

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
12 December 2002 (12.12.2002)

PCT

(10) International Publication Number
WO 02/099635 A1

(51) International Patent Classification⁷: **G06F 9/44**

(21) International Application Number: PCT/US02/16838

(22) International Filing Date: 30 May 2002 (30.05.2002)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/295,201 1 June 2001 (01.06.2001) US

(71) Applicant (for all designated States except US): **THE JOHNS HOPKINS UNIVERSITY** [US/US]; 34th and Charles Streets, Baltimore, MD 21218 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **SCHEIDT, David, H.** [US/US]; 2709 Thyme Drive, Edgewater, MD 21037 (US). **WATSON, David, P.** [US/US]; 17 Ridge Road, Severna Park, MD 21146 (US). **ALGER, David, L.** [US/US]; 1410 Rio Grande Court, Severn, MD 21144 (US). **MC-CUBBIN, Christopher, B.** [US/US]; 1205 Birch Avenue,

Baltimore, MD 21227 (US). **VICK, Shon, D.** [US/US]; 2805 Bartol Avenue, Baltimore, MD 21209 (US).

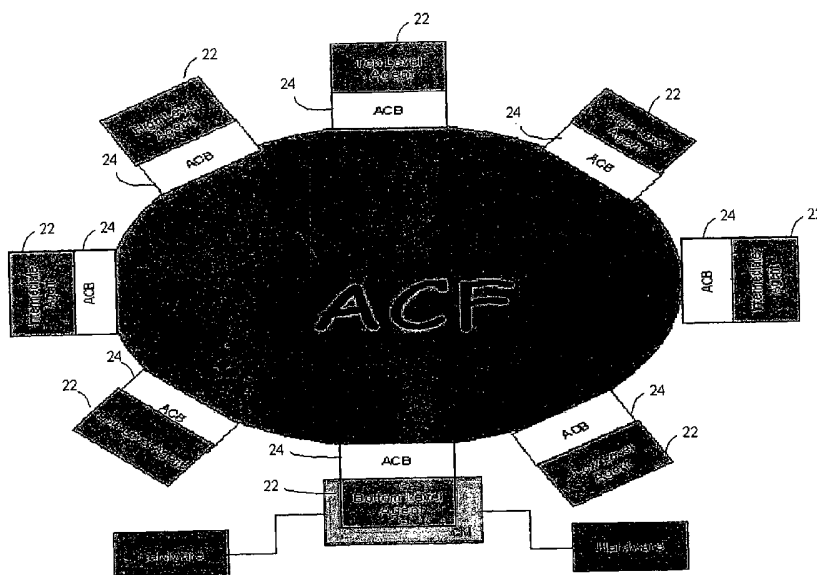
(74) Agents: **GRAF, Ernest, R.** et al.; The Johns Hopkins University, Applied Physics Laboratory, 11100 Johns Hopkins Road, Laurel, MD 20723-6099 (US).

(81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZM, ZW.

(84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

[Continued on next page]

(54) Title: SYSTEM AND METHOD FOR AN OPEN AUTONOMY KERNEL (OAK)



(57) Abstract: The Open Autonomy Kernel (OAK) addresses critical infrastructure requirements for next generation autonomous and semi-autonomous systems (24), including performance tracking, anomaly detection, diagnosis, fault recovery, and plant "safing". OAK combines technologies in automated planning and scheduling, control agent-based systems (22), and model-based reasoning to form a portable software architecture (26), knowledge-base, and open Application Programming Interface (API) to enable integrated auxiliary subsystem autonomy.



WO 02/099635 A1



Published:

— with international search report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

SYSTEM AND METHOD FOR AN OPEN AUTONOMY KERNEL (OAK)**CROSS-REFERENCE TO RELATED APPLICATIONS**

[0001] This application claims the benefit of U.S. provisional application no. 60/295,201, filed on June 1, 2001, which is hereby incorporated by reference in its entirety.

STATEMENT OF GOVERNMENTAL INTEREST

[0002] This invention was made with U.S. Government support under the Office of Naval Research, Arlington, Virginia under contract number N00014-00-C-0050. The U.S. Government has certain rights in the invention.

BACKGROUND OF THE INVENTION

[0003] Efforts to automate the control of complex connection-based systems such as, for instance, engineering plants aboard naval vessels have emphasized the infrastructure and diagnostic aspects of plant management, i.e., monitoring subsystems via sensors and presentation of the sensor data to human operators. Interpretation of and response to the data remain largely manual tasks. This interpretation and response function, especially in damage control scenarios, is a significant factor in determining manpower levels. If the incident assessment and response loop can be closed with a reliable autonomous reasoning process, significant relief in overall manpower levels can be realized. The best automation efforts to date have been based on expert diagnostic knowledge in the form of coded rules or procedures that are interpreted by the system at runtime to detect, predict, or diagnose fault conditions. However, even the best automation efforts require significant amounts of human diagnosis and input.

SUMMARY

[0004] The open autonomy kernel (OAK) architecture of the present invention extends an automated reasoning paradigm in a number of important ways. First, OAK is based upon the belief that the next evolutionary step in complex system automation involves goal-directed commanding at the system component level. This shifts the control paradigm from one of sending commands to subsystems, to sending goals and resources to intelligent subsystem management control agents that require minimal, if any, direct operator interaction. Secondly, subsystem management control agents are loosely coupled across a distributed, networked infrastructure. This results in a dynamic and adaptable coordination capability, leading to a more survivable overall system architecture. Finally, OAK uses declarative model-based reasoning as an extension to current rule and procedure based formalisms in the control agent control loop. Qualitative model-based reasoning provides the unique potential for performing real-time detection, identification, and diagnosis of unanticipated fault conditions.

[0005] OAK is based upon an open control agent communication infrastructure that includes a plurality of control agents coupled with respective subsystems. Control agents communicate via a messaging protocol that permits them to share their subsystem status and subsystem goals with other control agents. Status and goal data allows control agents to collaborate in diagnosis, prediction, or event response scenarios. The OAK infrastructure further supports the notion of a command element that can control and establish goals for the control agents. The command element may be a human operator, an outside control agent, an external software system, or the control agent community.

BRIEF DESCRIPTION OF THE DRAWINGS

- [0006] **FIGURE 1** illustrates one example of a logical OAK hierarchy.
- [0007] **FIGURE 2** illustrates a sample OAK communication network of control agents.
- [0008] **FIGURE 3** illustrates the architecture of a control agent within OAK.

DETAILED DESCRIPTION

[0009] The open autonomy kernel (OAK) provides a goal-directed coordinated intelligence approach with respect to complex connection-based system architectures. Connection-based systems can be thought of as physical or logical systems that are comprised of well defined components whose behavior is influenced by other well defined components along one dimensional paths. Examples of connection-based systems include, but are not limited to, electrical systems, plumbing systems, and computer networks. In addition, connection-based systems can be combinations of functionally distinct but interdependent system components such as an electrical system, a plumbing system, and a computer network that are components of, for instance, an overall engineering plant management system.

[0010] The term "components" includes both physical hardware components such as, but not limited to, a valve, a circuit, or a switch, as well as a cohesive collection of such hardware components, such as, but not limited to, a pumping station, an electrical substation, or a power grid. Components may also refer to a hardware/software combination that behaves as a distinct unit (e.g., a valve and its controller). Components may also comprise other more detailed components. A

key distinction is that the behavior of an individual component is easily defined from the perspective of peer components interacting with the individual component.

[0011] OAK provides intelligent support for three key features pertaining to overall system control: estimation, communication, and control. The primary building blocks of an OAK system architecture include control agents (CA), model-based reasoning engines (MBRE), control agent communications brokers (ACB), and control mediators (CM).

[0012] OAK utilizes two key technologies to perform the aforementioned functions. First, control is distributed through the use of control agent-based software. Secondly, OAK implements intelligent reasoning within each control agent using a technology known as Model-Based Reasoning (MBR). A primer of control agent based software and model-based reasoning as related to the present invention is included to facilitate a better understanding of the present invention.

Control agent Based Software

[0013] A control agent is an independently executing controller that is coupled to and responsible for an identifiable subsystem within a larger system, that is, at least in part, controlled by other peer control agents. A control agent is capable of perceiving changes in the subsystem, identifying the state of the subsystem, planning actions in accordance with the current subsystem state and/or desired subsystem goals, and executing planned actions.

[0014] OAK control agents are intrinsically permanent, stationary, exhibit both reactive and deliberative behavior, and are declaratively constructed. Control agents are reactive in their ability to reconfigure the subsystems within their control in the context of an existing task plan. Control agents are deliberative in their ability to create a task plan in response to observed states and defined goals. OAK control agents' extrinsic characteristics include proximity to the controlled

subsystem, social independence, and both awareness of and cooperativeness with goals and states of other control agents. OAK control agents are comprised of nearly homogeneous control agents, and are independently executed yet contain unique models of the subsystem for which the control agent is responsible. OAK control agents are environmentally aware and behavior of the environment is predictable through each control agent's model.

Model-based Reasoning

[0015] OAK control agents employ Model-based Reasoning (MBR) to deduce the condition of their underlying subsystems. Model-based reasoning, as applied to OAK, can be thought of as:

reasoning about a subsystem's behavior from an explicit model of the mechanisms underlying that behavior. Model-based techniques can very succinctly represent knowledge more completely and at a greater level of detail than techniques that encode experience, because they employ models that are compact axiomatic systems from which large amounts of information can be deduced.

[0016] In OAK, a particular flavor of MBR called "Qualitative Model-Based Reasoning" is used. This branch of MBR focuses upon using symbolic (rather than quantitative) representations of system behavior in order to simplify the modeling process and enable reactive-level response times. First, Model-based programming techniques are used to develop a set of qualitative, *first-principles* component models for elements of the controlled system. Multiple instances of these component models are combined to produce larger aggregate models of an entire system. Then, these aggregate models are used at runtime by an MBR kernel to produce system diagnoses. This kernel is able to resolve complex system-level interactions between components to produce a best estimate of system state. By placing the responsibility of resolving these complex interactions

upon the MBR kernel, the human modeler is freed from having to explicitly reason through and encode all possible diagnoses a priori.

First Principles Modeling

[0017] First principles models are interconnected finite state automata (FSA) in which interactions among connected FSA are explicitly enumerated by associating attributes used to define states. Model-based reasoning (MBR) is state estimation or planning of first principles models based upon resolving conflicts between state definitions and the values of attributes used to define states.

[0018] Discrete event system theory shows that systems may be modeled using an automaton $G = (X, E, f, I, x_o, X_m)$, in which X is the set of discrete states, or state space; E is the finite set of events; f is the transition function; I is the active event function mapping X to E ; x_o is the initial state and X_m is the set of marked states. Control of G is provided by a control policy that includes a set of control actions.

[0019] Automata that are memoryless are considered *Markov Processes*. Association of control actions, costs and probabilities with state transitions allows us to derive Markov Decision Processes (MDP). MDP are defined by the tuple (X_S, C, E_C, f_T, R) in which: X_S is the finite set of states of the system being tracked; C is the set of possible commands of which c is an element ($c \in C$); E_C is a finite set of commanded events described in the form (x_I, c, x_F) where x_I is the initial state and x_F is the final state; f_T is a state transition model of the environment which is a function mapping $X_S \times E_C$ into discrete probability distributions over X_S , and R is the cost function over E_C . The actions are non-deterministic, so we write $f_T(x, e)$, in which $(x \in X_S)$ and $(e \in E_C)$, for the probability that transition e will occur given the current state x .

[0020] OAK is required to reason with the possibility of component failures, so the commanded event set E_C is replaced by the set of possible events E_P that

includes the commanded set and the failure events E_F . ($E_P = E_C \cup E_F$) Note that failure events may be commanded or spontaneous, and are assumed permanent and non-recoverable. The state transition function f_T is extended to $X_S \times E_P$, providing us with a failure sensitive MDP (X_S, C, E_P, f_T, R).

[0021] By way of example, a failure sensitive model of a simple circuit breaker may be described as follows:

M(breaker):

X_S : {open, shut, stuck-open, stuck-shut}

C : {open, shut, none}

E_P : {(open, shut, shut), (shut, open, open), (open, shut, stuck-open), (open, none, stuck-shut), (shut, open, stuck-shut), (shut, none, stuck-open)}

f_T : $\rho(\text{open, shut, shut}), \rho(\text{shut, open, open}) \cong 1$,
 $\rho(\text{open, shut, stuck-open}), \rho(\text{shut, open, stuck-shut}), \cong 0.01$;
 $\rho(*, *, \text{stuck-open}), \rho(*, *, \text{stuck-shut}), \cong 0.0001$;

R : { $r_{\text{open}}, r_{\text{shut}} = 1, r_{\text{none}} = 0$ }

[0022] The behavior a system exhibits in state x , where $x \in X_S$, is defined by attributes. Attributes are physical properties of the system. Propositional statements containing these attributes are used to express system behavior in each state. The behavior of our circuit breaker may be defined through the following propositions:

$(\text{breaker.state} = (\text{open} \vee \text{stuck-open})) \rightarrow \text{conduct}$

$(\text{breaker.state} = (\text{shut} \vee \text{stuck-shut})) \rightarrow \neg \text{conduct}$

[0023] Because internal sensors are subject to faults, and the measure of a system is its effect on the systems it services, OAK is interested in the behavior of our system as it relates to the outside world. External influences are modeled as a special class of attributes known as interfaces. The circuit breaker has two such attributes: $current_{in}$ and $current_{out}$. Using these interfaces behavior definition propositions are rewritten as:

$$(breaker.state = (open \vee stuck-open)) \rightarrow (current_{in} = current_{out})$$

$$(breaker.state = (shut \vee stuck-shut)) \rightarrow (\neg current_{in} \wedge \neg current_{out})$$

[0024] A system containing multiple related components may be modeled by associating component model interfaces in propositional statements declared in the system model. A system model is represented by the tuple: $M_{system} = (M_{comp}, A, P)$ in which M_{comp} is a set of failure sensitive component models, A is a set of attributes defined by the system, and P is a set of propositions associating system defined attributes A and M_{comp} interfaces. A small system comprised of two circuit breakers connected in series may be represented as follows:

$M(system):$

$M(breaker_1), M(breaker_2): M(breaker);$ //defines $breaker_1$ and $breaker_2$ to be "breakers"

$breaker_1.current_{out} = breaker_2.current_{in}$

[0025] Partially Observable Markov Decision Processes (POMDP) are MDP that have been extended to include a finite set of observations. Failure sensitive POMDP are represented by the tuple $M = (X_S, C, E_P, O, f_T, R)$. in which O is the observation function that maps X_S to the finite set of observations. However, states have been defined as abstractions represented in terms of attributes. As

abstractions, states are not directly observable. The ability to observe a state is associated with the ability to observe the attributes that define the state. A state is observable if all of the attributes in any clause of a disjunctively formed state behavior definition are observable. A state may be considered partially observable if it is not observable and at least one attribute is observable. The observation function O in traditional POMDP is replaced by an observation function O_A that maps the attribute set A to the finite set of observations. The probability of making an observation o from the attribute a is denoted as $O(o,a)$.

Model Based Estimation

[0026] Estimation methods have been demonstrated for systems of components modeled as failure sensitive POMDP. The first step in estimation is to determine a belief state for the system. Assuming an initial state X_I of the components, and a set of commands $c \subseteq C$ applied to the system components model-based estimation makes the naïve assumption that the most probable component transitions in E_C have occurred. The belief state X_B for all components is easily derived. A single propositional statement that describes the believed values of attributes within the system may be generated by conjoining the propositional clauses associated with these belief states, general propositional clauses showing component relationships and known attribute values for interfaces external to the system. For example, consider the simple two circuit breakers in series model:

```

M(system): {
    M(breaker1), M(breaker2): M(breaker);
    breaker1.currentout = breaker2.currentin;
    breaker1.currentin = external stimulus; }
breaker1.initial_state = open;
breaker2.initial_state = closed;

```

Given the command $\text{breaker}_2.\text{command} = \text{open}$ and an externally provided flow ($\text{external stimulus} = \text{true}$) then the state set $\{\text{breaker}_1.\text{state} = \text{open}; \text{breaker}_2.\text{state} = \text{open}\}$ is realized and the following propositional statement presents itself:

$$(\text{external stimulus} = \text{true}) \wedge (\text{breaker}_1.\text{flow}_{\text{in}} = \text{external stimulus}) \wedge (\text{breaker}_1.\text{current}_{\text{out}} = \text{breaker}_1.\text{current}_{\text{in}}) \wedge (\text{breaker}_2.\text{current}_{\text{out}} = \text{breaker}_2.\text{current}_{\text{in}}) \wedge (\text{breaker}_1.\text{current}_{\text{out}} = \text{breaker}_2.\text{current}_{\text{in}}).$$

[0027] One can see how a value for an attribute, in this case our *external stimulus* propagates through the system inferring values on the $\text{current}_{\text{in}}$ and $\text{current}_{\text{out}}$ attributes in breakers one and two.

[0028] Observed attribute values generate additional observed propositional clauses (e.g., $\text{breaker}_2.\text{current}_{\text{out}} = \text{true}$). If the belief state is incorrect then a conflict between the observed clause and one or more of the belief state clauses will be generated. If no conflict is generated then the belief state is confirmed. A solution to a conflict, or set of conflicts, is defined as a set of alternate transitions E that generate propositional statements that do not conflict with observation clauses. For example, in the breaker example if a single transition e : (shut, open, open) is applied to the above initial state and $\text{breaker}_2.\text{current}_{\text{out}} = \text{false}$ is subsequently observed a conflict will be generated. If the transition set is replaced by either of the sets

$$E = \{\text{breaker}_2.\text{transition}(\text{shut}, \text{open}, \text{stuck-shut})\} \text{ or}$$

$$E = \{\text{breaker}_2.\text{transition}(\text{shut}, \text{open}, \text{open}), \text{breaker}_1.\text{transition}(\text{open}, \text{none}, \text{stuck-shut})\}$$

these transitions will generate a belief state that does not conflict with the observation clause. Using the transitional probabilities found in the transition function f_T we can apply Bayes' rule to the probabilities of our two candidate solutions to determine their relative probabilities. The generation of possible solution sets can be accomplished using *Conflict-directed Best First Search* (CBFS).

[0029] When multiple candidate solutions satisfy a set of conflicts, the most likely solution is selected as a basis for reconfiguration and subsequent estimation. To protect the estimation engine from permanently adopting an incorrect solution, a Truth Maintenance System is used to track likely solutions. If future observations generate clauses that provide support to alternative solutions, the TMS generates a historical revision of the belief states.

Encapsulation and Abstraction

[0030] Large systems require models that, due to their complexity, are difficult to process using CBFS and TMS. This complexity is mitigated through encapsulation, decomposition and distribution. A system may be decomposed by encapsulating portions of the system model into subsystems of modest size (<25 components). Subsystems inherit component and system semantics. The components of the system, the subsystem attributes, and the propositional statements used by the subsystem to associate components are represented by (M_{comp}, A, P) . Subsystem automata $(X_S, C, E_P, O_A, f_T, R)$ are associated with subordinate component automata M_{comp} by defining the subsystem states X_S using propositional statements over subordinate component attributes. The entire subsystem is represented by the rather large tuple $M_{sub} = (X_S, C, E_P, O_A, f_T, R, M_{comp}, A, P, A_I)$ in which A_I is the interface set for the subsystem. This structure provides for subsystems to exhibit automaton-like behavior at an abstract level; however, it is not necessary that they do so. Subsystems that contain null sets for

any or all of the abstract automata elements (X_S , C , E_P , O_A , f_T , R) are acceptable and often desirable.

[0031] If a subsystem is believed to be in a state, then the propositional statement associated with its state is *applicable*. Likewise if a subsystem is believed to be in a state, then the propositional statements that define the subordinate component states are also applicable. The set of applicable statements may be *conjoined and reduced to create a single belief statement for the subsystem*. Using the combined belief statement, CBFS and TMS may be applied to perform estimation and reconfiguration within the scope of the subsystem.

[0032] Because subsystems inherit the properties of component models, a subsystem model may be used as a component of another subsystem. This parent-child relationship permits building topologies of subsystems featuring numerous models at varying levels of abstraction. Parent-child relationships are expressed by associating the behavior, in terms of the child's interfaces and goals, within propositions internal to the parent.

[0033] Control with an independent CBFS and TMS for each subsystem is complete if and only if the subsystem behavior represented by the subsystem's proposition statement set is independent. Unfortunately this is not usually the case. *A subsystem's interface attributes are, by definition, constrained to the world outside of the subsystem*. This problem is limited by enforcing design constraints that remove some of the more difficult cross-subsystem dependencies. Design constraints are focused on eliminating the possibility of conflicts being generated by observations and propositional clauses that are difficult to resolve across subsystems, and also eliminating the possibility of goals being directed to a subsystem that cannot be directly addressed within the subsystem. While these constraints are limiting, experience has shown that robust, useful models are still a

possibility. The design constraint is described in terms of conflicts generated during the estimation process. Because our reconfiguration strategy is a mirror of our estimation strategy, equivalent design criteria may be used to simplify reconfiguration. Conflicts may be divided into two distinct types: *independent* conflicts are identifiable by observation-proposition combinations that are contained within a single subsystem; *dependent* conflicts are generated by behavior that is not identifiable by observation-proposition combinations within a single subsystem.

[0034] Collaboration between independent subsystems is straightforward. Observations are made by independent subsystems that resolve the conflict. After the conflict is resolved the effects of the observation are propagated to the system and other subsystems via interfaces.

[0035] As previously stated, OAK is a distributed, multicontrol agent system. An OAK system can have a varying topology based on the overall system application. OAK has two major use-cases that almost fully describe the operation of the system: OAK's reaction to user-input goals; and OAK's reaction to system state change. The primary intelligent components that enable OAK to accomplish these use-cases are a model-based reasoning engine and a planner.

[0036] To perform the diagnostic phase of the control cycle, OAK control agents continually update their states using the model-based reasoning engine, and pass these state updates to other control agents that are interested so that these control agents may update their states. In response to these states or to the system's environment, an external actor or an OAK control agent will provide goals to the OAK system, which are distributed for further processing.

[0037] **FIGURE 1** illustrates one example of a logical OAK hierarchy. The circles **10** represent control agents, and the diamonds **12** represent hardware. The connecting lines represent parent-child control agent relationships **14** and

connections to hardware **16**. User-defined goals flow down the hierarchy while component states generally flow upwards in the hierarchy.

[0038] While a hierarchical topology is the topology that is presented herein, the OAK architecture does not preclude other topologies. Other topological possibilities include: peer-to-peer, where each control agent communicates facts and goals to any other; multi-hierarchy, where there exist multiple loosely coupled hierarchies; and a star topology, or one-level hierarchy. For the rest of this discussion a hierarchical topology is assumed.

[0039] Each control agent has an associated Control agent Communication Broker (ACB), which is responsible for handling all of the control agent's communication with a control agent Communication Framework (ACF). The ACB maintains a queue of messages coming into the control agent. Additionally, any control agent that has direct communication with hardware has a control mediator (CM) to handle the hardware level goals that are generated by these control agents for the associated hardware, and to receive updates about this hardware. These messages are not handled by the ACF. By using this layered approach, control agents are decoupled from the ACF.

[0040] The Control Agent Communication Language (ACL) of OAK provides several message templates, including messages for queries, state updates, subscription requests, goals, exceptions, and control agent coordination. To protect information, each control agent's ACB may optionally have an information access matrix that determines which control agents are allowed to subscribe to that control agent's events. The ACF allows any control agent to communicate directly with any other control agent. Thus, communication between control agents is not restricted to any logical hierarchical framework. The ACF also supports message logging.

[0041] One of the major use-cases of OAK is to react to goals entered by an external actor. These are system-level commands which have the potential of transitioning the entire multicontrol agent system from one state to another.

[0042] OAK communicates with external actors via a graphical interface, which is also decoupled from the rest of the system. Goals that are entered from an external actor, such as a human operator, through this interface are sent directly to a root level control agent using a *goal message*. This control agent develops a plan with goals that apply to the domains of its child control agents. Goals have a priority associated with them, which is used for goal preemption.

[0043] After the root level control agent develops a plan and directs a goal to one of its child control agents, the goal is received by the child control agent's ACB, sorted into its queue, and eventually accepted by the control agent for processing. This control agent develops a plan to implement the goal. Since this control agent is a root of its own tree, the goals developed by its planner are passed to its children control agents. This propagation continues until leaf control agents at the hardware level receive goals for their specific domains.

[0044] Once goals are received at the leaf level, a similar process occurs in that a plan is developed and goals are passed out of the control agent. The only difference is that the goals are now passed to the control agent's control mediator (CM), which translates the goal into commands that a hardware driver can understand. The CM has a queue of such commands in case the control agent is able to generate commands more quickly than the CM is able to deliver them. Since the CM is the only component that has direct interaction with the hardware drivers, it is the only component that has to be updated when hardware itself is changed or when hardware drivers are updated.

[0045] Successful goal implementation implies a state change, so a control agent does not have to set up callbacks with the hardware to confirm that a command was successful. Leaf control agents are already required to monitor the hardware they control for changes in order to accomplish the second major use-case of OAK. Therefore, the leaf control agent waits for a reaction from the hardware monitors to indicate that the command has been successful. The control agent is then free to pass out goals that were temporally dependent on the goal just implemented. Since state changes are propagated up the hierarchy (as well as to unrelated control agents, potentially), control agents at higher levels are also informed that their goals were implemented and they can then pass out goals that had to be put in a wait state. To an implementer of OAK, this means that the incoming goal use-case and the state change use-case, which comprise the two major functions of OAK, are decoupled.

[0046] To appropriately handle changes in the state of the system, OAK uses a model-based reasoning engine (MBRE). Recall that in OAK, any control agent may subscribe to another control agent's events (which are triggered by state changes) assuming that it has permission. The control agent most likely to be interested, however, is the control agent's parent, because the state of its children are reflected as variables in its model. For this reason, parents will subscribe to most of their children's state change events. The external actor's control agent may also subscribe to any control agent's state change events, so that the user can be advised of a state change at any level in the hierarchy.

[0047] State change events are transmitted through the use of a *fact message*. This message contains a representation of the knowledge contained in a control agent. Once a fact message is received, the control agent, using the MBRE, determines whether or not the change is important enough to warrant a state change. If the state changes, all subscribed control agents are informed, and

propagation of state changes begins as discussed above. Note that since many control agents may subscribe to an event, state changes may be propagating in several subtrees at any given time.

[0048] One of OAK's strengths is realized in a control agent's ability to determine the state of its model, compare that state to a knowledge base, and reactively plan. Thus, a control agent can autonomously control its domain until a control agent that is higher in the hierarchy (or in the case of the root control agent, the external actor's control agent) preempts its control. The component of OAK that controls reactive planning is called the *reactive manager*.

[0049] There are two types of information in the reactive manager: persistent goals and emergency conditions. Persistent goals are simply goals that are desired true for the duration of the control agent. An emergency condition is defined as a state transition that cannot be reversed and requires OAK to act immediately to protect the resident system. When OAK detects an emergency, it will preempt the external actor's goal and go to a predetermined goal that will minimize damage to the system being modeled. This goal will be implemented until the actor enters a goal with 'preempt' priority. From this point on, goals from the user are implemented as completely as possible based on the damage to the system.

[0050] As mentioned above, when a control agent receives a goal it must use a planning mechanism to create a viable plan to achieve that goal. The plan is in a format of an ordered sequence of fragments. Each fragment is comprised of one or more subgoals. The idea is that within a fragment each subgoal may be accomplished in parallel, while subgoals in a fragment prior to a given fragment must be completed before the current fragment may be attempted.

[0051] Several different planners may be appropriate for different control agents depending on how complex they are, what domain they are planning for, etc.

Therefore, the planner is instantiated at run-time differently for each control agent from a group of developed planners. To date, two planners have been developed: the *Scripted Planner* and the highly specialized *Graph-Based Planner*.

[0052] The scripted planner is extremely simple but useful for simple control agents, such as so-called leaf control agents. The scripted planner originally matched the incoming goal with a pre-defined list of incoming goals, and output a pre-defined plan. A later extension to the scripted planner allowed matching on the incoming goal and a propositional logic expression about the current world-state. For example, Goal A along with $(variable1 = value1) \wedge (variable2 \neq value2)$ would generate plan A. Many different propositional expressions, and therefore plans, could be associated with each incoming goal. Also, since the scripts are checked in a specific predefined order, a simple priority of outputted plans can be imposed. Therefore, with these simple, essentially rule-based scripts, complex behavior could be created.

[0053] The graph-based planner has been designed specifically for the test domain described below. Test domain planning comprised determining how to move flow from a source to several sinks through a dynamic pipe network, with many operational constraints. This planner represented the target domain as a loadable model. The model was represented internally as a digraph, with weights on each edge set according to the operational constraints. The planner operated by performing Prim's Minimum Spanning Tree algorithm on the graph to determine how to get flow to as many of the desired sinks as possible. Along the way, the software determined the actions the control agent would need to take to align the system in the manner that Prim's algorithm output. Finally, the planner would generate a plan based on the actions determined.

[0054] OAK is a system of collaborative model-based control agents used for control. OAK's greatest strength is realized in a control agent's ability to determine

the state of its model, compare that state to its current goal, and *reactively* plan. Thus, a control agent autonomously controls its domain.

[0055] Referring again to **FIGURE 1**, each control agent performs reasoning and planning at various levels of abstraction on its own virtual machine. This allows for logical distribution of time-intensive operations like planning. A higher-level control agent, **A₁** for example, will develop a relatively abstract plan for accomplishing some goal without having to worry about its implementation. Instead, each goal in that plan is passed to lower level control agents that further decompose the details of that goal. This propagation continues with plans becoming increasingly detailed until they reach control agents (**A₁₁₁**, **A₁₁₂**, **A₁₁₃**) that can translate the goals directly into hardware commands. It is important to note that while this architecture implies a hierarchical planning structure, the planning at individual levels is not restricted to the Hierarchical Task Planning approach. This design also allows for a high degree of modularity, since control agents that have the same interface can replace each other.

[0056] **FIGURE 2** illustrates a sample OAK communication network of control agents. Communication is not limited to the logical channels shown in **FIGURE 1**. Rather, any control agent **22** may communicate to any other control agent **22** via its Agent Communication Broker (ACB) **24**, through the Agent Communication Framework (ACF) **26**. Thus, any control agent may subscribe to any other control agent's notifications. Control agents are associated with subsystems or high level models. Each type of control agent may operate with data and goals provided by other control agents. Within OAK, each control agent associated with a subsystem provides closed-loop subsystem management using sensed parameters to infer system state. The control agents can adjust control commands to achieve a desired operating profile and/or respond to external events or component failure. Each control agent can communicate state information and goals to other control agents using a "publish and subscribe" mechanism to limit bandwidth usage.

[0057] Each control agent in an OAK architecture continually performs a cycle of mode estimation, planning, and execution that is influenced by observed subsystem state and received goals. This cycle is implemented using various *manager* components in each control agent.

[0058] The architecture of a control agent **22** is illustrated in **FIGURE 3**. Control agent architecture includes managers for each major step in the functioning of the control agent.

[0059] The executive manager **32** acts as a gateway (analogous to OSI level 6) for the control agent **22**. It is responsible for decoding and executing incoming messages, as well as encoding outgoing messages. As such, it is the only manager that communicates directly with the ACB **24**. The ACB **24** communicates through the ACF **26** to other control agents **22**.

[0060] OAK control agents attempt to handle large amounts of information in an efficient manner. While the planner is planning for a goal, the control agent can still accept facts and resolve queries. As this is happening, there may be a change in state that may preempt a plan already in progress. Multi-tasking within control agents is handled by making each OAK message object a separate thread.

[0061] One type of message object is the query object. Query objects are created when the control agent receives a state query. It interfaces with an MBRE manager, which stores the current state of the model. The query object causes creation of a fact object addressed to the query originator if the queried state is recognizable or an exception if it is not.

[0062] Fact objects are another type of message object. In addition to the scenario above, a fact object is created when state updates from other control agents or from associated hardware are received. State updates from other control

agents are received via the ACB, and state updates from associated hardware are received via the Control Mediator (CM).

[0063] The third type of message object is the goal object. Goal objects are sent from higher level control agents to communicate their desire for this control agent's subsystem. Goal objects interface with the MBRE manager to get the control agent's current state, the Planner to receive the proposed plan, and the plan implementation manager to execute the plan.

[0064] The fourth type of message object is the subscription object. These objects contain a reference to a control agent, and a particular state variable that that control agent is interested in. After receiving a subscription object, the receiving control agent will send fact messages to the subscribing control agent whenever the information that it was interested in changes. To this end, the subscription object interfaces with the control agent's subscription manager.

[0065] A fifth type of message object is the administrative command object. These objects contain commands that deal with the control agent software's behavior outside the scope of the intelligent control system: things like stopping a control agent, printing debugging messages, or resetting the control agent.

[0066] In addition to message objects, there are also seven managers within an OAK control agent. The executive manager **32** routes all of the incoming and outgoing messages for a control agent. The MBRE manager **34** exists solely to store the current state of the model and to interface the query, fact, and goal objects with the model-based reasoning engine (MBRE) **36**. The MBRE manager **34** is the component that handles the mode estimation portion of the OAK control cycle. The ACL manager **38** translates each incoming message from the ACL into an object that the control agent can use.

[0067] Each incoming goal that results in a successful plan must also communicate with the plan implementation manager (PIM) **40**, which ensures that no step of a plan is implemented before each of the previous steps is successful. The PIM **40** corresponds to the execution phase of the OAK control cycle. This manager will have access to a matrix that stores the average transition time for each goal, a multiple of which is added to a constant time to determine the amount of time that the control agent will wait for an action to be performed. This makes the control agents adaptive and can help detect problems with the system early on.

[0068] The subscription manager **42** handles subscription requests from other control agents. It stores lists of control agents that are interested in various state changes, and automatically generates fact messages to each of the subscribed control agents when a given state variable changes.

[0069] The planner **44** is responsible for deducing plans from the current input goal and the current estimated state. The planner **44** corresponds to the planning phase of the OAK control cycle.

[0070] Finally, the reactive manager **46** continually compares the control agent's states to a vector of emergency states and a persistent goal to determine whether or not the control agent's model has entered a state that the developer has identified as abnormal or hazardous. In the event that the reactive manager **46** identifies that the control agent has entered such a state, it pushes out a new goal for the control agent to implement in order to prevent serious damage to the system.

[0071] The reactive manager **46** works by having a list of configurations of the system in which the reaction will occur. Each configuration is a propositional expression of facts, and can use the AND, OR, and NOT operators as well as arbitrarily nested parenthesized clauses. Associated with each proposition is a goal. When the control agent's model is in a state that is consistent with one of the

configurations, the reactive manager **46** will fire and give the control agent the goal associated with that configuration.

[0072] To determine whether or not the system is consistent with one of the listed configurations, the reactive manager **46** will compare the system state to each of the propositions in the list each time the system's state changes. The configurations in the list will have a strict ordering, so some propositions are guaranteed to be checked before others later in the list. If a proposition is consistent with the system state, the associated goal is sent to the executive. This goal will be treated the same as any incoming goal, including preempting the current goal if the precedence of the current goal is sufficiently low. Upon receiving a reactive goal, the executive may send an exception-type message into the ACF.

[0073] In addition to emergency states, the reactive manager **46** also stores a persistent goal. This goal will be given by the developer at startup, but may be changed by the system user via special configuration messages during run time. Consistency with the persistent goal is checked before the reactive manager **46** checks for emergency conditions, and if the state is found to be inconsistent, the control agent reacts as described above. The user should not, however, consider leaving the persistent goal's state to be catastrophic, as it is assumed to be in the emergency case.

[0074] Mode Estimation (ME) is the process of deducing system states based upon partial or incomplete information. Within the context of OAK, Mode Estimation involves inferring the states of various system components using: first-principles logic-based models, observations gathered from the hardware, and knowledge of past commands issued to the hardware. Internal to a control agent, this capability is provided by the REManager (Reasoning Engine Manager), which uses a single underlying reasoning engine to produce diagnostic estimates or "candidates". These candidates are then converted into fact objects, which are the atomic units of

knowledge representation in OAK. Fact objects are passed between control agents, and consequently, between reasoning engines.

[0075] Each time a command set is issued to the hardware, or new observations arrive from the hardware in the absence of an explicit command, a state transition occurs and the reasoning engine attempts to determine the most likely candidates (resultant system states). The process of generating candidates comprises a number of steps. First, the reasoning engine generates an expected next state for each component based upon any command information that may be available. Next, the reasoning engine considers all recent observations from the system. A consistency-checking algorithm is then run against the plant model given the new observations. The output from this algorithm is a set of ranked candidates, or state estimates. It is then the responsibility of the external actor who is using the reasoning engine to decide which candidate it wishes to select or “believe”.

[0076] The reasoning engine supports a number of parameters that are set when the engine is first initialized. These parameters are used to configure the estimation process, and different settings can potentially result in different diagnoses being reached. **Table 1** illustrates an example of reasoning engine settings to be used with OAK control agents.

TABLE 1

Parameter	Setting
Search engine type	CBFS (Conflict-directed Best First Search)
Number of candidates	3
Maximum number of candidates to search over	10000
History length	10
Maximum number of trajectories	10
Progress style	Full
Trajectory tracker type	Extend

[0077] OAK receives high level goals from an external actor and creates an implementation plan for these goals. The implementation is carried out by a plurality of control agents in a distributed manner utilizing model-based reasoning techniques. Moreover, OAK reacts autonomously to changes in the system. The distributed nature of OAK permits the aforementioned functions to be realized in a time efficient manner. In addition, OAK can respond to catastrophic changes in the system environment including damage to the OAK system itself. All subsystems of an OAK system, including the intelligent components, are decoupled from one another wherever possible allowing for many different such subsystems to be utilized. Ultimately, OAK controlled systems reduce manpower requirements which is especially significant in repetitive or dangerous domains.

[0078] In the following claims, any means-plus-function clauses are intended to cover the structures described herein as performing the recited function and not only structural equivalents but also equivalent structures. Therefore, it is to be understood that the foregoing is illustrative of the present invention and is not to be

construed as limited to the specific embodiments disclosed, and that modifications to the disclosed embodiments, as well as other embodiments, are intended to be included within the scope of the appended claims. The invention is defined by the following claims, with equivalents of the claims to be included therein.

CLAIMS

1. A system for controlling a connection-based system in which a plurality of independently executing control agents are distributed across an agent communication framework, each control agent being associated with a subsystem of the connection-based system wherein each control agent is responsive to other control agents such that connection-based system goals and connection-based system states can be propagated throughout the distributed network of control agents and acted upon to autonomously control the connection-based system, said system comprising:

an agent communication framework that serves as a communication network for the plurality of control agents, wherein

each control agent is comprised of:

an agent communication broker that provides a communication interface between a control agent and the agent communication framework;

a model based reasoning engine communicable with the agent communication broker that uses a first principles modeling technique to estimate subsystem states based on data gathered from subsystem hardware components and knowledge of past commands issued to said hardware components; and

a planner communicable with the agent communication broker and the model based reasoning engine that generates a plan for the implementation of a specific goal based input including the specific goal and the current subsystem state as determined by the model based reasoning engine.

2. The system of claim 1 wherein at least one control agent is communicable with an external actor that provides a system level goal to initiate connection-based system-wide control.

3. The system of claim 2 wherein a control agent is communicable with at least one hardware component within the control agent's subsystem via a control mediator that provides an interface between the control agent and the hardware's device controller.

4. The system of claim 3 wherein said control agent is further comprised of:

an executive manager communicable with the agent communication broker that routes incoming and outgoing messages for the control agent;

an ACL manager communicable with the executive manager for translating incoming messages into message objects, wherein message objects include a query object having data pertaining to subsystem status queries, a fact object having data pertaining to subsystem state updates, a goal object having data pertaining to the goal of a control agent, and a subscription object having data pertaining to another control agent that is interested in the status of the present control agent;

a model-based reasoning engine manager communicable with the executive manager, said model-based reasoning engine manager for storing the current state of a subsystem model; and communicating query, fact, and goal objects between the model-based reasoning engine and executive manager;

a plan implementation manager communicable with the executive manager for implementing a plan generated by the planner;

a subscription manager communicable with the executive manager for maintaining a list of other control agents authorized to receive messages from the present control agent handling requests from other control agents and generating fact message objects pertaining to subsystem state changes to said other control agents.

5. The system of claim 4 wherein said control agent is further comprised of a reactive manager communicable with the executive manager for monitoring and comparing the current control agent's state against a set of emergency states and a persistent goal to determine whether the current control agent's model has entered an abnormal or hazardous state, wherein when the reactive manager determines that a control agent's model has entered an abnormal or hazardous state, said reactive manager generates a

new pre-emptive goal for the control agent to implement in order to prevent damage to the connection-based system.

6. A system for autonomously controlling a connection-based system in which a plurality of independently executing control agents are distributed across an agent communication framework, each control agent being associated with a subsystem of the connection-based system wherein each control agent is responsive to other control agents, said system comprising:

means for developing a strategy for the propagation of goals across the distributed network of control agents;

means for propagating said goals across the distributed network of control agents;

means for developing a strategy for the propagation of inferred beliefs pertaining to subsystem states across the distributed network of control agents; and

means for propagating said inferred beliefs pertaining to subsystem states across the distributed network of control agents; and

model-based reasoning means using a first principles modeling technique to perform subsystem state estimation to determine inferred beliefs pertaining to subsystem states.

7. A method of autonomously controlling a connection-based system in which a plurality of independently executing control agents are distributed across an agent communication framework, each control agent being associated with a subsystem of the connection-based system wherein each control agent is responsive to other control agents, said method comprising:

(a) assigning, via an external actor, a goal to a control agent;

(b) in the control agent assigned a goal by the external actor, developing a goal oriented plan comprised of subgoals;

(c) propagating the current subgoals to child control agents based on the subsystems within the child control agents' domains wherein the child control agents

further decompose the original goal oriented plan by developing a new goal oriented plan of subgoals;

(d) repeating step (c) until the original goal oriented plan has been fully decomposed to the point where the subgoals can be translated into instructions for hardware that is within a control agent's domain;

(e) causing a hardware component to execute an instruction set in accordance with a control agent's plan;

(f) determining the state of a control agent's subsystem using a model-based reasoning means; and

(g) propagating the state of a subsystem to any interested control agents that require certain subsystem state information to develop their goal oriented plan.

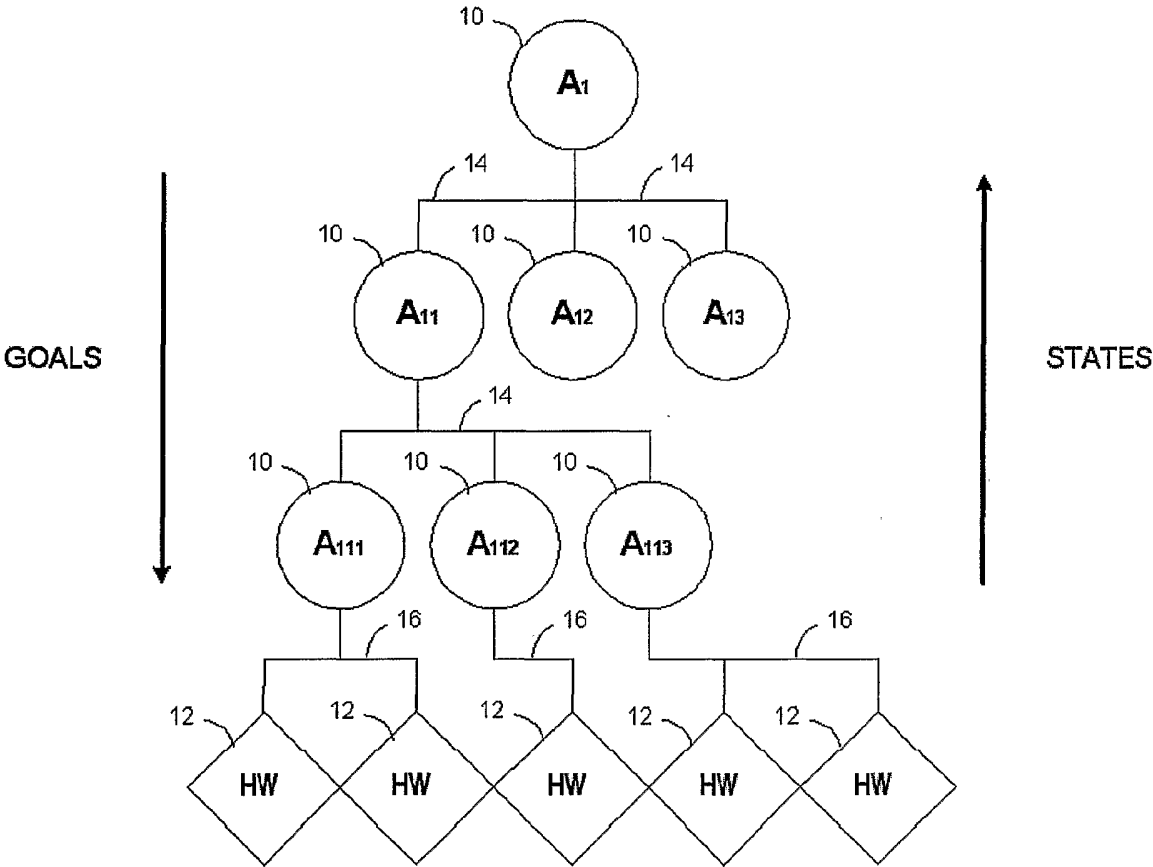


FIGURE 1





FIGURE 2

FIGURE 3



INTERNATIONAL SEARCH REPORT

International application No.

PCT/US02/16838

A. CLASSIFICATION OF SUBJECT MATTER												
IPC(7) : G06F 9/44												
US CL : 709/317												
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. : 709/317, 223, 225												
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)												
C. DOCUMENTS CONSIDERED TO BE RELEVANT												
Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.										
A	US 6,049,819 A (BUCKLE et al) 11 April 2000 (11.04.2000), the whole document.	1-5										
A	US 5,958,010 A (AGARWAL et al) 28 September 1999 (28.09.1999), the whole document.	1-5										
A	US 5,822,585 A (NOBLE et al) 13 October 1998 (13.10.1998), the whole document.	1-5										
A	US 5,655,081 A (BONNELL et al) 05 August 1997 (05.08.1997), the whole document.	6-7										
A	US 5,963,447 A (KOHN et al) 05 October 1999 (05.10.1999), the whole document.	6-7										
<input type="checkbox"/> Further documents are listed in the continuation of Box C. <input type="checkbox"/> See patent family annex.												
* Special categories of cited documents: <table border="0"> <tr> <td>"A" document defining the general state of the art which is not considered to be of particular relevance</td> <td>"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</td> </tr> <tr> <td>"E" earlier application or patent published on or after the international filing date</td> <td>"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</td> </tr> <tr> <td>"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)</td> <td>"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</td> </tr> <tr> <td>"O" document referring to an oral disclosure, use, exhibition or other means</td> <td>"&" document member of the same patent family</td> </tr> <tr> <td>"P" document published prior to the international filing date but later than the priority date claimed</td> <td></td> </tr> </table>			"A" document defining the general state of the art which is not considered to be of particular relevance	"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	"E" earlier application or patent published on or after the international filing date	"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	"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)	"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	"O" document referring to an oral disclosure, use, exhibition or other means	"&" document member of the same patent family	"P" document published prior to the international filing date but later than the priority date claimed	
"A" document defining the general state of the art which is not considered to be of particular relevance	"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											
"E" earlier application or patent published on or after the international filing date	"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											
"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)	"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											
"O" document referring to an oral disclosure, use, exhibition or other means	"&" document member of the same patent family											
"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		Date of mailing of the international search report										
23 August 2002 (23.08.2002)		16 SEP 2002										
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 LARRY DONAGHIE <i>James R. Matthews</i> Telephone No. 703 305-3665										