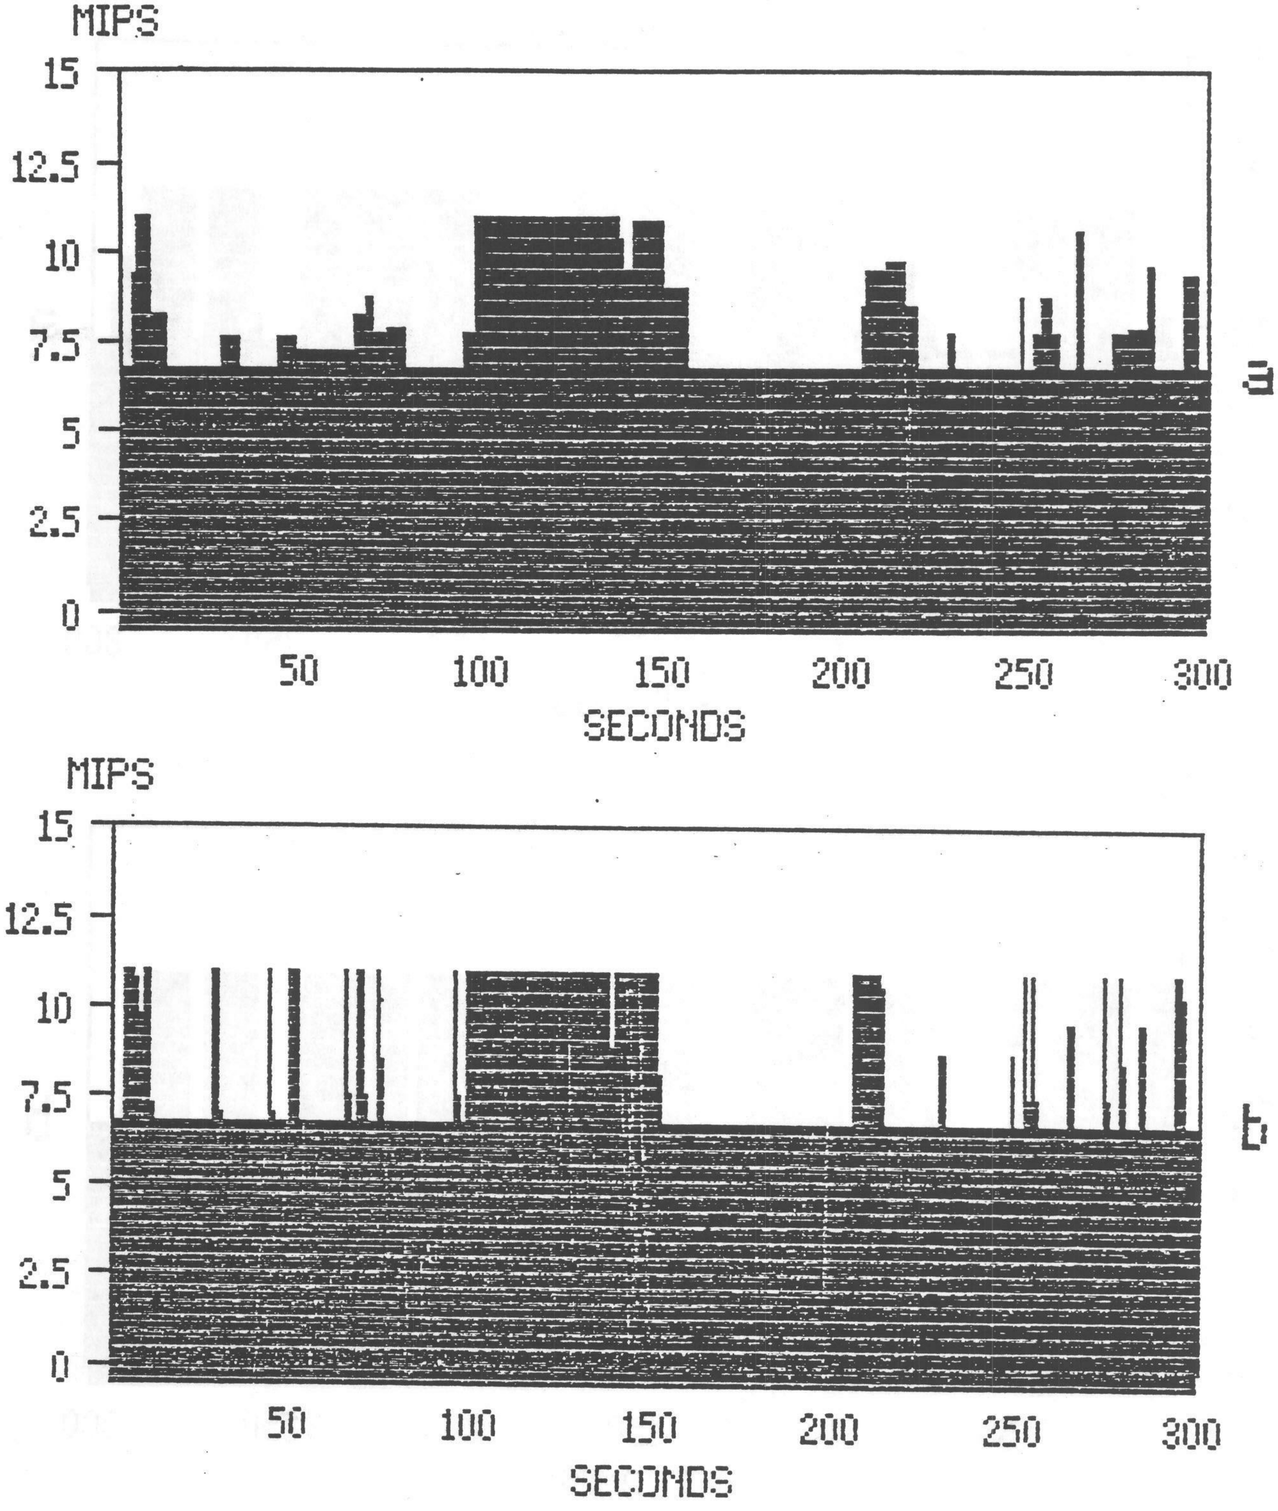

Figure 4. Same as Fig. 3, except for Scenario No. 2.

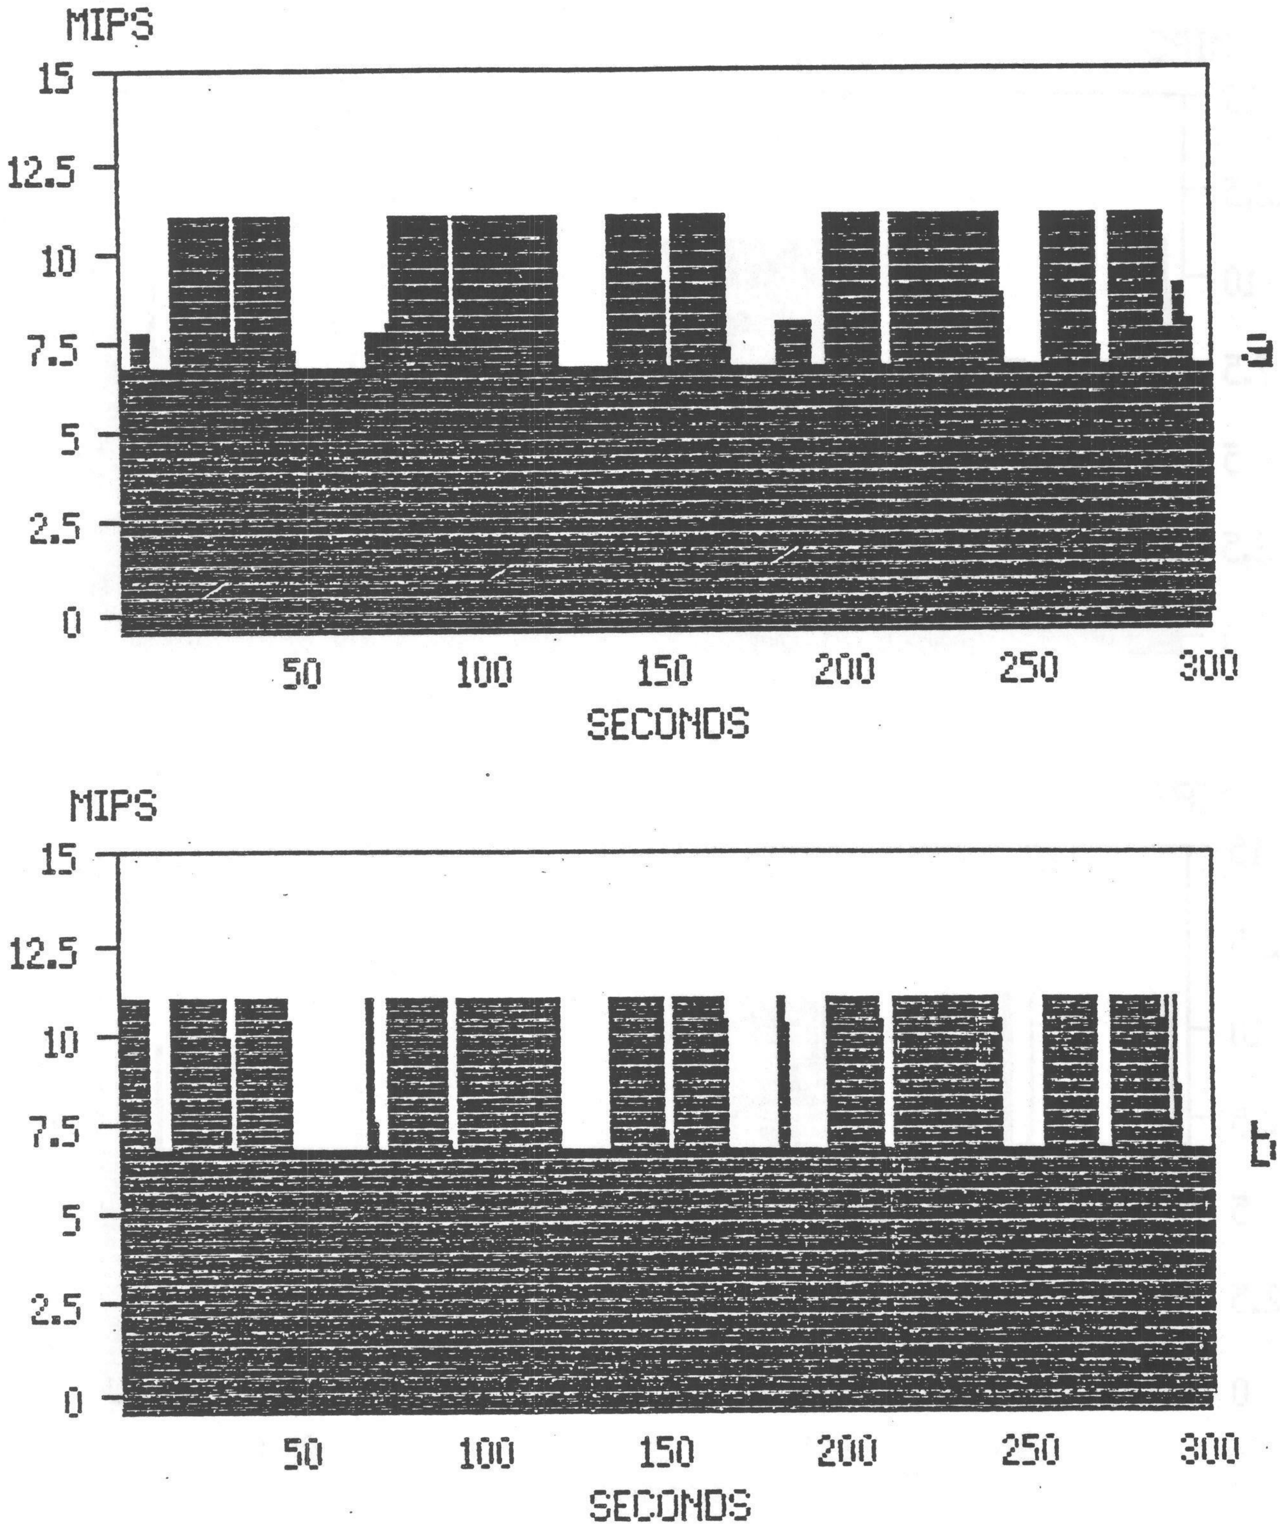

Figure 5. Same as Fig. 3, except for Scenario No. 3.