

HAZARD DETECTION FOR MULTI - AGENT REACTIVE ROBOTIC SYSTEM USING ON - LINE SOFTWARE ANALYSIS OF SAFETY CONSTRAINTS

A. Chennakesava Reddy

Associate Professor

Department of Mechanical Engineering

JNTU College of Engineering

Kukatpally, Hyderabad - 500 072

B. Rami Reddy

Registrar

Jawaharlal Nehru Technological University

Kukatpally, Hyderabad - 500072

Abstract: Hazard situations in safety-critical systems are typically complex, so there is a need for means to detect complex hazards and react in a timely and meaningful way. This paper addresses the problem of hazard detection through the development of an on-line analysis tool. The approach allows the user to specify complex multi-source hazards using a query-like language, uses both synchronous and asynchronous on-line checking approaches to balance efficiency and expressiveness, accommodates dynamic applications through dynamic constraint addition, and supports distributed and parallel applications running in heterogeneous environments.

Keywords: hazard detection, online analysis, multi source systems, dynamic applications

1. INTRODUCTION

Public awareness of safety issues involving computers is growing as incidents resulting in loss of life or near loss receive wide publicity. A *safe system* is one that is free from accidents or unacceptable losses. There is important research and development being done in providing intrinsically safe systems [1, 2]. But the goal of an intrinsically safe system is difficult to achieve for some of today's complex, dynamic applications running in parallel or distributed environments. Adopting guidelines established by system safety engineers [3], if one cannot guarantee an intrinsically safe system, the next preferred approach is a technique that prevents, minimizes, or detects the presence of hazards. A *hazard* is a state or condition of the system that combined with some environmental conditions can lead to an accident or loss event. Automatic pressure relief valves, lockins, lockouts, and interlocks are common hardware hazard prevention approaches. An example of a software approach is a trip computer in a nuclear power plant that initiates procedures to shutdown the plant when operating conditions are hazardous [4].

This paper addresses the problem of enhancing software safety through hazard detection through on-line hazard analysis. It is understood that the hazard situations can and do occur, and are often complex, involving multiple sources. So there is a need for a mechanism to detect complex, multi-source hazards and react in a timely and meaningful way. This paper focuses such a detection mechanism through C-net, an on-line analysis tool that supports the specification of complex hazards using a query-like language, uses both synchronous and asynchronous checking approaches to balance efficiency and expressiveness, accommodates dynamic applications through dynamic constraint addition, and supports distributed and parallel applications running in heterogeneous environments. The detection approach is applied to a set of autonomous robots that can navigate toward a goal across unmapped terrain, *reacting* to stimuli in the environment.

2. SOLUTION STRATEGY

The approach consists of a language for specifying complex hazards, a library for building executable versions of the hazard descriptions, and a run-time environment to handle execution. In particular, hazards are represented as constraints specified on application behavior. These constraints, called *safety constraints* are specified with a rule language. The compiler generates a list of commands to the library that in turn builds an executable entity, or node,

for each constraint as a collection of selection, projection, and join operations. The nodes are linked together at runtime as a Directed a Cyclic Graph (DAG). As shown in Fig.1, the compiler accepts sensor and constraint specifications. From a sensor specification, it generates sensor definitions that are used to instrument the application code. Information about sensor definitions is also used in calls to the library. The safety constraint specifications are compiled into a sequence of Tcl commands that are executed by the Tcl interpreter resident in the executable [5]. The interpreter is used initially when the graph is being created and thereafter only when a constraint is dynamically added to C-net, so the performance impact generally associated with interpreters is minimized. The dispatcher is responsible for event handling, event distribution, action execution, and user command handling. The user interface provides a vehicle through which commands to add, modify, and remove constraints are issued.

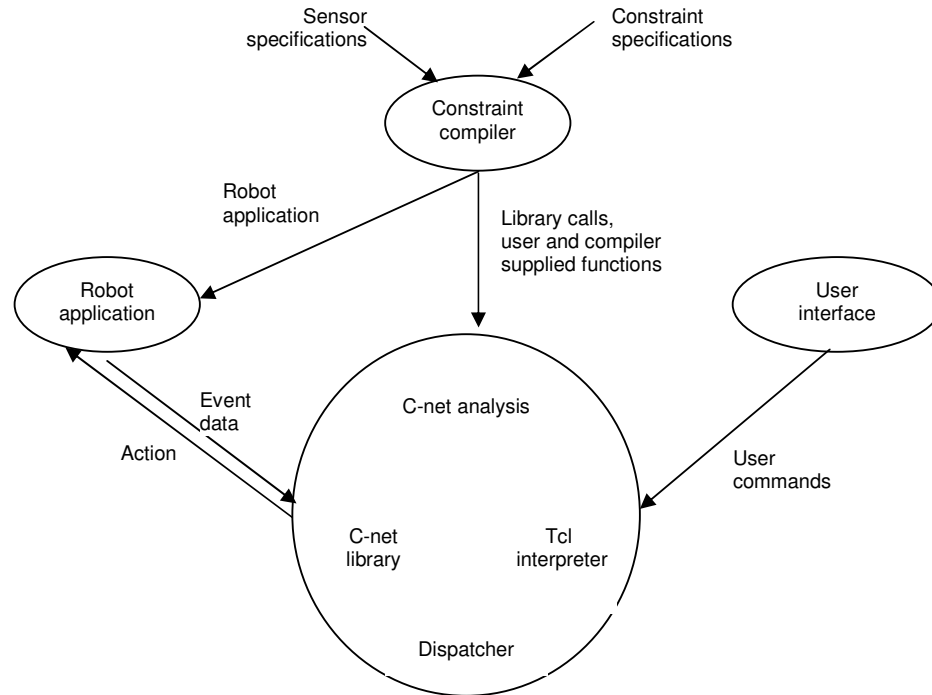


Fig.1 The system architecture

The hazard detection:

- accommodates dynamic applications through dynamic constraint addition, enabling, and disabling.
- provides synchronous and asynchronous constraint checking with predictable performance and low perturbation.
- supports specification of complex, multi-source hazards.

In contrast to previous static analysis approaches, C-net accommodates dynamic applications by allowing dynamic constraint addition. Additional constraints can be specified over a newly created application process, task, or object as long as the application entity is instrumented. Synchronous and asynchronous constraint checking is provided by flexible placement of constraints. This allows constraints with localized sensor data-needs to be evaluated in the sensor itself. Predictable performance guarantees can be made because of the DAG restriction on the constraint graph. Low perturbation is achieved through the use of a separate thread to perform monitoring and synchronous constraint checking. Perturbation could be further reduced by employing a technique used in software fault isolation of embedding monitoring related instructions in open slots in the instruction stream.

A multi-source hazard is a hazard that occurs when a combination of conditions exist: a failed valve indicator, condensate pump out of order, and a relief valve failing to close, for example. It is likely that each condition occurs

in a separate piece of controlling software. To enforce such a hazard, the analysis tool must be located such that it can gather information from each source. This precludes a completely synchronous approach. Additionally, the more complex the hazard, the stronger the need for a language in which to express complex relationships between conditions of a hazard and between hazards themselves. There is ample justification for implementing the analysis tool outside the application. The external nature of the tool makes possible the specification of complex constraints.

3. AUTONOMOUS ROBOTS

The application used in this work is a multi-agent reactive robotic system simulation. The work was carried with robots in the simulated environment with each robot running as a separate thread and employing shared memory as a communication medium. The robots perform one of three tasks: forage, consume, and wander. During *forage*, a robot wanders in an environment, looking for attractors. Upon encountering an attractor, it moves toward the attractor, attaches itself, and returns the object to a specified home base. During *consume*, a robot wanders the environment, looking for attractors. Unlike forage, after attachment a robot performs work on the object in place. In the final task, *graze*, discrete attractors are not involved; the object is to completely cover, or visit, the environment (akin to mowing the lawn). A robot searches for an area not grazed, moves toward it, then grazes until the entire environment (or some percentage thereof) has been covered. The robots are implemented as a three state finite state machine with the state selection dependent on the current task. For example, in the forage task a robot can be in *wander*, *acquire*, or *deliver* state. The robot roams freely in wander state. It transitions to acquire state when it has detected an attractor and to deliver state when it is returning the object to the home base. Once the object has been delivered, transition is made to back wander.

There are meaningful constraints that can be specified over the behavior of the group of robots just described. For instance, suppose a group of robots are tasked with digging up drums filled with radioactive material that is buried over a large area and moving them to a central site on higher ground. In retrieving a drum, one of the robots is splashed by the thick, murky contents and begins emitting radiation. The user might be greatly concerned to keep the other robots away from the contaminated robot. So in devising a constraint, the user specifies a hazard based on the concept of a danger zone, a region around the contaminated robot in which a robot approaching to help is in a danger but not imminently so. The user then wants to be notified when an approaching robot is within 10ft of the contaminated robot even though danger of picking up debris from the contaminated robot does not occur until the approaching robot is within 5ft. With this constraint in place and given a work area large enough and the number of attractors great enough, the user may be confident enough to allow work to continue even after a robot becomes contaminated. The constraint ensures the user will be notified when a robot has been contaminated, and will be notified again, with enough time to respond, when another robot has entered the **danger zone**. The danger zone constraint could trigger an alert to the operators to halt the approaching robot or it could invoke a user-defined handler in the approaching robot to change its state so it no longer seeks to assist the contaminated robot. The constraint for the scenario is: if the radioactivity level of a robot exceeds 200 roentgens per hour, then a violation occurs if a robot approaches within 10ft of the radiating robot more than once.

The constraint specifies that if any robot becomes radiated, a violation occur if any other robot comes within 10ft of the radiated robot. The second sample constraint is gathering of statistical information for evaluating performance: if during forage task, the average amount of time spent by robots in the wander state exceeds a threshold while progress toward a goal is less than some minimum, a violation has occurred.

That is, the user may be interested in knowing when the average amount of time a robot spends wandering as a function of the total amount of time working exceeds some reasonable estimate. As long as the value is reasonable, the user is willing to let the event pass without being notified. Why not simply modify the robot application to enforce the constraint behavior? Embedding constraint checking in the robot application suffers the same limitations as other embedded approaches: first, a single robot knows at most state and goal information about another robot. Global, statistical knowledge cannot be known or computed by any one robot. Second, constraints that dynamically modify parameters such as attraction/repulsion forces could be added to the application but at the cost of recompiling and relinking the application code.

4. THE RULE LANGUAGE

The language that is adopted is based on the active database rule definition language of the Starburst system. The language consists of five commands: **create rule**, **alter rule**, **drop rule**, **activate rule**, and **deactivate rule**. The **create rule** is used to define a new constraint. The syntax of the command is:

```
create rule name on event-type
if condition
then action-list
```

The *name* titles the rule and each rule is defined on a set of event-type. The **if** clause specifies the condition of the rule. The condition is the constraint that is checked when the rule is triggered by an incoming event. The language uses the temporal query language ATSQL2 for specifying the condition. ATSQL2 is a variant of TSQL2 and is currently being proposed for incorporation into SQL3. The condition can be any valid ATSQL2 query. The **then** clause specifies the rule's actions. An action is executed when the rule is triggered and its condition is true. Fig.2 illustrates a **create rule** command. The command creates a rule named C:1.

```
CREATE RULE C:1 ON robotRad, robotState
IF
SELECT radiatedRobot r.ID, r.rad
FROM robotRad as r, robotState as s
WHERE
s.task = FORAGE and s.state = DELIVER and
r.rad >= 200 R and r.ID = s.ID
THEN
STEER disableRobot r.ID
```

Fig.2. the creation of a rule named C:1

The event sources needed by this rule are *robotRad* and *robotState*. Event sources may originate from a sensor in the application or as the result of another query. The **if** statement delineates the rule's condition. The condition can be any ATSQL2 query. The query is composed of a SELECT statement, a FROM statement, and a WHERE statement. The SELECT statement builds a new relation or event type from attributes of one or more existing event types. The derived event, radiatedRobot, will contain three attributes: a robot ID and roentgen level taken from the robotRad relation and a timestamp derived from the timestamps of the tuples satisfying the condition. The FROM statement defines variable names *r* and *s* that will represent the event types *robotRad* and *robotState*, respectively. The WHERE clause specifies a predicate on the explicit attributes that selects those events that will contribute toward the new event type. The **then** statement delineates the action list. In this example there is a single action, a *STEER* command. The first parameter to *STEER* names a function in the robot application to be invoked. The second parameter identifies the task to be affected by the action.

4.1 Application of Query Language for Online Monitoring

A relational database query language is used to specify queries that are evaluated against event streams. The significant difference between evaluating queries against a database and evaluating them against an event stream is that in the latter constraints must be evaluated against a *conceptual* database rather than an *actual* database. That is, no database *per se* exists. Instead, each constraint must have sufficient storage to maintain the application state it needs. Hence, a given event may exist in multiple nodes at any moment; the length of time an event remains in the node depends on the attributes upon which the query is based and the complexity of the query.

There is an issue of efficiency with which one must deal when executing queries in an on-line analysis environment. Instead of a query being executed periodically or upon user demand, and a set of tuples satisfying the query returned, the query is in essence executed every time an event arrives. What keeps this characteristic from being wholly inefficient is that a constraint by its nature will reject the majority of the events it receives.

4.2 Action Statements

The user has control over the set of executing constraints in two ways. The first is through issuing commands (e.g., **alter rule**) through the user interface. The second is through action statement of the rule. An action statement is a command listed in the action part of a rule. The set of allowable actions must be basic enough and broad enough such that when taken alone or combined, they allow the user to effect a desired behavior. Three action statements are supported to:

- invoke a user-defined function in node
- invoke a handler in the application
- enable or disable a constraint.

Invoking a user-defined function in a node may do something simple like causing a bell to ring or a message to be printed to an operator console or something more complex like collecting statistics. Invoking a handler in the application causes a steering command to be issued to the application that results in the execution of a user-defined function residing in the application. The mechanisms for such steering are discussed in the next section.

4.3 Enabling and Disabling Constraints

When a constraint is disabled, it no longer processes events although its code is still resident in the system. Enabling, the default mode, reverses the disabling action and allows the constraint to process events once again. Constraint enabling and disabling has multiple useful application. This approach allows dynamic constraint addition must deal with the eventuality that an added constraint will conflict with an existing one. The effect of such conflicting constraints is that one or the other will continuously be violated; perhaps not a desirable behavior from the point of view of the user. To obviate the problem, the **DISABLE** clause is provided as a means for the user to manage the conflict. The enable and disable clauses are also useful for loosely hierarchical error recovery. For example, when a robot encounters an obstacle in its path, its first response could be to wait some amount of time in the hope that the obstacle will move. If this simple error recovery fails, its second response would be to determine a new route.

5. LIBRARY AND DISPATCHER

The rule language provides a means for specifying constraints. The library and run-time environment, on the other hand, provide the means for the specified constraints to be transformed into individual executable entities and the mechanism to execute the constraints against the incoming event stream. The **library** is a collection of functions that build two types of components: operations and nodes. An *operation* is a component that implements one of selection, projection, or join. Control flows between operations by procedure calls. A *node* is a collection of operations with a number of methods defined for it. A *node* can be created and can accept connections. Additionally, it possesses general information about itself so it can respond to questions as to its state (active, inactive), it can be called upon to activate or deactivate itself, or it can return a list of its input events and output events. Control flows between nodes under the control of the dispatcher.

The **dispatcher** controls net execution. At startup, it awaits nodes to register their existence and event needs. The dispatcher links those nodes together having data dependencies as shown by their event lists. During execution, the dispatcher accepts events from the application and routes them to the interested nodes and accepts and executes commands from the user.

5.1 From Rules to Executable Entities

The transformation of a constraint from a rule to a node begins with the rule compiler. The rule compiler parses a constraint, and converts it to a relational algebraic expression in conjunctive normal form. From the relational algebraic expression an abstract syntax tree is constructed and it is from this that the optimizer performs compiling time optimization before generating a sequence of Tcl commands. So roughly for every select, project, and Cartesian product operation in the relational algebraic expression, there is a corresponding Tcl command in the script file. Also included in the script file are Tcl commands to build the node, the entity to which the operations belong, and to link the operations in the proper order. The initial set of queries becomes executable nodes when the analysis tool is first executed. The dispatcher accepts the name of a script file as an argument, and invokes the Tcl

interpreter, passing it the name of the script file. The interpreter executes the commands in the file, each command resulting in a call to the library. Through a sequence of calls the node is built. The script file can contain any number of queries in any order. Linking queries takes place when a node registers itself with the dispatcher.

5.2 Multi-Source Hazards

Multi-source hazards, hazards which can be described as consisting of events from multiple sources, make up an important and substantial subset of hazard descriptions [4]. Implementing detection of such hazards requires making tradeoffs between latency, perturbation, and ease of use. Latency is far more of an issue in external approaches to hazard detection, when constraints are embedded directly in the application code. But embedded approaches suffer from increased perturbation and decreased breadth of potential event sources. The external approach trades decreased perturbation and the ability to specify multi-source hazards for increased latency. More complex hazard descriptions also require a more general language, such as a query language, to distinguish events from multiple sources and describe complex relationships between events in a natural way.

To achieve efficient communication between multiple sources and the analysis tool, a communication infrastructure, Data Exchange is employed. Data Exchange provides for binary IO of event data between the multiple sources and the analysis tool. The full features of Data Exchange, in its ability to forward data to multiple clients based on event type, can be utilized in the version of Cnet underway where the analysis tool is itself distributed.

5.3 Dynamic Applications

Support for dynamic applications occurs in part by the design concept of *event types*. All events possessing the same set of attributes belong to the same event type, regardless of the sensor from which they originate. For example, every event consisting of a robot ID, location, and current time stamp is of the *robotID* event type, regardless of which of the many robots generated the event. A constraint node registers its interest in event types, so any dynamically created application task generating events of a known event type will automatically be included in the constraint checking done by nodes accepting events of that type. Support for dynamic applications based solely on event types would only partially solve the problem. Also needed is the ability to add new constraints. With constraints descriptions encoded as Tcl scripts, dynamic constraint addition becomes straightforward. As shown in Fig.3, when an *ADD* command arrives at the dispatcher (from the user interface), the dispatcher invokes the Tcl interpreter, passing it the script file name as an argument. The interpreter executes the script, the execution of which results in a series of calls to the C-net library to build a node. The new node, once built, automatically registers its existence with the dispatcher, providing it a list of input and output events of which the dispatcher adds to its internal lists. Control then returns to the dispatcher and event processing resumes. Events arriving at the dispatcher for which the new node is interested will immediately be forwarded to the new node.

5.4 Synchronous Constraint Checking

Asynchronous constraint checking has been explained in some detail in the general discussions of C-net. Less has been said, however, about synchronous constraint checking. Synchronous constraint checking is checking performed in the application data space. It is suitable for filtering data to reduce the volume of events flowing to the analysis tool. Synchronous placement is restricted to constraints having a single input event type. The restriction is necessary because the constraint node is invoked from within the sensor generating the event type. The constraint is evaluated prior to the event being forwarded to the analysis tool. The constraint compiler provides support for handling synchronous constraint checking by identifying those constraints suitable for synchronous placement. A Tcl interpreter creates the synchronous nodes in a separate thread in the application data space.

6. CONCLUSION

This paper addresses the problem of improving software safety through hazard detection. The approach consists of a query-like language and compiler for specifying hazards, a library for creating operations and nodes, and a run-time tool to dispatch arriving events, trigger node execution, and handle dynamic node addition. Hazards are often described by a number of events occurring simultaneously. Any realistic approach to hazard detection must

accommodate the complex hazard as easily as it accommodates the simple one. This general language approach allows for the specification of complex constraints specified over distributed components.

For a hazard detection approach to be realistic, however, it must be responsive. Hazard descriptions often include components outside the software system. The hazard detection approach should be able to accommodate a description that includes state from these components as well. This approach also accommodates these components by adding shadow objects. These shadow objects would then be the source of state information for the analysis tool. Finally, the purpose of hardware hazard detection is often as a safeguard. It is the mechanism to which an engineer turns for the extra measure of safety. Such a device must at all times provide the extra measure of safety, not decrease the overall safety of the device on which it is placed.

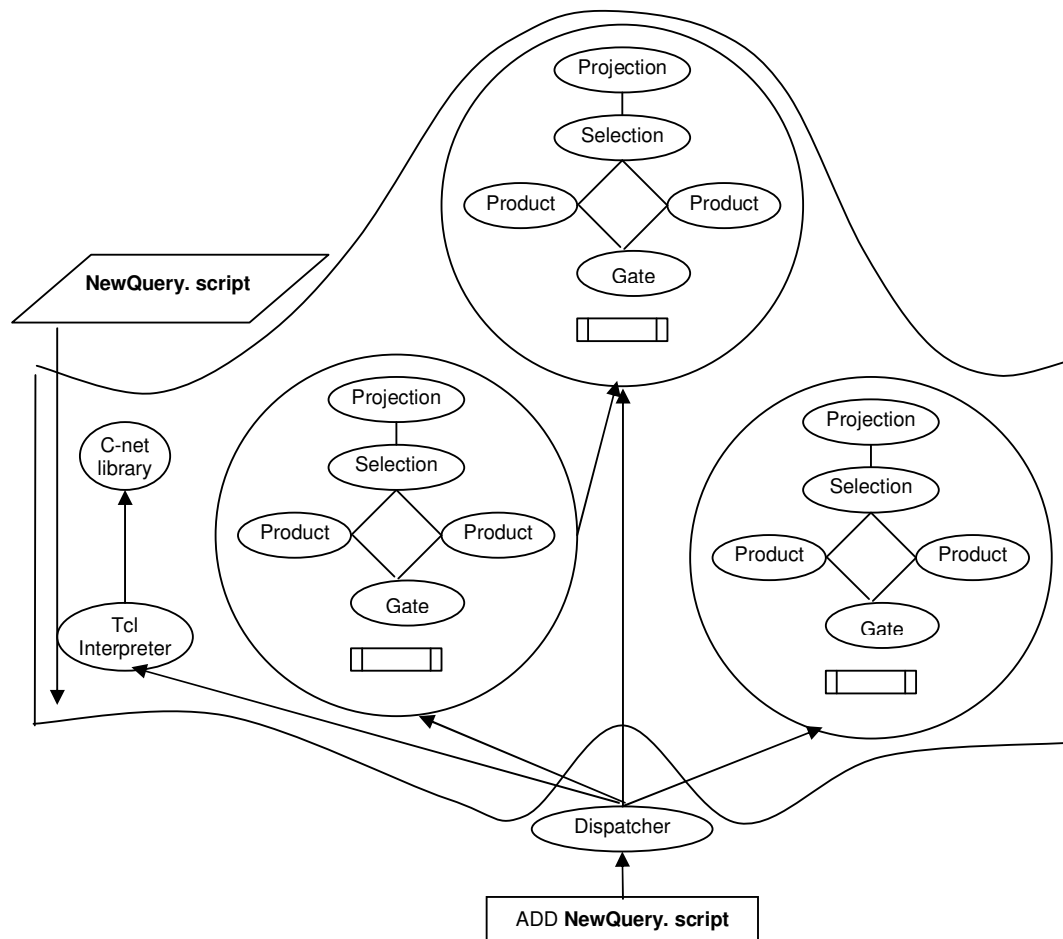


Fig.3 The C-net analysis

References

1. M. P. E. Heimdahl and Nancy G. Leveson. Completeness and consistency checking of software requirements. *IEEE Transactions on Software Engineering*, 22(6), June 1996.
2. Louise E. Moser and P.M. Melliar-Smith. Formal verification of safety-critical systems. *Software – Practice and Experience*, 20(8): 799–821, August 1990.
3. W. Hammer. *Handbook of system and product safety*. Prentice Hall, 1972.

4. Nancy G. Leveson. Software safety in embedded computer systems. *Communications of the ACM*, pages 34–46, February 1991.
5. John Ousterhout. *Tcl and the Tk toolkit*. Addison- Wesley, 1995.
6. Chennakesava Reddy A, Robot and its engineering field applications, *Engineering Advances*, Vol.14, No.1, pp.33-35, 2002.
7. Rajasekhara Reddy B and Chennakesava Reddy A, Obstacle avoidance of planar redundant manipulators for pick-and-place operations using real-coded genetic algorithm, National Conference on Advances in Manufacturing Technology, Palghat, 15-16th February, 2003.