Superelement Management In

A Network Environment

Keith H. Redner
President

Carl W. Hennrich
Consultant

DANFORD ASSOCIATES
350 West Fifth Street
San Pedro, CA 90731
(213) 514-2368

The 1990 MSC/NASTRAN World User Conference
Los Angeles - 26 - 30 March 1990

1.0 INTRODUCTION

The last few years have seen a veritable explosion in the performance capabilities of the computation components used by engineers in the MCAE arena. Hardware capabilities include vector processors, parallel architectures, electronic input/output devices and networks of non-homogeneous computing elements. Ever larger memories and ever faster clock cycles mean that the engineer can feasibly solve large complex problems that were considered impossible or impractical only a few years ago.

Operating systems such as UNIX are sophisticated and more complex than the operating systems of a few years ago. Structural analysis codes such as MSC/NASTRAN are providing a rapidly increasing set of capabilities including non-linear analysis, layered composite materials, design optimization, shape optimization, a new Executive System in Version 66 and unlimited size capacity. These factors provide the engineer with an awesome capability potential but, at the same time, present him with the requirement for orchestrating this sophisticated hardware and software to successfully complete his analysis. The potential of such an environment is illustrated by a recent paper by Foster and Clifford (Ref. 1).

Since the computing environment of hardware, software and model are increasingly complex and sophisticated, the typical engineer finds himself needing the help of a manager to control the use of these resources and keep track of what happens as the result of his activities in obtaining solutions to his engineering problems. The concept behind the Danford Project Application Manager (D/PAM) is that the engineer should not have to deal with details of computer resource allocation, record keeping, job status inquiry, and a host of other computing environment details, but should be left to concentrate on the engineering aspects of the analysis activity.

2.0 D/PAM OVERVIEW

D/PAM is a distributed application which runs in a heterogeneous network. Figure 2.1 depicts the information flow within D/PAM. Figure 2.2 identifies the principal components and their functions. These components are:

1. Workstation Client. The engineer/user communicates with D/PAM through the Workstation Client. There is one Workstation Client for each possible concurrent active user. The Workstation Client runs on a workstation which is attached to the network via TCP/IP. An interactive, icon driven, pop-up window style interface is provided.

   In the general scenario, the project manager (or lead engineer) defines a project and establishes a general set of defaults for jobs which will be run. These defaults specify (1) directories where input files are to be found, output results are stored, and database files are maintained; (2) the application to be run (D/PAM currently supports MSC/NASTRAN but will include other applications in the future); and (3) the processor on which this application will be run.

   An individual job definition inherits defaults from those defined for the project. The user can change any of these defaults for a specific run. A software component called the Application Expert (see 5. below) is available to assist the user in estimating the resources required for his run. Based on installation and/or project defined value parameters, a "cost" for a specific run is computed for each processor which has been defined. The user can use these "costs" as a guide in selecting the processor on which the run is to execute.

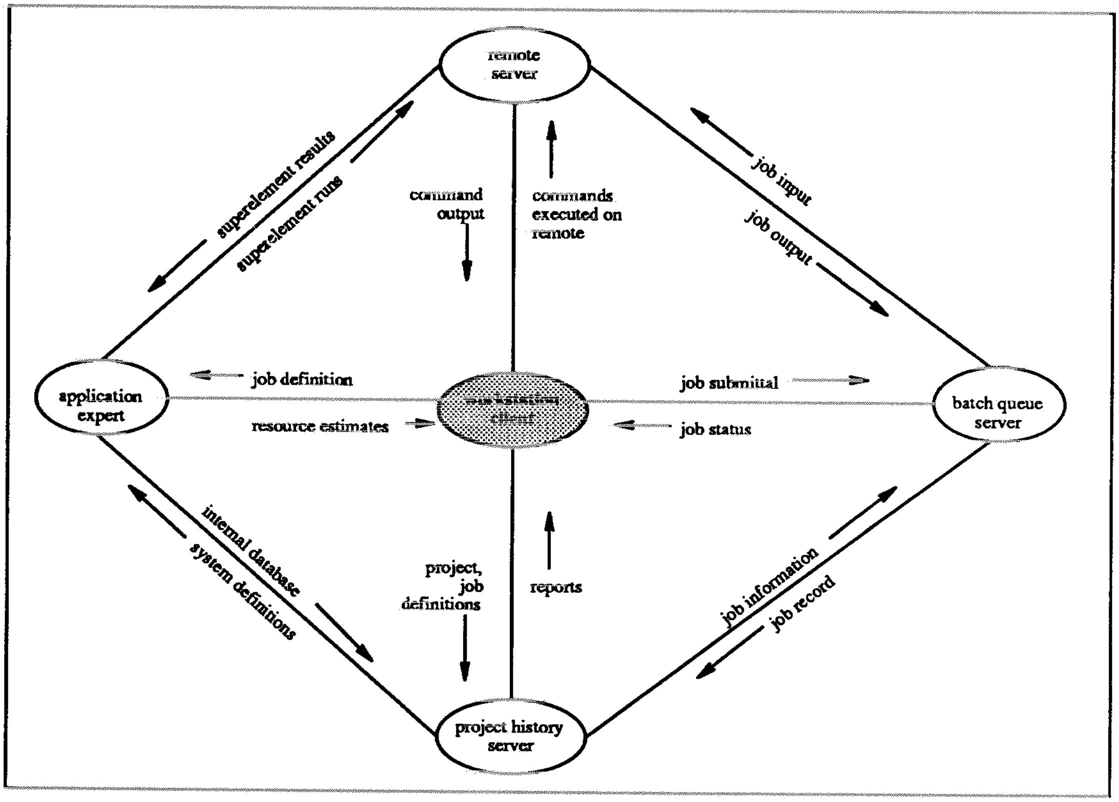

Figure 2.1 Information Flow in D / PAM

Once a job has been defined, the user initiates a run through D/PAM. This menu allows the user to override some of the job definition parameters (e.g., printing output files). He may also schedule the time at which the run is to be started.

D/PAM provides a Batch Queue Server which manages non-application-specific run control with dependencies and related file transfer. The Batch Queue Server (see 2. below) communicates with the Remote Servers (see 3. below) which invoke the requested or implied applications. Supported applications (i.e., NASTRAN) utilize Application Expert services for application-specific features such as superelement dependencies. The Remote Server notifies the Batch Queue Server when a run is complete and the Batch Queue Server sends collected results and statistics of the run to the Project History Server which updates the appropriate project history file. Through the Workstation Client, the user may request the status of his run(s). He may make queries to the project history file and/or cause reports to be generated. The user can add comments to individual runs. The project history fulfills two important needs: (1) it provides the project manager with summaries of all runs which have been made during the project; and (2) it provides the individual engineer/user with an accurate history of his changes and results so that costly reruns can be avoided.

2. Batch Queue Server. Requests for individual run submittals are sent to the Batch Queue Server. The Batch Queue Server routes requests to the specified Remote Server. The Remote Server notifies the Batch Queue Server when the request has been completed and its status. The Batch Queue Server sends request results to the appropriate Project History Server for inclusion in the project history file. The Batch Queue Server can run on any node in the network.

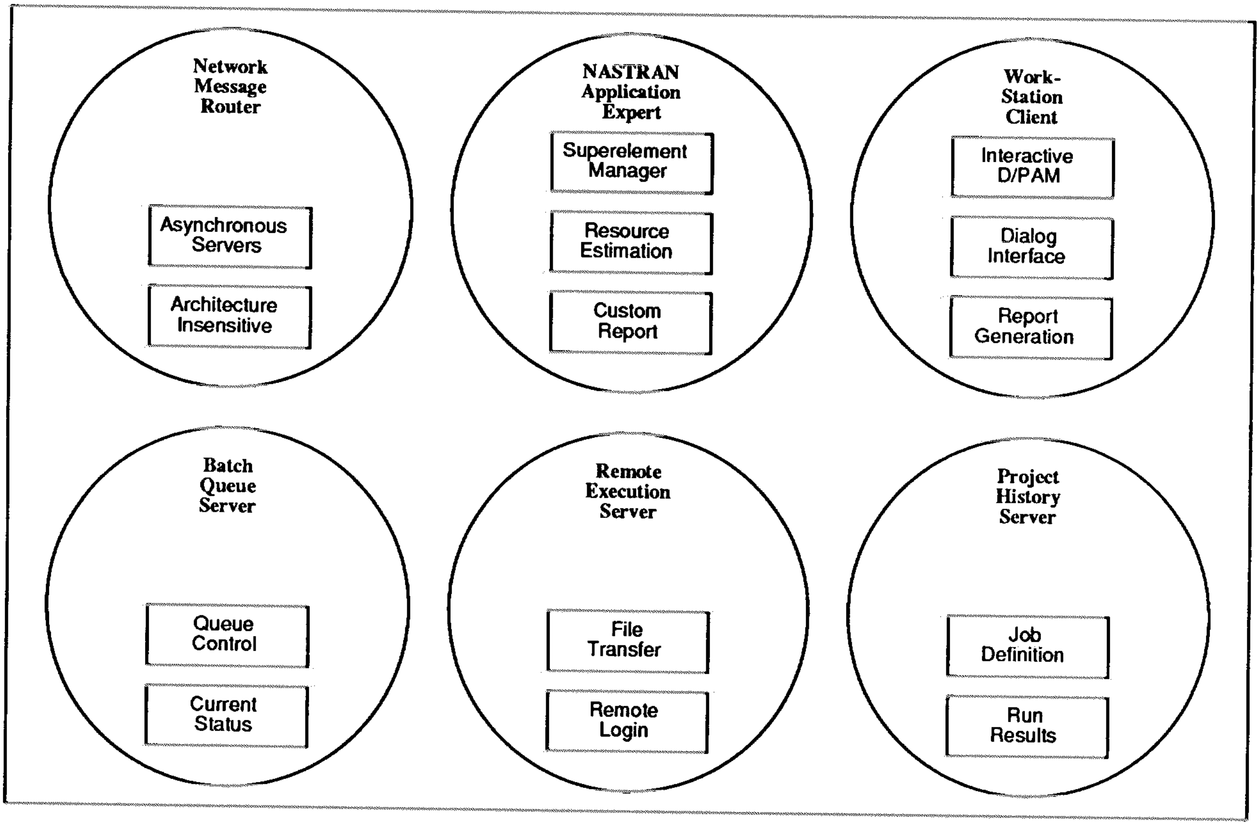

Figure 2.2 D/PAM Components and their Functions

3. Remote Server. A Remote Server runs on each node in the network on which supported applications may execute. It responds to requests from either a Workstation Client or the Batch Queue Server; communicates with the operating system on the node on which it is running in order to execute the request and monitor progress of the request; and sends results of the request to the initiating process. Execution of a command on a supercomputer or mainframe computer will typically be channeled through a local queue.

4. Project History Server. The function of the Project History Server is to process requests to store and/or retrieve project history information. Workstation Clients make requests to both store and retrieve information from a project history file. Batch Queue Servers make requests to store information. Each Project History Server is capable of accessing multiple project history files from multiple concurrent users.

5. Application Expert. For each supported application there is a server to provide application-specific services. These include resource estimation at the time a user defines a run, special runtime support, and specialized report format generation. The NASTRAN Expert scans the information which defines a run (input data file, run parameters, etc.), estimates computer resources (CPU time, disk space, costs), and sequences superelement run tasks and restarts.

6. Network Router. The Network Router is responsible for processing all message traffic between the various clients and servers in D/PAM. These include the Workstation Client, the Batch Queue Server, the Remote Server, the Project History Server, and the Application

Expert. Usually, there is one Network Router running in the system although, for greater reliability, a second Network Router may be active but dormant. The Network Router is much like the central telephone exchange. It passes along messages which have a specific destination (like a telephone call when the number is known). The Network Router also has a broadcast capability. For example, the Workstation Client wishes to know all the active Remote Servers. A message is sent to the Network Router which, in turn, broadcasts this message (much like calling a number which rings in many places at once).

The Network Router is not actually depicted in the information flow figure since it is transparent to each of the client and server processes. All data formats and protocols used throughout D/PAM are completely independent of machine and operating system architecture. Like remote-procedure-call networking models, D/PAM is built directly on TCP/IP sockets, but uses true asynchronous servers exclusively.

Figure 2.3 shows the D/PAM components as they might be configured in a typical engineering network.

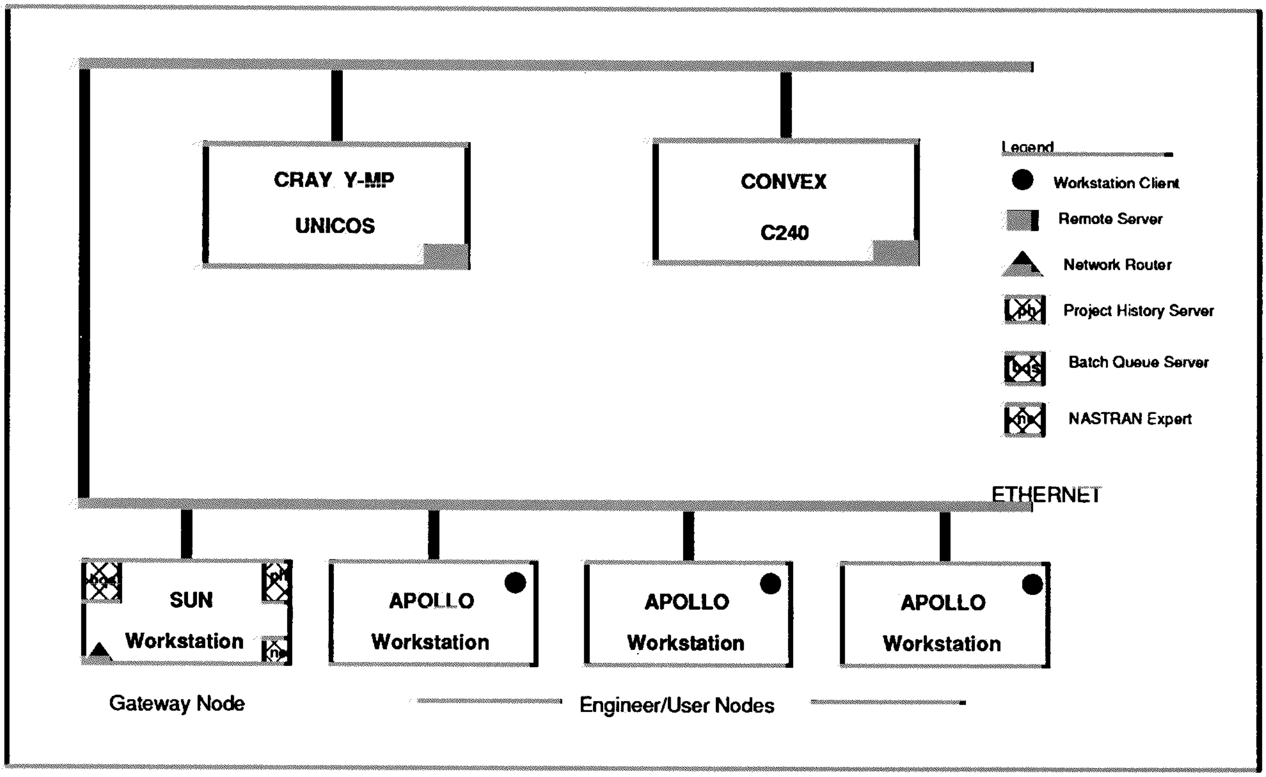

Figure 2.3 D/PAM in a Typical Engineering Network

3.0 SUPERELEMENT MANAGEMENT

Redner and Hennrich (Ref. 2) have indicated several MSC/NASTRAN applications of D/PAM. Superelement solutions in MSC/NASTRAN provide an ideal type of computer application for the D/PAM concept. In general, the models are large. This translates to considerable amounts of computer time to perform the solution and the data base files can consume large amounts of secondary storage. Often, there is more than one engineer working on the model and solution.

3.1 Concurrent Superelement Processing

Large superelement models are often analyzed in several runs with each run processing a collection of superelements called a group. This procedure allows control over the CPU consumed per run, the peripheral storage space required to hold the data base for the group, and other resource and organizational constraints. Several of the group runs can be made during the same time period on a fast, multiprocessing computer since most of the groups, if properly defined, are independent of each other. If the computer has more than one CPU, this leads to the natural extension of high level parallel processing. Deuermeyer and Clifford (Ref. 3) have demonstrated the performance potential of such a computer to significantly reduce the elapsed time of a superelement analysis. By implication, a network of computational processors sharing disk resources, sometimes called a cluster, is also capable of concurrent/parallel processing. Carrying the concept to networks of computers is conceptually simple but logistically difficult for the user to utilize in the sense that large data base files must be indirectly accessed, transferred, and/or translated to be generally usable.

For any of the concurrent schemes described above, the user must set up several runs involving a large number of data base files. These runs must be submitted in a prescribed order and the data base files must be off-loaded and re-loaded appropriately if the process is to be successful. This run/resource management problem is a bookkeeping burden for the engineer user but is essential to cost effective utilization of the available computational resources.

When analyzing a large superelement model with D/PAM, the user defines a single "run" to the D/PAM software system which, in turn, organizes and submits the required "tasks" to the appropriate available computational elements of the user's computer network. The generation and execution of the tasks, data base files, dependency and concurrent/parallel relationships, report generation, and re-run activities are all automatically handled by D/PAM.

3.2 The MSC/NASTRAN Application Expert Superelement Manager

D/PAM supports superelement solution sequences 61 and 63 in MSC/NASTRAN Version 66 (and 66A) through a Superelement Manager which is a part of the NASTRAN Expert. The principal functions performed by the Superelement Manager are as follows:

1. Run Definition. A superelement solution submittal (called a run in D/PAM terminology) is divided into a number of individual computer tasks with corresponding individual data base files. Figure 3.1 shows the data and execution flow for a simple multilevel superelement model. The initial task scheduled by the Superelement Manager is the SEMAP task which produces the superelement map and creates the basic data base which will be read by all other tasks. The Superelement Manager reads the output file produced by the SEMAP task and extracts various information including the superelement tree, time estimates, data base size estimates and parameters for each superelement. Based on this information, the individual tasks for the run are constructed. Smaller tip superelements are often grouped and scheduled to be processed together in order to provide tasks of more or less equivalent CPU time (e.g., superelements 101 and 102 in Figure 3.1).

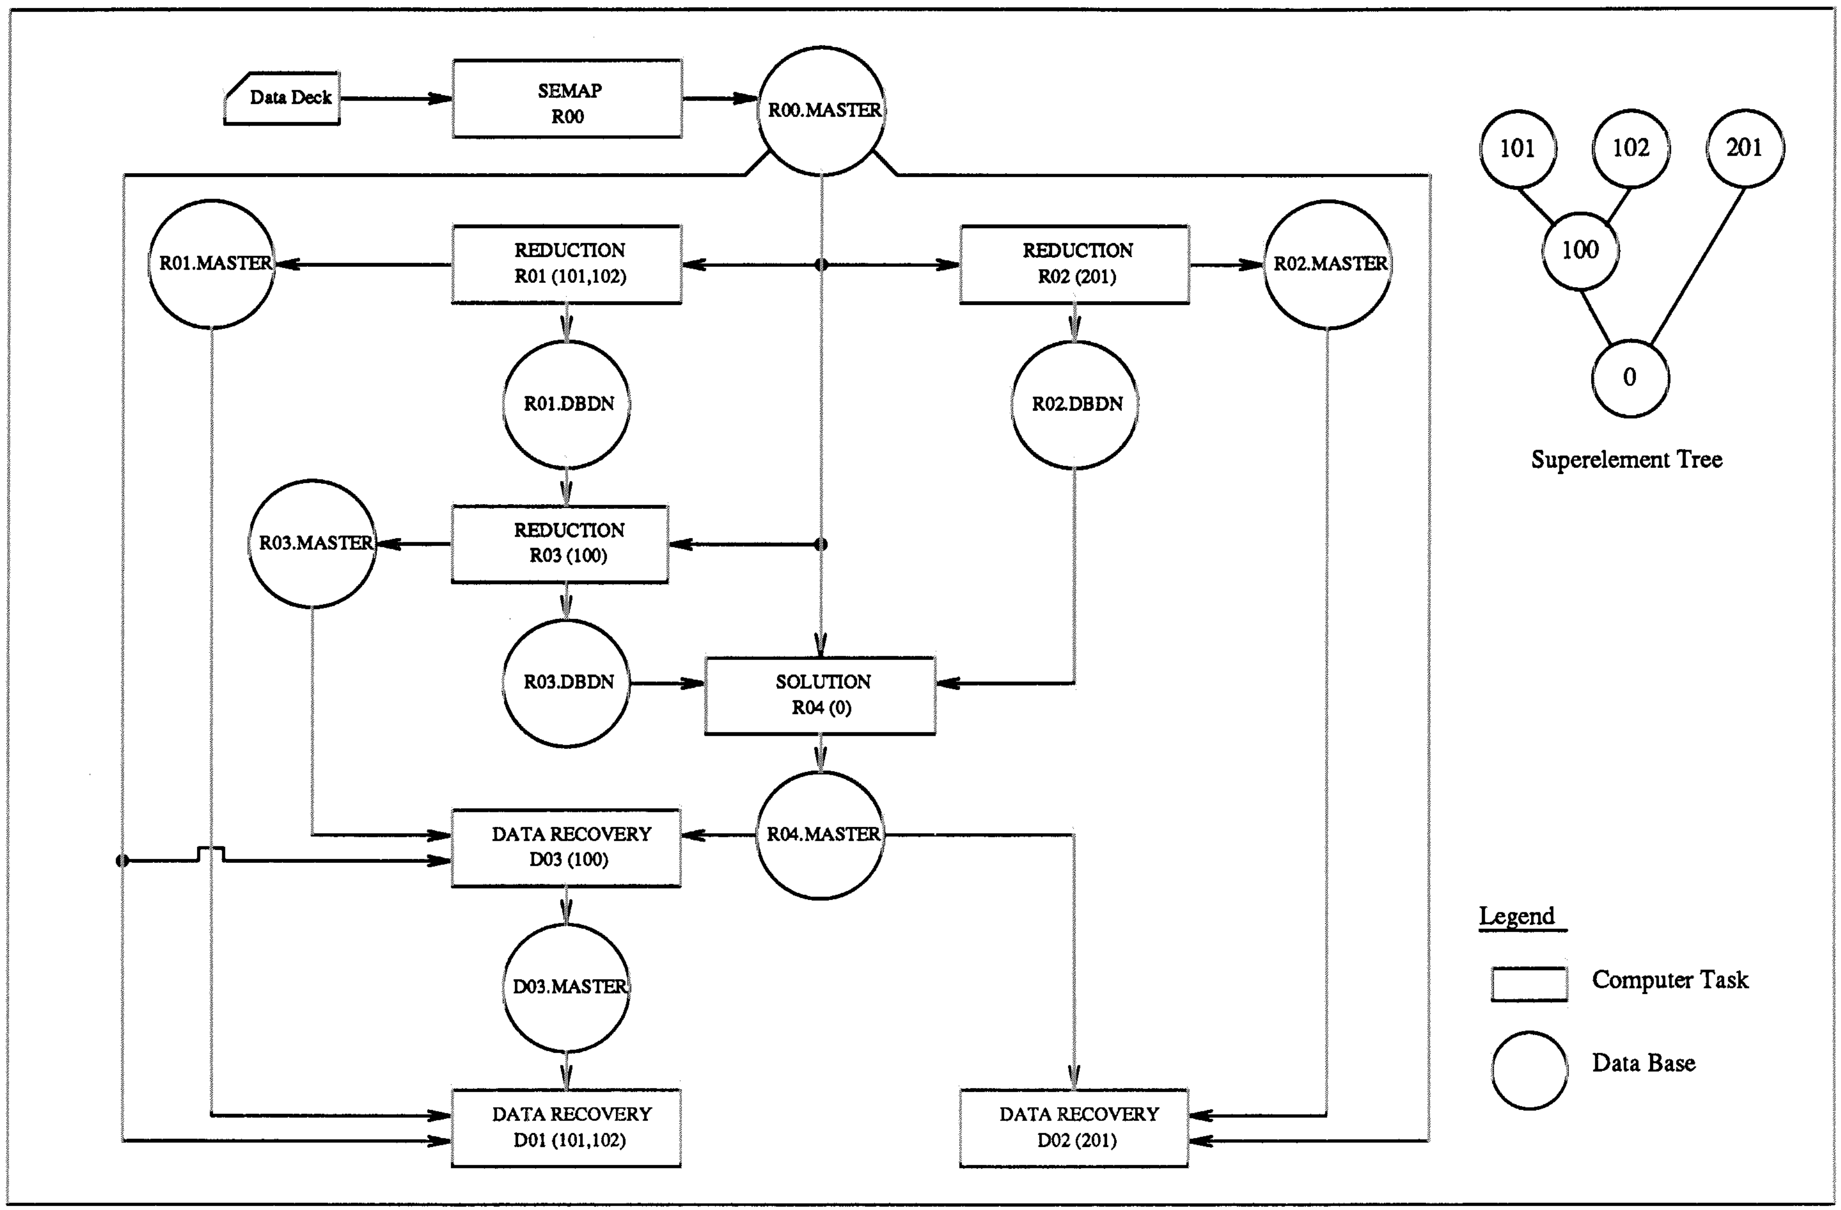

Figure 3.1 Data and Task Flow in a Superelement Solution

2. Task Sequencing. A number of tasks will be independent of each other (R01, R02 and D01, D02 in the example) and thus may be processed in parallel on networks with multiple processors and shared data storage. More computer tasks, each requiring fewer computer resources than a single run task, are easier to schedule and will generally provide shorter elapsed time to final solution (since a number of these tasks will be independent). D/PAM knows the dependencies of the individual tasks from the information which was assembled following the SEMAP task. Independent tasks are submitted by the Superelement Manager to the network. When an individual task completes, the Superelement Manager is notified (by a step which follows the NASTRAN step in the execution script). If the task completed successfully, and if all other required tasks have also successfully completed, all dependent tasks are submitted to the network. Through the project history, the Superelement Manager knows the status of all tasks. Tasks which have completed successfully are not reexecuted on a restart (unless a data change was made which invalidates the results).

3. Data Base File Management. Because the many data base files of a superelement model are often very large, the user is faced with the time consuming task of allocating them to secondary storage devices with sufficient space and insuring that the appropriate data base files are available for each run. Mistakes made performing these tasks can be time consuming and costly (particularly when runs fail to complete due to insufficient space allocation). The Superelement Manager assists the user in this area in the following ways:

   a. The File Management Section statements for each task are automatically generated. The Manager knows which data base files are required for each task.

   b. When a data base file is initially created, the Manager matches estimated space required with that actually available at the time of task submission. Appropriate ASSIGN and INIT statements are created to divide the space across more than one storage device (based on directory information which is provided by the user to D/PAM at project definition and job definition time).

4. Status Inquiry. The user may request the status of a job at any time. The Superelement Manager will respond by reporting the status of every superelement (or an individual superelement) with one of the following:

   The superelement is inactive and has not been run;
   The superelement is inactive and has been successfully run;
   The superelement is inactive and has failed (a rerun is required);
   The superelement is currently active (a run task has been submitted).

   The user may also request the status of a run at any time. The Superelement Manager will respond in this case by reporting the status of each task (superelement group) with one of the following:

   The task has been submitted and is queued;
   The task has been submitted and is currently active;
   The task has completed (with its completion status).

5. Report Generation. In addition to the general (application independent) report generation capability provided within D/PAM, the Superelement Manager provides several report formats which are tailored to superelement jobs.

The user interface with D/PAM for superelement runs takes the following form:

Input Data Deck. The user prepares his input data deck in the usual manner. The NASTRAN bulk data is pointed to by one or more INCLUDE statements. This is required in order to keep the size of the input data file to a minimum since it will be copied with extensions and modifications. INCLUDE statements can appear in other sections of the data file (e.g., Executive Control Section, Case Control Section). These INCLUDE statements will be expanded into the actual data when the task version is created. If the INCLUDE statements in the Bulk Data section rèfer to files which are not directly accessible by the computer on which the execution is to be made, these statements will be expanded when the task version is created. A local copy (task version) is created for purposes of task execution but this copy is not retained (by default).

Run Definition. Through the interactive, icon-based user interface provided in the Workstation Client, the user provides information about the superelement run to be processed including the pathname of the input data deck file, user directory names on the target computer for basic run files and data base files, and the name of a shared pool directory (for temporary overflow). When the user is satisfied with his run definition, he activates the run by selecting the appropriate icon.

Run Status. Through the Workstation Client, the user can request the status of his run at any time. He can request the status be reported by superelement or by task.

As individual tasks complete (with success or not), status information is automatically recorded in the project history file.

4.0 REPORT GENERATION IN D/PAM

The user invokes the Custom Report Generator through an icon in the Workstation Client user interface and selects the form to be used. The form determines which information is to appear in the report and how it is to be displayed. The user can select one of several forms provided within D/PAM or create his own form with a text editor. Consequently, a vast array of reports is possible. Forms, like the one below, require only a short time to create and can display all the information contained in a project history file.

```
sort    date
pgbreak def Totals for Job: {job LJ 8}      {subcpu RJ 7}  {subio RJ 8}
length  60
width   100
heading D A N F O R D / P R O J E C T   A P P L I C A T I O N   M A N A G E R  (D / P A M)
heading
heading                     R U N   R E P O R T --  {NOW}
heading
heading          Job: {job LJ 8}   Title: Superelement Concurrency Prototype
heading
heading      Input: {wscdir LJ 20}   Output: {rmtdir RJ 20}   Data-base: {dbdir RJ 20}
heading
heading              Processor: {proc LJ 19}   MSC/NASTRAN Version: {vs LJ 4}
heading Page{PAGE RJ 3}
heading
heading Date                 Task          CPU       I/O    Status   Superelements
subhead ----------------------------------------------------------------------------------------------------
run     {date LJ 25} {task LJ 4}    {cpu RJ 7}  {io RJ 8}   {status RJ 6}  {se RJ 6}
errmsg
```

Figure 4.1 Example of a Custom D/PAM Report Generation Form

When a custom report is initiated, the Workstation Client retrieves the appropriate records from the Project History Server and passes them to the Custom Report Generator. The report produced is stored as a text file which can be viewed at the terminal or sent to a printer to produce hardcopy. The form above can be used to produce a report complete with pagebreaks, subtotals, and the definition of NASTRAN exit status codes as illustrated in Figure 4.2.

```
        D A N F O R D / P R O J E C T   A P P L I C A T I O N   M A N A G E R  (D / P A M)

                    R U N   R E P O R T -- Fri Jan 19 09:15:34 1990

                 Job: c2x      Title: Superelement Concurrency Prototype

 Input: /danford/pam/data/input   Output: dn10000:/dpam/data   Data-base: dn10000:/dpam/dbs

                 Processor: Apollo/DN10000      MSC/NASTRAN Version: 66A
Page    1
                                                                              MSC/NASTRAN
Date                         Task        CPU         I/O     Status   Superelements
------------------------------------------------------------------------------------------
Wed Jan 17 10:00:13 1990     R00         103          94        0
Wed Jan 17 10:03:15 1990     R01         251         313       -2      101    102
Wed Jan 17 10:03:15 1990     R02          80          63        1      201
Wed Jan 17 10:12:00 1990     R03         102         333        0      100
Wed Jan 17 13:18:10 1990     R00         106          94        0
Wed Jan 17 13:21:17 1990     R02         195         292        0      201
Wed Jan 17 13:27:21 1990     R04         207         311        0        0
Wed Jan 17 13:33:38 1990     D03          78         115        0      100
Wed Jan 17 13:37:02 1990     D02          35          57       -2      201
Wed Jan 17 13:37:02 1990     D01          89         107        0      101    102
------------------------------------------------------------------------------------------
Totals for Job: c2x                     1246        1779


        -----------------------------------------------------------------------
        | Definition of Status Codes                                          |
        |                                                                     |
        |    -2      Warning message(s) present in output file                |
        |    -1      Unable to read output files                              |
        |     0      Normal completion                                        |
        |     1      Fatal error message(s) in output file                    |
        |     2      No EXIT BEGN message in output file                      |
        |     3      No EXIT BEGN and fatal error(s) in output file           |
        |     4      Error message in log file                                |
        |     5      Error message(s) in both log file and output file        |
        |     6      Error message in log file and no EXIT BEGN               |
        !     7      Error message in log file, no EXIT BEGN and fatal error(s) |
        -----------------------------------------------------------------------
```

Figure 4.2 Sample Output from D/PAM Report Generator

5.0 CONCLUSION

Installation and use of the D/PAM software will provide the following benefits:

1. Existing computing hardware is leveraged by D/PAM through its management of computer resources.

2. Engineering hours are saved since D/PAM's management of computer resources in superelement analysis relieves the engineer/user of performing this time consuming task.

3. Costly reruns (or redundant runs) in a superelement analysis are avoided since D/PAM has a complete history of previous runs and generates only the minimum set of tasks required to complete an analysis request.

4. D/PAM provides a consistent user interface which transparently manages a heterogeneous environment, improving user efficiency and productivity.

5. D/PAM is extensible to other applications.

References:

1. Foster N. and Clifford G., "MSC/NASTRAN and a Supercomputer Engineering Environment", Proceedings of the Sixteenth MSC/NASTRAN European Users' Conference, 1989.

2. Redner K. and Hennrich C., "The Danford Project Application Manager", Proceedings of the Sixteenth MSC/NASTRAN European Users' Conference, 1989.

3. Deuermeyer D. and Clifford G., "MSC/NASTRAN on a Multiple CPU Supercomputer", Proceedings of the Fifteenth MSC/NASTRAN European Users' Conference, 1988.