




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 15/177		A1	(11) International Publication Number: WO 96/28784
			(43) International Publication Date: 19 September 1996 (19.09.96)
(21) International Application Number: PCT/US96/03406 (22) International Filing Date: 13 March 1996 (13.03.96) (30) Priority Data: 403,418 13 March 1995 (13.03.95) US (71) Applicant: TRILOGY DEVELOPMENT GROUP [US/US]; Suite 130, 6034 West Courtyard Drive, Austin, TX 78730 (US). (72) Inventors: GHATATE, Bhalchandra; 803 Single Oak Cove, Austin, TX 78746 (US). LIEMANDT, Joseph; 4711 Spicewood Springs #244, Austin, TX 78759 (US). PRICE, Andrew; 1227 Creekford Circle, Sugar Land, TX 77478 (US). (74) Agents: HECKER, Gary, A. et al.; Hecker & Harriman, Suite 1600, 2029 Century Park East, Los Angeles, CA 90067 (US).		(81) Designated States: AU, BR, CA, JP, KR, NO, NZ, European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE). Published <i>With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>	

(54) Title: FLASH CONFIGURATION CACHE

(57) Abstract

The present invention provides a constraint-based configuration (10) system using structural hierarchy. The structural aspects of the model provide the ability to define a model element as being contained in, or by, another model element. To configure a system, the present invention accepts input in the form of requests (202), such as an expression of a need for a desktop computer system to be used in a CAD environment. Using this information, the present invention configures a system by identifying the resource (210) and component needs, constraints imposed on or by the resources or components identified, and structural aspects of the system. In one embodiment, a flash configuration cache (10) is utilized, and in another embodiment a bundling cache is used to speed up the process of configuring an end product, such as a user computer.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AM	Armenia	GB	United Kingdom	MW	Malawi
AT	Austria	GE	Georgia	MX	Mexico
AU	Australia	GN	Guinea	NE	Niger
BB	Barbados	GR	Greece	NL	Netherlands
BE	Belgium	HU	Hungary	NO	Norway
BF	Burkina Faso	IE	Ireland	NZ	New Zealand
BG	Bulgaria	IT	Italy	PL	Poland
BJ	Benin	JP	Japan	PT	Portugal
BR	Brazil	KE	Kenya	RO	Romania
BY	Belarus	KG	Kyrgyzstan	RU	Russian Federation
CA	Canada	KP	Democratic People's Republic of Korea	SD	Sudan
CF	Central African Republic	KR	Republic of Korea	SE	Sweden
CG	Congo	KZ	Kazakhstan	SG	Singapore
CH	Switzerland	LI	Liechtenstein	SI	Slovenia
CI	Côte d'Ivoire	LK	Sri Lanka	SK	Slovakia
CM	Cameroon	LR	Liberia	SN	Senegal
CN	China	LT	Lithuania	SZ	Swaziland
CS	Czechoslovakia	LU	Luxembourg	TD	Chad
CZ	Czech Republic	LV	Latvia	TG	Togo
DE	Germany	MC	Monaco	TJ	Tajikistan
DK	Denmark	MD	Republic of Moldova	TT	Trinidad and Tobago
EE	Estonia	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	UG	Uganda
FI	Finland	MN	Mongolia	US	United States of America
FR	France	MR	Mauritania	UZ	Uzbekistan
GA	Gabon			VN	Viet Nam

FLASH CONFIGURATION CACHE

BACKGROUND OF THE INVENTION

This is a continuation in part of United States patent application entitled "Method and Apparatus for Configuring Systems," Serial number 08/039,949, filed March 29, 1993, and assigned to the assignee of the present application.

1. FIELD OF THE INVENTION

This invention relates to the field of computer-based system configuration.

2. BACKGROUND ART

Configuring a system refers to the process of selecting and connecting components to satisfy a particular need or request. If a system is based on a limited number of components, the process of configuring the system can be relatively straightforward. For example, the purchase of an automobile requires a salesperson to configure a system (automobile and assorted options) to satisfy a customer's request. After selecting from a plurality of models, the salesperson completes the transaction by selecting options to configure and price an automobile. The configuring of such a simple system can be accomplished with a pencil and paper.

As system specifications become more customized and varied, configuration alternatives increase and the task of configuring a system becomes more complex. This increased complexity has resulted in a need for

computer-based assistance with the configuration process. Early computer-based systems expand independently-generated configuration orders for systems into manufacturing orders. They do not address the actual need for computer-based tools prior to the order expansion. That is, they do not
5 address the actual generation of a system configuration based on needs and/or request input.

An example of a complex system is a desktop computer system. The available configuration alternatives of a computer system are numerous and
10 varied, including alternatives available when choosing the microprocessor, motherboard, monitor, video controller, memory chips, power supply, storage devices, storage device controllers, modems, and software.

Configuring a desktop computer system requires that a selected
15 component is compatible with the other components in the configured system. For example, a power supply must be sufficient to supply power to all of the components of the system. In addition, the monitor must be compatible with the video controller (e.g., resolution), and the storage device must be compatible with its controller (e.g., SCSI interface). A motherboard
20 must have enough slots to handle all of the boards installed in the system.

The physical constraints of the cabinet that houses the system's components are also considered. The cabinet has a fixed number of bays available for storage devices (e.g., floppy disk drives, hard disk drives, or tape
25 backup units). These bays have additional attributes that further define their use. For example, the bay may be located in the front of the cabinet and provide access from the front of the cabinet. Another bay may be located

behind the front-accessible bays, and be limited to devices that do not need to be accessed (e.g., hard disk drive). Bays may be full-height or half-height. Before a storage device can be added to the configuration, a configuration system must identify a bay into which the storage device will be housed. This
5 requires that at least the accessibility and height of the storage device must be examined to determine compatibility with an available cabinet bay.

The connection between a storage device and its controller must be determined based on the location of each. The cable that connects the storage
10 device and its controller must provide compatible physical interfaces (e.g., 24-pin male to a 24-pin female).

A method of establishing a communication pathway in a computer system is known as daisy chaining. Daisy chaining provides the ability to
15 interconnect components such that the signal passes through one component to the next. Determining whether a daisy chain may be established requires that the available logical (e.g., IDE or SCSI) and physical interfaces (e.g., 24-pin) of all elements in a daisy chain be known. In addition, it is important to know whether conversions from the source datatype to the destination
20 datatype are allowed. When a daisy chaining candidate is added to the system, the interconnections and conversions between existing components may be checked to determine whether the new component should be an element of the daisy chain.

25 The power supply and storage device component examples illustrate the need to define the structural interrelationships between components (i.e., physical and spatial relationships). To further illustrate this notion, consider

placing components requiring electrical power such as computer, telecommunication, medical or consumer electronic components into two cabinets. Further, each cabinet has an associated power supply that supplies electrical power to the components inside the associated cabinet. To account
5 for electrical power consumption and the requirement that no power supply is overloaded, the model must comprehend the specific cabinet in which each component is placed and update the consumed power for each cabinet. While the total power available in the two cabinets may be sufficient for all of the components to be placed in both of the cabinets, a component cannot be
10 included in a cabinet if its inclusion would cause the cabinet's power supply to overload. Therefore, the physical placement of the component in a cabinet must be known to make a determination if the subsequent placement of a component is valid. Similarly, any physical connections between these components must be taken into account. Each component's position in the
15 structural hierarchy is used to determine minimal or optimal lengths for the connecting components.

Early computer-based configuration systems employed an approach referred to as the rule-based approach. Rule-based configuration systems
20 define rules (i.e., "if A, then B") to validate a selection of configuration alternatives. Digital Equipment Corporation's system, called R1/XCON (described in McDermott, John, "R1: A Rule-Based Configurer of Computer Systems", Artificial Intelligence 19, (1982), pp. 39-88) is an example of a rule-based configuration system. R1/XCON evaluates an existing independently-
25 generated system order and identifies any required modifications to the system to satisfy the model's configuration rules. The rules used to perform the configuration and validation processes are numerous, interwoven, and

interdependent. Before any modification can be made to these rules, the spider's web created by these rules must be understood. Any changes to these rules must be made by an individual that is experienced and knowledgeable regarding the effect that any modifications will have to the entire set of rules.

5 Therefore, it is difficult and time-consuming to maintain these rules.

A possible solution to the problems associated with rule-based systems is a constraint-based system. A constraint-based system places constraints on the use of a component in a configuration. For example, a hard disk drive
10 cannot be added to the configuration unless a compatible storage device controller is available for use by the request storage device. The requirement of a controller is a "constraint" on the hard disk drive.

While existing constraint-based systems address some of the
15 shortcomings of rule-based systems, they do not provide a complete configuration tool. Pure constraint-solving systems do not employ a generative approach to configuration (i.e., they do not generate a system configuration based on needs, component requests, and/or resource requests). Existing constraint-based systems use a functional hierarchy that does not
20 address structural aspects associated with the physical placement of a component in a configuration (e.g., memory chip on motherboard or memory expansion board, storage device in cabinet bay, or controller in motherboard slot).

25 Bennett et al., United States Letters Patent No. 4, 591, 983 provides an example of a constraint-based system that employs a recognition or verification approach to system configuration instead of a generative

approach. That is, Bennett merely validates an independently-configured system. In essence, an order is generated by an independent source such as a salesperson, and Bennett is used to verify that the system contained in the order does not violate any constraints. Bennett does not generate a system
5 configuration based on needs or component requests (i.e., a generative approach). Thus, Bennett does not provide the capability to interactively configure a system by interactively selecting its components.

A model consists of all of the elements that may be included in a
10 configured system. In Bennett, the model elements are grouped into an aggregation hierarchy. An aggregation hierarchy creates hierarchical levels that represent a group of elements. Branches from one entry in the current level expand the entry, and the entry is "composed of" the elements in the lower level branches. For example, a desktop computer system is "composed
15 of" a keyboard, a monitor, and a system box. A system box is "composed of" a power supply, motherboard, cards, and storage devices. The "composed of" relationship merely describes the elements that comprise another element. However, the "composed of" relationship does not define the structural relationships between the model elements. The "composed of" relationship
20 does not describe the physical, structural relationships among the elements such as "physically contained inside," "physically subordinate part of," and "physically connected to." Using the desktop computer system previously described, it cannot be determined whether or not a monitor is "physically contained inside" a desktop computer system. A system box is "composed of"
25 storage devices, however it cannot be determined whether one or more of the storage devices are "physically contained inside" the system box.

A functional hierarchy organizes the components of a model based on the purpose or function performed by the components in the model. Each entry in the hierarchy can be further broken down into more specific functional entries. Thus, an entry's parentage defines its functionality, and
5 progression from one level to the next particularizes the functionality of a hierarchy entry.

As used in current configuration systems, a functional hierarchy does not define the structural interrelationships or the physical and spatial
10 interconnections among elements. A functional hierarchy cannot place a storage device in a cabinet bay, a controller card in a particular slot on the motherboard, or a memory chip in a slot on the memory expansion board.

Figure 2 illustrates an example of a functional hierarchy.

15 HardwareComponent 30 is the root element of the hierarchy. The next level below HardwareComponent 30 (i.e., the second level 49) identifies general functions in the model. For example, ROM 31, Processor Unit 31, Processor 32, Memory 34, Cage 35, Board 36, Connector 37, and Storage Device 38 all perform the function of Hardware Component 30 in addition to their own
20 specialized functions. Processor 33 can be specialized to the function of a SpecialPurpose 40 or GeneralPurpose 41. SpecialPurpose 40 can be specialized to ArithmeticProcessor 51.

Referring to Figure 2, it can be seen that a functional hierarchy does
25 not provide the ability to define the structural aspects of the system. For example, there is no capability to determine the contents of Cage 35. The physical and spatial location of MotherBoardSlot 54 descending from Slot 46,

in turn descending from Connector 37 cannot be determined from the functional hierarchy. There is no way of determining that MotherBoardSlot 54 is contained by the motherboard. It is not clear from the functional hierarchy definition whether ArithmeticProcessor 51 is located on the
5 MotherBoard 44 or another model element. It cannot be determined whether MemoryChip 42 and ROM 31 are located on MotherBoard 44, MemoryBoard 52, or another model element.

A functional hierarchy does not provide the ability to define actual
10 interconnections between configured instances or the data transfer. That is, that one component is connected to another with compatible logical datatypes (e.g., serial interface) and compatible physical interconnections (e.g., 24 pin). A functional hierarchy only defines the function that a component performs.

15

Because it does not define the actual connections between the components selected for a configuration, it cannot establish a daisy chain between configured components . Referring to Figure 2, a functional hierarchy defines Connector 37, Storage Device Controller 53, Floppy Drive
20 48, and Hard Drive 49 as types of components. To conserve resources, a user may wish to configure a system such that an occurrence of Floppy Drive 48 is daisy chained to an occurrence of Storage Device Controller 53 through Hard Drive 49. However, the functional hierarchy can only reflect that fact that a configured system may contain the functionality provided by Storage Device
25 Controller 53, Hard Drive 49, and Floppy Drive 48. It cannot reflect the fact that an occurrence of Floppy Drive 48 is connected to an occurrence of Storage Device Controller 53 through an occurrence of Hard Drive 49.

Therefore, a functional hierarchy can not traverse a connection pathway to identify structural interrelationships among configured instances. Thus, a functional hierarchy cannot establish a daisy chain. Therefore, a
5 functional hierarchy can not provide the ability to daisy chain components.

Another example of a constraint-based system using a functional hierarchy is provided in the following articles: Mittal and Frayman, "Towards a Generic Model of the Configuration Task," in Proceedings of the
10 Ninth IJCAI (IJCAI-89), pp. 1395-1401; and Frayman and Mittal, "COSSACK: A Constraints-Based Expert System for Configuration Tasks," in Sriram and Adey, Knowledge-Based Expert Systems in Engineering: Planning and Design, September 1987, pp. 143-66.

15 The Cossack system employs a functional hierarchy-based configuration system. According to Cossack, a system using a functional hierarchy must identify a configured system's required functions. Once the required functions are identified, Cossack must identify some particular component, or components, that are crucial, or key, to the implementation of
20 these required functions. The Cossack representation does not make structure explicit. Further, Cossack does not provide mechanisms for reasoning about or with structural information. Therefore, Cossack cannot make any structure-based inferences. For example, the internal data transfer paths within components are not represented. Therefore, there is no ability
25 to trace data transfer within a component, and no ability to establish a data connection with another element.

A configuration system, whether used to configure a computer system or other system, should provide a tool to interactively: define and maintain a model; define and maintain (i.e., upgrade) a configured system; generate marketing bundles; generate a graphic representation of the physical and spatial locations of the components of the configured system; use the graphic representation to modify or upgrade a configured system; and generate configuration reports (e.g., failed requests, quotations, and bill of materials). Such a system must define the components of a system, the structural relationships among the components (i.e., spatial and physical locations), the actual physical and spatial interconnections of the components, and the constraints imposed by each component.

SUMMARY OF THE INVENTION

The present invention employs a generative approach for configuring systems such that a system may be configured based on component or
5 resource requests, or input in the form of need. The present invention provides a constraint-based configuration system using a functional hierarchy that comprehends hierarchical and non-hierarchical structure, and associated constraints that can reason about and generate structural relationships. The structural aspects of the model provide the ability to
10 define a model element as being contained in, or by, another model element. In addition, the structural model provides the ability to identify logical datatype and physical interconnections between elements and establish connections between elements.

15 To configure a system, the present invention accepts input in the form of requests (e.g., component or resource) or needs, such as an expression of a need for a desktop computer system to be used in a CAD (i.e., computer-aided design) environment. Using this information, the present invention configures a system by identifying the resource and component needs,
20 constraints imposed on or by the resources or components identified, and the structural aspects of the system.

The system configuration can be based on a general definition of a system (i.e., desktop computer system to operate in a CAD environment), or
25 at any level of increased specificity (e.g., disk drive by manufacturer and model number). The system configuration can be based on specific component requests (e.g., laser printer), or by need (e.g., printing capability).

Once the system is configured, the configured system can be bundled into products, and a quote can be generated. The bundling process may include the specification of heuristics to control the product-to-component mapping. For example, the product that covers the largest number of components can
5 be selected over other possible product selections that cover a lesser amount of components.

The functional, structural hierarchy of the present invention provides the ability to define the structure of the configuration model and the systems
10 configured from the model. The structural hierarchy includes a container structure. A container provides the ability to specify that one component is contained by, or in, another component. Thus, it is possible, for example, to identify that a component request for a disk drive cannot be satisfied because there are no empty cabinet bays in the cabinet specified to contain the
15 component requested.

The structure hierarchy notion provides the ability to pool resources. Explicit representation of structure, specifically hierarchical structure, provides the ability to define and access inherited resources. For example,
20 computer, telecommunication, medical, or consumer electronic components can be placed in a cabinet that provides power to those components. These individual components can inherit the electrical power resource from a structural superior (i.e., a hierarchical entry that resides one or more levels above the components in the model hierarchy). Further, the structural
25 superior can pool resources and provide an homogeneous resource to its structural inferiors (i.e., a hierarchical entry tht resides one or more levels below the structural superior in the model hierarchy). For example, a cabinet

might contain more than one electrical power source, however, the resource is presented to structurally inferior components as a single resource pool.

Thus, if a component requires a particular resource, this resource can be supplied by a resource pool. For example, if a desktop computer system's cabinet contains multiple power supplies, a disk drive component may draw from resource pool without any knowledge that the resource need is satisfied by multiple power sources.

In addition, the structural specification provides the ability to specify the connections between components of a configured system. As components are added to a configuration, the physical and logical interconnections that are required to assemble the system components may be verified. For example, before adding a printer with a serial logical connection and a 24 pin physical connection to the configuration, a serial port must be available in the configured system. In addition, a physical connection must be made between the printer and a serial port. If the serial port is a 9-pin female physical connection and the printer has a 24-pin female connection, a cable must be available to physically connect the printer and the serial port. In addition, the actual connection is created in the configuration and can be examined in subsequent connection processing . Connection processing provided the ability to identify any criteria for satisfying a connection request. For example, connection criteria may include the cheapest, longest, or optimal throughput connection.

Connection processing may also be used to optimize the use of the configured system's resources. For example, a controller's resources can be optimized by daisy chaining other components together. By connecting one

component to another via multiple intervening components, multiple components may be connected to a single component via a single port or connection.

5 In the present invention, a modeling language is used to define a model hierarchy. The model hierarchy is structural and functional. The modeling language provides the ability to define a Product Base that may be grouped into Product Lines. The structural hierarchy model includes the Component, Composite, Container, Port, and Connector base classes. These
10 base classes may branch into derived classes (i.e., system-specific classes) and terminate at leaf-descendants. Leaf-descendants define the type of components in the functional, structural hierarchy model. Attributes, datatypes, resources, and constraints further define the model.

15 A model language provides the format for defining the elements, the constraints placed on the elements, and the structure of the model. The model language may be used directly, or generated based on input from an interactive model maintenance system used to facilitate the creation and maintenance of the model.

20 The maintenance system graphically displays the model, and provides the interface for the selection of model elements to be updated. Once the desired updates have been made, the maintenance system provides the ability to test the new model, or verify that the new model can be successfully
25 compiled.

Once a model has been successfully defined, the present invention provides the ability to configure a system using the functional, structural hierarchical model. An interactive interface provides the ability to express a configuration in terms of a model element (i.e., components) request,
5 resource request, and/or needs (i.e., requirements) request. A configuration engine is invoked to satisfy these requests.

The configuration engine accesses the Product Base to satisfy the requests in a defined priority. A request is processed by adding components
10 to the configuration, or identifying existing components that can satisfy the request. Further, the interconnections, data transfer pathways, and dynamically-determined structural relationships are defined. When a request is successfully processed, the configuration modifications are "committed." Failed requests are reported.

15

A graphical depiction illustrates the configured system and its structural characteristics. The elements of the configured system are illustrated in terms of their physical and spatial location relative to other elements. Elements are contained in other elements, comprised of other
20 elements, or connected to each other. This graphical depiction further provides an interface to modify and maintain elements of the configured system.

The configured system's elements are bundled into available
25 marketing and manufacturing packages for system quotation and manufacturing purposes. The bundling process performs a product-component mapping based on product definitions.

In one embodiment, a flash configuration cache is utilized to speed up the process of configuring a user computer. This is performed by taking advantage of the fact that the invention uses a structured set of requests to
5 configure the user computer. The host computer first makes an initial determination as to which set of requests take greater time to configure than the time taken to recall the resulting configuration from the host computer cache. Those requests are stored in the cache and are arranged in the form of a tree structure. When a new set of requests is obtained, the sets of old requests
10 in the request tree are methodically searched to find a matching set of old requests. The configuration corresponding to the matching set of old requests is then recalled from the cache.

In other embodiments, the invention's flash configuration cache is
15 used to speed up the process of configuring a variety of end products. The end products are, for example, electronic systems such as voice mail systems, PBX systems, central office switches, and handheld communication devices. The present invention's flash configuration is also used to configure end products such as airplanes where a variety of power system options, landing
20 system options, and interior system options need by configured in an efficient and thorough manner. Other end products configured by the flash configuration cache of the invention are trucks, test equipment, and chemical processes. The flash configuration cache is also used to configure vacation packages where each package involves a number of transportation options,
25 lodging options, and recreational options.

In another embodiment, a bundling cache is used to speed up the process of bundling, namely, the process of mapping components required for a user computer configuration, or other end product configuration, into actual commercial products.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a block diagram of the configuration computer system.

5 Figure 2 illustrates a functional hierarchy.

Figure 3 illustrates the functional, structural hierarchy comprised of the five base classes, derived classes, and component types.

10 Figure 4 is the functional, structural hierarchy for a model to configure computer systems.

Figure 5 illustrates component interconnections with multiple intervening components and data types.

15

Figure 6 illustrates the Configuration Engine process flow.

Figure 7 illustrates the SatisfyResourceRequest process flow.

20 Figure 8 illustrates the SatisfyContainerConstraint and SatisfyComponentConstraint process flow.

Figure 9A illustrates the SatisfyConnectionConstraint process flow.

25 Figure 9B illustrates the CandidatePorts process flow.

Figure 10 illustrates the EstablishSetCover process flow.

Figure 11 illustrates a system window for a desktop computer system configuration.

5 Figure 12 is a flow diagram illustrating the functional operation of the Configuration System.

Figures 13A-13C illustrate a flow diagram of the algorithm used to construct the search tree in the flash configuration cache.

10

Figure 14 is illustrates an example of how the search tree of the flash configuration cache is constructed.

Figures 15A-15C illustrate a flow diagram of the algorithm used to
15 search the search tree of the flash configuration cache.

Figure 16 is a diagram showing products mapped into available options.

20 Figure 17 is a diagram showing available options mapped into products.

DETAILED DESCRIPTION OF THE INVENTION

A method and apparatus for configuring systems is described. In the following description, numerous specific details are set forth in order to
5 provide a more thorough description of the present invention. It will be apparent, however, to one skilled in the art, that the present invention may be practiced without these specific details. In other instances, well-known features have not been described in detail so as not to obscure the invention.

10 The present invention provides a tool for configuring systems that has application to a wide range of domains including the following: computer hardware, computer software, computer networks, telecommunication systems (e.g., PBX and voice mail), copiers, medical imaging systems, vehicles (e.g., fire trucks and construction equipment), electronic control systems,
15 buildings, modular furniture, manufacturing equipment, manufacturing systems, consumer electronic equipment, and electronic systems.

Figure 1 is block diagram of the configuration system of this invention. The configuration system 10 is comprised of the Model Maintenance
20 Subsystem 12, the Configuration Generation and Reporting Subsystem 14, the Bundling/Quotation Subsystem, Communications Bus 18, Input/Output 20, Memory 22, Central Processing Unit 24, and Mass Storage 26.

Figure 12 is flow diagram illustrating the functional operation of the
25 Configuration System. At block 600, a model base is read. The Configuration System uses a model base that contains information about all of the elements

available to configure a system (e.g., components and resources) This model base is referred to as a Product Base.

- A model language is used to create the Product Base. The model
- 5 language provides the syntax, or statements, used to define the model elements, the constraints placed on the model elements, and the structure of the model. At processing block 604, the model definition can be entered using the model language and model definition processing is ended at 606.
- 10 Model Maintenance - The process of defining a model can be facilitated if the Model Maintenance Subsystem is chosen at decision block 602 (i.e., "use Model Maintenance Subsystem?"). At block 608, the model, either new or existing, is displayed. At block 610, the model can be edited. The Model Maintenance Subsystem 12 provides the ability to test the validity of and
- 15 debug the modified model at decision block 612 (i.e., "write integrity, ProductBase integrity, or Debugger?"). A "write integrity" selection determines the integrity of the parse file (i.e., subsets of the Product Base) with the addition of the modifications. a "ProductBase integrity" selection determines the integrity of the Product Base with the addition of the
- 20 modifications.

- If the "Debugger" is chosen, benchmark system configuration requests are read from a file at block 618. At block 14, the Configuration Generation and Report System 14 is invoked to configure a system using the modified
- 25 model and the benchmark configuration requests. A trace of the processing of these requests by the Configuration Generation and Reporting System 14 may be made to examine the configuration process.

If there are additional modifications to the model at decision block 622 (i.e., "modify model?"), a graphic representation of the model is displayed at 608, and the modification process continues at block 610. If there are no other
5 modifications, the model definition is generated at block 624, and the Model Maintenance Subsystem ends at block 606.

Configuration and Reporting System - The Configuration and Reporting System 14 uses the model definition to generate a system
10 configured according to the user-specified requests and needs. The resulting configuration is graphically depicted. Reports are generated to provide information regarding the configuration. If it is determined that an existing configuration is being upgraded at decision block 630 (i.e., "upgrading existing system?"), the existing system is read and its elements marked as existing in
15 block 632. If a new system is being configured, a blank system instance is created at block 634. The forms used to input element requests or needs is displayed at 636. If input is not complete at decision block 638 (i.e., "requests completed?"), processing continues at block 636.

20 Configuration Engine - Once all of the request and need input is completed, ConfigurationEngine is invoked to generate a configured system based on the input at 640. A graphical representation of the configuration is displayed at 642. The configuration may be modified, reports may be generated, or the components of the configuration may be bundled and a
25 quotation generated. If modifications are intended at decision block 644 (i.e., "configuration modification?"), processing continues at decision block 652 (i.e., "filter model?"). If a filtered model is chosen at decision block 652, a

subset of the model is generated at block 654. The model subset includes those model elements that can be selected given the current configuration. Processing continues at 636 to display input forms. If a filtered model is not used, processing continues at 636.

5

After a system is configured, the elements of the configuration can be bundled into marketing, or manufacturing, products. Bundler 660 maps the configuration components to products. Quoter 662 generates a cost quotation for the configured system. At 664, the quotation is displayed. If there are no
10 configuration modifications at decision block 666 (i.e., "configuration modification?"), processing ends at 606. If there are modifications to the configuration, the Configuration Generation and Reporting Subsystem 14 is invoked at block 668.

STRUCTURAL HIERARCHY

The Configuration System of the present invention is a constraint-based scheme using a functional, structural hierarchy. Figure 3 illustrates the functional, structural hierarchy and five intrinsic base classes. The functional, structural hierarchy contains a class hierarchy comprised of five intrinsic base classes 70 that define the basic types of model objects. These five base classes are: Component 60, Composite 62, Connector 64, Container 66, and Port 68. The Component 62 is the base class from which all other classes and component types are derived. From Component 62, each branch of the hierarchical tree begins with an intrinsic base class and branches into system-specific classes called derived classes 88. Derived classes 88 are definitions of broad component categories, such as storage devices, power supplies, and peripheral cards. Multiple generations of derived classes can descend from the base classes. Each branch terminates with "leaf descendants," or Component Types 90. Component Types 90, represent actual components that can be instantiated and configured.

The Composite class 62 is a static structure (i.e., elements that have substructure). Elements in this class have, or are, subcomponents of a composition. The Connector class 64 branches from the Composite class 62. This class defines the model elements that connect elements. An element in the Container class 66 indicates that the element may contain other elements. Elements in the Port class 68 provide the port alternatives and define a port's datatype. Elements derived from the Port class 68 can be physically connected with other components derived from the Port class 68.

The present invention provides the ability to represent within a structural hierarchy how components of a particular system exist spatially and physically. Within the structural hierarchy, there are three type of substructures: composite hierarchies, container hierarchies, and port relationships. Composite hierarchies identify components as part of other components. For example, a chassis has eight card slots. Container hierarchies identify components as being contained in other components. A Container hierarchy is a dynamic structure in that the structure is dynamically created when a configuration is generated. For example, a CPU card is placed in slot 0 of the chassis). Port relationships identify components that connect to other components. A connection, or port, relationship is dynamically created when a configuration is generated. The relationships between generations within these substructures are expressed by the keywords "childOf," "containedBy," and "connectsWith."

15

The "childOf" keyword indicates that a component is part of a component that is descended from class Composite. The "containedBy" keyword indicates that a component is contained within a component that is descended from the Container base class. The "connectsWith" keyword indicates that a component connects to a component that is descended from the Port Class.

20

Container hierarchies typically exhibit an alternating relationship with Composite hierarchies. That is, a container is often a "childOf" a composite component, and the composite component is "containedBy" another container. Each substructure type has a root member that is also a descendant of the base class of the same name (i.e., Composite, Container, or Port).

25

Members of a substructure can be of any class defined in the Class Hierarchy. For example, a component of class bay, descended from Container Class might contain a component of class storage-device (descended from Component Class) or of class card_chassis (descended from Container Class).

5

Figure 4 illustrates a structural hierarchy with the five base classes, derived classes, leaf descendants, and substructure relationships. The structural relationships further define the structural aspects of the model. For example, Slot 114 is a "childOf" Cabinet 110. Therefore, Slot 110 is a subcomponent of the composite component, Cabinet 110. Further, Cabinet 110 is a "ChildOf" System 116. Second occurrences of Card 118 (i.e., 118A) and Slot (i.e., 114A) illustrate the substructural relationship between Card and Slot. Card 118A is "containedBy" Slot 114A. Similarly, StorageDevice 120A is "containedBy" Bay 122A, and DB25MaleDeviceOut 124A "connectsWith" DB25FemaleDeviceOut 126.

The structural aspects of the present inventions's model provides the ability to inherit and pool resources. For example, a container component, Cabinet, may consist of a chassis and two one-hundred watt power supplies, A and B. Each of the elements within the chassis container consume, or require some amount of power. If the chassis component contains two central processing units (CPUs) that together consume one-hundred and ten watts (e.g., fifty-five watts each), random access memory that consumes seventy watts, and multiple cards (e.g., controllers) that consume a total of twenty watts, neither of the power supplies independent of the other could supply sufficient power to the chassis and its elements.

However, because the two power supplies are contained in, and are a part of, the Cabinet container, the two power supplies can be pooled together to supply the elements within Cabinet. Therefore, when the resource requisitions are being processed for the elements in this example, one or the other may be used to satisfy the request. In addition, it is possible to satisfy the resource need for any one of the elements by using both power supplies. For example, if one CPU's resource needs are processed first using fifty-five watts of power supply A, and the resource processing for the RAM is processed next, the resource needs of the RAM maynot be satisfied by power supply A alone. However, it is possible to satisfy the RAM's resource needs by using 45 watts from power supply A and twenty-five from power supply B. Another resource that may use this resource pooling capability is a heat dissipation resource.

15 CONTAINERS

The structural hierarchy provides the ability to structure the model such that one model element, or group of model elements, may be contained by another. The use of the contained model element in a configuration will be constrained by the availability of a container model element in the configuration.

Figure 8 illustrates the SatisfyContainerConstraint and SatisfyComponentConstraint process flow. At decision block 500 (i.e., "required instance already available in configuration?"), if the required instance exists and is available to satisfy the constraint, the constraint is satisfied by this available instance and processing returns at block 526. If not,

the required instance is instantiated, and the Modifications List is updated at processing block 502. At decision block 504 (i.e., "any constraints to be processed?"), if there are no constraints on the new instance, the constraint is satisfied by the new instance, and processing returns at block 526.

5

If there are constraints to be processed, the next constraint is identified at block 508. If it is determined that it is a requiresContainer constraint at decision block 510 (i.e., "requiresContainer?"), processing continues at processing block 512 (i.e., "satisfyContainerConstraint") to satisfy the
10 requiresContainer constraint, and processing continues at decision block 522 (i.e., "constraint satisfied?").

If it is determined that it is not a requiresContainer constraint at decision block 510, but it is determined that it is a requiresConnection
15 constraint at decision block 514 (i.e., "requiresConnection?"), processing continues at processing block 516 (i.e., "satisfyConnectionConstraint") to satisfy the requiresConnection constraint, and processing continues at decision block 522 (i.e., "constraint satisfied?").

20 If it is not a requiresContainer constraint at decision block 510 and not a requiresConnection constraint at decision block 514 (i.e., "requiresConnection?"), processing continues at decision block 518 (i.e., "requiresComponent?"). If it is determined that it is a requiresComponent constraint at decision block 518 (i.e., "requiresComponent?"), processing
25 continues at processing block 520 (i.e., "satisfyComponentConstraint") to satisfy the requiresComponent constraint, and processing continues at decision block 522 (i.e., "constraint satisfied?"). At decision block 522 (i.e.,

"constraint satisfied?"), if the constraint was satisfied, processing continues at decision block 504 (i.e., "any constraints to be processed?"). If the constraint was not satisfied, the constraint is marked as not being satisfied by an existing instance or the new instance, and the new instance is removed from the

5 Modifications List at processing block 524. Processing returns at block 526.

CONNECTION PROCESSING

The use of a model element in a configuration may also be constrained

10 by the ability to establish a connection to another model element. The requiresConnection constraint requires that a physical connection exist between two components. Figure 9A illustrates the process flow for satisfying the requiresConnection constraint. At processing block 280, a target component is selected and a list of ports is created. At processing block 282,

15 the requested resources are allocated. At processing block 284, CandidatePorts(list) is invoked to identify unconnected ports that are accessible from the target component. At processing block 286, candidate local ports (i.e., those ports that are unconnected and have the appropriate datatype) are identified. At processing block 288, candidate connectors are

20 identified.

At decision block 290 (i.e., have all connectors been tested?), if all of the connectors have been tested, the request is marked as failed, and processing continues at block 306 (i.e., "return"). If not, the next connector is

25 selected at block 294. At decision block 296 (i.e., "can physical type of connector's port1 connect with physical type of target port?"), if port1 of the connector is not the same physical type (e.g., 25 pin) as the target port's

physical type, processing continues at decision block 290 (i.e., "have all connectors been tested?").

Otherwise, processing continues at decision block 298. At decision
5 block 298 (i.e., "can physical type of connector's port2 connect with physical
type of local port?"), if port2 of the connector is not the same physical type
(e.g., 25 pin) as the local port's physical type, processing continues at decision
block 290 (i.e., "have all connectors been tested?"). Otherwise, processing
continues at decision block 300. At decision block 300 (i.e., "does a transfer
10 path exist between port1 and port2?"), if a transfer path does not exist between
port1 and port2, processing continues at decision block 290 (i.e., "have all
connectors been tested?"). Otherwise, the requested resource is allocated at
block 302. At processing block 304, the target port is connected to the
connector's port2, and the local port is connected to the connector's port1.
15 Processing ends at block 306.

Candidate ports must be identified to satisfy a requiresConnection
constraint. Figure 9B illustrates the CandidatePorts(list) process flow.
Processing block 310 of CandidatePorts(list) set thePort variable to the next
20 port in the list. At decision block 312 (i.e., "is the port connected?"), if the
port is connected, processing continues at processing block 316. If not,
decision block 314 (i.e., "thePort the right datatype or are conversions
allowed?") determines if the datatypes are compatible. If not, processing
continues to block 310 and the next port is found.

25

If they are compatible, thePort is added to the port list, and processing
continues at block 310. If it is determined that thePort is already connected at

decision block 312, processing continues at processing block 316, and newPort is set to the port to which thePort is connected. At block 320, a new port list is created for all ports to which newPort transfers. At decision block 322 (i.e., "does newList contain a port of the requesting component?"), if the newList
5 contains one of the requesting component's ports, the connection is marked as already being in existence at block 326 and processing returns at block 328. If not, CandidatePorts(list) is invoked for the newList.

CONFIGURATION ENGINE

10

When the user has selected the components for the system to be modeled, the user requests the invocation of the configuration engine. The configurator accesses the Product Base to identify the object class. After certain validation checks are successfully performed, the configurator
15 instantiates (i.e., creates) a member of that class, called an object instance. The configurator only instantiates those objects required to configure the requested system.

The configuration engine processes component and resource requests
20 in the priority specified. As each request is processed, the existing configuration is modified by: (1) adding the requested component and other components required to support the component requested, or (2) identifying existing components and new components required to provide the requested resource. When a request is successfully processed, the configuration
25 modifications are "committed," and this configuration becomes the input configuration in processing the next request.

Figure 6 illustrates the Configuration Engine process flow. Processing block 202 creates a prioritized list of requests. If it is determined that all of the requests have been processed at decision block 204 (i.e., "all requests processed?"), processing ends at block 206. If not, the next request is selected
5 at processing block 208.

The request type is determined at decision block 210 (i.e., "request type?"). If the request is a component request, processing continues at processing block 212. At block 212, the component requested is instantiated
10 and posted to the Modifications List, and processing continues at decision block 216. If the request is a resource request, the component that can supply this resource is identified at processing block 214 (i.e., "SatisfyResourceRequest"), and processing continues at decision block 216. At decision block 216 (i.e., "Instantiation or allocation successful?"), if the
15 component instantiation or resource allocation is successful, processing continues at decision block 224 (i.e., "any constraints to be processed?"). If the component instantiation or resource allocation is not successful, processing continues at decision block 218 (i.e., "do any other alternatives exist to satisfy this request?").

20

If it is determined at decision block 218 (i.e., "do any other alternatives exist to satisfy this request?") that no other alternatives exist to satisfy the request, the request is identified as a failed request, and processing continues at decision block 204 (i.e., "all requests processed?"). If there are other
25 alternatives, the failed alternative's modifications are removed from the Modifications List at 220, the next alternative is posted to the Modifications

List at 222, and processing continues at decision block 224 (i.e., "any constraints to be processed?").

At decision block 224 (i.e., "any constraints to be processed?"), if there
5 are no constraints to be processed, the modifications are committed to the
configuration at processing block 244, and processing continues at decision
block 204 (i.e., "all requests processed?"). If there are constraints to be
processed, the next constraint is identified at block 226. If it is determined
that it is a requiresContainer constraint at decision block 228 (i.e.,
10 "requiresContainer?"), processing continues at processing block 230 (i.e.,
"satisfyContainerConstraint") to satisfy the requiresContainer constraint, and
processing continues at decision block 240 (i.e., "constraint satisfied?"). If it is
determined that it is not a requiresContainer constraint at decision block 228
but it is determined that it is a requiresConnection constraint at decision
15 block 236 (i.e., "requiresConnection?"), processing continues at processing
block 232 (i.e., "satisfyConnectionConstraint") to satisfy the
requiresConnection constraint, and processing continues at decision block 240
(i.e., "constraint satisfied?").

20 If it is not a requiresContainer constraint at decision block 228 and not
a requiresConnection constraint at decision block 236 (i.e.,
"requiresConnection?"), processing continues at decision block 238 (i.e.,
requiresComponent?). If it is determined that it is a requiresComponent
constraint at decision block 238 (i.e., "requiresComponent?"), processing
25 continues at processing block 234 (i.e., "satisfyComponentConstraint") to
satisfy the requiresComponent constraint, and processing continues at
decision block 240 (i.e., "constraint satisfied?"). At decision block 240 (i.e.,

"constraint satisfied?"), if the constraint was satisfied, processing continues at decision block 224 (i.e., "any constraints to be processed?"). If the constraint was not satisfied, processing continues at decision block 218 (i.e., "do any other alternatives exist to satisfy the request?").

5

The fact that resources are offered by individual component instances, and are not represented as global system entities, assists in the exploration of alternatives. Figure 7 illustrates the SatisfyResourceRequest process flow. At processing block 250, the next component that offers the required resource is
10 found. If, at decision block 252 (i.e., "any component instances found?"), it is determined that no component offers the resource, processing continues at processing block 262.

If a component is found, processing continues at decision block 254
15 (i.e., "has this resource been consumed?"). If the resource has been consumed processing continues at processing block 250 (i.e., "Find next component that offers the required resource"). If the resource has not been consumed, a check is made to determine whether class requirements and optional requirements are valid at decision block 256. If all of the checks are
20 valid, the current resource instance is chosen at processing block 258, and processing continues at processing block 264. If one of the checks is invalid, processing continues at decision block 260 (i.e., "have all resource instances been checked?"). If all of the resource instances have not be checked, processing continues at block 250 where the next component offering the
25 resource is found.

If all of the components offering this resource have been checked, or it is determined (at decision block 252) that no existing component offers this resource, processing continues at block 262, and a new component instance that offers the resource is created, the configuration modification is posted to the Modifications List, and processing continues at block 264. At block 264, an instance of the requested component type is assigned to the requesting component's returned instance variable. Processing continues at decision block 266 (i.e., "does the current instance satisfy query and test conditions?") to determine if all query and test functions are satisfied. If not, processing continues to processing block 250. If they are, processing ends at block 268.

MODEL LANGUAGE

The model language provides the ability to define a model (e.g., model elements, model constraints, and model structure). Using the syntax of the model language, statements may be entered to define the model base, or Product Base. The Product Base contains all of the information about a model. The Product Base contains the information used to configure a system.

20

The Product Base may also contain Hierarchical Product Lines. Product Lines allow a Product Base to be subdivided into groups. An example of such a grouping is marketing divisions, such as DesktopSystems. A DesktopSystem might contain all of the components that are commonly sold as parts of a desktop computer system such as operating system software, modem cards, microprocessor chips, etc. Only components that are part of the same product line can be configured together. However, each component

type can be part of several product lines. Product Lines hierarchies may also be declared. A child in a product line hierarchy inherits from the parent, and every component in the parent is inherited by the child. The format of a product line declaration is as follows (Note: reserved words are bold, double-underscores indicate repetitive portions, and portions contained in "<<>>" are required):

productLine <<ProductLineName>>;

Or, to declare product line hierarchies:

10

productLine <<ProductLineName1>>; <<ProductLineName2>>;

System models are stored in files, called parse files. Collectively, the parse files are known as the Product Base. Parse files contain information about a general category within a system model. Data representations of individual system parts are known as objects. Cabinets, storage devices and peripheral cards are examples of objects in a Product Base used to configure computer systems. A property provides attributes of an object. For example, in a computer systems' Product Base, capacity, power requirements, and connection interface are properties of a storage device object. Further, a property categorizes an object. That is, objects with similar properties are called a class of objects. Objects can inherit properties from other objects. That is, one class of objects acts as the parent of another class, and the child class exhibits all of the properties of the parent class in addition to others.

25

Attributes define the aspects of a component that must be considered to successfully configure a component. Examples of attributes of a power supply are the cabinet space required for the supply and the remaining power

available after power-consuming components are added to the configuration. Attributes can be assigned at the class level, and descendants of that class inherit the class attributes. In addition, attributes can be associated with particular component types. There is no limit to the number of attributes
 5 that can be assigned to a component or class.

Attribute values may be of type floating point, boolean, string, datatype, component, and resource. Attributes may be multivalued. That is, multivalued attributes can have more than one value. For example, with a
 10 component that can use either a full height internal bay or a front accessible bay, the attribute "attribute_Bay_type_required" can retain both values. An attribute is declared by the statement (Note: "|" indicates a choice):

AttributeType <<Attribute Name>>; |
 15 Multivalued AttributeType <<AttributeName>>;

An example of attribute declarations are:

20	Float	Position
	Float	throughput_available
	Float	load_consumed
	resource	space_type_required

A resource is a system commodity that is associated with component
 25 types. A resource may be assigned to multiple component types. Multiple resources may be assigned to a component. When a component is instantiated, the resource assigned to this component type is made available to the configuration. When a component's resource is consumed, only the

resource supplied by its associated component becomes unavailable. The availability of a resource of the same type that is offered by a second component is unaffected by the consumption of the first component's resource. Therefore, if the same resource type is available from a second
5 component, the consumption of the first component's resource does not consume all of this resource type in the modeled system.

Before a resource type can be assigned to a component type or used by a component instance, the resource type must be declared. A resource
10 declaration has the following format:

```
resource <<ResourceName>>;
```

An example of a resource declaration is as follows:

```
15 resource static_RAM_resource;
```

Datatype declarations define the types of interfaces and data transfer protocols available to connections in a modeled system. SCSI and IDE are
20 examples of datatypes. A datatype is declared as follows:

dataType <<DataTypeName>>;

A derived class is defined by the following statement (Note: the portion with the ";" symbol is optional):

```

5      Class <<ClassName>>: <<BaseClassName | SuperClassName>>
      {
      displayStatus: <<HIDDEN | LISTED | DRAWN>>
      ;attributes:
10     <<AttributeName = AttributeValue>>;
      ;dimensions [Xsize, Ysize];
      ;connectionOrigin <<TRUE | FALSE>>;
      }

```

15 The display status includes the values Hidden, Listed, and Drawn. Drawn allows the class member to be displayed in the graphical rendering of the configuration. Listed allows the class members to be listed on the additional components list. Hidden is used for members that are Hidden (i.e., not drawn), but have children that are Drawn. An attribute value may

20 be assigned at the time of declaration, but this is not necessary. Connection origin identifies whether or not instances of this class are to be used as starting points for cabling report generation. An example of a derived class declaration is as follows:

```

25      class Bay: Container
      {
      displayStatus: DRAWN;
      attributes:
30          front_accessible;
          height;
          half_height_compatible;
          position;
      }

```

In this example a derived class, bay, is created. It is a member of the Container base class. Therefore, it may contain other elements. Its attributes define its height, half_height compatibility, front_accessibility (i.e., is a component installed in this bay accessible from the front of a system cabinet), height, and position. These attributes will be inherited by each descendant of this derived class.

System components, or component types, are defined by the following declaration:

```

10      component <<ComponentTypeName>>: <<DerivedClassName>>
      {
      {productLines: <<ProductLineName>>;}
      {label: <<"LabelName">>;}
15      {description: <<"DescriptionString">>;}
      {resource: <<ResourceName i, IntegerValue i>>;}
      {dataType: <<DataTypeName>>;}
      {partNum: <<"PartNumString">>;}
      {subComponents: <<SubcomponentName>> V
20          <<SubcomponentName[Integer]>>;}
      {transfers: <<SubcomponentName[0] <-> SubcomponentName[1]>>;}
      {dimensions: [<<Xsize, Ysize>>];}
      {values: <<AttributeName = AttributeValue>> V
          <<AttributeName = {AttributeValue, ...}>>;}
25      {fillDirection: [ <<TB | BT | LR | RL>> ];}
      }

```

The label field defines the label given to the graphical representation of this component. The description field defines the description that is displayed or reported. The dataType field is used if the component type is descended from a port, and defines the type of data that may be transferred from this component type. The subComponents field defines the structural children of a Composite component type. The transfers field defines the

paths that data can travel through a Composite component. Transfers are a mechanism for expressing an internal data path within a Composite component. For example, a cable is represented as a component with two ports, and the cable is used to transfer data from one port to another. The values field provides the ability to establish a component's attributes, or properties. The fillDirection describes the order in which multiple components in a single container are drawn.

The following is an example of a component definition:

```

10      Component Cabinet1 : Cabinet
      {
      partNum: "001-001";
      Children: Slot1_1;
15          Slot1_2;
          Slot1_3;
          ...
          Slot1_9;
          Slot1_10;
20      CabinetBay {4};
      Values:
          position = 1;
          resources_provided = {10_Slot_Resource, CPU_Slot_Resource,
25      MCU_Slot_Resource, Mem_Slot_Resource, Bay_Resource};
      }

```

This example defines a component type, Cabinet1, within Cabinet and Composite classes. Figure 4 is the structural hierarchy for a model used to configure computer systems. Cabinet1 108 is descended from Cabinet 110 which is a descendant of Composite 112. Therefore, Cabinet1 108 is a composite component type. It has subcomponents, or children, Slot1_1 through Slot1_10 and CabinetBay{4}). The integer "4" indicates that there are four CabinetBay component types within Cabinet1.

The following is an example of a Composite component type that descends from a connector:

```
5      Component SCSIChainCable: Cable
      {
        description: "SCSI Chain Cable";
        partNum: "003-002";
        subComponents:
10          SCISICablePort_3;
          SCISICablePort_4;
        values:
          length = 2;
        transfers:
15          SCISICablePort_3 <-> SCISICablePort_4;
      }
```

The following is an example of a component type definition that provides a resource:

```
20      Component 16mbMemCard : Card
      {
        description: "16mb Memory Card";
        partNum: "004-016";
25      resource: Memory_Resource, 16;
        values:
          slot_resource_required = Mem_Slot_Resource;
      }
```

30 Constraints provide conflict resolution information used to determine whether components may be added to the configured system. Constraints can control such things as space allocation, space occlusion, and additional component requirements. Constraints are expressed as component qualifiers and component dependencies. Constraints test the attributes and lineage of

components and identify the components that are required for the successful instantiation of components.

```

5      constraint <<ConstraintName>> on <<ClassName>>
      {
        <<requiresComponent | requiresContainer>>
          (<<ClassName, ResourceName | ClassName | ComponentName>>,
            <<?ReturnedInstance>> , ?ReturnedInstance.AttributeName;
10      , Consumed; , Existing; , New;);
      }
      constraint <<ConstraintName>> on <<ClassName>>
      {
        <<requiresConnection ( ;StartingComponentName,;
          <<ClassName, ResourceName | ClassName | ComponentName>>,
15      <<DataType>>, <<?ReturnedInstance>>, <<%Path>>
          ,?ReturnedInstance.AttributeName;
          , Connector ( <<ClassName>>, <<?ConnectorInstance>>,
                        <<?ConnectorInstance.AttributeName>>);
          , Longest; , Consumed; , Existing; , New; , Conversions;);
20      }

```

The Constraint Name and the Class upon which the constraint may be applied are identified in the first line of the declaration. The requiresComponent, requiresContainer and requiresConnection expression

25 identifies additional items (i.e., components, container, or connections) that are required to configure the constrained component. The additional items needed may be identified by a derived class name and resource combination, a derived class name, or the name of the component type. When a request is satisfied during configuration, the configuration engine returns the instance

30 of the requested component type found. The ?ReturnedInstance variable identifies the variable that is associated to the instance of the requested component type found by the configuration engine. A request may further ask that the configuration engine make a choice based on attribute maximization. That is, make a choice that will maximize a given attribute.

Therefore, a ?ReturnedInstance.AttributeName declaration will return the requested item with the greatest amount of AttributeName. The attribute maximization option can also be an expression that refers to other returned instances created by previous component requests with the current constraint

5 and perform operations with them. A component instance is said to be consumed when it is unavailable to satisfy a constraint requirement. The Consumed keyword can be used to mark an instance returned by a request as unavailable. Once an instance is consumed, the configuration engine will exclude this instance in subsequent searches to satisfy another request. The

10 Existing keyword limits the search to existing instances. The New keyword requests that a new instance be created to satisfy a constraint requirement.

The requiresConnection constraint requirement has additional arguments that describe the requirements for an entire connection path that

15 can contain several different components. The requiresConnection constraint requirement has one requirement that is additional to and different from the requiresComponent and requiresContainer constraints. Like the other two constraint requirements, the requiresConnection requires that the request be satisfied. In addition, the requiresConnection constraint

20 requirement, requires that the constrained instance be connected to the satisfying instance.

The StartingComponentName field, or variable, refers to the starting component in the connection (i.e., where the connection will begin). If this

25 variable is not set, the starting component is assumed to be the constrained instance. The next line (i.e., "<<ClassName, ResourceName | ClassName | ComponentName>>") identifies the connection component.

The type of data that the connection will carry is specified by the `DataType` field. The `dataType` field specifies the data type requirements of a port of the requested instance. Further, the `dataType` field specifies the data
5 type requirements of a port of the constrained instance. Because the `dataType` field only requires that the constrained instance's port and the requested instance's port be of data type `dataType`, a connection constraint can be satisfied by a multiple stage connection. For example, it is possible to connect a SCSI device to a SCSI card through intervening components.

10

Figure 5 illustrates component interconnections with multiple intervening components and data types. `ConstrainedInstance` 161 has port 160 and port 162. Port 162 is connected to `Connector` 179 at Port 163. Port 164 of `Connector Block` 179 is connected to Port 165 of
15 `FirstInterveningComponent` 166. Port 167 of `FirstInterveningComponent` is connected to Port 168 of `Connector` 180. `MultipleInterveningComponents` 183 represents some number of intervening components that may be placed between `FirstInterveningComponent` 166 and `NthInterveningComponent` 173. `Connector` 180 and `Connector` 181 are positioned on either end of the
20 `MultipleInterveningComponents` 183. Port 171 of `Connector` 181 is connected to Port 172 of `NthInterveningComponent` 173. Port 174 is connected to Port 175 of `Connector` 182. Port 176 of `Connector` 182 is connected to Port 177 of `DiskDriveController` 178. Chain 184 represents the chained communication or connection path between `ConstrainedInstance` 161 and
25 `DiskDriveController` 178.

The ?ReturnedInstance and ?ReturnedInstance.AttributeName fields have the same functionality as in the requiresComponent and requiresContainer constraint expression. The %Path variable is bound to all of the instances used to make the connection. That is, all of the instances
5 involved in a connection are referred to as the connection path.

With respect to the ?ReturnedInstance.AttributeName and the ?ReturnedInstance instance variables, the maximization option is the same as for the requiresComponent and requiresContainer constraints. There are
10 two maximization options for the path instance variable. The first option is the connector the option. The ClassName field specifies the desired class of connector instances used to build the path. The ?ConnectorInstance field is bound to the returned connector instance, and the AttributeName is the connector instance attribute to be maximized. The request for
15 ?ConnectorInstance is maximized in the same way as the returned instances for requiresComponent and requiresContainer.

The second maximization option provided by requiresConnection is the path length option. This option provides the ability to prioritize choices
20 among paths from the requested component to the requesting component. The length of a path is defined as the number of component instances in the path, including instances of class Connector. The longest path may be specified by using the "Longest" keyword in the constraint declaration. If the longest path option is not chosen, the configuration engine selects the
25 shortest path.

The Consumed, Existing and New specifications of the requiredConnection constraint have the same functionality as in the requiresComponent and requiresContainer constraint declarations. The Conversions option provides the ability to specify that the requested instance

5 can have a datatype that is dissimilar to the constrained instance. That is, if this option is chosen, the requested-side port is no longer required to carry data of type DataType. The only requirement is that the datatype specified by the dataType variable be available at the requester-side port. This option expands the alternatives that the configuration engine is allowed to consider

10 in satisfying the connection request, since it does not have to choose the terminal component with the same datatype as the requester instance. Therefore, if a connection constraint allows conversions, satisfaction of a request for a SCSI connection need only deliver SCSI data to the requesting instance.

15

The following is an example of a constraint definition:

```

constraint Storage_device_constraint on StorageDevice
{
20  requiresConnection (SCSICard, SCSIDataType, ?card, %path,
      Connector (Cable, ?c, -?c.length, Longest));
  requiresContainer (Bay, Bay_Resource, ?bay.Consumed);
      *
      *
25      *
}

```

The requiresContainer constraint indicates that the StorageDevice component type requires a container (i.e., a bay). In addition, this constraint

30 definition imposes a constraint on the StorageDevice class of the model

hierarchy and all of its descendants. It requires the longest cable component type connection to a SCSI Card component type. The type of data that will be carried by this connection is of datatype SCSI Datatype. A port of the constrained instance must also be of this datatype. The datatype constraints
5 may be fulfilled with a multiple stage connection. Thus, the SCSI StorageDevice may be connected to the SCSI Card through intervening components. The variable ?card identifies the SCSI Card instance used. The %path variable contains information regarding the instances used to make the connection.

10

The model language provides the ability to perform further tests and queries to ensure that the configuration engine returns usable instances or instance sets. If a constraint contains a component request, these queries and tests are placed after that request. If the queries and tests are not satisfied, the
15 configuration engine continues to search for another alternative to satisfy the request. The following are examples of the tests provided in the model language:

	mathematical operators:
20	+ (addition)
	- (subtraction)
	* (multiplication)
	/ (division)
	ABS (absolute value)
25	SQRT (square root)
	relational operators:
	> (greater than)
	< (less than)
	== (equality)
30	>= (greater than or equal to)
	<= (less than or equal to)
	!= (not equal)
	boolean operators:

OR (logical inclusive or)
 AND (logical conjunction)
 NOT (logical negation)
 assignment operator:
 5 := (becomes; takes the value of)

For example, in configuring a computer system, a test may be performed when configuring a floppy disk drive for the computer system. A floppy disk drive requires a bay or slot within the system cabinet. Such a
 10 constraint would be expressed as a requiresContainer component request. This request would cause the configuration engine to search for a candidate instance to satisfy this request. Once the engine returns the candidate instance (i.e., ?bay), further testing can be done to determine whether the drive will fit in the returned instance. This can be tested by comparing the
 15 height attribute values of the candidate instance (i.e., ?bay) and the constrained instance (i.e., ?this) as follows:

?bay.height >= ?this.height

20 Intrinsic functions provide additional capability to perform tests and queries. Intrinsic functions can be grouped into query functions and predicate functions. The following are examples of query functions:

25 ceil Queries an attribute of type float, or any expression that evaluates to a floating point value, for the smallest integer value greater than or equal to the floating point value. Returns an integer.
 Syntax: ceil (<<Expression>>)

30 ClassName Queries a set variable for all instances in the set that belong to the specified class.
 Syntax: ClassName (<<%InstanceSet>>)

5	ComponentName	Queries a set variable for all instances in the set that belong to the specified component type (i.e., leaf class). Syntax: ComponentName (<<%ReturnedInstance>>)
10	Component	Queries a set variable for all instances that are not descended from class Connector. Syntax: Component (<<%InstanceSet>>)
15	component	Queries an instance for the component type (i.e., class hierarchy leaf class) from which it is descended. Returns the parent component type. Syntax: component (<<%ReturnedInstance>>)
20	COUNT	Queries a set variable for all instances in the set that belong to the specified class. Syntax: COUNT (<<ClassName ComponentTypeName>> <<(%InstanceSet)>>)

The following is an example of a constraint definition using query and predicate functionality:

```

25  constraint Storage_device_constraint on Storage_Device
    {
    requiresConnection (SCSICard, SCSIDatatype, ?card, %path,
        Connector (Cable, ?c, -?c.length, Longest);
    requiresContainer (Bay, Bay_Resource, ?bay, Consumed);
    ancestor (?bay, Cabinet) == ancestor (?card, Cabinet);
30  FORALL (?inst1, Storage_Device (CONNECTS(FIRST(%path))));
        ancestor (?inst1, Cabinet) == ancestor (?this, Cabinet));
    }

```

In this example, Storage_Device requires a connection to a component
35 of type SCSICard. The connection must be of datatype SCSIDatatype. The
component instance of type SCSICard is bound to the instance variable ?card,
and the components in the connection path are bound (as a set) to the set

variable %path. The connector component used to complete the connection must be of type Cable, and is bound to the instance variable ?c. Candidate cables are ordered from shortest to longest, and if alternative paths from the SCSI Card instance exist, the longest path (in terms of number of
5 components) is preferred.

This example further indicates that Storage_Device must be placed in a container component of type Bay. This instance of type Bay must supply Bay_Resource. The instance of Bay is bound to instance variable ?bay, and
10 the instance is marked as consumed (i.e., unavailable in subsequent requests for components of type Bay).

In the example, the phrase "ancestor (?bay, Cabinet) == ancestor (?card, Cabinet)" requires that the structural ancestor (of type Cabinet) of the instance
15 identified by ?bay must be the same instance as the structural ancestor (of type Cabinet) of the instance identified by ?card. In other words, the card and the bay must be in the same cabinet.

The "Forall" phrase used in the previous example indicates that all
20 component instances of type Storage_Device connected to the first cable in %path must be in the same cabinet as the constrained instance of Storage_Device.

Constraint relationships may be established either at the component
25 level or at the class level. At the component level, constraint relationships specify which component types are constrained by what constraints. The component designated in the constraint relationship may be any of the

component types that have been defined by a Component Type declaration. The constraint may be a constraint declared by a Constraint declaration. The following is the syntax for specifying a component level constraint:

```

5      <<ComponentTypeName>> constrainedBy <<ConstraintName1>>
      {<<OR | AND>> <<ConstraintName2>>}
      {, <<OR | AND>> <<ConstraintNameN>>};

```

Constraints may also be expressed at the class level. A class-level

10 constraint is evaluated as a conjunct in component-level constraint expressions for all component types derived from the constrained class. When a component-level constraint expression is evaluated, class-level constraints are appended to the beginning of the constraint expression and end with that constraint's request and predicate function expressions. If a

15 component inherits class level constraints from several levels in the Class Hierarchy, the constraints are ordered from the most primitive class (i.e., the root class Component) to the most system-specific class (i.e., the user-defined component type). The syntax for a class-level constraint relationship declaration is as follows:

```

20      constrain class <<ClassName>> with <<ConstraintName>>

```

The present invention provides the ability to represent within a structural hierarchy how components of a particular system exist spatially

25 and physically using three type of substructures: composite hierarchies, container hierarchies, and connection relationships. Composite hierarchies identify components as part of other components. Container hierarchies identify components as being contained in other components. Connection

relationships identify components that connect to other components. The relationships between generations within the structural hierarchy are expressed by the keywords "childOf," "containedBy," and "connectsWith." Structural relationships are declared as follows:

5

```
<<ClassName>> childOf <<ClassName>>  
<<ClassName>> containedBy <<ClassName>>  
<<ClassName>> connectsWith <<ClassName>>
```

10 MODEL MAINTENANCE

A model can be defined by providing statements that syntactically conform to the model language described above. In addition, an interactive facility, the Model Maintenance Subsystem, provides the ability to define, and
15 maintain a model, using a graphical user interface. The Model Maintenance Subsystem provides the ability to interactively define the Product Base using a graphical user interface. The semantic representations, class hierarchies, and structural hierarchies of the model may be viewed (i.e., browsed) and modified (i.e., edited) interactively using a graphical user interface. Further,
20 constraint input is verified. Testing and debugging capabilities are provided to identify problems in the model, and to test and optimize the performance of the modified model. For example, model definition syntax is parsed and verified, and sample requests may be executed. Diagnostics functions may be invoked to monitor the performance of the configuration requests with the
25 modified model.

The browsing capability of the maintenance system provides the ability to view graphic representations of the class and substructural components of

the model hierarchy. A Class Tree is used to represent objects descending from base classes within the model hierarchy (i.e., an object class hierarchy). The object class hierarchy is represented by five separate trees, one for each base class. Each branch may have multiple descendants.

5

A Component Tree is used to depict the Composite, Connector and Container Component substructural interrelationships. Composite Trees are listed first followed by Connector and Container Trees.

10 A hierarchy member may be selected for modification by double-clicking on the box that contains the hierarchy member. An editor window for the selected hierarchy member is displayed. A List menu may also be used to select the member to be edited. In the preferred embodiment, the List
15 menus are a series of pulldown menus that may be selected from a menu bar of the Maintenance window. The initial menu bar contains a selection for each general element of the ProductBase model (i.e., classes, component types, constraints, etc.). Once a general element is chosen, a new window is displayed that lists the model members of the general type selection. A model member may be chosen along with an operation (i.e., Comment,
20 View, New, or Edit). A Comment operation provides the ability to add a comment to the ProductBase after the selected member. A View operation provides the ability to view the settings for the selected model element. The model member may be modified by choosing either a New or Edit operation.

25 For example, to modify an attribute of a model member in the preferred embodiment, the attribute type is chosen from the List Menu. Once the attributes are displayed, a New or Edit operation may be chosen to add a

new attribute, or modify an existing attribute, respectively. An attribute selection must also be made, if the Edit operation is chosen. After these selections have been made, the Attribute Editor window is displayed. The fields of the window (e.g., name, attribute type, and multivalued) are

5 initialized to either blank or the default settings for a New operation, or initialized to the current attribute settings for an Edit operation. The attribute name field may be selected and modified. The type field may be modified by selecting from a list of valid attribute types. The multivalued field may be toggled on or off. After making modifications, the modifications may be

10 saved or cancelled.

Resources and Datatypes may be added or modified in a manner that is similar to the method for adding or modifying an attribute. Model elements that require relational definitions require additional designations. Examples

15 of these are derived classes, product lines (i.e., parent Product Line), constraints (i.e., constrained class), and component types.

In the preferred embodiment, adding a derived class requires an additional initial step to define the location of the new derived class within

20 the model hierarchy. At this point, the New and Edit operations have the same operational characteristics, including the ability to save or cancel. That is, the derived class field values (existing, default, or blank) are displayed in an Editor window. In addition, attributes may be added to all members of the derived classes and their component types; constraints may be specified at the

25 class level for the derived class; structural hierarchy relationships may be defined for the derived class; the System Window display status may be defined; the derived class may be selected as a connection origin (i.e., a

starting point of a cabling report); and the component distance (i.e., the average distance from members of the derived class to other objects that are a part of the same composite, and the distance from the member of the derived class to an external port on the composite) may be defined for children of
5 composite objects that are involved in connections.

To add a new component to the model, the class from which the new class is descended must be chosen. The subcomponent field provides the ability to specify the structural hierarchy (i.e., structural children) of a
10 composite component. The New or Edit operations further provide the ability to specify connectivity fields such as transfers (i.e., paths that data can travel through a Composite component), datatype, connection origin. In addition, the following field information may be specified: component type name, associated attributes, products lines (i.e., product lines that contain this
15 component), leaf-level constraints, resources, description, label, part number, fill direction, and display status.

The Maintenance system further provides the capability to test a modified model. The Write integrity option determines whether a ParseFile
20 (i.e.,) can be parsed, and a modified ParseFile written. The ProductBase Integrity option determines whether a ParseFile (i.e.,) can be parsed, and a modified ParseFile written. If not, syntax error messages are displayed. The Debugger (i.e., Configure) option reads component requests from a request file and attempts to configure those components using selected constraints in
25 the current ParseFile. The Debugger provides a tracer capability to provide constraint tracing. A deep trace generates trace output for a traced constraint

and all the constraints it spawns. A shallow trace generates a trace output for traced constraints.

NEEDS ANALYSIS

5

The process of translating customer requirements into specific components and configurations is called "Needs Analysis." The model language provides the ability to express a model in terms of customer needs and requirements.

10

With a needs analysis approach to modeling, a configuration may also be expressed in terms of capacities (e.g., minimum required response time) or throughput. The needs analysis configuration may be illustrated by a voice messaging system model. A configured voice messaging system may be
15 required to record some specific number of hours of voice data, and provide a response time of less than five seconds for accessing stored messages. To further illustrate, a telecommunications configuration may be specified in terms of traffic load supported and some maximum acceptable failure rate (e.g., dropped calls), or a computer system configuration may be required to
20 support certain processing loads, data storage requirements, and response times.

The model language provides the capability to express a needs analysis model in the configuration modeling language by: (1) interpreting customer
25 requirement quantities (e.g., voice message storage capacity), and (2) identifying associated quantities of configuration components and resources. This provides the ability to make modeling requests in terms of needs in

addition to component requests. Components can be identified as satisfying requirements or needs. That is, components may be identified as supplying some quantity of a resource (e.g., megabytes of storage capacity). When a user expresses a system, or some portion of a system, in terms of needs or
5 requirements, one or more components that satisfy the needs may be selected from the ProductBase.

INPUT FORMS

10 Input forms provide the capability to accept component requests from the user. Input forms allow the user to specify the types and quantities of components in the system to be configured. Input forms consist of standard windowing formats such as listboxes and pushbuttons. A third type of input form provides the ability to specify a quantity of a given component. The user
15 selections on the input forms are called component requests. Input forms provide the ability to associate a default priority for component requests. Default priorities may be overridden by a requestPriority. These priorities provide the ability to designate the order in which component requests are satisfied by the configuration engine.

20

PRODUCT-COMPONENT MAPPING

Product_component mapping defines discrete and composite components as parts and products in a sales inventory, and then maps those
25 parts and products (i.e., bundles) onto a set of all component instances in a configured system. The product-component map contains representations that define each part and product in terms of its required and optional

constituent components. These representations further specify how the products are displayed by the Quoter. A representation is comprised of a the following sections: a Product Header, an Optional Equipment List, and an Option Restriction List.

5

The Product Header section provides the product name as it appears in the ProductBase. This allows the Bundler to match components in a configured system to products and identify a set cover. This section also includes the following additional information: a Product Description String
10 that describes the product for use by other portions of this invention (e.g., the Quoter); a Product Number String; the Price (i.e., the price of the product); Product Lines String identifies the product lines of which the product is a member, and is used to narrow the set covering search; and a Required Components List that identifies components (i.e., by part number) or
15 products (i.e., by product number) that are required by this product.

The Optional Equipment List is a list of additional product packages that can be included in the base package (i.e., the product described in the Product Header). An Optional Equipment List entry contains: an Option
20 Unique ID to uniquely identify the entry; an Option Description that describes the entry; an Additional Cost to identify an additional cost associated with the inclusion of this entry; and a Constituent Product Number List identifies those products or components, by number, that comprise the entry.

25 The Option Restriction List is a list of groups of options that are interdependent or that must be chosen according to special criteria. Each entry in the Option Restriction List contains the following: a Group Unique

ID to uniquely identify the entry, a Quantity Specifier, and an Option Unique ID List. The Quantity Specifier field specifies the number of members of an option group that may or must be chosen. The Quantifier Specifier field may consist of bounds or the atLeastOne, atMostOne, or exactlyOne keywords.

- 5 The bounds are two integers (enclosed in parentheses and separated by a comma) that express the lower and upper bound. The atLeastOne keyword indicates that one member of the option group must be chosen. The atMostOne keyword indicates that only one member of the option group may be chosen, and that it is not required that any member be chosen. The
- 10 exactlyOne keyword indicates that at least one member of the option group must be chosen, but no more than one. The Option Unique ID List is a space-separated list of Option Unique ID's.

- An example of an entry in a product-component map for a model
- 15 configuring computer systems is as follows:

```
product base_system
{
  description: "Base System";
  productNumber: "001-001";
  cost: 10000;
  values:
    category1 = "System";
    category2 = "XXX";
  productLines: Tower;
  required: ("001-001" reference) "002-001" "002-002";
  options:
    COM1 "Comm Option 1" 1 "002-005";
    COM2 "Comm Option 2" 1 "002-006";
  optionGroups:
    g1 atMostOne Com1 Com2;
}
```

BUNDLER

The Bundler bundles components into product (i.e., marketing) packages. The Bundler uses the product-component map to establish a set
5 cover for a configured system. A set cover is a set of many-to-one mappings of component instances in a configured system to product packages in which each component instance is mapped to one product package.

Set covering is the process of covering a set of objects (e.g., component
10 instances in a configuration) with a set of covers (e.g., products). This process is used to associate components created for the current configuration with some grouping or cover (e.g., products). A common problem associated with the set covering process is that as the number of objects and set cover alternatives increase, the number of set covering alternatives explodes
15 exponentially. To limit the set covering alternatives, heuristics may be used to identify the minimum set of covers. The Lowest Cost Cover is an example, of a heuristic. Using this heuristic, covering is maximized and cost is minimized. That is, the products providing the most cover for the least amount of cost are selected.

20

Another heuristic is based on the structural context of the alternatives. That is, in some instances, a product will have structure, and that structure will define a physical unit or grouping of components. This may occur, for instance, when a reduction in manufacturing cost is incurred when
25 components are produced as a unit. This savings may be passed on to the purchaser of a system where the reduced-cost unit is actually being purchased. Therefore, it is necessary to examine the configured components

to determine their structure context, and then match these attributes with the structure context of the products. An example of this is a disk array in a computer configuration model. The disk array is physically configured, or manufactured, with a chassis, power supply, controller and five disk drives.

- 5 Therefore, it is necessary to examine the structure context of any disk drive component requests. The process of selecting instances as "covered" by the disk array product must include a determination that the "covered" instances were configured to be inside the chassis, or as a disk array unit.

- 10 Figure 10 illustrates the EstablishSetCover process flow. At processing block 450, the products that can cover some or all of the component instances in the current configuration are identified. At decision block 452 (i.e., "any products identified?"), if no products have been identified, processing ends at block 454. If products were identified, the products are prioritized based on
- 15 the number of instances that can be covered by the product at processing block 456. At decision block 458 (i.e., "any instances not covered?"), if all of the instances have been mapped to the current prioritized product list, a new product list is created that covers products in the current configuration at block 474, and processing continues at decision block 452 (i.e., "any products
- 20 identified?").

- If not, the next product is selected from the list at block 460. At decision block 462 (i.e., "do all required elements exist?"), if all of the elements of the product do not exist in the configured system, processing
- 25 continues at processing block 460. If they do exist, the instances that have not been previously mapped and that can be covered by the current product are identified at processing block 464. At decision block 466 (i.e., "any instances

identified?"), if no instances can be covered by the product processing continues at decision block 458 (i.e., "any instances not covered?").

If some instances were identified, it is determined whether any
5 product option restrictions can not be met at decision block 468 (i.e., "any product option restrictions that are not met?"). If there are, processing continues at decision block 458 (i.e., "any instances not covered?"). If not, processing continues at decision block 470 (i.e., "all structural contexts satisfied?"). If they are not, processing continues at block 460 and the next
10 product is obtained. If they are, the mapped component instances are marked as covered by the current product at block 472 and processing continues at decision block 458 (i.e., "any instances not covered?").

REPRESENTATION OF MODELED SYSTEM

15

Once a system has been configured based on the requests made, various reporting tools are employed to provide information regarding the configured system. In the preferred embodiment, these tools include a graphical depiction of the general layout of the system, a list of materials, a
20 list of spare parts, and a list of any component requests that could not be satisfied.

The present invention provides the ability to express a model in structural terms. That is, components are defined in terms of their structural
25 parents (i.e., containers), interconnections, and compositions. Therefore, the present invention has the ability to graphically display the configured system along with its structural characteristics.

The graphical depiction of the configured system and its structural characteristics, called the system window, provides a depiction of the general layout of the configured system. In the preferred embodiment, the system window for a model that configures computer systems shows the interior and front of all cabinets used in the system, and shows the placement of cards, power supplies, and storage devices. Figure 11 illustrates a system window for a desktop computer system configuration. System Window 540 illustrates the configured system's components and their relative locations within the system. Chassis 550 contains System Board 552, DriveCage 554 and Power Supply 556. Main Board 552A is a detailed depiction of System Board 552.

Main Board 552A illustrates the physical placement of other components on the system board and their relative positions. For example, EVGA Video Board 558 is located below CPU Board 560. Further, the placement of Network Card 562 and FAST SCSI 564 in slots relative to CPU Board 560 can be determined from System Window 540. Free slots 566 can be viewed as being open and the closest slots to CPU Board 560. Memory Expansion Board 568A is a detailed depiction of Memory Expansion Card 568. 1M Simm chips 570 are located on Memory Expansion Board 568A. Eight memory banks 572 remain unused. Drive Cage (Side View) 554A is a detailed depiction of the Drive Cage 554. 535 MB Hard Drive (SCSI) 574, 3.5" 1.44MB FD 576, and a 525MB Tape Backup (SCSI) 578 are contained within the Drive Cage 554. Front 580 indicates the location of the front side of Drive Cage (Side View) 554A. Therefore, 3.5" 1.44MB FD 576 and 525MB Tape Backup 578 have been configured to be front-accessible components. Bay 582 is a front-

accessible bay that does not contain any device. Bay 584 is a free bay located in the back of the Drive Cage 554.

The system window further provides the ability to interactively edit
5 the graphically rendered structures. The present invention provides the ability to modify the structural aspects of the configured system by adding, deleting or replacing components within a configured structure. The present invention further provides the ability to modify the configured structure by modifying the structural interconnections and compositions.

10

This capability to graphically display and edit can be used on a newly configured system, or an existing configuration, or system. That is, any upgrades to an existing, configured system may be performed graphically. A "freeze and fill" capability allows the user to freeze some portion of the
15 existing system, and fill, or modify the unfrozen portion. This "freeze and fill" capability further provides the ability to generate a quote for the new configuration that represents only those components added to the original configuration, and that incorporate any credit for the deleted or replaced components.

20

In the preferred embodiment, the list of materials, called the Bill of Materials (BOM) provides a list of all of the configured components and spare parts that are used in the system since the last request to configure the system. The part number and description is provided for each component and spare
25 part.

In the preferred embodiment, the parts list provides information regarding additional components (i.e., spare parts), resource totals, failed requests, and failed optional requests. Resource totals provides a total of all components and resources requested directly from the user. Failed Requests
5 and Failed Optional Requests are those component requests that could not be satisfied because of a lack of space, connector availability, etc.

QUOTER

10 The Quoter calculates the cost of the individual product packages and determines the cost of all product packages required to complete the system. The Quoter provides the ability to display the quote in various ways. For example, the quote may be displayed by product with the capability to expand or collapse the product information to show pricing for individual product
15 parts or for the entire package, respectively. The way in which products are presented or prices are calculated may be customized.

THE FLASH CONFIGURATION CACHE

20 An embodiment of the present invention includes method and apparatus for expediting configuration of a variety of "end products." Examples of the end products are computers, electronic systems such as voice mail systems, PBX systems, central office switches, and handheld communication devices. The present invention's flash configuration is also
25 used to configure end products such as airplanes where a variety of power system options, landing system options, and interior system options need be configured in an efficient and thorough manner. Other end products

configured by the flash configuration cache of the invention are trucks, test equipment, and chemical processes. The flash configuration cache is also used to configure vacation packages where each package involves a number of transportation options, lodging options, and recreational options.

5

In order to provide a specific example, the application of the flash configuration cache of the invention in configuring user computers is explained below. However, it is understood that the flash configuration cache applies as well to configure end products other than computers, examples of which end products were given above. Thus, in one implementation, the flash configuration cache of the invention is used for expediting configuration of a customer computer (also referred to as the user computer) in response to "new" customer requests. According to this embodiment, certain "old" customer requests for configuring the user computer are structured in the form of a tree (also called a "search tree") and saved in a "flash configuration cache" by the computer processing the requests (also called the host computer). A "branch" of the search tree is defined as a number of customer requests that originates from the root node of the tree and ends in a terminal node. For example, referring to Figure 14, customer requests R1, R2, R3, R4, R5, and R6 constitute a tree branch which originates from the root node and ends with a terminal node (i.e., the last customer request R6). A "path" of the search tree is defined as a number of customer requests that originates from a root node, but that does not end in a terminal node. In the example, customer requests R1, R2, and R3 constitute a tree path which originates from a root node and ends with a non-terminal node (i.e., the non-terminal node of customer request R3). From the above definition it is apparent that each branch consists of a number of paths. In

the example, the branch R1, R2, R3, R4, R5, and R6 consists of five paths, which are (1) R1; (2) R1 and R2; (3) R1, R2, and R3; (4) R1, R2, R3, and R4; and finally (5) R1, R2, R3, R4, and R5.

5 Each branch or path of the tree represents an "old configuration." Each old configuration represents a number of old customer requests. As each set of new customer requests is obtained, the flash configuration cache is searched to match the new requests with an old configuration stored in the flash configuration cache, namely, a path or branch of old customer requests
10 stored in the search tree. The longest path or branch of old requests, representing the maximum extent of matching requests, is then found and selected. The old configuration represented by the matching path or branch is then recalled from the cache. In this manner, new customer requests can be quickly configured if they at least partially match against previously
15 configured requests. In other words, instead of generating a computationally intensive new configuration, a preconfigured system corresponding to an old set of requests is recalled from the cache. This results in a speed advantage. A partially matched configuration is then partially configured and hence can be configured from that point on.

20

 In this embodiment, a request is either a component request or a resource request. Each component or resource request is associated with a number of constraints that are specified in the productbase. As each new request is compared to an old request, the constraints associated with the new
25 request are identified in order to match them against the constraints associated with the old request. However, it is understood that if the productbase has not changed, the constraints associated with the old and new

requests automatically match. Accordingly, if the productbase has not changed, the matching of the constraints is unnecessary. For matching the old and new component requests, the component name, quantity, priority, and attribute values should match. In matching old and new resource
5 requests, component class names, quantity, priority, and allocation types should match.

According to the flash configuration cache embodiment of the invention, the cache is initially void of requests and the corresponding
10 configurations. However, as customer requests are obtained, certain ones of the requests and their corresponding configurations are saved in the cache. The requests that do get stored are called "old" requests, and their corresponding configurations are called "old" configurations.

15 The host computer makes a decision as to which ones of the requests should be stored in the cache. The requests that are stored in the cache are those for which generating a new configuration is more time consuming than recalling the preconfigured requests from the cache. Accordingly, the host computer keeps track of the time that it takes to generate each
20 configuration in response to a given set of requests. The host computer then compares the time consumed in generating a configuration to a predetermined time. If the generation time is greater than the predetermined time, the configuration is stored in the cache. Otherwise, the configuration is not saved. The predetermined time is typically an estimate of the time that it
25 takes to recall a given configuration from the cache. By way of example, suppose that in response to customer requests for a monitor, a floppy drive, an IDE hard drive, a power supply cabinet, an IDE hard drive controller, a 486

CPU, and a memory board, the host computer generates a new user computer configuration. In this example, the generation of the configuration takes five minutes of the host computer time. The host computer compares this time (namely, the five minutes) with a predetermined time, for example one
5 second, that it takes to recall a new configuration from the cache. Since the generation of the configuration is more time consuming than recalling the configuration from the cache, this configuration gets stored in the cache as a configuration representing the specific customer requests that resulted in the configuration. This configuration is stored in the cache and represented by
10 the "old" set of requests that resulted in that configuration. The old set of requests is stored as one path or branch of the search tree. When a set of new customer requests is input to the host computer, the host computer methodically searches the tree to find a path or branch of old customer requests that matches the new customer requests. If a matching set of old
15 customer requests is found, the host computer recalls from the cache the old configuration associated with the set of old customer requests. This is because the host computer has already determined that the time required to generate a new configuration for the new customer requests is longer than the time required for recalling the matching old configuration from the flash
20 configuration cache.

Thus, the configurations that are saved in the cache are those that require a shorter time to be recalled from the cache than be regenerated. The sets of stored requests (namely, "old" requests) are organized in a "tree"
25 structure and methodically searched so that a new set of requests can be matched against an old set of requests. Each set of old requests is stored as a path or branch, and a number of branches constitute the search tree. The

algorithm and related flow diagrams to construct and to search the tree are explained below. For simplicity in the following discussion, the assumption is made that all of the request sets take longer to configure than they take for recall from the cache.

5

As stated above, the cache does not initially have any requests stored therein. As each set of customer requests for user computer (or other end product) configuration is obtained, the search tree is expanded. Thus, the cache size increases as time goes by. Accordingly, the chances that new
10 customer requests have already been stored in the cache increases as time goes by. Figures 13A-13C illustrate the algorithm used for building the search tree and adding requests to the cache. The algorithm for searching the request tree begins in step 101. Since a new request list is to be configured, the algorithm searches for the root of the search tree (step 103). The algorithm
15 then proceeds to search for the first request of those stored sets of requests that diverge directly from the root node of the tree (step 105). Then the host computer determines whether any of these stored first requests matches the first request in the new request list (step 107). If none of the stored first requests diverging from the root node matches the first request in the new
20 request list, the entire new request list is added to the tree (step 109). The new request list is added such that the entire request list diverges from the root node of the tree. The reason that the new request list is added as diverging from the root node is that the root node is the only common node between the new request list and the stored request lists. The host computer then
25 generates a new configuration and stores it in the cache. The new configuration is represented by the set of requests that was just added to the search tree (step 110). After adding the new set of requests to the flash

configuration cache, the algorithm for searching the request tree ends (step 113). Thereafter, the request list just added is considered as a list of "old" requests constituting one of the branches of the search tree stored in the flash configuration cache.

5

If, however, there is a match between the first request and one of the stored first requests directly diverging from the root node, the algorithm proceeds to determine whether the next request in the request list matches an old request diverging from the first stored request diverging from the root
10 node (step 111). In other words, the algorithm proceeds to match the next request in the new request list against a stored request diverging from the first matching stored request. In step 115, the algorithm determines whether a stored request diverging from the first matching request matches the next request in the new request list. If no match is found, the new request list is
15 added to the search tree (step 117). This is performed by adding the new request list in a manner such that the first matching request is shared with another branch of the search tree. As such, the new request list diverges from the first matching request of that other branch. In step 118, the new configuration represented by the new request list is generated and added to
20 the cache. After adding the list of new requests to the flash configuration cache, the algorithm for searching the request tree ends (step 121). Thereafter, the request list just added is considered as a list of "old" requests constituting one of the branches of the search tree stored in the flash configuration cache.

25

If, however, there is match between the stored request diverging from the first matching request and the next request in the new request list, the

algorithm proceeds to determine whether the next request in the new request list matches an old request diverging from the next stored request diverging from the root node (step 119). In other words, the algorithm proceeds to match the next request in the new request list against a stored request

5 diverging from the second matching stored request. In step 123, the algorithm determines whether a stored request diverging from the second matching request matches the next request in the new request list. If no match is found, the new request list is added to the search tree (step 124). This is performed by adding the new request list in a manner such that the second

10 matching request is shared with another branch of the search tree. As such, the new request list diverges from the second matching request of that other branch. In step 125, the new configuration represented by the new request list is generated and added to the cache. After adding the list of new requests to the flash configuration cache, the algorithm for searching the request tree

15 ends (step 126). Thereafter, the request list just added is considered as a list of "old" requests constituting one of the branches of the search tree stored in the flash configuration cache. The algorithm continues in this manner until all of the requests are processed (step 127). When all of the requests in the list of new requests are processed the algorithm ends (step 129). The following

20 example illustrates the operation of the algorithm.

Referring to Figure 14, suppose that the first request list to be processed is that consisting of R1, R2, R3, R4, R5, and R6. Following the algorithm of Figures 13A-13C and since there are no preexisting requests stored in the

25 cache, the host computer finds no matches between the first request (namely R1) and the contents of the cache. Accordingly, the entire request list consisting of R1, R2, R3, R4, R5, and R6 is added to the cache as shown in

Figure 14. The host computer then generates a new configuration represented by this request list and adds the new configuration to the cache.

Suppose that the next request list consists of R1, R2, R3, R4, R7, and R8.

5 In this case, the algorithm first finds a match for the first request (R1) which is already stored in the cache as a part of the stored request branch consisting of R1, R2, R3, R4, R5, and R6. The algorithm continues by searching for the next matching request that diverges from R1. In other words, the algorithm searches all of the second requests of the stored branches that had R1 as their
10 first request. The only branch is the R1, R2, R3, R4, R5, and R6 branch which is stored in the cache. Since there is an R2 diverging from R1 in this branch, the algorithm finds a match between the first and second requests (R1 and R2) of the new request list and the stored branch consisting of R1, R2, R3, R4, R5, and R6. The algorithm continues in this manner and finds that the first
15 four requests in the stored branch consisting of R1, R2, R3, R4, R5, and R6 match the first four requests of the new request list consisting of R1, R2, R3, R4, R7, and R8. The algorithm then finds that no matching request for R7 has been stored in the cache. Accordingly, as shown in Figure 14, the remainder of the new request list, namely R7 and R8 are added to the cache as
20 branching off from the first four requests of the request list already stored in the cache. Finally, the algorithm generates a new configuration represented by the new request list, i.e. R1, R2, R3, R4, R7, and R8, and stores the new configuration in the flash configuration cache.

25 The next request list is R1, R9, R10, R11, R12, and R13. The algorithm finds a match only for the first request in the list of new requests, namely R1. The remaining requests, R9, R10, R11, R12, and R13, are added to the cache as

a new branch diverging from R1 of the stored request list as shown in Figure 14. Then the algorithm generates a new configuration represented by the new request list, i.e. R1, R9, R10, R11, R12, and R13, and stores the new configuration in the cache. The next request list is R14, R15, and R16. For this
5 request list no matches are found. Thus, the algorithm saves the entire request list as branching off from the root node as shown in Figure 14. Thereafter, the algorithm generates a new configuration represented by the new request list, i.e. R14, R15, and R16, and stores the configuration in the flash configuration cache.

10

The assumption implicit in explaining the algorithm and in discussing the above example was that all of the request lists discussed take longer to configure than they take for recall from the cache. In actual operation, before storing any of the request lists, the host computer makes an initial
15 determination whether the configuration time exceeds a predetermined time. The requests are stored only when the configuration time exceeds the predetermined time. The predetermined time is typically the time that it takes to recall a configuration from the cache.

20

Thus far the procedure for building and organizing the contents of the flash configuration cache was discussed. Now the operation for recalling a stored configuration from the cache is discussed. To this end, the host computer must, partially or completely, match a new request list with an old request list stored in the cache and then recall the configuration represented
25 by the old request list. The matching operation is similar to the operation for storing the requests. As shown in Figures 15A-15C, the algorithm begins by searching for the root of the tree (steps 201 and 203). Then, a search is

performed on all stored paths for their first stored requests diverging from the root of the tree (step 205). The algorithm then determines whether any of the stored paths has a first request that matches the first request of the new request list (step 207). If there is no match, the host computer proceeds to

5 generate a new configuration for the user computer (or other end product) (step 209) and the operation ends (step 213). If there is a match, the algorithm checks the next new request diverging from the first request of the new request list. This is done by performing a search on all stored paths that had a successful match on their first requests. These paths are searched for their

10 second stored request diverging from their respective first requests (step 211). The algorithm then determines whether any of the stored paths has a second request that matches the second request of the new request list (step 215). If there is no match, the host computer recalls the configuration for the matching requests (step 216). At this point only the first request is matching.

15 Thus, the configuration corresponding to the first request is recalled from the flash configuration cache. The host computer proceeds to generate a configuration corresponding to the remaining (unmatched) requests in the new request list (step 217). Thus, the configuration of the end product now consists of the combination of an old configuration and a newly generated

20 configuration. After completion of the configuration process, the operation ends (step 221).

If there is a match, i.e. if one of the stored paths has a second request that matches the second request of the new request list, the algorithm

25 considers the next new request diverging from the second request of the new request list. This is done by performing a search on all stored paths that had successful matches on their first and second requests. These paths are

searched for their third stored request diverging from their respective second requests (step 219). The algorithm then determines whether any of the stored paths has a third request that matches the third request of the new request list (step 223). If there is no match, the host computer recalls the configuration
5 for the matching requests (step 224). At this point only the first and second requests are matching. Thus, the configuration corresponding to the first and second requests is recalled from the flash configuration cache. The host computer proceeds to generate a configuration corresponding to the remaining (unmatched) requests in the new request list (step 225). Thus, the
10 configuration of the end product now consists of the combination of an old configuration and a newly generated configuration. After completion of the configuration process, the operation ends (step 226). This process continues until all requests in the new request list are processed in the manner described above. The algorithm ends when all of requests in the set of new
15 requests have been processed (steps 227 and 229).

From the above description, it is apparent that there may in fact be a set of new requests for which there is a "partial match." This can happen in two ways. A first way is when the matching branch is "too short" for the new
20 request list. By way of example and referring to Figure 14, suppose that the set of new requests consists of requests R14, R15, R16, and R17. The matching set of requests is the branch consisting of R14, R15, and R16 as shown in Figure 14. Thus, following the algorithm described above, the invention recalls the configuration corresponding to the matching branch R14, R15, and
25 R16 which is only a "partial configuration" relative to the set of new requests consisting of R14, R15, R16, and R17. This "partial configuration" is combined with a configuration that the host computer generates for the

unmatched new request, namely R17. The combination of the partial configuration corresponding to R14, R15, and R16 and the new configuration corresponding to R17 is a complete configuration corresponding to the set of new requests R14, R15, R16, and R17.

- 5 In this manner, the necessity to generate an entirely new configuration for the new request list is obviated and the process of configuring the user computer (or other end product) in response to the new set of requests is expedited.

- 10 A second way in which a "partial match" occurs is when there is a matching path, as opposed to a matching branch, which is "too short" for the new request list. Referring to Figure 14 and by way of another example, suppose that the set of new requests consists of R1, R2, R3, R18, and R19. Following the above algorithm, the invention finds the matching path
- 15 consisting of R1, R2, and R3 belonging to the branch R1, R2, R3, R4, R5, and R6 in Figure 14. The host computer then recalls the configuration corresponding to the path R1, R2, and R3. This configuration is a "partial configuration" for the set of new requests R1, R2, R3, R18, and R19. The host computer then generates a new configuration for the remaining (unmatched)
- 20 requests, namely R18 and R19. The newly generated configuration is then combined with the old "partial configuration." The combination of the newly generated configuration and the old partial configuration makes up a complete configuration for the set of new requests R1, R2, R3, R18, and R19.

THE BUNDLING CACHE

Another embodiment of the invention includes a bundling cache. The bundling cache is used to speed up the process of converting a
5 configuration generated by the host computer into actual available commercial products. As with the flash configuration cache, the application of the bundling cache is not limited to user computers. However, in order to provide a specific example, application of the invention's bundling cache to user computers is discussed below. It is nevertheless understood that the
10 bundling cache of the invention is also used to expedite the process of converting the configurations generated by the host computer into actual commercial products for end products other than user computers. Examples of such end products were given above and include voice mail systems, PBX systems, central office switches, handheld communication devices, airplanes,
15 trucks, test equipment, chemical processes, and vacation packages.

Thus, the invention's bundling cache is used to expedite the process of converting, for example, user computer configurations into actual commercial products. In generating a configuration of the user computer,
20 the host computer requires a number of computer components. The computer components then must be converted into available commercial products. For example, suppose that as a part of the configuration of the user computer, the host computer requires a series of components such as certain types of monitors, floppy drives, IDE hard drives, power supply cabinets, and
25 IDE hard drive controllers, 486 CPU's, and memory boards. Often, many components such as the ones mentioned above are available as a group in a single commercial product. In other words, each commercial product is

comprised of a number of components. For example, a single commercial product may contain a monitor, a floppy drive, an IDE hard drive, a power supply cabinet, an IDE hard drive controller, a 486 CPU, and a memory board. By way of a specific example, a commercial product called 486-100 contains a

5 Tower Powerhouse 486/33, a 4MB system memory, a 3.5" 1.44 MB floppy drive, a 100 MB IDE disk drive, an IDE controller card, and a VGA monitor.

The host computer may, in addition to requiring a number of components for the user computer configuration, require certain options.

10 For example, the host computer may require memory upgrade options of 4MB, 8MB, and 12MB. A commercial product that consists of the components required by the host computer may or may not offer the required options. For example, the commercial product 486-100 offers three different memory upgrade options of 4MB, 8MB, and 12MB. Each memory upgrade

15 option offers a respective additional amount of memory of 4MB, 8MB, or 12MB beyond the 4MB system memory that is provided by the 486-100 commercial product.

The bundling cache is utilized to speed up the process of finding

20 products that match the required components. In the absence of the bundling cache, the host computer searches for products that offer the required components in various large data bases. This is a slow process. According to the bundling cache embodiment of the invention, a cache of products and their required part numbers is maintained.

25

The various products that offer the required components of the configuration generated by the host computer are obtained from various data

bases accessible to the host computer. Each of these products offers certain options as shown in Figure 16. For example, product P1 requires options O1, O2, O3 and O5. Product P2 requires options O2, O3, O4, and O5, and product P3 requires only option O5. A first hash table in the host computer correlates
5 each product to the required (i.e. necessary) options offered by that product.

A second hash table is built which correlates all the available options to their respective products in which they are required (i.e. necessary). As shown in Figure 17, option O1 is correlated to product P1 only. Option O2 is
10 correlated to products P1 and P2, option O3 to products P1 and P2, option O4 to product P2, and option O5 to products P1, P2, and P3. The products that do not offer the required options are not of interest to the host computer for building the bundling cache. Thus, the bundling cache is utilized to store information related to those products that have been identified as offering
15 the options required by the host computer. For example, if the host computer requires options O1 and O2 in generating a configuration for the user computer, using the second hash table the host computer determines that only P1 offers both these options. Thus, every time that a set of requests results in a configuration requiring options O1 and O2, product P1 is recalled
20 from the bundling cache and is utilized in implementing the configuration of the user computer. Products such as P1 that are stored in the bundling cache are compiled in a product data base that is used in implementing the configurations generated by the host computer.

CLAIMS

1. A method using a host computer for configuring an end product in response to a plurality of new requests, said method comprising the steps of:
- 5 arranging a plurality of old requests into a plurality of paths and branches constituting a tree, each path or branch representing a unique old configuration of said end product stored in said host computer;
- inputting said plurality of new requests;
- searching a cache of said host computer for finding a matching path or
- 10 branch, said matching path or branch having a plurality of old requests matching said plurality of new requests; and
- recalling from said cache one of said plurality of old configurations represented by said matching path or branch.
- 15 2. The method of claim 1 wherein said end product is a user computer.
3. The method of claim 1 or 2 wherein said plurality of old requests and said plurality of new requests are component requests.
- 20 4. The method of claim 1 or 2 wherein said plurality of old requests and said plurality of new requests are need requests.
5. The method of claim 1 or 2 wherein said plurality of old requests and said plurality of new requests are resource requests.

6. The method of claim 2 wherein said plurality of old requests and said plurality of new requests comprise container, connection and component constraints.

- 5 7. A method using a host computer for configuring a user computer in response to a plurality of new requests, said method comprising the steps of:
- defining in said host computer a plurality of old configurations of said user computer comprising a plurality of old requests;
- arranging said plurality of old requests into a plurality of branches
- 10 constituting a tree, each branch representing one of said plurality of old configurations;
- inputting said plurality of new requests;
- searching a cache of said host computer for finding a matching branch, said matching branch having a plurality of old requests matching said
- 15 plurality of new requests; and
- recalling from said cache one of said plurality of old configurations represented by said matching branch.

8. The method of claim 7 wherein said plurality of old requests and
- 20 said plurality of new requests are component requests.

 9. The method of claim 7 wherein said plurality of old requests and said plurality of new requests are need requests.

- 25 10. The method of claim 7 wherein said plurality of old requests and said plurality of new requests are resource requests.

11. The method of claim 7 wherein said plurality of old requests and said plurality of new requests comprise container, connection and component constraints.

5 12. A method using a host computer for configuring a user computer in response to a plurality of new requests, said method comprising the steps of:

(a) defining in said host computer an old configuration of said user computer comprising a plurality of old requests;

10 (b) arranging said plurality of old requests into a branch, said branch representing said old configuration;

(c) determining a total time by adding a first time period for execution of the defining step to a second time period for execution of said arranging step;

15 (d) storing said branch and said old configuration in a cache of said host computer if said total time is greater than a predetermined time;

(e) repeating steps (a)-(d) to store in said cache a plurality of branches and a respective plurality of old configurations;

(f) inputting said plurality of new requests;

20 (g) searching a tree comprised of said plurality of said branches for finding a matching branch, said matching branch having a plurality of old requests matching said plurality of new requests; and

(h) recalling from said cache one of said plurality of said old configurations represented by said matching branch.

25

13. The method of claim 12 wherein said plurality of old requests and said plurality of new requests are component requests.

14. The method of claim 12 wherein said plurality of old requests and said plurality of new requests are need requests.

5 15. The method of claim 12 wherein said plurality of old requests and said plurality of new requests are resource requests.

16. The method of claim 12 wherein said plurality of old requests and said plurality of new requests comprise container, connection and
10 component constraints.

17. A method using a host computer for configuring an end product, said method comprising the steps of:

(a) generating a configuration of said end product in response to a
15 request list, said configuration comprising a plurality of components;

(b) searching a product catalog to find a plurality of matching products, each of said plurality of matching products correlating to one of said plurality of components;

(c) inputting a respective list of required options for each of said
20 components; and

(d) storing in a cache each of said plurality of matching products which offers said respective list of required options for a corresponding component.

18. The method of claim 17 wherein said end product is a user
25 computer.

19. The method of claim 17 or 18 further comprising the steps of storing in said cache said respective list of required options for a corresponding one of said plurality of matching products.

5 20. The method of claim 17 or 18 wherein said list of required options is comprised of one or more options.

21. The method of claim 17 or 18 further comprising the step of storing in said cache each of said plurality of matching products when no options are
10 required for a respective one of said plurality of correlating components.

22. A method using a host computer for configuring a user computer, said method comprising the steps of:

(a) generating a configuration of said user computer in response to a request list, said configuration comprising a plurality of components;

5 (b) constructing a product catalog comprising the steps of:

(i) inputting a plurality of required options for each of said plurality of components;

(ii) generating a first hash table correlating each of said plurality of matching products to its plurality of required options;

10 (iii) identifying said plurality of required options among said plurality of required options;

(iv) generating a second hash table correlating each of said plurality of required options to respective ones of said plurality of matching products, thereby identifying a target group of said plurality of matching
15 products;

(iv) storing said target group of said plurality of matching products as said product catalog in a cache of said host computer; and

(c) searching said product catalog to find a plurality of matching products, each of said plurality of matching products correlating to one of said
20 plurality of components.

1/33




FIG. 1

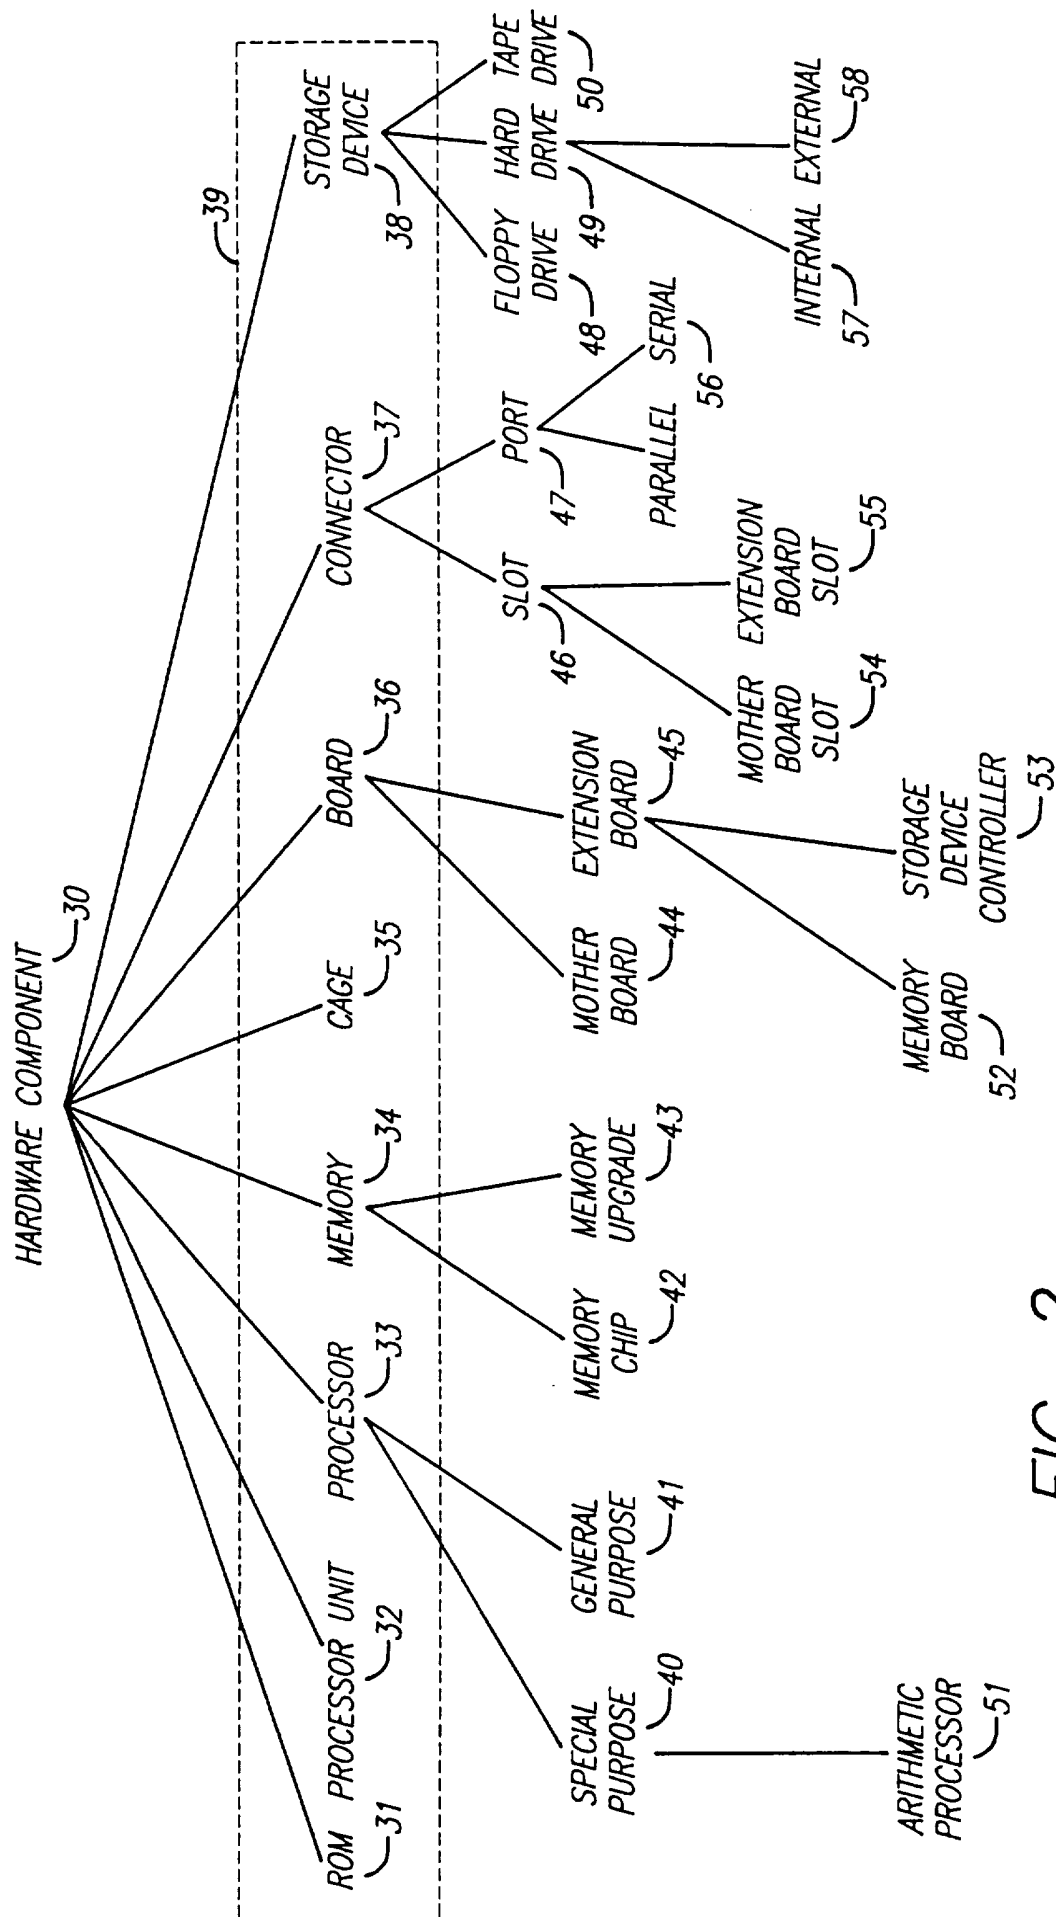


FIG. 2

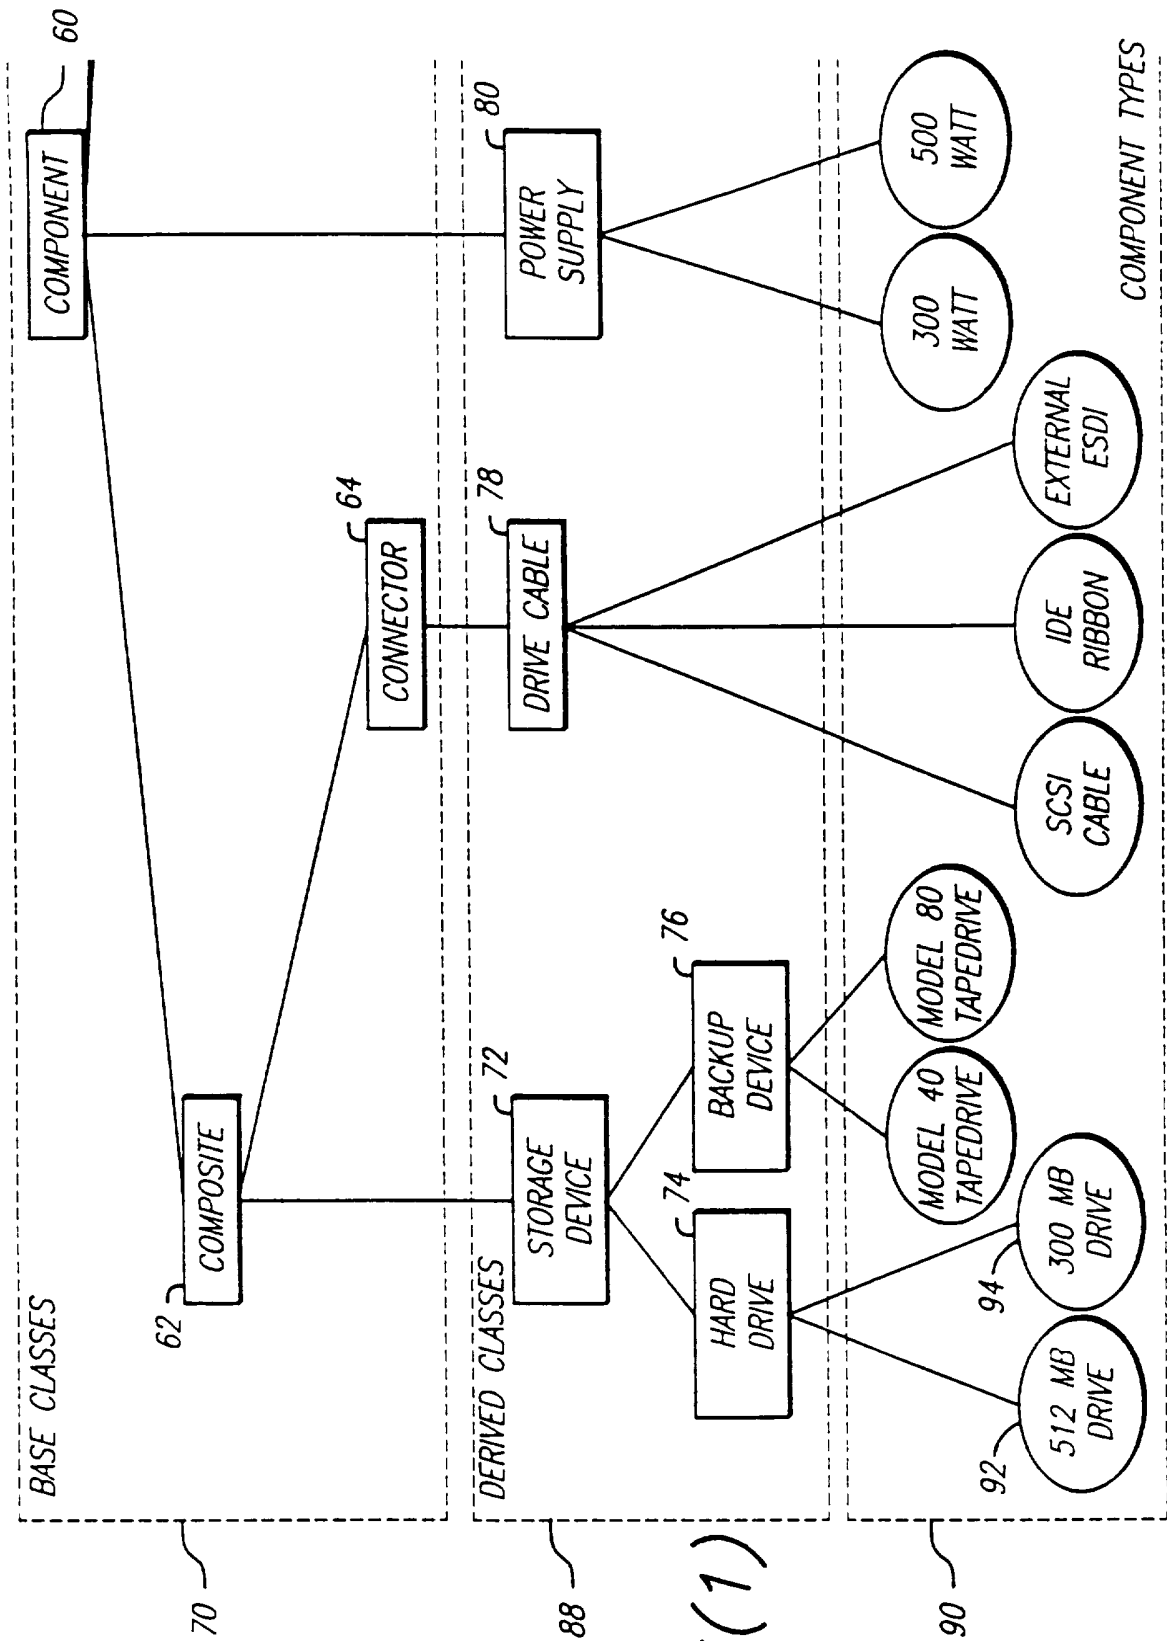



FIG. 3(1)

4 / 3 3

FIG. 3(2)



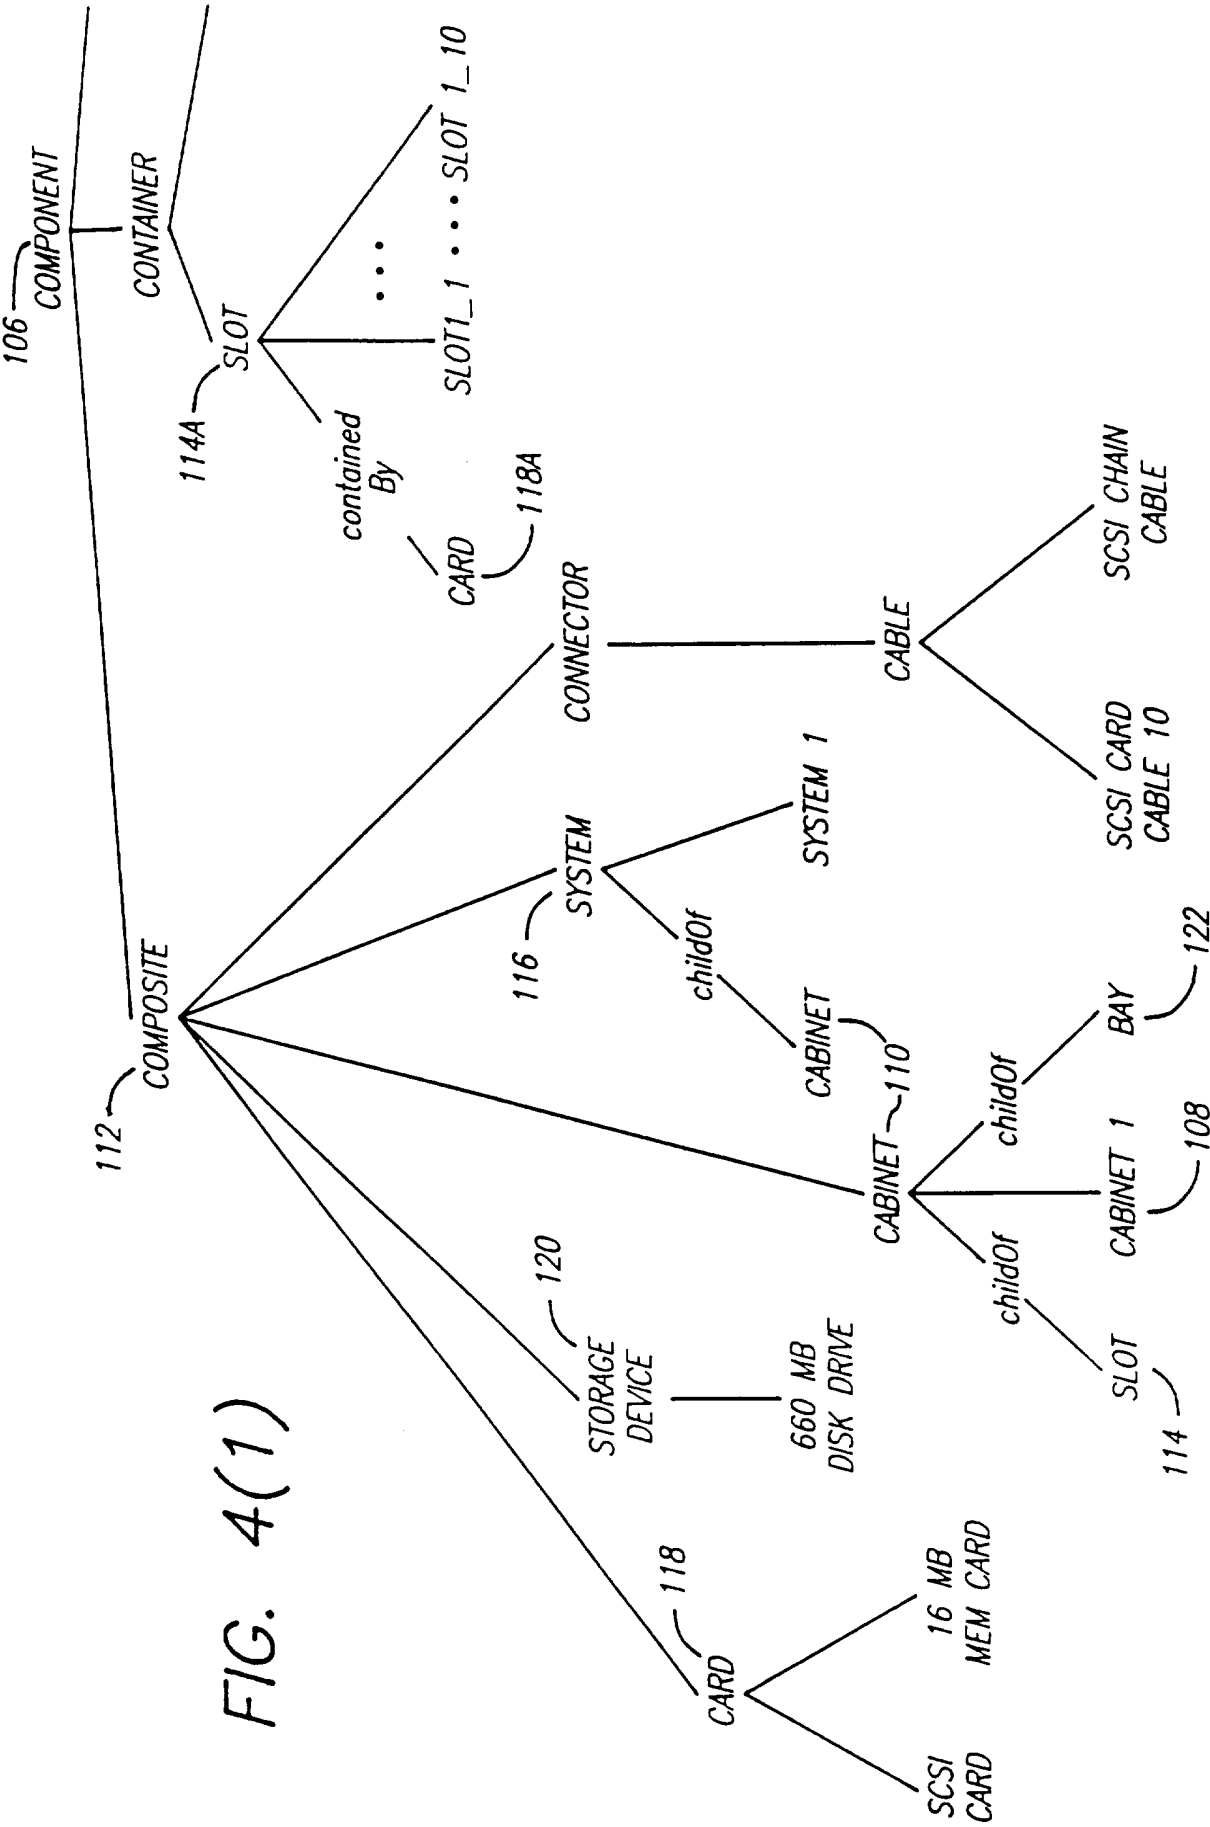


FIG. 4(1)

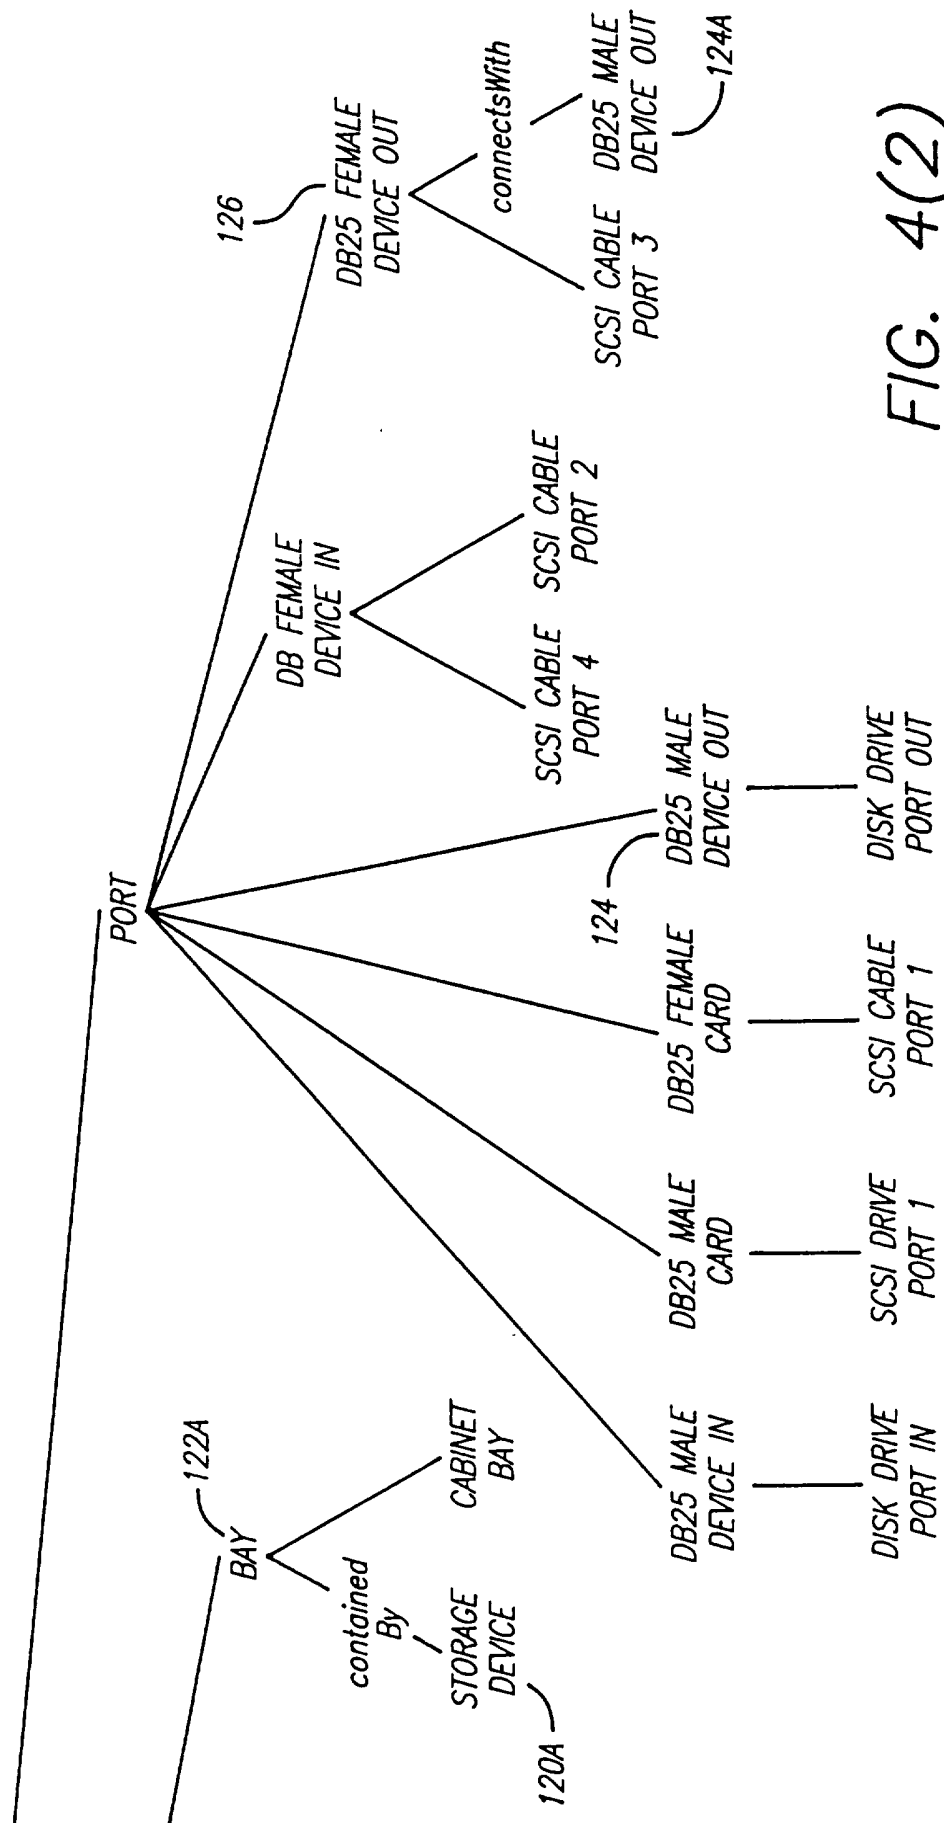


FIG. 4(2)

7 / 3 3


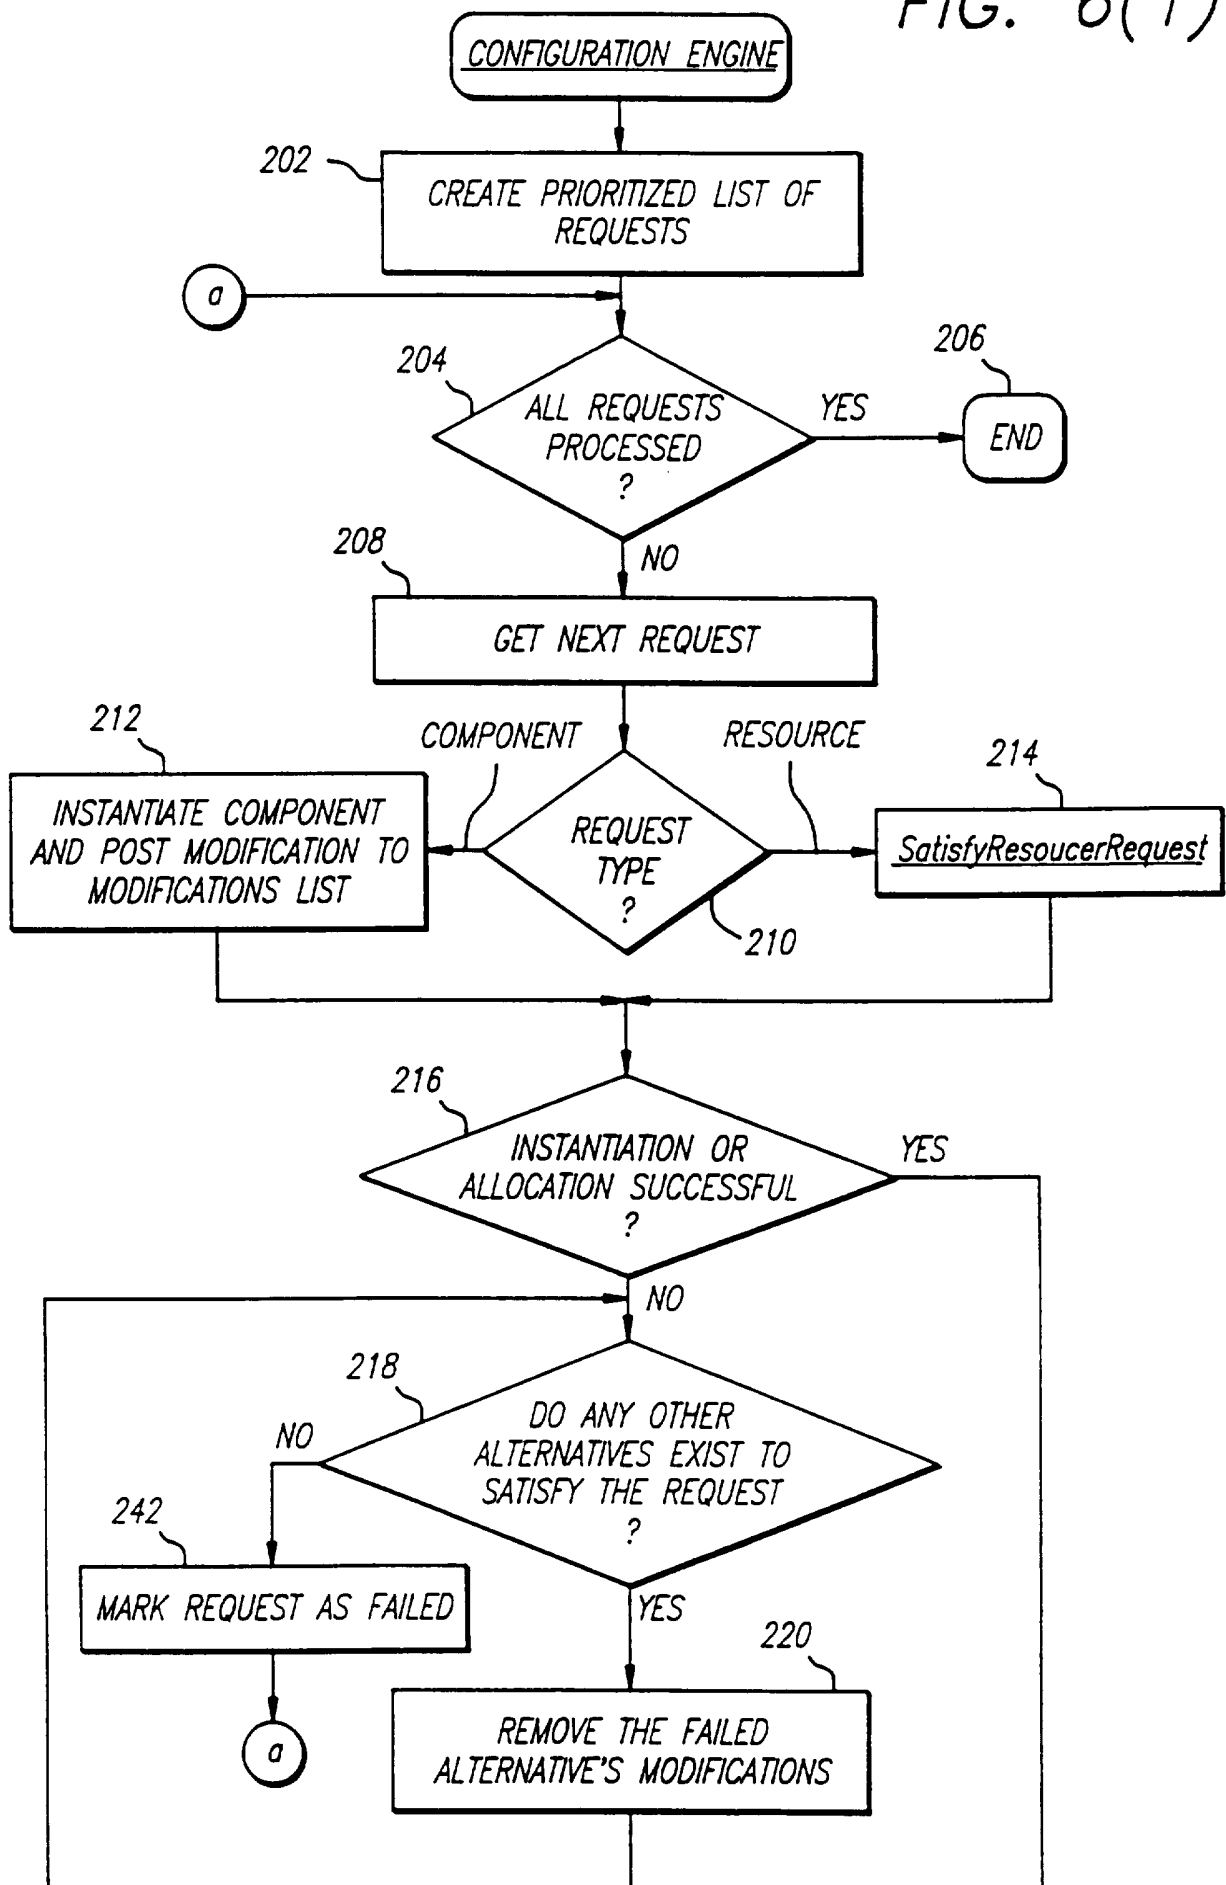


FIG. 5

8/33

FIG. 6(1)



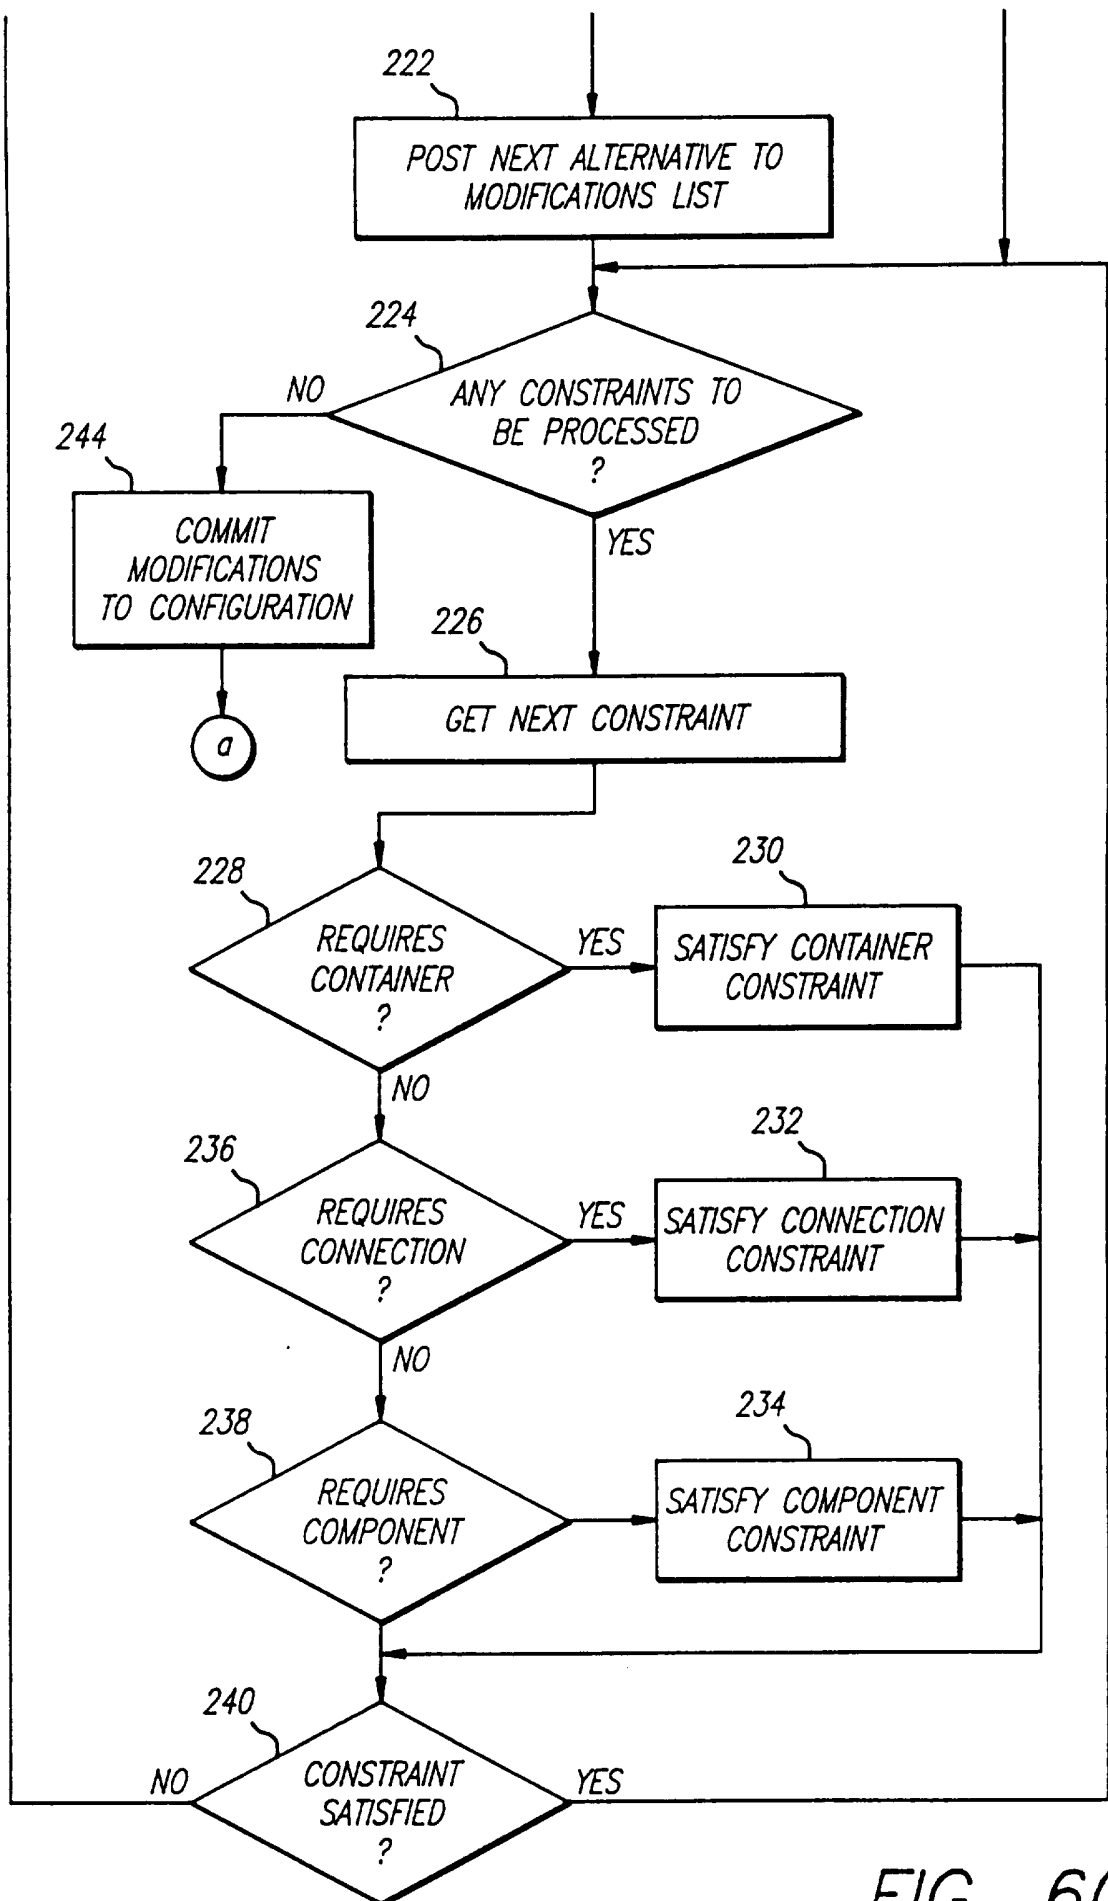
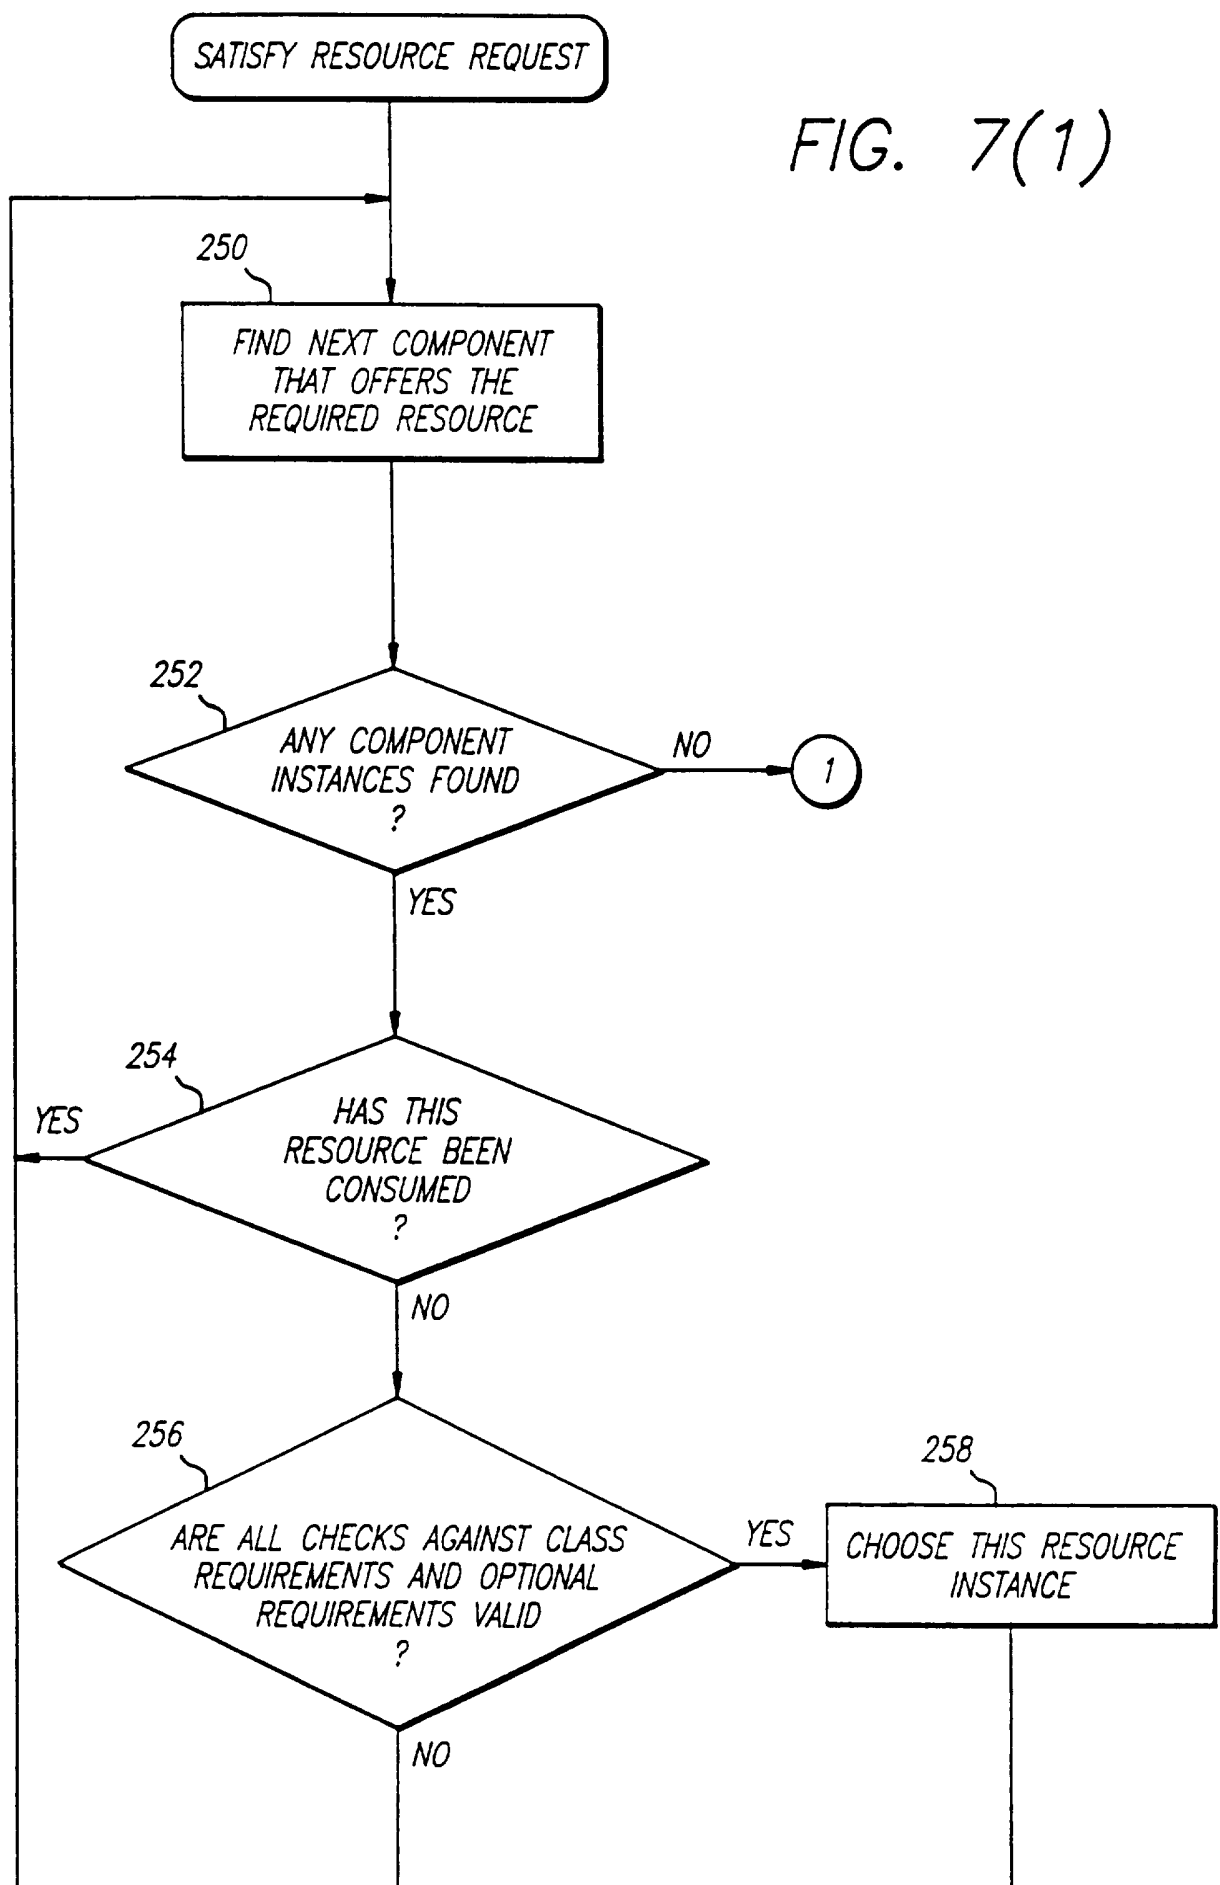


FIG. 6(2)

10/33

FIG. 7(1)



11/33

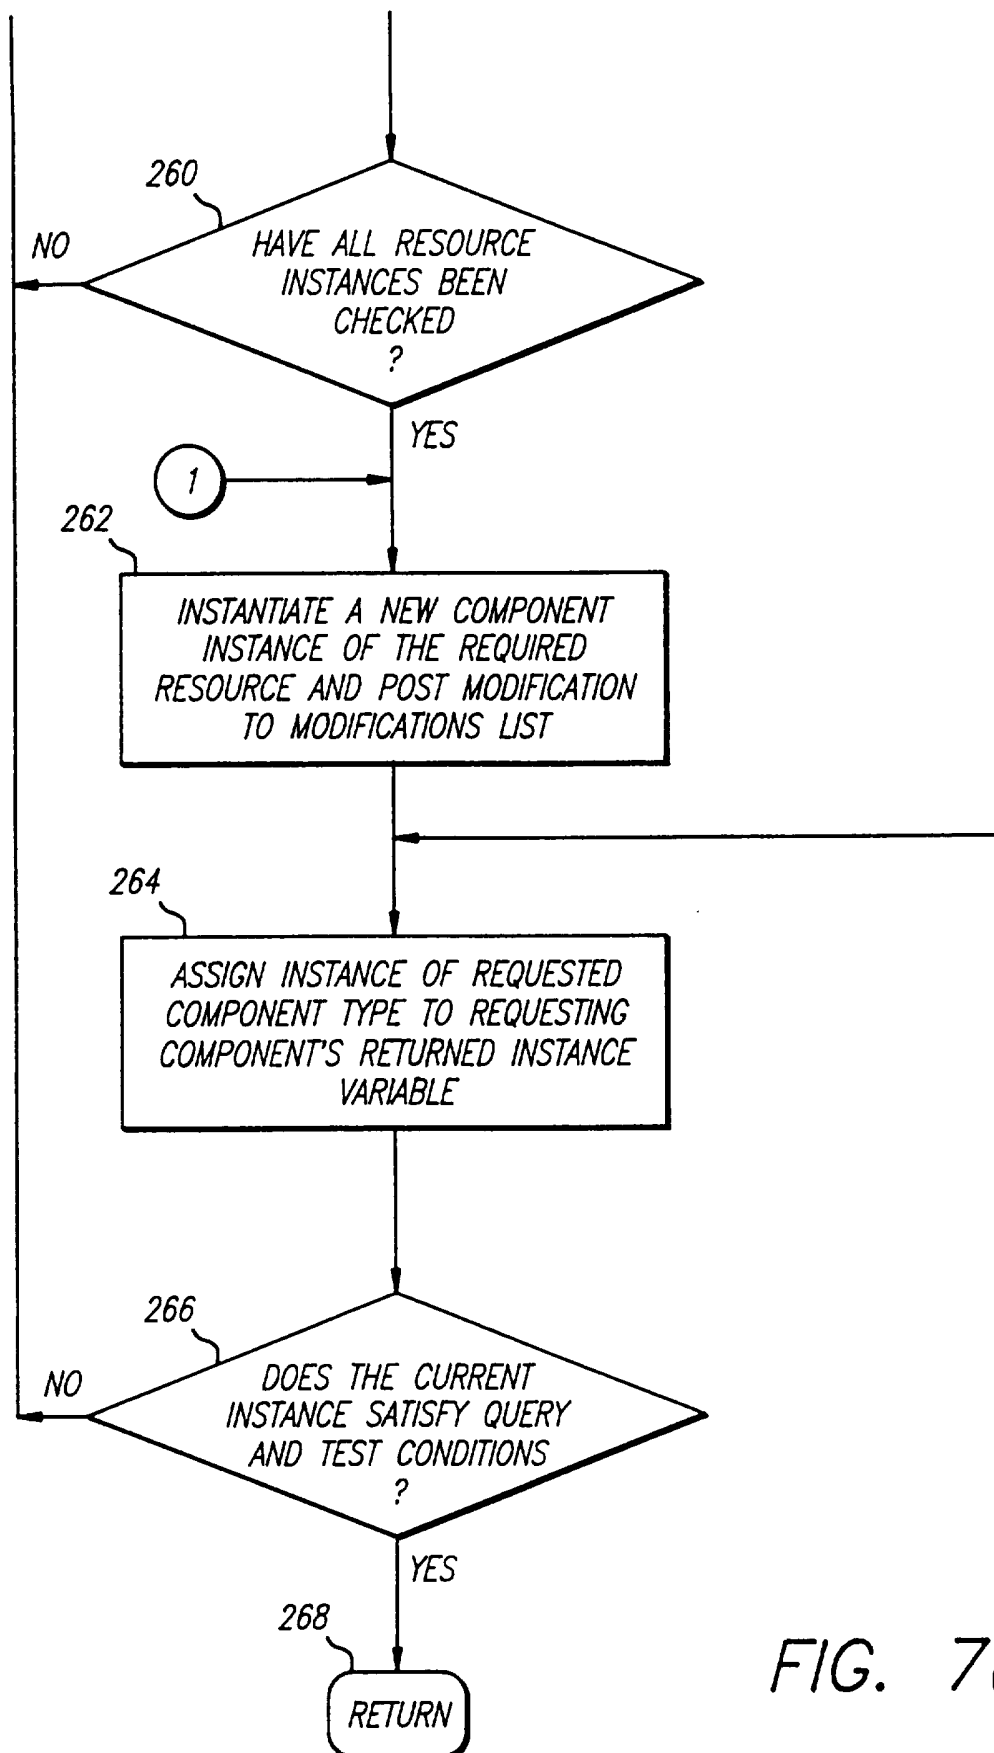
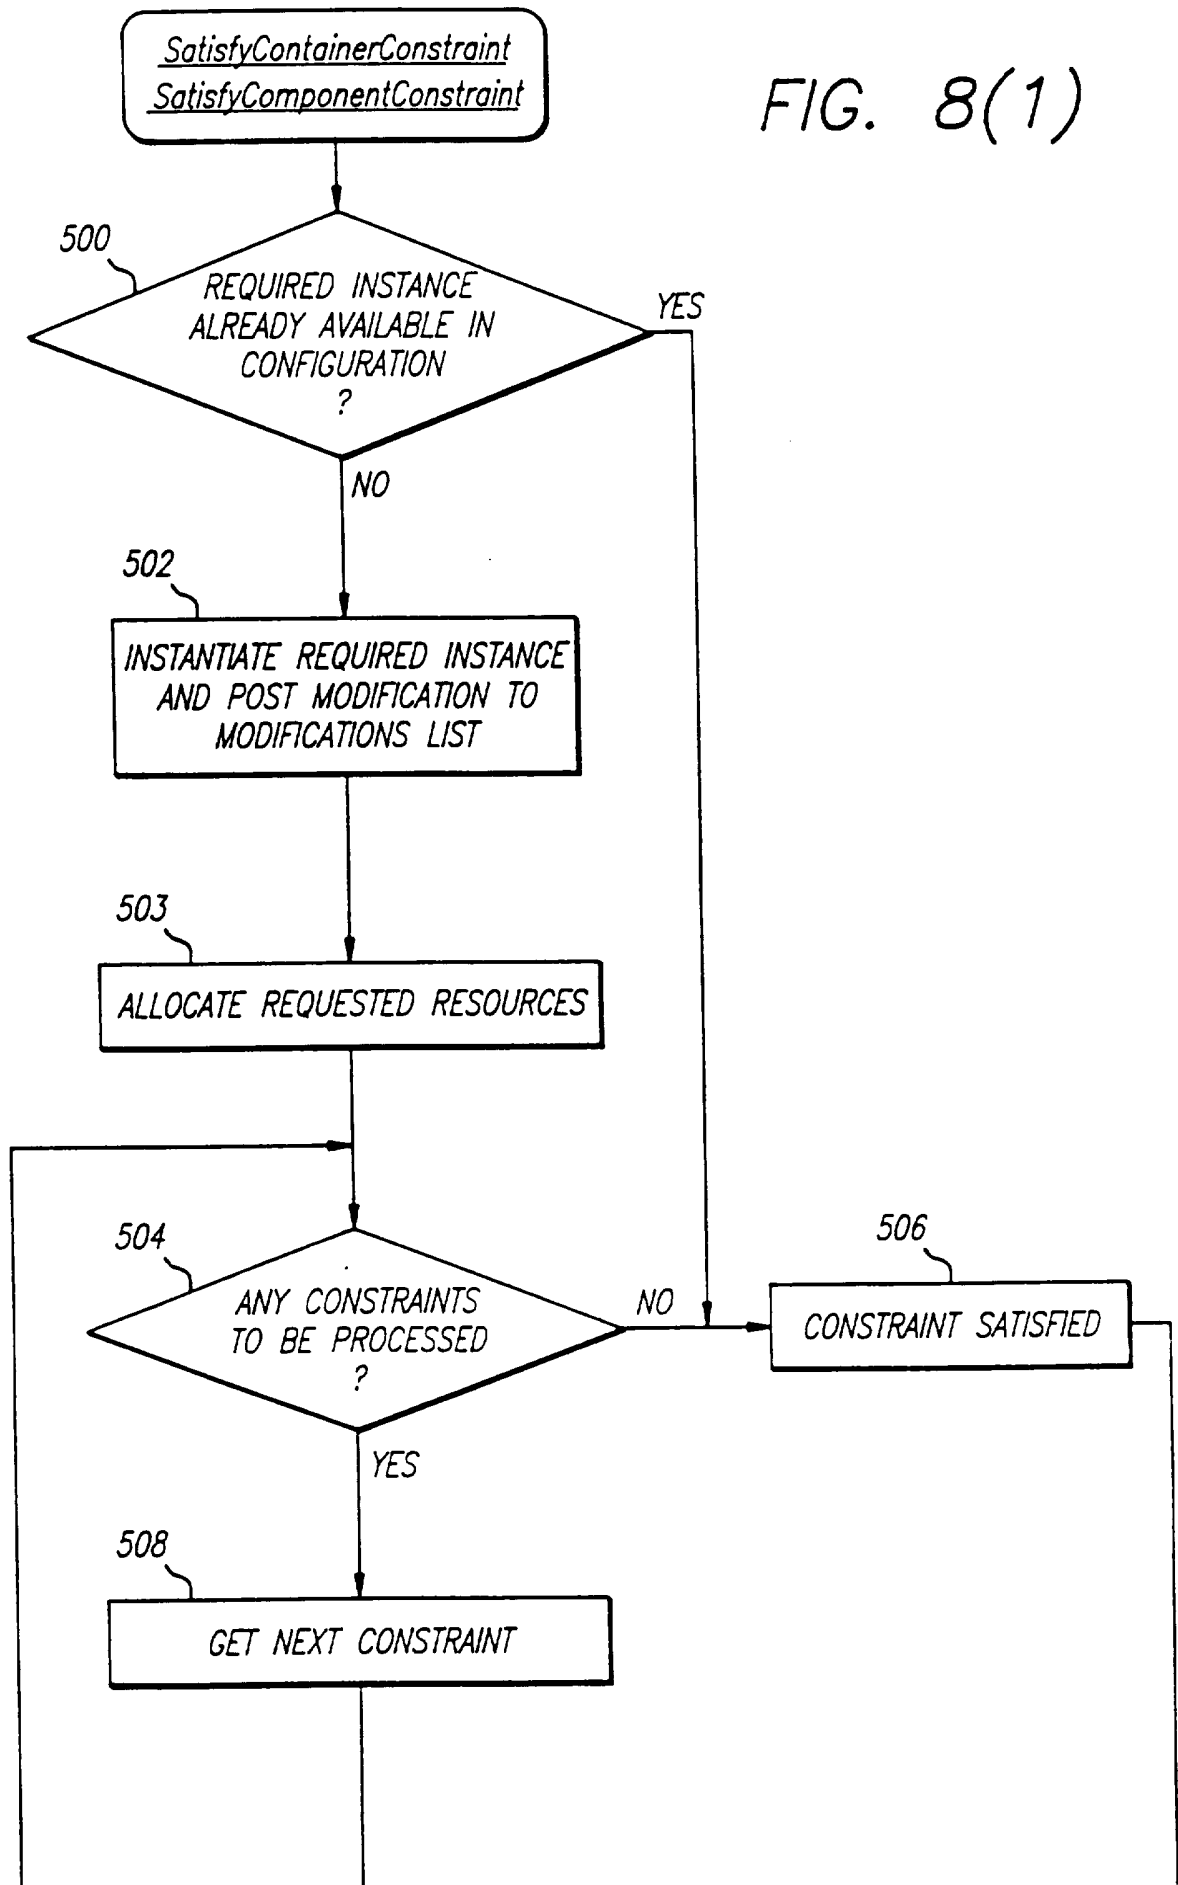


FIG. 7(b)

12/33

FIG. 8(1)



13/33

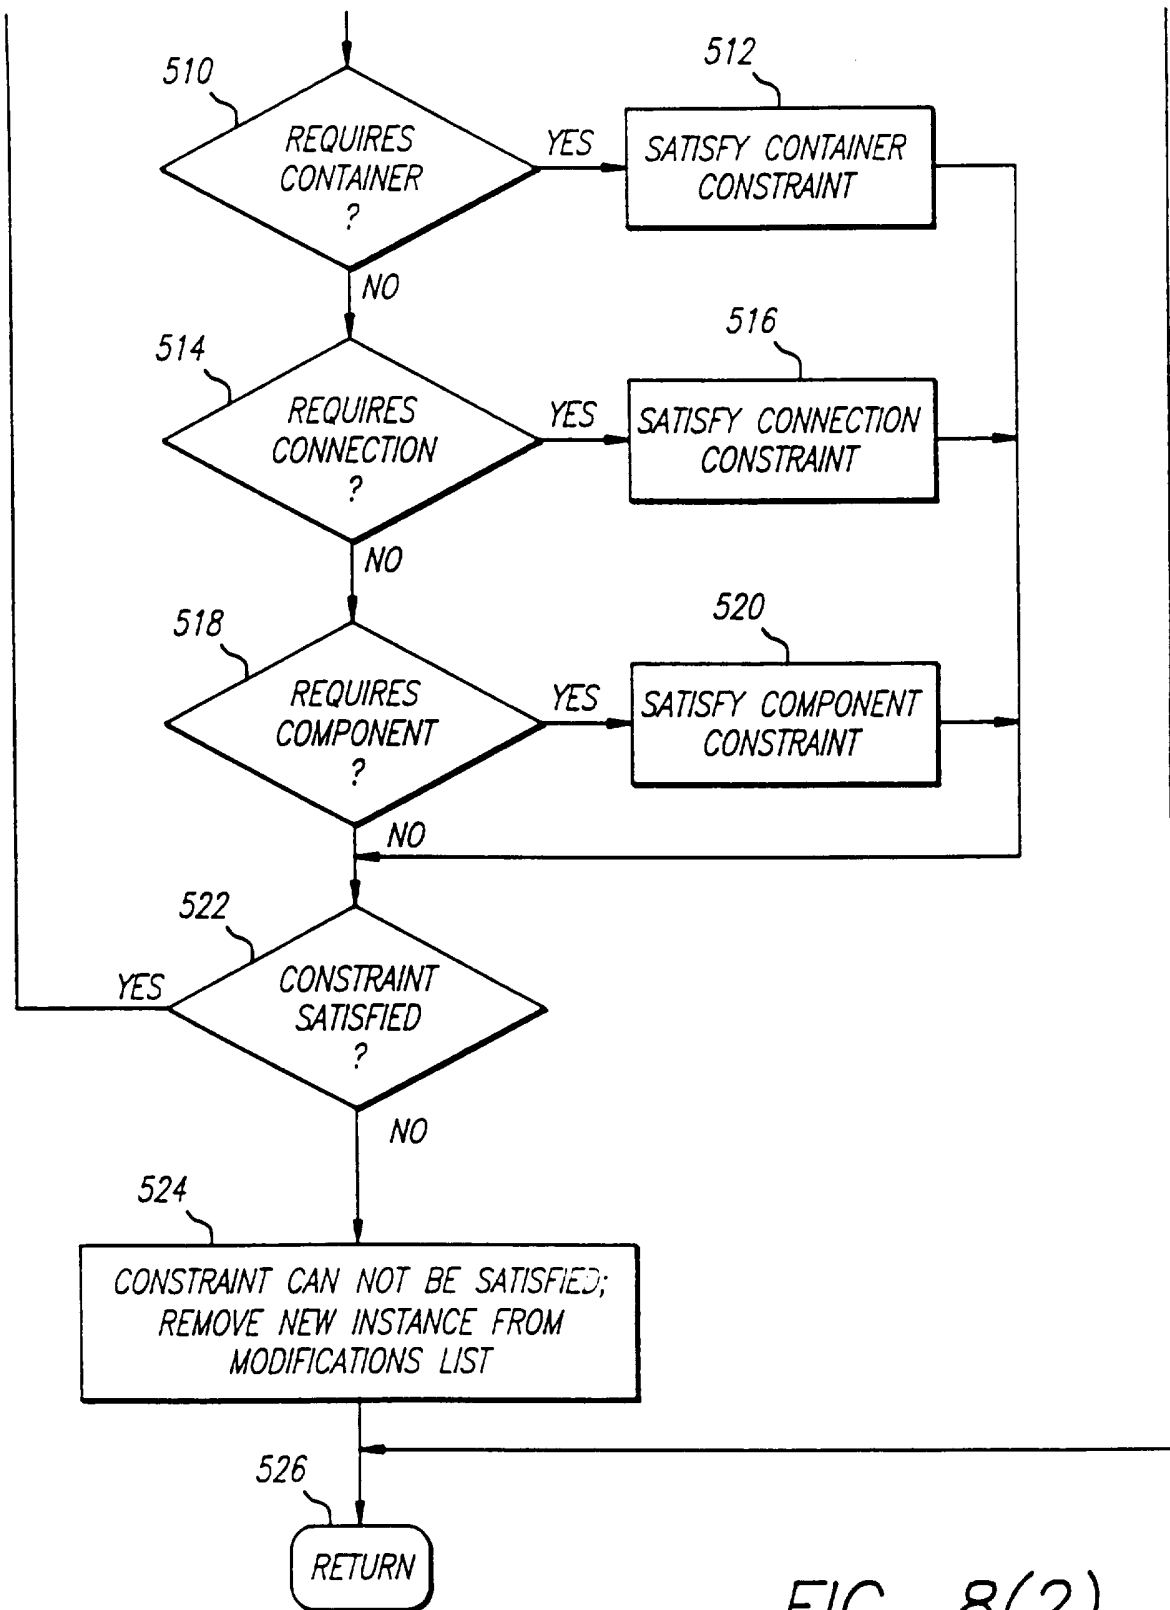
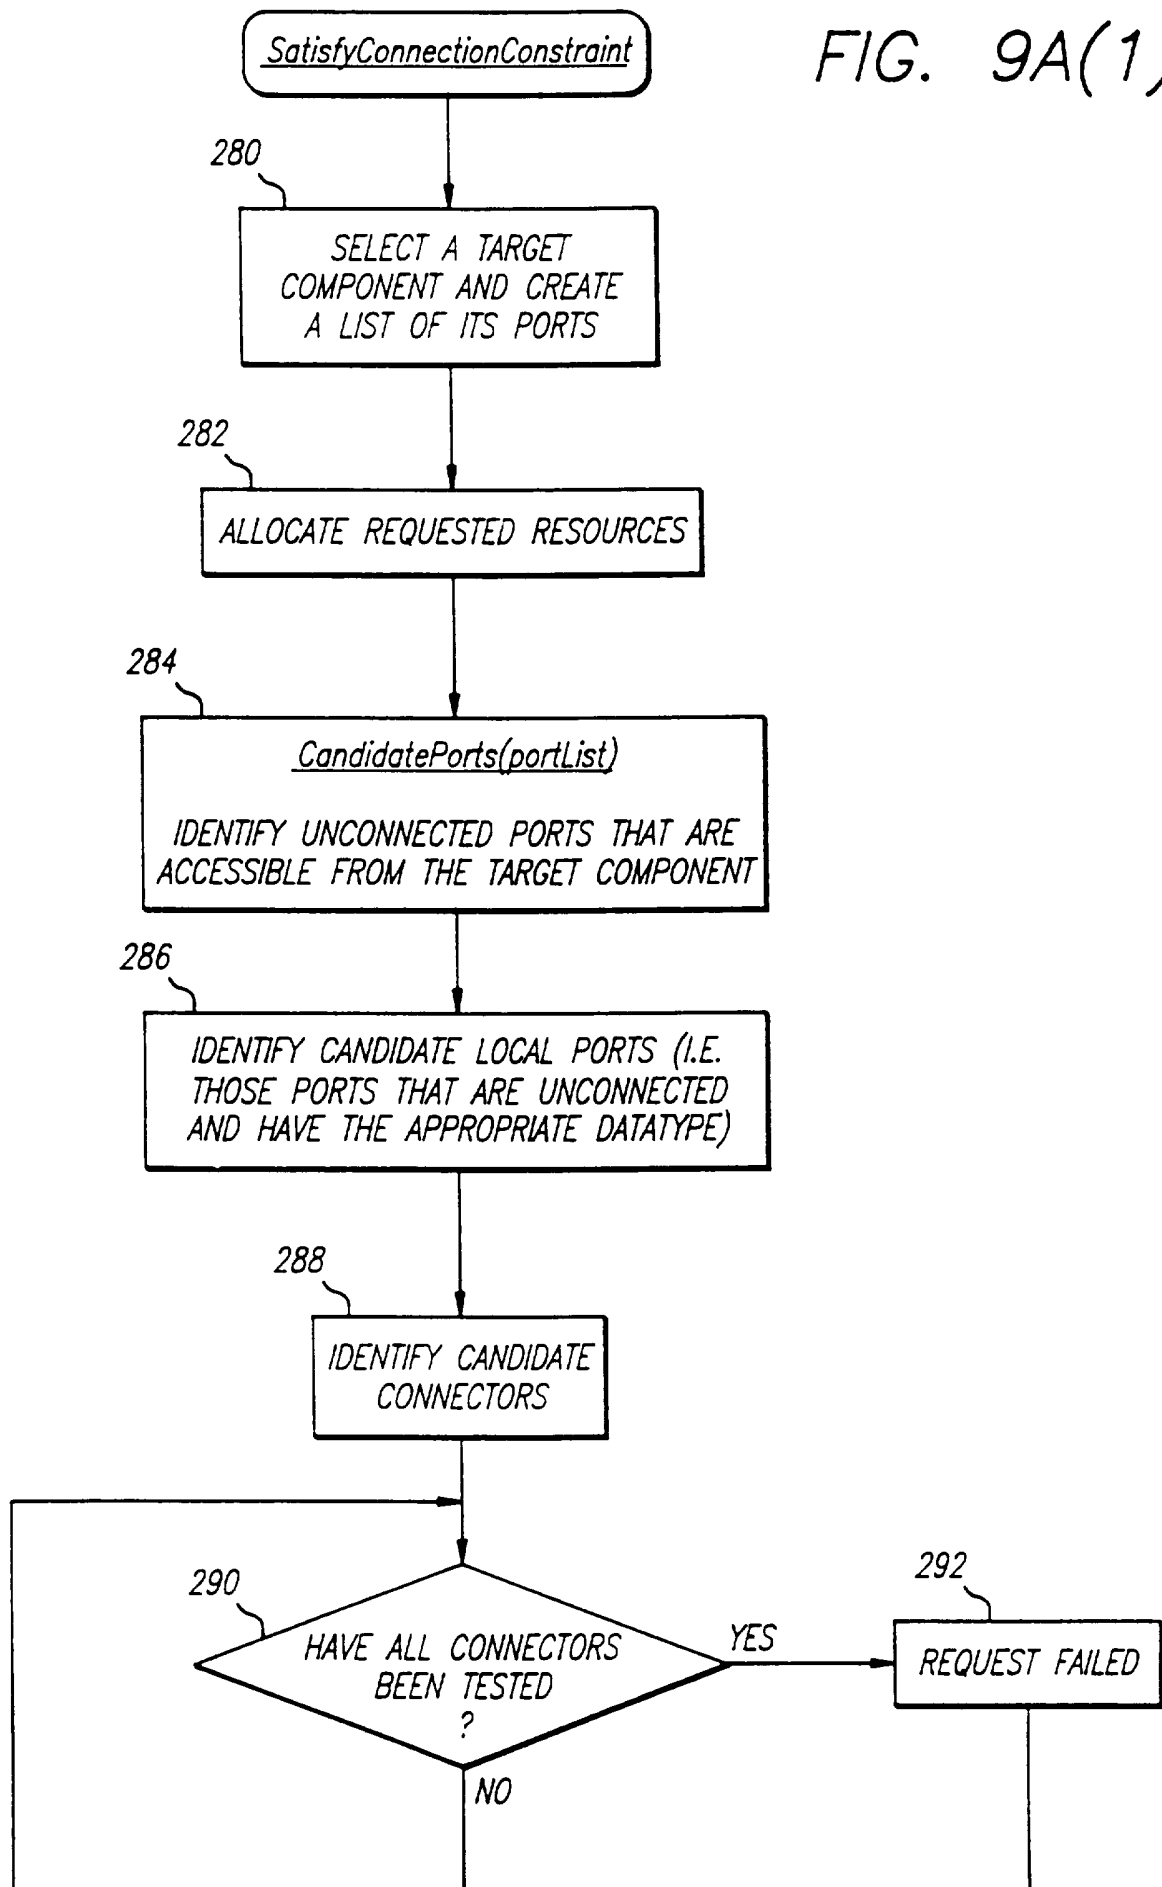


FIG. 8(2)

14/33

FIG. 9A(1)



15/33

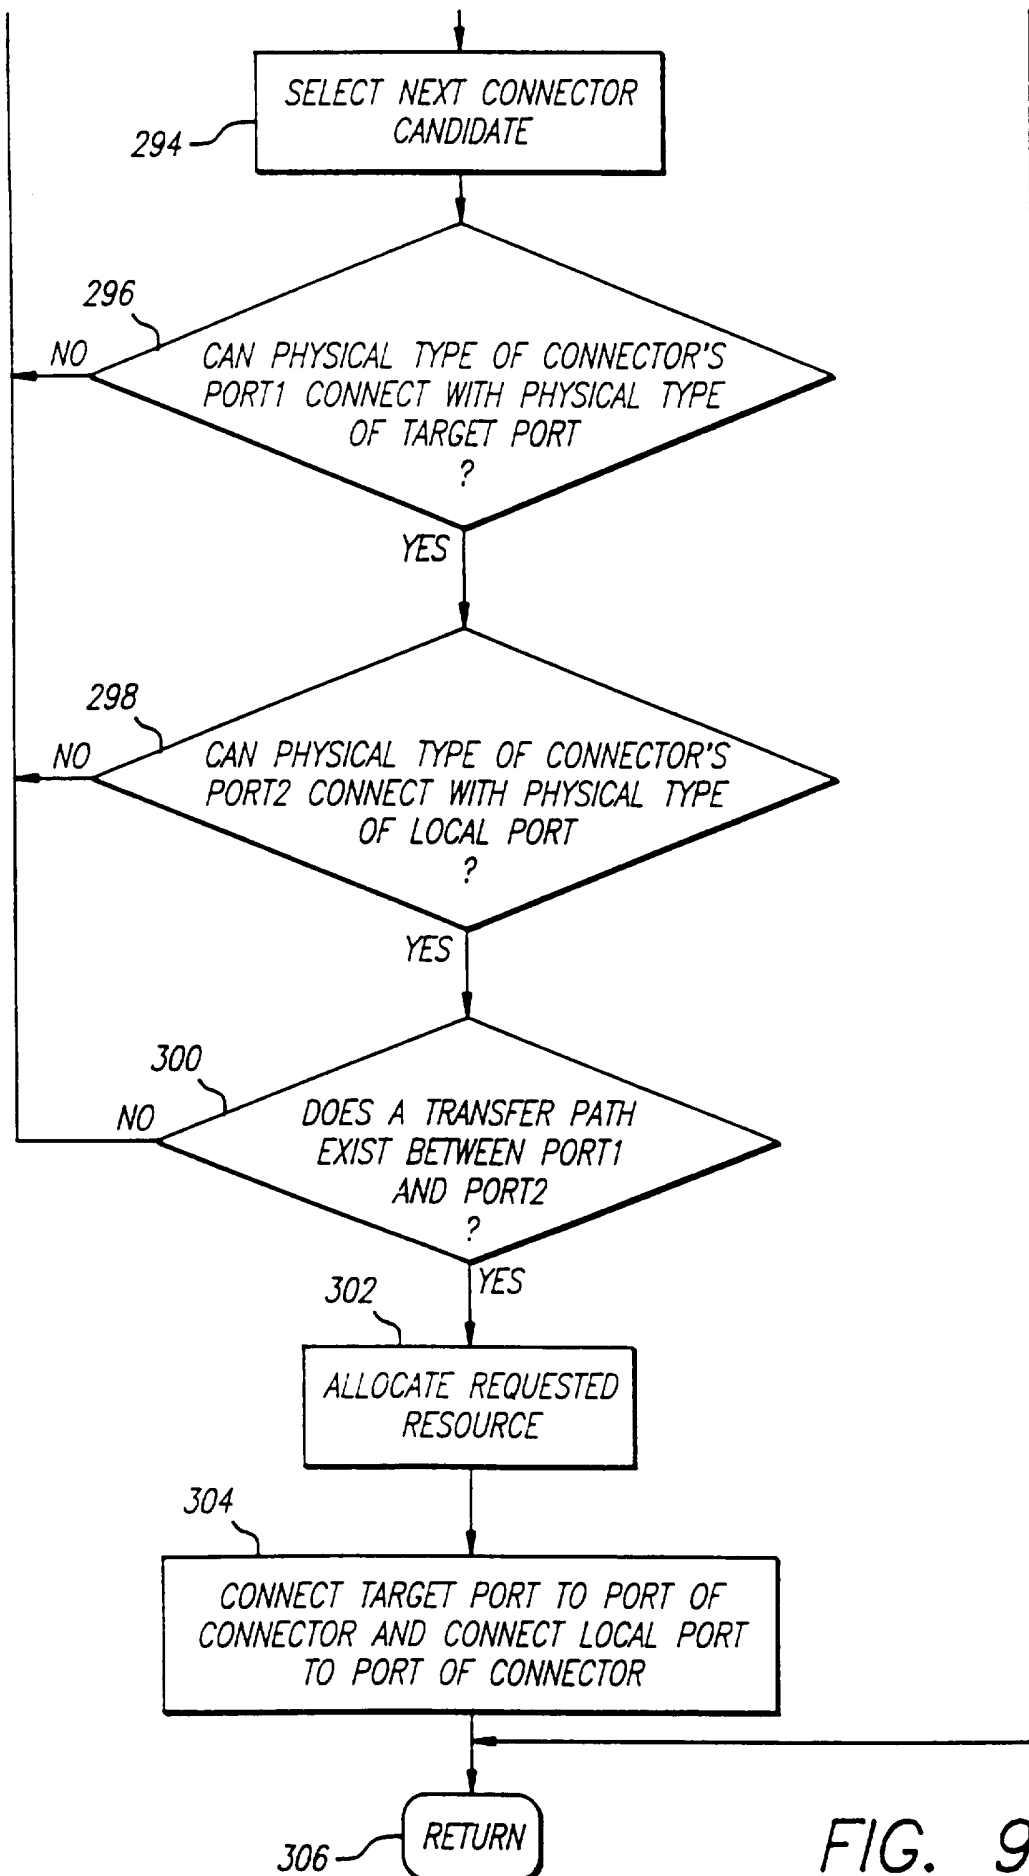


FIG. 9A(2)

16/33

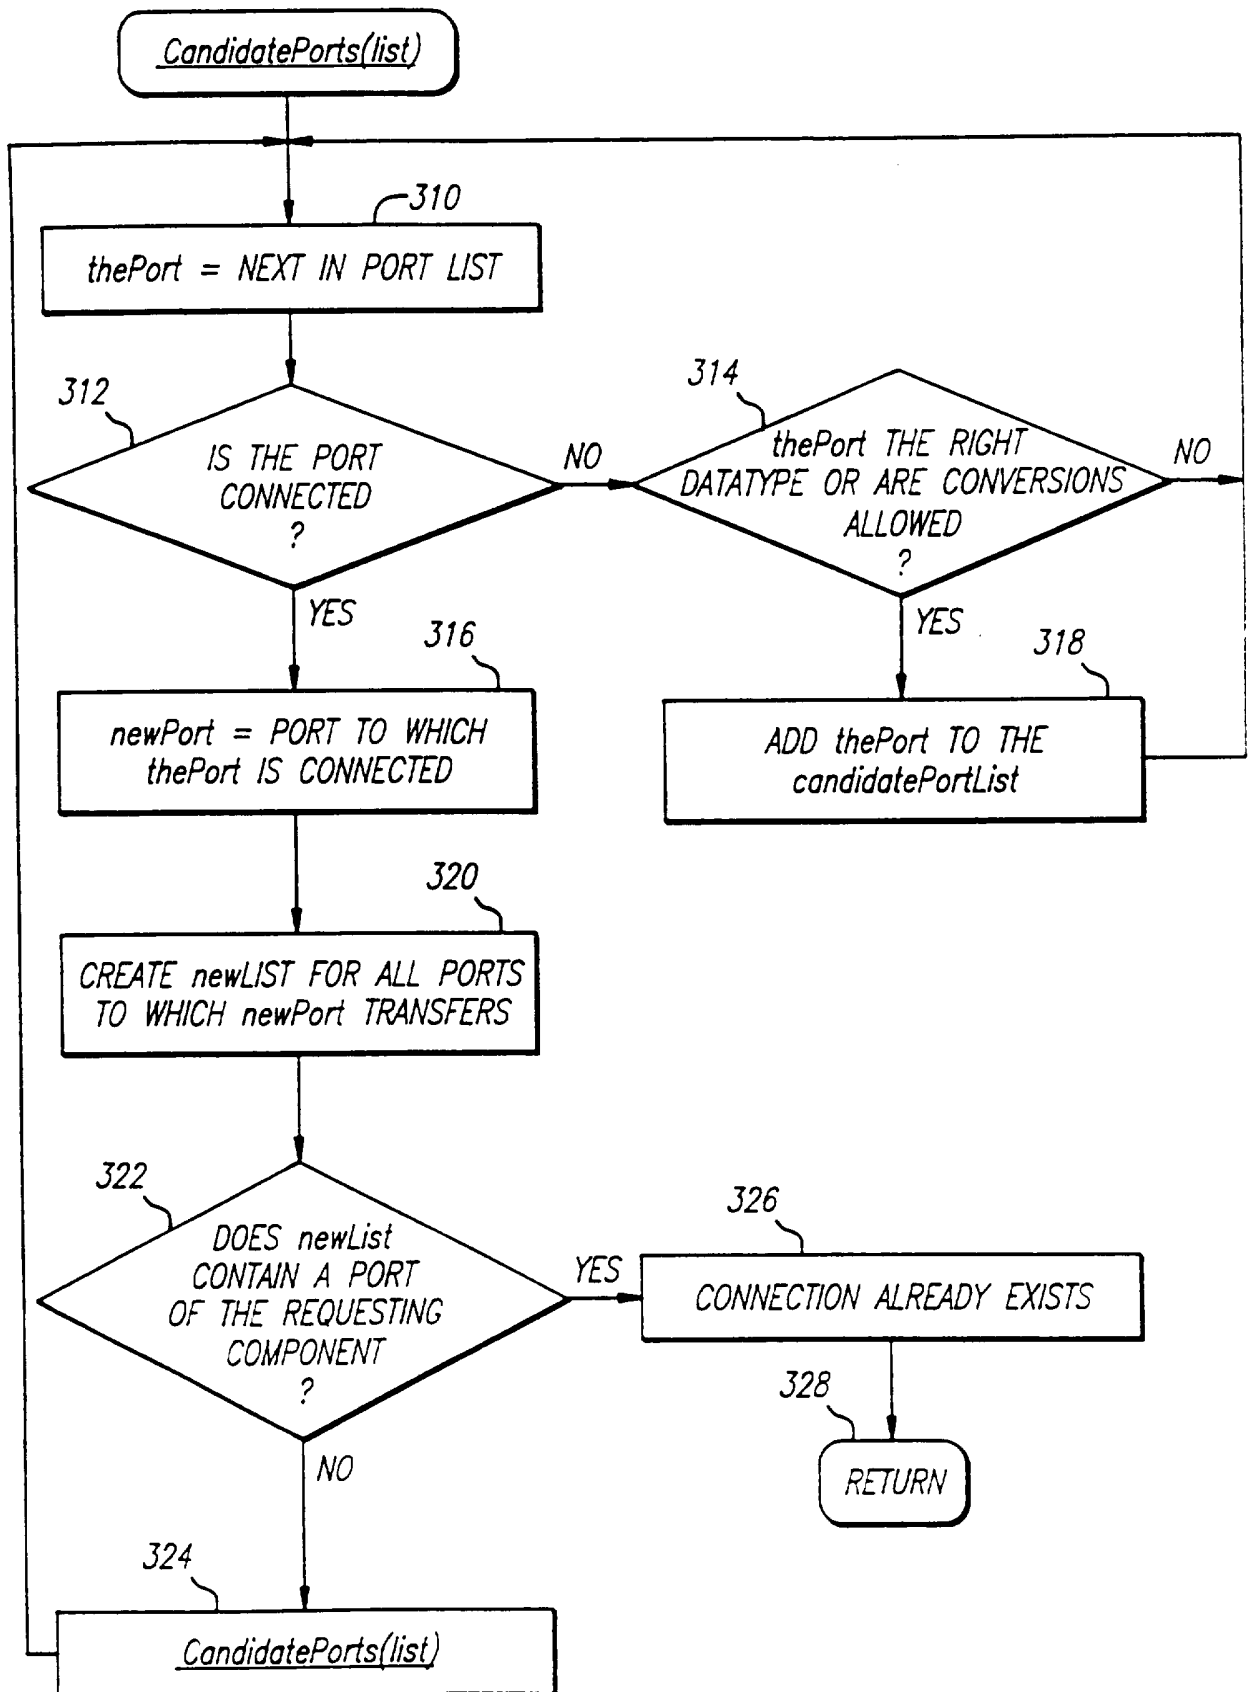
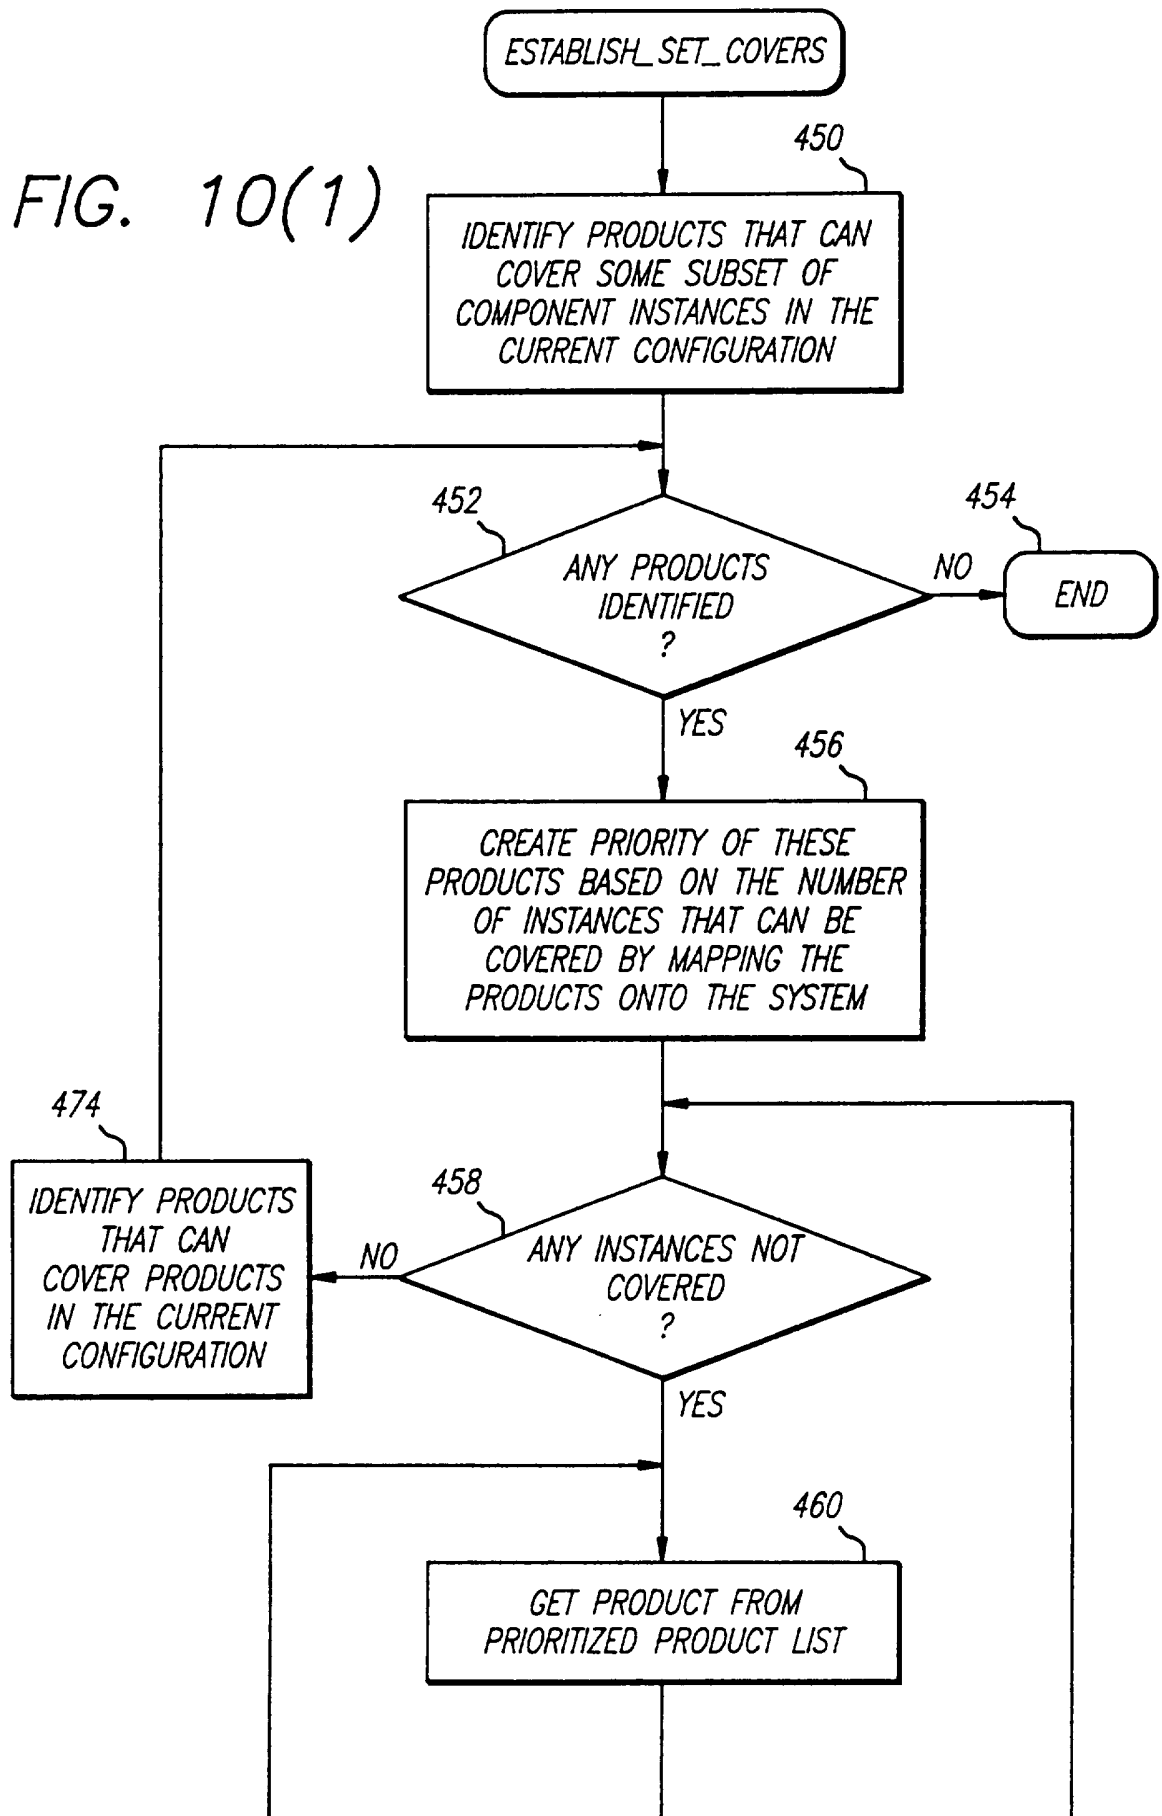


FIG. 9B

FIG. 10(1)



18/33

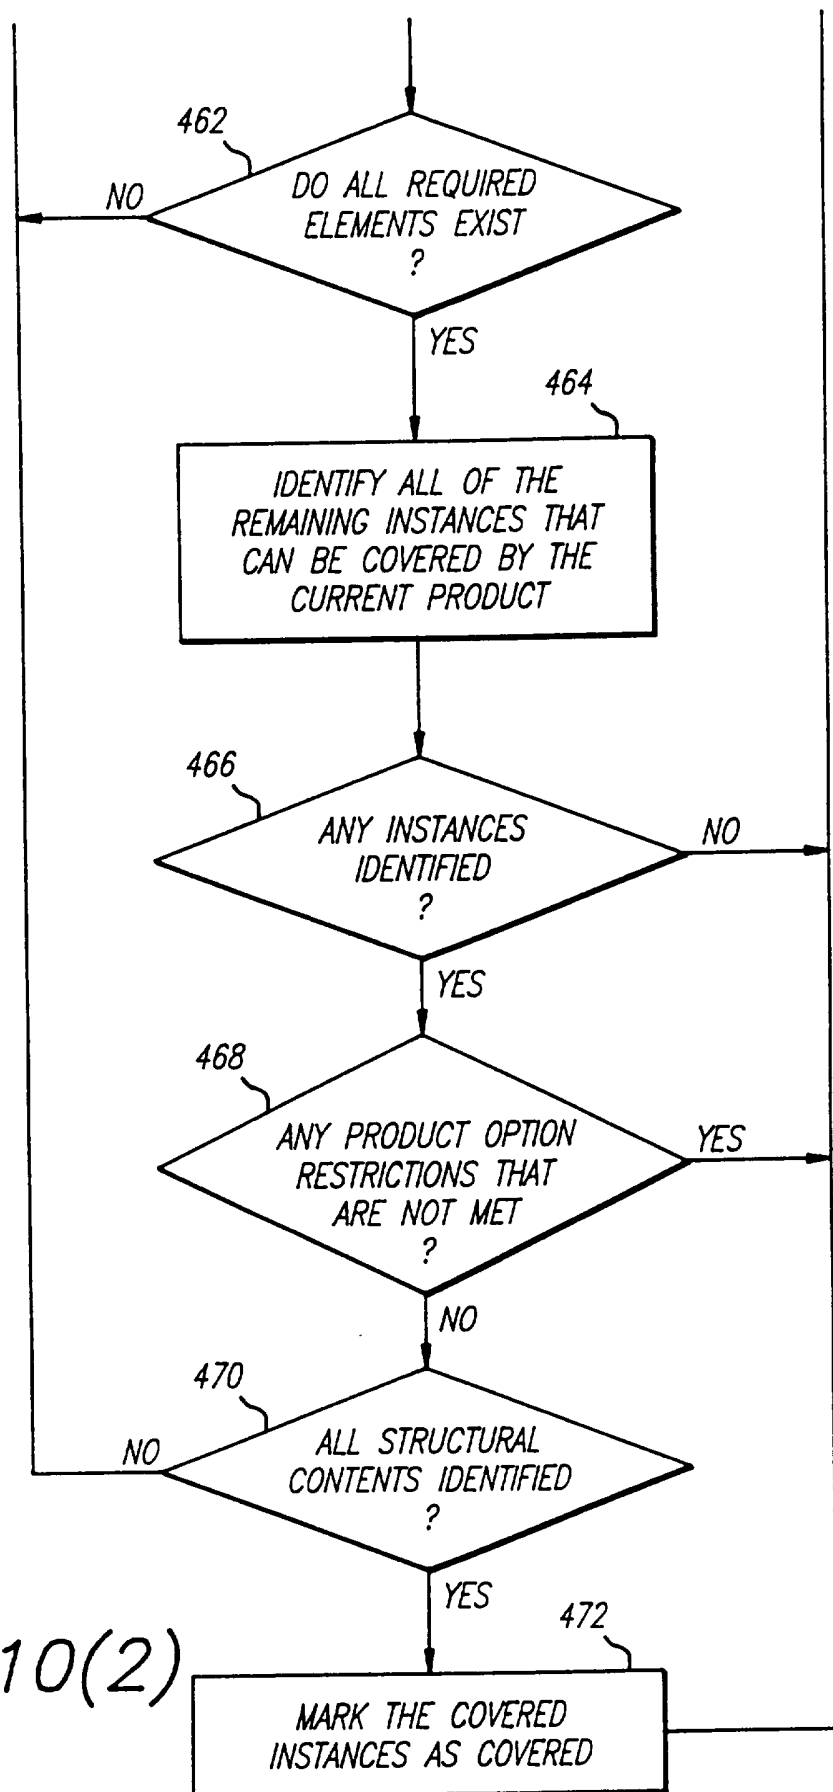


FIG. 10(2)


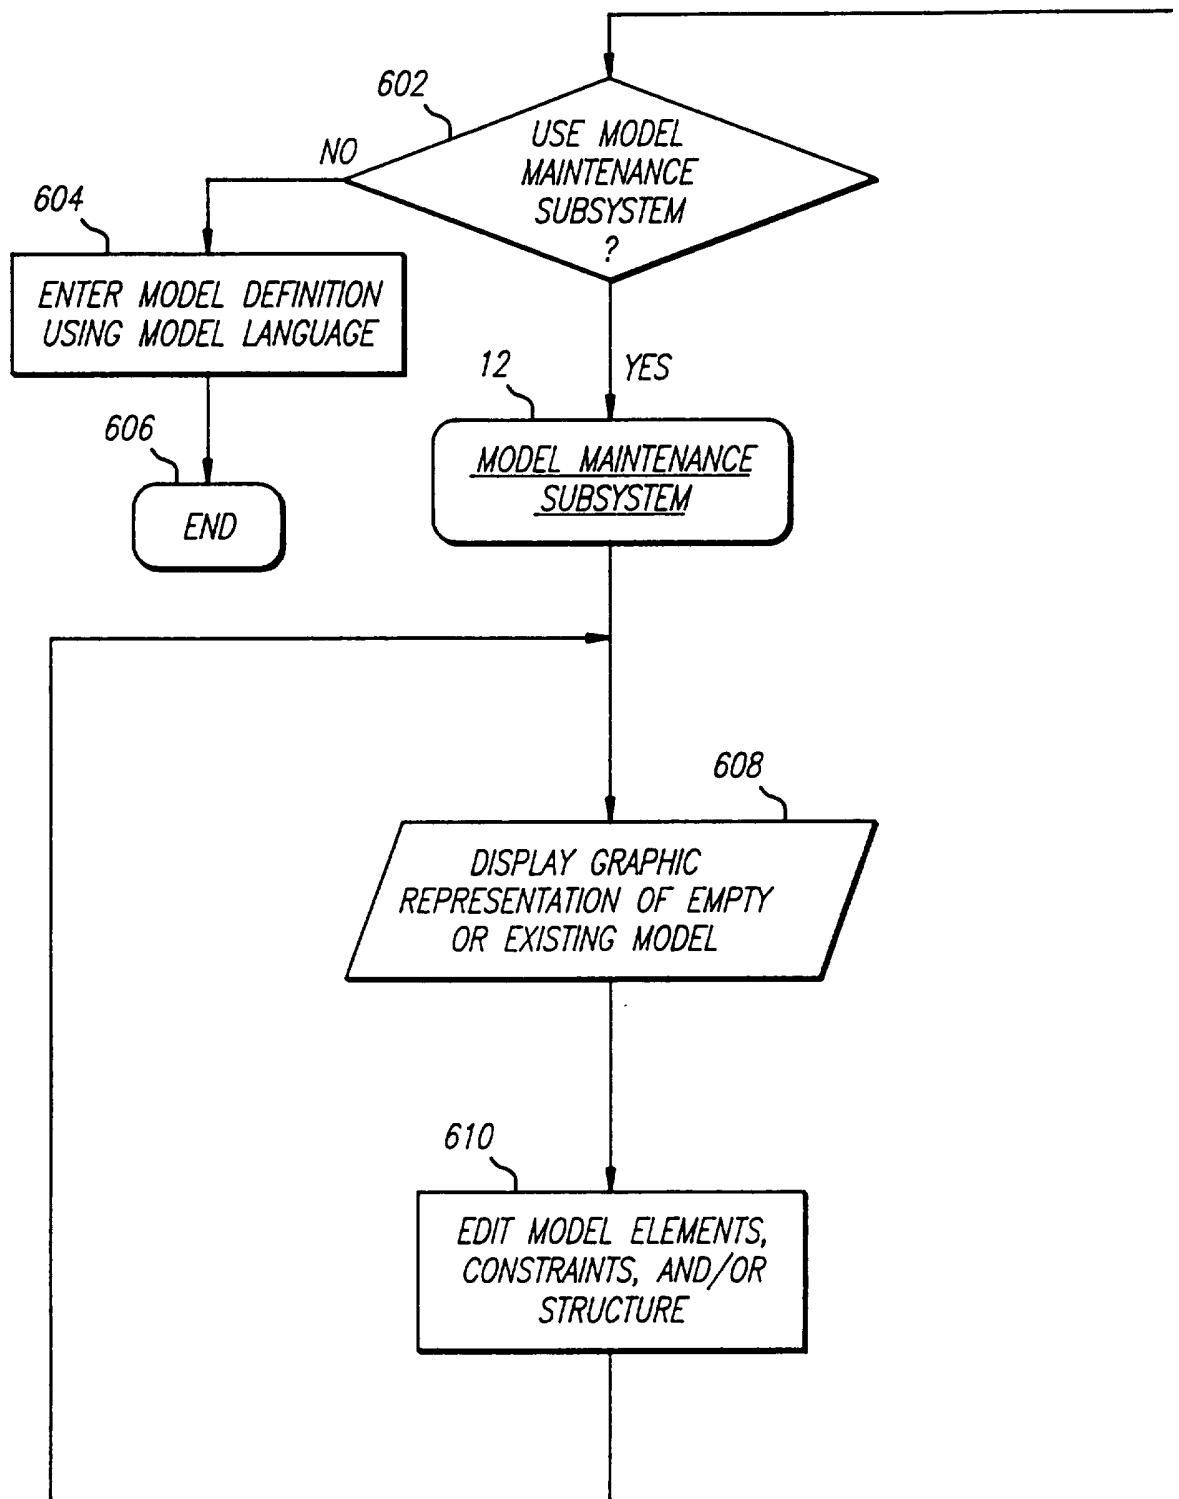
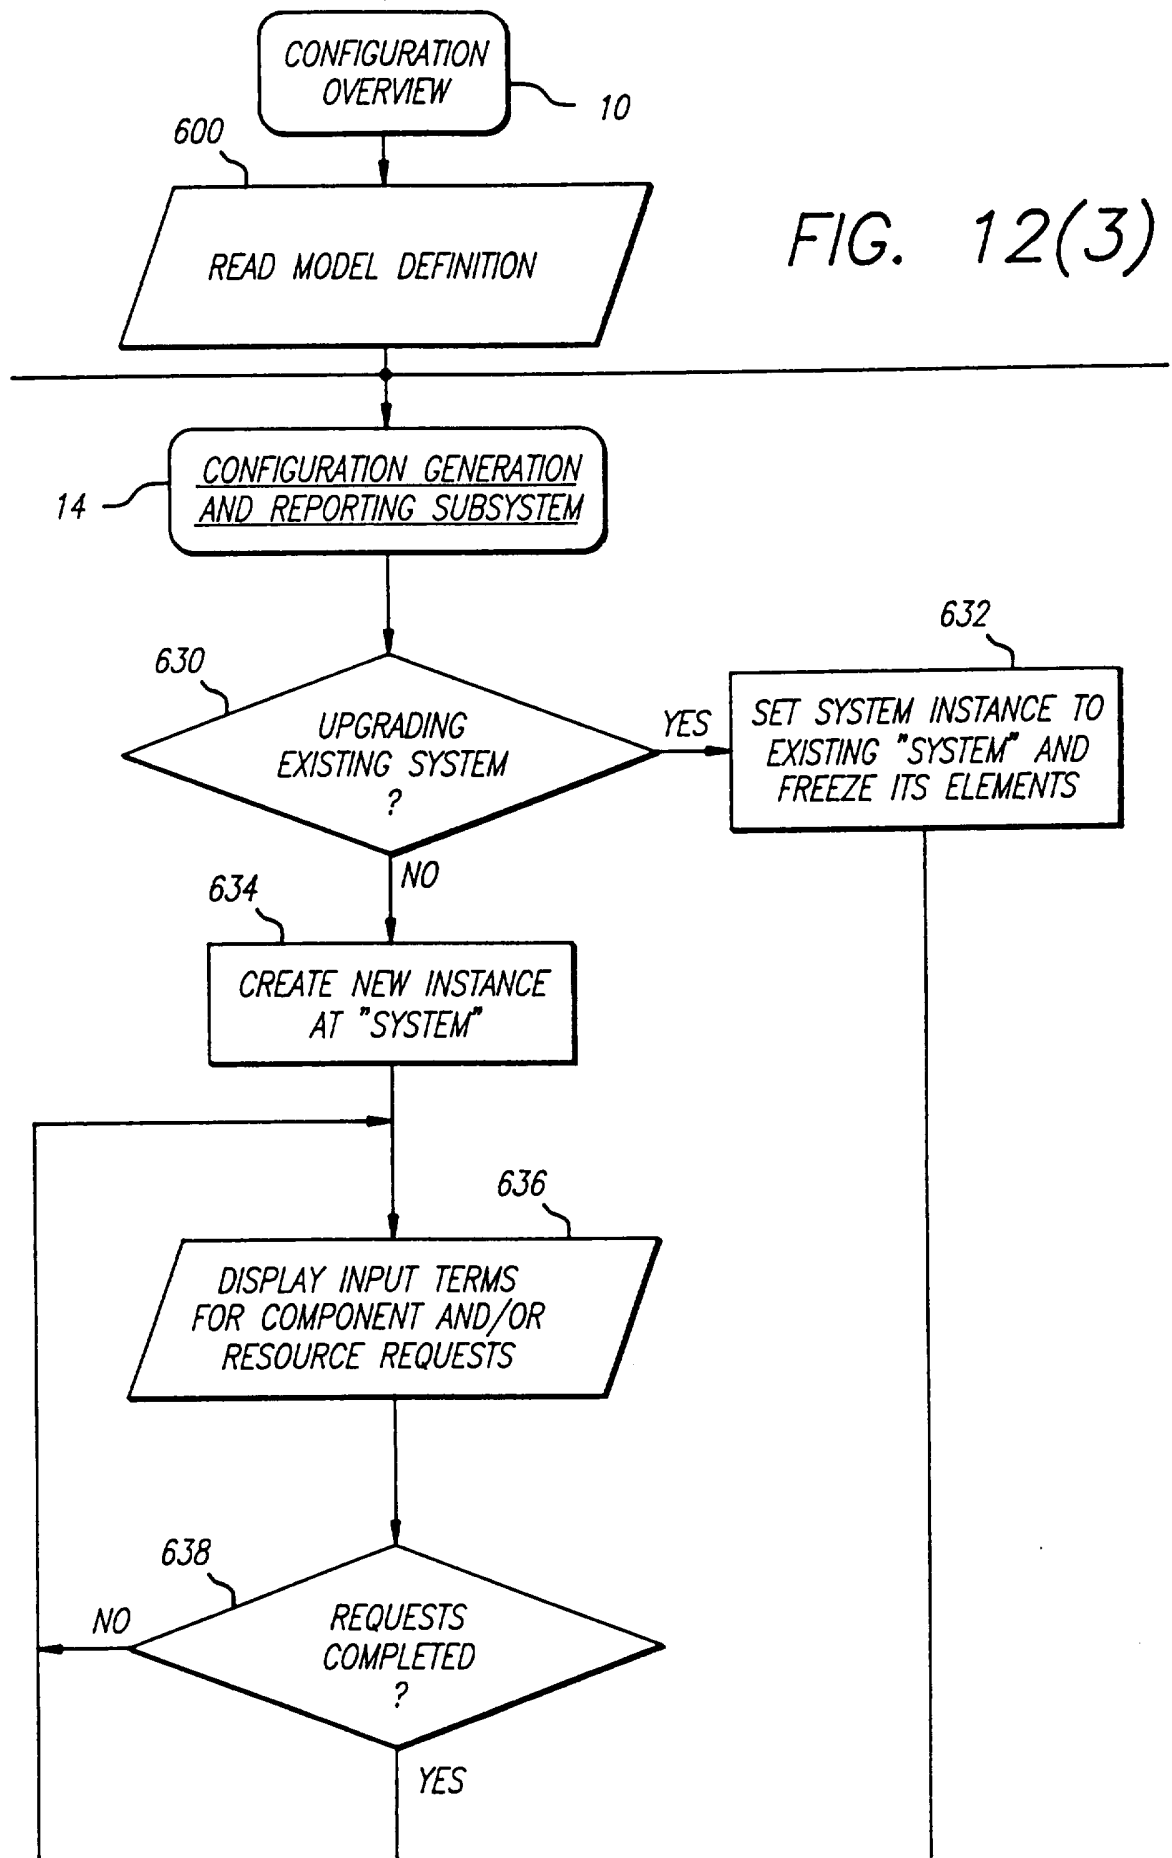


FIG. 11

20/33


FIG. 12(1)





22/33

FIG. 12(5)



23 / 3 3

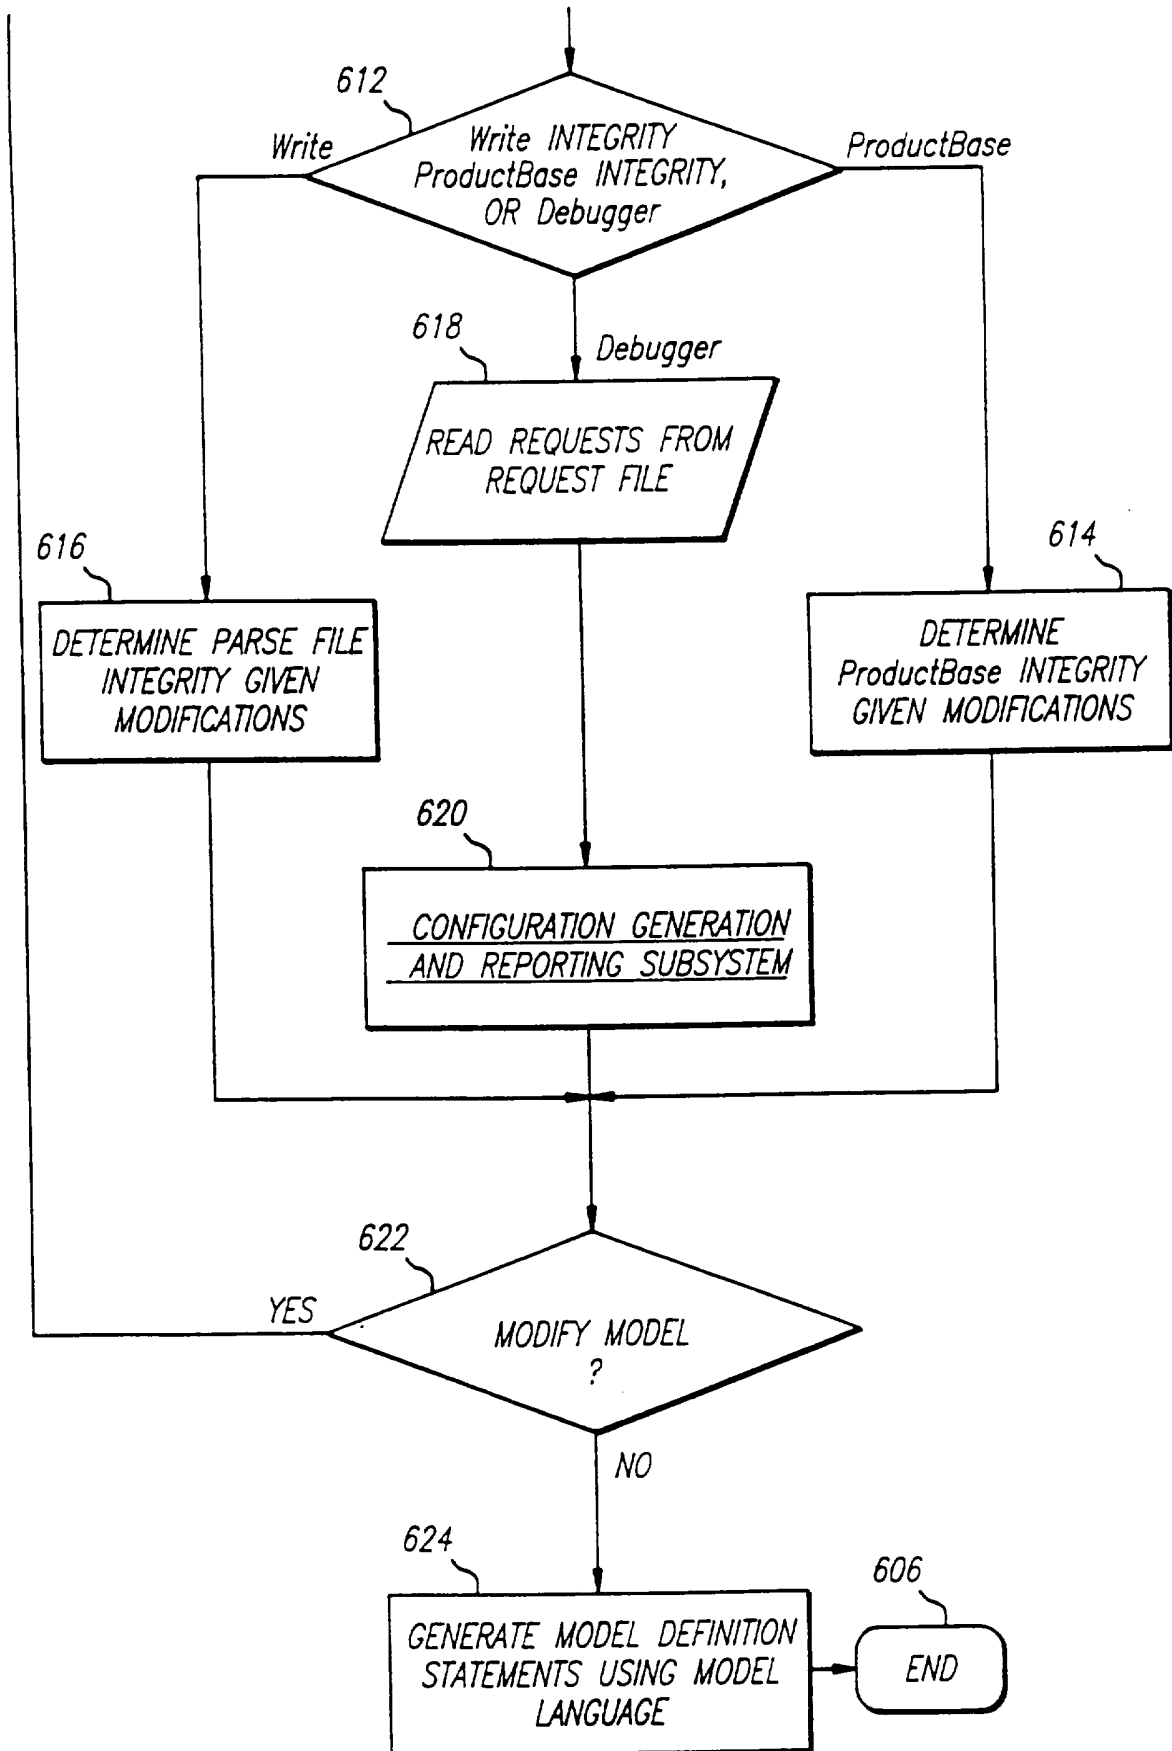


FIG. 12(2)

24/33

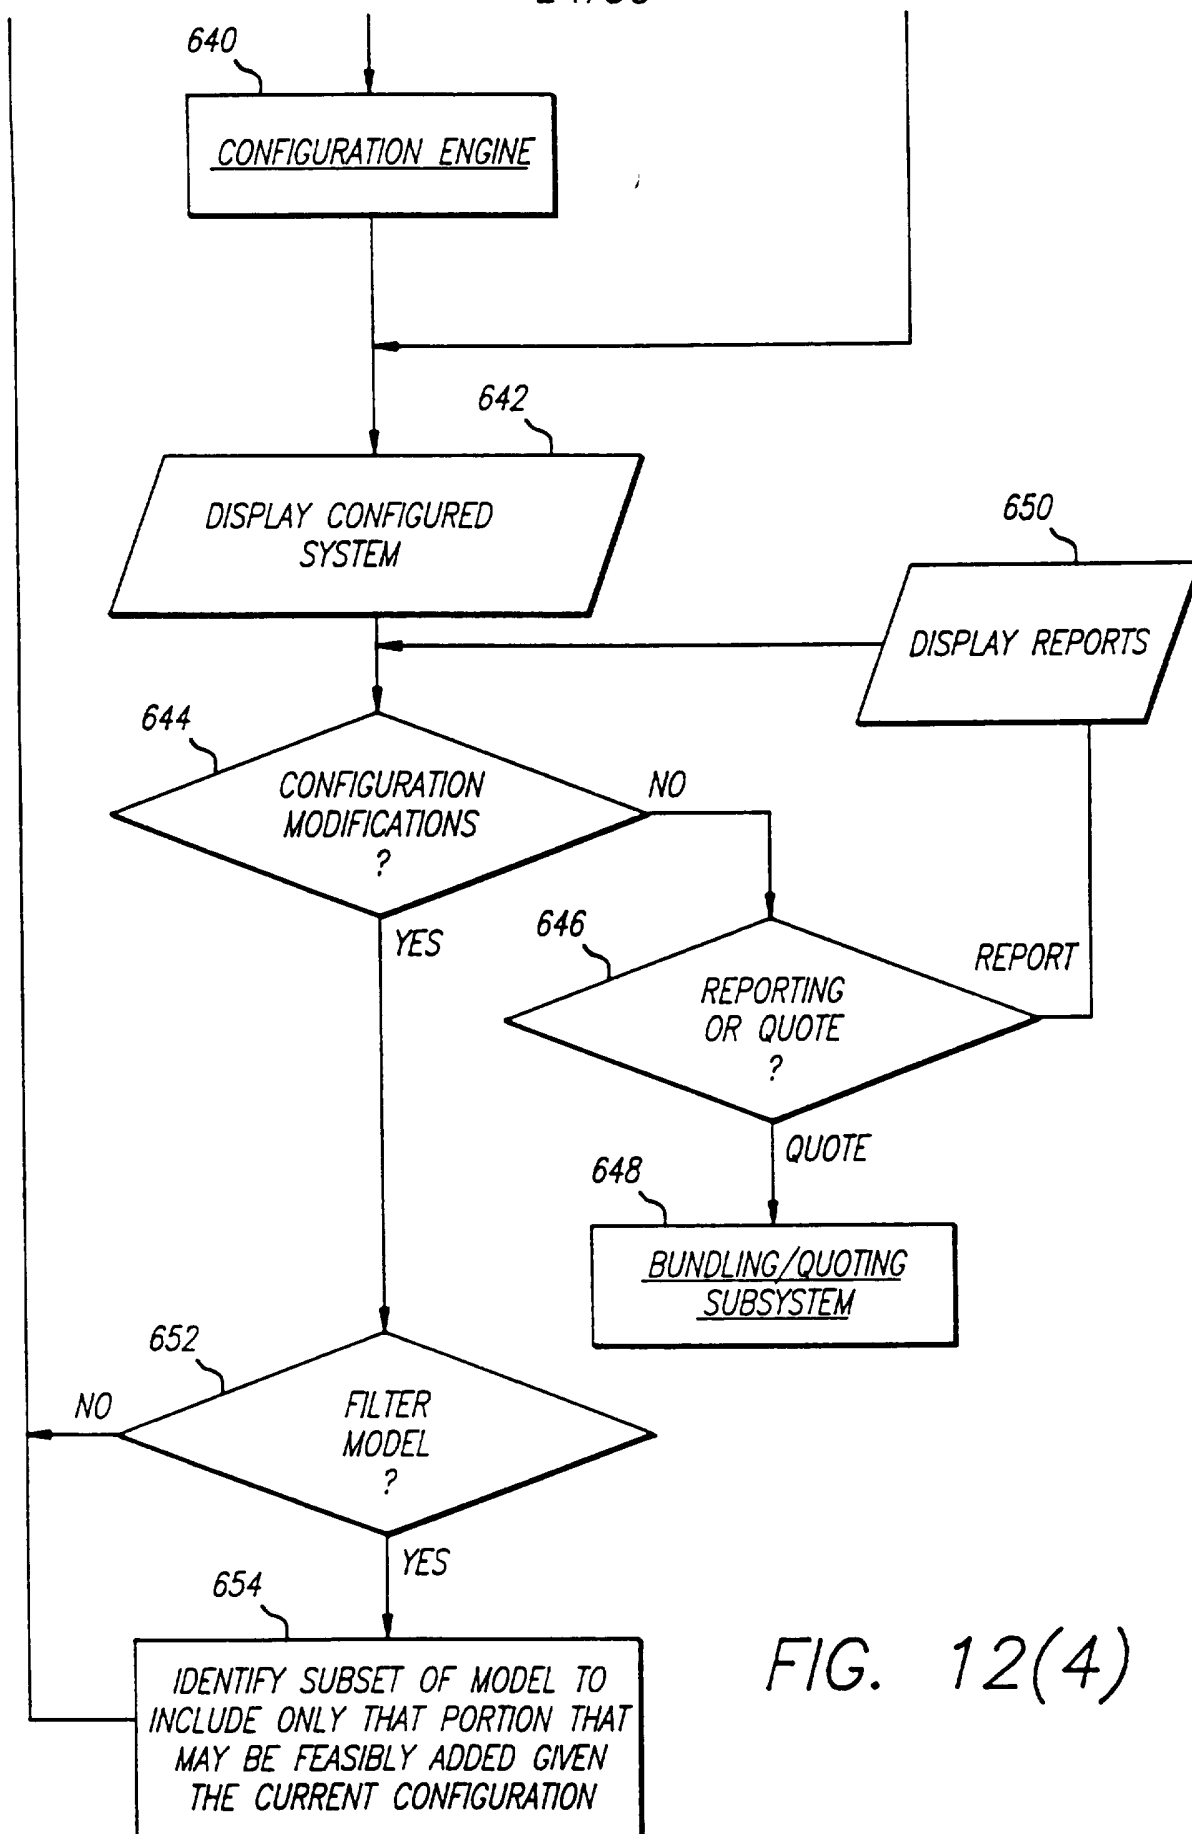
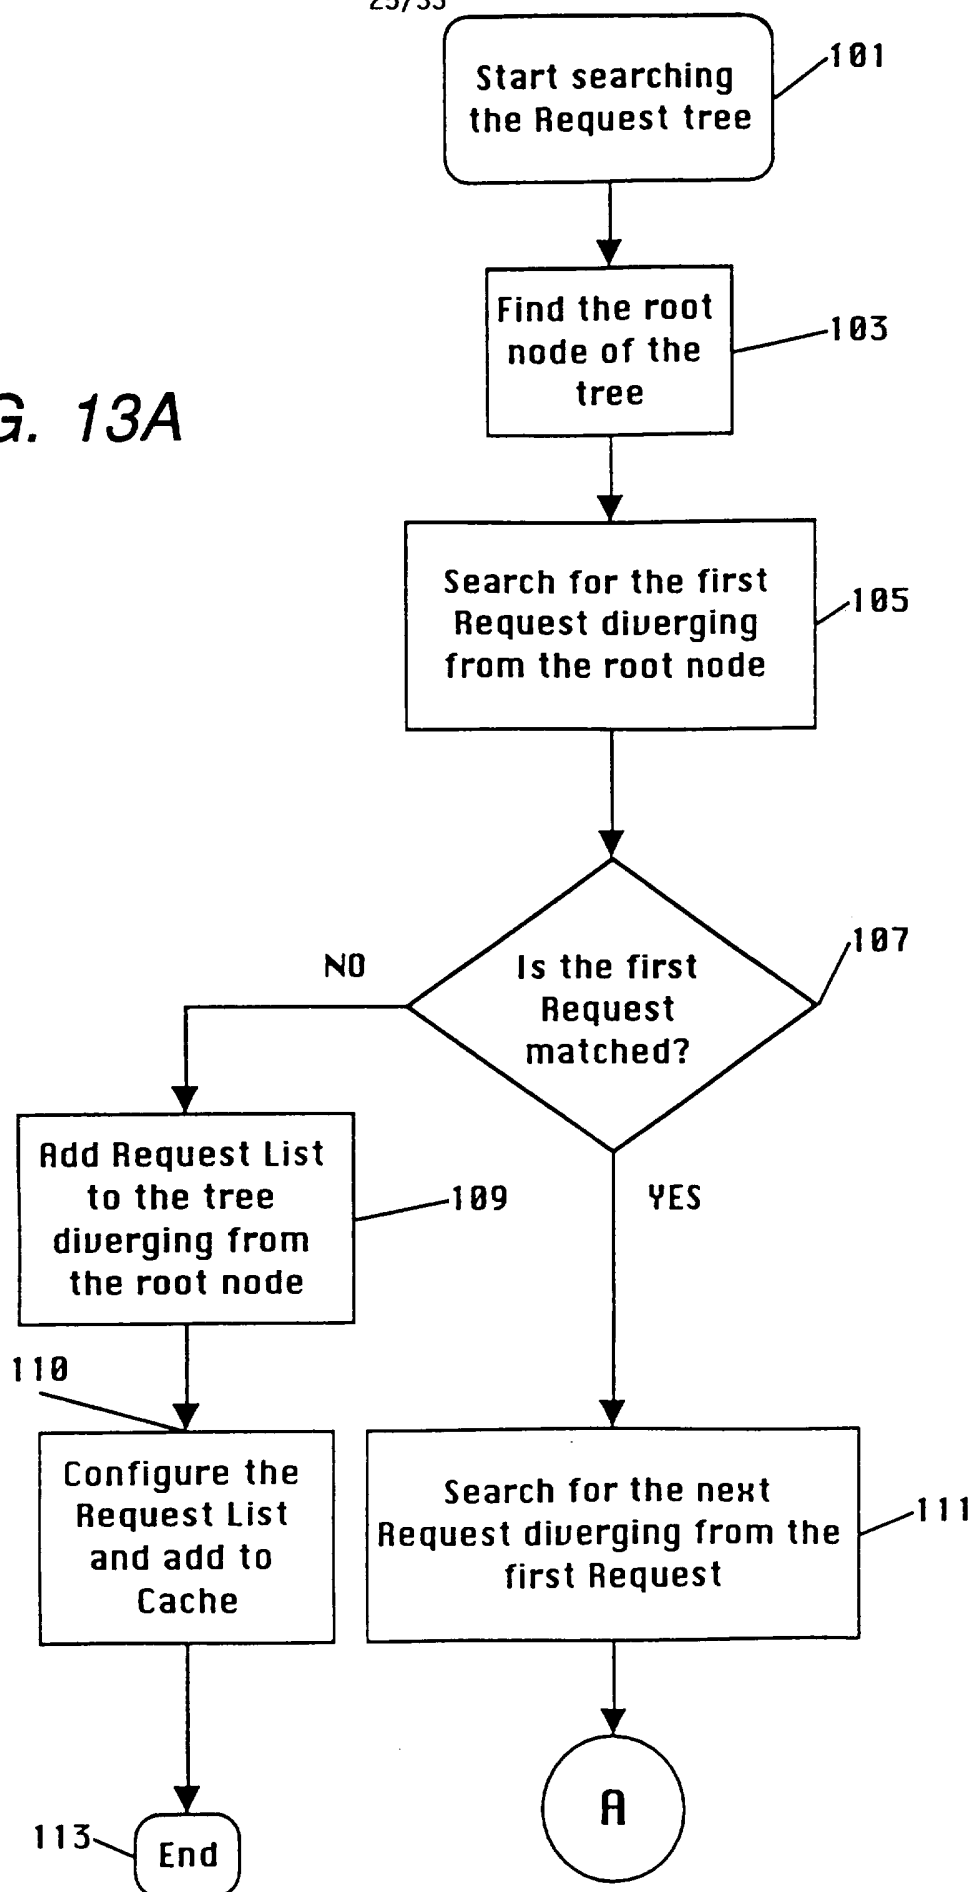


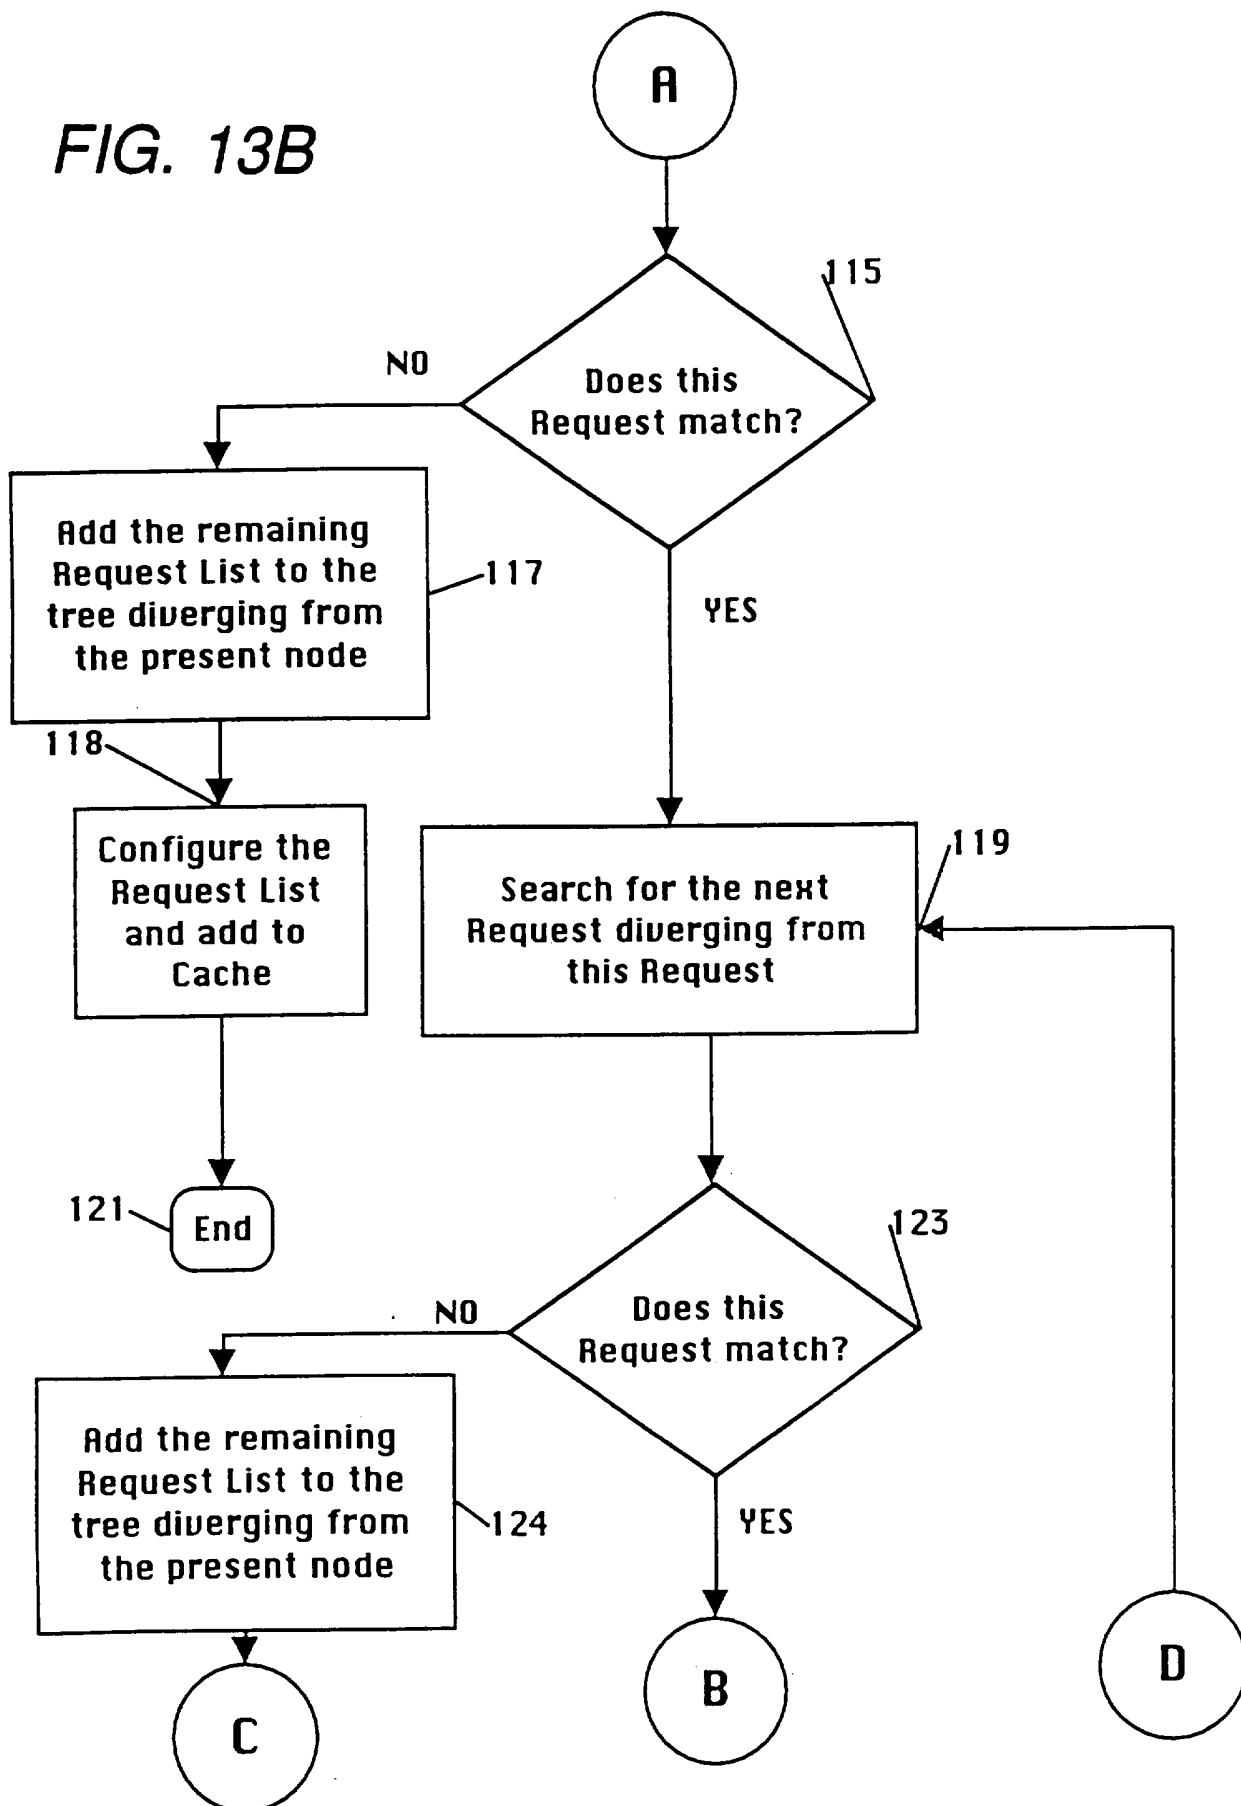
FIG. 12(4)

FIG. 13A

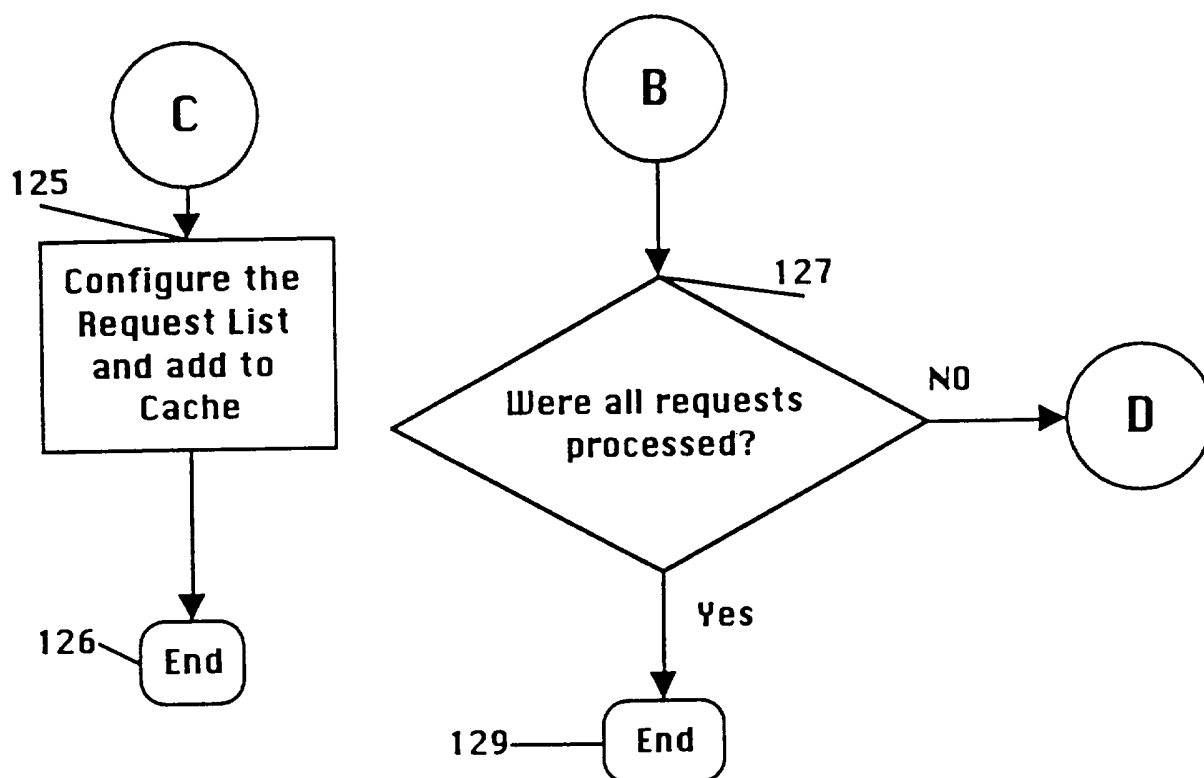


26/33

FIG. 13B



27/33

*FIG. 13C*

28/33




FIG. 14

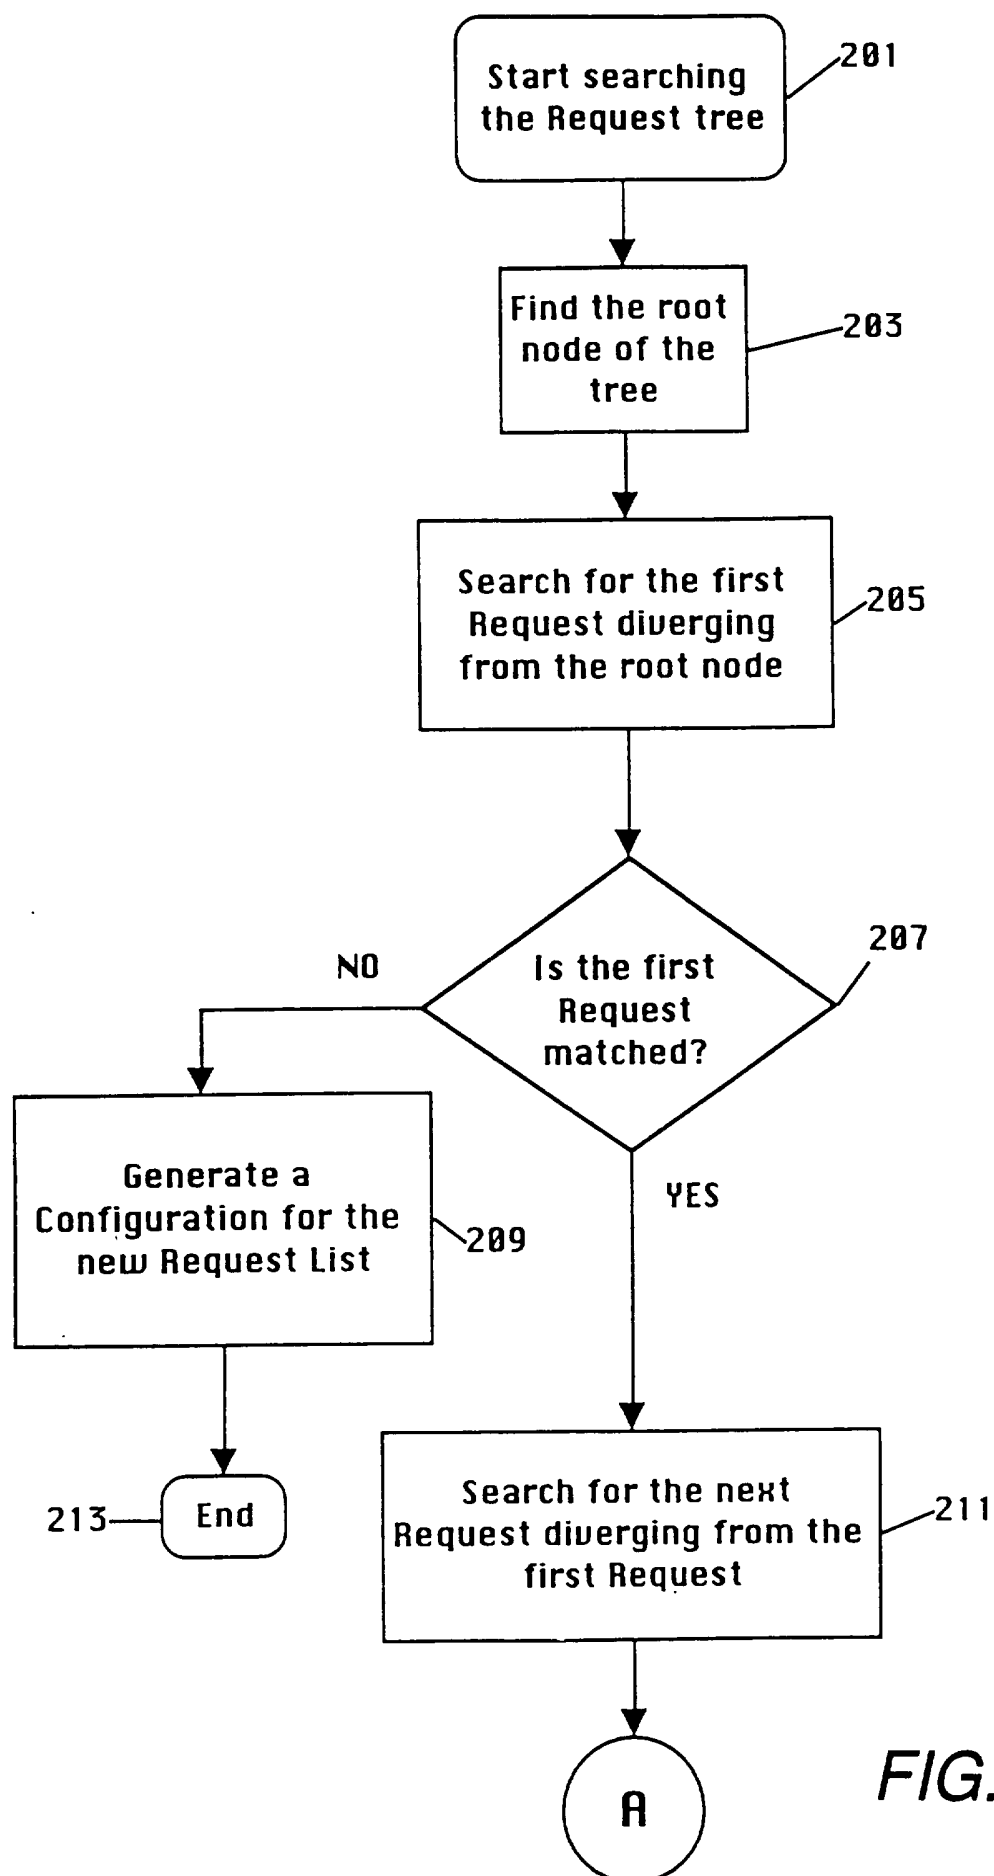
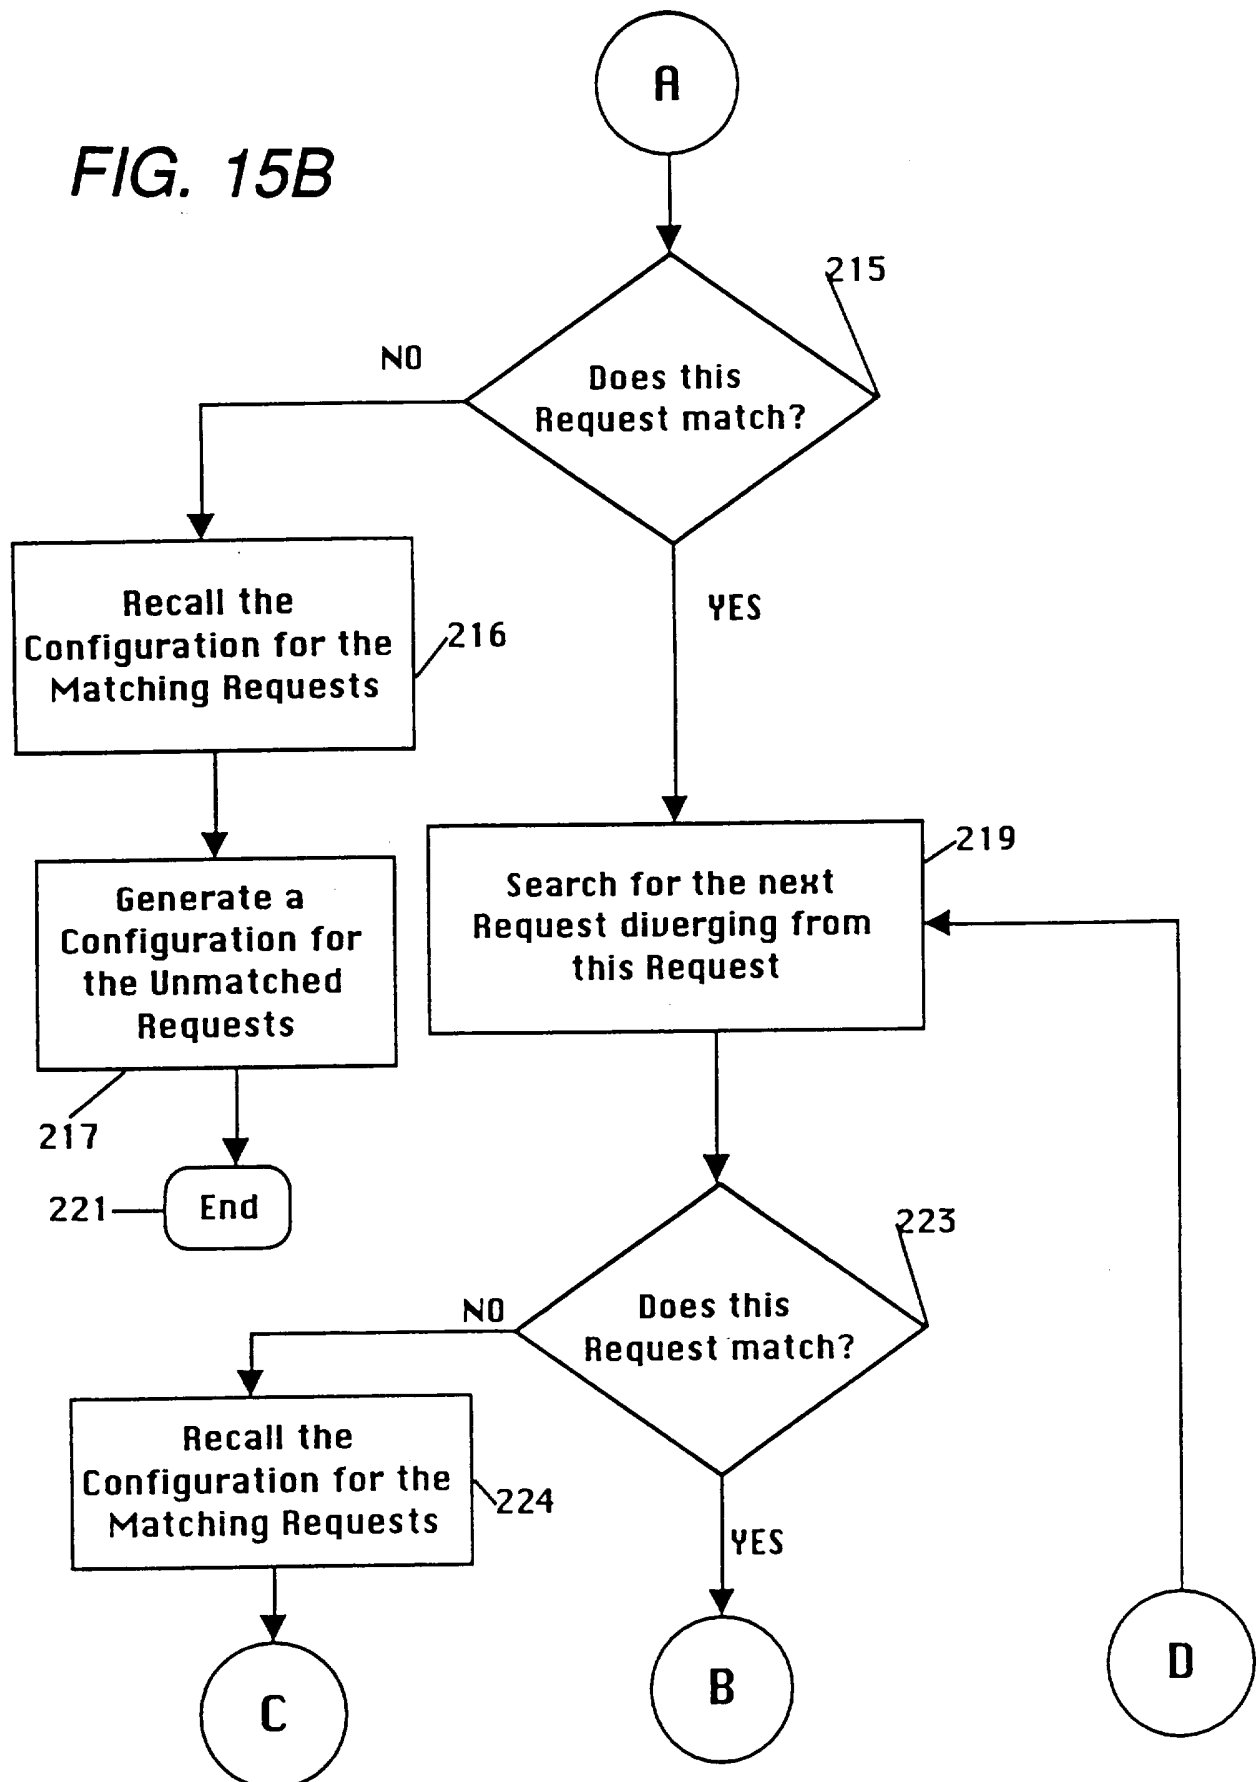


FIG. 15A

FIG. 15B

31/33

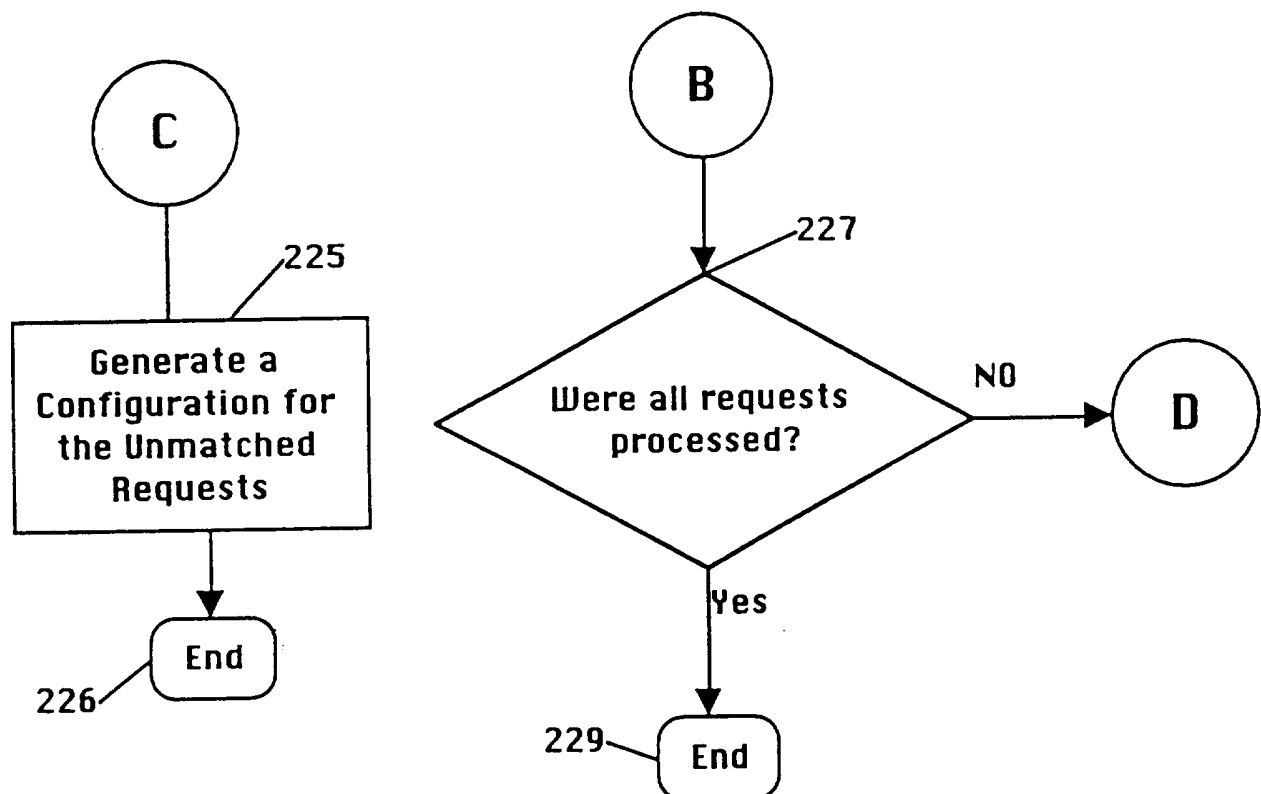



FIG. 15C

32/33

*FIG. 16*

33/33




FIG. 17

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US96/03406

A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) : G06F 15/177

US CL : 395/500

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 395/500

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

Please See Extra Sheet.

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US, A, 4,704,676 (FLANAGAN ET AL) 03 November 1987, See the entire Document.	1-22
A	US, A, 4,367,473 (CHU ET AL) 22 NOVEMBER 1994, See the Entire Document.	1-22
A	US, A, 5,367,635 (BAUER ET AL) November 22, 1994, See the Entire Document.	1-22



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents:	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"A" document defining the general state of the art which is not considered to be part of particular relevance	"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"E" earlier document published on or after the international filing date	"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents; such combination being obvious to a person skilled in the art
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"G" document member of the same patent family
"O" document referring to an oral disclosure, use, exhibition or other means	
"P" document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search

29 MAY 1996

Date of mailing of the international search report

09 JUL 1996

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

KYN MOHAMED

Telephone No. (703) 305-9694

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US96/03406

B FIELDS SEARCHED

Electronic data bases consulted (Name of data base and where practicable terms used):

APS

l1 = host(10a)configur####; l2 = configur####(10a)(pc or computer); l3 = request###; l4 = cache#; l5 = (path# or bus#)(30a)(branch### or tree#); l6 = match####(10a)(path# or bus##); l