



AD-A223 957

R1: A RULE-BASED CONFIGURER  
OF COMPUTER SYSTEMS

John McDermott

April, 1980

DEPARTMENT  
of  
COMPUTER SCIENCE



DTIC  
ELECTE  
JUL 16 1990  
S D  
GE

DISTRIBUTION STATEMENT A  
Approved for public release;  
Distribution Unlimited

Carnegie-Mellon University

90 07 16 460

DO NOT REMOVE  
\*Z0AAAAAA1906143\*

# R1: A RULE-BASED CONFIGURER OF COMPUTER SYSTEMS

John McDermott

April, 1980

Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By _____	
Distribution/ _____	
Availability Codes	
Dist	Avail and/or Special
A-1	



**Abstract.** R1 is a program that configures VAX-11/780 computer systems. Given a customer's order, it determines what, if any, modifications have to be made to the order for reasons of system functionality and produces a number of diagrams showing how the various components on the order are to be associated. The program is currently being used on a regular basis by Digital Equipment Corporation's manufacturing organization. R1 is implemented as a production system. It has sufficient knowledge of the configuration domain and of the peculiarities of the various configuration constraints that at each step in the configuration process, it simply recognizes what to do. Consequently, little search is required in order for it to configure a computer system. (KR)

APPROVED FOR PUBLIC RELEASE  
DISTRIBUTION UNLIMITED

The development of R1 was supported by Digital Equipment Corporation. The research that led to the development of OPS4, the language in which R1 is written, was sponsored by the Defense Advanced Research Projects Agency (DOD), ARPA Order No. 3597, and monitored by the Air Force Avionics Laboratory under Contract F33615-78-C-1151. The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of Digital Equipment Corporation, the Defense Advanced Research Projects Agency, or the U.S. Government. VAX, PDP-11, UNIBUS, and MASSBUS are trademarks of Digital Equipment Corporation.

## INTRODUCTION

R1<sup>1</sup> is a rule-based system that has much in common with other domain-specific systems that have been developed over the past several years [Amarel 77, Waterman 78]. It differs from these systems primarily in its use of Match rather than Generate-and-test as its central problem solving method; rather than exploring several hypotheses until an acceptable one is found, it exploits its knowledge of its task domain to generate a single acceptable solution. R1's particular area of expertise is the configuring of Digital Equipment Corporation's VAX-11/780 systems. Its input is a customer's order and its output is a set of diagrams displaying the spatial relationships among the components on the order; these diagrams are used by the technician who physically assembles the system.<sup>2</sup> Two inter-dependent activities must be performed in configuring a VAX system.

- The customer's order must be determined to be complete; if it is not, whatever components are missing must be added to the order.
- The spatial relationships among all of the components (including those that are added) must be determined.

The criterion of success for whether a configuration is complete does not reside in any simple test, but involves instead particular knowledge about all the individual components and their relationships. The criterion of successful spatial arrangement is more compact (reflecting the uniform character of geometric structure), but it too involves particular knowledge on a component by component basis. Thus, the task accomplishment is defined by a large set of constraints embodying a large amount of knowledge.

Although a significant portion of this paper is devoted to a description of precisely how R1 goes about doing the configuration task, I have tried to avoid letting the details of R1's inner workings overshadow the domain independent lessons that have emerged from this research. There are two important lessons:

- Recognition knowledge can be used to drive an expert system's behavior, provided that it is possible to determine locally (ie. at each step) whether taking some particular action is consistent with acceptable performance on the task.
- When an expert system is implemented as a production system, the job of refining and extending the system's knowledge is quite easy.

The paper is divided into three sections. The first section describes the VAX-11/780 configuration task and characterizes its difficulty. The second section describes R1 and discusses its evolution

---

<sup>1</sup> Four years ago I couldn't even say "knowledge engineer", now I ...

<sup>2</sup> R1's output for a sample order is shown in Appendix 2.

from a system with only the most limited capabilities to what might fairly be called, a true expert. The third section describes R1's current level of expertise and isolates the design decisions that made the building of R1 straightforward.

## 1. THE TASK

The VAX-11/780 is the first implementation of Digital Equipment Corporation's VAX-11 architecture. It is similar in many respects to the PDP-11, though its virtual address space is  $2^{32}$  rather than  $2^{16}$ . The VAX-11/780 uses a high speed synchronous bus, called the sbi (synchronous backplane interconnect), as its primary interconnect. The central processor, one or two memory control units, one to four massbus interfaces, and one to four unibus interfaces can be connected to the sbi. The massbuses and particularly the unibuses can support a wide variety of peripheral devices. Because the number of system variations is so large, the VAX configuration task is non-trivial.

### 1.1. THE SIZE OF THE TASK

A configurer must have two sorts of knowledge. First, he must have information about each of the components that a customer might order. For each component, the configurer must know the properties that are relevant to system configuration -- eg, its voltage, its frequency, how many devices it can support (if it is a controller), how many ports it has; I will call this knowledge *component information*. Second, he must have rules that enable him to associate components to form partial configurations and to associate partial configurations to form a functionally acceptable system configuration. These rules must indicate what components can (or must) be associated and what constraints must be satisfied in order for these associations to be acceptable; I will call this knowledge *constraint knowledge*.

The difficulty of the VAX configuration task is a function of the amount of component information and the amount of constraint knowledge required to perform the task. It is fairly easy to estimate the amount of component information that is needed. On the average, a configurer must know eight properties of a component in order to be able to configure it appropriately. Currently about 420 components are supported for the VAX.<sup>3</sup> Thus there are over 3300 pieces of component information that a VAX configurer must have access to.

Before R1 was developed, it would have been difficult to estimate accurately the amount of constraint knowledge required for the configuration task. Much of the required knowledge was not

---

<sup>3</sup>Of the 420 components, about 180 are actually bundles composed of various subsets of the remaining 240 components.

written down anywhere and thus the only source of estimates would have been individual human experts. But the experts find the task of quantifying their constraint knowledge foreign. As I extracted this knowledge from them, it became clear that their knowledge takes two forms: (1) The experts have a sparse but highly reliable picture of their task domain. When asked to describe the configuration task, they do so in terms of the subtasks involved and the various temporal relationships among these subtasks. (2) They also have a considerable amount of very detailed knowledge that indicates the features that particular partial configurations and unconfigured components must have in order for the partial configurations to be extended in particular ways. Both sorts of knowledge are easily expressible as rules. I extracted 480 rules. Of these, 96 define situations in which some subtask should be initiated. The other 384 rules define situations in which some partial configuration should be extended in some way.

## 1.2. THE CONSTRAINTS

This subsection provides two examples of specific subtasks that can arise within the configuration task and indicates for each (1) the constraint knowledge involved, (2) the informational demands imposed by that constraint knowledge, and (3) the extent to which the subtask presupposes other subtasks. The first subtask is to place unibus modules into backplanes; the second is to assign massbus devices to massbuses.

**Example: Placing unibus modules in backplanes.** Whenever more than one unibus option is ordered for a VAX, it is necessary to place the modules on the unibus in an acceptable sequence. It is straightforward to determine the optimal sequence for the modules; the modules are sorted on the basis of their interrupt priority and within that on the basis of their transfer rate. Before a module can be placed on the unibus, it is necessary to select a backplane. Several constraints come into play. Backplanes come in two sizes (4-slot and 9-slot) and can have any of several pinning types. The backplane selected must be of the pinning type required by the unibus module. To determine the size of the backplane to be selected, it is necessary, first, to determine whether the size is constrained by the box that the backplane will be placed in. A box can accommodate five 4-slot backplanes. In most cases a 9-slot backplane may be used in place of two 4-slot backplanes; the exception is that a 9-slot backplane may not occupy the space reserved for the second and third 4-slot backplanes.<sup>4</sup> Assuming that either a 4-slot or a 9-slot backplane would be acceptable, the next constraint to come into play is that a 9-slot backplane should not be selected unless the next N modules in the optimal sequence all require a backplane of the same type and will not all fit in a 4-slot backplane. Once a backplane is

---

<sup>4</sup>The box that contains unibus modules has two +5 volt regulators. One of these regulators supplies power to the first two 4-slot backplanes (or to the first 9-slot backplane); the second supplies power to the other backplanes. All of the modules in a backplane must draw power from the same regulator.

selected, the board or boards comprising the next module in the optimal sequence can be placed in the backplane. However, the first and last slots of a backplane cannot accommodate a full width board; thus the configurer must make sure that the boards will physically fit into the backplane. There are several other constraints. For example, the total amount of power that can be drawn from a regulator is limited; also, if the length of the unibus exceeds a certain limit or if the load on the unibus exceeds a certain limit, a backplane containing a unibus repeater must be placed in the box.

After a module has been placed in a backplane, there is frequently room for additional modules, but the next module in the optimal sequence may require a backplane of a different type or more space than is available in the backplane. At that point the configurer must decide whether to deviate from the optimal sequence or to leave some of the backplane slots empty; the decision is based on the total amount of box space available and the seriousness of the deviation. If there is sufficient box space to accommodate the modules when they are in the optimal sequence, the optimal sequence should be preserved (even if this entails adding additional backplanes to the order). If there is not sufficient space, the seriousness of the deviation must be determined; there are some less than optimal sequences that are acceptable. If the decision is that the deviation from the optimal would impair the functionality of the system, then the configurer must add another box (and possibly a cabinet as well) to the order; if the decision is that an acceptable suboptimal sequence can be found, then some module other than the next one in the optimal sequence is placed in the backplane. In general, it is acceptable to add a module that is not next in the optimal sequence if it meets all of the constraints mentioned above and has a transfer rate that is lower than that of any other module with the same interrupt priority that it will precede.

There are reasons, other than lack of slot space or power, why the configurer must consider deviating from the optimal sequence. If the module that is to be added is a multiplexer, for example, then in addition to the space and power constraints, there is the added constraint that sufficient panel space must be available in the cabinet containing the box that will contain the multiplexer. If there is not enough panel space in that cabinet, the configurer must decide whether to deviate from the optimal ordering or to put all of the remaining modules in the remaining cabinets. Again, the decision must be made on the basis of the total space available and the seriousness of the deviation. If there are no other cabinets on the order, one must be added to the order.

This description of how unibus modules are configured brings to light most of the constraints relevant to that task for a simple, single unibus system. But it does not make very clear what the demands for component information are or the extent to which the task presupposes other tasks. The component information that the rules require is, for a module, the module type, the number and size of each board in the module, transfer rate and interrupt priority, the pinning type required, the power drawn by the module, and the load it puts on the unibus. For a backplane, the information required is

the pinning type, the size, the power drawn by the modules that have been placed in it, and the slots still available. For a box, the information is the box type, the amount of backplane space still available in the box, and the length of and load on its unibus. For a cabinet the information is the cabinet type and amount of box and panel space still available in the cabinet. Some of the tasks that this task presupposes were mentioned: determining the optimal sequence, selecting a backplane, assigning the backplane to a box, selecting some module other than the next one in the optimal sequence, verifying (when the module is a multiplexer) that there is sufficient panel space in the cabinet. There are several others: the unibus adaptors have to have been configured (so unibus length can be computed); boxes have to have been assigned to cabinets and also to unibuses (so unibus length can be computed and so the amount of usable panel space will be known); the unibus cables that connect backplanes must be selected (so unibus length can be computed); before a module that is a laboratory peripheral can be put in a backplane, it must be determined that there is sufficient panel space in the cabinet and sufficient space and power in the box for the backplane that will contain the laboratory peripheral options.

**Example: Assigning massbus devices to massbuses.** Whenever an order contains massbus devices, those devices must be assigned to massbuses. There are a number of rules that constrain how these assignments can be made: Up to eight disk drives and master tape drives can be assigned to each massbus. *If there are enough massbuses so that disk drives and tape drives can be assigned to separate massbuses, that should be done.* If not, the order of assignment should be such that the disk drives precede the tape drives on the massbus containing both types of devices (even though in the spatial layout, some of the tape drives will be closer to the cpu cabinet -- and thus to the massbus adaptor -- than will the disk drives). Unless there are more massbus adaptors than massbus devices, each massbus should be assigned at least one massbus device. Disk drives are either single port or dual port; each dual port disk drive must be assigned to two massbuses. Dual port disk drives should precede single port disk drives on the massbus. Up to seven slave tape drives can be assigned to each master tape drive. If the ratio of slave tape drives to master tape drives on an order is greater than 7 to 1, a formatter must be added to one of the slave tape drives to make it a master.

If the number of massbus adaptors on the order is too few to accommodate the disk drives and tape drives, then additional massbus adaptors must be added to the order. To determine whether all of the massbus devices can be configured, it is necessary to determine whether the number of adaptors exceeds the number permitted. Up to four massbus adaptors are permitted. If the total number of unibus adaptors and massbus adaptors on the order is greater than three, a cpu expansion cabinet is required. If the total number of unibus adaptors and massbus adaptors is greater than seven and there is only one memory controller on the order, a second cpu expansion cabinet is required. If the total number of unibus adaptors and massbus adaptors is greater than five and there

are two memory controllers, a second cpu expansion cabinet is required.

The component information required by the rules is the type of each massbus device (disk drive, master tape drive, or slave tape drive), the number of ports (for disk drives), the number of devices assigned to each massbus and the type of each device assigned, the number of slave tape drives assigned to each master, and the space available for massbus adaptors in the cpu and cpu expansion cabinets. The task of assigning slaves to masters must be done before either disk drives or master tape drives are assigned so that the number of massbus adaptors required and the distribution of devices among massbuses can be determined. The tasks of assigning memory and unibus adaptors to the cpu and cpu expansion cabinets must be done before the massbus devices are assigned in order to determine whether the number of massbus devices exceeds the maximum permitted and to determine whether additional cabinet space will be required.

## 2. THE SYSTEM

This section focuses on how to represent the knowledge required for the VAX configuration task so that the resulting system can perform the task expertly and efficiently and can easily acquire additional knowledge about the domain. The architecture in which R1 is embedded is described. Issues of search are discussed. The content and use of R1's knowledge is analyzed. Finally, some design and implementation history is provided in order to show the extent to which the development of R1 was an evolutionary process.

### 2.1. THE PRODUCTION SYSTEM ARCHITECTURE

R1 is implemented as a production system [Newell 77]. The particular production system language used is OPS4. Since detailed descriptions of this language have been provided elsewhere [Forgy 79, Forgy 77, McDermott 78], only a brief indication of the basic features of the language will be given in this paper. An OPS4 production system consists of a set of productions held in *production memory* and a set of data elements (eg, state descriptions) held in *working memory*. A *production* is a conditional statement composed of conditions and actions; a production has the form:

$$P_i (C_1 C_2 \dots C_n \rightarrow A_1 A_2 \dots A_m)$$

Actions typically modify working memory by deleting, adding, or modifying a data element; users may, however, define application specific actions. Conditions are templates; when each of the conditions in a production can be matched by an element in working memory, the production is said to be instantiated. An *instantiation* is an ordered pair of a production and a set of elements from working memory that satisfy the conditions of the production. The OPS4 interpreter operates within a control framework called the *recognize-act* cycle. During the recognition part of the cycle, it finds the instantiation to be executed; during the act part, it performs the actions. The recognize-act cycle is



repeated until either no production can be instantiated or an action explicitly stops the processing. Recognition can be divided into *match* and *conflict resolution*. During match, the interpreter finds the set of all instantiations of productions that are satisfied on the current cycle. During conflict resolution, it determines which instantiation to execute.

Each of R1's productions (rules) embodies a piece of constraint knowledge. The production's conditions typically look for situations in which a particular type of extension to a particular type of partial configuration is permissible or required; the actions then effect that extension. R1's rules are such that on almost every cycle several rules can be instantiated -- often in a number of different ways. From R1's point of view, it often makes no difference which of these instantiations is executed; in these cases, how OPS4 determines which instantiation to execute is irrelevant.<sup>5</sup> R1 does, however, rely heavily on one of OPS4's conflict resolution strategies, the *special case strategy*, and so this strategy needs to be understood. Given two instantiations, one of which contains a proper superset of the data elements contained by the other, OPS4 will select the instantiation containing more data elements on the assumption that it is specialized for the particular situation at hand. OPS4's cycle time, though it is essentially independent of the size of both production memory and working memory [Forgy 80], depends on particular features of the production system (eg, the number and complexity of the conditions and actions in each production). The average cycle time for OPS4 interpreting R1 is about 150 milliseconds.<sup>6</sup>

As we saw in the previous section, there is a considerable amount of information which has to be known about each component that can appear on an order. To provide R1 with access to this information, OPS4's two memories have been augmented, for this application, with a third memory. This memory, the *data base*, contains descriptions of each of the 420 components currently supported for the VAX. Each entry in the data base consists of the name of a component and a set of attribute/value pairs that indicate the properties of that component that are relevant for the configuration task. Every component has a *type* attribute and a *class* attribute; the class of a component determines what other attributes are relevant. There are 15 classes: bundle, cabinet, sbi module, sbi device, box, backplane, massbus device, unibus device, unibus module, panel, power supply, software, cable, document, and accessory. Each component description consists, on the average, of eight attribute/value pairs; there are only about 50 distinct attributes, several of which are common to all (or most) of the classes. Figure 2-1 shows five of the entries in the data base. The RK711-EA is a bundle of components; it contains a 25 foot cable (70-12292-25), a disk drive

---

<sup>5</sup>For example, a rule that bears on configuring some particular type of component will have more than one instantiation if more than one such component is available; any of these instantiations could be executed. Or if R1 is filling a cabinet, it might make no difference which part of the cabinet is filled first.

<sup>6</sup>OPS4 is implemented in MACLISP; R1 is run on a PDP-10 (model KL) and loads in 412 pages of core.

**RK711-EA**  
 CLASS: BUNDLE  
 TYPE: DISK DRIVE  
 SUPPORTED: YES  
 COMPONENT LIST: 1 070-12292-26  
                   1 RK07-EA\*  
                   1 RK611

**RK07-EA\***  
 CLASS: UNIBUS DEVICE  
 TYPE: DISK DRIVE  
 SUPPORTED: YES  
 FLOOR RANK: 8  
 DEPTH: 28 INCHES  
 WIDTH: 24 INCHES  
 HEIGHT: 42 INCHES  
 UNIBUS MODULE REQUIRED: RK611\*  
 PORTS: 1  
 VOLTAGE: 120 VOLTS  
 FREQUENCY: 60 HERTZ  
 CABLE TYPE REQUIRED: 1 070-12292 FROM A DISK DRIVE UNIBUS MODULE  
                           OR 1 070-12292 FROM A DISK DRIVE UNIBUS DEVICE

**RK611**  
 CLASS: BUNDLE  
 TYPE: DISK DRIVE  
 SUPPORTED: YES  
 COMPONENT LIST: 3 G727  
                   1 M9202  
                   1 070-12412-00  
                   1 RK611\*

**070-12412-00**  
 CLASS: BACKPLANE  
 TYPE: RK611  
 SUPPORTED: YES  
 NUMBER OF SYSTEM UNITS: 2  
 LENGTH: 2.0 FEET  
 NUMBER OF SLOTS: 9  
 SLOT TYPES: 3 SPC (1 TO 3)  
                   6 RK611 (4 TO 9)

**RK611\***  
 CLASS: UNIBUS MODULE  
 TYPE: DISK DRIVE  
 SUPPORTED: YES  
 PRIORITY LEVEL: BUFFERED NPR  
 TRANSFER RATE: 212  
 NUMBER OF SYSTEM UNITS: 2  
 SLOTS REQUIRED: 6 RK611 (4 TO 9)  
 BOARD LIST: (HEX A M7904) (HEX A M7903) (HEX A M7902) (HEX A M7901) (HEX A M7900)  
 DC POWER DRAWN: 15.0 .175 .4  
 UNIBUS LOAD: 1  
 NUMBER OF UNIBUS DEVICES SUPPORTED: 8  
 CABLE TYPE REQUIRED: 1 070-12292 FROM A DISK DRIVE UNIBUS DEVICE

Figure 2-1: Some representative items from the data base

(RK07-EA\*), and a bundle of components (RK611) which itself consists of three continuity boards (G727), a unibus jumper cable (M9202), a backplane (70-12412-00), and a disk drive controller (RK611\*). The RK07-EA\* is a single port disk drive; it is a unibus device that requires an RK611\* as its controller, and it is connected to that controller either directly or via another disk drive with a 70-12292 cable of some length. The 70-12412-00 is a 9-slot backplane whose "electrical length" is two feet; its first three slots are for boards that require spc (small peripheral controller) pinning, and

the remaining six slots are for the RK611\*. The RK611\* is a disk drive controller which, because of its interrupt priority and transfer rate, is typically located toward the front of the unibus. The module is comprised of five hex boards each of which start in lateral position "A"; it draws 15.0 amps of +5 volt current, .175 amps of -15 volt current, and .4 amps of +15 volt current, and it generates 1 unibus load. It can support up to eight disk drives and is connected to the first of these with a 70-12292 cable of some length.

In addition to containing descriptions of VAX components, the data base also contains a few cabinet templates. A *cabinet template* describes what space is available in a particular cabinet type. These templates serve two purposes: (1) they enable R1 to know, at any point in the configuration process, what container space is still available, and (2) they enable R1 to assign a specific location (ie, coordinates) to each component that it places in a cabinet. Figure 2-2 shows the templates for the cpu and unibus expansion cabinets. The components that may be ordered for the cpu cabinet are sbi modules, power supplies, and an sbi device. The template for the cpu cabinet contains descriptions of the space available for each of these classes of components and specifies what can be put where. For example, up to six sbi modules fit into a cpu cabinet; each cabinet contains a cpu module and some memory; in addition there are three "slots" for options that occupy 4 inches of space and one slot for an option that occupies 3 inches of space. The description "cpu nexus-2 (3 5 23 30)" indicates that the cpu module must be associated with nexus 2 of the sbi; the numbers in parentheses indicate the top left and bottom right coordinates of the space that can be occupied by a cpu module. The components that may be ordered for the unibus expansion cabinet are boxes and panels. Note that multiplexer panels and (half-size) laboratory peripheral panels occupy the same space; one piece of R1's constraint knowledge is that two panels cannot occupy the same space at the same time.

Initially, working memory is empty. It grows, during the course of configuring a system, to contain descriptions of the components ordered, and as various components are associated, to contain descriptions of partial configurations as well as other component information required to do the configuration task. A component is represented in working memory as a *component-token* with associated attribute/value pairs. R1 retrieves information from the data base as the need for such information arises. There are five actions that R1 can perform that provide it access to the data base. Three of these functions, *generate-tokens*, *find-token*, and *find-substitute-token*, retrieve specified information about a component (or list of components) from the data base, create a component-token, and then add a partial description of the component to working memory. The other two, *get-attributes* and *get-template*, augment the description of an already existing component-token. *Generate-tokens* takes a list of component names and a set of attribute names as its arguments, and returns, for each component on the list, a component-token and the value of each of the specified attributes. *Find-token* takes a partial description (a set of attribute/value pairs) and a

**CPU-CABINET**

```

CLASS: CABINET
HEIGHT: 60 INCHES
WIDTH: 52 INCHES
DEPTH: 30 INCHES
SBI MODULE SPACE: CPU NEXUS-2 (3 5 23 30)
                   4-INCH-OPTION-SLOT 1 NEXUS-3 (23 5 27 30)
                   MEMORY NEXUS-4 (27 5 38 30)
                   4-INCH-OPTION-SLOT 2 NEXUS-5 (38 5 42 30)
                   4-INCH-OPTION-SLOT 3 NEXUS-5 (42 5 46 30)
                   3-INCH-OPTION-SLOT NEXUS-6 (46 5 49 30)
POWER SUPPLY SPACE: FPA NEXUS-1 (2 32 10 40)
                   CPU NEXUS-2 (10 32 18 40)
                   4-INCH-OPTION-SLOT 1 NEXUS-3 (18 32 26 40)
                   MEMORY NEXUS-4 (26 32 34 40)
                   4-INCH-OPTION-SLOT 2 NEXUS-5 (34 32 42 40)
                   CLOCK-BATTERY (2 49 26 62)
                   MEMORY-BATTERY (2 46 26 49)
SBI DEVICE SPACE: IO (2 52 50 56)

```

**UBX-CABINET**

```

CLASS: CABINET
HEIGHT: 60 INCHES
WIDTH: 28 INCHES
DEPTH: 30 INCHES
BOX SPACE: UBX 1 (2 39 26 48)
           UBX 2 (2 13 26 22)
PANEL SPACE: MUX 1 1 BACK (14 27 26 36)
             MUX 2 2 BACK (2 2 26 11)
             MUX 3 3 FRONT (2 27 14 36)
             LPA 1 1 BACK (14 27 26 31)
             LPA 1 2 BACK (14 32 26 36)
             LPA 2 3 BACK (2 2 26 6)
             LPA 2 4 BACK (2 7 26 11)
             LPA 3 5 FRONT (2 27 14 31)
             LPA 3 6 FRONT (2 32 14 36)

```

Figure 2-2: Two sample templates

set of attributes as its arguments, finds a component in the data base matching that partial description, and returns a component-token and the value of each of the specified attributes. Find-substitute-token takes a component name, a partial description, an exception list, and a set of attribute names as its arguments, finds a component in the data base that is like the original component except that it satisfies the partial description and may differ with respect to the attributes on the exception list, and returns a new component-token and the values of the specified attributes. Get-attributes takes a component, a component-token, and a set of attribute names as its arguments and returns the values of the specified attributes. Get-template takes a template name and a component-token as its arguments and returns the attribute/value pairs of that template.

In addition to containing component descriptions, working memory contains three other types of elements:

- Elements that define partial configurations.
- Elements that indicate the results of various sorts of computations.

- Context symbols.

An element that defines a partial configuration contains a description of the relationships among two or more components. Typically, these elements indicate either that one component is to be connected to another by means of a cable or, in the case of a component that is a container, the spatial relationship between the container and each of the components it contains. An element that indicates the result of some computation contains a symbol identifying the computation and one or more values indicating the result. The component descriptions, together with the elements that define partial configurations and the elements that indicate the results of various computations, constitute the component information. A context symbol contains a context (subtask) name and an indication of whether or not the context is active.

Production memory contains constraint knowledge -- all of R1's permanent knowledge about how to configure VAX systems. R1 currently has 772 rules that enable it to perform the task. An English translation of a sample rule is shown in Figure 2-3. The first condition indicates that the context in which this rule is relevant is the distributing of massbus devices among massbuses. The other five conditions specify one of the sets of constraints that must be satisfied within this context in order for a disk drive to be assigned to a massbus. When an instantiation of this rule is executed, one of the single port disk drives on the order is assigned to one of the massbuses.

**DISTRIBUTE-MB-DEVICES-3**

**IF: THE MOST CURRENT ACTIVE CONTEXT IS DISTRIBUTING MASSBUS DEVICES  
 AND THERE IS A SINGLE PORT DISK DRIVE  
       THAT HAS NOT BEEN ASSIGNED TO A MASSBUS  
 AND THERE ARE NO UNASSIGNED DUAL PORT DISK DRIVES  
 AND THE NUMBER OF DEVICES THAT EACH MASSBUS SHOULD SUPPORT IS KNOWN  
 AND THERE IS A MASSBUS THAT HAS BEEN ASSIGNED AT LEAST ONE DISK DRIVE  
       AND THAT SHOULD SUPPORT ADDITIONAL DISK DRIVES  
 AND THE TYPE OF CABLE NEEDED TO CONNECT THE DISK DRIVE  
       TO THE PREVIOUS DEVICE ON THE MASSBUS IS KNOWN**

**THEN: ASSIGN THE DISK DRIVE TO THE MASSBUS**

**Figure 2-3: A sample rule**

## 2.2. CONTEXTS

The configuration task can be viewed as a hierarchy of subtasks that have strong temporal interdependencies. The actions required within each subtask are highly variable; they depend completely on the particular combination of components that appear on an order and the way in which sets of those components have been configured up to the point when the new subtask becomes appropriate. It is possible, however, to indicate what actions (including the action of initiating a new subtask) are appropriate within a subtask in terms of a relatively small number of rules that are each sensitive to a few salient features of the current situation.

R1's approach to exploiting the temporal relationships among subtasks is straightforward. The function of several of R1's rules is to recognize, on the basis of the component information in working memory, when a new subtask should be initiated. When one of these rules fires, it adds a context symbol to working memory. Each context symbol contains a context name, an indication of whether the context is active or not-active, and an indication of when (ie, how recently) the context was made active. Each rule contains two condition elements that are sensitive to context symbols. Together these conditions insure that only those rules that bear on the most current active context will fire. Which of these rules fire depends on what descriptions of components, partial configurations, and computation results are in working memory. Typically, a few of the rules associated with a context contain the constraint knowledge ordinarily relevant within that context. Other rules associated with the context are special case rules; these rules, when their more stringent conditions are satisfied, fire before (and often instead of) the ordinary case rules. Each context has an associated rule containing only the two condition elements sensitive to a context symbol. This rule deactivates the current context. Since each deactivation rule is a general case of all the other rules sensitive to its context, it fires only after the other satisfied rules have fired. R1, then, is a recognition-driven system that relies on its knowledge of the structure of the configuration task as well as on information about the set of components it is configuring to determine what to do. When it has several courses of action open to it, it falls back on general (non-task-specific) strategies for selecting among alternatives. But it needs only two such strategies: (1) it uses special case rules in preference to more general ones, and (2) it does all that it can within a context before leaving that context.

The contexts that R1 makes use of in order to do the configuration task were not arrived at through a rigorous analysis of the demands of the configuration task. Rather, they reflect the way in which experts who do the task actually approach it. As might be expected, as R1 evolved it became apparent that some modifications and reordering of the contexts would make the processing easier, and so these changes were made. But basically, the approach that R1 takes to the task is the same as that of humans who do it. At the top level, the task divides up into 6 major subtasks:

1. Determine whether there is anything grossly wrong with the order (eg, mismatched items, major pre-requisites missing).
2. Put the appropriate components in the cpu and cpu expansion cabinets.
3. Put boxes in the unibus expansion cabinet and put the appropriate components in those boxes.
4. Put panels in the unibus expansion cabinets.
5. Lay out the system on the floor.
6. Do the cabling.

The following paragraphs give a rough picture of what each of these subtasks involves; the purpose of this excursion into the guts of the configuration task is to provide some concreteness on which to build the subsequent discussions of the effectiveness (and adequacy) of R1's problem solving method and of the use that it makes of its configuration knowledge.

**Subtask # 1** (196 rules). The first subtask is to determine whether there are major problems with the order and to rectify them if possible. The work that R1 does during this stage is considerably more complex than that done by human experts during this stage. Humans tend to assume that an order will be unproblematic and wait until a problem actually arises before dealing with it. The advantage of the human approach is that it saves an unnecessary first step when configuring unproblematic orders; its disadvantage is that if problems do arise, they may impact earlier decisions and thus require redoing part of the configuration. R1 first retrieves partial descriptions of each of the components on the order. If a component is a bundle, it retrieves a partial description of each of the components on that bundle's component list. Then R1 checks to see whether it has been instructed to treat any of the components in an exceptional way -- eg, to leave some components (presumably spares) unconfigured. It next determines whether all of the components on the order have compatible voltage and frequency requirements, and if not, substitutes components of the appropriate voltage and frequency. Finally R1 determines, to the extent that it can before actually doing the configuration, whether any of the massbus or unibus devices on the order have pre-requisite components that are not on the order. It makes sure that there is at least one master tape drive for every seven slaves, that there are enough adaptors for the massbus devices, and that there is enough cabinet space for the adaptors. It makes sure that there is at least one unibus adaptor and that there are enough controllers for the unibus devices. It also makes sure that there are enough controllers and enough cabinet space for the memory ordered. If any pre-requisite components are missing, it adds them to the order.

**Subtask # 2** (87 rules). The second subtask involves putting whatever components belong in the cpu and cpu expansion cabinet into those cabinets; in performing this subtask, R1 relies heavily on the templates for the cabinets. It augments the description of whatever cpu cabinet is on the order with the information in the template for the cpu cabinet. R1 then finds a component in working memory whose class is sbi module and whose type is the type of one of the template elements for the sbi. It adds the component (and the coordinates it will occupy) to the list of components that are to be put in the cpu cabinet and deletes the template element. It repeats this step until no template elements remain that can be paired with a component on the order. There are, of course, a number of decisions that have to get made along the way. When putting in the cpu module, R1 must check to see if either the floating point accelerator or writeable control store options are on the order, and if so, put the boards for these options in the appropriate place in the cpu backplane. When putting in the

memory, R1 must know how many memory controllers and how many adaptors have been ordered in order to know whether or not to interleave the memory. As it puts in the massbus adaptors, R1 must determine the relationships of the massbus devices to one another so that it will know which devices go with which adaptors; this task is non-trivial since there are a number of rules having to do with how to distribute massbus devices. As R1 puts modules on the sbi, it must also put the appropriate power supplies into the cabinet; in order to do this, it needs to know what regulators are required for what modules. If there are cpu expansion cabinets on the order, R1 fills them using the same rules that it uses for the cpu cabinet; the only difference between the two types of cabinets is that the expansion cabinet has no cpu, and may or may not contain memory. When all of the sbi modules have been placed in cabinets, R1 puts the sbi terminator in the appropriate place in the final cabinet and adds module simulators to any option slots that have not been filled.

**Subtask #3** (256 rules). The third subtask is to put boxes into the unibus expansion cabinets, and to put unibus modules into the boxes. Given a limited amount of box space, the information that is needed to determine whether a module has been configured acceptably is not available until after all of the modules have been configured. Thus R1 sometimes has to generate a number of candidate configurations before it finds one that is acceptable. There are a number of independent constraints on the placing of unibus modules:

1. Each module must be put in a backplane slot of the appropriate pinning type.
2. The position of each backplane in a box must be such that its modules draw power from a single set of regulators.
3. There is a limit on the amount of power that the modules in a backplane can collectively draw from any regulator.
4. If a module requires panel space, that panel space must be in the cabinet containing the module.
5. If a module requires other supporting modules either in the same backplane or in the same box, space must be available for those supporting modules.
6. The modules should be placed on the unibus in a sequence that is as close to the optimal sequence as possible.

If only the first five constraints applied, R1 would generate only acceptable configurations; but the addition of the sixth constraint, since it is elastic, makes that impossible. In order to limit the amount of search it has to do to configure the unibus modules, R1 interprets the sixth constraint somewhat liberally. It defines three equivalence classes of sequences: optimal (any ordering that is optimal), almost-optimal (any less than optimal ordering such that no module whose interrupt priority is  $i$  and whose transfer rate is  $j$ , occurs before a module whose interrupt priority is  $i$  and whose transfer rate is less than  $j$ ), and suboptimal (any other ordering).



To configure a set of unibus modules, R1 first estimates the amount of space required to configure the unibus modules optimally and the amount of space required to configure the modules suboptimally; it then determines the optimal sequence. If the amount of box space available is greater than or equal to the space required for an optimal configuration, it tries to place the modules on the unibus in that sequence. If it fails (or if the amount of box space available is less than the space required for an optimal configuration, but greater than that required for a suboptimal configuration), it retries the subtask, modifying the sequence whenever such a modification would save space and result in an almost-optimal sequence. If this attempt fails, it retries the subtask again, but this time modifies the sequence whenever such a modification would save space. If this attempt fails or if the amount of box space available is less than that required for a suboptimal configuration, R1 adds another box to the order and retries the subtask.

With this background, the subtask of configuring unibus cabinets, boxes and modules can now be described. R1's first step is to assign each box to a unibus and each box to a cabinet. To determine whether a particular unibus module can be placed in some backplane, the distance from the beginning of the unibus (the adapter) to that backplane must be known. For R1 to have this information, the box assignments have to have been made. Once R1 has assigned each box to a cabinet, it starts to fill the first box on each unibus. Filling a box involves first selecting a module (and a box to put it in if the system has more than one unibus). If R1 is attempting to configure the modules in a way that will preserve the optimal sequence, module selection is straightforward. If R1 is willing to accept an almost-optimal or a suboptimal configuration, it selects the next module in the optimal sequence unless that module requires a backplane that will occupy more space than remains in the box; if there is insufficient space for such a backplane, R1 considers each of the remaining modules to see if there is one that will satisfy the relaxed sequencing constraint and requires a backplane that will fit in the box. After selecting a module, R1 selects (or adds to the order) a backplane that has pinning and a size that will accommodate the module; if there are two such backplanes on the order, one of which is a 4-slot backplane and the other a 9-slot, R1 selects the 4-slot backplane unless the next N modules in the optimal sequence would not all fit in the 4-slot backplane. After selecting a backplane, R1 assigns it to a box, and then generates elements that serve as a backplane template.

At this point, R1 attempts to put the module in the backplane. Although it knows that there is sufficient space for the module in the backplane, it does not know whether the other constraints are satisfied. If adding this module would result in the power-drawn limit being exceeded, or if the unibus-load limit would be exceeded, or if required panel space is not available, or if there is not room in the backplane or box for supporting modules if any are required, then the attempt to add the module will fail. Here again if R1 is willing to accept a less than optimal configuration, it will try a different module. As long as space remains in the backplane and there are modules that have not

been configured, R1 will try to put additional modules in the backplane. The module selection process is identical to the one described above, except that a backplane is already available. When R1 has put all of the modules that it can into a backplane, there are a number of things it must do before it is finished with that backplane. Each of the unfilled spc slots in the backplane must be filled with a continuity card. Then R1 must note that the length of the unibus has increased by the length of the backplane, and must check to see whether (either because of the unibus length or because of the unibus load) it needs to put a unibus repeater in the next position on the unibus. R1 then adds a two foot jumper cable to connect the backplane it just filled to the next (as yet unspecified) backplane. If one of the modules in the backplane has an associated lpa-backplane, it adds the lpa-backplane to the box.

When R1 has finished filling a box, it checks to see if there are more modules or backplanes that need to be put in a box; if there are, it replaces the two foot jumper cable that it has associated with the last backplane in the box with a longer jumper cable (the length being determined by the distance between the box just filled and the next box to be filled); if there are no unconfigured modules or backplanes, R1 replaces the two foot jumper cable with a unibus terminator. If R1 is able to configure all of the unibus modules in the available box space, it goes on to the fourth subtask. If it cannot, it either relaxes the sequencing constraint or adds another box to the order and tries again.

**Subtask # 4 (30 rules).** The fourth subtask is to assign panels to cabinets and to associate those panels with unibus modules and with whatever devices the modules serve. Since by this point all of the unibus modules have been configured and since panel space has been assigned to each module requiring it, all that R1 has to do is select a panel of the appropriate size and line type (or add one to the order) and assign that panel to the panel space set aside for the module. Then R1 assigns those devices that must connect to a module via a panel to a panel of the appropriate type; if more than one such panel is available, R1 distributes the assignments evenly across panels.

**Subtask # 5 (61 rules).** The fifth subtask is to generate a floor layout for the system (ie, to specify the spatial relationships among all of the cabinets and free standing devices). To do this subtask adequately, it is necessary to have information about the site at which the system will be installed (eg, room dimensions, locations of obstructions) that is not currently available to R1. At the moment, therefore, all that R1 does is group together those components that must be spatially proximate and then lay the devices out in a line in the appropriate order. This subtask is relatively easy since all of the inter-device assignments have been made and information about how close a particular device should be to the cpu cabinet relative to other devices is in R1's data base. The only additional knowledge that R1 needs for this subtask is the knowledge of when the spatial ordering of devices should be different from their electrical ordering. It needs to know, for example, that though tape drives always come after disk drives when both are assigned to the same massbus, tape drives are

sometimes placed nearer to the cpu cabinet (ie, are bolted to the last unibus expansion cabinet).

**Subtask # 6** (36 rules). The final subtask is to specify what cables are to be used to connect each device to the other devices to which it has been assigned. Given the inter-device assignments, all that R1 has to do is determine the distance between each pair of devices that must be connected and find (or add) a cable of the type and length required. Because R1 does not currently lay out systems in the way they will be laid out at the installation site, it does not do an adequate job of determining the precise cable lengths required. However, since much of the cabling is between devices that have a site-independent spatial relationship, the job that R1 does is more helpful than might first appear. All of the cabinets must be bolted together, and devices associated with the same controller must be physically proximate. Thus, the only cable lengths that R1 cannot determine are those involving cables that connect a controller to the first device it serves. In order to make these cabling assignments, R1 selects the controller/device pair that is farthest apart in the linear layout generated during the fifth subtask and selects for that pair the longest cable on the order that is of the appropriate type. Once R1 has finished the cabling task, the system is configured; at that point, R1 generates output describing the configuration.

### 2.3. SEARCHING THE SPACE OF POSSIBLE CONFIGURATIONS

The configuration task performed by R1 requires finding an acceptable configuration within the space of possible configurations. The basic operations that R1 uses to explore this space are creating and extending partial configurations. Other of R1's operations, such as retrieving component information from the data base or counting the number of components of a particular type, prepare the way for the actual acts of configuration. It would appear that performing any task for which there are many constraints on an acceptable solution requires a heuristic search (ie, a combinatorial search in which candidate partial solutions are constructed and their potential evaluated). This has been the experience of AI in all sorts of tasks [Nilsson 80] and in particular in real-world domains where the methods used have been characterized as forms of Generate-and-test [Feigenbaum 77]. R1, however, by and large does without search. It is useful to understand just what the structure of R1's method is and what permits search to be avoided.

R1's method is essentially a generalized form of matching; that is, it is the method normally used to match a form (ie, a symbolic expression containing variables) against an exemplar in order to instantiate the variables, thus making the form identical to the exemplar. Matching is usually not considered to be search. In typical AI heuristic search programs, it is taken to be a part of the computation performed within a state to determine which operators are applicable. However, Match is clearly also a search technique, analagous to Generate-and-test, Means-ends Analysis, etc. The Match method can be found in some of the earliest heuristic search programs [Newell 63], and later

Newell included a generalized form of Match as one of his *weak methods* [Newell 69].

The search space for Match is the space of all instantiations of the variables in a form. Each state in the space is a partially instantiated form. The form is the body of domain specific knowledge that defines acceptable sets of instantiations; thus the form holds the constraints for task satisfaction. This knowledge serves two functions: (1) it enables the form to be brought into correspondence with the exemplar, and (2) it permits local tests (comparisons) to be made between the form and the exemplar at each point of correspondence. The primary condition on Match, then, is

- *The Correspondence Condition:* Match can be successful only if each constituent of the form can take on a locally determined value (through the comparison of the form to the exemplar at the point of correspondence).

A consequence of the correspondence condition is that Match never requires backtracking. If there is at least one solution state, Match will find a path from the initial state to a solution state without generating any false paths. This condition abstracts from all of the details of how expressions are put in correspondence. It states only the essential condition that permits the operators to be selected and applied in a single pass.

The decision that determines the next step in the space (ie, the next instantiation) is local. This does not imply, however, that the space is decomposable into a set of completely independent subtasks. Instantiating a variable can have global consequences for the final solution and hence for subsequent instantiations (because the substitution must take place in all occurrences). However, none of these consequences affect what has already been matched. Indeed, the result of putting the form into correspondence with the exemplar is to impose an order on the instantiations (ie, the search through the space) so that all global consequences are pushed into the unmatched portion of the form. In typical string matching this requires only that matching start from one end of the string or the other. This can be stated as another general condition:

- *The Propagation Condition:* A partial ordering on decisions must exist such that the consequences of applying an operator bear only on aspects of the solution that have not yet been determined.

Applying this analysis of Match to R1, the initial state is the set of descriptions of the components ordered. The intermediate states are sets of descriptions of partial configurations and the as yet unconfigured components. At each decision point, the constraint knowledge about what next step can be taken (ie, what next partial configuration to produce) is provided by R1's rules. R1 has enough knowledge to satisfy the correspondence condition stated above. Thus, there is no need for backtracking, since its knowledge is sufficient to determine an acceptable next step.

R1's rules can be divided into three categories based on their role in this generalized Match

method. (1) *Operator* rules take the actual next step in creating or extending a partial configuration. (2) *Sequencing* rules determine the order in which decisions need to be made to satisfy the propagation condition. These rules are primarily those involved in the sequencing of contexts, and this sequencing is thus seen to be an essential part of R1's method, not simply an additional bit of "programming". (3) *Information-gathering* rules access the data base or perform various computations in order to provide the information needed for operator and sequencing rule selection. This information has an important subsequent use that has no analogue in the simple Match task. It provides a justification for decisions (and their consequences) that is necessary in order to preserve the propagation condition.

The power and limitations of Match are exemplified by two interesting phenomena. The first one is that humans often do not solve the configuration task by Match, but rather by a more general (weaker) heuristic search; that is, they engage in backtracking. This is true of novices and to a lesser extent also true of experienced configurers. Novices in particular frequently make decisions in the wrong order or neglect to preserve information needed for future decisions. The second phenomenon is that, for R1, Match is not in fact sufficient for the complete task. The subtask of placing modules on the unibus is formulated essentially as a bin-packing problem -- namely how to find an optimal sequence that fits within spatial and power-load constraints. No way of solving this problem without search is known (short of table lookup, if the configurations are limited enough to be enumerated). R1 does not use a fully adequate search method for this subtask, but does a simple Generate-and-test, using Match as a heuristic guide.

The desirability of using a generalized Match method is hardly in doubt; it is good to avoid search, especially when this can be done at low cost. The two phenomena above shed some light on when it is possible to use Match. As the bin-packing subtask shows, it depends in part on the specific nature of the task environment (ie, whether its structure is sufficiently interlocking). Bin packing has been studied enough to attribute the necessity of search to the structure of the task. However, as the search behavior of humans shows, tasks almost invariably become more tractable as they become better understood. The role of experience with a task is to permit the acquisition of enough knowledge, both of decision order and of decision content, to satisfy the two conditions stated above, thus permitting Match to be used.

R1 appears to be the first of the "knowledge-engineered" domain-specific AI systems to employ Match as its central problem solving method. As mentioned above, the more typical approach has been to develop systems whose central method is Generate-and-test. Such systems contain knowledge that enables them to generate one or more hypotheses that explain (or otherwise give structure to) the phenomena proper to their domains; these hypotheses are then tested against some particular collection of empirical data. The systems differ from one another in the degree to which

they make use of the data in generating hypotheses. Some systems generate hypotheses independently of the data to be explained; others generate partial hypotheses that are then modified and extended on the basis of the data. MYCIN [Davis 77], for example, has a collection of rules that allow it to infer the likely causes of a bacterial infection in the blood. Hypothesis generation is accomplished by backward chaining of the rules. If a rule's conditions are not known to be true or false (on the basis of patient data), rules that bear on the truth or falsity of its conditions must be examined. Ultimately this process leads back to the data; the data either confirms or disconfirms the hypothesis. Though other domain-specific systems make more use of the data in selecting hypotheses to generate, they also use the data to test hypotheses.

The feature that distinguishes R1 most clearly from these other systems is that it attempts to generate only a single hypothesis -- the solution. In R1, the knowledge that other systems would use to test hypotheses is part of the generator. As we have seen, except in the case of unibus module configuration, this strategy has been successful. Clearly its success is due in part to the structure of the configuration domain. It will be interesting to see to what extent this strategy can be successfully applied in other domains.

#### 2.4. THE CONTENT OF R1'S RULES

R1's rules can be distinguished from one another in terms of the functions that they perform and the extent to which they embody domain knowledge. As Figure 2-4 shows, only 480 of R1's 772 rules contain knowledge that is directly related to the configuration task. The other 292 rules contain more general knowledge. About a third of this more general knowledge is used by R1 to generate its output; this knowledge is not used until after the configuration task has been finished. Another third consists of rules that deactivate contexts; essentially these rules contain only the knowledge that if one is in a context and there is nothing left to do, one should exit from the context.<sup>7</sup> The final third of the general knowledge is about evenly divided between rules whose function is to do various kinds of counting tasks, and rules that generate "empty" data structures for the domain knowledge rules to use. Of the rules that embody domain knowledge, about a fourth generate new contexts, another fourth deal with missing pre-requisites (mostly by adding whatever component is missing to the order), and another fourth create or extend partial configurations. The final fourth is about evenly divided between rules that retrieve partial descriptions of components from the data base, and rules that do various sorts of computations. The classification of these domain specific rules is somewhat rough since many of these rules have dual functions.

---

<sup>7</sup> These rules are, embarassingly enough, completely unnecessary since their function could (and soon will) be taken over by a single general rule.

Domain-specific rules	General rules
Context generation (96 rules)	Output generation (106 rules)
Pre-requisites (127 rules)	Context deactivation (84 rules)
Component association (156 rules)	Counting (54 rules)
Retrieval (54 rules)	Set-up (48 rules)
Computation (47 rules)	

**Figure 2-4:** The distribution of R1's knowledge

The knowledge encoded in individual rules is, of course, just the knowledge needed to make use of the component information and context symbols that can appear in working memory. Thus all of R1's constraint knowledge falls into the following four classes:

- Knowledge of the significance of the various attributes of VAX components and of how to retrieve component information from the data base.
- Knowledge of how to recognize and of how to construct partial configurations.
- Knowledge of how to make and make use of various kinds of computations.
- Knowledge of what actions are appropriate within various contexts and of what contexts to enter from other contexts.

R1 has essentially no knowledge of the defining characteristics of its contexts; it simply has names for contexts, and these names are bare symbols with no associated descriptions. It recognizes 84 context names;<sup>8</sup> each context has, on the average, about eight rules associated with it.

Table 2-1 shows the relative frequency with which the various sorts of constraint knowledge occur as conditions and actions. The first column provides this breakdown for all of R1's rules; the second

<sup>8</sup>This excludes the 18 context names recognized by the 106 rules that generate output.

	Mean number in each rule	Mean number in each domain rule
<b>Components</b>		
conditions	2.02	2.79
actions	1.24	1.82
assertions	0.61	0.97
modifications	0.30	0.43
deletions	0.33	0.42
<b>Partial configurations</b>		
conditions	1.15	1.55
actions	0.52	0.68
assertions	0.20	0.31
modifications	0.25	0.29
deletions	0.07	0.08
<b>Results of computations</b>		
conditions	1.06	1.20
actions	0.72	0.82
assertions	0.31	0.35
modifications	0.27	0.32
deletions	0.14	0.15
<b>Contexts</b>		
conditions	2.04	2.07
actions	0.35	0.32
assertions	0.19	0.28
modifications	0.16	0.04
deletions	0.00	0.00
<b>Total</b>		
conditions	6.27	7.62
actions	2.83	3.64
assertions	1.30	1.91
modifications	0.98	1.09
deletions	0.55	0.64

**Table 2-1: The composition of R1's knowledge**

column provides the breakdown for just the domain-specific rules. In order for the numbers in Table 2-1 to have any significance, it is necessary to understand how much information each piece of knowledge contains. In general, each working memory element contains from three to six pieces of information that R1 can use. Elements of type component each contain a component-token and, on the average, two attribute/value pairs. Elements of type partial configuration almost always contain at least two component-tokens, a symbol indicating their interrelationship, and often other symbols that further specify the relationship. When the partial configuration element is a list of the contents of a cabinet or some other container, it may contain 20 or more pieces of information. But R1 typically



does not look very far inside these lists; it merely builds them up for its output routines. Elements of type *computational result* contain a symbol identifying the purpose of the computation and usually two or three values. Elements of type *context* contain a context name, an indication of whether or not the context is active, and an indication of how recently the context was asserted. The amount of information in a rule is proportional to the amount of information in these working memory elements. Conditions are patterns that are instantiated by the working memory elements; typically a pattern will contain some constants, some variables that occur only once in the condition part of the rule, and some variables that occur more than once in the condition part of the rule and thus function to constrain the combinations of elements that can instantiate the rule. Actions are either functions that retrieve information from the data base, modify or delete specified elements in working memory, or perform various arithmetic calculations, or they are patterns that are added to working memory after being instantiated with values bound in the conditional part of the rule.

Given this background, the numbers in Table 2-1 provide some insight into the amount of constraint knowledge that R1 has. Since the primary interest is in understanding the amount of domain knowledge required for the task, I will focus on the numbers in the second column; they do not differ significantly from the numbers in the first column except they show that R1's domain rules are more discriminating than its more general rules. Each of the task-specific rules has at least two conditions that are *sensitive to the context*. One of these conditions specifies that a particular context must be active; the other that there not be a more recent active context. The number is actually slightly greater than 2 because a few rules use the fact that a context is no longer active to determine that an action is appropriate. Beyond the context information, there are about 5.5 conditions, each of which can draw on from three to six pieces of information. This suggests that by and large, each rule captures only a very small part of the knowledge in the domain. The ratio of conditions to actions is about 2 to 1.

Rather free English translations of four of the nine rules associated with the context of assigning power supplies to sbi modules are shown in Figure 2-5. The first rule adds a power supply to the order if one is needed and if all of the power supplies ordered have already been configured. The second rule configures a power supply -- ie, indicates what sbi module the power supply should be connected to and where in what cabinet it should be put. The fourth rule has the same function as the second, and is fired instead of the second when both are satisfied because it is a special case. It contains the extra knowledge that when the sbi module that needs a power supply is a unibus adaptor, a particular regulator (the H7101) must be associated with the power supply. The third rule fires if all of the conditions required to satisfy the fourth rule are satisfied except the availability of a regulator; it adds the appropriate regulator to the order. The character of these rules is fairly representative of the rules in each of the contexts. In general, a context has a few rules associated

**ASSIGN-POWER-SUPPLY-1**

**IF: THE MOST CURRENT ACTIVE CONTEXT IS ASSIGNING A POWER SUPPLY  
 AND AN SBI MODULE OF ANY TYPE HAS BEEN PUT IN A CABINET  
 AND THE POSITION IT OCCUPIES IN THE CABINET (ITS NEXUS) IS KNOWN  
 AND THERE IS SPACE AVAILABLE IN THE CABINET FOR A POWER SUPPLY FOR THAT NEXUS  
 AND THERE IS NO AVAILABLE POWER SUPPLY  
 AND THE VOLTAGE AND FREQUENCY OF THE COMPONENTS ON THE ORDER IS KNOWN  
 THEN: FIND A POWER SUPPLY OF THAT VOLTAGE AND FREQUENCY AND ADD IT TO THE ORDER**

**ASSIGN-POWER-SUPPLY-2**

**IF: THE MOST CURRENT ACTIVE CONTEXT IS ASSIGNING A POWER SUPPLY  
 AND AN SBI MODULE OF ANY TYPE HAS BEEN PUT IN A CABINET  
 AND THE POSITION IT OCCUPIES IN THE CABINET (ITS NEXUS) IS KNOWN  
 AND THERE IS SPACE AVAILABLE IN THE CABINET FOR A POWER SUPPLY FOR THAT NEXUS  
 AND THERE IS AN AVAILABLE POWER SUPPLY  
 THEN: PUT THE POWER SUPPLY IN THE CABINET IN THE AVAILABLE SPACE**

**ASSIGN-POWER-SUPPLY-6**

**IF: THE MOST CURRENT ACTIVE CONTEXT IS ASSIGNING A POWER SUPPLY  
 AND A UNIBUS ADAPTOR HAS BEEN PUT IN A CABINET  
 AND THE POSITION IT OCCUPIES IN THE CABINET (ITS NEXUS) IS KNOWN  
 AND THERE IS SPACE AVAILABLE IN THE CABINET FOR A POWER SUPPLY FOR THAT NEXUS  
 AND THERE IS AN AVAILABLE POWER SUPPLY  
 AND THERE IS NO H7101 REGULATOR AVAILABLE  
 THEN: ADD AN H7101 REGULATOR TO THE ORDER**

**ASSIGN-POWER-SUPPLY-7**

**IF: THE MOST CURRENT ACTIVE CONTEXT IS ASSIGNING A POWER SUPPLY  
 AND A UNIBUS ADAPTOR HAS BEEN PUT IN A CABINET  
 AND THE POSITION IT OCCUPIES IN THE CABINET (ITS NEXUS) IS KNOWN  
 AND THERE IS SPACE AVAILABLE IN THE CABINET FOR A POWER SUPPLY FOR THAT NEXUS  
 AND THERE IS AN AVAILABLE POWER SUPPLY  
 AND THERE IS AN H7101 REGULATOR AVAILABLE  
 THEN: PUT THE POWER SUPPLY AND THE REGULATOR IN THE CABINET IN THE AVAILABLE SPACE**

**Figure 2-5: Sample rules from a single context**

with it whose function is to perform the basic (ordinary) actions appropriate in that context. Other rules associated with the context insure that if one or more of these basic rules are not satisfied because some required component is not on the order, the component will be added to the order so that the appropriate basic rules can fire. Still other rules associated with the context handle exceptional situations; these are typically special cases of the basic rules and fire before (and often instead of) the basic rules if their conditions are satisfied.

The rules shown in Figure 2-6 complement the sample rules in Figure 2-5; they give an idea of the kinds of cues that cause R1 to generate a new context and provide a somewhat broader understanding of how component information can be used to select among competing actions. The first three rules shown are all associated with contexts that become relevant within the first major subtask. The first rule checks to see if there are components on the order with incompatible voltage

**CHECK-VOLTAGE-AND-FREQUENCY-1**

**IF:** THE MOST CURRENT ACTIVE CONTEXT IS CHECKING VOLTAGE AND FREQUENCY  
 AND THERE IS A COMPONENT THAT REQUIRES ONE VOLTAGE OR FREQUENCY  
 AND THERE IS ANOTHER COMPONENT THAT REQUIRES A DIFFERENT VOLTAGE OR FREQUENCY  
**THEN:** ENTER THE CONTEXT OF FIXING VOLTAGE OR FREQUENCY MISMATCHES

**VERIFY-SBI-AND-MB-DEVICE-ADEQUACY-3**

**IF:** THE MOST CURRENT ACTIVE CONTEXT IS VERIFYING SBI AND MASSBUS DEVICE ADEQUACY  
 AND THERE ARE MORE THAN TWO MEMORY CONTROLLERS ON THE ORDER  
**THEN:** MARK THE EXTRA CONTROLLERS AS UNSUPPORTED (IE, NOT TO BE CONFIGURED)  
 AND MAKE A NOTE TO THE SALESPERSON  
 THAT ONLY TWO MEMORY CONTROLLERS ARE PERMITTED PER SYSTEM

**ASSIGN-UB-MODULES-EXCEPT-THOSE-CONNECTING-TO-PANELS-4**

**IF:** THE MOST CURRENT ACTIVE CONTEXT IS ASSIGNING DEVICES TO UNIBUS MODULES  
 AND THERE IS AN UNASSIGNED DUAL PORT DISK DRIVE  
 AND THE TYPE OF CONTROLLER IT REQUIRES IS KNOWN  
 AND THERE ARE TWO SUCH CONTROLLERS NEITHER OF WHICH HAS ANY DEVICES ASSIGNED TO IT  
 AND THE NUMBER OF DEVICES THAT THESE CONTROLLERS CAN SUPPORT IS KNOWN  
**THEN:** ASSIGN THE DISK DRIVE TO EACH OF THE CONTROLLERS  
 AND NOTE THAT THE TWO CONTROLLERS HAVE BEEN ASSOCIATED  
 AND THAT EACH SUPPORTS ONE DEVICE

**SELECT-BOX-AND-UB-MODULE-FOR-NEXT-SU-2**

**IF:** THE MOST CURRENT ACTIVE CONTEXT IS SELECTING A BOX AND A MODULE TO PUT IN IT  
 AND THE NEXT MODULE IN THE OPTIMAL SEQUENCE IS KNOWN  
 AND THE NUMBER OF SYSTEM UNITS OF SPACE THAT THE MODULE REQUIRES IS KNOWN  
 AND AT LEAST THAT MUCH SPACE IS AVAILABLE IN SOME BOX  
 AND THAT BOX DOES NOT CONTAIN MORE MODULES  
 THAN SOME OTHER BOX ON A DIFFERENT UNIBUS  
**THEN:** TRY TO PUT THAT MODULE IN THAT BOX

**PUT-UB-MODULE-6**

**IF:** THE MOST CURRENT ACTIVE CONTEXT IS PUTTING UNIBUS MODULES IN BACKPLANES IN SOME BOX  
 AND IT HAS BEEN DETERMINED WHICH MODULE TO TRY TO PUT IN A BACKPLANE  
 AND THAT MODULE IS A MULTIPLEXER TERMINAL INTERFACE  
 AND IT HAS NOT BEEN ASSOCIATED WITH ANY PANEL SPACE  
 AND THE TYPE AND NUMBER OF BACKPLANE SLOTS IT REQUIRES IS KNOWN  
 AND THERE ARE AT LEAST THAT MANY SLOTS AVAILABLE  
 IN A BACKPLANE OF THE APPROPRIATE TYPE  
 AND THE CURRENT UNIBUS LOAD ON THAT BACKPLANE IS KNOWN  
 AND THE POSITION OF THE BACKPLANE IN THE BOX IS KNOWN  
**THEN:** ENTER THE CONTEXT OF VERIFYING PANEL SPACE FOR A MULTIPLEXER

**CHECK-FOR-UB-JUMPER-CHANGES-6**

**IF:** THE MOST CURRENT ACTIVE CONTEXT IS CHECKING  
 FOR UNIBUS JUMPER CABLE CHANGES IN SOME BOX  
 AND THE BOX IS THE SECOND BOX IN SOME CABINET ON SOME UNIBUS  
 AND THERE IS AN UNCONFIGURED BOX ASSIGNED TO THAT UNIBUS  
 AND THE JUMPER CABLE THAT HAS BEEN ASSIGNED TO THE LAST BACKPLANE IN THE BOX  
 IS NOT A BC11A-10  
 AND THERE IS A BC11A-10 AVAILABLE  
 AND THE CURRENT LENGTH OF THE UNIBUS IS KNOWN  
**THEN:** MARK THE JUMPER CABLE ASSIGNED TO THE BACKPLANE AS NOT ASSIGNED  
 ASSIGN THE BC11A-10 TO THE BACKPLANE  
 INCREMENT THE CURRENT LENGTH OF THE UNIBUS BY TEN FEET

Figure 2-6: Sample rules from different contexts

or frequency requirements; if so, the rule generates the context of fixing the mismatches (which is done by isolating the set of components of the "wrong" voltage or frequency and replacing them on the order with components of the right voltage and frequency). The second rule is one of several rules that makes sure that various system limits are not exceeded; this particular rule is satisfied if there are more than two memory controllers on the order. The third rule assigns unibus devices to controllers; it is a special case rule since it deals with the case in which the device being assigned is a dual port device. The fourth, fifth, and sixth rules are all associated with contexts that become relevant while the unibus expansion cabinets are being filled. The fourth rule is one of several whose function is to select the next unibus module to configure and determine what box to put it in. The fifth rule tests for an exceptional case in the context of putting modules in a backplane; if the module that is to be put in the backplane is a multiplexer, then the context of verifying that there is panel space available in the cabinet is generated. The sixth rule checks to make sure that the jumper cable that connects two backplanes is of the appropriate length; this rule looks for the case in which the cable has to connect the second box in one cabinet with the first box in another.

## 2.5. R1'S USE OF ITS RULES

This subsection first provides some information from 20 sample runs to show how R1's use of its knowledge varies from order to order. Then the degree of conditionality in the task is analyzed. Finally, R1's use of its knowledge across the 20 runs is considered; this provides an indication of how much of R1's knowledge is "core" knowledge and how much of it is "exceptional" knowledge.

Each of the first 20 rows in Table 2-2 provide information about R1's performance on an order<sup>9</sup>; the final row shows the mean performance over the 20 orders. The column headed "Number of components ordered" shows the number of components that R1 configured. The column headed "Number of cycles" shows the number of rule firings required to configure these components. The first number includes the rule firings that generated output; the number in parentheses shows the number of cycles required excluding output generation. On 18 of the 20 orders, R1 had to generate only one candidate unibus module configuration. On the other two (orders 2 and 15), it generated three candidates for each; 281 rule firings were required, on the average, to generate each unacceptable candidate. The next two columns show the degree to which R1 uses the knowledge it has available. The first of these columns shows the ratio of the number of distinct rules that fired to

---

<sup>9</sup>The trace of the run that configured one of these orders (order 7) is given in Appendix 1.

Order number	Number of components ordered	Number of cycles	Percent of rules used	Percent of domain rules used	Run time in cpu minutes
1	97	1113 (877)	.46	.44	3.27
2	99	1864 (1622)	.47	.39	3.90
3	102	1083 (829)	.45	.42	2.73
4	129	1570 (1249)	.53	.51	4.75
5	65	665 (502)	.39	.35	1.28
6	64	812 (657)	.43	.41	1.48
7	142	1591 (1249)	.53	.51	5.20
8	54	613 (476)	.38	.35	1.10
9	63	702 (541)	.40	.37	1.35
10	73	770 (643)	.43	.41	2.13
11	78	928 (743)	.43	.40	1.73
12	115	1167 (921)	.43	.41	3.07
13	107	1151 (905)	.44	.41	2.90
14	78	915 (706)	.45	.42	1.80
15	99	1900 (1644)	.48	.41	3.80
16	105	1065 (907)	.47	.44	2.28
17	59	670 (517)	.40	.37	1.23
18	67	768 (592)	.42	.37	1.50
19	83	845 (642)	.43	.41	1.87
20	84	940 (734)	.44	.41	2.13
Mean	88	1056 (847)	.44	.41	2.48

Table 2-2: Statistics for individual runs

the total number of rules;<sup>10</sup> the other shows the ratio of the number of distinct domain specific rules that fired to the number of domain specific rules. The ratio of domain specific knowledge used to general knowledge used is about 2 to 1, which is approximately the ratio of domain specific knowledge available to general knowledge available. It should be noted that the number of distinct productions used does not grow directly with the number of cycles. On the average, for the orders that take the smallest number of cycles to configure, each rule is used twice; for the orders that take the largest number of cycles to configure, each rule is used five times. Only about 40 to 50 percent of the knowledge that R1 has available is required on any particular order (about 200 domain specific rules and about 100 general rules). The column that shows run time in cpu minutes includes the time it takes to generate output. The average amount of time required to generate output is .72 minutes;

<sup>10</sup>Since the 106 rules that generate output are invoked only after the configuration task has been done (and after all of the necessary domain knowledge has been used) they are excluded from consideration here and in Tables 2-3 and 2-4 as well.

thus the average run time required to configure an order is 1.76 minutes.<sup>11</sup> The average working memory size during a run is about 500 elements. Although some elements remain in working memory after their usefulness has ended, R1 deletes elements from working memory whenever it is clear that the information contained in those elements is no longer needed; thus mean working memory size provides a reasonably good estimate of the amount of information that it is useful to have readily available.

The fan-in and fan-out of R1's rules provide a measure of the degree of conditionality in the configuration task. The *fan-in* of a rule is the number of distinct rules that could fire immediately before that rule; the *fan-out* is the number of distinct rules that could fire immediately after the rule. The task of determining precisely the average fan-in and fan-out of R1 rules is more or less impossible. Since the elements asserted by a rule ordinarily partially instantiate a large number of other rules (in many different contexts), in order to determine fan-in and fan-out, one would have to trace the effects of all possible inputs (orders). It is, however, possible to establish an upper bound on the fan-in and fan-out of R1's rules by taking into account the consequences of R1's use of contexts. And it is possible to establish a lower bound by determining fan-in and fan-out for a subset of the possible inputs.

R1's use of contexts severely limits fan-in and fan-out. Since a production can fire only if it contains a condition element that matches the most current active context, the productions that can fire immediately before some production, *p*, are those that are associated with *p*'s context and those that assert *p*'s context. The only productions that can fire immediately after *p* are those that are associated with *p*'s context or, if *p* asserts a new context, those that are associated with that new context. Each context, on the average, has eight productions associated with it and fewer than two rules assert that context. Thus 10 is an upper bound on average fan-in, and 8 is an upper bound on average fan-out. The actual average fan-in and fan-out are, of course, considerably less than this. Since at least some of the rules associated with a context are mutually exclusive and since some have a special-case/general-case relationship, it is not the case that every rule associated with a context can fire immediately before and immediately after every other rule in that context.

Table 2-3 shows, for the 20 sample orders, the actual number of rules that, on the average, fired just before and just after each rule. The first five rows of the table indicate the fan-in and fan-out on five different runs; the other three rows indicate the fan-in and fan-out when more than one run is taken into account at one time. Throughout, the fan-in and fan-out are approximately the same. On the individual runs, fan-in and fan-out range between 1.21 and 1.42; about three-fourths of the rules have a fan-in equal to 1 and three-fourths of the rules have a fan-out equal to 1 (ie, these rules are

---

<sup>11</sup>The 20 sample runs were made on a PDP-10 (model KL).

preceded by only one rule and are succeeded by only one rule). As one would expect, as more runs are taken into account, the fan-in and fan-out increase. When 20 runs are analyzed together, both the fan-in and fan-out are about 2, and only about half of the rules have fan-in and fan-out equal to 1. A lower bound on fan-in and fan-out, then, is 2. A significant number of productions fired only once, and a significant number did not fire on any of the 20 runs; thus there is reason to suspect that the actual fan-in and fan-out is considerably greater than 2.

Number of runs	Percent of rules with fan-in = 1	Mean fan-in	Max fan-in	Percent of rules with fan-out = 1	Mean fan-out	Max fan-out
1	.73	1.35	6	.74	1.35	5
1	.72	1.40	7	.73	1.41	8
1	.68	1.42	7	.69	1.42	6
1	.85	1.21	4	.84	1.21	4
1	.79	1.27	5	.80	1.28	9
5	.57	1.67	9	.59	1.67	12
10	.49	1.86	9	.53	1.86	12
20	.46	1.98	10	.49	1.98	12

Table 2-3: fan-in and fan-out

If we were to assume that the average fan-in and fan-out of the rules is 3, and that about a third of the rules have a fan-in of 1 and a third have a fan-out of 1, then the rule network that emerges has the following characteristics: It has 666 nodes, one for each of the rules (excluding the 106 output generation rules). Perhaps half of these nodes have a single edge coming into them and/or a single edge going out; the other nodes have an average of four edges coming in and/or four edges going out. It should be clear that unless the selection of which edge to follow can be highly constrained, the cost (in nodes visited) of finding an adequate configuration (an appropriate path through the rule network) will be enormous. It is in this context that the power of the Match method used by R1 should become apparent. When R1 can configure a system without backtracking, it finds a single path through the rule network; this path consists, on the average, of about 800 nodes. When R1 must backtrack, it visits an additional N nodes, where N is the product of the number of unsuccessful unibus module sequences it tries (which is rarely more than 2) and the number of nodes that must be expanded to generate a candidate unibus module configuration (which is rarely more than 300).

Table 2-4 shows the extent to which R1's knowledge is used on the 20 runs. As the first row in the table shows, about a third of R1's knowledge was not used in any of the sample runs. The bottom row shows that another third of R1's knowledge was used for every order. This is the core knowledge in the domain. The use of the other third of R1's knowledge is fairly evenly distributed along the "core to exceptional" continuum. I suspect that if I were to perform the same analysis on 200 runs, that the

number of unused rules would be very small, but that the distribution of the non-core knowledge would remain quite even. The fact that so much of the domain knowledge is exceptional is not particularly surprising and is certainly consistent with the fan-in/fan-out analysis. The rule network is quite branchy at two-thirds of its nodes because within each context a wide variety of often mutually exclusive alternatives arise. But a third of the nodes in the network are connected by a single edge; the state transitions represented by these edges are presumably actions that are necessary in order to configure any system.

Number of runs	Percent of rules used on exactly N runs	Percent of domain rules used on exactly N runs
0	31.9	34.9
1	5.4	5.9
2	6.3	5.0
3	2.4	3.2
4	3.3	3.6
5	3.3	3.4
6	.9	.9
7	2.2	2.0
8	2.4	2.7
9	1.0	.7
10	.5	.7
11	.7	.9
12	.7	1.1
13	.5	.7
14	.7	.9
15	.6	.9
16	.1	.2
17	.3	.5
18	2.5	2.9
19	2.4	1.6
20	<u>31.9</u>	<u>27.3</u>
	100.0	100.0

**Table 2-4:** Knowledge use over 20 sample orders

## 2.6. DESIGN AND IMPLEMENTATION HISTORY

Up to this point in section 2, the focus has been on how R1 exploits the many constraints imposed by the configuration task to minimize (indeed, almost eliminate) the search required to perform the task. This subsection focuses on the engineering issues that underly R1's performance. R1 evolved. In a period of less than a year, it went from an idea, to a demonstration system that had most of the basic knowledge required in the domain but lacked the ability to deal with the intricacies of unusual orders, to a system that possesses true expertise. Its development parallels, in many respects, the development of the several domain-specific systems engineered by Stanford University's Heuristic



### Programming Project [Feigenbaum 77].

R1's implementation history divides quite naturally into two stages. During the first stage, which began in December of 1978 and lasted for about four months, I spent five or six days being tutored in the basics of VAX system configuration, read and reread the two manuals that describe many of the constraints on how VAX systems can be configured, and implemented an initial version of R1 (consisting of fewer than 200 domain rules) that could configure the simplest of orders correctly, but made numerous mistakes when it tried to tackle more complex orders.<sup>12</sup> The second stage, which lasted for another four months, was spent in asking experts in the VAX configuration task to examine R1's output, point out R1's mistakes, and indicate what knowledge R1 was lacking. R1 was sufficiently ignorant that finding mistakes was no problem. During this stage, R1's domain knowledge almost tripled.

During the first stage, the initial problem was extracting knowledge from human experts in the absence of a system that could display errorful behavior (ie, in the absence of a system that could be refined). It was clear, almost from the start, that the experts had a clear idea of (ie, could articulate) the structure of the task. There are a basic set of subtasks that must be performed when configuring any order; the relationship among these subtasks is easily described by the experts at an abstract level. It soon became clear that within a subtask, the decision about what to do was almost invariably exception driven. Put another way, the manner in which the configuration task is talked about has the flavor of "When you're performing subtask x, do y, unless z". Thus the mode of extracting information that I adopted was to present a subtask, ask what sorts of actions might be appropriate, and then, given an action, push for the exceptional cases. This maps nicely into a rule-based representation since there is some ordinary case rule, typically with little information required except knowledge of a few values of germane components, and then a collection of other rules that are special cases of that ordinary case rule. The additional conditions in these special case rules look for the "unless features" that signal an exceptional situation.

The problem with this approach to knowledge extraction is that humans do not seem to be particularly good at producing the exceptional cases on demand. I naively assumed, after implementing the basic system, that I had managed to extract most of the knowledge relevant to the configuration task. But as soon as the initial system was able to process orders, it became apparent that the knowledge was nowhere near adequate. Experts would look at the output, look at me strangely, and ask how R1 could have done that. As we discussed the problem, it became apparent that crucial conditions had simply not emerged during the initial knowledge extraction process. By and large, these conditions were not esoteric; as people explained to me why particular conditions

---

<sup>12</sup>During this first stage, R1's name was XCON.

had to hold, it usually seemed "obvious". As I tried to extract knowledge, I had pushed hard to determine if any relevant factors were being ignored, and typically I was assured that all of the relevant knowledge was out on the table. Perhaps I just don't know how to push. But it is more likely, I think, that configurers store their knowledge in a way that makes it extremely difficult for them to access conditions (constraints) given only actions.

Once the need for significant additions to R1's domain knowledge became apparent, the rule-based nature of the representation again demonstrated its worth. The second stage was characterized by what could be called "rule splitting" (complemented occasionally by what might be called "context splitting"). Given an inappropriate action on R1's part, it was quite easy (since the context in which the inappropriate action occurred was usually obvious) to find the offending rule. I would then simply ask the expert what he would have done differently and how he would have known to do that different thing. Sometimes a known feature of the situation could be used to signal the different action. More often, though, there was additional information that R1 did not know about that had to be taken into account. In the first case, all that was required to make R1's performance acceptable was to copy the offending rule and add a condition to it. In the second case, not only did the one rule become two, but information gathering rules had to be added to production memory. In both cases the refinement process was straightforward. When there were negative side-effects from a change, the effects were usually confined to the rules associated with the changed rule's context. But since R1's rules are mostly driven by positive features of a situation, such negative interactions were rare. Thus encoding knowledge in the form of rules appears to have facilitated the refinement process.

### 3. RESULTS

R1 has achieved a certain success as a configurer, but hopefully it has also contributed something to our understanding of domain-specific systems. This section briefly characterizes R1's current level of expertise. Then the features of R1 that have an interest beyond the configuration domain are isolated.

During October and November of 1979, R1 was involved in a formal validation procedure. The purpose of the validation stage was to determine whether R1 was expert enough in the configuration task to be used in place of human experts. Over the two month period, R1 was given 50 orders to configure. A team of six experts examined R1's output, spending from one to two hours on each order. In the course of examining the configurations, 12 pieces of errorful knowledge were uncovered. The rules responsible for the errors were modified and the orders were resubmitted to R1 and were all configured correctly. At the end of this validation stage, R1 was judged to be expert enough to be used in place of human experts. R1 then began to be integrated into the actual

manufacturing environment; various processes were put in place to enable it to be used on a regular basis and to control the additions of new product descriptions to its data base and the additions or modifications to its rules implied by these new products. R1 is currently used as part of the regular manufacturing process to produce configuration descriptions of each VAX system before it is assembled. These descriptions are then used by the technicians as they physically assemble the systems. Since the amount of time that ellapses between the placing of an order and assembly can be as much as six to nine months, it is important that each order be screened as soon as it is received to make sure that it is configurable (ie, does not required additional components in order to be a functional system). Thus R1 is also beginning to be used to configure each order booked at several regional field offices on the day it is booked. At the moment, these two uses of R1 are independent, but the expectation is that they will be merged, at some point in the future, to form the basis for a coherent order processing environment. Since the end of the validation stage, R1 has been run on over 300 orders and its output has been examined by experts; all but nine of the configurations that it has generated have been acceptable, and the rules responsible for the unacceptable configurations were easily fixed.

An important remaining question is whether anything has been learned from building R1 that can help guide the design of expert systems in other domains. I think that there are four ideas that have emerged that can provide some assistance:

- There are real-world domains in which Match can be used.
- There are real-world domains with sufficient structure to make it possible to recognize what to do at any time on the basis of a handful of situational cues.
- There are real-world domains in which production system languages can be used to advantage.
- The production system language, OPS4, is an appropriate tool for the construction of domain-specific systems.

I will consider each of these ideas in turn.

Because Match is a powerful method, it should be considered for all tasks that can support its use. In order to support Match, a task must be decomposable; it must involve moving from an initial state, through a number of intermediate states, to some desired state. If there are several alternative desired states, they must be equally acceptable. It must be possible, at any intermediate state, to determine whether that state is on a solution path: this implies that whatever information is required in order to determine whether a state is on a solution path can be available when it is needed. Even if a task does not meet these requirements, Match may be able to be used in conjunction with other, less powerful, methods. In the configuration task, the desired states (acceptable configurations) are not

all equally acceptable since unibus modules can be configured more or less optimally. Nevertheless, Match can be used as a stand-alone method to configure everything except unibus modules and as an embedded method to generate candidate unibus module configurations.

The fact that Match can be used for some task does not imply that its use is warranted. Match may be usable, but impractical. How practical Match is, depends, at least in part, on how difficult it is to order the intermediate states in such a way that no state can be reached before all of the information required to determine if that state is on a solution path has become available. In task domains with little structure, this ordering may be almost impossible to achieve since each state would have to test explicitly for the availability of that information. In task domains with considerable structure, on the other hand, the ordering may be quite easy to achieve since the fact that some other state or set of states have already been visited may guarantee the availability of much of that information and thus only a few explicit tests would be required.

Production system languages (or at least a subset of them) have two properties which make them particularly well-suited for domains like the configuration domain: (1) they make it easy to implement recognition-driven systems, and (2) they make it easy to implement systems incrementally. It is difficult in the case of the configuration task to specify in advance what particular pieces of component information will be relevant at any point in the configuration process. Since component information has no structuring principle, it is, at least conceptually, just a big set of information and is most naturally stored in a single global memory. Because a production system language provides such a memory, it is unnecessary to determine a priori what particular features of situations need to be attended to. Rather, each rule can simply watch for a set of features it recognizes. A production system language also makes it easy to incrementally build up a system's expertise. As we have seen, R1's shortcoming during its apprenticeship stage was that it was insufficiently discriminating. If a system's knowledge is represented as a set of rules that associate actions with the conditions under which they are appropriate, then the unit of representation is one which enables this shortcoming to be easily overcome. Given a criticism of some aspect of the configuration by an expert, all that was necessary in order to refine R1's knowledge was to find the offending rule, ask the expert to point out the problem with the condition elements in the rule, and then either modify the rule or split it into two rules that would discriminate between two previously undifferentiated states.

Though various production system languages could have been used to implement R1, OPS4 has three characteristics that make it a particularly appropriate vehicle. First, OPS4 is an efficient implementation of a recognize-act architecture; it uses techniques that make cycle time essentially independent of the size of both production memory and working memory. Second, it provides a special case conflict resolution strategy which makes it easy to deal with knowledge of exceptional situations. Third, OPS4 is a fairly powerful language for expressing patterns. This makes it easy to

describe a wide variety of different sorts of constraints. There is, however, a drawback: an OPS4 rule looks only remotely like its English equivalent.<sup>13</sup> As a consequence, people for whom OPS4 is not a second language require a human translator in order to understand the rules. This makes the refinement process less straightforward than it otherwise might be.

## CONCLUDING REMARKS

R1 has proven itself to be a highly competent configurator of VAX-11/780 systems. The configurations that it produces are consistently adequate, and the information that it makes available to the technicians who physically assemble systems is far more detailed than that produced by the humans who do the task. In addition to being used to provide guidance to the technicians, R1 also screens orders up front to insure that a customer will not be unpleasantly surprised, sometime after placing his order, to hear that it is unconfigurable. The function that R1 currently serves is a valuable one, and once a site-specific floor layout capability has been added, R1 will be fully competent in its little domain.

There are, however, some obvious ways in which to enlarge its domain to make it a more helpful system. Work has already begun on augmenting R1's knowledge to enable it to configure other computer systems manufactured by DEC. In addition, we plan to augment its knowledge so that it will be able to help with the scheduling of system delivery dates. We also plan to augment R1's knowledge so that it will be able to provide interactive assistance to a customer or salesperson that will allow him, if he wishes, to specify some of the capabilities of the system he wants and let R1 select the set of components that will provide those capabilities. Ultimately we hope to develop a salesperson's assistant, an R1 that can help a customer identify the system that best suits his needs.

## ACKNOWLEDGEMENTS

Many people have provided help in various forms. Jon Bentley, Scott Fahlman, Charles Forgy, Betsy Herk, Jill Larkin, Allen Newell, Paul Rosenbloom, and Mike Rychener gave me considerable encouragement and many valuable ideas. Dave Barstow, Bruce Buchanan, Bob Englemore, Penny Nii, Ted Shortliffe, and Mark Stefik contributed their knowledge engineering expertise. Finally, Jim Baratz, Alan Belancik, Dick Caruso, Sam Fuller, Linda Marshall, Kent McNaughton, Vaidis Mongirdas, Dennis O'Connor, and Mike Powell, all of whom are at DEC, assisted in bringing R1 up to industry standards.

---

<sup>13</sup>The OPS4 rules that are rendered in English in Figures 2-5 and 2-6 are shown in Appendix 3.

## REFERENCES

- [Amarel 77] Amarel, S., J. S. Brown, B. Buchanan, P. Hart, C. Kulikowski, W. Martin, and H. Pople.  
Reports of panel on applications of artificial intelligence.  
In *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, pages 994-1006. MIT, 1977.
- [Davis 77] Davis, R., B. Buchanan, and E. Shortliffe.  
Production rules as a representation for a knowledge-based consultation program.  
*Artificial Intelligence* 8(1):15-45, 1977.
- [Feigenbaum 77] Feigenbaum, E. A.  
The art of artificial intelligence.  
In *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, pages 1014-1029. MIT, 1977.
- [Forgy 77] Forgy, C. L. and J. McDermott.  
OPS, A domain-independent production system language.  
In *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, pages 933-939. MIT, 1977.
- [Forgy 79] Forgy, C. L.  
*The OPS4 user's manual*.  
Technical Report, Carnegie-Mellon University, Department of Computer Science, 1979.
- [Forgy 80] Forgy, C. L.  
*RETE: A fast algorithm for the many pattern/many object pattern match problem*.  
Technical Report, Carnegie-Mellon University, Department of Computer Science, 1980.
- [McDermott 78] McDermott, J. and C. L. Forgy.  
Production system conflict resolution strategies.  
In D. A. Waterman and F. Hayes-Roth (editors), *Pattern-Directed Inference Systems*, pages 177-199. Academic Press, 1978.
- [Newell 63] Newell, A., J. C. Shaw, and H. A. Simon.  
Empirical explorations with the Logic Theory Machine.  
In E. A. Feigenbaum and J. Feldman (editors), *Computers and Thought*, pages 109-133. McGraw-Hill, 1963.
- [Newell 69] Newell, A.  
Heuristic programming: ill-structured problems.  
In J. S. Aronofsky (editor), *Progress in Operations Research*, pages 361-414. John Wiley and Sons, 1969.
- [Newell 77] Newell, A.  
Knowledge representation aspects of production systems.  
In *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, pages 987-988. MIT, 1977.
- [Nilsson 80] Nilsson, N.

*Principles of Artificial Intelligence.*  
Tioga Publishing Co., 1980.

[Waterman 78] Waterman, D. A. and F. Hayes-Roth (editors).  
*Pattern-Directed Inference Systems.*  
Academic Press, 1978.

## APPENDIX 1

This appendix shows a sample order and the trace (at the context level) of the run that configured the order. The number preceding each context name is the cycle on which that context was entered. A context that is invoked by a rule in another context is indented in order to show the invoking context (ie, to make the relationships among contexts apparent). Appendix 2 contains the output that R1 produced for this sample order.

```

COMPONENTS ORDERED:
  1 SV-AXHHA-LA      [packaged system]
  1 FP780-AA        [floating point accelerator]
  1 OW780-AA        [unibus adaptor]
  1 BA11-KE         [unibus expansion cabinet box]
  6 MS780-DC        [memory]
  1 MS780-CA        [memory controller]
  1 HQ602-HA        [cpu expansion cabinet]
  1 H7111-A         [clock battery backup]
  1 H7112-A         [memory battery backup]
  1 REPO5-AA        [single port disk drive]
  4 RP05-BA         [dual port disk drive]
  1 TEE16-AE        [tape drive with formatter]
  2 TE16-AE         [tape drive]
  8 RK07-EA         [single port disk drive]
  1 DR11-B          [direct memory access interface]
  1 LP11-CA         [line printer]
  1 DZ11-F          [multiplexer with panel]
  1 DZ11-B          [multiplexer]
  2 LA36-CE         [hard copy terminal]

```

```

1. MAJOR-SUBTASK-TRANSITION
2.   SET-UP
3.     UNBUNDLE-COMPONENTS
53.    NOTE-CUSTOMER-GENERATED-EXCEPTIONS
56.    NOTE-UNSUPPORTED-COMPONENTS
67.    CHECK-VOLTAGE-AND-FREQUENCY
104.   CHECK-FOR-TYPE-OR-CLASS-CHANGES
110.   VERIFY-SBI-AND-MB-DEVICE-ADEQUACY
111.   COUNT-SBI-MODULES-AND-MB-DEVICES
126.   GET-NUMBER-OF-BYTES-AND-COUNT-CONTROLLERS
137.   FIND-UBA-MBA-CAPACITY-AND-USE
146.   VERIFY-MEMORY-ADEQUACY
148.   PARTITION-MEMORY
160.   ASSIGN-UB-MODULES-EXCEPT-THOSE-CONNECTING-TO-PANELS
177.   VERIFY-UB-MODULES-FOR-DEVICES-CONNECTING-TO-PANELS
178.   FIND-ATTRIBUTE-OF-TYPE-IN-SYSTEM
179.   VERIFY-COMPONENT-OF-SYSTEM
207.   NOTE-POSSIBLY-FORGOTTEN-COMPONENTS
213.   CHECK-FOR-MISSING-ESSENTIAL-COMPONENTS
216. MAJOR-SUBTASK-TRANSITION
216.   DELETE-UNNEEDED-ELEMENTS-FROM-WM
236.   FILL-CPU-OR-CPUX-CABINET
240.   ADD-UBAS
246.   ASSIGN-POWER-SUPPLY
251.   ADD-MBAS
262.   DISTRIBUTE-MB-DEVICES
260.   ASSIGN-SLAVES-TO-MASTERS
269.   ASSIGN-POWER-SUPPLY
272.   FILL-MEMORY-SLOTS
278.   SHIFT-BOARDS
298.   ADD-MEMORY-MODULE-SIMULATORS
306.   ASSIGN-POWER-SUPPLY
312.   FILL-CPU-SLOTS
318.   ASSIGN-POWER-SUPPLY

```



322. ADD-NECESSARY-SIMULATORS  
 326. DELETE-TEMPLATES  
 340. DELETE-UNNEEDED-ELEMENTS-FROM-WM  
 363. FILL-CPU-OR-CPUX-CABINET  
 366. ADD-MBAS  
 369. ASSIGN-POWER-SUPPLY  
 362. ADD-UBAS  
 364. FILL-MEMORY-SLOTS  
 369. SHIFT-BOARDS  
 389. ADD-MEMORY-MODULE-SIMULATORS  
 396. ASSIGN-POWER-SUPPLY  
 399. TERMINATE-SBI  
 402. ADD-NECESSARY-SIMULATORS  
 406. DELETE-TEMPLATES  
 416. MAJOR-SUBTASK-TRANSITION  
 416. SET-FILL-MODE  
 417. GENERATE-OPTIMAL-SEQUENCE  
 436. ASSIGN-UBAS-TO-BOXES-TO-CABINETS  
 436. ASSIGN-UBAS-TO-BOXES  
 441. DISTRIBUTE-BOXES-AMONG-CABINETS  
 442. SET-UP-FOR-BOX-ASSIGNMENTS  
 446. ASSIGN-BOXES-TO-CABINETS  
 452. COMPUTE-DISTANCES-FROM-UBAS-TO-BOXES  
 458. SET-SEQUENCING-MODE  
 462. FILL-BOXES  
 465. FILL-HALF-BOXES  
 466. SELECT-BOX-AND-UB-MODULE-FOR-NEXT-SU  
 470. ASSIGN-BACKPLANE-TO-BOX  
 474. GENERATE-SLOT-TEMPLATES  
 478. PUT-UB-MODULE  
 482. LEAVE-BACKPLANE  
 486. AUGMENT-UB-LENGTH  
 488. GET-UB-JUMPER  
 491. CHECK-NEED-FOR-UB-REPEATER  
 497. SELECT-BOX-AND-UB-MODULE-FOR-NEXT-SU  
 501. ASSIGN-BACKPLANE-TO-BOX  
 506. GENERATE-SLOT-TEMPLATES  
 510. PUT-UB-MODULE  
 518. ADD-SUBOPTIMAL-UB-MODULE  
 527. LEAVE-BACKPLANE  
 540. AUGMENT-UB-LENGTH  
 543. GET-UB-JUMPER  
 547. CHECK-NEED-FOR-UB-REPEATER  
 568. LEAVE-HALF-BOX  
 569. CHECK-FOR-UB-JUMPER-CHANGES  
 561. CHECK-TERMINATION-CONDITIONS  
 568. SELECT-BOX-AND-UB-MODULE-FOR-NEXT-SU  
 571. ASSIGN-BACKPLANE-TO-BOX  
 576. GENERATE-SLOT-TEMPLATES  
 580. PUT-UB-MODULE  
 581. ASSOCIATE-MULTIPLXER-WITH-PANEL-SLOT  
 590. ASSOCIATE-MULTIPLXER-WITH-PANEL-SLOT  
 598. ASSOCIATE-MULTIPLXER-WITH-PANEL-SLOT  
 604. ADD-SUBOPTIMAL-UB-MODULE  
 608. ASSOCIATE-MULTIPLXER-WITH-PANEL-SLOT  
 616. ADD-SUBOPTIMAL-UB-MODULE  
 617. LEAVE-BACKPLANE  
 626. AUGMENT-UB-LENGTH  
 629. GET-UB-JUMPER  
 633. CHECK-NEED-FOR-UB-REPEATER  
 643. LEAVE-HALF-BOX  
 644. CHECK-FOR-UB-JUMPER-CHANGES  
 646. CHECK-TERMINATION-CONDITIONS  
 657. SELECT-BOX-AND-UB-MODULE-FOR-NEXT-SU  
 660. ASSIGN-BACKPLANE-TO-BOX  
 663. GENERATE-SLOT-TEMPLATES  
 667. PUT-UB-MODULE  
 668. ASSOCIATE-MULTIPLXER-WITH-PANEL-SLOT  
 677. ASSOCIATE-MULTIPLXER-WITH-PANEL-SLOT  
 690. LEAVE-BACKPLANE  
 711. AUGMENT-UB-LENGTH  
 714. GET-UB-JUMPER  
 718. CHECK-NEED-FOR-UB-REPEATER

732. LEAVE-HALF-BOX  
733. CHECK-FOR-UB-JUMPER-CHANGES  
736. CHECK-TERMINATION-CONDITIONS  
738. ASSIGN-UB-JUMPER-CABLES-TO-BOX  
749. LEAVE-HALF-BOX  
750. CHECK-FOR-UB-JUMPER-CHANGES  
762. CHECK-TERMINATION-CONDITIONS  
766. ASSIGN-UB-JUMPER-CABLES-TO-BOX  
769. ACCEPT-UNIBUS-CONFIGURATION  
832. MAJOR-SUBTASK-TRANSITION  
833. ASSIGN-TERMINALS-TO-LINES  
834. PUT-PANELS-IN-UBX-CABINET  
848. MAKE-TERMINAL-ASSIGNMENTS  
854. MAJOR-SUBTASK-TRANSITION  
855. LAY-OUT-SYSTEM  
857. FIND-FLOOR-RANKINGS  
882. DETERMINE-FLOOR-POSITIONS  
886. DETERMINE-FLOOR-POSITIONS-OF-CABINETS  
893. DETERMINE-FLOOR-POSITIONS-OF-DEVICES  
900. DETERMINE-FLOOR-POSITIONS-OF-SLAVES  
908. DETERMINE-FLOOR-POSITIONS-OF-DEVICES  
920. DETERMINE-FLOOR-POSITIONS-OF-DEVICES  
934. DETERMINE-FLOOR-POSITIONS-OF-DEVICES  
942. DETERMINE-FLOOR-POSITIONS-OF-DEVICES  
973. MAJOR-SUBTASK-TRANSITION  
974. COMPUTE-CABLE-LENGTHS  
1021. FIND-LENGTHS-OF-CABLES-ON-ORDER  
1135. ASSIGN-CABLES  
1179. FIND-LENGTHS-OF-CABLES-ON-ORDER  
1183. FIND-LENGTHS-OF-CABLES-ON-ORDER  
1187. FIND-LENGTHS-OF-CABLES-ON-ORDER  
1192. NOTE-POSSIBLY-FORGOTTEN-COMPONENTS  
1196. GENERATE-COMPONENT-NUMBERS-FOR-CABLES  
1248. GENERATE-OUTPUT

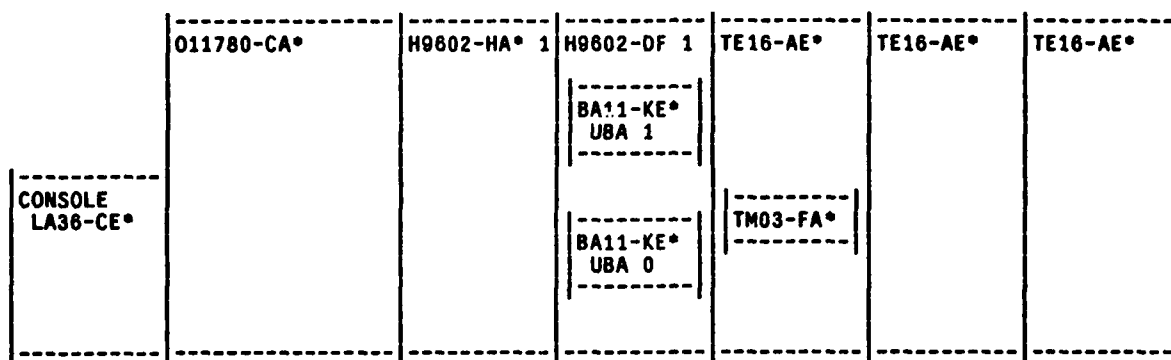
## APPENDIX 2

This appendix contains the output generated by R1 on the sample run shown in Appendix 1. The information on this page provides an overview of the configured system. The next two pages give information about the composition of the system; note that bundles are "exploded" in order to show their constituent components. The remaining pages of Appendix 2 display the spatial relationships among various subsets of the components in more detail.

DEC-NUMBER VO-7 12-1-79

### CABINET LAYOUT

### COMPONENTS ORDERED



- 1 SV-AXHHA-LA
- 1 LA36-CE
- 1 070-1577-02
- 1 QE001-AV
- 1 BA11-KE
- 1 DR11-B
- 1 DW780-AA
- 1 DZ11-B
- 1 DZ11-F
- 1 FP780-AA
- 1 H7111-A
- 1 H9602-HA
- 2 LA36-CE
- 1 LP11-CA
- 1 MS780-CA
- 6 MS780-DC
- 1 REP06-AA
- 8 RK07-EA
- 4 RP06-BA
- 2 TE16-AE
- 1 TEE16-AE

### O11780-CA\*

F	D	MS780-CA*	D	R
P	W	3328 KB	W	H
7	7	MS780-DC	7	7
8	8	MS780-DC	8	8
0	0	MS780-DC	0	0
		MS780-DA		
	0		1	0
KA780-AA*				
KC780-AA		H7112-A		

### H9602-HA\*

MS780-CA*	R
3328 KB	H
MS780-DC	7
MS780-DC	8
MS780-DC	0
MS780-DA	
	1

### BOX 1 (CABINET 1)

END	9
	8 DZ11-B*
	7 DZ11-B*
IN	6
OUT	4 ----
	3 ----
	2 ----
IN	1 DR11

### BOX 2 (CABINET 1)

END	19
	18
	17
	16
	16
	14 LP11
	13 DZ11-D*
	12 DZ11-D*
IN	11
OUT	9 ----
	8 ----
	7 ----
	6 ----
	5 ----
	4 RK611*
	3
	2
IN	1

### FREE STANDING OPTIONS

RP06-AA*	RP06-BA*	RK07-EA*	LA36-CE*	LP14-CA
1	4	8	2	1

## COMPONENTS ORDERED

1 BA11-KE  
   1 BA11-KE\*  
   1 BC11A-10  
 1 DR11-B  
   1 DR11  
   1 070-08329-00  
 1 DW780-AA  
   1 DW780  
   1 H7100-A  
 1 DZ11-B  
   1 BC06W-16  
   1 DZ11-B\*  
 1 DZ11-F  
   1 DZ11-C  
     1 DZ11-D  
       1 BC08S-16  
       1 DZ11-D\*  
       1 H317-F  
   1 DZ11-D  
     1 BC08S-16  
     1 DZ11-D\*  
 1 FP780-AA  
   1 FP780-AA\*  
   1 H7100-A  
 1 H7111-A  
 1 H9602-HA  
   1 H9602-HA\*  
   6 017-00087-03  
 2 LA36-CE  
   1 BC06F-16  
   1 LA36-CE\*  
 1 LP11-CA  
   1 LP11  
   1 LP14-CA  
   1 070-11212-25  
 1 MS780-CA  
   1 H7100-A  
   1 H7102  
   1 H7103  
   1 MS780-CA\*  
   1 MS780-DA  
 6 MS780-DC  
 1 REPO6-AA  
   1 BC06S-25  
   1 RH780-AA  
     1 H7100-A  
     1 RH780  
   1 RPO5-AA  
     1 BC06S-03  
     1 RPO5-AA\*  
 8 RK07-EA  
   1 070-12292-08  
   1 RK07-EA\*  
 4 RPO5-BA  
   2 BC06S-03  
   1 RPO5-BA\*  
 1 SV-AXHHA-LA  
   1 LA36-CE  
     1 BC06F-16  
     1 LA36-CE\*  
   1 070-15777-02  
     1 H9602-DA  
       1 BA11-KE  
         1 BA11-KE\*  
         1 BC11A-10  
       3 BC06L-16  
       1 DZ11-A  
         1 DZ11-B  
           1 BC06W-16  
           1 DZ11-B\*  
         1 H317-E  
       8 G727  
       1 H9602-DF  
       1 M9014  
       1 M9302

1 070-11164  
 1 011780-CA  
     1 DW780  
     1 H7111-A  
     1 KA780-AA  
         2 H7100-A  
         1 H7101  
         1 KA780-AA\*  
     1 KC780-AA  
     1 M9043  
     1 MS780-CA  
         1 H7100-A  
         1 H7102  
         1 H7103  
         1 MS780-CA\*  
         1 MS780-DA  
     1 011780-CA\*  
 1 RK07-EA  
     1 070-12292-08  
     1 RK07-EA\*  
 1 RK711-EA  
     1 070-12292-26  
     1 RK07-EA\*  
     1 RK611  
         3 G727  
         1 M9202  
         1 070-12412-00  
         1 RK611\*  
     1 QE001-AV  
 2 TE16-AE  
     3 BC06R-10  
     1 TE16-AE\*  
 1 TEE16-AE  
     3 BC06R-20  
     1 RH780-AA  
         1 H7100-A  
         1 RH780  
     1 TE16-AE  
         3 BC06R-10  
         1 TE16-AE\*  
     1 TM03-FA  
         3 BC06R-10  
         1 TM03-FA\*

**SUBSTITUTIONS**  
NONE

**COMPONENTS ADDED**

1 H7101  
 1 M9014  
 1 M9202  
 1 M9302  
 1 070-11528

**THE FOLLOWING COMPONENTS WERE NOT CONFIGURED**

2 RK07-EA\* (POSSIBLY-FORGOTTEN-PREREQUISITE DISK-DRIVE CONTROLLER)  
 1 H7111-A (NOT-NEEDED)  
 1 H7100-A (NOT-NEEDED)  
 6 BC06R-10 (NOT-NEEDED)  
 2 070-12292-08 (NOT-NEEDED)  
 1 BC06F-15 (NOT-NEEDED)  
 2 BC11A-10 (NOT-NEEDED)

**POSSIBLY FORGOTTEN COMPONENTS**

3 BC06L-15 TO CONNECT UBA 0 TO BOX 1 OF CABINET 1  
 1 RK611\* TO SUPPORT 2 UNUSED DISK-DRIVE

**UNUSED CAPACITY**

THE MEMORY CONTROLLER (MS780-CA\*) IN THE CPU CABINET COULD SUPPORT 768 K BYTES MORE MEMORY  
 THE MEMORY CONTROLLER (MS780-CA\*) IN THE CPU CABINET COULD SUPPORT 768 K BYTES MORE MEMORY  
 DZ11-D\* 1 IN BOX 2 OF CABINET 1 COULD SUPPORT 7 MORE MA20 LINES (THROUGH PANEL 2)  
 DZ11-D\* 2 IN BOX 2 OF CABINET 1 COULD SUPPORT 7 MORE MA20 LINES (THROUGH PANEL 2)  
 DZ11-B\* 2 IN BOX 1 OF CABINET 1 IS UNUSED (IT CAN SUPPORT 8 EIA LINES (THROUGH PANEL 1))  
 DZ11-B\* 1 IN BOX 1 OF CABINET 1 IS UNUSED (IT CAN SUPPORT 8 EIA LINES (THROUGH PANEL 1))  
 MBA 0 COULD SUPPORT 2 MORE MB-DEVICES AND ITS 1 MASTER TAPE DRIVES COULD SUPPORT 6 MORE SLAVES  
 MBA 1 COULD SUPPORT 4 MORE MB-DEVICES

THE OPTIMAL ORDERING OF THE MODULES IS (DR11 RK611\* DZ11-B\* DZ11-B\* DZ11-D\* DZ11-D\* LP11)  
 THE ORDERING ON UBA 0 IS (DR11 DZ11-B\* DZ11-B\*)  
 THE ORDERING ON UBA 1 IS (RK611\* DZ11-D\* DZ11-D\* LP11)

## FLOOR-PLAN (TOP VIEW)

LA36-CE* CONSOLE	O11780-CA* CPU-CABINET	H9602-HA* CPUX-CABINET 1	H9602-DF UBX-CABINET 1		
TE16-AE* (AND TM03-FA*) MB-DEVICE 6 MBA 0	TE16-AE* SLAVE 1 MB-DEVICE 6 MBA 0	TE16-AE* SLAVE 2 MB-DEVICE 6 MBA 0	RP06-BA* MB-DEVICE 0 MBA 0 MB-DEVICE 0 MBA 1		
RP06-BA* MB-DEVICE 1 MBA 0 MB-DEVICE 1 MBA 1	RP06-BA* MB-DEVICE 2 MBA 0 MB-DEVICE 2 MBA 1	RP06-BA* MB-DEVICE 3 MBA 0 MB-DEVICE 3 MBA 1	RP06-AA* MB-DEVICE 4 MBA 0		
RK07-EA* 0 RK611* 1 BOX 2 CABINET 1 UBA 1	RK07-EA* 1 RK611* 1 BOX 2 CABINET 1 UBA 1	RK07-EA* 2 RK611* 1 BOX 2 CABINET 1 UBA 1	RK07-EA* 3 RK611* 1 BOX 2 CABINET 1 UBA 1	RK07-EA* 4 RK611* 1 BOX 2 CABINET 1 UBA 1	RK07-EA* 6 RK611* 1 BOX 2 CABINET 1 UBA 1
RK07-EA* 6 RK611* 1 BOX 2 CABINET 1 UBA 1	RK07-EA* 7 RK611* 1 BOX 2 CABINET 1 UBA 1	LA36-CE* 0 SLOT 8 PANEL 2 CABINET 1 UBA 1	LA36-CE* 1 SLOT 0 PANEL 2 CABINET 1 UBA 1	LP14-CA LP11 1 BOX 2 CABINET 1 UBA 1	

## CABLE ASSIGNMENTS

BC06R-20 1 FROM MB-DEVICE 5 (TE16-AE\*) ON MBA 0 TO ITS TM03-FA\*  
LENGTH: 20.0 (ESTIMATED-LENGTH-REQUIRED: 10.0)

BC06R-20 2 FROM MB-DEVICE 5 (TE16-AE\*) ON MBA 0 TO ITS TM03-FA\*  
LENGTH: 20.0 (ESTIMATED-LENGTH-REQUIRED: 10.0)

BC06R-20 3 FROM MB-DEVICE 5 (TE16-AE\*) ON MBA 0 TO ITS TM03-FA\*  
LENGTH: 20.0 (ESTIMATED-LENGTH-REQUIRED: 10.0)

BC06R-10 1 FROM MB-DEVICE 5 SLAVE 1 (TE16-AE\*) ON MBA 0 TO SLAVE 2 (TE16-AE\*)  
LENGTH: 10.0 (ESTIMATED-LENGTH-REQUIRED: 9.0)

BC06R-10 2 FROM MB-DEVICE 5 SLAVE 1 (TE16-AE\*) ON MBA 0 TO SLAVE 2 (TE16-AE\*)  
LENGTH: 10.0 (ESTIMATED-LENGTH-REQUIRED: 9.0)

BC06R-10 3 FROM MB-DEVICE 5 SLAVE 1 (TE16-AE\*) ON MBA 0 TO SLAVE 2 (TE16-AE\*)  
LENGTH: 10.0 (ESTIMATED-LENGTH-REQUIRED: 9.0)

BC06R-10 4 FROM MB-DEVICE 5 (TE16-AE\*) ON MBA 0 TO SLAVE 1 (TE16-AE\*)  
LENGTH: 10.0 (ESTIMATED-LENGTH-REQUIRED: 9.0)

BC06R-10 5 FROM MB-DEVICE 5 (TE16-AE\*) ON MBA 0 TO SLAVE 1 (TE16-AE\*)  
LENGTH: 10.0 (ESTIMATED-LENGTH-REQUIRED: 9.0)

BC06R-10 6 FROM MB-DEVICE 5 (TE16-AE\*) ON MBA 0 TO SLAVE 1 (TE16-AE\*)  
LENGTH: 10.0 (ESTIMATED-LENGTH-REQUIRED: 9.0)

BC06L-15 1 FROM UBA 1 TO BOX 2 OF CABINET 1  
LENGTH: 15.0 (ESTIMATED-LENGTH-REQUIRED: 25.0)

BC06L-15 2 FROM UBA 1 TO BOX 2 OF CABINET 1  
LENGTH: 15.0 (ESTIMATED-LENGTH-REQUIRED: 25.0)

BC06L-15 3 FROM UBA 1 TO BOX 2 OF CABINET 1  
LENGTH: 15.0 (ESTIMATED-LENGTH-REQUIRED: 25.0)

070-12292-25 1 FROM RK611\* 1 IN BOX 2 OF CABINET 1 TO RK07-EA\* 0  
LENGTH: 25.0 (ESTIMATED-LENGTH-REQUIRED: 36.6)

070-12292-08 1 FROM RK07-EA\* 6 ON RK611\* 1 IN BOX 2 OF CABINET 1 TO RK07-EA\* 7  
LENGTH: 8.0 (ESTIMATED-LENGTH-REQUIRED: 8.0)

070-12292-08 2 FROM RK07-EA\* 5 ON RK611\* 1 IN BOX 2 OF CABINET 1 TO RK07-EA\* 6  
LENGTH: 8.0 (ESTIMATED-LENGTH-REQUIRED: 8.0)

070-12292-08 3 FROM RK07-EA\* 4 ON RK611\* 1 IN BOX 2 OF CABINET 1 TO RK07-EA\* 5  
LENGTH: 8.0 (ESTIMATED-LENGTH-REQUIRED: 8.0)

070-12292-08 4 FROM RK07-EA\* 3 ON RK611\* 1 IN BOX 2 OF CABINET 1 TO RK07-EA\* 4  
LENGTH: 8.0 (ESTIMATED-LENGTH-REQUIRED: 8.0)

070-12292-08 5 FROM RK07-EA\* 2 ON RK611\* 1 IN BOX 2 OF CABINET 1 TO RK07-EA\* 3  
LENGTH: 8.0 (ESTIMATED-LENGTH-REQUIRED: 8.0)

070-12292-08 6 FROM RK07-EA\* 1 ON RK611\* 1 IN BOX 2 OF CABINET 1 TO RK07-EA\* 2  
LENGTH: 8.0 (ESTIMATED-LENGTH-REQUIRED: 8.0)

070-12292-08 7 FROM RK07-EA\* 0 ON RK611\* 1 IN BOX 2 OF CABINET 1 TO RK07-EA\* 1  
LENGTH: 8.0 (ESTIMATED-LENGTH-REQUIRED: 8.0)

BC06S-25 1 FROM MBA 0 TO MB-DEVICE 0 (RP06-BA\*)  
LENGTH: 25.0 (ESTIMATED-LENGTH-REQUIRED: 29.6)

BC06S-03 1 FROM MB-DEVICE 3 (RP05-BA\*) ON MBA 0 TO MB-DEVICE 4 (RP05-AA\*)  
LENGTH: 3.0 (ESTIMATED-LENGTH-REQUIRED: 9.0)

BC06S-03 2 FROM MB-DEVICE 2 (RP05-BA\*) ON MBA 0 TO MB-DEVICE 3 (RP05-BA\*)

LENGTH: 3.0 (ESTIMATED-LENGTH-REQUIRED: 9.0)  
 BC06S-03 3 FROM MB-DEVICE 2 (RP05-BA\*) ON MBA 1 TO MB-DEVICE 3 (RP05-BA\*)  
 LENGTH: 3.0 (ESTIMATED-LENGTH-REQUIRED: 9.0)  
 BC06S-03 4 FROM MB-DEVICE 1 (RP05-BA\*) ON MBA 0 TO MB-DEVICE 2 (RP05-BA\*)  
 LENGTH: 3.0 (ESTIMATED-LENGTH-REQUIRED: 9.0)  
 BC06S-03 6 FROM MB-DEVICE 1 (RP05-BA\*) ON MBA 1 TO MB-DEVICE 2 (RP05-BA\*)  
 LENGTH: 3.0 (ESTIMATED-LENGTH-REQUIRED: 9.0)  
 BC06S-03 6 FROM MB-DEVICE 0 (RP05-BA\*) ON MBA 0 TO MB-DEVICE 1 (RP05-BA\*)  
 LENGTH: 3.0 (ESTIMATED-LENGTH-REQUIRED: 9.0)  
 BC06S-03 7 FROM MB-DEVICE 0 (RP05-BA\*) ON MBA 1 TO MB-DEVICE 1 (RP05-BA\*)  
 LENGTH: 3.0 (ESTIMATED-LENGTH-REQUIRED: 9.0)  
 BC06S-03 8 FROM MBA 1 TO MB-DEVICE 0 (RP05-BA\*)  
 LENGTH: 3.0 (ESTIMATED-LENGTH-REQUIRED: 25.3)  
 BC06S-03 9 FROM MB-DEVICE 4 (RP05-AA\*) ON MBA 0 TO TM03-FA\* IN MB-DEVICE 6 (TE16-AE\*)  
 LENGTH: 3.0 (ESTIMATED-LENGTH-REQUIRED: 27.0)

O70-11212-25 1 FROM LP11 1 IN BOX 2 OF CABINET 1 TO A LP14-CA  
 LENGTH: 25.0 (ESTIMATED-LENGTH-REQUIRED: 58.6)

BC06F-16 1 FROM SLOT 8 IN PANEL 2 OF CABINET 1 TO A LA36-CE\*  
 LENGTH: 16.0 (ESTIMATED-LENGTH-REQUIRED: 45.6)

BC06F-16 2 FROM SLOT 0 IN PANEL 2 OF CABINET 1 TO A LA36-CE\*  
 LENGTH: 16.0 (ESTIMATED-LENGTH-REQUIRED: 47.3)

O17-00087-03 1 FROM THE CPU CABINET TO CPU EXPANSION CABINET 1  
 LENGTH: 1.6 (ESTIMATED-LENGTH-REQUIRED: 1.6)

O17-00087-03 2 FROM THE CPU CABINET TO CPU EXPANSION CABINET 1  
 LENGTH: 1.6 (ESTIMATED-LENGTH-REQUIRED: 1.6)

O17-00087-03 3 FROM THE CPU CABINET TO CPU EXPANSION CABINET 1  
 LENGTH: 1.6 (ESTIMATED-LENGTH-REQUIRED: 1.6)

O17-00087-03 4 FROM THE CPU CABINET TO CPU EXPANSION CABINET 1  
 LENGTH: 1.6 (ESTIMATED-LENGTH-REQUIRED: 1.6)

O17-00087-03 6 FROM THE CPU CABINET TO CPU EXPANSION CABINET 1  
 LENGTH: 1.6 (ESTIMATED-LENGTH-REQUIRED: 1.6)

O17-00087-03 6 FROM THE CPU CABINET TO CPU EXPANSION CABINET 1  
 LENGTH: 1.6 (ESTIMATED-LENGTH-REQUIRED: 1.6)

BC08S-16 1 FROM DZ11-D\* 2 IN BOX 2 OF CABINET 1 TO PANEL 2  
 LENGTH: 16.0 (ESTIMATED-LENGTH-REQUIRED: 4.0)

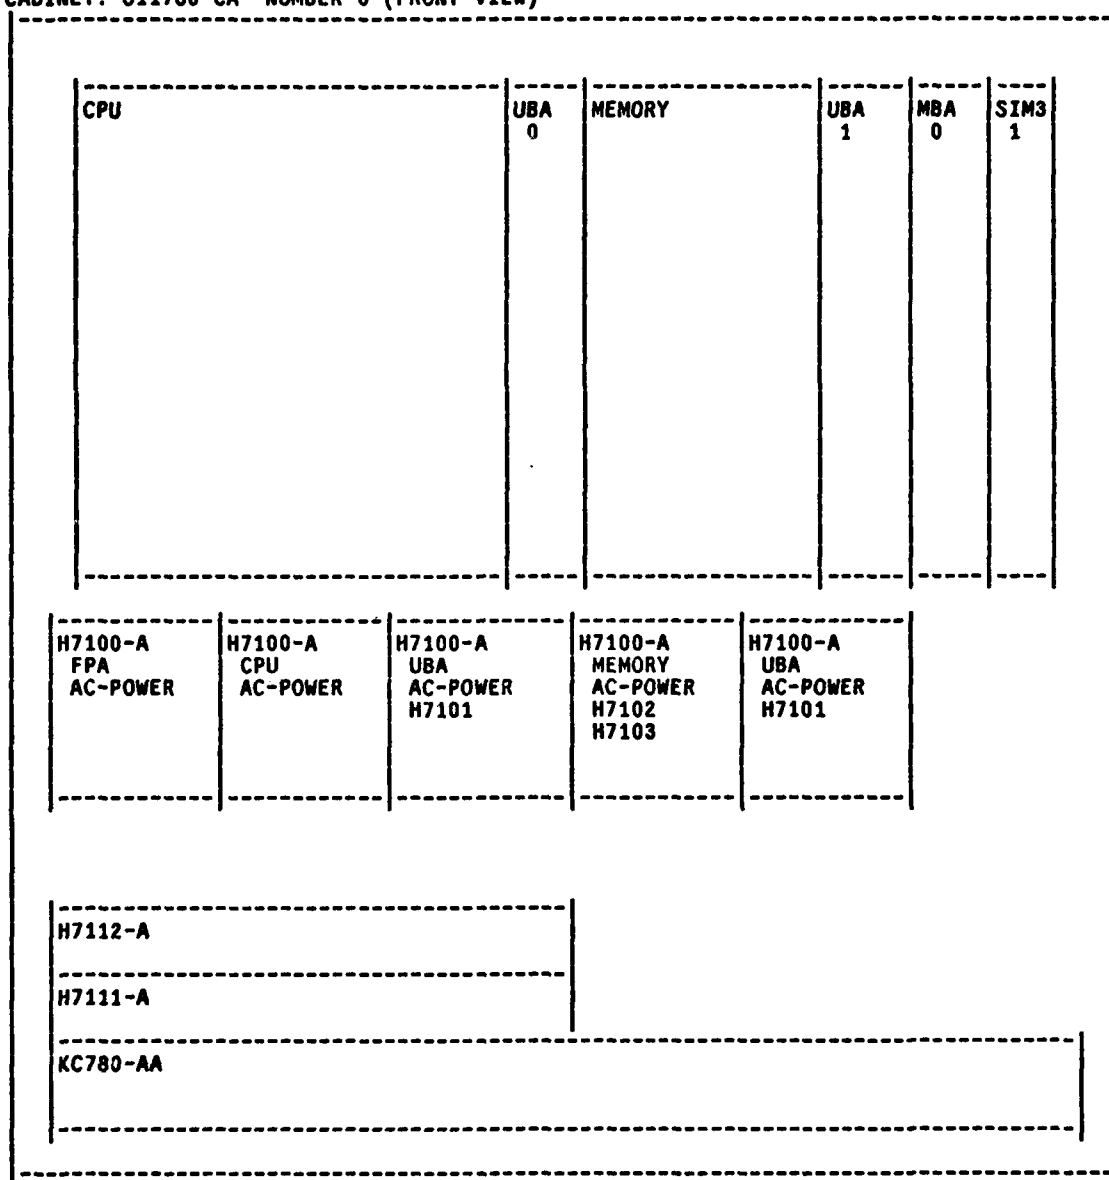
BC08S-16 2 FROM DZ11-D\* 1 IN BOX 2 OF CABINET 1 TO PANEL 2  
 LENGTH: 16.0 (ESTIMATED-LENGTH-REQUIRED: 4.0)

BC06W-16 1 FROM DZ11-B\* 2 IN BOX 1 OF CABINET 1 TO PANEL 1  
 LENGTH: 16.0 (ESTIMATED-LENGTH-REQUIRED: 4.0)

BC06W-16 2 FROM DZ11-B\* 1 IN BOX 1 OF CABINET 1 TO PANEL 1  
 LENGTH: 16.0 (ESTIMATED-LENGTH-REQUIRED: 4.0)



## CABINET: 011780-CA\* NUMBER 0 (FRONT VIEW)



## CPU (BACKPLANE: 070-13628-00)

KA780-AA\* ((29 M8236) (23 M8236) (22 M8234) (20 M8233) (16 M8232) (16 M8231)  
 (14 M8230) (13 M8229) (12 M8228) (11 M8227) (10 M8226) (9 M8225) (8 M8224)  
 (7 M8223) (6 M8222) (5 M8221) (4 M8220) (3 M8219) (2 M8218) (1 M8237))

## SIMULATOR-MODULES ((18 70-14103))

FP780-AA\* ((28 M8289) (27 M8288) (26 M8287) (25 M8286) (24 M8285))

## UBA 0 (BACKPLANE: 070-13626-00)

DW780 ((6 070-14103-00) (5 070-14103-00) (4 M8273) (3 M8272) (2 M8271) (1 M8270))

## MEMORY (BACKPLANE: 070-13626-00)

MS780-CA\* ((20 M8214) (19 M8213) (18 M8212))  
 MS780-DC ((17 M8210) (16 M8210) (15 M8210) (14 M8210))  
 MS780-DC ((13 M8210) (12 M8210) (11 M8210) (10 M8210))  
 MS780-DC ((9 M8210) (8 M8210) (7 M8210) (6 M8210))  
 MS780-DA ((5 M8210))

## SIMULATOR-MODULES ((4 70-14103) (3 70-14103) (2 70-14103))

## UBA 1 (BACKPLANE: 070-13626-00)

DW780 ((6 070-14103-00) (5 070-14103-00) (4 M8273) (3 M8272) (2 M8271) (1 M8270))


## MBA 0 (BACKPLANE: 070-13627-00)

RH780 ((6 070-14103-00 M9041) (5 070-14103-00) (4 M8278) (3 M8277) (2 M8276) (1 M8275))

## SIM3 1 (BACKPLANE: NONE)

OPTION-SIMULATOR ((74-18976))

## CABINET: H9602-HA\* NUMBER 1 (FRONT VIEW)



MEMORY (BACKPLANE: 070-13626-00)  
 MS780-CA\* ((20 M8214) (19 M8213) (18 M8212))  
 MS780-DC ((17 M8210) (16 M8210) (15 M8210) (14 M8210))  
 MS780-DC ((13 M8210) (12 M8210) (11 M8210) (10 M8210))  
 MS780-DC ((9 M8210) (8 M8210) (7 M8210) (6 M8210))  
 MS780-DA ((6 M8210))  
 SIMULATOR-MODULES ((4 70-14103) (3 70-14103) (2 70-14103))

MBA 1 (BACKPLANE: 070-13627-00)  
 RH780 ((6 070-14103-00 M9041) (5 070-14103-00) (4 M8278) (3 M8277) (2 M8276) (1 M8275))

END (BACKPLANE: NONE)  
 SBI-TERMINATOR (M9043 (74-18973))

SIM3 2 (BACKPLANE: NONE)  
 OPTION-SIMULATOR ((74-18975))

## CABINET: H9602-DF NUMBER 1 (FRONT VIEW)

H317-F (BACK)  
MUX-PANEL 2  
DZ11-D\* 1 IN BOX 2  
DZ11-D\* 2 IN BOX 2

BA11-KE\*  
BOX 2  
UBA 1

H317-E (BACK)  
MUX-PANEL 1  
DZ11-B\* 1 IN BOX 1  
DZ11-B\* 2 IN BOX 1

BA11-KE\*  
BOX 1  
UBA 0

BOX: BA11-KE\* NUMBER 1 IN CABINET NUMBER 1 (USA 0)  
MODULES

BACKPLANES

	A	B	C	D	E	F	
DR11 1	1	M9014 XXXXXXXXXXXX	M7219 XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX				070-08329-00
	2	M9680 XXXXXXXXXXXX	M208 XX   M208 XX   M7821 X   M796 XX				
	3	M9882 X   M206 XX   M611 XX   M611 XX   M112 XX   M113 XX					
	4	M9202 XXXXXXXXXXXX	M957 XX   M957 XX   M116 XX   M239 XX				
	5	////////////////////					
	6	UNIBUS-IN XXXXXXXX		G727 XX			070-11628
DZ11-B* 1	7	M7819 XX					
DZ11-B* 2	8	M7819 XX					
	9	M9302 XXXXXXXXXXXX		G727 XX			
	10	////////////////////					
	11	////////////////////					
	12	////////////////////					
	13	////////////////////					
	14	////////////////////					
	15	////////////////////					
	16	////////////////////					
	17	////////////////////					
	18	////////////////////					
	19	////////////////////					
	20	////////////////////					
	21	////////////////////					
	22	////////////////////					
	23	////////////////////					
	24	////////////////////					

BOX: BA11-KE\* NUMBER 2 IN CABINET NUMBER 1 (UBA 1)  
MODULES

BACKPLANES

	A	B	C	D	E	F	
	1	M9014 XXXXXXXXXXXX		G727 XX			070-12412-00
	2			G727 XX			
	3			G727 XX			
RK611* 1	4	M7904 XXX					
	5	M7903 XXX					
	6	M7902 XXX					
	7	M7901 XXX					
	8	M7900 XXX					
	9	M9202 XXXXXXXXXXXX					
	10	////////////////////////////////////					
	11	UNIBUS-IN XXXXXXXX		G727 XX			070-11164
DZ11-D* 1	12	M7814 XXX					
DZ11-D* 2	13	M7814 XXX					
LP11 1	14		M7258 XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX				
	15			G727 XX			
	16			G727 XX			
	17			G727 XX			
	18			G727 XX			
	19	M9302 XXXXXXXXXXXX		G727 XX			
	20	////////////////////////////////////					
	21	////////////////////////////////////					
	22	////////////////////////////////////					
	23	////////////////////////////////////					
	24	////////////////////////////////////					

### APPENDIX 3

The rules shown here are the OPS4 equivalents of the rules in Figures 2-4 and 2-5. It will be apparent, even to the most casual peruser, that my English renderings of the rules do not capture all of the information that the rules contain. In order to read the OPS4 rules, one must know that a symbol prefixed with "=" is a variable; it matches any single (subelement within a) working memory element. If a variable occurs multiple times within the conditional part of a rule, all occurrences must match equal objects. The symbol "=" by itself (ie, unprefixed) is a "don't care" symbol. A symbol prefixed with "#" is a not-variable (so #x will match any symbol except the symbol bound to =x). A symbol prefixed with ">" is a greater-than-or-equal-to variable, and a symbol prefixed with "<" is a less-than-or-equal-to variable. A variable which is preceded by "!" is a segment variable; it must be the last symbol in a list and matches everything up to the end of (the corresponding list in) the working memory element. The pattern-and, "&", specifies that the structures to its immediate left and right must match the same working memory element. A condition can be negated with "-"; in order for a production with a negated condition to be satisfied, there must be no element in working memory that matches it.

#### ASSIGN-POWER-SUPPLY-0

```
( (CONTEXT-ACTIVE =CABINET ASSIGN-POWER-SUPPLY =TIME)
  -(CONTEXT-ACTIVE = = >TIME & #TIME)
  (COMPONENT =SBI-MODULE STCI (CONFIGURED =TYPE SBI-MODULE =))
  (COMPONENT =SBI-MODULE NEXUS =NEXUS)
  -(COMPONENT = STCI (NOT-CONFIGURED AC-POWER POWER-SUPPLY =))
  (TEMPLATE =CABINET POWER-SUPPLY (= (= =NEXUS ! =)))
  (RESULT-GLOBAL VOLTAGE-AND-FREQUENCY =VF)
  -->
  (<FIND-TOKEN> POWER-SUPPLY (STCI2 AC-POWER VOLTAGE-AND-FREQUENCY =VF SUPPORTED YES)
    (STCI COMPONENT-LIST SUPPORTED)
    (ORIGIN (ADDED-BY-RULE ASSIGN-POWER-SUPPLY-0))) )
```

#### ASSIGN-POWER-SUPPLY-1

```
( (CONTEXT-ACTIVE =CABINET ASSIGN-POWER-SUPPLY =TIME)
  -(CONTEXT-ACTIVE = = >TIME & #TIME)
  (LIST-INCOMPLETE =CABINET CABINET ! =LIST) & =L
  (COMPONENT =SBI-MODULE STCI (CONFIGURED =TYPE SBI-MODULE =))
  (COMPONENT =SBI-MODULE NEXUS =NEXUS) & =C1
  (COMPONENT =POWER STCI (NOT-CONFIGURED AC-POWER POWER-SUPPLY =ITEM)) & =C2
  (TEMPLATE =CABINET POWER-SUPPLY (= (= =NEXUS =CO-ORDINATES =))) & =T
  -->
  (<DELETE> =L)
  (<ADD> (LIST-INCOMPLETE =CABINET CABINET ! =LIST
    ((=ITEM) =CO-ORDINATES ((=TYPE) (AC-POWER))))))
  (<DELETE> =T) (<DELETE> =C1)
  (<REPLACE> 4 =C2 (CONFIGURED AC-POWER POWER-SUPPLY =ITEM)) )
```

## ASSIGN-POWER-SUPPLY-6

```
( (CONTEXT-ACTIVE =CABINET ASSIGN-POWER-SUPPLY =TIME)
-(CONTEXT-ACTIVE = = >TIME & #TIME)
(LIST-INCOMPLETE =CABINET CABINET I =LIST)
(COMPONENT =UBA STCI (CONFIGURED UBA SBI-MODULE =))
(COMPONENT =UBA NEXUS =NEXUS)
(COMPONENT =POWER STCI (NOT-CONFIGURED AC-POWER POWER-SUPPLY =ITEM))
(TEMPLATE =CABINET POWER-SUPPLY (= (= =NEXUS =CO-ORDINATES =)))
-(COMPONENT =REGULATOR STCI (NOT-CONFIGURED REGULATOR POWER-SUPPLY H7101))
-->
(<GENERATE-TOKENS> (H7101) (STCI COMPONENT-LIST)
(ORIGIN (ADDED-BY-RULE ASSIGN-POWER-SUPPLY-6))) )
```

## ASSIGN-POWER-SUPPLY-7

```
( (CONTEXT-ACTIVE =CABINET ASSIGN-POWER-SUPPLY =TIME)
-(CONTEXT-ACTIVE = = >TIME & #TIME)
(LIST-INCOMPLETE =CABINET CABINET I =LIST) & =L
(COMPONENT =UBA STCI (CONFIGURED UBA SBI-MODULE =))
(COMPONENT =UBA NEXUS =NEXUS) & =C1
(COMPONENT =POWER STCI (NOT-CONFIGURED AC-POWER POWER-SUPPLY =ITEM)) & =C2
(TEMPLATE =CABINET POWER-SUPPLY (= (= =NEXUS =CO-ORDINATES =))) & =T
(COMPONENT =REGULATOR STCI (NOT-CONFIGURED REGULATOR POWER-SUPPLY H7101)) & =C3
-->
(<DELETE> =L)
(<ADD> (LIST-INCOMPLETE =CABINET CABINET I =LIST
((=ITEM) =CO-ORDINATES ((UBA) (AC-POWER) (H7101))))))
(<DELETE> =T) (<DELETE> =C1)
(<REPLACE> 4 =C2 (CONFIGURED AC-POWER POWER-SUPPLY =ITEM))
(<REPLACE> 4 =C3 (CONFIGURED REGULATOR POWER-SUPPLY H7101)) )
```

## CHECK-VOLTAGE-AND-FREQUENCY-1

```
( (CONTEXT-ACTIVE SYSTEM CHECK-VOLTAGE-AND-FREQUENCY =TIME)
-(CONTEXT-ACTIVE = = >TIME & #TIME)
(COMPONENT =X1 VOLTAGE-AND-FREQUENCY =VF1)
(COMPONENT =X2 & #X1 VOLTAGE-AND-FREQUENCY =VF2 & #VF1)
-(COMPONENT =X1 VOLTAGE-AND-FREQUENCY =VF2)
-(COMPONENT =X2 VOLTAGE-AND-FREQUENCY =VF1)
-->
(<ADD> (CONTEXT-ACTIVE SYSTEM FIX-VOLTAGE-MISMATCH (<BIND>))) )
```

## VERIFY-SBI-AND-MB-DEVICE-ADEQUACY-3

```
( (CONTEXT-ACTIVE SYSTEM VERIFY-SBI-AND-MB-DEVICE-ADEQUACY =TIME)
-(CONTEXT-ACTIVE = = >TIME & #TIME)
(COMPONENT =CONTROLLER STCI (NOT-CONFIGURED MEMORY-CONTROLLER SBI-MODULE =ITEM)) & =C
(RESULT-GLOBAL MEMORY-CONTROLLER =NUMBER & (@GREATER 2)) & =D
-->
(<REPLACE> 3 =D (@- =NUMBER 1))
(<REPLACE> 4 =C (NOT-SUPPORTED MEMORY-CONTROLLER SBI-MODULE =ITEM))
(<ADD> (COMPONENT =CONTROLLER RATIONALE
(SINCE ONLY 2 MEMORY CONTROLLERS ARE PERMITTED PER SYSTEM))) )
```

## ASSIGN-UB-MODULES-EXCEPT-THOSE-CONNECTING-TO-PANELS-4

```
( (CONTEXT-ACTIVE SYSTEM ASSIGN-UB-MODULES-EXCEPT-THOSE-CONNECTING-TO-PANELS =TIME)
-(CONTEXT-ACTIVE = = >TIME & #TIME)
(COMPONENT =MODULE1 STCI (= DISK-DRIVE UB-MODULE =ITEM1))
(COMPONENT =MODULE1 NUMBER-OF-UB-DEVICES-SUPPORTED =NUMBER)
(COMPONENT =MODULE2 & #MODULE1 STCI (= DISK-DRIVE UB-MODULE =ITEM2))
(COMPONENT =MODULE2 NUMBER-OF-UB-DEVICES-SUPPORTED =NUMBER)
(COMPONENT =DEVICE STCI (= DISK-DRIVE UB-DEVICE =))
(COMPONENT =DEVICE PORTS 2)
(COMPONENT =DEVICE UB-MODULE-REQUIRED =ITEM1) & =C1
(COMPONENT =DEVICE UB-MODULE-REQUIRED =ITEM2) & =C2
-(ASSIGNMENT-PERMANENT = UB-DEVICE-TO-UB-MODULE =MODULE1 =)
-(ASSIGNMENT-PERMANENT = UB-DEVICE-TO-UB-MODULE =MODULE2 =)
-(ASSIGNMENT-PERMANENT =DEVICE UB-DEVICE-TO-UB-MODULE ! =)
-->
(<DELETE> =C1) (<DELETE> =C2)
(<ADD> (RESULT-GLOBAL DUAL-UB-MODULES =MODULE1 =MODULE2))
(<ADD> (RESULT-LOCAL NUMBER-OF-DEVICES-ON-MODULE =MODULE1 1))
(<ADD> (RESULT-LOCAL NUMBER-OF-DEVICES-ON-MODULE =MODULE2 1))
(<ADD> (ASSIGNMENT-PERMANENT =DEVICE UB-DEVICE-TO-UB-MODULE =MODULE1 0))
(<ADD> (ASSIGNMENT-PERMANENT =DEVICE UB-DEVICE-TO-UB-MODULE =MODULE2 0)) )
```

## SELECT-BOX-AND-UB-MODULE-FOR-NEXT-SU-2

```
( (CONTEXT-ACTIVE SYSTEM SELECT-BOX-AND-UB-MODULE-FOR-NEXT-SU =TIME)
-(CONTEXT-ACTIVE = = >TIME & #TIME)
(ASSIGNMENT-PERMANENT =MODULE POSITION-IN-OPTIMAL-SEQUENCE =POSITION)
-(ASSIGNMENT-PERMANENT = POSITION-IN-OPTIMAL-SEQUENCE <POSITION & #POSITION)
(COMPONENT =MODULE STCI (PENDING = UB-MODULE =ITEM))
(COMPONENT =MODULE SUS (= =SUS))
(TEMPLATE =BOX SUS-REMAINING >SUS)
-(RESULT-LOCAL NEXT-UB-MODULE ! =)
(RESULT-LOCAL BACKPLANE-BOX-UNUMBER-UPOSITION-AND-BPNUMBER
NOT-ASSIGNED =BOX =UNUMBER = =BPNUMBER)
-(RESULT-LOCAL BACKPLANE-BOX-UNUMBER-UPOSITION-AND-BPNUMBER
NOT-ASSIGNED #BOX = = <BPNUMBER & #BPNUMBER)
-->
(<ADD> (RESULT-LOCAL NEXT-UB-MODULE =MODULE OPTIMAL =BOX)) )
```

## PUT-UB-MODULE-8

```
( (CONTEXT-ACTIVE =BOX PUT-UB-MODULE =TIME)
-(CONTEXT-ACTIVE = = >TIME & #TIME)
(RESULT-LOCAL NEXT-UB-MODULE =NAME = =BOX)
(COMPONENT =NAME STCI (PENDING MULTIPLEXER-TERMINAL-INTERFACE UB-MODULE =ITEM))
(COMPONENT =NAME SLOTS-REQUIRED (=SLOTS-REQUIRED =BACKPLANE-TYPE ! =))
(TEMPLATE =BACKPLANE SLOTS-AVAILABLE (=ID >SLOTS-REQUIRED =BACKPLANE-TYPE =))
(TEMPLATE =BACKPLANE CURRENT-UB-LOAD =CURRENT-LOAD)
(TEMPLATE =BOX BACKPLANE-POSITION =POSITION)
(ASSIGNMENT-PERMANENT =NAME POSITION-IN-OPTIMAL-SEQUENCE =)
-(RESULT-LOCAL POSSIBLE-MULTIPLEXER-AND-BOX-POSITION =NAME ! =)
-->
(<ADD> (CONTEXT-ACTIVE =BOX ASSOCIATE-MULTIPLEXER-WITH-PANEL-SLOT (<BIND>))) )
```

## CHECK-FOR-UB-JUMPER-CHANGES-6

```
( (CONTEXT-ACTIVE =BOX CHECK-FOR-UB-JUMPER-CHANGES =TIME)
-(CONTEXT-ACTIVE = = >TIME & #TIME)
(TEMPLATE =BOX HALF 2)
(ASSIGNMENT-TEMPORARY =BOX BOX-TO-UBX-CABINET = = 2 =UNUMBER =UPOSITION)
(ASSIGNMENT-PENDING = BOX-TO-UBX-CABINET = = =UNUMBER >UPOSITION & #UPOSITION)
(RESULT-LOCAL JUMPER-CABLE =BOX =ITEM & (@NOTANY BC11A-10) =SLOT =LATERAL) & =D1
-(RESULT-LOCAL JUMPER-CABLE =BOX = >SLOT & #SLOT =)
(LIST-TENTATIVE =BOX BOX ! =LIST)
(COMPONENT =NAME STCI (TEMPORARILY-CONFIGURED UB-JUMPER CABLE =ITEM)) & =C1
(COMPONENT = STCI (NOT-CONFIGURED UB-JUMPER CABLE BC11A-10)) & =C2
(RESULT-GLOBAL UB-LENGTH =UNUMBER =LENGTH) & =D2
-->
(<REPLACE> 4 =D1 BC11A-10)
(<REPLACE> 4 =D2 (@+ 10.0 =LENGTH))
(<REPLACE> 4 =C1 (NOT-CONFIGURED UB-JUMPER CABLE =ITEM))
(<REPLACE> 4 =C2 (TEMPORARILY-CONFIGURED UB-JUMPER CABLE BC11A-10)) )
```